

Cisco Encrypted Traffic Analytics

Introduction

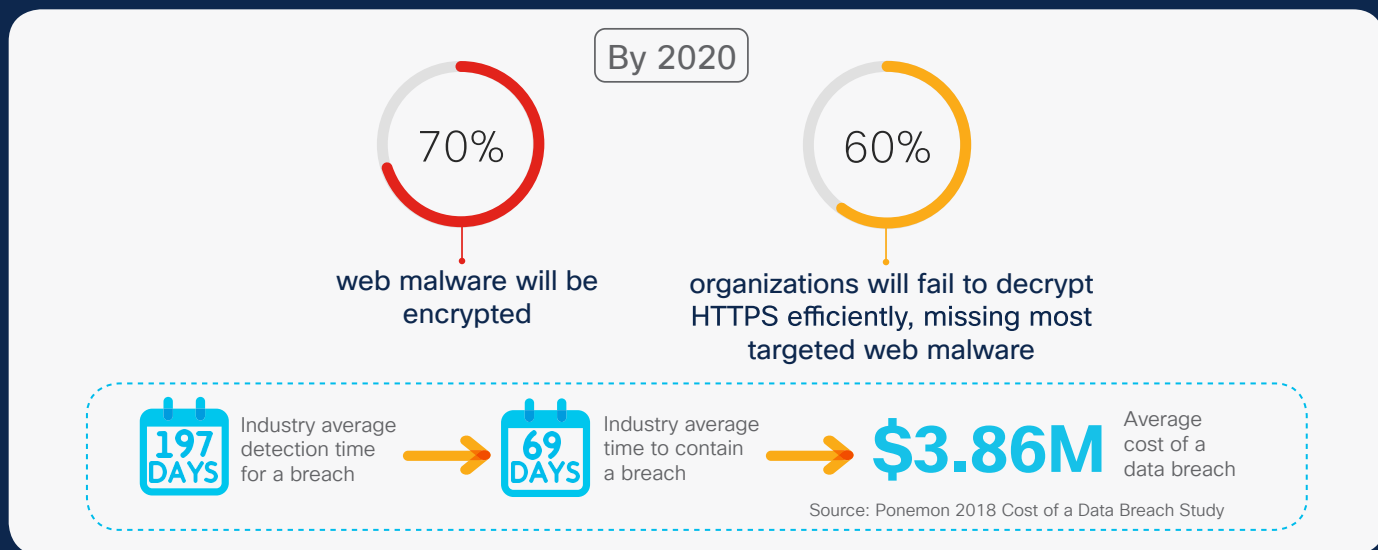
The rapid rise in encrypted traffic is changing the threat landscape. As more businesses become digital, a significant number of services and applications are using encryption as the primary method of securing information. Gartner estimates that more than 80% of enterprises' web traffic is encrypted in 2019.¹ In fact, as of May 2019, 94% of all Google web traffic is encrypted.² And nearly 80% of web pages loaded by Firefox use HTTPS.³

Encryption technology has enabled much greater privacy and security for enterprises that use the Internet to communicate and transact business online. Mobile, cloud, and web applications rely on well-implemented encryption mechanisms, using keys and certificates to ensure security and trust. However, businesses are not the only ones to benefit from encryption. Threat actors have leveraged these same benefits to evade detection and to secure their malicious activities.

Visibility across the network is getting increasingly difficult, and our traditional means of detection cannot assume that data is available for inspection. We need to be able to simultaneously assess how much of our digital business is protected and unprotected by encryption while also assessing what traffic is malicious and what is benign.

More than 70% of malware campaigns in 2020 will use some type of encryption to conceal malware delivery, command-and-control activity, or data exfiltration.¹ And 60% of organizations will fail to decrypt HTTPS efficiently, missing critical encrypted threats.¹

Figure 1. Economic impact of malicious attacks



Contents

Introduction

Challenges of encrypted traffic security

Overview of Cisco Encrypted Traffic Analytics

Encrypted Traffic Analytics—New data elements for encrypted traffic

Encrypted Traffic Analytics - Components

Enhanced NetFlow

Cisco Stealthwatch

Cryptographic assessment

Efficacy results and real-world deployment

Conclusion

Appendix A

References

Table 1. New threat vectors based on nature of encrypted traffic

Uninspected encrypted traffic	Threats
Employees' web browsing over HTTPS	<ul style="list-style-type: none">Malware infectionCovert channel with the command-and-control serverData exfiltration
Employees on an internal network connecting securely to network edge (DMZ) servers	Lateral expansion from infected hosts
Internet users connecting to the enterprise's public servers using encrypted protocols	Reduced defense-in-depth, with only one protection technology inspecting incoming traffic

Challenges of encrypted traffic security

The majority of organizations today do not have a solution to detect malicious content in encrypted traffic. They lack the security tools and resources to implement a solution that can be deployed throughout their network infrastructure without slowing down the network.

Traditional threat inspection with bulk decryption, analysis, and re-encryption is not always practical or feasible, for performance and resource reasons. Also, it compromises privacy and data integrity.

On any given day, no one knows how much of their digital business is in the clear versus encrypted. If traffic is encrypted, the encryption is typically done to meet compliance requirements that mandate specific security policies.

Overview of Cisco Encrypted Traffic Analytics

Traditional flow monitoring provides a high-level view of network communications by reporting the addresses, ports, and byte and packet counts of a flow. In addition, intraflow metadata, or information about events that occur inside of a flow, can be collected, stored, and analyzed within a flow monitoring framework. This data is especially valuable when traffic is encrypted, because deep-packet inspection is no longer viable. This enhanced intraflow metadata is derived by using new types of data elements or telemetry that are independent of protocol details, such as the lengths and arrival times of messages within a flow. These data elements have the attractive property of applying equally well to both encrypted and unencrypted flows.

Using these data elements or enhanced intraflow telemetry to identify malware communication in encrypted traffic means Cisco® Encrypted Traffic Analytics can maintain the integrity of the encrypted flow without the need for bulk decryption (Figure 2). Table 2 lists the benefits of using Encrypted Traffic Analytics.

Figure 2. Encrypted Traffic Analytics – technical solution overview

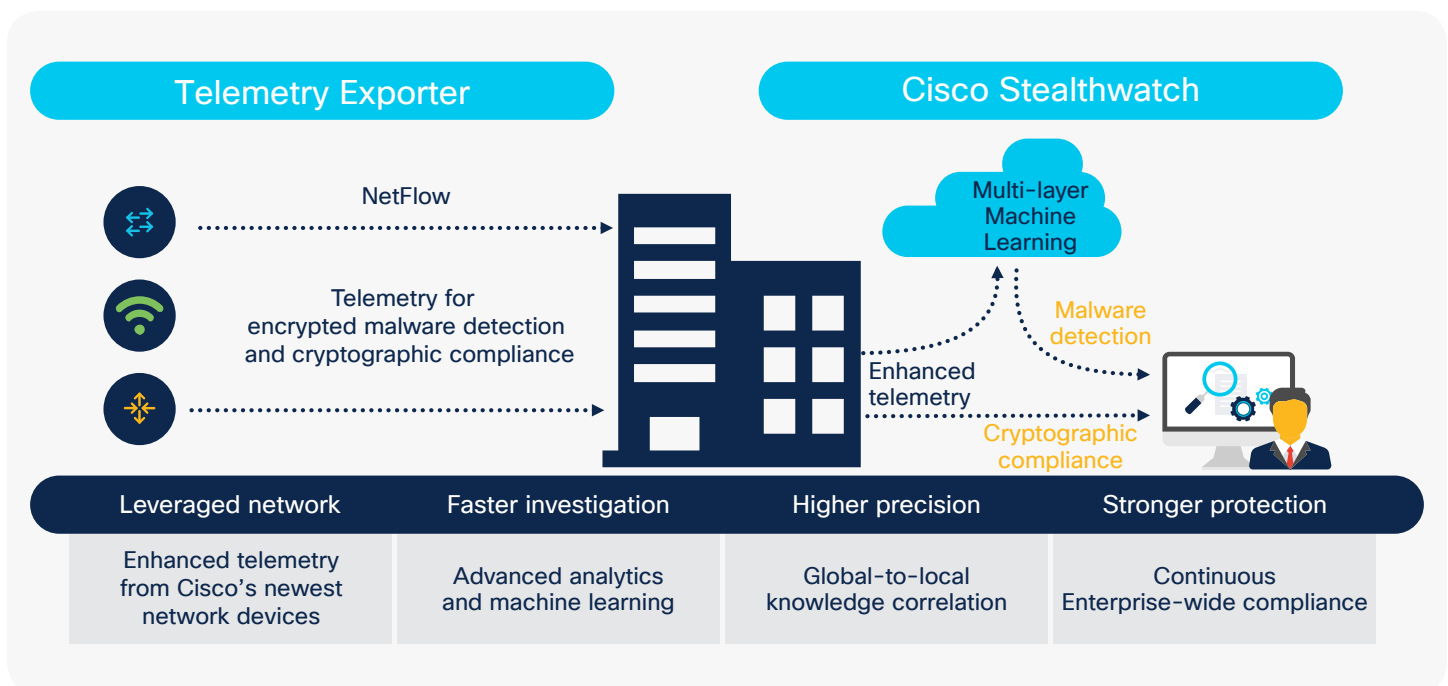


Table 2. Benefits of using Encrypted Traffic Analytics

Benefits
<ul style="list-style-type: none"> • Security visibility: Gain insight into threats in encrypted traffic using network analytics. Obtain contextual threat intelligence with real-time analysis correlated with user and device information. • Cryptographic assessment: Ensure enterprise compliance with cryptographic protocols and visibility into and knowledge of what is being encrypted and what is not being encrypted on your network. • Faster time to response: Quickly contain infected devices and users. • Time and cost savings: Use the network as the foundation for the security posture, capitalizing on security investments in the network.

Encrypted Traffic Analytics—New data elements for encrypted traffic

Encrypted Traffic Analytics⁴ focuses on identifying malware communications in encrypted traffic through passive monitoring, the extraction of relevant data elements, and a combination of behavioral modeling and machine learning with cloud-based global visibility.

Transport Layer Security (TLS) is a cryptographic protocol that provides privacy for applications. TLS is usually implemented on top of common protocols such as HTTP for web browsing or Simple Mail Transfer Protocol (SMTP) for email. HTTPS is the use of TLS over HTTP. This is the most popular way of securing communication between a web server and client and is supported by most major web servers.

Encrypted Traffic Analytics extracts four main data elements: The initial data packet, the sequence of packet lengths and times, the byte distribution, and TLS-specific features. Cisco's unique Application-Specific Integrated Circuit (ASIC) architecture provides the ability to extract these data elements without slowing down the data network.

- **Initial Data Packet (IDP):** IDP is used to obtain packet data from the first packet of a flow. It allows extraction of interesting data such as an HTTP URL, DNS hostname/address, and other data elements. The TLS handshake is composed of several messages that contain interesting, unencrypted metadata used to extract data elements such as cipher suites, TLS versions, and the client's public key length.

- **Sequence of Packet Lengths and Times (SPLT):** SPLT conveys the length (number of bytes) of each packet's application payload for the first several packets of a flow, along with the interarrival times of those packets.

SPLT can be represented as an array of packet sizes (in bytes) along with an array of times (in ms) representing the time since the previous packet was observed.

- **Byte distribution:** The byte distribution represents the probability that a specific byte value appears in the payload of a packet within a flow. The byte distribution of a flow can be calculated using an array of counters. The major data types associated with byte distribution are full byte distribution, byte entropy, and the mean/standard deviation of the bytes. For example, using one counter per byte value, an HTTP GET request, "HTTP/1.1.", can be calculated by incrementing the corresponding counter once for the "H," then incrementing another counter twice for the two consecutive "T's" and so on. Although the byte distribution is maintained as an array of counters, it can easily be turned into a proper distribution by normalizing by the total number of bytes.

Appendix A shows a detailed table of new data elements.

Encrypted Traffic Analytics - Components

Enhanced NetFlow

In the NetFlow architecture, data is transmitted from exporter to collector in sets of records. Each record in a data set has the same format, which is specified by its template. The data record consists of a series of NetFlow information elements or "fields," and a specific ID value is assigned to each field. The ID values for information elements may be globally defined and archived by the Internet Assigned Numbers Authority (IANA), or they may be enterprise specific and defined by individual organizations.

NetFlow templates use several globally defined elements administered by IANA. Some of the global elements, such as IP addresses and Layer 4 port numbers, form a familiar 5-tuple that is used as a unique flow identifier (flow key). Additional elements are used to report basic packet/octet statistics and timestamps.

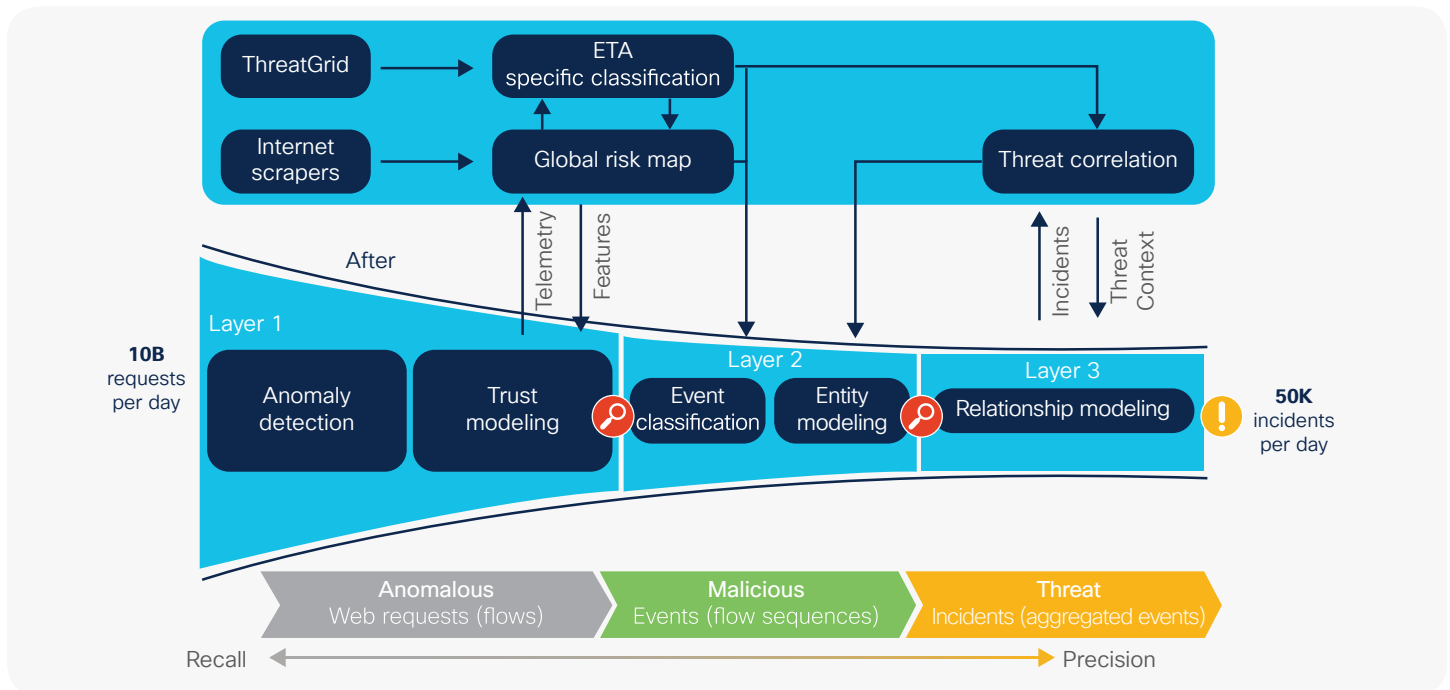
These globally defined elements are enhanced with vendor-specific (Cisco vendor ID) data elements described earlier and in Appendix A. The vendor-specific data elements provide insights into threats and vulnerabilities in encrypted traffic using Cisco Stealthwatch®.

Cisco Stealthwatch

[Cisco Stealthwatch](#) uses NetFlow, proxy servers, endpoint telemetry, policy and access engines, and traffic segmentation as well as behavioral modeling and machine learning to establish baseline “normal” behavior for hosts and users across the enterprise. Stealthwatch can correlate traffic with global threat behaviors to automatically identify infected hosts, command-and-control communication, and suspicious traffic.

Stealthwatch maintains a global risk map—a very broad behavioral profile about servers on the Internet, identifying servers that are related to attacks, may be exploited, or may be used as a part of an attack in the future (Figure 3). This is not a blacklist, but a holistic picture from a security perspective. Stealthwatch analyzes the new encrypted traffic data elements in enhanced NetFlow by applying machine learning and statistical modeling. The global risk map and Encrypted Traffic Analytics data elements reinforce using advance security analytics. Rather than decrypting the traffic, Stealthwatch uses machine learning algorithms to pinpoint malicious patterns in encrypted traffic to help identify threats and improve incident response.

Figure 3. Stealthwatch multi-layer machine learning



The security insight dashboard on the Stealthwatch Management Console (SMC) provides a view of affected users identified by risk type. An expanded dashboard provides detailed information regarding the top risk escalations and relative threat exposure. Table 3 lists some high-risk threats that use encrypted command-and-control communications.

Figure 4. Stealthwatch security insight dashboard

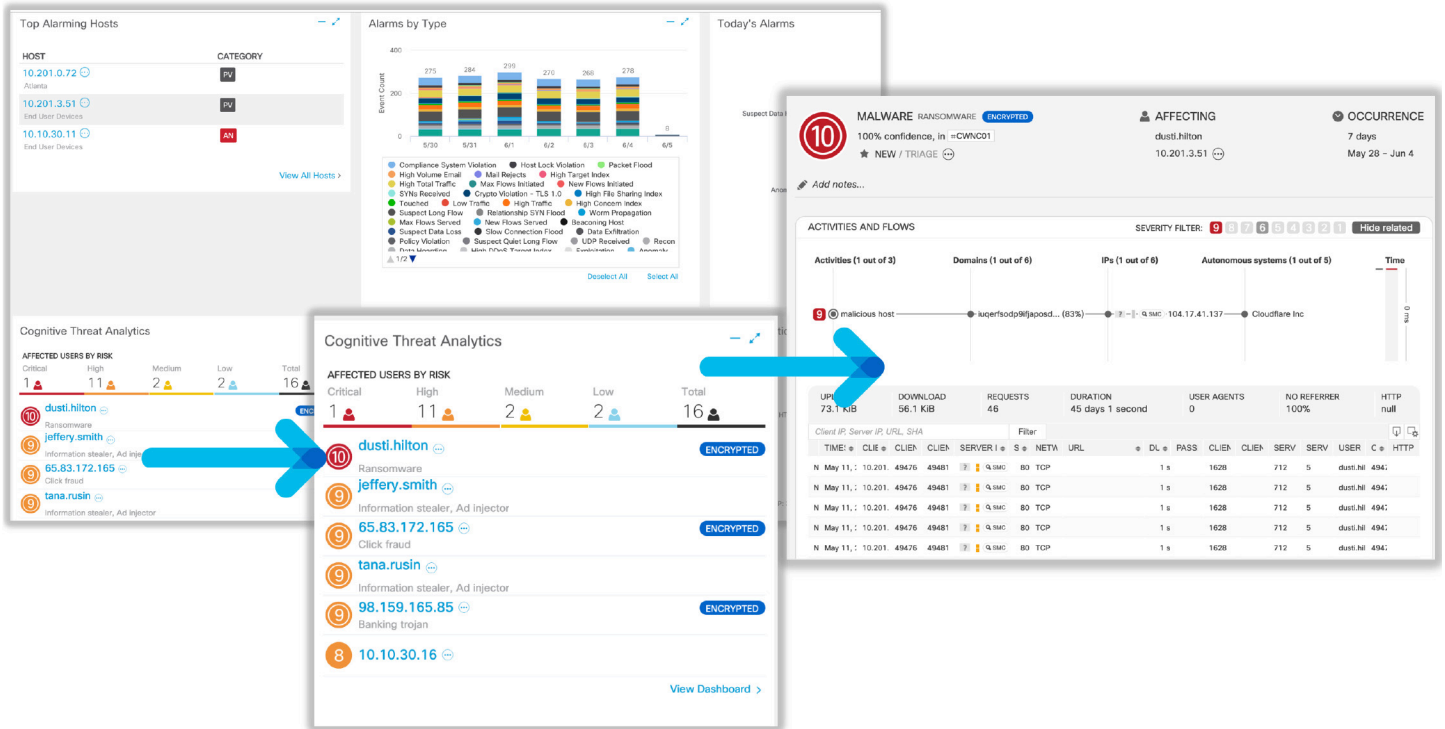


Table 3. Examples of high-risk threats using encrypted command and control

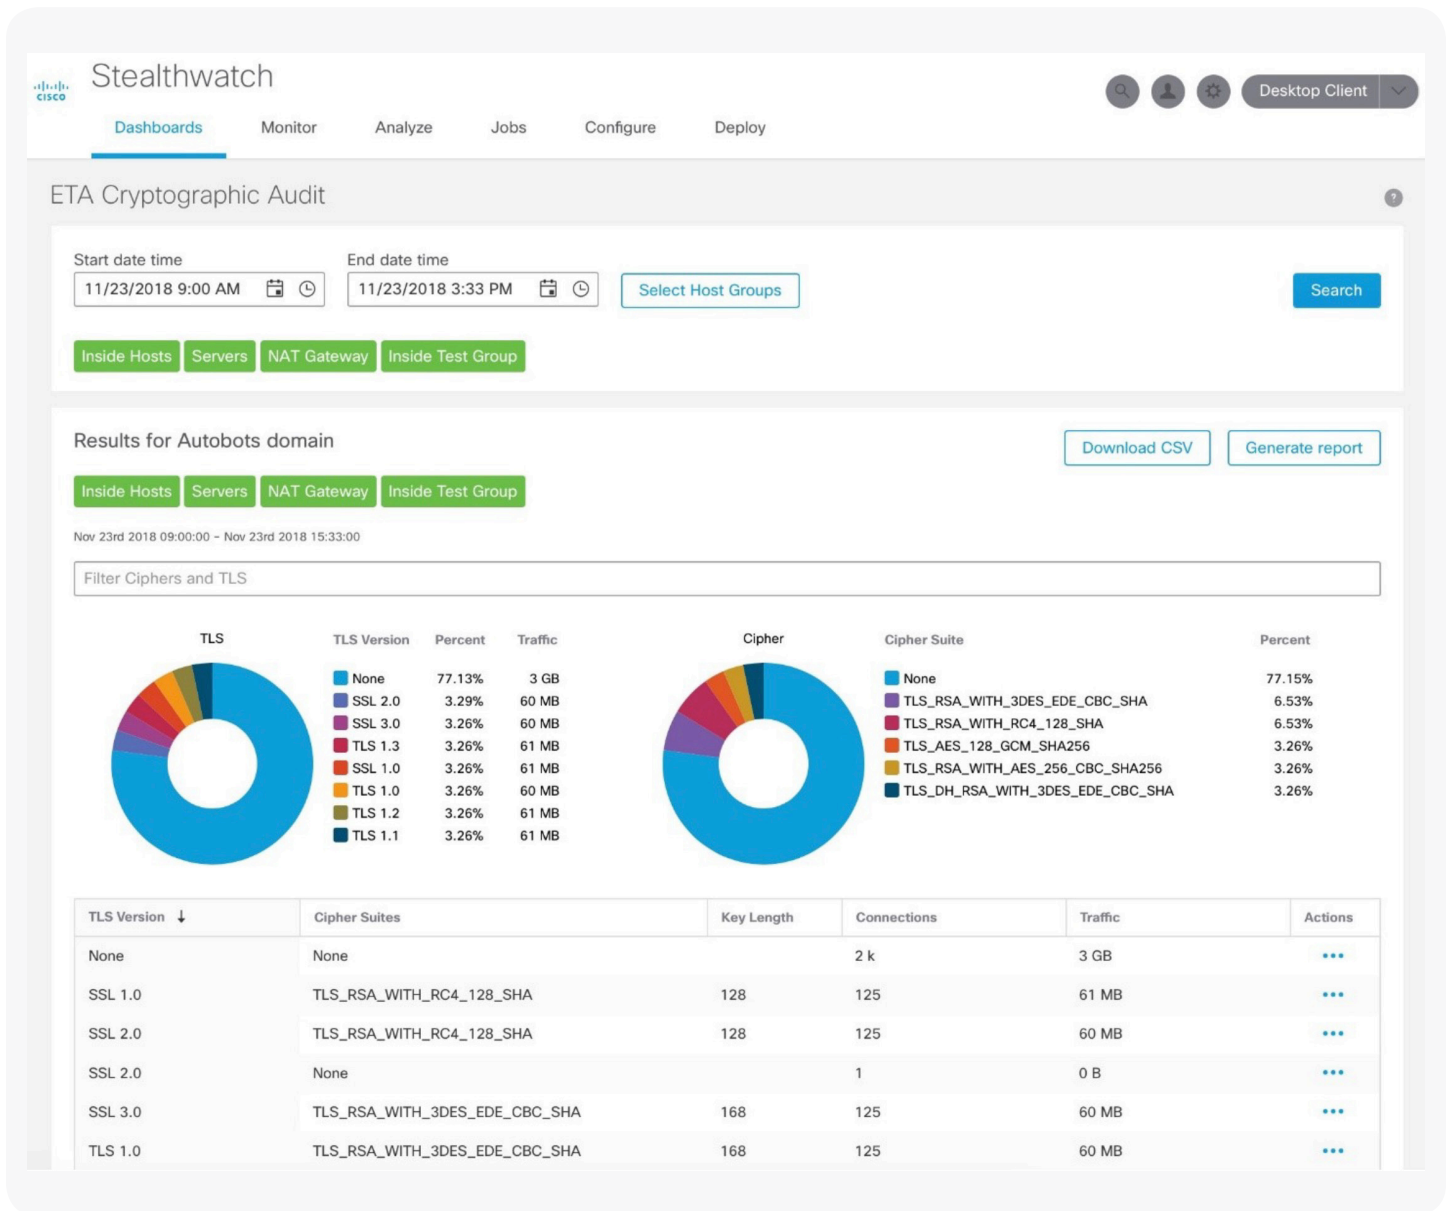
Name	Type
Gamarue/Andromeda	Modular botnet
Salinity	File infector, modular botnet
Necurs	Information stealer, backdoor, botnet
Rerdom	Click-fraud, botnet

Upon discovery, a malicious encrypted flow can be blocked or quarantined by Stealthwatch. Policy-driven remediation actions via pxGrid using Cisco Identity Services Engine (ISE) with Cisco TrustSec® and Software-Defined Access (SD-Access) simplify and accelerate network security operations.

Cryptographic assessment

Encrypted Traffic Analytics also identifies encryption quality instantly from every network conversation, providing the visibility to ensure enterprise compliance with cryptographic protocols. For example, using SSL or early TLS violates PCI DSS (Payment Card Industry Data Security Standard) compliance. It delivers the knowledge of what is being encrypted and what is not being encrypted on your network so you can confidently claim that your digital business is protected. This cryptographic assessment is displayed in Stealthwatch and can be exported via APIs to third-party tools for monitoring and auditing of encryption compliance (Figure 5).

Figure 5. Cryptographic assessment



Feature support

Cisco's newest networking equipment, starting with Cisco IOS® XE 16.6, will support an enhanced NetFlow with Encrypted Traffic Analytics capability.

1. Compatible Cisco equipment supporting enhanced NetFlow telemetry with Encrypted Traffic Analytics:
 - **Switches:** Cisco Catalyst® 9300 Series (starting with the Cisco IOS XE Software Release 16.6.1) and the 9400 Series (starting with the Cisco IOS XE Software Release 16.6.2)
 - **Routers:** Cisco ASR 1001-X, ASR 1002-X, ASR 1001-HX, ASR 1002-HX, ASR1000 RP2, ASR1000 RP3, ASR1000 ESP-40, 4221 ISR, 4321 ISR, 4331 ISR, 4351 ISR, 4431 ISR, 4451-X ISR, and ISR 1000 series routers, Cisco Integrated Services Virtual Router (ISRV), including the 5000 Enterprise Network Compute System, and Cisco Cloud Services Router (CSR) 1000V (starting with the Cisco IOS XE Software Release 16.6.2)
 - **Wireless controllers:** Cisco Catalyst 9800 Series (starting with Cisco IOS XE Software Release 16.10.1)
 - **Stealthwatch Flow Sensors:** Installed on a mirroring port or network tap to generate telemetry based on the observed traffic. Available as hardware or virtual appliances (starting with Stealthwatch Software Release 7.1).
2. Cisco Stealthwatch network traffic analysis solution with machine learning and behavioral modeling capabilities (starting with Stealthwatch Software Release 6.9.2) to analyze enhanced NetFlow for Encrypted Traffic Analytics.

Efficacy results and real-world deployment

Cisco engaged technology testing and certification firm [Miercom](#) to evaluate the efficacy and performance of Encrypted Traffic Analytics. The results were exceptional, earning Encrypted Traffic Analytics the Miercom Performance Verified certification.

You can read the [full report here](#). Here are a few highlights:

Cisco Encrypted Traffic Analytics showed as much as 36 percent faster rates of detection, finding 100 percent of threats in three hours

Miercom tested a variety of malware, both encrypted and unencrypted, along two network paths. Both paths utilized Cisco networking equipment and Stealthwatch, but one path had Encrypted Traffic Analytics enabled and the other did not. The malware tested included exploits such as Trojans, botnets, ransomware, and keyloggers, and more than two-thirds of these threats used encrypted communications. Without Encrypted Traffic Analytics, this malicious activity would have persisted undetected.

On the path using Encrypted Traffic Analytics, threats were detected up to 36 percent faster than the path that did not use Encrypted Traffic Analytics. Furthermore, all threats were detected in three hours. Faster, more complete detection is something your attackers do not want you achieve.

Encrypted Traffic Analytics detected 100 percent of malicious flows within three hours

Immediate detections—those in under five minutes—and learning capabilities over three hours showed impressive performance. In under five minutes, Encrypted Traffic Analytics detected nearly two-thirds of all malicious flows, almost double of the non-Encrypted Traffic Analytics path. Even with low volumes of 0 to 20 flows, Encrypted Traffic Analytics showed higher detection results, while 2000+ flows allowed for 100 percent detection.

Threat findings from the field

The technology continues to have great success in customer environments, and we can share the results of monitoring the [Cisco Live USA 2018](#), [Mobile World Congress 2018](#), and [Mobile World Congress 2019](#) conferences. The categories of threats that the technology has detected in these and other environments include, but are not limited to:

- Illicit cryptomining
- Android OS Trojans
- Ad Injectors
- SALITY malware
- Malware using SMB (Server Message Block) service discovery
- Potentially unwanted applications such as Tor and BitTorrent

“Autodesk’s security team is working closely with Cisco to evaluate and implement Encrypted Traffic Analytics (ETA) in our environment. I’m excited about the potential of ETA powered by Stealthwatch and the visibility it gives us into encrypted traffic anomalies without the overhead and legal implications of SSL decryption. We are already seeing some interesting findings from the solution and hope to fully incorporate the results and alerting into our normal detection and response workflows this year.”

Drew Miller, Lead Enterprise Security Architect, Autodesk

Conclusion

In summary, the network is now an even more advanced security sensor, capable of detecting threats in encrypted traffic. An infrastructure readied with Cisco Digital Network Architecture turns the network into an end-to-end sensor and enforcer that detects, contains, and prevents emerging, sophisticated security threats.

Appendix A

Data Elements Extracted by Encrypted Traffic Analytics.

Data element name	Description
IDP	The content of the first packet of this flow that contains actual payload data, starting at the beginning of the IP header.
SPLT	An array of LENGTH values followed by an array of INTERARRIVAL TIME values describing the first N packets of a flow that carry application payload. Each LENGTH is encoded as a 16-bit integer to form a 20-byte array. Immediately following this, each INTERARRIVAL TIME is encoded as a 16-bit integer to form another 20-byte array.
Byte distribution	A histogram giving the frequency of occurrence for each byte value or (range of values) in the first N bytes of application payload for a flow. Each “frequency of occurrence” is represented as a 16-bit integer.

Data element name	Description
TLS records	An array of LENGTH values, followed by an array of INTERARRIVAL TIME values, followed by an array of CONTENT TYPE values, followed by an array of HANDSHAKE TYPE values. These arrays describe the first N records of a TLS flow.
TLS record lengths	A sequence of record lengths for up to the first N records of a TLS flow.
TLS record times	A sequence of TLS interarrival times for up to the first N records of a TLS flow.
TLS content types	A sequence of ContentType values for up to the first N records of a TLS flow.
TLS handshake types	A sequence of HandshakeType values for up to the first N records of a TLS flow.
TLS cipher suites	A list of up to N cipher suites offered by the client, or selected by the server in a TLS flow.
TLS extensions	An array of LENGTH values followed by an array of EXTENSION TYPE values describing the TLS extensions observed in the Hello message for a TLS flow.
TLS extension lengths	A list of extension lengths for up to the first N TLS extensions observed in the TLS Hello message for a flow.
TLS extension types	A list of extension types for up to the first N TLS extensions observed in the TLS Hello message for a flow.
TLS version	The TLS version number observed in the TLS Hello message for a flow.
TLS key length	The length of the client key observed in the TLS ClientKeyExchange message.
TLS session ID	The session ID value observed (if any) in the TLS Hello message for a flow.
TLS random	The random value observed in the TLS Hello message for this flow.

References

1. Gartner, Predicts 2017: Network and Gateway Security, December 13, 2016
2. [Google Transparency Report](#).
3. [Let's Encrypt Stats](#).
4. [Identifying Encrypted Malware Traffic with Contextual Flow Data](#), Blake Anderson and David McGrew, AISEC '16