SANS

Analyst Program

Whitepaper

SANS 2021 Ransomware Detection and Incident Response Report

Written by **Matt Bromiley**

November 2021

ANOMALI®

Introduction

Having a plan in place is critical for an effective response to a cyber incident. Incident response (IR) plans help guide an organization when an adversary has taken a foothold in their environment. They are designed to help the security team successfully identify, scope, eradicate, and recover from threats to their environment. For many, their plans work well—if the adversary's goal is to maintain a long-term foothold and remain relatively stealthy.

However, what if the adversary was not interested in a long-term foothold, and their goal was not to slowly extract data out of the environment? Instead, what if the adversary was interested in financial gain and extortion, and cared little about the actual data they have access to? How do our incident detection and response capabilities stack up against these types of attacks? In recent years, we've seen a significant rise in exactly these types of attacks: ransomware.

Ransomware attacks are often quick and can cripple an organization, leading to various types of losses: financial, reputation, and trust. By detecting and responding to these incidents sooner, we can minimize these losses and remove the power from the threat actor.

In this report, we will address ransomware attacks head-on. These are a different type of attack and thus require that we approach incident detection and response differently. Security teams cannot afford to wait for the ransom note to be placed on a file server or for users to complain that they cannot access resources. We must catch the adversary earlier—as early as possible—to prevent any additional damage to the environment.

As you work your way through this report, you may come across issues that have not been addressed in your organization. While many organizations have basic security structures in place, it is possible that those plans were not created with an attack as pervasive as ransomware in mind. Regardless, we encourage you to ask your security teams or providers the "tough questions" about ransomware response. Some of these questions may include:

- Has your organization considered how it will respond to a ransomware attack, and are you prepared for such an event?
- If you have your own IR team, do you have a separate response plan for ransomware?
- If you do not have your own team, but outsource security capabilities, ask your managed security service provider (MSSP) if it has specific ransomware response actions.

This paper is also designed to help remind you how security teams should detect and respond to ransomware attacks, which differ from other threats. We have summarized our tips into a checklist at the end of this paper and encourage you to use that checklist with your security team to assess your current ransomware preparedness.

Defense and Detection

The first step to combatting a ransomware attack in your organization is to evaluate your detection capabilities. Detection is the primary area where resourceful security teams can gain an advantage over a ransomware threat actor. As we will examine, many techniques used during a ransomware attack are not too different from a non-ransomware attack. Oftentimes, it is the end goal of the adversary—the locking up of files and extortion of funds—that is the major difference. However, it is this key difference that forces us to detect an incident sooner rather than later.

Ransomware attacks are not *just* security issues. They often envelop the entire organization, including customer and reputation management, legal and finance involvement, along with securing the organization's digital assets and getting back to normal. Detecting these attacks helps everyone.

Once an organization's files are encrypted, ransomware threat actors control the "clock," as it were. They set deadlines that force the organization to move at an undesirable pace, often making rash decisions in attempts to control the situation. Even worse, ransomware threat actors are notoriously vocal about their feats, meaning that the attack often becomes public, and the organization has even more battles to fight. Therefore, the goal for detecting a ransomware attack as early as possible is not only a security issue but also an effort to protect the entire organization.

To understand how we can detect a ransomware incident, let us first look at the high-level components of a ransomware attack (see Figure 1).

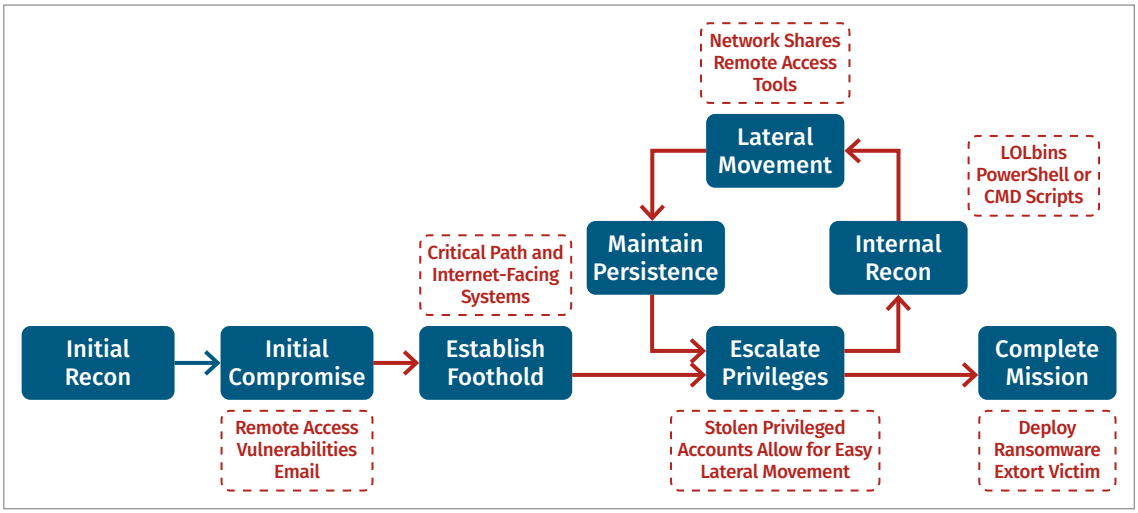

Figure 1. A Basic Ransomware Attack

As mentioned earlier, ransomware attacks are not entirely different from other intrusions. Thus, implementing endpoint detection and prevention capabilities should serve as a basic requirement for any organization. Note that this does not have to be a complex endpoint detection and response (EDR) capability. Even basic antivirus or OS-included malware defense can stop low-level attacks or those using legacy techniques.

Detection Tip #1: Endpoint detection, even if using OS-included antivirus, is a necessary first step to preventing low-level threats. Some ransomware attacks would fail to get past free-tier A/V defenses, earning the security team an advantage.

Admittedly, many ransomware threat actors are skilled at thwarting low-level defenses on *critical path systems*. Systems on this path are crucial for adversaries to establish a foothold, escalate privileges, or reach the rest of the domain via lateral movement. Let's analyze various steps within the attack life cycle to find opportunities for enhanced detections.

Critical Path Systems

Recognizing the key systems within your environment is paramount to any cyber defense strategy. Key systems are those that may contain sensitive data or allow an adversary to gain control of an environment. These often include, but are not limited to, domain controllers, database servers, and file servers. Ransomware threat actors also take particular interest in:

- Systems providing remote access to the organization, including Remote Desktop Protocol (RDP)-enabled systems and VPN providers (covered more later in this report)
- Backup servers and virtual machine hypervisors that can be used to restore the organization back to normal operations

Detection Tip #2: Identify critical systems within your environment with respect to a ransomware attack. Ensure that you are receiving telemetry—both endpoint *and* network—from these devices, that they are fully patched, and that they are clearly identified for the security team.

Given their nature, critical systems will likely have permanent network addresses and hostnames within the environment. Ensure that the security team has a list of these systems, and that list is updated as frequently as necessary. Going forward, these systems should have a priority in ongoing information security efforts, such as patching and auditing.

Initial Compromise and Foothold Establishment

Identifying an entry vector and gaining initial access into an environment is a step found in every cyberattack. Thus, we would assume that this is where defenders want to focus most of their efforts and resources. However, this is often where security teams run into the most difficulty. Identifying and defending an organization's entire attack surface could be a full-time job (and sometimes is!) for large organizations.

We cannot cover every single entry vector in this whitepaper. However, we can look to ransomware trends to help organizations determine where they can focus efforts and gain a significant advantage over typical ransomware attacks.

Detection Tip #3: Remote access tools have their place in enterprise environments, but they also can create a significant security gap for adversaries to easily abuse. The security team *must* know of remote access tools in the environment and how they are deployed and used.

Remote Access

One of the most common entry vectors for ransomware actors is the use of open (and often unsecured) remote access tools. This can include Windows-native protocols such as RDP or third-party applications such as LogMeIn or TeamViewer. And while remote access programs certainly have their place for legitimate usage within an organization, they are often not deployed securely and are easy to identify during an adversary's reconnaissance stage. Analysis from Unit42 in July 2021[1] found that in 50% of cases analyzed, RDP was the initial entry vector.

[1] "Diagnosing the Ransomware Deployment Protocol (RDP)," www.paloaltonetworks.com/blog/2021/07/diagnosing-the-ransomware-deployment-protocol

A security team *must* be aware of the remote access deployments within their environment. Endpoint monitoring and visibility can be useful in detecting these applications, and each should be either removed or documented with a legitimate, "must-have" business case. Admittedly, the recent COVID-19 pandemic has increased the *legitimate* need for remote access tools. Remote access capabilities cannot be deployed and simply "forgotten." They must be monitored and eliminated when no longer required. Network data is also an excellent source to identify remote access into the environment—especially if endpoint telemetry has been tampered with! When coupled, endpoint and network data provide a one-two punch that adversaries have a hard time working around.

To help navigate through the COVID-19 pandemic and deal with changing workforces, many organizations have turned to remote access tools. As operations continue and the organization determines the path forward, constantly reassess the need for remote access tools and disable them if no longer necessary.

Email

Email phishing campaigns are another popular entry vector used by ransomware threat actors. A clicked link or downloaded file gives an adversary an account and a foothold in one fell swoop, which can be used for further exploitation in the environment. Furthermore, adversaries can launch massive email campaigns that target multiple users simultaneously, increasing their chance of success.

Email threats can be combatted in two ways: user education and email security. User education will pay off tenfold for an organization, as educated users will help detect and defend against various types of phishing emails, not just ransomware.

Detection Tip #4: Adversaries love email phishing campaigns because a trusted user is the one clicking the link. User education is paramount in defending against this type of threat, coupled with email defenses, if available.

Known Vulnerabilities

Finally, recent trends also have shown ransomware threat actors quickly taking advantage of vulnerabilities in a widespread manner. Whereas security teams may have had some "breathing room" between the time a vulnerability was announced and an exploit was released, this timeframe is now often a matter of minutes or just a few hours. Realizing that they can automate and batch target *entire IP ranges*, adversaries have built infrastructure to allow for rapid exploitation and ransomware deployment.

However, the true problem here is the time it takes for organizations to patch vulnerable systems. For example, in early 2021, adversaries targeted unpatched Fortigate VPN servers to deploy "Cring" ransomware.[2] The adversaries achieved success at multiple facilities from a vulnerability that was patched in 2019. We're not here to shame any organization—patching often takes time and approval. However, when defending against ransomware attacks, it is best not to leave any known doors open.

Detection Tip #5: Threat actors love unpatched systems, especially if the respective vulnerabilities can be used to access an environment, run malicious code, or obtain credentials for subsequent access. Patching vulnerabilities is a key step to preventing widespread, automated ransomware deployments. Unpatched systems should be rotated into the aforementioned Critical Path Systems list until they are patched (if they do not have a permanent place on the list!).

[2] "Ransomware crooks are targeting vulnerable VPN devices in their attacks," www.zdnet.com/article/ransomware-crooks-are-targeting-vulnerable-vpn-devices-in-their-attacks

Escalate Privileges

In a correctly configured environment, a normal user account cannot be used to access critical servers or move laterally throughout an environment. Thus, adversaries seek to escalate privileges as quickly as possible to execute subsequent steps of an attack. Furthermore, privileged accounts are less likely to be changed or locked like user accounts, so threat actors will use them as a form of persistence as well.

Detection Tip #6: **The security team *must* audit domain and privileged accounts frequently. Adversaries gain an unnecessary advantage when accounts are given too many privileges or highly privileged accounts are used for domain operations. Privileged accounts should be monitored and have importance like the Critical Path Systems list.**

We cannot simply ban escalated privileges from an environment. Escalated accounts are necessary for many operations in a corporate environment. However, what can be limited is how "far and wide" these accounts are utilized. Users should not have privileges above what is necessary for their daily activities. Services should be run with least-privilege accounts, and network access should be granted to non-user accounts only when necessary.

Pro Tip: **If you have a Windows domain, there are multiple account protections built into Windows Server 2016 domain environments. These mechanisms can help protect accounts during remote access and service execution, and protect credential theft attacks against designated attacks. If applicable, leverage these mechanisms to wrap additional security around protected accounts.**

Even after correctly securing accounts and limiting privilege abuse, adversaries may still gain control of a privileged account. This creates a unique opportunity for security teams with account visibility and logging: A "stolen" account will quickly begin to display anomalous behavior. We would not expect a domain administrator account to be launching services, browsing the internet, or downloading files. These simple displays of "non-privileged" behavior allow for easy identification of anomalous behavior that must be investigated.

Detection Tip #7: **There are things accounts should and should not do. If a privileged or domain administrator account begins to display anomalous behavior or appears to be "acting as a user," the security team must investigate—ransomware or not.**

Internal Recon

With a foothold in the environment and escalated privileges acquired, adversaries must next work to understand the environment they are in. Often referred to as "internal reconnaissance," this is the stage at which an adversary will utilize built-in tools or custom scripts to discover additional systems or users within the environment. This activity may occur multiple times, each giving a new chance for detection.

Pro Tip: **Note that internal reconnaissance may occur multiple times as adversaries "feel out" the environment they are in. Anytime an adversary repeats behavior, it means security teams have increased chances of detection!**

The hallmark signs of internal reconnaissance are often easy to spot. Adversaries will use commands to discover systems, retrieve basic networking information, and enumerate user and domain groups. Figure 2 provides a snippet of discovery commands as observed from an incident that began with IcedID malware and eventually resulted in a ransomware infection.

IcedID initial Environment Discovery

Several discovery commands executed from IcedID after the initial execution:

```
ipconfig /all
cmd.exe /c chcp >&2
WMIC /Node:localhost /Namespace:\\root\SecurityCenter2 Path AntiVirusProduct Get * /Format:List
systeminfo
net config workstation
nltest /domain_trusts
nltest /domain_trusts /all_trusts
net view /all /domain
net view /all
net group "Domain Admins" /domain
```

Figure 2. Snippet of Discovery Commands Used in an Attack Involving IcedID Malware[3]

The commands present in Figure 2 are not inherently malicious. On their own, or a few in tandem, these commands could represent legitimate system administrator activity. However, this is where our previous knowledge of critical systems and privileged accounts comes into play. A complex, multisource detection can help teams determine the (1) source system and (2) user responsible for the account, and quickly determine if they are observing legitimate or malicious activity.

Detection Tip #8: **With endpoint visibility, execution of internal scanning and reconnaissance tools should be relatively straightforward. Associate these commands and event log entries (Windows 10, 2016 or later) with an account to be able to quickly determine if activity is legitimate or not and neutralize it immediately.**

Move Laterally

After a ransomware threat actor has "discovered" their victim environment, the attack begins to move much faster. With credentials in hand and knowledge of the environment, the next step is to begin moving laterally and deploying malware. Detecting lateral movement can seem daunting to some teams; however, it is one of the "noisiest" activities that an adversary can undertake.

Detection Tip #9: **Lateral movement detection exists at both the endpoint and network layers. Account authentication logs and host artifacts, coupled with network port access and data transfer, create a situation in which adversaries cannot remain silent. Integrate these two telemetry sources to quickly identify lateral movement activity.**

Lateral movement activity is recorded in abundance at both the endpoint and network levels. Let's examine artifacts provided by each source (see Table 1).

Table 1. Comparison of Endpoint and Network Artifacts

Endpoint	Network
• Authentication logs	• Directionality (internal <-> internal)
• Remote process creation	• Protocol and ports used
• Process thread instantiation	• Data transferred
• Protocol usage (SMB, RDP, etc.)	• One-to-many transfers
• File Share access	• Packet introspection to decode traffic (if applicable)
• Windows event logs	

[3] "IcedID to XingLocker Ransomware in 24 hours," https://thedfirreport.com

Despite a plethora of lateral movement artifacts, many security teams have issues correlating what is malicious lateral movement vs. expected or "normal." This is where your critical list of systems will come into play. Use that list to help with false-positive elimination and contact system administrators to help explain potentially legitimate activity.

Pro Tip: **Detection "languages" such as Sigma or Suricata are excellent examples of the security community working together to disrupt adversaries. When you implement new tools, be sure to inquire what types of detections the tool can use and ensure your team writes their own!**

Each telemetry source also comes with powerful detection capabilities that teams can easily take advantage of. Aside from log contents, powerful detection mechanisms such as Sigma or Suricata can be easily applied on top of Windows event logs and network traffic, to quickly fire alerts for lateral movement. Figure 3 provides an example Sigma rule to detect remote admin share creation over SMB.

```yaml
title: SMB Create Remote File Admin Share
id: b210394c-ba12-4f89-9117-44a2464b9511
description: Look for non-system accounts SMB accessing a file with write (0x2) access mask via administrative share (i.e C$).
status: experimental
date: 2020/08/06
author: Jose Rodriguez (@Cyb3rPandaH), OTR (Open Threat Research)
tags:
    - attack.lateral_movement
    - attack.t1021.002
references:
    - https://github.com/OTRF/ThreatHunter-Playbook/blob/master/playbooks/WIN-201012004336.yaml
    - https://securitydatasets.com/notebooks/small/windows/08_lateral_movement/SDWIN-200806015757.html?highlight=create file
logsource:
    product: windows
    service: security
detection:
    selection:
        EventID: 5145
        ShareName|endswith: 'C$'
        AccessMask: '0x2'
    filter:
        SubjectUserName|endswith: '$'
    condition: selection and not filter
falsepositives:
    - Unknown
level: high
```

Figure 3. Snippet of a Sigma Rule Used to Detect SMB Create Remote File Admin Share[4]

Network traffic is an equally powerful detection source, especially if the security team has packet inspection. The following Suricata is one of many sample rules that can be used to detect potential lateral movement (Figure 4).

```
alert smb any any -> $HOME_NET any (msg:"ET ATTACK_RESPONSE Possible Lateral Movement -
File Creation Request in Remote System32 Directory (T1105)"; flow:established,to_server;
content:"|05 00|"; offset:16; depth:2; content:"|00|W|00|i|00|n|00|d|00|o|00|w|00|s|00 5c
00|S|00|y|00|s|00|t|00|e|00|m|00|3|00|2|00|"; fast_pattern; classtype:attempted-user;
sid:2027267; rev:2; metadata:affected_product Windows_XP_Vista_7_8_10_Server_32_64_Bit,
created_at 2019_04_23, deployment Internal, former_category ATTACK_RESPONSE,
performance_impact Low, signature_severity Major, tag T1105, tag lateral_movement, tag
remote_file_copy, updated_at 2019_04_23;)
```

Figure 4. Sample Suricata Code to Detect Lateral Movement[5]

Detection Tip #10: **Your security team does not need to create its own detection platform to detect malicious activity. Leverage powerful, open source tools such as Sigma and Suricata to gain a collective advantage over adversaries using the same techniques in multiple environments.**

[4] "SigmaHQ/sigma," GitHub, https://sigma/win_smb_file_creation_admin_shares.yml at 8beb70e970b814d0ab60625206ea0d8a21a9bff8

[5] 2027267 < Main < EmergingThreats, https://doc.emergingthreats.net/bin/view/Main/2027267

Complete Mission

The final "digital" stage of a ransomware attack is trifold:

1. Remove access to backups and other data that could be used to restore encrypted data.
2. Encrypt critical data and key servers, leaving ransom notes behind.
3. If applicable, remove forensic artifacts to "hide" evidence of an attack.

At this stage of ransomware detection and response, if a team is waiting for a ransomware note to detect an incident, they have obviously waited too long. Our goal is to detect earlier in the attack, preventing an adversary from ever getting to this stage. However, even with the best-laid plans, it is possible an adversary might compromise a network and manage to encrypt files or systems. In such a situation, activating an incident response (IR) plan tailored toward ransomware will help get operations back to normal as quickly as possible.

Incident Response

When an organization detects a ransomware incident in its environment, it is paramount that it moves swiftly to deal with the threat. This, of course, does not imply that an organization would respond to a threat slowly. However, the clock is ticking once a ransomware note has been received. In the previous section, we identified areas and techniques for early ransomware detection. Next, we will examine how the six-step IR process can be impacted by a ransomware breach as well.

Figure 5 diagrams the six-step IR process, which serves as a guideline for responding to any incident.

Using this process as a template, we can design an IR approach suited toward a ransomware breach. For the better part of this whitepaper, we explored ways to identify (detect) and scope a ransomware incident, by using the attack life cycle to help guide ransomware-centric detections. When responding to an incident, ransomware may have more influence on the containment and eradication steps, which we will examine next.

Pro Tip: While detecting ransomware earlier than later is better, organizations can look to integrate and utilize ransomware threat intelligence to enable effective response and fine-tune detection capabilities. Even better, threat intelligence can help an organization prepare for a ransomware attack *before* an adversary even has a chance to set off alarms.

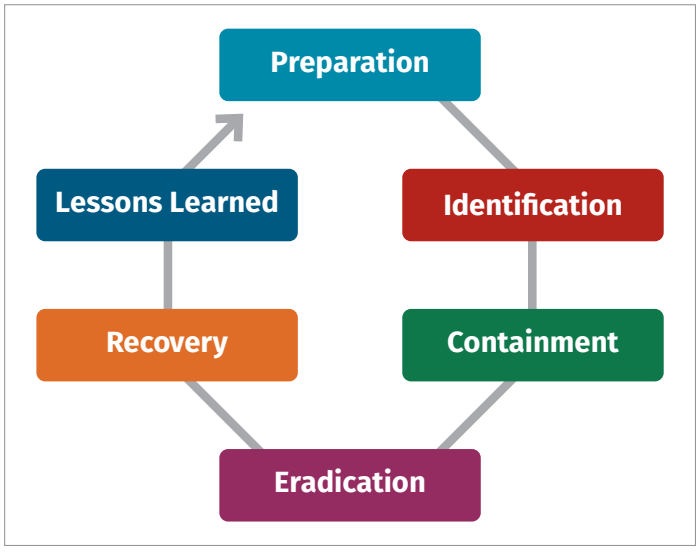

Figure 5. The SANS Six-Step IR Process

Pro Tip: The identification stage is where the security team has the advantage over the adversary. If the environment suffers data encryption and/or partial or full ransom, the advantage can quickly shift to the attacker. Use an IR process to ensure your team stays focused.

Containment

When activity has been detected, containing the incident is a critical step in any situation, regardless of adversary or objective. However, ransomware increases the criticality, as the adversary's endgame is harmful to the environment. Whereas containment may sometimes serve as a step to gather intelligence, in this case it is quite the opposite. As such, we want to adjust this step of the IR plan to include the following:

- Systems with detections should be cut off from the network to prevent further spread.
- Accounts that have been detected as malicious should be changed immediately.
- If ransomware is confirmed, segregate Critical Path Servers (from our earlier list) from the network as quickly as possible, if necessary.
- Utilize endpoint and network security products to block malicious processes and traffic immediately.

While some of these adjustments may be applicable in any type of incident, oftentimes organizations implement "red tape" or blockers that require extended periods of time to change an account or take a system offline. Unfortunately, that luxury goes away with a ransomware attack.

> **Pro Tip:** **Backup servers should be "stored" or hosted elsewhere, and network segmentation should be the default. We don't want adversaries getting access to the systems too easily!**

Eradication/Recovery

The eradication stage, during which a security team is removing an adversary and their tools from an environment, is just as critical as the recovery stage, during which the team is (hopefully) repairing wounds and is no longer under an active threat. Ransomware threat actors have been known to retarget organizations, especially those with unpatched systems or easy re-entry.

During these two stages, the team must use its detections and investigative results to clean up the environment and prevent the same incident from re-occurring. For example, if the entry vector was Remote Desktop Protocol, and the business has determined that remote access functionality is required, then the security team must retool around that business requirement. An adversary should never be able to compromise an environment via the same means twice—we learn from the adversary about how to make our networks stronger!

Conclusion

At the top of everyone's mind over the past few years—and keeping security teams up at night—ransomware attacks have risen to global infamy and show little sign of slowing down. Industry-agnostic and to the tune of billions of dollars,[6] far too often we are seeing organizations getting blindsided by a ransomware attack. Furthermore, the list of victim organizations is not limited to defenseless industries with no security team; global organizations with security teams have also fallen victim.

In this report, we defined how ransomware attacks differ from other types of cyberattacks. These key differences lead to opportunities to detect ransomware threat actors early in the attack life-cycle process, ensuring that it is harder for them to reach the final stage of encrypting files and locking up systems. Security teams that recognize these opportunities will be able to secure their networks and detect malicious activity easier.

Ransomware Detection Checklist

- ☑ **Endpoint defenses, even if free-tier A/V or OS-included, can stop low-level or legacy technique attacks from occurring.**
- ☑ **Identify critical systems within the environment, i.e., systems crucial to business operations but also common ransomware targets. Ensure these systems are fully patched and the security team has a real-time list to augment detections.**
- ☑ **Identify and remove all *unnecessary* remote access capabilities. If they are deemed necessary, document the use case and add this system(s) to the critical systems list.**

[6] "Suspected Ransomware Payments Nearly Doubled This Year, Treasury Says," www.wsj.com/articles/suspected-ransomware-payments-for-first-half-of-2021-total-590-million-11634308503

About the Author

Matt Bromiley is a SANS digital forensics and incident response instructor, teaching FOR508: Advanced Incident Response, Threat Hunting, and Digital Forensics and FOR572: Advanced Network Forensics: Threat Hunting, Analysis, and Incident Response. He is a principal consultant at a global incident response and forensic analysis company, combining his experience in digital forensics, log analytics, and incident response and management. His skills include disk, database, memory and network forensics; incident management; threat intelligence; and network security monitoring. Matt has worked with organizations of all shapes and sizes, from multinational conglomerates to small, regional shops. He is passionate about learning, teaching and working on open source tools.

Sponsor

SANS would like to thank this paper's sponsor:

ANOMALI®