SOPHOS
Security made simple.

What's New in Sophos Intercept X

January 2018

New Deep Learning Malware Detection

Deep learning file scanning model detects new unseen malware and potentially unwanted applications.

The model is less than 20mb and requires infrequent updates.

During training, the model identifies important attributes automatically, resulting in a more accurate decision boundary between malware and benign files.

- **New and enhanced exploit prevention techniques**
  - **Malicious process migration** – This detects a remote reflective DLL injection used by adversaries to move laterally between processes running on the system.
  - **Process privilege escalation** – This prevents a low-privilege process from being escalated to a higher privilege, a tactic often used by an active adversary to gain system access rights.
- **New Active Adversary Mitigations**
  - **Credential theft protection** – Preventing theft of authentication passwords and hash information from memory, registry, and off the hard disk.
  - **Code cave utilization** – Detects the presence of code deployed into another application, often used for persistence and antivirus avoidance.
  - **APC protection** – This detects abuse of Application Procedure Calls (APC) often used as part of the Atom Bombing code injection technique and more recently used as the method of spreading the WannaCry worm and NotPetya wiper via EternalBlue and DoublePulsar. Adversaries can abuse these calls to get another process to execute their code.
- **Enhanced Application Lockdown**
  - **Browser behavior lockdown** – Intercept X prevents the malicious use of PowerShell from browsers as a basic behavior lockdown.
  - **HTA application lockdown** – HTML applications loaded by the browser will have the lockdown mitigations applied as if they were a browser.
- **New registry protections**
  - **Application verifier protection** – Intercept X prevents the replacement of application verifier DLLs that would allow the adversary to circumvent antivirus and other normal process start-up behavior.

Technique Detected	User Notification	RCA	Admin action required	Security Health State	Process Terminated
Credential theft	YES	YES	YES (ALERT)	RED	YES
Code cave	YES	YES	NO (Event)	GREEN	YES
Remote reflective DLL injection	YES	YES	YES (ALERT)	RED	YES
Privilege escalation	YES	YES	YES (ALERT)	RED	YES
APC protection	YES	YES	NO (Event)	GREEN	YES
Application verifier	NO	NO	NO	GREEN	N/A
Lockdown (Browser PowerShell)	YES	YES	NO (Event)	GREEN	YES
Lockdown (ATA from browser)	YES	YES	NO (Event)	GREEN	YES

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.

Advantages of deep learning over traditional machine learning approaches:

- Deep learning automatically identifies what's important in raw data and in so doing yields better accuracy
- Deep learning is "big data native," scaling easily such that it can "memorize" the broad threat landscape and generalize from it to novel threats
- Deep learning is the dominant technology trend in artificial intelligence, meaning that Sophos' deep learning strategy benefits from innovation from the major industry players
- Deep learning yields better detection rates, lower false positives, and dramatically lower footprints than other machine learning detection systems

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 select the attributes of software that they determine most closely correspond to malware. 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 this is on a suppression list. We talk about false positive suppression below, but for now, know that the suppression list allows us to run an extremely 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 detection of malicious activity occurs Sophos Clean will be told to perform a directed removal of 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, if this file is detected it, the block will be suppressed and it will execute as designed.

In addition to In addition to the customer-specific false positive supression list Sophos maintains a global suppress capability. The Sophos false positive suppression 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 reason we have a global false positive suppression capability is to enable 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 this 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.

Credentials Theft Prevention

Intercept X detects when an adversary-controlled process is attempting to extract user and administrator authentication credentials from a device. An adversary attempting 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 lots 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.

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 hide from 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 instead 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. You can read more about AtomBombing here.

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. The 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 the 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.

Improved Process Lockdown (browsers and HTML applications)

Intercept X already includes process lockdown, where we prevent various malicious behaviors of identified process types. With Intercept X we are extending the lockdown capability to prevent web browsers from launching PowerShell and extending the browser lockdown capability to HTML applications that are run by the browser (HTA applications).

How does Intercept X prevent applications running from the browser?

Intercept X will automatically classify an application as a browser and when it does so, classified application lockdown is there to prevent malicious behaviors like PowerShell. Lockdown leverages the ability of Intercept X to monitor the application's activity at the kernel and is always running with the application.

What happens when an attack is detected?

When Intercept X detects an application misbehaving in this way it is prevented from completing the activity and the user is notified.

An event is also generated for the administrator to review.

Detection of this type of malicious behaviour will also request the generation of a root cause analysis for review by the admin.

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

2018-1-9 (GH)

SOPHOS