SYNCHROWORKS CONSULTING

Solution Brief

Introduction

Sophos Intercept X is the world's most comprehensive endpoint protection solution. Built to stop the widest range of attacks, Intercept X has been proven to prevent even the most advanced ransomware and malware by leveraging a unique combination of next-generation techniques. This includes the ability to detect never-before-seen malware with deep learning, stop ransomware with Sophos anti-ransomware technology, and deny attacker tools with signatureless exploit prevention. Intercept X also includes root cause analysis to provide insight into threats, and instant malware removal to ensure no attack remnants remain.

Intercept X

Deep learning malware detection – Executable files are checked pre-execution, powered by machine learning.

Exploit Prevention – Detects and stops over 25 exploit methods used to compromise vulnerable applications (see our whitepaper on Exploits Explained).

CryptoGuard – Detect and roll back malicious file encryption activity caused by ransomware.

Application lockdown – Preventing malicious behaviors of applications, like a weaponized Office document that installs another application and runs it.

Credential theft protection – Prevents dumping of credentials from Windows credential store on disk, registry, from memory. This technique is seen in attacks that use tools such as Mimikatz.

Process Protection – Prevents use of techniques such as code cave and AtomBombing often used by adversaries looking to take advantage of the presence of legitimate applications.

Registry Protections – Protects sensitive areas of the registry often used by adversaries to manipulate application and system behavior.

Active adversary mitigations – A group of technical mitigations designed to detect and disrupt an attack that has established a presence on the host.

Safe Browsing – Monitoring a web browser's crypt, presentation, and network interfaces to detect man-in-the-browser attacks that are common in many banking trojans.

Root cause analysis – Provides an explanation of what happened when malicious activity is detected.

Sophos Clean – Provides a robust malware removal capability that will restore tampered Windows OS files and registries.

Compatibility with other vendor antivirus solutions – Sophos Intercept X is designed to work alongside competitive antivirus products, while Intercept X Advanced is designed to be deployed as a single, integrated agent that replaces your existing endpoint protection.

Sophos Intercept X Advanced combines the modern techniques of Intercept X with the foundational techniques of Sophos Central Endpoint Protection into a single product, and a single agent.

Synchronized Security – Collaboration with other Sophos Synchronized Security-enabled products to share contextual threat information and to respond automatically to detected threats.

Some common questions

What is the performance impact?

Intercept X, when deployed with all features enabled, consumes <1% CPU utilization on a typical system. This can spike when malicious activity is detected and files are being restored, a root cause analysis is being performed, and Sophos Clean is triggered to remove the attacking software. In those scenarios, CPU usage can spike to 1 core for several seconds.

What is the memory utilization?

The full Intercept X product consumes approximately 150MB of runtime memory.

Deep Learning Malware Detection

With the new deep learning model, we are able to perform a signatureless pre-execution evaluation of any executable file and determine if it is malware, potentially unwanted software, or a legitimate application.

At Sophos we've taken a unique approach to our security machine learning capabilities: we've invested heavily in deep neural network technology over more prevalent methods that, while still dominant in the security industry, are being rapidly abandoned by the machine learning computer science community.

How does Intercept X detect malicious executable files?

Instead of performing a signature and heuristic scan as traditional antivirus does, deep neural networks are able to extract millions of features from a file and determine if it is malicious before the program executes. The deep learning model learns what to look for in the code, how adversaries evade detection, how they build their software, and how the software plans to deploy and run. This information is evaluated by a multi-stage deep learning algorithm to determine how similar the software is to malware or potentially unwanted software. Depending on the score it is classified as malicious, potentially unwanted, or legitimate. It does all of this in about 20 milliseconds with a model that is under 20MB in size.

What happens when an attack is detected?

When malware is detected by the deep learning model, Intercept X will check if the file is on a suppression list. We talk about false positive suppression below, but for now, know that the suppression list allows us to run a very aggressive model to detect malware and still maintain an extremely low false positive detection rate.

Software detected as malicious will be put into quarantine and a root cause analysis will be triggered. If the detection was in error an administrator can release the sample by simply adding it to their local allowed application list.

The endpoint will be in a green security health state as the malware was prevented from executing.

What should an admin do?

The attack was detected prior to execution, but the admin may want to check the RCA report to determine how it reached the device so they can take actions prevent further infection attempts.

If the administrator determines the detection was in error, they can add the application to the allowed application list for their site directly from the detection event. This will automatically restore the application to where it was detected on all affected devices and will suppress future detections based on the file hash, signing certificate, or file path and name.

False Positive Suppression

A new quarantine has been created to hold convicted malware. When a detection of malicious activity occurs, Sophos Clean will be told to wipe the file and any of its associated registry entries, links, and files. The information is placed in quarantine and can be released by the administrator directly from the detection event in Sophos Central.

Releasing a detected malware or PUA file will add it to a site-wide Allowed Application list and restore the file on the endpoints affected. When adding a file to the allowed application list, the administrator can select the file hash identity, signing certificate, or file name and path. In the future, detections of this file will be suppressed and it will execute as designed.

In addition to this customer-specific false positive suppression list, Sophos maintains a global suppress capability. The Sophos false positive suppression list? is automatically checked when Live Protection is enabled, and Sophos will release small data updates to the endpoint when it is connected to the network. The global false positive suppression capability enables the deep learning malware and potentially unwanted application detection model to be extremely aggressive in detecting malware. This provides us the ability to have a forward-leaning detection model that has extremely low false positive detection.

The other huge advantage of providing robust false positive suppression is customers can deploy and start taking advantage of machine learning without having to go through weeks of tuning and configuration as many other vendors require. Instead, the model works on day one of deployment.

Exploit Prevention

Unlike other endpoint protection products, Intercept X addresses exploit detection with a comprehensive approach. While almost all other vendors can claim some degree of exploit prevention, they are often only targeting a small subset of the exploit methods available to adversaries. Sophos comprehensive exploit prevention addresses all exploit methods available. With over 20 exploit detection and prevention techniques included, Intercept X provides the most robust exploit prevention product on the market today. See our Exploits Explained whitepaper for a description of each technique and a vendor comparison chart.

How does Intercept X prevent vulnerabilities being?

Intercept X monitors classes of applications at the kernel level. This injection into the process allows close and continuous monitoring of activity in the process, including memory access, disk, network access, DLLs loaded, and other process interactions.

As part of Intercept X the agent will classify applications based on how they are registered in the system to understand what types of protection should be applied. Applications that interact with the end user or the internet are automatically classified as such and exploit protection is provided. Applications are classified into the following profiles.

- **Web browsers:** Works with all commercial web browsers.
- **Browser plugins:** These are toolbar helpers and other applications that directly integrate into the browser.
- **Java applications:** Java-based applications are identified and protected.
- **Media applications:** This classification includes image readers, editors, audio players, and other media players.
- **Office applications:** These constitute document creation and reader applications like Word, Excel, Adobe, and a variety of PDF readers and editors.

What happens when an attack is detected?

When exploit activity is detected the exploited application will be terminated, the user notified of the detected activity, a Clean scan will be triggered, and a root cause analysis will be requested.

What should an admin do?

Exploit activity can be from a drive-by attack where the browser is interrogated by an adversary site and compromised. In this case, the browser is shut down and the adversary never successfully breaches the system. In other cases, the adversary may have gained access to the device and launches the attack from another process. In this situation, the device penetrated by the adversary and other non-exploit actions could have been performed by the adversary. The administrator should review the incident report, RCA visualization, and artifact list to get an understanding of the root cause of the exploit. Understanding how the adversary penetrated the device is important. Often it will be found that the user downloaded or authorized an application that granted the adversary access and training will be advised for safe browsing.

CryptoGuard

Ransomware protection is a tough problem to solve, and most vendors currently detect ransomware attacks by the same old methods, detecting the specific malware variants that perform the attack. Given how easy it is to create new software that encrypts or otherwise renders unusable valuable documents, it is not surprising that these types of attacks have flourished. With CryptoGuard, Intercept X is monitoring the file data for rapid change, and when it's detected CryptoGuard suspends the offending process and evaluates if it is a legitimate tool like Sophos SafeGuard, a file and folder based encryption product. If it is not a legitimate business tool performing the encryption the process will be terminated and the just-encrypted files will be restored.

By observing the behavior performed by ransomware attacks instead of trying to detect the infinite variety of software that can be written to perform these attacks, CryptoGuard has been able to detect over 99% of the new ransomware variants unleashed on businesses and consumer users without having to change the behavior-monitoring model. For those rare variants we miss with behavior monitoring, we often prevent the attack with the other security layers, like exploit prevention and deep learning.

Key features

False positive suppression: If anti-exploit detection happens for a desired application, the administrator has three options to suppress the detection.

Fine-grained exclusion controls: The recommended action is to suppress the specific exploit method detected and to continue to monitor the process for other exploit activity. The admin can check the user/device policy and examine the scanning exclusions, exclusion type 'Detected Exploits.' From here the administrator will see a list of exploit detections for users and devices and can add suppression for the specific exploit detection and application.

Course grained exclusion controls: Intercept X also has a broader exclusion capability where an identified application will be exempt from all exploit protection. This control is available in the global settings (exploit mitigation exclusions). The administrator will be provided a list of all detected applications that have been monitored to date by the Intercept X agent and can identify the application that should have all exploit detections suppressed.

Disable exploit policy: Administrators can turn off exploit notifications for entire application profiles like web browsers, plug-ins, java applications, media applications, and Office applications. Disabling exploit notifications for the application profile is a broad exemption and not recommended but available for when administrators require it.

How does Intercept X stop ransomware?

Intercept X monitors over 70 file types that are often targeted in a ransomware attack. These files are the type that adversaries the adversaries have determined individuals and organizations are likely to pay to recover and are primarily focused on productivity documents, images, audio files, and the like. Most ransomware attacks are careful to not encrypt the operating system's core components. They want the device to still function so that they can easily provide instructions to the end user and can eventually decrypt files when a ransom payment is made.

CryptoGuard monitors process activity that interacts with the designated file types and will take a copy of a file prior to any modification. These files are cached on the device using a Sophos-designated recovery folder, and when the files are encrypted the process performing the activity is suspended and examined.

What happens when an attack is detected?

The process is examined to determine if it is a legitimate business application like a file/folder encryption product. If the process is not a legitimate business application, the process is terminated and the files are recovered to their pre-modification state. The end user will be notified of the detection and a root cause analysis and an incident report will be generated and made available for the admin to understand the origin of the attack so they can determine if additional actions should be taken.

The detection event will also trigger a Sophos Clean scan to identify any other latent malware on the device.

On termination of the ransomware process, the device is restored to a green health state. The attack has been mitigated.

What should an admin do?

The attack was detected at runtime so the adversary was not able to interact with the device. Administrators should review the root cause analysis data and confirm that no other actions are required.

Depending on how the ransomware process was deployed and if other activity associated to activity, it may be appropriate to remove the device from the network and perform a more thorough investigation.

Key features

What is the size limit of protected files?

CryptoGuard caches files under 75MB. It is our experience that most files targeted by ransomware are significantly smaller than this limit and as we are able to detect a ransomware attack after observing just a handful of file changes we will detect and terminate the attack before larger files have been impacted.

Is there a difference in what files are protected for a Windows desktop, server, or mac?

CryptoGuard for Windows desktop and servers works the same. In addition, CryptoGuard for Mac will be available shortly.

How long does it take to recover files?

File recovery is fairly straight forward. The cached file cached file has not been encrypted so all CryptoGuard has to do is move it back to where it was prior to the attack. Typically, ransomware attacks are detected after only a handful of files have been modified and so recovery takes less than a second.

What happens to files put in the Sophos cache location?

Files arrive in the cache automatically as part of CryptoGuard protection and are removed once a process is determined to to not be ransomware or when the automatic restore has completed.

How much disk space does the CryptoGuard file recover space use?

We do not set a size limit on how much space CryptoGuard can consume when protecting against ransomware. the file space used to cache files typically stays under 100MB, but this could vary depending on The number of processes being monitored and the rate of file changes.

Application Lockdown

Application Lockdown stops attacks that do not typically rely on software bugs in applications, but instead abuse legitimate capabilities to perform the attack or deploy malware.

How does Intercept X protect applications?

The agent will classify applications based on how they are registered in the system to understand what types of protection should be applied. Applications that interact with the end user or the internet are automatically classified as such and exploit protection is provided.

Application Lockdown monitors an application's activity. It automatically terminates the application when it maliciously attempts to run new code introduced by the application.

For example: Macros in documents are potentially dangerous, as they are created in the Visual Basic for Applications (VBA) programming language, which includes the ability to download and run binaries from the web and also allows the use of powershell and other trusted applications. This unexpected feature (or logic-flaw exploit) offers attackers an obvious advantage as they do not need to exploit a software bug or find a way to bypass code and memory defenses to infect computers. They simply abuse standard functionality offered by a widely-used trusted application and only need to use social engineering to persuade the victim to open the specially-crafted document.

What happens when an attack is detected?

Sophos Intercept X with Application Lockdown will automatically terminate a protected application based on its behavior. For example: when an Office application is leveraged to launch PowerShell, run a macro to install arbitrary code, or manipulate critical system areas, Sophos Intercept X will block the malicious action, even when the attack doesn't spawn a child process.

Upon detection, the user is notified, Sophos Clean is triggered to detect potential other malware components, and a root cause analysis incident report is requested and made available to the administrator.

What should an admin do?

The attack was detected at runtime as a prohibited behavior. Often what was detected was an activity being controlled by a document, webpage, or media file. The incident report and root cause analysis data should identify the file or activity that is the root cause of the attack. Administrators should review the RCA data and confirm that no other actions are required.

Key features

What behaviors are prohibited by lockdown?

The primary behaviors prohibited involve an application attempting to launch powershell unexpectedly, download and install other software, or modify auto-start registry and folder locations.

Can I configure the prohibited behavior controls?

Prohibited behaviors are controlled by Sophos and are not available for customization by the administrator.

Folder Restrictions vs. Application Lockdown?

Some next-gen solutions employ a list of folders that applications can use to run code from (trusted/whitelisted folders). Many solutions, for example, block the temp folder from being used by attackers, whereas Application Lockdown monitors the activity of applications and automatically blocks attacks without IT administrators having to maintain allow or deny lists.

Attackers can easily bypass folder restrictions by using a folder not listed in the blacklist. Due to the behavior-based nature of Application Lockdown, Sophos Intercept X offers a more robust and automated approach to preventing attackers from abusing legitimate applications for a malicious purpose.

Credential Theft Prevention

Intercept X detects when an adversary-controlled process is attempting to extract user and administrator authentication credentials from a device. An adversary credential theft can target multiple operating system components to steal the password or the hashed passwords of users and administrators for the device. Dozens of different tools are available for the adversary to achieve this, but the most commonly used include Mimikatz, a credential extraction tool that targets LSASS (Local Security Authority Subsystem Service) memory, and hashdump, a credential theft tool that extracts the hashed password from the SAM (Security Account Manager) database.

How does Intercept X prevent credential theft?

Instead of targeting the specific tools used by adversaries (and there are a lot of them) Intercept X instead looks for unauthorized interactions with the LSASS runtime memory, the SAM DB registry, and direct extraction of credential data from the hard disk. As a prevention technique, we have tested with a variety of malware and penetration and hacking tools and found the mitigation to be extremely effective without generating false positive alerts for legitimate software that interacts with the LSASS and SAM DB.

What happens when an attack is detected?

When Intercept X detects an adversary attempting credential theft, the process performing the attack will be terminated and a notification will be presented to the end user. This will also initiate a root cause analysis and will alert the administrator of the activity so it can be investigated. The endpoint will be in a red security health state until the administrator clears the alert notification after investigating.

What should an admin do?

The attack was detected at run time, and though the attacking process was terminated, the initial penetration technique could be repeated or the attacker may still have access to the device. Penetration of the device often involves tricking the end user into authorizing the installation of malicious software, or enabling macros or other actions, but in some instances the penetration involved no direct authorization by the end user. Detection of an attack will generate an alert to inform the administrator that a credential theft attempt was detected and further examination of the incident is warranted. To aid in the investigation, this detection will also request the generation of an incident report using Intercept X's root cause analysis capability.

Prevent Child Process vs. Application Lockdown?

Some next-gen solutions require IT administrators to specify which applications should not spawn child processes. They need to chart the workings of each and every application, typically 500 or more, in use by the organization.

However, in today's threat landscape, attackers inject and run arbitrary code from within trusted applications without spawning new child processes. So, unlike other solutions, Sophos Intercept X's behavior-based Application Lockdown monitors an application's activity. When the protected application runs code that was already present on the system, these new processes inherit the lockdown features of the parent process.

Also unlike any other solution, a lockdown is not limited to the protected application. Lockdown is system-wide, meaning attackers cannot leverage the Windows Registry, Windows Command Prompt or other trusted applications to run the new code.

Process Protection (code cave)

Code cave utilization is a technique used by adversaries where they modify what is likely legitimate software so that it contains an additional application. This additional application is inserted into what is called a code cave, a section of the target application's file that is unused by the program. Code caves exist in most applications and adding code to these sections should not break the behavior of primary application. Often the execution code inserted into a code cave is simply a remote shell launcher; these can be very small and simply grant the adversary access to the box where they can perform other actions. This type of attack requires the adversary to have established a presence on the device so they can deploy the software or to trick the user to download and install an application that has the code cave already exploited. One of the primary reasons adversaries use code caves is to avoid detection by the general user and administrators. The expected application still works fine, but the inserted application is also running. If the application that has been modified is a legitimate business tool that the administrator expects to be on the device they are less likely to consider it malware if traditional antivirus detects a problem. Administrators may simply add it to the exemption list, assuming the antivirus engine has generated a false positive. In this way, the adversary establishes persistence on the endpoint and may have even tricked the admin to allow their inserted application to run.

How does Intercept X prevent use of the code cave technique?

A number of tools exist that can use the code cave technique to embed software into another application, and most traditional antivirus solutions simply looks for tell-tale indicators or signatures these tools leave behind when they insert code into the code cave. For Intercept X, we did not want to follow that approach and instead evaluate applications for any code cave utilization. This is done at initial execution of the software, and when we detect the presence of an additional application residing in a code cave we terminate the application.

What happens when an attack is detected?

Upon detection of the use of a code cave the application will be terminated and the user notified. This will also initiate a root cause analysis, and will alert the administrator of the activity so it can be investigated. Sophos Clean will then remove the malware from the device.

What should an admin do?

Upon detection of a code cave utilization, the administrator should check the root cause analysis to determine how the infected application was deployed to the device. It may be that the adversary had already compromised the device by another means and was simply deploying the code cave to ensure persistence on the device. With this attack blocked, the adversary is likely looking for other avenues of attack and persistence. If this was an end user who was tricked into downloading an application with a code cave it is likely the attack has been prevented, but understanding how they attempted to penetrate the device will help determine what training is required or if additional policy controls need to be put in place.

Process Protection (malicious migration – remote reflective DLL injection)

Process migration is a common technique performed by an adversary when they first establish their presence on a device and want to move to another process to either escalate privileges or gain more enduring access. The adversary does not want to lose control when the end user simply closes their browser or terminates a process that has been compromised, so migrating to a system process is desired. Migration techniques can leverage a remote reflective DLL injection. For more information on DLL injections in general, MITRE provides a great resource. A remote reflective DLL attack is similar, but harder to address; the adversary has already compromised one process and from there they are manipulating another process to load DLLs, and run arbitrary code.

How does Intercept X prevent malicious migration?

Intercept X monitors process activity for the behavior allocating memory in a remote process and the injection of DLLs into that process. This behavior is not something that should be happening, and when Intercept X detects this behavior we have high confidence it is malicious and indicates an active adversary or malware script running on the compromised system.

What happens when an attack is detected?

When Intercept X detects an adversary attempting to migrate to another process in this way the attacking process will be terminated and a notification will be presented to the end user. This will also initiate a root cause analysis, and will alert the administrator of the activity so it can be investigated. The endpoint will be in a red security health state until the administrator clears the alert notification after investigating.

What should an admin do?

Because the attack was detected at run time, it is possible that an adversary is still active on the device, and though the attacking process was terminated, the initial penetration technique could be repeated or the attacker may still have access from another process. The detection will also generate an alert to inform the administrator that process migration with remote reflective DLL injection was detected and further examination of the device is warranted. To aid in the investigation, this event will also request the generation of an incident report using Intercept X's root cause analysis capability.

Process Protection (privilege escalation)

When an adversary has gained access to a system, they are often not running at the privilege level they want or need to complete the rest of their attack. A number of methods exist for the adversary to elevate privileges from credential theft to process migration, but with these doors now locked by Intercept X the adversary has to use other techniques. One that comes to mind is stealing the authentication token of a privilege process and inserting it into another process to elevate privileges. All processes running on the device have an authentication token that the operating system uses to determine the privileges of the process. With this technique, the adversary is likely looking to steal the authentication token of a system process. If an adversary can steal the authentication token of a process with system privileges and use it, they have what they want and didn't need to crack the admin user password or perform a process migration to get it. By taking advantage of known system kernel vulnerabilities in unpatched Windows devices, the adversary has a number of well-documented techniques to capture a privileged token from a process and use it for their own purposes. Given the number of methods available for privileged token theft, it is likely more yet-unknown vulnerabilities in the operating system and kernel remain.

How does Intercept X prevent token theft?

Instead of trying to protect from the numerous known vulnerabilities that allow privileged token theft, Intercept X is looking for when a process has a privileged authentication token inserted into it to elevate privileges. This behavior is simply not used by legitimate software and when spotted we can be fairly sure it is an active adversary attack. By detecting this escalation of privileges, Intercept X is able to protect against this technique regardless of what vulnerability, known or unknown, was used to steal the authentication token in the first place.

What happens when an attack is detected?

We will terminate the process and notify the end user. This will also initiate Sophos Clean to remove the malware. Upon detection, a root cause analysis will be generated to determine how the attacking process started and what else may have been happening on the device that is related to the root cause or detected escalation. The endpoint will be put into a red security health state, as this attack indicates an adversary has likely penetrated the device and more investigation is recommended.

What should an admin do?

Like similar exploit prevention detections, administrators should review the root cause analysis report to determine how the attack unfolded and where it came from. Once the investigation is complete the administrator can clear the alert to allow normal operation of the device.

Process Protection (malicious APC use – AtomBombing)

AtomBombing is a technique used by adversaries to trick another application into running malware or other code. The technique is fairly complex and new and involves abuse of the operating systems ATOM tables and asynchronous procedure calls.

How does Intercept X prevent AtomBombing?

Intercept X is looking for abuse of APC calls. Like many of the exploit protection methods already available in Intercept X, the product is able to monitor process activity at the kernel level and, as far as we can tell, this type of behavior is never good.

What happens when an attack is detected?

We will terminate the exploiting application and notify the end user. This will also initiate Sophos Clean to remove the malware and trigger a root cause analysis evaluation to determine how the attacking process started and what else may be happening.

What should an admin do?

Like similar exploit prevention detections, administrators should review the root cause analysis report to determine how the attack initiated and if other actions are required.

Registry Protection (application verifier mitigation – DoubleAgent)

This is another registry trick that adversaries have available in their toolbox. The attack involves modification of the registry to identify software that should run whenever an application is started. This feature from Microsoft is intended to enable developers to monitor and diagnose application activity, but when used by an adversary it is often to ensure that they have access to the box and can circumvent the protection capabilities of the application being run. This attack made the news in 2017 when it was noted that many antivirus products were susceptible to having the registry for the antivirus software modified to run an adversary's application as well. In reality, the attack is much broader than just targeting antivirus products; an application verifier registry change can be used for any application on the operating system. See this Sophos Naked Security article for more information.

How does Intercept X prevent registry modification?

Intercept X will enforce the authorized Windows DLLs when application verification is used. This way, even if an adversary managed to tamper with the registry and set it to launch their attack, the application will instead ignore these illegitimate registry changes. It is important to note that when Intercept X is deployed alongside a competitor's antivirus, we will protect that antivirus product from attacks that use the DoubleAgent (application verifier) technique.

What happens when an attack is detected?

We do not notify the end user or generate an alert when the registry has been modified to launch another application on application verifier activation; we simply ensure the authorized Microsoft Windows utility is launched.

Protect Critical Functions in Web Browsers

This protection is intended to warn when a browser is compromised by a man-in-the-browser attack (MITB) and was developed to defend users from banking trojans by informing them when their browser may be compromised.

How does Intercept X protect the browser?

Intercept X detects when the browser's presentation to the user may not match what is posted by picking up when critical browser functions get hooked by malicious code, and monitors for presentation, networking, and cryptographic functions of the browser. This protection is provided to all applications that register as a browser, including commercial browsers and some homegrown ones.

What happens when an attack is detected?

Intercept X will notify the user that they should close the browser session as it appears compromised. Upon detection, a Sophos Clean scan is run, and an incident report generated with root cause analysis information for the administrator to review.

What should an admin do?

Given that a potential MITB attack was detected, the administrator is advised to utilize the incident report to identify the IP/URL connection that was associated with the attack and determine if that is a location that should be added to a blacklist in the corporate firewall. Because the attack reached this point before it was detected, if web protection was enabled the site is not currently classified as malicious and blocked earlier in the attack.

If the user was providing authentication passwords as part of the session, they should be advised to change their passwords. If they were in fact trying to connect to their financial institution, they are advised to notify the institution to confirm no abnormal activity on the account has occurred.

Root Cause Analysis (RCA)

When malicious activity is detected on a device it is critical that the administrator has more information so that they can understand the nature of the incident, how it happened, what lead to the detection, and what actions that should be taken to prevent similar incidents in the future. This ability to perform forensics on an attack has been the work of security operations centers armed with SIEM and device forensic tools. Unfortunately, the volume of malicious activity detections in a typical organization can easily overwhelm understaffed IT security administrators. To address this challenge, Sophos Intercept X includes automatic root cause attribution with recommended next steps.

Root cause analysis core components

Sophos data recorder – Intercept X includes a data recorder that tracks activity on the device. Information collected on process, memory, network, disk, and registry changes are recorded locally on the device and made available for RCA generation when an incident is detected. Activity over the last 30 days on the device is stored and consumes about 100MB of disk space.

Root cause attribution – Detection of malicious activity, including malware, exploits, ransomware detections, and blocking activity will trigger the RCA algorithm to examine the contents of the data recorder. The objective of the algorithm is to take an identified malicious action, called a beacon event, and then track events associated with the beacon back to an origin or root cause. The root cause can be things like opening an email attachment, inserting a USB drive, browsing to website, and other activity. Once we have traced back to the root cause, the algorithm then moves forward from that event and identifies associated activity. This collection of actions is automatic and will screen out all other activity on the device that was not associated to the beacon or root cause event. By automatically tracking to the root cause and collecting associated data, we are able to perform much of the work that a manual forensic examination would need to do. This dramatically reduces the time required to understand the origin of malicious activity, and by discarding non-associated activity on the device we are able to present the most critical information to the administrator. This is provided in an incident report that is automatically made available at the administration console.

In addition to collecting the associated cause and effect information for each malicious activity, the algorithm also makes an initial determination of the severity of the incident. This will automatically set the priority as low for events that are simple block actions and raise the priority for incidents that may have interacted with multiple user-generated files, or involved multiple processes that may themselves be suspect.

Artifact list information – Artifacts collected during RCA determination are further evaluated to determine additional information about the individual artifact. How long was the process running? What is the reputation of the artifact? Is this a Windows-protected resource being used or is this an unknown or low-reputation executable? What is the SHA256 hash identity of the process involved? Was the network connection made to a known classified website?

Top RCA questions

What is the performance impact of running the data recorder?

The Sophos data recorder consumes less than .5% of a typical machines CPU.

What will trigger an RCA?

A detection event through Intercept X or Sophos Endpoint Protection will trigger a request to generate RCA data.

Will I get an incident report for third-party detections?

When deployed alongside an antivirus product, Sophos Intercept X will not be notified to generate the RCA data for detection events from the third-party software. If the third-party software fails to detect the threat and Intercept X intervenes, an RCA will be requested.

Can this integrate with my SIEM?

Detection event details are accessible to SIEM and other platforms via the API.

Visualization of an attack – To facilitate attack analysis, a visualization graph is available to the administrator that shows how each process in the activity chain interacted with the file system, registry hive, network, and other processes. Administrators can select each node of the graph and receive additional information to further clarify the activity.

Recommended Actions – As part of the incident report we include recommended actions for the administrator. These can be as simple as recommending the administrator evaluate the visualization, or as specific as recommending that device control features be enabled to prevent unauthorized use.

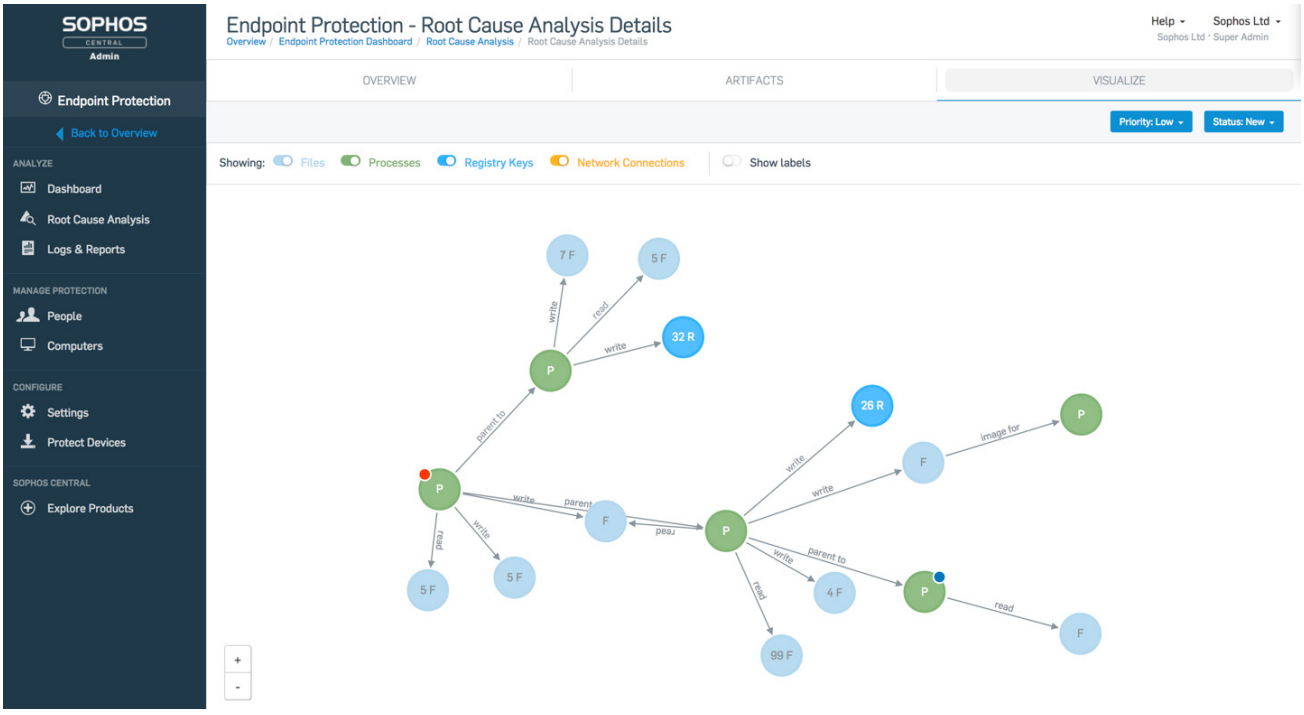

Sophos Clean

Sophos Clean is a component that is called upon whenever malware is detected. Its objective is to collect the detected file, confirm it is not in fact a critical system file that is being abused, and to collect associated elements to the file, like registry keys, links, and other files installed with it. For most malware files, the collection is quite small, but for potentially unwanted applications the collection of files, registry settings, and other components can be extensive. Once collected, the components are removed from the device.

Sophos Clean requires an active internet connection to enable the product to evaluate files, registry, and other remnants of malware as they are collected and to confirm if the component is malicious or not.

When Sophos Clean encounters a tampered-with Windows resource-protected file it will restore the file to correct version and remove the malicious one. This ability to restore resource-protected files on the device allows Clean to successfully take a machine with multiple infections and automatically recover it to a fully working and safe state.

Sophos Clean can detect and remove malware files, cookies, unwanted applications, and registry-based malware and provides detailed results in a log file. When run alone, Sophos Clean provides a detailed report on what was found and the actions taken during the removal and restore process.

Sophos Endpoint Protection

There is no one silver bullet to security, no one security product that will protect you from every potential attack or security risk you may face. Sophos Intercept X has been built to run alongside antivirus products that may provide additional protections or mitigations. Should you chose to run Intercept X alongside Sophos Endpoint Protection, the two products will share the same core agent components, avoiding the need to run multiple agents or perform multiple installations to protect a single device. Both Intercept X and Endpoint Protection can be managed and deployed from the Sophos Central management platform, delivering a single agent and single management interface.

Sophos Endpoint Protection features include:

Control

Peripheral Control – Malware can easily be introduced through removable media. Plugging in an unknown USB could have devastating consequences. Peripheral Control gives admins policy controls to explicitly allow or deny the use of removable media when a new device is detected.

Web Control – Allows administrators to create a list of permitted or denied websites based upon site content category.

Application Control – Block/allow defined applications to further protect endpoints from running unauthorized applications.

Data Loss Prevention – Content scanning includes a comprehensive set of sensitive data type definitions to enable immediate protection of your sensitive data.

Pre-Execution

Web Security – SophosLabs real time lookups prevent devices from communicating with or browsing known bad websites or command and control hosts. Endpoint Protection will also scan the content of sites accessed to look for malicious redirection code or compromised components on the page like malicious Flash objects or Javascript.

Known bad file detection – Signature- or genotype-matching of matching of known malware delivers an efficient and highly effective way to detect malware variants providing baseline for security detections.

File and code analysis – HIPS analysis examines files to look for snippets of code that may indicate it will delete other files, make registry changes, install other files, use encrypted execution code, and perform other malicious actions.

Download Reputation – When we observe the download of a highly suspicious file that has not already been classified as malicious or benign we can ask the user if they want to continue or not. The level of suspicion for a downloaded executable is determined by SophosLabs.

Detection

Application Behavior Detection – SophosLabs curated behavior rules use the HIPS engine to detect malicious behavior from running processes.

Malicious Traffic Detection – Available in both Intercept X and Endpoint Protection. Network behavior monitoring has traditionally been a technique applied only at the firewall or through the aggregation of network data into a SIEM for analytics processing. By performing network behavior monitoring on the endpoint with Malicious Traffic Detection, Sophos Endpoint Protection and Intercept X can observe the process level communications to external devices and detect communications to suspect command and control or other malware delivery servers both when the endpoint is on the corporate network and when the device is roaming off the network. Sophos uses malicious traffic detection to trigger additional analysis of the process that generated the traffic to convict malware that has avoided other detection techniques.

Synchronized Security – Available in both Intercept X and Endpoint Protection. By automating threat discovery, investigation, and response, Synchronized Security revolutionizes threat detection. Incident response times are reduced exponentially and tactical resources can be refocused on strategic analysis. Synchronized Security allows next-generation endpoint and network security solutions to continuously share meaningful information about suspicious and confirmed bad behavior across an entire organization's extended IT ecosystem. Leveraging a direct and secure connection called the Sophos Security Heartbeat, endpoint and network protection act as one integrated system, enabling organizations to prevent, detect, investigate, and remediate threats in near real time, without adding any staff. As an example, when the Sophos XG Firewall detects an advanced threat or an attempt to leak confidential data, it can automatically utilize the Sophos Security Heartbeat to take a series of actions across both the network and endpoint to mitigate risk and stop data loss instantly.

Similarly, if a protected endpoint is discovered to be compromised, synchronized Security allows automated and near instantaneous isolation of this endpoint.

Synchronized App Control: Available in both Intercept X and Endpoint Protection. Endpoint Protection and Intercept X can automatically share detail on which applications are generating network traffic. This allows the Sophos XG Firewall to automatically identify all previously unknown applications enabling you to easily block the apps you don't want and prioritize the ones you do.

Summary

There are many definitions of what a next-generation endpoint security product is or does exist today. This makes selecting the right technology a complex task. With an ever increasing risk surface and growing complexity and volume of attacks, combined with small teams and very tight labor markets, we face a very challenging world for IT security teams.

Multiple point product approaches introduce more problems alongside the challenges they are trying to solve. We must implement new solutions that are simple, yet effective, automated, and synchronized. To learn more and see how Sophos Intercept X and Endpoint Protection can better protect your business, visit sophos.com/intercept-x.

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

Oxford, UK

2018-08-06 SB-NA (MP)