SOPHOS

Managed Threat Response

Welcome Guide

Version: 1.4

Welcome

"It's dangerous to go alone! Take this." So read the first lines of dialogue in the 1986 Nintendo classic The Legend of Zelda when your character, a young adventurer, is bequeathed a sword to fend off the forces of evil he will face on his impending quest. It's a reference that may not be as familiar to you as it is to me, but the sentiment should still hold true as we begin our partnership: You are no longer alone. From this point forward, our team has your back.

As your security partner, our team—the Sophos Managed Threat Response (MTR) team—will be working alongside you as a true extension of your organization. Our goal is to help you not only achieve your security goals but surpass them.

After reading this guide, you will have a clear understanding of how the MTR Operations team detects, investigates, and takes action to neutralize the most sophisticated cyber threats. While we've done our best to clearly document each piece in the pages that follow, we are always happy to answer any lingering questions you may have.

Email: mtr-ops@sophos.com

Who We Are and Why We're Here

We refer to our team as "MTR Operations," or "MTR Ops" for short. But you will get to know many of our analysts by name. Here's a little bit about us:

We are a team of security professionals: analysts, engineers, ethical hackers, specialists, and inventors. Our backgrounds are comprised of armed forces, law enforcement, intelligence, and public and private enterprise.

Our team works 24/7 to hunt and neutralize threats that cannot be detected or neutralized by technology solutions alone.

We're not merely watching over your IT environment – we're actively defending your business. Our objective isn't to clean up the damage following an attack; it's to stop attacks before they even start.

We're here to drive continuous improvement. In addition to neutralizing threats, we will provide detailed recommendations for improving your overall security posture.

Getting Set Up

This is the part in most "Getting Started" guides where we'd have a long list of things you need to download and install, but we wanted to keep things as quick and easy as possible. In fact, if you're reading this guide, our MTR Ops team is already actively defending every MTR-enabled device in your environment. All that we ask at this stage is that you check a few things:

Be on the lookout for your Security Health Check: One of the first things we do is evaluate your current Intercept X Advanced with EDR settings and provide a readout of recommended configuration changes to optimize the MTR service.

Make sure MTR is deployed everywhere you need it: Please continue deploying MTR to your devices throughout your environment. The more visibility we have, the better. If you are unsure how to deploy the MTR license in Central, please refer to the Onboarding Guide here.

Set your escalation contacts and Response Mode: You are in complete control of how potential threats are escalated, what response actions (if any) you want us to take, and who should be included in those communications. Please review your escalation contacts and Response Mode in Central to ensure things are set up how you want them. And remember, you can change these preferences at any time. If you are unsure how to set your MTR preferences in Central, please refer to the Onboarding Guide here.

MTR Advanced Customers

If you're an MTR Advanced customer, we'll also reach out to schedule a 30-minute orientation call. This call gives us an opportunity to review the results of your Security Health Check (as noted above), establish any next steps, and answer any questions you may have.

If you're a Standard MTR customer, don't worry. You can still reach our team at any time by emailing mtr-ops@sophos.com, and if you ever believe that you're experiencing an Active Threat, you can call us directly in North America at 888-201-7672, or globally through your Sophos support number in your local region.

Working With the MTR Ops Team

While other managed detection and response (MDR) services simply send notifications for potential threats or suspicious events – leaving it up to you to verify and respond to threats – Sophos MTR arms you with an elite, 24/7 team of threat hunters and response experts who take targeted actions on your behalf to neutralize even the most sophisticated threats. The work our MTR Ops team includes:

- Proactively hunting for and validating potential threats
- Using all available information to determine the scope and severity of threats
- Applying the appropriate business context for valid threats
- Providing actionable advice for addressing the root cause of recurring threats
- Taking actions on your behalf to disrupt, contain, and neutralize threats

Below, you'll find more information on the different ways to work with our team. This will help you determine the best way for our team to work alongside yours.

Response Modes: What They Are and How They Work

We know that teams responsible for managing IT security vary greatly in terms of their size, capabilities, and needs, so we made Sophos MTR customizable with different ways to engage our MTR Ops team based on the unique and constantly evolving needs of your team and broader organization.

We call these customizable options "Response Modes," and there are three from which to choose regardless of whether you selected MTR Standard or Advanced. Here's a description of each Response Mode (complete with aviation similes):

Notify: Select this mode if you only want to receive notifications of observed activities that include recommendations to help you prioritize and manage response efforts. Selecting Notify means you don't want the MTR Ops team to take any response actions on your behalf. (We're like an air traffic control tower serving as an extra set of eyes and alerting you, the captain, to potentially important events.)

Collaborate: Select this mode if, in addition to notifications and corresponding recommendations, you also want the MTR Ops team to perform some (but not all) response actions on your behalf. Selecting Collaborate gives you the option to have some response actions performed by the MTR Ops team and others to be performed by your team or another partner (e.g. an IT managed service provider). Please note that in this mode, the MTR Ops team must receive written authorization before performing response actions. (We're like your co-pilot and you're the captain.)

Authorize: Select this mode if you want the MTR Ops team to proactively manage all containment and neutralization actions on your behalf and inform you of the action(s) taken. Selecting Authorize means you want us to handle as much workload as possible and only escalate things that require specific actions from you or your team. ("We are the captain now.")

Response Modes in Action

Now that you have a better handle on what Response Modes are, let's explore how they work using a practical example. Let's say the MTR Ops team identified an active ransomware attack in your environment. Here's how we would engage with you based on each Response Mode:

Notify: The MTR Ops team calls each of your listed escalation contacts until at least one of those contacts answers. If all escalation contacts are unreachable by phone, the MTR Ops team will contact you via email. Whether the MTR Ops team reaches your escalation contacts by phone or email, they will provide detailed information about the detection and recommended response actions that need to be performed by your team to neutralize the ransomware attack.

Collaborate: The MTR Ops team follows the same phone and email notification process as noted above. But perhaps you're unable to perform the required response actions because you're on holiday with no laptop or you're at home caring for a sick kid and you need us to take those actions for you. In situations like these, all we need is your permission to perform those actions and we manage things from there.

Authorize: The MTR Ops team rapidly executes response actions to neutralize the ransomware attack. Once the threat is neutralized, the MTR Ops team contacts you via phone and/or email and provides detailed information on the threat and the action(s) taken to neutralize it.

Cases, Threat Hunts, and Other Important Terms

Our main objective is to identify and investigate potentially malicious activity in your environment. We do this via two methods: 1) Investigation of MTR detections, and 2) analyst-led threat hunts.

If the MTR Ops team concludes that a detection or activity requires further evaluation, a case is created, and our operators conduct a full investigation (this process is outlined in more detail on page 6). Only cases that require customer input or action will be escalated via email or over the phone. Here are some succinct definitions of these terms:

- **Detections:** Technology-generated indicators of potential threats
- **Threat hunts:** Analyst-led investigations to identify attacks that cannot be detected or stopped by existing tools
- **Cases:** Detections or activities identified through threat hunts that require an analyst investigation
- **Escalations**: Cases that require customer input or action that cannot be performed by MTR Ops
- **Incidents:** Confirmed malicious activities that require immediate response

The Investigative Framework

When we say "analyst investigation," what do we mean specifically? Our proprietary Investigative Framework provides structure to guide analysts while investigating Cases. The framework enables our MTR Ops team to construct an attack narrative which aids them in concluding whether malicious activity is present within a customer environment (provided data coverage and data quality are at their maximum). The process follows the iterative nature of the OODA Loop (Observe, Orient, Decide, Act).

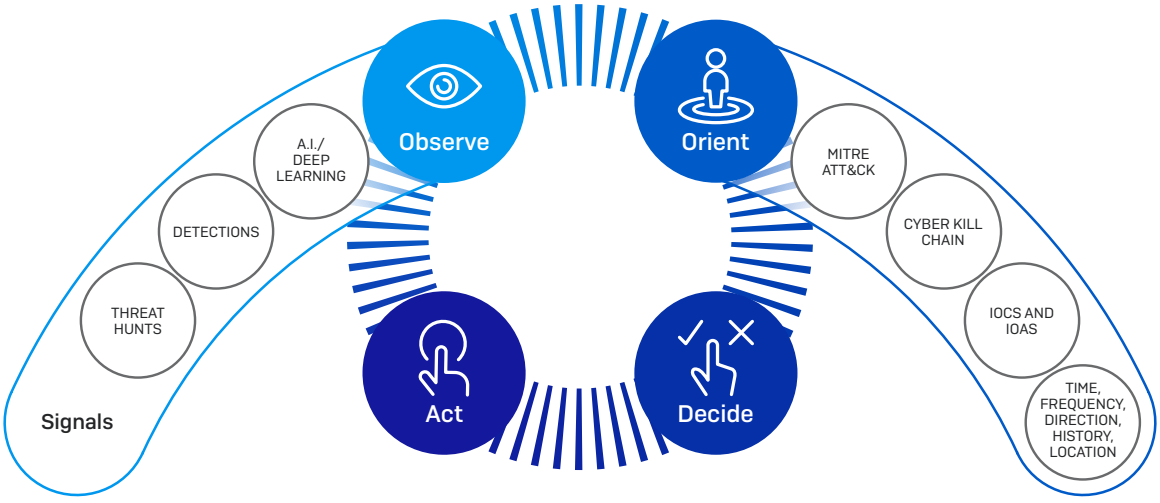

1. The objective of the Observe phase is to select key points of data that help establish a logical narrative of activity that is occurring on a customer device or within a customer's environment. Each chosen data point is a logical observation that has the potential to indicate malicious activity.
2. During the Orient phase, analysts validate observables which can create indicators. Validation is performed by applying the data points to the MITRE ATT&CK Matrix, the Cyber Kill Chain, and an analyst's tribal knowledge. The result is a logical narrative of activity. If enough observables are validated into indicators the activity will create an attack narrative.
3. During the Decide phase, the analyst will iterate through his or her previously compiled data points to determine what is required in the Act phase. If the analyst does not come to a confident decision in the identification of malicious activity, he or she will progress through the act phase and start the OODA Loop over again.
4. During the Act phase, the analyst will act based upon the conclusion of the investigation. If the analyst did not accrue enough information to adequately answer the questions in the Decide phase, then a new OODA Loop is initiated with the previously gathered indicators serving as the new basis for the investigation. If enough information has been gathered, enabling the analyst to answer the key questions in the Decide phase, then the analyst will move forward and take the necessary actions.

Threat Hunting

"Threat hunting" is one of the latest cybersecurity terms to earn platinum-level buzzword status alongside "AI," "machine learning," and countless others. Everybody is talking about it, few actually do it, and many don't understand what it really means (and at this point are too afraid to ask). So, let's start with our definition of threat hunting.

> **Threat Hunting**
> A human-led investigation of causal and adjacent events (weak signals) to discover new Indicators of Attack (IoA) and Indicators of Compromise (IoC) that cannot be detected or stopped by existing tools.

To put it another way, a threat hunt is when an analyst conducts an investigation to detect attacks that tools don't issue an alert about and which only a human can find. And while threat hunting is often described in absolute terms, there are actually three different categories of threat hunts: automated, lead-driven, and lead-less.

Automated: This type of threat hunt uses automation and/or machine learning to surface potentially malicious activity that may require further investigation by human analysts. While this is what many service providers are referencing when they say they do "managed threat hunting," this is what is programmatically handled by Intercept X Advanced with EDR.

Lead-driven: This type of threat hunt involves a manual (human-led) identification and investigation of events and activities (leads) that do not generate an alert but could be indicative of new attacker behavior. The MTR Ops team performs lead-driven hunts for all Standard and Advanced tier MTR customers.

Lead-less: This type of threat hunt combines threat intelligence, data science, and knowledge of attacker behavior with what's known about the customer's environment (e.g. company profile, high-value assets, high risk users, etc.) to anticipate new attacker behaviors and validate detection and response capabilities. This category is sometimes called "methodology hunting," and very few service providers have the ability to perform this kind of threat hunt. The MTR Ops team performs lead-less hunts for all Advanced tier MTR customers.

Performing Threat Hunts

As soon as Sophos MTR is enabled on a device (whether it's an endpoint or server), data from that device is continuously collected and analyzed in the MTR platform. Depending on availability, the following data can be used as the foundation from which the MTR Ops team performs threat hunts:

Device data

- Process execution: Contains information on processes run on specific hosts
- Registry data: Contains data related to registry objects, including key and value metadata
- File artifacts: Information on stored files and artifacts kept on a local host

Network data

- Session data: Information regarding network connections between hosts
- DNS logs: Data related to DNS resolution
- IDS data*: Traffic information from all devices on the network
- Firewall logs*: Connection data on the edge of the network (allowed and blocked)

* Available to MTR Advanced customers with Sophos XG Firewall v18

Security data

- Security alerts: Automated alerts from security tools (e.g. Intercept X Advanced with EDR, Firewall XG IPS and ATP rules, etc.)
- Threat intelligence: Data that contains indicators and known tactics, techniques, and procedures (TTPs) used by attackers
- CVE vulnerability disclosures

Equipped with this information, the MTR Ops team takes the following steps to hunt for new Indicators of Attack (IoAs) and Indicators of Compromise (IoCs). Keep in mind that threat hunting requires a great deal of critical and creative thinking. People tend to think of threat hunting (or cybersecurity in general) as a purely analytical discipline. But the reality is that security operations is equal parts art and science. Each situation our MTR Ops team encounters can be viewed from multiple perspectives, and all of them can be right. This notion, which is infused into our process, is part of the reason why our team is able to conduct threat hunts with the highest level of efficacy.

1. A threat hunt case is created to track and document work. Endpoint, network, and security data is combined with threat intelligence and refined into actionable information for the MTR Ops team to perform the hunt.
2. Analysts create a hypothesis based upon the contextual and actionable information available to them and begin to identify and collect the data needed to investigate potential adversarial actions.
3. The Investigative Framework (outlined earlier) is then used to hunt for adversarial activity within the collective data.
4. If malicious and/or adversarial activity is identified, an incident is created and categorized.
5. If the threat hunt is successful (i.e. a new IoC or IoA is identified), findings will be used to create automated detections for future use. Analysts should never have to perform the same threat hunt twice.

Activity Reporting

The MTR Ops team constantly reviews alerts, investigates anomalous activity, and responds to confirmed threats with speed and precision based upon your selected preferences (i.e. Response Mode). In addition to providing comprehensive assessments of attacker activity and corresponding response actions as they occur, the MTR Ops team also provides monthly and weekly activity reports summarizing case activities and providing the context needed to understand threats, assess organizational risk, and prioritize actions.

Just as our investigation and threat hunting methodologies are aligned to the adversarial tactics and techniques described in the MITRE ATT&CK Matrix, so too are our monthly activity reports. You will find a guide on the following pages for understanding and leveraging your monthly and weekly activity reports.

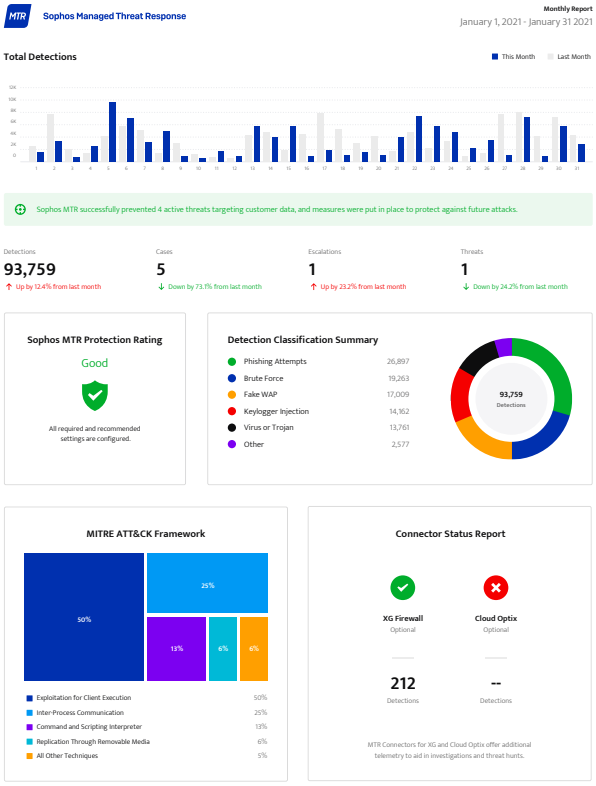

Not familiar with MITRE ATT&CK?

MITRE ATT&CK is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations of cyberattacks. They're displayed in matrices that are arranged by attack stages, from initial system access to data theft or machine control.

ATT&CK stands for adversarial tactics, techniques, and common knowledge. Tactics are the why of an attack technique. Techniques represent how an adversary achieves a tactical objective by performing an action. Common knowledge is the documented use of tactics and techniques by adversaries.

Understanding Your Monthly Report

Total Detections

Are the total number of threat detections in your environment increasing, decreasing, or staying the same? This section provides insight into total alert volume observed throughout the course of a month and helps the MTR Ops team identify inflection points in potential adversary activity.

MTR detections are broken down by day and compared with the daily total from the previous month. The highlighted green text indicates the number of active threats that were prevented over the course of the month. The report also provides the total number of detections, cases, escalations, and threats detected with comparisons to the previous month.

The MTR team is constantly improving our detection capabilities which could naturally cause fluctuations in the volume of detections seen in a report. These adjustments could be for tuning out detections that have provided limited value in identifying threats or adding to our scope and visibility to identify new and emergent threats.

Definitions:

- Detections: Any indicator of suspicious, but not necessarily malicious, activity is classified as a detection. In most cases, these data points are purely informational and do not result in the creation of a case on their own. Detections often include items such as command executions, open network sockets, authentication events, running applications, and events generated by Sophos products. However, detections are not confirmed as suspicious or malicious until further analysis is completed. Detections that do not result in case creation are primarily used to supplement new or ongoing cases, to provide context during investigations, and to aid in the creation of new detections.
- Cases: Human-led investigations by MTR operators to determine if a detection is an indicator of an attack or compromise.
- Escalations: Cases that are determined to be potential incidents and may require customer attention.
- Threats: Confirmed indicators of attack or compromise observed within a customer network.

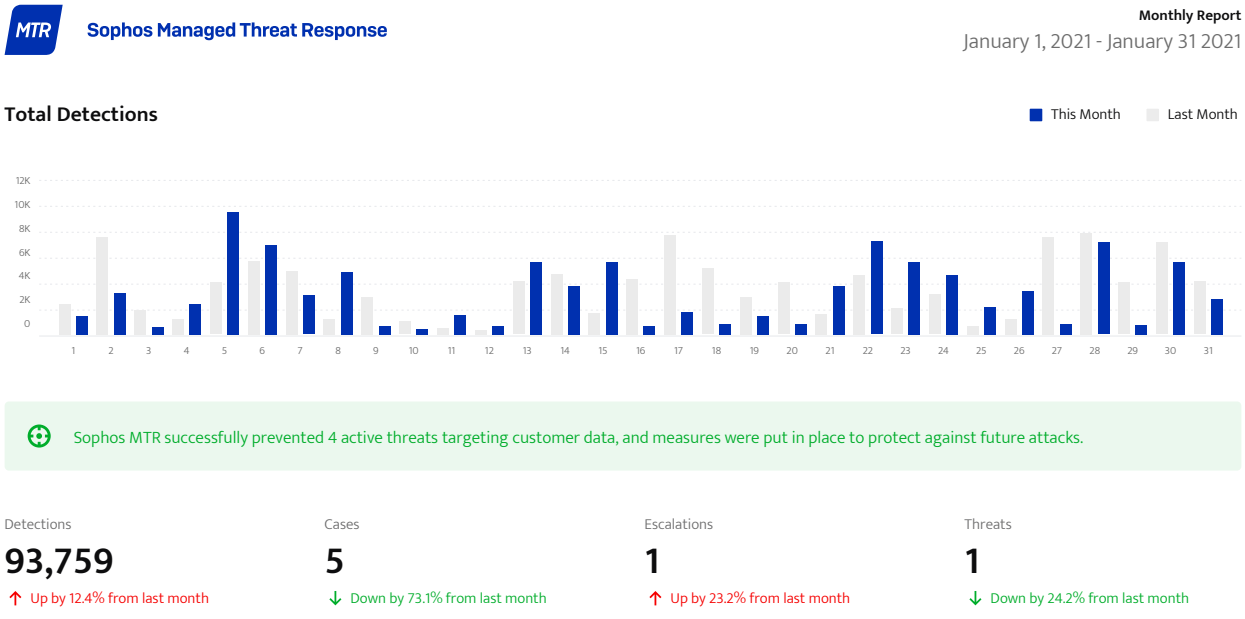

Protection Rating

Your Sophos MTR Protection Rating is an aggregate analysis of the security posture improvement recommendations that have been implemented versus those that have not been implemented.

These Health Check recommendations can include things like enabling anti-exploitation features to protect against credential theft protection or privilege escalation or enabling malicious traffic detection to hinder communication to command and control servers. Health Checks serve to proactively improve your security posture and remedy weaknesses that can adversely affect your security capabilities.

Your Overall Protection Rating will be either Green (you are in full compliance with recommended best practices), Yellow (some of your configurations increase risk, but they will not affect the efficacy of the MTR service), or Red (you have several high-risk configurations that prohibit the MTR service from functioning at an optimal level).

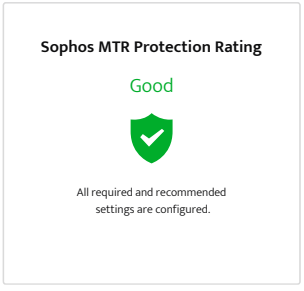

Detection Classification Summary

MTR detections are classified into high-level categories to aid in understanding the overall types of detections observed in your network. Examples include common attack tools, PowerShell execution, and persistence. As with all detections, they are not inherently indicative of suspicious or malicious activity and could be related to benign data that was collected.

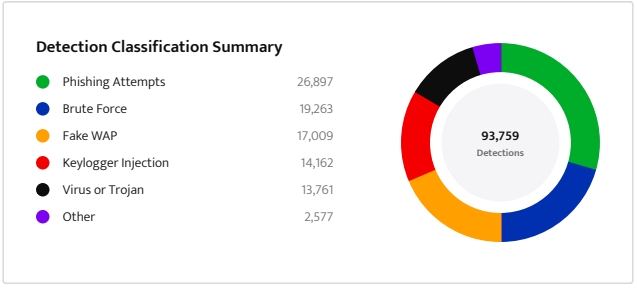

MITRE ATT&CK Framework

MTR detections are mapped to specific techniques in the MITRE ATT&CK framework, a widely used knowledge base of adversary behaviors based on real-world observations. You will see the breakdown of detections, by percentage, in this section of the monthly report.

As with all detections, these are not necessarily malicious and benign behavior may align to adversarial tactics and techniques. It is also important to note that the total number of MTR Cases may not be equal to the total number of adversarial tactics observed. Multiple adversarial tactics can be observed in one MTR Case, resulting in the number of tactics being greater than the total number of MTR cases. Conversely, MTR Cases may be created that are not associated with adversarial (health check cases, for example), resulting in the total number of MTR Cases being greater than the total number of adversarial tactics.

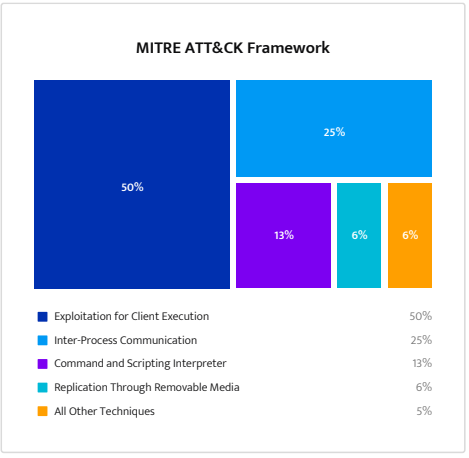

Connector Status Report

MTR Connectors were designed to ensure MTR operators have the most crucial data at their fingertips ensuring attackers have fewer places to hide. The Connector Status Report indicates whether the MTR Connector is active or disabled and shows the number of detections that were generated due from the enhanced telemetry.

Sophos MTR Advanced customers can fully deploy Sophos XG Firewall across their environment or deploy XG Firewall in tap mode while utilizing a non-Sophos firewall. Customers must manage their XG Firewalls in Sophos Central and use XG Central Firewall Reporting. The Sophos Firewall MTR Connector generates MTR detections from

the following network security events: ATP (Command & Control), IPS, Sophos AV (email, web, FTP), and Sophos Sandstorm (sandbox).

The Sophos Cloud Optix MTR Connector provides Sophos MTR operators with the visibility needed to quickly identify critical cloud security events used in breach attempts across Amazon Web Services, Microsoft Azure, and Google Cloud Platform environments. Events from Sophos Cloud Optix generate MTR detections, including anomalous IAM user login activity, outbound network traffic connections, and other high-risk activity. Additional threat detections can be added via integration with the Amazon GuardDuty service, which analyzes CloudTrail, DNS and VPC flow logs.

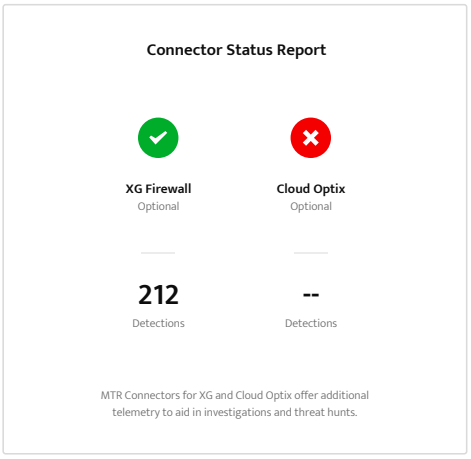

Understanding Your Weekly Report

In addition to some of the sections described in the monthly report, the weekly report highlights some immediate action items.

Action Required Status

An alert will be displayed at the top of the report when additional information or action is needed in order for the MTR analysts to continue an investigation. If action is required, please contact the MTR team as soon as possible.

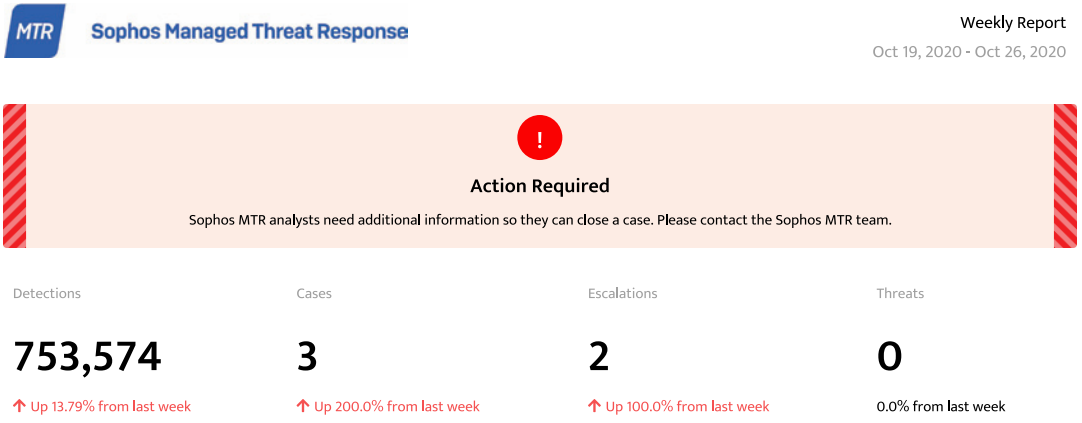

Case Activity

A detailed account of weekly cases will be displayed. Status includes “in progress” indicating that the case is still be investigated, “resolved” indicating that the case is now closed, and “waiting on client” indicating that additional information is being requested before the case can be resolved.

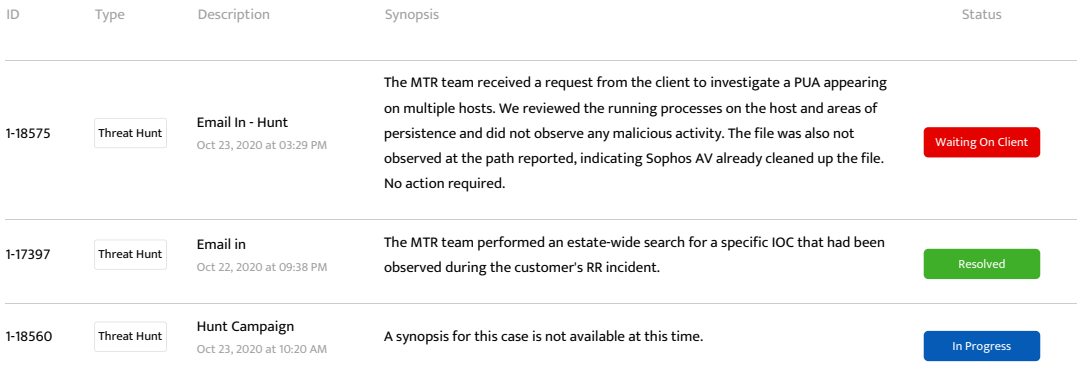

MITRE ATT&CK Framework Definitions

Initial access: The adversary is trying to get into your network

Consists of techniques that use various entry vectors to gain their initial foothold within a network. Techniques used to gain a foothold include targeted spear phishing and exploiting weaknesses on public-facing web servers. Footholds gained through initial access may allow for continued access, like valid accounts and use of external remote services, or may be limited-use due to changing passwords.

Execution: The adversary is trying to run malicious code

Consists of techniques that result in adversary-controlled code running on a local or remote system. Techniques that run malicious code are often paired with techniques from all other tactics to achieve broader goals, like exploring a network or stealing data. For example, an adversary might use a remote access tool to run a PowerShell script that does Remote System Discovery.

Persistence: The adversary is trying to maintain their foothold

Consists of techniques that adversaries use to keep access to systems across restarts, changed credentials, and other interruptions that could cut off their access. Techniques used for persistence include any access, action, or configuration changes that let them maintain their foothold on systems, such as replacing or hijacking legitimate code or adding startup code.

Privilege escalation: The adversary is trying to gain higher-level permissions

Consists of techniques that adversaries use to gain higher-level permissions on a system or network. Adversaries can often enter and explore a network with unprivileged access but require elevated permissions to follow through on their objectives. Common approaches are to take advantage of system weaknesses, misconfigurations, and vulnerabilities. Examples of elevated access include: SYSTEM/root level; local administrator; user account with admin-like access; user accounts with access to specific system or perform specific function. These techniques often overlap with persistence techniques, as OS features that let an adversary persist can execute in an elevated context.

Defense Evasion: The adversary is trying to avoid being detected

Consists of techniques that adversaries use to avoid detection throughout their compromise. Techniques used for defense evasion include uninstalling/disabling security software or obfuscating/encrypting data and scripts. Adversaries also leverage and abuse trusted processes to hide and masquerade their malware. Techniques of other tactics are cross-listed here when those techniques include the added benefit of subverting defenses.

Credential Access: The adversary is trying to steal account names and passwords

Consists of techniques for stealing credentials like account names and passwords. Techniques used to get credentials include keylogging or credential dumping. Using legitimate credentials can give adversaries access to systems, make them harder to detect, and provide the opportunity to create more accounts to help achieve their goals.

Discovery: The adversary is trying to figure out your environment

Consists of techniques an adversary may use to gain knowledge about the system and internal network. These techniques help adversaries observe the environment and orient themselves before deciding how to act. They also allow adversaries to explore what they can control and what's around their entry point in order to discover how it could benefit their current objective. Native operating system tools are often used toward this post-compromise information-gathering objective.

Lateral Movement: The adversary is trying to move through your environment

Consists of techniques that adversaries use to enter and control remote systems on a network. Following through on their primary objective often requires exploring the network to find their target and subsequently gaining access to it. Reaching their objective often involves pivoting through multiple systems and accounts. Adversaries might install their own remote access tools to accomplish lateral movement or use legitimate credentials with native network and operating system tools, which may be stealthier.

Collection: The adversary is trying to gather data of interest to their goal

Consists of techniques adversaries may use to gather information and sources of information that are relevant to following through on their objectives. Frequently, the next goal after collecting data is to steal (exfiltrate) the data. Common target sources include various drive types, browsers, audio, video, and email. Common collection methods include capturing screenshots and keyboard input.

Command and Control: The adversary is trying to communicate with compromised systems to control them

Consists of techniques adversaries may use to communicate with systems under their control within a victim network. Adversaries commonly attempt to mimic normal, expected traffic to avoid detection. There are many ways an adversary can establish command and control with various levels of stealth depending on the victim's network structure and defenses.

Exfiltration: The adversary is trying to steal data

Consists of techniques adversaries may use to steal data from your network. Once they've collected data, adversaries often package it to avoid detection while removing it. This can include compression and encryption. Techniques for getting data out of a target network typically include transferring it over their command and control channel or an alternate channel and may also include putting size limits on the transmission.

Impact: The adversary is trying to manipulate, interrupt, or destroy your systems and data

Consists of techniques adversaries use to disrupt availability or compromise integrity by manipulating business and operational processes. Techniques used for impact can include destroying or tampering with data. In some cases, business processes can look fine, but may have been altered to benefit the adversaries' goals. These techniques might be used by adversaries to follow through on their end goal or to provide cover for a confidentiality breach.

Need Additional Support?

We are serious about serving as a true extension of your team and doing so in a way that best supports and extends your internal capabilities. If you ever have questions about the MTR service or an MTR Case that has been escalated to you, you can always reach our 24/7 MTR Ops team by emailing mtr-ops@sophos.com. And if you ever believe you're experiencing an active threat, you can call us. All our regional numbers are found on sophos.com/support at the bottom of the page.

For help with other questions or addressing technical support issues, we've created a guide to help you get the answers you need quickly and easily. Please note that MTR Advanced customers also have direct, 24/7 call-in access to our Security Operations Center (SOC), regardless of whether they are experiencing an active threat.

	MTR ADVANCED	MTR STANDARD	SOPHOS SUPPORT
You have questions related to an MTR case that's been escalated to you or to one of your escalation contacts	✓	✓	
You believe you're experiencing an active threat	✓	✓	
You have questions about your monthly activity report	✓	✓	
You want to schedule an Ops review with the MTR team	✓		
You have questions about a recommended change to your configurations or architecture	✓		
You or someone in your organization has a general question related to your ability to detect and respond to threats	✓		
You need help resolving a product issue (not MTR)			✓
You're experiencing endpoint or server resource degradation			✓
You have Sophos product configuration changes related to business continuity			✓

Feedback

While we've done our best to clearly document the key elements of the Sophos MTR service, we are always happy to answer any lingering questions you may have. So please don't hesitate to reach out. We look forward to working alongside you and defending your business.

Contacting MTR Ops

Email us at mtr-ops@sophos.com or find your regional telephone support number at sophos.com/support

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

North American 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

21-02-24 (DD)

SOPHOS