SANS

Analyst Program

Whitepaper

Anatomy of a Ransomware Operation

Written by **Jake Williams**

September 2022

BlackBerry Cybersecurity

Introduction

When you hear about yet another ransomware attack, do you truly understand what's happened to the victim organization? Ransomware has changed significantly, even over the past few months. If you're still thinking about ransomware as a drive-by download with a single-system impact, that's an outdated perception. Ransomware deployments today are disruption operations designed to inflict maximum pain on the victim organization. They typically have a single goal: getting the victim to pay by any means necessary. In this paper, we discuss the anatomy of a ransomware operation and the techniques used by ransomware threat actors to provide maximum disruption to your operations. We also discuss what you can do to detect the ransomware operation—before the business disruption begins.

Ransomware Evolution

When ransomware first became widespread, most contemporary reporting described how businesses were affected by a malicious payload detonating on a single user's machine. The consequences of losing a single machine are certainly not trivial, particularly in smaller organizations (and threat actors could hit the jackpot by encrypting the right machine). Some ransomware would also target network drives, leading some organizations to instruct users to dismount network drives that were not in active use. Ransoms back then were measured in the low tens of thousands of dollars.

These stories from what now seem like the good old days of ransomware bear little resemblance to the state of affairs today. Ransomware operations over the past few years have employed human penetration testing teams (sometimes referred to as HumOR, or human-operated ransomware) to move through the network, elevate privileges, find and exfiltrate data of strategic value, and encrypt network devices. Threat actors also hunt for backups, with the goal of rendering them unusable for restoration. Once the encryption operation begins, most organizations can do painfully little to limit its impact, which speaks to the critical need for detection before the encryption begins.

Ransomware threat actors often specifically search for regulated or customer data, with core groups coaching their affiliate teams on specific search terms to use, such as "confidential" and "liability." Ransomware operators threaten organizations that if they don't pay the ransom, their data will be released publicly (typically referred to as "double extortion"). Some operators advertise the impending release of the data on their blogs, complete with countdown timers and descriptions of the data they intend to release. When that doesn't work, some ransomware operators have been observed contacting business partners and even customers of the victim organization, trying to gain additional advocates for paying the ransom. One final tactic used to increase the odds of payment is to perform DDoS operations against the victim organization while it is trying to recover. The takeaway from these activities is that the stakes have been raised and ransomware operators have refined their techniques to maximize payoffs.

Stages of a Ransomware Operation

In this section, we discuss the phases of a ransomware operation. These phases can broadly be categorized as shown in Figure 1.

The next section expands on the analysis of these phases and explores opportunities for detection. First, let's discuss the flow of a ransomware operation from the threat actor's viewpoint.

Initial Access

Initial access is usually obtained through brute-forcing access to internet-accessible services such as Remote Desktop Protocol (RDP), exploiting a known vulnerability, or using commodity malware already in the environment. The increased use of commodity malware for initial access is alarming. Today there is no such thing as a commodity malware infection. Security teams should treat every malware detection as a potential ransomware precursor (because many, in fact, are). Security teams should be aware that although commodity malware may be used for initial access, the backdoors dropped for the ransomware threat actor to have access are usually one-off deployments created specifically for the target environment, typically leading to lower detections.

In addition to backdoors and trojans, another commodity malware family used by threat actors for initial access consists of info stealers. These dump credentials, exfiltrate them to threat actors, and often do not maintain persistence on a compromised target. The credentials gained from info stealers may also be sold to ransomware threat actors and used to gain access through internet-accessible servers and services, all while appearing to be a legitimate access method.

Although commodity malware may be used for initial access, the backdoors dropped for the ransomware threat actor to have access are usually one-off deployments created specifically for the target environment, typically leading to lower detections.

Command and Control

After the threat actor gains execution on a target system, they must establish command and control (C2) of the malware deployed. Many red teams focus on esoteric C2 methods that evade most out-of-the-box detection strategies, but ransomware threat actors typically use HTTPS beaconing on standard ports—because it works.

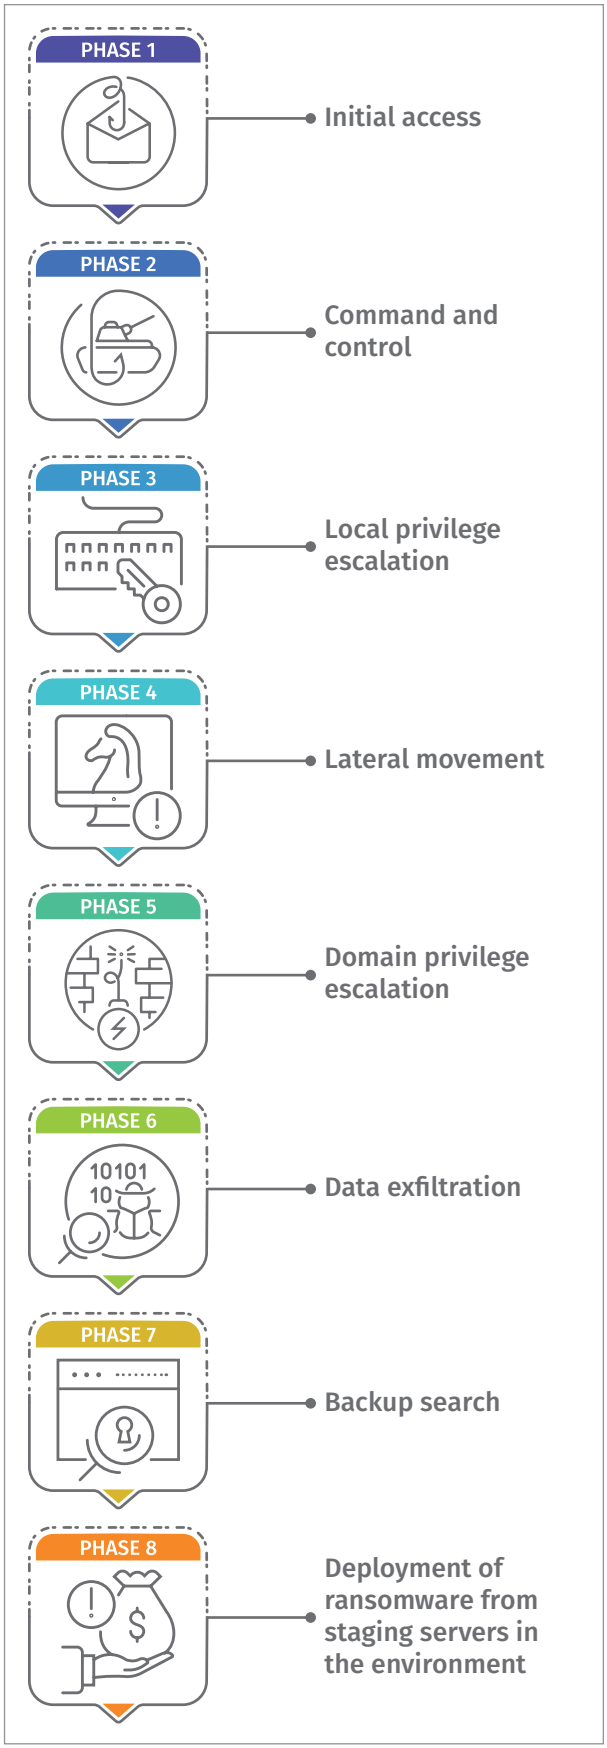

Figure 1. Phases of a Ransomware Operation

Local Privilege Escalation

Before moving laterally, the threat actor typically attempts to escalate privileges on the local system. Local privilege escalation usually abuses insufficient filesystem or registry permissions or exploits an unpatched vulnerability on the system. Unfortunately, most operations teams struggle with patching third-party applications, many of which execute in privileged contexts on end user systems. This weakness leaves the threat actor plenty of opportunities to gain privileged execution contexts on most systems. For organizations still allowing users to log in with local administrative privileges, it should be noted that User Account Control (UAC) is not considered a security boundary by Microsoft, and many bypasses around UAC are well known to threat actors.

Lateral Movement

Threat actors will need to move laterally in most networks to find a machine where a domain administrator is logged in. This movement may be achieved by using a tool such as SharpHound to use the currently logged-in domain user to generate maps showing how to compromise a domain admin account using credentials already known to the threat actor. In other cases, techniques such as SMBrelay may be used.

When systems are configured with the same local administrator password, threat actors may use pass-the-hash techniques to log in to remote systems where privileged domain accounts are logged in. Alternatively, they may move laterally simply to expand their footprint in the network, potentially searching for data to exfiltrate.

Domain Privilege Escalation

To inflict maximum disruption and maximize the likelihood of ransom payment, threat actors will need to elevate their privileges to domain admin (or rough equivalent). This elevation typically occurs by:

- Exploiting an unpatched vulnerability (via PrintNightmare or ZeroLogon) on a server with logged-in domain admins
- Targeting cached domain credentials that are cracked offline
- Waiting for a domain admin to log in to a system where the threat actor already has local admin
- Performing various Kerberos attacks (such as Kerberoasting)[1]

Regardless of the method used to escalate privileges to domain admin, organizations should recognize that domain admin is not the goal of the threat actor—it is merely a means to an end. In several ransomware cases, the threat actor deployed ransomware without gaining domain administrator permissions by using a domain account that was added to the local administrators' group on every endpoint in the domain.

Regardless of the method used to escalate privileges to domain admin, organizations should recognize that domain admin is not the goal of the threat actor—it is merely a means to an end.

[1] "Kerberoasting Attacks," Crowdstrike, June 13, 2022, www.crowdstrike.com/cybersecurity-101/kerberoasting/

Data Exfiltration

Ransomware threat actors exfiltrate data prior to encryption to perform the aforementioned double-extortion operations. The intent of this activity is to threaten the public release of sensitive data for organizations that might otherwise recover without payment. Many ransomware threat actors specifically target regulated data for exfiltration, hoping to maximize the pain of any subsequent release.

In some extreme cases, threat actors have used stolen data to communicate with business partners and even customers of the victim organization. In these cases, the threat actor tells the third party that the data they entrusted with the victim organization has been stolen and will be released publicly if the victim organization doesn't pay (see Figure 2). This tactic is extremely effective in influencing the pay/no pay decision calculus for many organizations because the third parties often threaten to discontinue their relationships with the victim organization if their data is made publicly available.

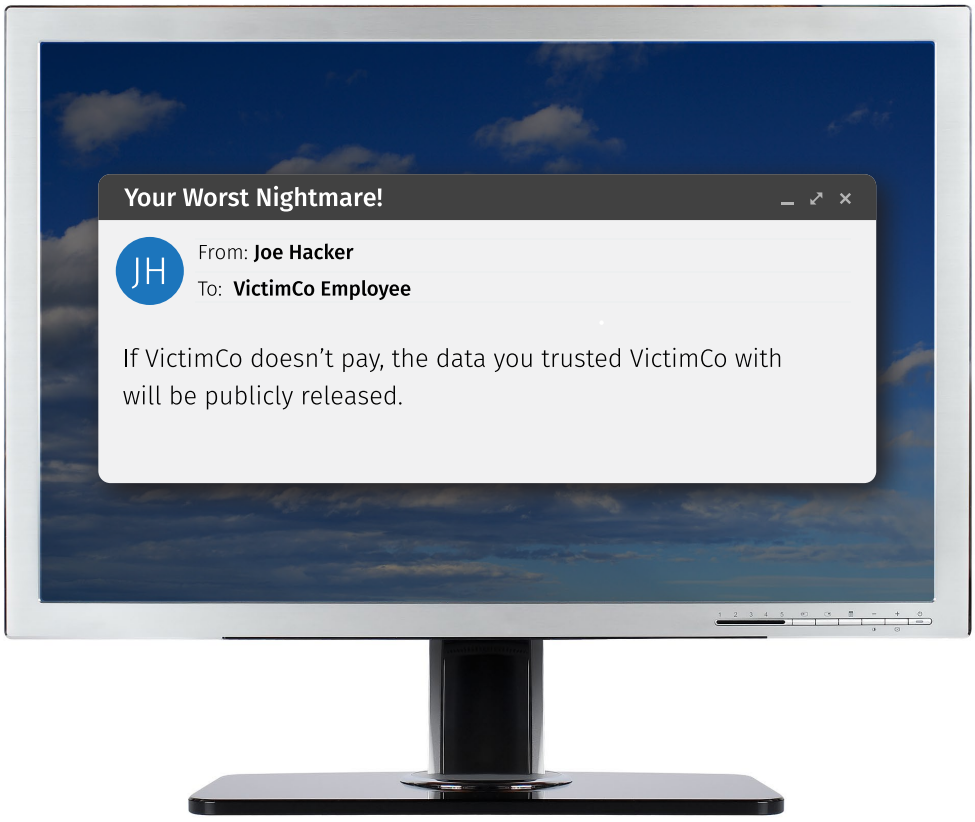

Figure 2. Example of Double Extortion Attempt

Lastly, some ransomware-adjacent threat actors are engaging in extortion-only operations. In these operations, everything is usually identical to their ransomware counterparts, with the notable exception that encryption is never performed. Because the data exfiltrated is the only leverage the threat actors bring to bear against the victim organization, the quality and quantity of data exfiltrated are usually greater than in ransomware operations. We only mention extortion-only operations in this paper because the initial TTPs (tactics, techniques, and procedures) are so similar, and both types of attack are financially motivated.

Backup Search

Threat actors wish to deliver maximum interruption with their ransomware operations, and one of the best ways to hinder recovery is by ensuring backups are unusable. Threat actors will search for backup servers and attempt to render the backups useless for recovery. Some threat actors have been observed completely deleting volumes on storage filers when they contain only backups. In other cases, backup files have been deleted. Even in cases where immutable backups are employed, organizations should note that the threat actor's malware is already in the backups for some systems, including the zero patient and any servers they laterally move to and establish persistence on prior to beginning the encryption operation. Even when backups are not accessed by the threat actor, one does not simply "restore from backup" to recover from a ransomware event.

Ransomware Deployment from Staging Servers in the Environment

At the final stage of the operation, the threat actor deploys file encryption ransomware from staging servers in the environment. These staging servers are typically chosen because they are positioned in the network to be able to talk to practically any endpoint using SMB. The vast majority of ransomware operations use SMB for distribution. Even those that use Background Intelligent Transfer Service (BITS) as the transfer mechanism typically authenticate over SMB. End user workstations are typically not chosen for this activity because they are in subnets that may limit their SMB communication. This situation is doubly true of networks where end users are remote workers connecting over VPN.

When deploying ransomware from a staging server, the ransomware operator typically uses either a PowerShell or a batch script to loop through a list of target endpoints. The list of target endpoints is usually provided as a text file containing IP addresses or hostnames and is normally supplied as a command line argument for the script. For each target endpoint, the threat actor copies the ransomware to the machine and ultimately executes it.

Detection Opportunities

In this section, we explore opportunities for detection prior to the start of the encryption operation. Although the threat actor ideally would never be permitted to operate in the network at all, it is important to consider how to mitigate an operation in progress. Anything less leaves us with a single line of defense: a classic "hard crunchy outside, soft gooey inside" situation.

Initial Access

If the threat actor exploits a known vulnerability, advanced endpoint security controls can likely detect many generic methods used by threat actors, including heap spraying and return-oriented programming (ROP). Although network IDS used to be a primary detection method, ubiquitous encryption has hampered network detection in most environments.

Regardless of the method used to obtain initial access, the threat actor will deploy a beachhead for remote access. In most recently observed ransomware operations, this beachhead is a Cobalt Strike Beacon custom build unique to the victim network. Because the backdoor is custom built for the target environment, hash and static signature matching employed by traditional antivirus is worthless. Advanced endpoint solutions will often provide detection by observing the activities performed by the backdoor or through scanning memory for the original executable content as it is deobfuscated into memory.

C2

Many ransomware operators use familiar C2 frameworks, such as Cobalt Strike or Brute Ratel. NDR (network detection and response) tools can be used to identify periodic beaconing from inside the network to the C2 server. However, because many of these tools provide jitter (varying their callback periodicity while introducing some randomness), additional detection strategies are needed. By tuning NDR parameters to cast a wide net for beacons, analysts will necessarily discover false positives. When supplemented with other detection methods, however, the number of false positives is entirely manageable and may yield unique detections before the encryption operation begins.

By enriching DNS logs with registration information, analysts can readily identify those domains most likely to be used by threat actors and investigate the endpoints communicating with them.

DNS logs provide another key method for detection. Threat actors often utilize domains that have been registered relatively recently for their C2. Even when domains have been given time to age before being operationalized, they typically receive little legitimate traffic, making them stand out with the right cyber threat intelligence (CTI) tooling. Manually checking individual domains for their age and reputation doesn't scale. By enriching DNS logs with registration information, however, analysts can readily identify those domains most likely to be used by threat actors and investigate the endpoints communicating with them.

Threat intelligence can provide information about compromised domains, expired domains, dynamic DNS domains, and TOR exit nodes, each of which can be used for detecting C2. Threat actors may compromise legitimate domains and use them for C2 operations. In the classic case, domain operators using known-vulnerable software (such as unpatched WordPress installations) are compromised. Some threat actors target domain registrar control panel accounts so they can directly modify DNS and send C2 traffic to machines under their complete control. Dynamic DNS domains, although less common than in years prior, continue to be used for C2. Because these domains are regularly blocked by enterprise network admins, dynamic DNS operators add new domains on a fairly regular basis (see afraid.org). Tracking these domains manually would be impossible, but threat intelligence can be used to automatically update blocking (and alerting) lists.

Local Privilege Escalation

Ransomware operators primarily exploit unpatched vulnerabilities to gain local admin. Although there are certainly opportunities for local privilege escalation through abusing misconfigurations, the operational pace precludes searching for these in most victim environments. Detecting (and in many cases, blocking) local privilege escalation exploits will usually require monitoring of API calls, something best suited to endpoint detection and response (EDR) software.

Threat actors gaining initial access through commodity malware may already have local admin privileges. These privileges, however, are gained by exploiting the same unpatched vulnerabilities that would otherwise be exploited manually. Exploitation of a misconfiguration (usually taking advantage of weak filesystem or registry permissions) will usually not be detected by EDR. Security posture management tooling, however, can be used to ensure systems aren't deployed in these vulnerable configurations and that any drift from an established security baseline is quickly detected and remediated.

Domain Admin Escalation

Each of the methods used for elevation to domain admin can be detected with the appropriate instrumentation on servers, but it is most important to monitor domain controller logs for abnormal behavior. ZeroLogon and Kerberoasting each has a unique signature in domain controller logs. In the case of ZeroLogon, look for an authentication record (Security Log Event ID 4624) with the machine account of the domain controller coming from a remote IP address. Machine accounts rarely (if ever) authenticate from remote locations. Kerberoasting attacks can be detected by monitoring Kerberos service tickets, particularly those using RC4 encryption. This detection tactic requires enabling "Audit Kerberos Service Ticket Operations" and monitoring Security Log Event ID 4769 for service tickets using nonstandard (and weak) encryption methods. Other privilege escalation techniques will still result in logs on the domain controller that can be used for detection. Pay special attention to any logins that should not be accompanied by Security Event ID 4672, indicating that special privileges were assigned to the login.

Some organizations don't deploy EDR software on their servers, fearing service interruptions or thinking the security benefits aren't aligned with availability risks. Unfortunately, they make that choice at their peril. As noted previously, EDR and endpoint security software will often detect the exploitation of vulnerabilities. They are also likely to detect either the malware payload delivered or the actions taken by the malware payload.

For example, once threat actors gain domain admin privileges, they will usually target the Active Directory (AD) database ntds.dit on a domain controller. Because this database is locked open during operation, the threat actors will often use volume shadow copy or raw disk access, both of which are trivially detected by EDR. Most EDR platforms have rules for the execution of vssadmin.exe, used for creating volume shadow copies used by threat actors.

Data Exfiltration

Before ransomware threat actors can exfiltrate sensitive data, they must locate it. If commands such as “find” or “grep” are used, the search strings used will be revealed to analysts who have any command line detection. PowerShell is also often used to scan files for sensitive content. Search strings can be detected using advanced PowerShell logging features, such as script block logging.

Regardless of the methods used for locating sensitive data, it is usually aggregated into a staging location before being archived and exfiltrated. Filesystem monitoring tools, such as those offered by EDR, can be used to detect file writes in the new staging location. Data loss prevention (DLP) tools can identify file contents and may be configured to alarm on files containing sensitive data patterns stored in unauthorized locations, though this will require configuration by the organization.

Before exfiltrating data, the threat actor will often create archive files, such as zip, rar, or 7-zip. These archives are usually created using a command line tool, such as 7za.exe. Although archives are common in most operating environments, creating them from the command line is not. EDR software can be used to alert on command line execution of archive utilities. Analysts get a bonus of learning where the staging directories used by the threat actor are located, as well as other TTPs, such as:

- Naming convention for the archive file
- Method of compression used
- Passwords used to encrypt the archive (if any)

As a practical matter, threat actors create archives for two reasons:

- Simplifying exfiltration (fewer individual files to exfil)
- Compressing data (less data on the wire)

So how do ransomware threat actors exfiltrate victims’ sensitive data? They typically use tools such as pscp.exe (Putty SCP), file transfer tools such as FileZilla or rclone, or file synchronization tools such as Yandex Disk. Of particular interest is rclone. The rclone.exe binary has been observed in an extremely high number of ransomware operations. The fact that most organizations do not legitimately use rclone makes it a high-confidence detection. Instructions for using rclone have even been included in leaked ransomware threat actor playbooks.

Once the data exfiltration begins, it can be detected on the network in a variety of ways. One is by inspecting the volume of data transferred. When FTP, SFTP, or SCP are used, there is typically a small number of outbound transfers (sometimes only one) with relatively high volume. When file synchronization tools such as Dropbox are used, these transfers should be automatically triaged as high risk unless the specific service used by threat actors is also officially used by the organization. Even in cases where the threat actor and the organization use the same tool (Dropbox, for example), endpoint controls such as DLP can identify the non-corporate account in use. When tools such as rclone are used, the threat actors often send files to Mega, a network destination with no legitimate business use in the vast majority of organizations. NDR and DNS logs rule the day here.

Note that Windows 11 supports systemwide DNS over HTTPS; if this is enabled (as recommended in some security checklists), centralized DNS logging will be missing entries from these endpoints. Although DNS over HTTPS is enabled by default for many modern web browsers, this is less of an issue for capturing threat actor traffic, such as C2 and data exfiltration (which do not commonly utilize a browser).

Backup Search

Backup servers are typically discovered by threat actors using one of three methods: searching for backup server hostnames, examining backup client software configurations on compromised systems, and port scanning.

When searching for backup server hostnames, threat actors typically use a tool such as adfind.exe. This tool is not commonly used by anyone other than system administrators, so any use can be considered suspect. When the target of the search is strings such as "backup" or common backup tools such as Veeam, most organizations should consider this a likely indicator of a ransomware operation in progress until confirmed otherwise.

Threat actors also use port scanning to discover backup servers. Scanning for common ports, such as TCP port 9392 used by Veeam or TCP port 6106 used by Backup Exec, helps threat actors discover these servers. Detection can be achieved by looking for execution of port scanning programs on endpoints and looking for port scanning activity using NDR. Note that NDR is often a more reliable detection mechanism because, by this stage, threat actors have often disabled endpoint protection software (or scan from endpoints that don't have it installed).

Security teams should note that tools such as adfind.exe are LOLBins and as such are not likely to be alerted on by EDR using out-of-the-box configurations due to false positives. Because adfind is used relatively infrequently in most environments, tuning should be easy. Any adfind.exe execution with Veeam as a command line argument should be configured to alert, especially if the organization doesn't use Veeam as a backup solution.

When threat actors examine configuration of backup solutions, this means reading either registry values or configuration files associated with the backup agent. Few processes should access these files or registry keys, and those that do can be easily baselined using most endpoint monitoring tools (including Sysmon or EDR). After baselining, any unexpected process examining these configuration entries can generate an alert that points the analyst to the processes being used by the threat actor on the compromised machine.

Ransomware Deployment from Staging Servers in the Environment

Detecting ransomware deployment in the environment is fairly easy when systems are appropriately baselined. Unfortunately, because this represents the beginning of the encryption operation, organizations should use this detection as a last hope (a safety net of sorts), rather than relying on it as a primary detection method.

On most systems that are ideally positioned for malware, there should be little variation in the scripts (PowerShell or batch) that are executed in normal operations. This makes baselining and alerting on any unknown .bat or .ps1 script relatively easy with endpoint controls such as EDR. If threat actors try to use encoded PowerShell commands for this activity, they should be trivially detected (and alerted on).

To copy the ransomware executable to the target machine, a network share may be mapped using the net.exe command. Because the deployment operates in a loop across the target list, it is normal to see dozens (or even hundreds) of sequential executions of net.exe (and net1.exe) in endpoint monitoring systems. Network monitoring systems can also detect this activity relatively easily. Most systems used for ransomware deployment receive far more connections than they make. The ransomware deployment turns this model upside down, with the system establishing far more connections than it makes. When considering that these connections are almost exclusively on TCP port 445, the detection opportunities are obvious. NDR systems can be configured to alert on excessively high numbers of outbound SMB connections. Although this requires tuning and is not appropriate for all systems, it is nonetheless an effective detection for most. If file share access auditing is enabled in Windows for target systems, Security Log Event ID 5140 can be used to track access to the file share (including the source IP of the staging server and the account used to authenticate).

The threat actor must also trigger the malware to execute. Detections will generally depend on the method used, but generally, tracking command execution is easiest from the staging server. Endpoint detection controls can be configured to look for multiple executions of tools such as:

- Schtasks.exe (scheduled tasks)
- Bitsadmin.exe (triggers execution after transfer)
- Psexec.exe (Sysinternals systems administration utility)
- Wmic.exe (look specifically for “process call create” and “/node” arguments)

Of course, this is only a partial list of tools used to trigger remote execution. Threat actors may use other tools, and each of these examples can be renamed to avoid detection. However, renamed files will still have the same digital signatures (where applicable) and file hashes. Because they are executed against multiple systems (potentially hundreds), detections are aided by some threshold of execution in a given time period (more than 10 executions in a five-minute period, for example).

On target systems, new service creations may be noted for PSExec using System Event ID 7045 or Security Log Event ID 4697. Similarly, creation of new scheduled tasks can be detected by Security Log Event ID 4698. Process execution can be examined on the target system using Security Log Event ID 4688, but this is the beginning of encryption and therefore is not at all ideal for detection. Like so many other useful events, this is not enabled by default (highlighting the need for third-party tooling).

Next Steps

In this paper, we discussed how modern ransomware operations function, from gaining initial access through data exfiltration, all the way to beginning encryption. Additionally, we examined detection opportunities throughout the ransomware operation. Armed with this knowledge, organizations should examine their own technical controls and determine how many detection opportunities they are currently missing. We recommend that organizations that identify detection gaps examine the detection opportunities presented in this paper and either tune existing controls or deploy new ones. Lastly, regardless of their current detection posture, security teams should circulate this paper to IT operations teams to educate them on the current state of ransomware operations. Armed with appropriate understanding, IT teams can operationalize the concepts in this paper to more easily recognize existing alerts as possible portions of a ransomware operation (and act accordingly).

Sponsor

SANS would like to thank this paper's sponsor:

BlackBerry®

Cybersecurity