ExtraHop

WHITE PAPER

Network Traffic Analysis Meets the MITRE ATT&CK™ Framework for Enterprise

How Reveal(x) Enables Greater TTP Coverage

The MITRE ATT&CK™ Framework has rapidly become popular among security teams looking to take a structured and proactive approach to improving threat detection. For many security professionals, using the ATT&CK™ Framework means taking a close look at each of the hundreds of tactics, techniques, and procedures (TTPs) and trying to figure out which tool in their patchwork of solutions is most likely to detect or block any given threat. MITRE provides an evaluation framework for Endpoint Detection and Response (EDR) platforms to test their standard deployments against a subset (56) of the TTPs listed. However, no such evaluation yet exists for network traffic analysis (NTA) products. This document will provide a high-level view of how enterprise NTA with ExtraHop Reveal(x) detects and enables investigation of a broad range of the TTPs catalogued by MITRE ATT&CK™.

TABLE OF CONTENTS

What Is the MITRE ATT&CK™ Framework?

MITRE ATT&CK™ is a globally-accessible knowledge base of adversary tactics and techniques based on real-world observations. MITRE ATT&CK™ was started in 2013 to catalogue observed tactics, techniques, and procedures (TTPs) in use by advanced persistent threats (APTs) around the world. Many of the TTPs included in the framework are in use by far less sophisticated attackers as well, and the structure of the framework is usable by organizations of all sizes and security postures for identifying gaps in security coverage. Since sophisticated TTPs that work well tend to enter the mainstream attack vernacular, the ATT&CK™ matrix offers enduring value for companies looking to vet and improve their detection and investigation coverage.

What Is Network Traffic Analysis (NTA)?

Network traffic analysis (NTA) is the passive monitoring and real time examination of data in flight across a network. By definition, NTA requires protocol parsing and traffic decryption for visibility into application transactions, for the purpose of detecting and investigating adversary behaviors and attack TTPs. This is a nascent category whose criteria and terminology are still evolving, to the point where some industry analysts use the term Network Detection and Response (NDR) to discuss the same set of products. For the purposes of this document, we will use NTA, in alignment with Gartner's 2019 Market Guide for Network Traffic Analysis.[1]

Why NTA Is the Best Approach for Detecting Many MITRE ATT&CK™ TTPs

We contend that NTA is the best overall approach for detecting TTPs used by advanced adversaries, especially in middle and later stages of an attack. Because NTA passively gathers, reassembles, and analyzes traffic that crosses the wire, it offers richer context and is not subject to the same blind spots as solutions that rely on application logs or simply analyze packet headers rather than full content. Additionally, NTA does not require agent instrumentation on each device that it monitors, allowing much broader coverage in places where it is impossible or prohibitively labor-intensive to install agents. In modern, multi-stage attacks, there is little way to avoid communicating on the network at some point. While EDR can offer important detection capabilities for individual actions, NTA offers the best chance of seeing all the links in the chain of an attack, especially in the later stages from lateral movement through exfiltration and impact, and being able to put together the complete picture of what is happening in time to mount a meaningful response.

NTA complements EDR and log-based solutions by covering unmanaged devices and those that cannot be instrumented with an EDR agent. In addition, NTA resists counter-incident-response (IR) activities by attackers who target the endpoint agent itself. Because NTA is passive, operates in real time, and is able to see transaction payloads and decrypt traffic, it is the ideal solution for detecting many of the TTPs listed here, often without adversaries even knowing they are being watched.

[1] Gartner, "Market Guide for Network Traffic Analysis," by Lawrence Orans, Jeremy D'Hoinne, and Sanjit Ganguli. February 28, 2019. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

How Reveal(x) Enables MITRE ATT&CK™ Coverage

Reveal(x) network traffic analysis can fill visibility gaps and provide investigative coverage for some TTPs in all twelve categories of the MITRE ATT&CK™ Framework—including Initial Access, Persistence, Privilege Escalation, Defense Evasion, Discovery, Lateral Movement, Execution, Collection, Exfiltration, Command & Control, and Impact—with special focus on the late-stage attack campaign categories. The framework lists many exploits and behaviors that are likely to generate subtle, yet unusual network traffic patterns on the east-west corridor inside an enterprise network, as well as traversing north-south through the nominal perimeter.

Methodology note: MITRE doesn't yet offer a formal evaluation for Network-centric products, so this coverage has been validated through ExtraHop customer and POC experiences, product engineering, threat research, and internal testing. To assure alignment with MITRE's criteria, we took into account their recommended data sources, detection methods, and mitigation steps for each TTP in determining whether Reveal(x) could provide any of the recommended coverage.

The capabilities that differentiate Reveal(x) and enable it to go above and beyond other NTA products in detecting and investigating MITRE ATT&CK™ TTPs are:

- Out-of-band, passive processing of network traffic at scale (up to 100Gbps). Many vendors top out at 40Gbps, which is not enough for today's enterprises.
- Instant access to application transaction contents at Layer 7 (application details), enabling rapid detection and investigation of suspected threats.
- Real-time detection of threats based on machine-learning driven behavioral analysis to catch unknown unknowns in ways that rules-based detection can't.
- Decryption capabilities, including for Perfect Forward Secrecy (PFS), providing access to concrete evidence of TTPs in use that would otherwise escape detection by concealing themselves in genuine, legitimate traffic. See Appendix B for a deeper dive on how decryption enables confident detection of TTPs that would otherwise remain hidden.

To learn more about network traffic analysis as a category, download a complimentary copy of the Gartner Market Guide for Network Traffic Analysis,[1] which lists ExtraHop as a representative vendor.

Reveal(x) MITRE ATT&CK™ Coverage by Category

Initial Access

Following preliminary reconnaissance, attackers attempt to gain a foothold in the network without being detected. Reveal(x) can provide the MITRE-recommended detection capabilities for the following Initial Access TTPs:

Note: All quotes in TTP listings are taken from the MITRE ATT&CK Framework web page unless otherwise noted.

Drive-by Compromise (T1189): Reveal(x) is able to detect and alert on cross-site scripting, a common mechanism for delivering malicious code in a drive-by compromise.

Exploit Public-Facing Application (T1190): Reveal(x) is able to detect SQL injection, cross-site scripting, and other mechanisms of exploiting public-facing applications.

[1] Gartner, "Market Guide for Network Traffic Analysis," by Lawrence Orans, Jeremy D'Hoinne, and Sanjit Ganguli. February 28, 2019. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

Hardware Additions (T1200): NTA is the most effective way to discover new devices on a network. Reveal(x) can detect when new devices connect to a network, and can tell whether their behavior is consistent with that network's norms, and with the behavior of that device's peers. Reveal(x) can also be configured to apply extra scrutiny against the behaviors of new devices for a set time period after the devices are discovered, so that any malicious behavior is more likely to be caught.

Supply Chain Compromise (T1195): In many circumstances, Reveal(x) can provide telemetry when functional components of software have been tampered with or swapped out for malicious scripts.

Trusted Relationship (T1199): Mitre recommends to "establish monitoring for activity conducted by second and third party providers and other trusted entities that may be leveraged as a means to gain access to the network." Many B2B partner relationships involve intentionally granting some level of network access that may then be misused. Reveal(x) is able to see this activity, and provide precise monitoring and ML-based threat detection on traffic from these trusted third parties.

Valid Accounts (T1078): Reveal(x) can detect when hosts start behaving abnormally compared to their own historical behavior and that of their peers. Reveal(x) can also see when a user has logged into a particular device, enabling analysts and threat hunters to investigate when valid accounts are being used in malicious ways across the network.

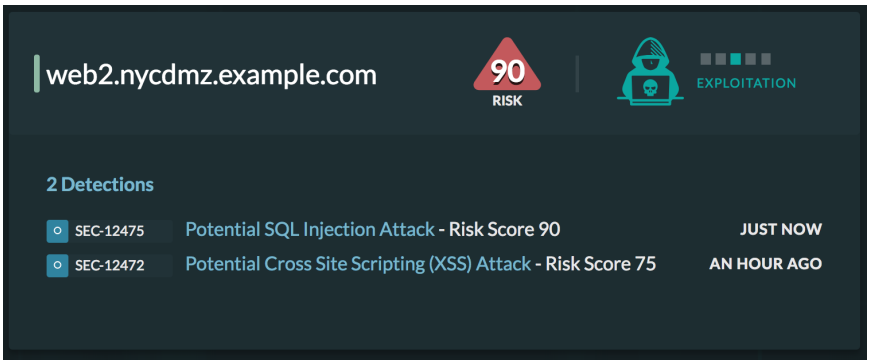

Reveal(x) analyzes SQL and HTTP payloads to detect exploits of public-facing applications.

Execution

This category covers TTPs that allow attackers to execute malicious code on hosts inside the target network. This step makes nearly every other step in the attack chain easier. Attackers may use this ability to pursue persistence, to exfiltrate data, to evade defenses, to establish command and control channels, and more.

Many execution TTPs either rely on payloads that are delivered across the network, or rely on network-based remote management tools to execute malicious actions on compromised endpoints. Reveal(x) is able to detect PowerShell, PSexec, Windows Management Instrumentation and many other mechanisms often deployed by adversaries at this stage of an attack.

Through visibility into this behavior, Reveal(x) provides either detection or telemetry into these Execution TTPs:

Dynamic Data Exchange (T1173): Reveal(x) may provide telemetry into this behavior if executed in ways that generate network traffic.

Graphical User Interface (T1061): Reveal(x) can detect unusual RDP behavior that indicates an adversary manipulating a compromised endpoint via a GUI.

PowerShell (T1086): Reveal(x) detects abnormal usage of remote PowerShell.

Windows Management Instrumentation (T1047): Reveal(x) can detect abnormal Windows Management Instrumentation activity.

Third-party Software (T1072): Reveal(x) may provide telemetry into this behavior if executed in ways that generate network traffic.

Windows Remote Management (T1028): Reveal(x) can detect unusual usage of Windows Remote Management.

Persistence

Once an attacker has gained initial access, they commonly try to create mechanisms that will allow them to maintain access, or regain it through various channels, so that being discovered in any one channel doesn't end their entire operation.

Persistence activities frequently require communication across the network. This creates activity evidence that Reveal(x) will record, detect, and alert on. Attackers may use the remote access capabilities established in the Initial Access step to transmit malware, insert code into local processes, or to gain access to legitimate remote access services such as VPNs or virtual desktop/thin client systems to maintain access to the target network through multiple channels. Examples of Persistence activities that can be detected by Reveal(x) include, but are not limited to:

External Remote Services (T1133): Reveal(x) can detect unusual access patterns and protocol behavior by external remote services.

Browser Extensions (T1176): Malicious browser extensions can be configured to exfiltrate data from the host browser to an attacker's system. Reveal(x) analyzes network traffic to detect data exfiltration, and has successfully detected the C2 channels established by a malicious browser extension. BONUS CONTENT: A real story of Reveal(x) detecting data exfiltration by a malicious Chrome extension.

Windows Management Instrumentation Event Subscription (T1084): If WMI services are accessed remotely, as is often the case, Reveal(x) can see that activity on the network.

Privilege Escalation

Attackers often gain first access to a target network through a low-privilege user account. From there, they will attempt to either increase the privilege level of that account, or use that account to gain access to further accounts with administrative/root access privileges.

Privilege escalation can be detected by monitoring for behavioral changes and changes in interactions with assets that require higher privileges.

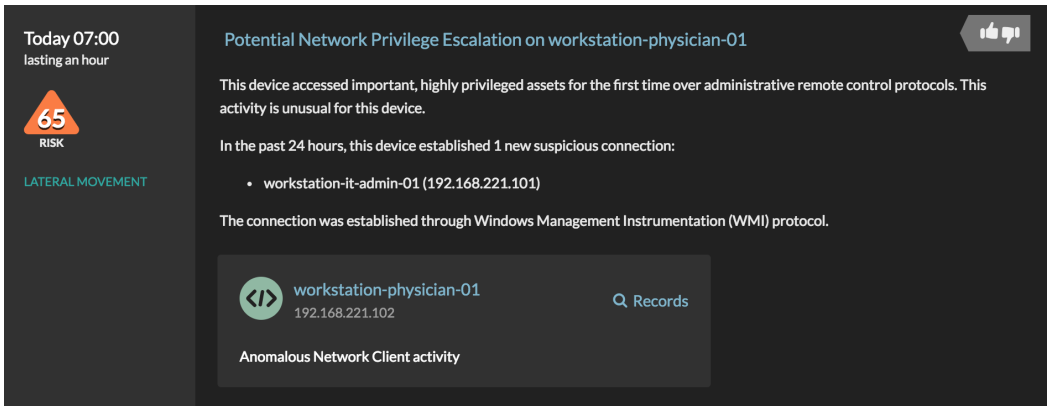

Reveal(x) builds privilege models for each device on the network to identify privilege escalation activity in real time.

Defense Evasion

Attackers are well aware of the common controls and countermeasures likely to be in place on target networks, and will proactively work to hide from them, as well as hide their tracks afterwards.

Reveal(x) can help with the detection and investigation of the following defense evasion TTPs:

BITS Jobs (T1197): Mitre suggests to "monitor and analyze network activity generated by BITS," which uses HTTPs and SMB for remote connections. Reveal(x) detects HTTPs and SMB and analyzes traffic patterns on these protocols, and can therefore detect suspicious activity and even data exfiltration via BITS jobs.

Redundant Access (T1108): Reveal(x) sees all traffic to and from hosts on the network, and can detect a broad range of activity across many channels where adversaries may attempt to establish redundant access.

Network Share Connection Removal (T1126): MITRE recommends being able to detect and parse SMB traffic on the network associated with establishing and removing remote shares, as well as detecting network share session and file transfer activity. Reveal(x) provides visibility into these behaviors.

Web Service (T1102): MITRE recommends monitoring for uncommon flows and decrypting SSL/TLS traffic for analysis to detect this TTP. Reveal(x) can do these activities in real time and parse the Layer 7 traffic for visibility into the actual content of the transactions to see whether malicious code or other unwelcome transmissions have occurred.

Port Knocking (T1205): MITRE recommends detecting this TTP by recording network packets sent to or from the system and examining it for unusual behavior. Reveal(x) observes all traffic crossing the network and can offer telemetry into this behavior.

DCShadow (T1207): Reveal(x) is able to detect DCShadow behavior in real-time.

Install Root Certificate (T1130): Reveal(x) offers telemetry into certificates being used on the network. This includes information about when certificates will expire, sessions using self-signed certificates, and more, but does not include certificate validation. These capabilities enable rapid investigation of this TTP.

Exploitation for Defense Evasion (T1211): This is an extremely broad TTP that essentially covers any vulnerability in software that already exists on the system. In cases where this exploitation results in anomalous use of network ports, traffic, or credentials across the network, Reveal(x) can likely provide telemetry and investigative capability into the outcome of this TTP.

File Deletion (T1107): If executed on networked file shares or executed remotely on a host using remote access tools such as PowerShell, Reveal(x) provides telemetry into the activity of file deletion.

File Permissions Modification (T1222): Reveal(x) can provide telemetry into attempts to change file access permissions if they are executed remotely using PsExec or other operations that generate SMB/CIFS traffic across the network.

Indicator Blocking (T1054): Since Reveal(x) monitors all network traffic to and from each host in its purview, it can provide visibility when a host stops communicating, or experiences a behavioral change in communications with a central logging or event handling system.

Credential Access

Accessing and abusing legitimate credentials is a standard procedure for attackers. Network traffic analysis can help detect and investigate the following credential access TTPs:

Account Manipulation (T1098): Reveal(x) can parse Kerberos traffic and offer visibility into suspicious behavior that could indicate that an account has been compromised.

Brute Force (T1110): Reveal(x) can detect repeated failed login attempts on an asset by parsing Layer 7 traffic to see the login requests and rejection responses, and can tell based on context (such as the timing of the attempts and other metadata) whether it is likely to represent a brute force attack.

Credential Dumping (T1003): Many mechanisms of credential dumping involve compromising domain controllers and other actions that create suspicious network activity. Reveal(x) can detect DCSync attacks, several Mimikatz techniques, GPP file compromise, and more.

Network Sniffing (T1040): Reveal(x) can see excessive ARP broadcasts and other network traffic anomalies likely to indicate attempts at network sniffing, as recommended by MITRE. Reveal(x) can also detect LLMNR/NBT-NS Poisoning & Relay, another TTP that may be a precursor, or related activity, to network sniffing.

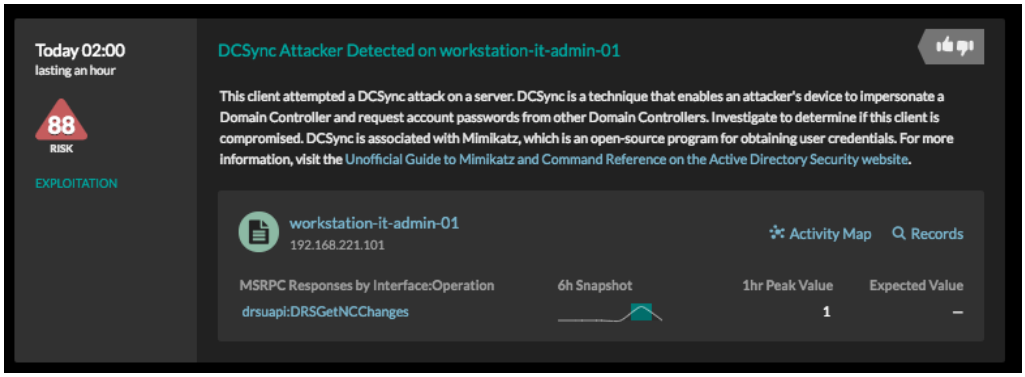

Reveal(x) analyzes MSRPC Responses to detect credential dumping activity, such as DCSync attacks.

Credentials in Files (T1081): Reveal(x) is able to detect web directory and file share scans that would be required for a remote adversary seeking credentials in files. MITRE also recommends monitoring credential usage for anomalous behavior to find if credentials have been compromised. Reveal(x) offers visibility into credential usage across the network.

Exploitation for Credential Access (T1212): Reveal(x) is able to detect unusual access patterns used by specific credentials, and can therefore offer visibility into the impact of this TTP.

Forced Authentication (T1187): Reveal(x) provides visibility into this TTP by monitoring SMB traffic and unusual activity on TCP ports 139, 445, and UDP port 137.

Input Prompt (T1141): In cases where a user changes context on a Windows or Active Directory environment, it generates a kerberos transaction that is observed by Reveal(x). In many enterprises, input prompts of this nature would be noticed and investigated.

Kerberoasting (T1208): Reveal(x) monitors and parses Kerberos traffic and is able to detect this TTP in real time.

LLMNR/NBT-NS Poisoning and Relay (T1171): Reveal(x) is able to detect LLMNR/NBT-NS Poisoning in real time.

Discovery

Discovery is the process an attacker goes through as they attempt to scan a target network to learn about its structure, where valuable data is kept, user names, applications, and other information that a savvy attacker can use to navigate laterally and take actions on objectives.

While many discovery activities are conducted directly on an endpoint, using local command line tools or other local utilities, these activities are also often conducted remotely in ways that require communication across a network. Since Reveal(x) automatically discovers all assets communicating across the network, as well as identifying what type of asset they are and which protocols they're using to communicate, Reveal(x) can often detect these activities when they are executed remotely. Examples of discovery activity into which Reveal(x) can provide detection or telemetry include, but are not limited to:

Account Discovery (T1087): Reveal(x) can detect when local admin accounts are enumerated via Active Directory.

Network Service Scanning (T1046): Reveal(x) is able to detect port scans and can identify known vulnerability scanners active on the network.

Network Share Discovery (T1135): Reveal(x) is able to detect unusual access to network shares.

File and Directory Discovery (T1083): Reveal(x) is able to detect unusual access to files and directories in network shares.

Password Policy Discovery (T1201): Reveal(x) can provide telemetry into this behavior if it is executed across the network via remote access tools.

Permission Groups Discovery (T1069): Reveal(x) can detect unusual LDAP requests that could indicate this TTP in action.

Query Registry (T1012): Reveal(x) is able to detect remote registry enumeration, and can offer telemetry into remote registry queries run over PsExec of Windows Remote Registry Protocol.

Remote System Discovery (T1018): Reveal(x) can detect ICMP scans indicative of remote system discovery.

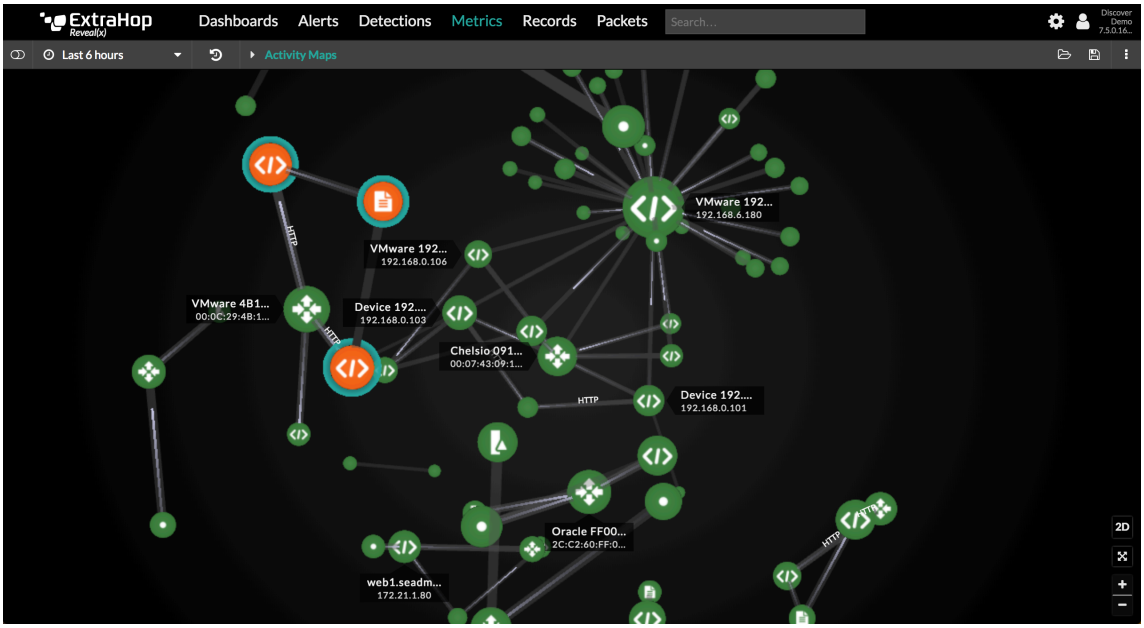

Reveal(x) maps devices communicating on the network, automatically classifying them according to their observed role. Scanning activity can be easily discerned with these live activity maps.

Lateral Movement

The lateral movement phase of an attack is when the attacker moves outward from their original beachhead or compromised host and starts to locate valuable assets to steal or destroy. This may involve using the original compromised host to remotely access other devices in the network, steal admin credentials, or otherwise increase access to internal assets in the target network.

Because NTA focuses on analyzing internal traffic and detecting behavioral anomalies, it is well suited to detecting lateral movement. Examples of MITRE TTPs that Reveal(x) can detect including:

Remote File Copy (T1105): Remote file copying is used both to bring malware and C2 tools into a target network, and to exfiltrate data using protocols and tools such as FTP, rsync, and SMB. Reveal(x) can offer detection and telemetry against this behavior.

Remote Desktop Protocol (T1076): Attackers use RDP to manipulate systems whose credentials have already been compromised, or to expand access by using one compromised system to remotely control another inside the target network. Reveal(x) can detect suspicious RDP activity as well as network privilege escalation.

Shared Webroot (T1051): If a web property's root directory is in a shared location, attackers can upload malicious content to the root directory and then browse to that content, causing it to execute malicious code. Reveal(x) can provide telemetry into this behavior.

Windows Admin Shares (T1077): Windows systems have hidden network shares that are accessible only to administrators and provide the ability for remote file copy and other administrative functions. Attackers use this to expand access once they have compromised a valid account on a machine with administrative shares access. Reveal(x) can detect abnormal access to IPC$, indicating that this TTP is being used.

Windows Remote Management (T1028): This is both a service and a protocol that allows users to remotely access a system. Reveal(x) can see when this protocol is being used in ways indicative of malicious intent.

Remote Services (T1021): An attacker who has access to a valid account will often use it to remotely access systems via Telnet or SSH and execute commands using stolen user credentials. Reveal(x) provides visibility into this behavior. Reveal(x) is also able to detect network privilege escalation by which attackers gain access to administrative privileges.

SSH Hijacking (T1184): MITRE recommends "... Monitor[ing] for user accounts logged into systems they would not normally access or access patterns to multiple systems over a relatively short period of time..." to detect this TTP. Reveal(x) can detect network privilege escalation and can provide real-time visibility into SSH hijacking.

Application Deployment Software (T1017): MITRE recommends to "monitor application deployments from a secondary system" to detect this TTP. Reveal(x) can potentially provide telemetry into suspicious network traffic generated by malicious use of application deployment processes.

Distributed Component Object Model (T1175): Reveal(x) is monitors the DCE/RPC traffic and can provide telemetry into this TTP.

Exploitation of Remote Services (T1210): Remote services often exploit known vulnerabilities in common network services such as SMB and RDP. Reveal(x) detects the use of vulnerable versions of SMB, and unusual RDP activity.

Pass the Hash (T1075): MITRE recommends auditing all logon and credential use events to detect this TTP. Reveal(x) parses LDAP and Kerberos traffic in real time on the wire and can provide visibility into this behavior.

Pass the Ticket (T1097): Reveal(x) monitors Kerberos traffic on the wire so it can detect anomalous behavior and enable forensic investigation of this TTP.

Taint Shared Content (T1080): MITRE indicates that excessive writes or overwrites of files to a network shared directory may indicate this TTP in action. Reveal(x) detects this type of behavior and enables investigation into whether the behavior is malicious. The same behavior can indicate ransomware in effect.

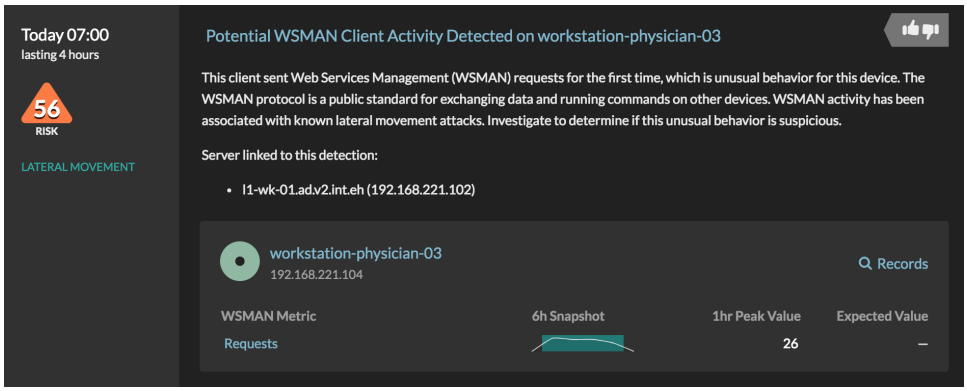

Reveal(x) detects suspicious use of remote management services such as WSMAN and PSexec.

Collection

The collection stage is when attackers begin accessing the data they plan to steal. This may involve moving the data into staging areas for exfiltration, copying data into more easily accessible locations, or establishing new paths to accessing data that will eventually be stolen.

Many, but not all TTPs in this stage involve moving data across the network in ways that would look suspicious to Reveal(x). Collection TTPs that can be detected and investigated with Reveal(x) include, but are not limited to:

Data from Network Shared Drive (T1039): Reveal(x) analyzes traffic patterns against network shared drives and can detect data smuggling activity in this traffic.

Data Staged (T1074): Attackers may gather data into a central directory, break it up into chunks across directories, or otherwise prepare data to be exfiltrated secretly. Reveal(x) can provide telemetry for files created in network shares in this process.

Automated Collection (T1119): Rather than manually combing through data, attackers often automate data collection from network shares or other locations deemed likely to contain valuable data. Reveal(x) can detect scans of web directories and file shares.

Data from Information Repositories (T1213): Attackers may steal data from central repositories like Atlassian Confluence or Microsoft Sharepoint, since these repos are often used to store valuable information about an organization, and even intellectual property like code or summaries of business plans, roadmaps, and HR information. Reveal(x) can detect data smuggling across various protocols used to access these repositories.

Man in the Browser (T1185): This type of 'browser pivoting" behavior may be accomplished through a DNS rebinding attack, which Reveal(x) can detect.

Exfiltration

The Exfiltration stage is when attackers move data off of the target network. Attackers often try to hide this behavior by transferring small amounts of data over long periods of time and by using unexpected protocols and encrypted channels to transfer data.

Exfiltration almost always involves sending data across the network, often to never-before-seen destinations. Reveal(x) can detect when data transfers from inside the network to outside are occurring, and whether the protocols, traffic patterns, and even contents are suspicious. MITRE TTPs in the Exfiltration category that can be detected by Reveal(x) include, but are not limited to:

Exfiltration over Alternative Protocol (T1048): This technique involves sending data across a different protocol than the one being used for command & control. Commonly used protocols for exfiltration include HTTP/HTTPS, DNS, and FTP. Reveal(x) monitors these protocols and uses sophisticated predictive modeling to detect anomalous activity that indicates exfiltration.

Exfiltration over Other Network Medium (T1011): Reveal(x) parses all of the network traffic it can see, so if exfiltration occurs from a part of the network being watched by Reveal(x), it will be detected, even if across a different network medium than that being used for command and control. This method will not detect exfiltration across non-IP network media such as Bluetooth.

Automated Exfiltration (T1020): Once target data has been collected by the attacker, they may use an automation script to exfiltrate the data in batches to avoid detection. Reveal(x) uses sophisticated predictive modeling techniques to detect anomalous traffic patterns indicative of data exfiltration.

Exfiltration over Command and Control Channel (T1041): This technique involves using the same protocols and workstations that are conducting command & control to also exfiltrate data off the network. Reveal(x) uses sophisticated predictive modeling to detect anomalous behaviors that indicate data exfiltration. This method works even if the C2 channel is used.

Data Transfer Size Limits (T1030): Since many security systems will sound the alarm if a large file transfer occurs, attackers often break stolen data into smaller chunks before exfiltrating, to fly below the radar. Reveal(x) can detect unusual protocol behavior and file transfers characteristic of this TTP.

Data Encrypted (T1022): Reveal(x) uses sophisticated predictive modeling techniques to detect anomalous traffic patterns indicative of data exfiltration. This method of detection is effective regardless of whether the data is encrypted.

Scheduled Transfer (T1029): MITRE recommends monitoring network traffic for file access patterns and unusual behavior. Reveal(x) uses sophisticated predictive modeling techniques to detect anomalous traffic patterns indicative of data exfiltration.

Data Compressed (T1002): Reveal(x) detects data exfiltration by conducting traffic analysis, which is effective regardless of whether the data has been compressed.

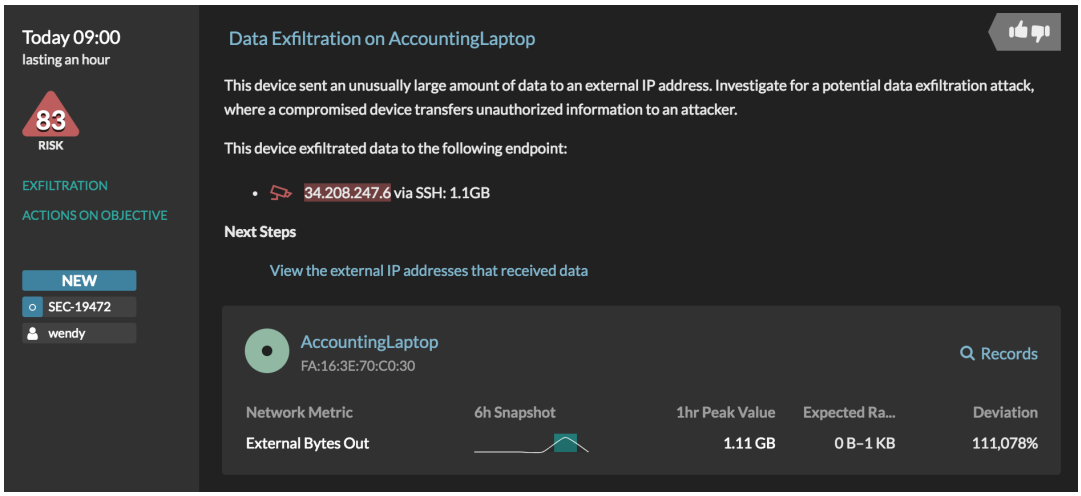

Reveal(x) automatically detects data exfiltration by analyzing network connections and applying behavioral analysis and matching known-bad IPs against threat intelligence feeds.

Command & Control (C2)

Command & Control is the stage when attackers cultivate channels for manipulating data and using compromised endpoints to expand their reach and establish a persistent toehold in the target network. This may involve installing malware on a compromised host that can execute commands on that host or control other rootkitted hosts on the network to coordinate attack activities across multiple hosts.

Many common C2 TTPs require communication across the network, often using common ports. With Reveal(x) watching the traffic on those ports, it becomes much easier to detect C2 activity. For example, an attacker using a commonly open port for C2 may be

transmitting using a protocol that is not usual for that port. By detecting which ports and protocols are in use and decrypting and analyzing the behavior patterns in this traffic, Reveal(x) is able to detect C2 activity in spite of the many tactics adversaries have developed to cover their tracks.

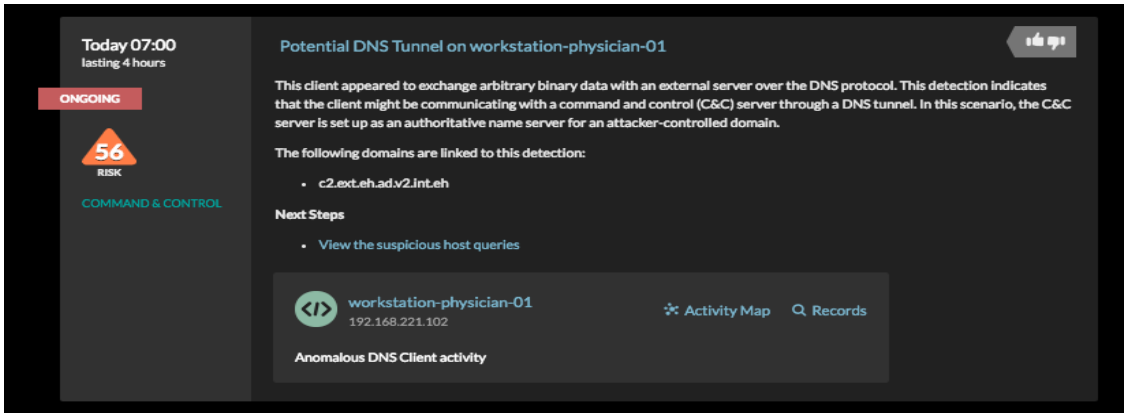

Reveal(x) detects DNS tunneling activity, a subtle method that attackers use to hide their C2 communications.

C2 TTPs that can be detected and investigated by Reveal(x) include, but are not limited to:

Commonly Used Port (T1043): MITRE says that "Adversaries may communicate over a commonly used port to bypass firewalls or network detection systems and to blend with normal network activity to avoid more detailed inspection." Reveal(x) can provide visibility and enable investigation of this behavior.

Connection Proxy (T1090) A proxy is just an intermediary for network communications. Attackers may use proxies to piggyback on existing trust relationships, or to obscure their tracks so it is more difficult to trace the true source of the attack. Reveal(x) analyzes network traffic patterns to detect this and other types of C2 behavior.

Custom Command & Control Protocol (T1094): Attackers may develop their own protocols for command and control instead of using standard protocols. These custom protocols often mimic the behavior of standard application layer protocols, or specifically attempt to mask their behavior by taking advantage of traits of TCP/IP and the network stack. Reveal(x) can detect unusual protocol behavior in real time.

Custom Cryptographic Protocol (T1024): Reveal(x) can detect C2 traffic via behavioral analysis without decrypting the traffic. The use of a custom cryptographic protocol has no bearing on this capability.

Data Encoding (T1132): Adversaries may encode their C2 traffic using standard encoding schemes to evade detection. Reveal(x) detects C2 behavior by analyzing traffic patterns, and is therefore able to detect C2 regardless of the encoding scheme used.

Data Obfuscation (T1001): Reveal(x) can detect when malformed or anomalous data passes across a connection, indicating that the connection is being used for unapproved and potentially malicious purposes.

Domain Fronting (T1172): Reveal(x) can offer telemetry into this behavior.

Domain Generation Algorithms (T1483): Algorithmically generated domains can be used by malicious actors to thwart IP or domain-based perimeter blocking of C2 activity. Reveal(X) uses advanced models to analyze DNS host queries and detect domain names that are likely algorithmically generated and indicative of compromise.

Fallback Channels (T1008): Since Reveal(x) examines all network traffic and conducts behavioral analysis, it can detect when network communication channels are being used abnormally.

Multi-Stage Channels (T1104): Detecting multi-stage C2 channels can be achieved through network protocol analysis the same way that detecting single-stage C2 is achieved.

Multi-Hop Proxy (T1188): Reveal(x) can detect multi-hop proxy C2 behavior through the same advanced traffic analysis that is used to detect other forms of C2 traffic. The addition of more hops does not reduce the effectiveness of this detection.

Multiband Communication (T1026): Reveal(x) parses over 50 enterprise protocols and can detect abnormal usage of those protocols.

Multilayer Encryption (T1079): Reveal(x) detects C2 behavior through traffic analysis regardless of the encryption scheme being used.

Remote Access Tools (T1219): MITRE says that "Adversaries may communicate over a commonly used port to bypass firewalls or network detection systems and to blend with normal network activity to avoid more detailed inspection." Reveal(x) can provide visibility and enable investigation of this behavior.

Standard Application Layer Protocol (T1071), Standard Cryptographic Protocol (T1032), Uncommonly Used Port (T1065), and Standard Non-Application Layer Protocol (T1095): These four TTPs all involve abusing normal protocols and ports for malicious ends. Reveal(x) uses machine-learning driven behavioral analysis and full Layer 2 through Layer 7 visibility to detect anomalous behavior on the network without relying on port numbers for protocol classification. This means that even attackers using standard protocols, ports and cryptographic schemes cannot evade detection.

Web Service (T1102): MITRE recommends monitoring for uncommon flows and decrypting SSL/TLS traffic for analysis to detect this TTP. Reveal(x) can do these activities in real time and parse the Layer 7 traffic for visibility into the actual content of the transactions to see whether malicious code or other unwelcome transmissions have occurred.

Port Knocking (T1205): MITRE recommends detecting this TTP by recording network packets sent to or from the system, looking for "extraneous packets that do not belong to established flows." Reveal(x) observes all traffic crossing the network and can offer telemetry into this behavior.

Remote File Copy (T1105): Remote file copying is used both to bring malware and C2 tools into a target network, and to exfiltrate data using protocols and tools such as FTP, rsync, and SMB. Reveal(x) can offer detection and telemetry against this behavior.

Impact

This category represents “techniques whose primary objective directly reduces the availability or integrity of a system, service, or network.” According to MITRE, these techniques “may represent an adversary’s end goal, or provide cover for a breach of confidentiality.” Essentially, these techniques are at the far-right end of the attack chain, and their occurrence is highly likely to affect business or operations. Of the 14 techniques added to the MITRE ATTACK Framework in this category, three are well suited to detection via network traffic analysis.

Reveal(x) can offer alerts or telemetry for the following TTPs in the Impact category:

Data Encrypted for Impact (T1486): Reveal(x) detects suspicious READ and WRITE actions against network shares using the CIFS protocol. Reveal(x) can also detect EternalBlue. These detection methods can provide visibility and alerts against behavior that indicates the presence of WannaCry, SamSam, Petya/NotPetya, and many other ransomware variants.

Endpoint Denial of Service (T1499): Reveal(x) continuously monitors and analyzes network traffic and can provide visibility into SYN Floods, HTTP Floods, and many other volume- and protocol-based DoS attacks.

Resource Hijacking (T1496): Reveal(x) is able to detect the use of Stratum and other cryptomining protocols on the network, a strong signal that network resources have been hijacked for cryptomining.

Conclusion

Reveal(x) offers great coverage for many of the TTPs in the MITRE ATT&CK™ Framework, representing at least some coverage of each of the attack categories. Several capabilities, including full-stream reassembly of L7 transactions at scale, real-time TLS 1.3 decryption, and guided investigations featuring direct links to the relevant MITRE TTP listings for some detections, make it unique in the NTA product space. These capabilities enable Reveal(x) to detect more MITRE ATT&CK™ TTPs with fewer false positives and more rapid, confident investigations available for each detection.

Reveal(x) provides especially strong coverage of late stage TTP categories, with 89% coverage of the TTPs in the Lateral Movement, Command & Control, and Exfiltration stages, and 100% coverage of TTPs labeled as "Requires Network." Although the MITRE ATT&CK Framework is currently quite endpoint focused, our assessment is that Reveal(x) can provide at least telemetry, and often alerts and enrichment for many of the TTPs listed in these categories based on their behavior and actions on the network. Representing defense-in-depth for endpoint solutions and incremental detection to offset typical blind spots and gaps in network security, this powerful visibility and detection linked to our guided investigations will help Reveal(x) play a crucial role in your security architecture.

Appendix A: Reveal(x) Coverage Matrix

MITRE ATT&CK™ Category	TTPs For Which Reveal(x) Provides Detection, Telemetry, or Enrichment
Initial Access	Drive-by Compromise, Exploit Public-Facing Application, External Remote Services, Hardware Additions, Supply Chain Compromise, Trusted Relationship, Valid Accounts
Execution	Dynamic Data Exchange, Graphical User Interface, PowerShell, Third-party Software, Windows Remote Management, Windows Management Instrumentation
Persistence	BITS Jobs, Browser Extensions, External Remote Services, Port Monitors, Port Knocking, Redundant Access, Windows Management Instrumentation Event Subscription
Privilege Escalation	Port Monitors, Valid Accounts
Defense Evasion	BITS Jobs, DCShadow, Exploitation for Defense Evasion, File Deletion, File Permissions Modification, Indicator Blocking, Install Root Certificate, Network Share Connection Removal, Redundant Access, Web Service, Port Knocking
Credential Access	Account Manipulation, Brute Force, Credential Dumping, Credentials in Files, Exploitation for Credential Access, Forced Authentication, Input Prompt, Kerberoasting, LLMNR/NBT-NS Poisoning, Network Sniffing
Discovery	Network Service Scanning, Network Share Discovery, File and Directory Discovery, Password Policy Discovery, Permission Groups Discovery, Query Registry, Remote System Discovery
Lateral Movement	Application Deployment Software, Distributed Component Object Model, Exploitation of Remote Services, Pass the Hash, Pass the Ticket, Remote File Copy, Remote Desktop Protocol, Remote Services, Shared Webroot, SSH Hijacking, Taint Shared Content, Windows Admin Shares, Windows Remote Management
Collection	Data from Network Shared Drive, Data Staged, Automated Collection, Data from Information Repositories, Email Collection, Man in the Browser
Exfiltration	Automated Exfiltration, Data Transfer Size Limits, Data Compressed, Data Encrypted, Exfiltration over Alternative Protocol, Exfiltration over Command & Control Channel, Exfiltration over Other Network Medium, Scheduled Transfer
Command & Control	Commonly Used Port, Connection Proxy, Custom Command & Control Protocol, Custom Cryptographic Protocol, Data Encoding, Data Obfuscation, Domain Fronting, Domain Generation Algorithm, Fallback Channels, Multi Stage Channels, Multi-Hop Proxy, Multiband Communication, Multilayer Encryption, Port Knocking, Remote File Copy, Remote Access Tools, Standard Application Layer Protocol, Standard Cryptographic Protocol, Uncommonly Used Port, Standard Non-Application Layer Protocol, Web Service
Impact	Endpoint Denial of Service, Resource Hijacking, Data Encrypted for Impact

Appendix B: TTP Coverage Enabled by Decryption

In some cases, decrypting traffic makes the difference in whether you can confidently detect a particular TTP in action. Here are a few examples of TTPs that Reveal(x) excels at detecting, where other vendors cannot, because Reveal(x) decrypts traffic for analysis.

ATT&CK™ TTP	Why Layer 7 Visibility & Decryption Matter
SQL Injection (See T1190)	SQL Injection involves transmitting SQL commands via user input fields transmitted across HTTPS. An NTA tool that cannot decrypt will only see a totally legitimate looking transaction. Reveal(x), by decrypting the traffic, can see that a user has typed a SQL command into the input box instead of the intended content such as their username or password.
Brute Force (T1110)	Brute force attacks across the network create a large number of transactions that can be detected. However, in large, high-traffic environments, that may not be enough of a signal to trigger an alarm. By decrypting the traffic for analysis, Reveal(x) is able to detect brute force attacks more confidently by looking at transaction details, rather than just transaction volume used by many tools.
Database Log Deletion (See T1089)	After an attacker steals data out of a database, they're likely to try deleting the audit logs, if any exist. If the traffic is encrypted, most solutions will not be able to tell that the method to delete the audit table was issued. Reveal(x) can decrypt the traffic and see the methods being used, and allow analysts to investigate the full chain of events leading up to, and following the suspicious event.
Drive-by Compromise (T1189)	Reveal(x) is able to detect and alert on cross-site scripting, a common mechanism for delivering malicious code in a drive-by compromise.

ExtraHop Networks, Inc.
520 Pike Street, Suite 1600
Seattle, WA 98101 USA

http://www.extrahop.com/
info@extrahop.com
T 877-333-9872
F 206-274-6393