NOVETTA

Novetta Cyber Analytics
The Top 10 Built-in Investigative Analytics

Novetta Cyber Analytics is a network security situational awareness solution that substantially increases the effectiveness of security analysts and current security infrastructure. Advanced analytical capabilities enable security analysts to 'see' a complete, truthful, near real-time picture of their entire network, then ask and get answers to subtle and complex questions at the speed of thought. The solution supercharges the effectiveness of incident responders, network security analysts, and current security infrastructure, using a state of the art data processing platform that has proven its speed and scalability on one of the largest and most attacked enterprise networks in the world—the U.S. Department of Defense.

Today, security analysts, incident responders, and network hunters spend most of their time looking for exactly what is described in the below analytics. But most of their time is spent wrangling data from multiple systems with queries that either do not provide enough information, or provide too much, creating queries that take hours to days to provide an answer. The following analytics run on near-real-time data and usually provide answers in seconds, even when run against the metadata of the largest data stores. Novetta Cyber Analytics, for the first time, enables analysts to think and react as fast as their attackers – immeasurably improving the security posture of any enterprise or organization.

The following list describes the top 10 built-in investigative analytic searches. Any analytic can be used as-is or customized for specific use cases.

BEACON

Description:

The Beacon analytic finds beaconing behavior originating from an organization's network. Beaconing is the practice of sending short and regular communications to an external host to inform the external host that the client is alive, functioning, and ready for instructions. Beacons often originate from infected hosts (e.g. bots or zombies) and are sent to command and control (C2 or C&C) servers outside the enterprise network. This "phone home" communication strategy allows botnet administrators to automatically track, manage, and control hundreds of thousands of infected hosts.

This analytic is useful because beaconing behavior is one of the first network-related indications of a malware infection. Typically after malware gets a foothold on a host it quickly determines the host environment and calls out to its C2 infrastructure. If an analyst could spot this activity before the malware can download additional tools, or worse upload sensitive data, they can minimize the damage the malware is able to inflict. The analytic provides the analyst with controls such as beaconing window interval (e.g. one day), beaconing frequency (e.g. three beacons per day), and a persistence percentage to allow for beacon randomization. Since Novetta Cyber Analytics is service agnostic (it does not rely on service/port relationships), it can identify this type of behavior across any service/port pair.

Example:

This analytic is particularly useful when applied

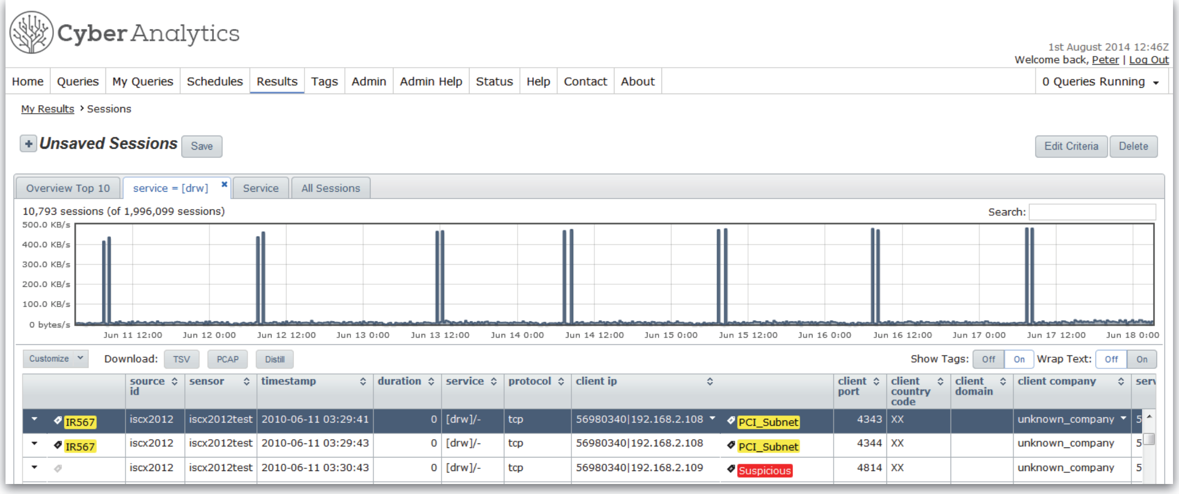

10,793 sessions (of 1,996,099 sessions)

	source id	sensor	timestamp	duration	service	protocol	client ip		client port	client country code	client domain	client company	serv
IR567	iscx2012	iscx2012test	2010-06-11 03:29:41	0	[drw]/-	tcp	56980340|192.168.2.108	PCI_Subnet	4343	XX		unknown_company	5
IR567	iscx2012	iscx2012test	2010-06-11 03:29:43	0	[drw]/-	tcp	56980340|192.168.2.108	PCI_Subnet	4344	XX		unknown_company	5
	iscx2012	iscx2012test	2010-06-11 03:30:43	0	[drw]/-	tcp	56980340|192.168.2.109	Suspicious	4814	XX		unknown_company	5

Beacons being shown on the Novetta Cyber Analytics analyst's interface

to traffic types that by necessity are allowed by enterprise firewalls, such as HTTP (port 80), HTTPS (port 443), and DNS (port 53). If organizations want web browsing to function for employees then these communication channels must be open, and attackers take advantage of this by tunneling their beaconing through these service types. By applying the Beacon analytic to historic web and DNS traffic data an analyst can uncover slow beaconing behavior hidden within the network noise and spot malware infections before they can do any real damage.

Potential False Positives:

False positives for this analytic include regular authorized updates through Windows Update and anti-virus software. These software packages are programmed to actively beacon out at set intervals to check for updates, so they fit the definition of a beacon but they are benign. In Novetta Cyber Analytics, however, it is very easy to identify and remove false positives from result sets through the addition of custom filters and free-text tags. These features enable analysts to iteratively analyze and adjust to hone in on malicious activity.

DISTANT ADMIN

Description:

The Distant Admin analytic finds network sessions between two end points where (a) the service/application being used is administrative in nature and (b) the end points are geographically far apart. By default the distance measure is set to 300 nautical miles or greater, and this value is configurable by the user.

The purpose of the analytic is to uncover remote unauthorized access to enterprise servers and workstations. Although it is possible for infrastructure to be managed from halfway across the world, in most current enterprise environments servers and workstations are serviced by administrators who can physically get to machines if required. This means that administrative activity from a remote and unfamiliar location is cause for investigation.

Example:

An example of network behavior found by this analytic is Remote Desktop Protocol (RDP) traffic between a client in Japan and a server in Canada. If there is no administrator living or traveling in Japan, then there should be no remote access from that location. Potential False Positives: False positives for this analytic include legitimate authorized access from traveling administrators.

HTTP(S) EXFILTRATION

Description:

The HTTP(S) Exfiltration analytic finds unencrypted (HTTP) and encrypted (HTTPS) web traffic where the traffic ratio between the client and the server indicates a data upload to the server. Normal web browsing traffic has more traffic being provided to the client by the server than vice versa – large uploads to servers are uncommon.

The purpose of the analytic is to uncover exfiltration of data from enterprise servers and workstations. Uploads of files to remote servers, especially above certain thresholds (e.g. 100MB), are rare and could be cause for investigation depending on the client performing the upload and depending on the destination server. Although detection of this activity means that damage has already been done, it is still beneficial to detect exfiltration as early as possible. Attackers have been known to steal large volumes of data from enterprises over periods of weeks or months.

Example:

An example of network behavior found by the analytic is stealth data theft using internet file sharing sites as drop points. Attackers commonly use free file sharing or dump sites such as Dropbox to anonymously transfer stolen files out of corporate networks.

Potential False Positives:

False positives for this analytic include (a) photo and video upload sites, (b) personal file storage sites like Dropbox being used for innocent purposes, and (c) remote storage and file sharing sites such as Amazon S3 being used for legitimate corporate file exchange purposes.

PORT SCANNERS

Description: The Port Scanners analytic finds client IP addresses that are performing port scans on the monitored network. It accomplishes this by looking for clients that are sending 0-byte TCP and UDP packets to more than 100 (configurable) distinct server ports on the network.

The purpose of the analytic is to identify network scanning, which is part of an attacker's active reconnaisance activities. Attackers look for open ports and exposed/vulnerable services that they can exploit. If an analyst is able to identify port scanning early they will benefit by (a) identifying potential attackers as early as possible and (b) seeing what responses are sent back to the scanning attempts as this will help the analyst identify weaknesses.

Example:
An example of network behavior found by the analytic is slow randomized port scans from outside or inside a network. Attackers know that automated defense systems are watching for port scans, so they slow down their activities and randomize which ports they touch in an effort to avoid tripping alarms on devices like IDS systems. Port scans typically come from external sources, but if a trusted insider is looking to gain unauthorized access within a network they could launch port scans as well in search for vulnerable internal services.

Potential False Positives:
Network analysts periodically port scan their own networks looking for security weaknesses. This activity would be returned because the analytic does not know the intention of the client, only that they match the profile for the search. So it is possible to see some benign activity. For a properly-configured Port Scanners analytic search, however, false positives should be very rare since benign clients don't typically have a reason for sending 0-byte (empty) packets to many server ports. Once authorized internal IP addresses are whitelisted, the results should only include attacker activity.

PROTOCOL ABUSE

Description:
The Protocol Abuse analytic finds network sessions where there is a mismatch of the service being used and the communication channel (port) between the two hosts. Traffic of this type is said to be abusing a protocol because it is creating network traffic that does not conform to established standards.

The purpose of the analytic is to uncover covert communication channels created by attackers. After a successful intrusion into a machine, attackers routinely set up backdoors or hidden access paths that give them direct and undetected access. A common technique is to tunnel communication through a common service port, such as port 80 (HTTP), because these ports are allowed by firewalls and other network security devices.

Example:
An example of network behavior found by the analytic is custom encryption that does not adhere to SSL protocols, an example of which is custom RC4-encrypted reverse shell activity. A reverse shell is created when an attacker opens a command line shell connection from the victim machine to the attacking machine. It is called a reverse shell because the normal direction is usually the opposite – the client creates a connection to the server. This is effective because firewalls typically focus on blocking incoming traffic and allowing all outbound traffic. If an attacker manages to compromise a machine and starts an encrypted reverse shell, especially on a common port (port 80 for web traffic), this activity often goes unnoticed since it is lost within the network noise. Novetta Cyber Analytics can detect that SSL traffic is being transmitted over port 80, can quickly highlight this malicious traffic, and can provide the observed network traffic to the analyst for forensic analysis.

Potential False Positives:
False positives for this analytic include benign activity such as HTTP over the HTTPS port 443 and streaming video protocols over web ports 80, 8080, and 443.

RELAY FINDER

Description:
The Relay Finder analytic finds internal and external hosts that were used by attackers as relays (also known as proxies, pass-throughs, and hops) while attacking a network. It takes as input a known relay and finds other relays based on the assumption that a chain of relays are used concurrently. The analytic does this by first computing distinct IP addresses that the known relay talks to (the potential victims), then computing the distinct IP addresses that those potential victims talked to (the potential next relays), and returning the potential next relays that interact with at least 80% (configurable) of the potential victims.

The purpose of the analytic is to retrace the path an attacker took by analyzing relationships between hosts in network traffic. This is a simple form of graph analysis in which the system attempts to connect dots between nodes.

Example:
A popular form of relay for disguising lateral movement pivots is a Netcat backpipe relay. Netcat

is a networking service for reading directly from and writing directly to network connections. Attackers use this service to chain together hosts and forward traffic across a network without anyone noticing, because it gets lost in the network noise. Through the analysis of network traffic, however, one can detect these anomalous chained connections by traffic relationships and detect a stealthy attacker.

Potential False Positives:

False positives for this analytic include high traffic sites (e.g. Google and Facebook) that interact with a very large number of hosts.

RDP KEYBOARD LAYOUT

Description:

The RDP Keyboard Layout analytic summarizes Remote Desktop Protocol (RDP) sessions by the layout of the keyboard being used by the client. Administrators of corporate resources typically use keyboard layouts (e.g. US English) that are consistent with the primary locations for the enterprise. If non-standard keyboard layouts are observed, this could indicate unauthorized access to infrastructure by a foreign attacker.

Example:

An example of activity found by the analytic is RDP interaction between a corporate database server and a client using a keyboard layout of Simplified Chinese. It's important to note that this is not just a hypothetical scenario – this behavior is observed in real world attacks. A blog post by FireEye explains that "in 1,849 of the 1,905 (97%) APT1 Remote Desktop sessions [FireEye] observed in the past two years, the keyboard layout setting was 'Chinese (Simplified) — US Keyboard.'" 1

Potential False Positives:

False positives for this analytic include authorized remote administrators accessing corporate resources. Some organizations that use the "follow-the-sun" systems support workflow see administrative activity from all over the world. If, however, these keyboard layouts and network locations are whitelisted they can be ruled out as suspicious activity.

SUSPICIOUS ADMIN TOOLS

Description:

The Suspicious Admin Toolkits analytic finds sessions where the client is using a Remote Administration Toolkit (RAT) to interact with the server. There are many RATs that are often used by attackers to streamline or automate malicious actions. Novetta Cyber Analytics detects toolkits such as Poison Ivy, Radmin, and Gh0st RAT, and is continually adding additional logic and filters to detect other RATs.

Example:

An example of activity found by the analytic is traffic related to the Poison Ivy RAT. Poison Ivy bypasses normal security mechanisms to secretly control programs, computers, and network connections. It gives an attacker nearly complete control over the infected computer and enables the following functionality: file upload and modifications, Windows registry changes, current process control, service control, remote shell execution, keylogging, screen grabbing, and password dumping. The tool is popular because it makes controlling a compromised machine easy.

Potential False Positives:

If configured properly this analytic should not produce false positives, but if the analyst includes common admin services such as Remote Desktop Protocol (RDP) and Secure Shell (SSH), then benign traffic may be returned.

TWO DEGREES OF SEPARATION

Description:

The Two Degrees of Separation analytic takes two known IP addresses or CIDR blocks as input and returns all IP addresses that communicate with both. This powerful analytic is useful for determining relationships between IP addresses based on common activity with other hosts.

Example:

This analytic is useful for many traffic intersection searches, including discovering infrastructure used by attackers after multiple hosts within a network have been compromised. Attackers commonly spread their activities across many external and internal hosts, and use the same command and control infrastructure across multiple compromised internal networks. By being able to intersect traffic between hosts or networks an analyst can rapidly narrow down infrastructure that talks to two or more internal hosts or networks, which can flush out the full depth of a compromise across an enterprise.

UNKNOWN SERVICE

Description:

The Unknown Service analytic returns sessions where the service being used by the client and/or server could not be recognized. Examples of known services are HTTP (web browsing) and FTP (file transfer). An unknown service means that an application-specific service is being used or the traffic is abnormal and doesn't match a known application.

The purpose of the analytic is to give the analyst visibility of network traffic that is uncommon, suspicious, and potentially malicious. Attackers frequently attempt to hide their activities by using unknown or custom services on open ports within a network. This analytic will set aside the normal traffic and reveal the abnormal.

Example:

An example of network behavior found by the analytic is reverse shell activity. A reverse shell is created when an attacker opens a covert command line shell connection from the victim machine to the attacking machine. It is called a reverse shell because the normal direction is usually the opposite – the client creates a connection to the server. This is effective because firewalls typically focus on blocking incoming traffic and allow all outbound traffic. If an attacker manages to compromise a machine and starts a reverse shell, especially on a common port (port 80 for web traffic) using an unknown (custom) service type, this activity often goes unnoticed since it is lost within the network noise.

Potential False Positives:

False positives for this analytic can include BitTorrent file sharing and data channels that are separate from control channels (such as in FTP file transfers). These traffic types can be filtered from the analytic so that the analyst can focus on more malicious activity.