TX Protect

Overview

Version 3.20, 2025-01-10

Table of Contents

Introduction

The ThreatX™ Managed API and Application Protection platform is a Protection as a Service offering that protects your APIs and applications from the full breadth of Layer 7 security threats, including traditional OWASP attacks, bots, malicious automation, DDoS, and API-specific attacks.

This guide provides a detailed description of the ThreatX platform, its features, and how it interacts with your environment.

The ThreatX platform delivers complete protection and threat visibility by combining attacker behavior profiling, collective threat intelligence, and advanced analytics. The ThreatX managed service combines threat hunting with 24×7 access to the security experts in the ThreatX Security Operations Center (SOC), significantly reducing the operational overhead and maintenance burdens for enterprises.

Instead of relying solely on application baselines and matching known attack signatures, the ThreatX platform focuses on the attacker and suspicious behaviors. By tracking attacker movement through the cyber kill-chain, and combining and corroborating many indicators of suspicious activity, the ThreatX platform builds a progressive risk profile of an attack and threat intent. The result is deep visibility into varying attack types and techniques, confirmation that it legitimately needs to be blocked, and rapid identification of known and unknown vulnerabilities in the application environment.

Profiling Techniques

The ThreatX platform uses the following active and passive techniques to evaluate suspicious behaviors and deliver broad spectrum attack support.

Dynamic Application Profiling

The platform creates real-time profiles, or baselines, of acceptable application inputs, responses, usage, query times and more. As the application is modified and as more client engagements occur, the profile is dynamically enhanced to improve anomaly detection and reduce false positives.

TLS fingerprinting

Multiple interactive techniques test the IP/user to determine human vs machine and fingerprint key user meta-data to identify attacks coming from multiple IPs.

Cookie challenge

Challenge an offending entity with a cookie and attempt to fingerprint the user-agent.

DDoS protection

Advanced L7 DOS behavior detection and risk-based response include automated tarpit or rate limiting and active blocking.

Bot management

Combination of inputs and interrogation actions to identify automated users and malicious intent.

SQLI/XSS Analyzer

Deep analysis of input parameters to identify and test potentially malicious SQL injection or cross-site scripting attempts.

Entity profiling

Automated behavior recognition and dynamic rule generation to provide visibility on attack profiles and targeted application weaknesses. Several dimensions are considered, including intensity, TOR exit nodes, IP reputation, open proxy IP, error count over time, URI disbursement, geographic location, U.S. embargo status, and kill-chain state, which are combined with inputs from multiple context signatures.

Attack profiling and threat intelligence

A continuously expanding Data Warehouse of identified attack patterns, classifications and correlations from suspicious entities collected across all ThreatX protected web applications and customers. The analytics engine correlates the latest intelligence and behavior over time to create a highly accurate decision and response engine to reduce workloads for security teams.

Dynamic behavior rules

Determines attacker intent (based on application and entity profiling data) and generates behavior rules (not signatures) that can watch, temporarily block, permanently block, interrogate, or tarpit suspicious traffic. You can work with the ThreatX SOC to create custom behavior rules to meet the needs of your environment.

Caching and resource optimization

Static and dynamic caching along with advanced page optimization.

SSL/TLS inspection

Full decryption and re-encryption of requests for deep inspection.

> *You can use the data the platform gathers to better understand your sites, traffic usage, and behavior, and uncover potential issues such as exposure of sensitive data and vulnerabilities.*

---

s

Protection as a Service

The ThreatX platform is managed API and application protection that secures APIs and applications without the complexity of installing, configuring, and managing a massive security system. The ThreatX platform is hosted and managed by the ThreatX Security Operations Center (SOC).

The ThreatX SOC is responsible for the initial setup and maintenance of the ThreatX platform. More importantly, the SOC monitors your API traffic to proactively prevent malicious attacks and adjust the ThreatX platform accordingly. The SOC also watches industry wide traffic to understand potential threats and update the ThreatX rules before you and other clients are attacked.

Architecture

The following figure illustrates how the ThreatX platform is organized and how it relates to your environment.

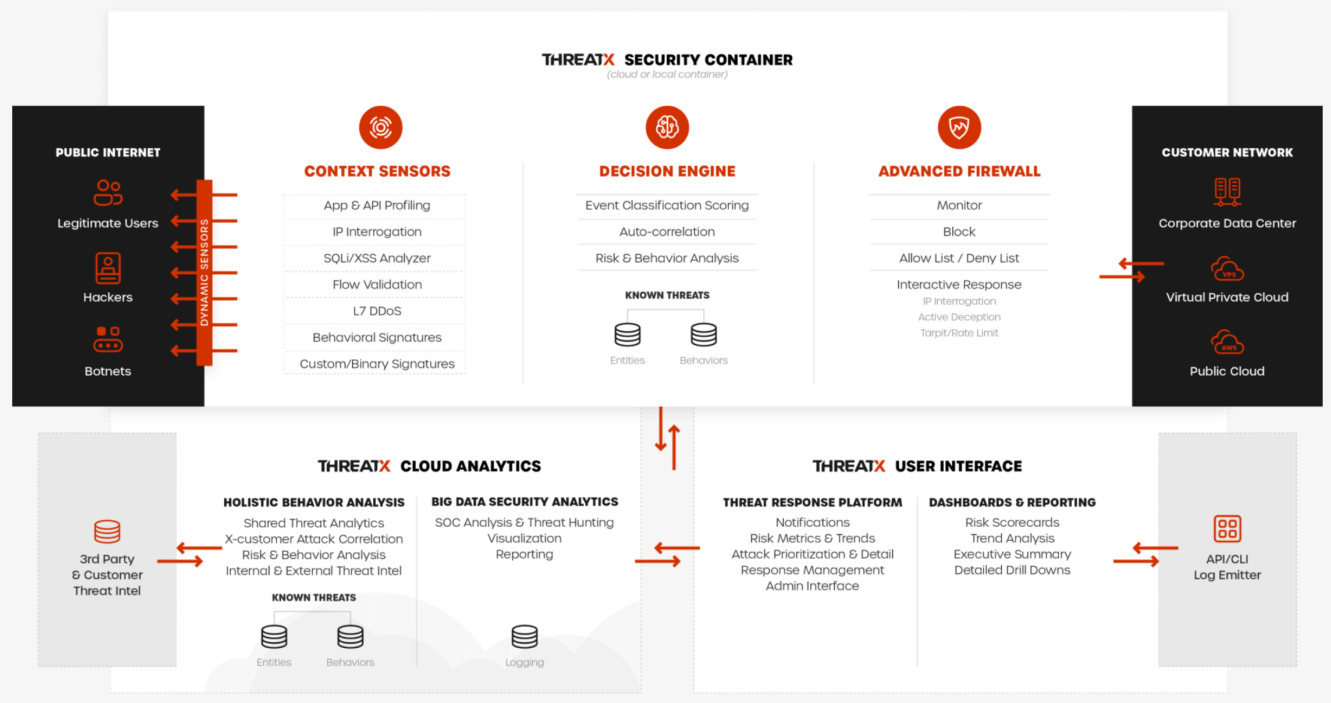

ThreatX Security Container

The ThreatX Security Container monitors your ingress API traffic and performs the initial risk analysis and response. *API traffic* is traffic that includes HTTP and HTTPS messages containing programmatic content sent between the site and client applications.

Context Sensor

The ThreatX Security Container includes one or more sensors. Sensors are decoupled from the ThreatX platform so they can be hosted in the ThreatX environment or deployed to your local environment.

A WAF sensor is a reverse proxy-based Web Application Firewall (WAF). The sensor monitors bi-directional web-based (HTTP and HTTPS) traffic flows for malicious and legitimate activity. The sensor inspects and cleanses user traffic that terminates on customer web applications or API endpoints.

The sensor intercepts traffic from web clients through the configuration of your DNS CNAME pointers. The sensor scrutinizes the traffic, and decides whether to allow, tarpit, interrogate, or block traffic directed at customer origin servers. Additionally, the sensor collects meta-data about web clients. The meta-data is then used to create entity profiles and feed the ThreatX Attacker-Centric Behavioral Risk model.

The risk model first profiles web client entities via a combination of IP address, TLS Fingerprint, and user agent information. It then scrutinizes entity behavior to detect risky behavior. A proprietary risk score is incremented and tracked for any given entity. The sensor blocks traffic from that entity if this risk score crosses a configurable threshold of risk tolerance. You have the option to overrule which entities are in the blocked or disallowed lists.

The sensors have local rules which they use to determine whether to pass, block or flag suspicious traffic. It also sends data about threats to the ThreatX platform for further analysis. The analytical engine updates the sensors with current security intelligence.

The sensor is based on the NGINX project, with modifications written in C++ and Rust. ThreatX backend

services are written in Rust to allow secure sub-millisecond transactional performance, even under load. The ThreatX web application is written in typescript (React, Angular).

Decision Engine

The Decision Engine analyses suspicious entities and techniques against known entities and techniques. An *entity* is a specific IP address or IP group. A suspicious entity is a threat. A *technique* is suspicious behavior tracked over time and across multiple applications. The platform uses these indicators to track malicious or suspicious users across many IP addresses as they use various evasion techniques and modify attack parameters.

Decision Engine Classification and Scores

Risk Score

Number between 0 and 100. It is associated with a single activity of a threat and is signature specific. The higher the score, the greater the risk.

Risk Level

Associated with all activities of a threat. The level is calculated from many inputs including Risk Score. One input is the kill chain model that classifies the attacker behavior and methods used to attempt to gain unauthorized access or control. The higher on the kill chain, the greater the severity of the threat.

Classification

Describes the type of attack which a rule assigns to a threat.

Advanced Firewall

The Advanced Firewall uses behavioral rules with advanced analyzers, IP interrogation techniques, and a combination of other detection capabilities working in parallel to determine the response. A ThreatX *rule* is a set of Boolean conditions that, when True, implement the rule's defined action and risk level. A True state is known as a match. The conditions are based on threat, request, or response attributes.

Rule Actions

Track

Begin or continue tracking a risk score for the offending entity, based on the risk assigned to this rule and other factors. This is the default and recommended action for most custom rules.

Block

Immediately block the request and track a risk score for the offending entity. Blocking rules are best used to stop known malicious behavior, "virtually patch" known vulnerabilities, and more.

Tarpit

Limit the speed at which the offending entity receives responses and track a risk score for the entity. Tarpit rules are best used to discourage scanning or scraping behavior without immediately blocking the traffic.

Interrogate

Challenge an offending entity with a cookie and attempt to fingerprint the user-agent. Interrogation allows a custom rule to explicitly invoke anti-bot mitigations for an entity.

The ThreatX blocking modes are designed to block malicious requests and deter suspicious entities from attacking your sites while allowing benign traffic and real users through.

Blocking Modes

Request Blocking

Blocks block malicious traffic at the request level when an attack such as SQL injection, XSS, or another malicious request is detected.

Manual Action Blocking

Users can manually block specific IP addresses. Users can also add entity IP addresses to the deny list for permanent blocking

Risk-Based Blocking

The ThreatX behavioral analytics engine evaluates each unique entity and blocks persistently malicious entities based on their behavior over time.The ThreatX behavioral analytics engine blocks persistently malicious threats when the threats' behavior surpasses the Risk-Based Blocking threshold. The analytics engine automatically places a threat on the permanent list after it is blocked three times.

You can leverage all three blocking modes when on-boarding a new application to help prevent potential false positives or unwanted impacts to your sites then gradually expand blocking levels.

Additionally, you can configure the platform to not allow traffic from embargoed countries and Tor Exit Nodes.

As needed, you can request the ThreatX SOC group to create new rules or modify any rule in the ThreatX platform to meet the specific needs and behavior of your environment.

ThreatX Cloud Analytics

The ThreatX Cloud Analytics performs an in-depth risk analysis and response, which is provided to the Security Container. The events are tracked in real-time and available to your SOC in the ThreatX User Interface.

The Cloud Analytics is a single analytics engine that receives feeds from different detection techniques.

Types of Cloud Analysis

Holistic behavior analysis

Uses shared threat intelligence with other organizations, both internally and externally, to maintain the latest threat intelligence. As shown in the architecture diagram, this includes HTTP custom header, also known as X-customer header, attack correlation. The analysis includes threat entities and threat behavior.

Big data security analytics

The ThreatX SOC analyzes and studies threats and behavior to improve rules that can prevent attacks. When the ThreatX SOC detects and analyzes new threat behavior for one customer, updated rules are deployed to all customers.

The ThreatX Cloud Analytics uses Attacker-Centric Behavioral Analytics (ACBA), which is an approach that identifies critical elements of an attack, responds to them before any damage is done, and maintains protection even if attackers change or obfuscate their attack pattern to avoid detection.

ACBA continuously monitors all users as they interact with an application or API. It looks for key indicators of suspicious behavior and tracks risk over time and across multiple applications. It observes risky behavior that is not obviously malicious on the basis of a single request but exhibits a pattern of risky

behavior known to be associated with malicious actors. This data is correlated in the Actor activity logs and can be displayed in the ThreatX Dashboard. The ThreatX Cloud Analytics also provides a visualization of the threats in the form of charts and graphs to the ThreatX Dashboard for your analysts.

You can import your own threat intelligence. The ThreatX platform can use deny lists from threat intelligence solutions by integrating with SOAR solutions or by scripting using the ThreatX API.

Log Emitter

The ThreatX Log Emitter allows efficient and secure, real-time export of event logs from the ThreatX platform to your log receiver and SIEM. The details contained in these event logs can be leveraged in your investigations and used to trigger events in your chosen log management solution. Logs are pushed in JSON lines format over a TCP connection that is encrypted (and optionally authenticated) via TLS.

Log Emitter forwards full details for all:

- Block Events
- Match Events
- Audit Events

In the event a Log Emitter subscription becomes suspended, the Log Emitter service queues your logs for delivery upon successful re-connection, and periodically attempts to re-establish a connection.

Tenants and Channels

Your ThreatX platform configuration is organized by at least one tenant, where a tenant is an organizational unit containing your users and sites. Your users can view protected sites, attack heuristics, real-time data, and other configuration information in the ThreatX dashboard within the tenant.

Alternatively, you can have your ThreatX platform organized by channels, where a channel can contain multiple tenants.

Analysis

The ThreatX platform analyzes HTTP traffic then extracts identifying metadata, including IP address, user agent, TLS fingerprint, and other characteristics to create a profile and identifier for each attacker.

The data is presented within the Dashboard using various pages and tables with a special emphasis on key attributes to further help identify trends and patterns. The tables provide different perspectives of how the data relates to each other, which can help in your analysis.

Metrics

Sites

The ThreatX platform displays all the sites under the ThreatX protection, the API profile for each site, and every endpoint for each site.

Threats

The tables in the ThreatX dashboard provide analytical data about the threat, including status, IP address, last seen, location, and attack class. For location, the Threat Map provides an interactive map that identifies how many unique attackers are acting from each geographical location.

Risk Score associated with a single activity of a threat

The ThreatX platform displays Risk Score as a number between 0 and 100. The higher the score, the greater the risk.

Risk Level associated with all activities of a threat

The level is calculated from many inputs including Risk Score and kill chain model that classifies the attacker behavior and methods used to attempt to gain unauthorized access or control.

Rule activity

Number of times the rule was matched by requests.

Time range

Allows the user to view data within a specific time frame.

You can use the data for various analytical tasks, including:

- Review traffic trends including unexpected usage patterns.
- Monitor threats, including those that matched rules and were blocked.
- Discover if sites contain sensitive data or vulnerabilities.
- Upload schemas for your endpoints and determine if there are any discrepancies between the schema and observed behavior.
- Verify that all expected sites are included.

Additional Features

You can customize the ThreatX platform to meet the needs of your environment. The following sections describe the features that you can add or modify.

Risk-Based Blocking feature

With the ThreatX Risk-Based Blocking feature, the ThreatX platform can add a threat automatically to the Blacklist or Blocklist based on the threat's behavior. The ThreatX behavioral analytics engine blocks persistently malicious threats when the threats' behavior surpasses the Risk-Based Blocking threshold. The analytics engine automatically places a threat on the permanent Blacklist after it is blocked three times.

Sensitive Data

The ThreatX Sensitive Data feature monitors API responses to detect various sensitive data. Sensitive data includes authentication credentials, credit card (PCI-DSS), and Personally Identifiable Information (PII).

The ThreatX platform reports only sensitive data that is in plain text. It does not report partial or obfuscated data, such as *--1234. The ThreatX platform does not correlate sensitive data with rules or threats or store sensitive data due to security and compliance reasons.

Edge Caching

Edge Caching is available to ThreatX customers who wish to take advantage of the performance and speed improvements commonly associated with caching, but who do not have a caching solution in place. The benefits of Edge Caching include:

- Faster page load times for end-users.
- Lower latency.

- Increased load capacity and reduced application server load.
- Better ratings from search engines such as Google.

By default, ThreatX Edge Caching follows Cache-Control headers defined by the origin servers.

Supported Edge-Caching

Static Caching

Caches static elements such as images, CSS and JavaScript. Static caching does not store HTML pages and as a result does not enhance performance if the origin server becomes unresponsive.

Dynamic Caching

Provides a higher level of performance, allowing caching and optimization of dynamic content. In some cases, cached content can be delivered even if the origin servers are unresponsive. The ThreatX platform caches all responses to requests made with HTTP GET, and HEAD methods. To avoid caching dynamic pages that are rarely accessed, ThreatX sensors cache dynamic pages only after they are requested at least three times. Subsequent requests are served from the cache until the cache expiration defined in the Cache-Control occurs, or for 30 minutes for responses where the expiration is not defined. Dynamic caching requires an add-on license.

Rate Limiting

By default, the ThreatX platform offers rate limiting capability by the rules in the common rule set. For example, one rule, 10 404s in 10s, assigns risk to an entity that receives more than ten 404 responses within 10 seconds.

AWS Shield Standard is also deployed by default on all AWS hosted infrastructure to assist in mitigating DDoS attacks.

Additionally, the ThreatX SOC can create custom rate limiting rules tailored for your environment. A typical use of this would be to assign risk to entities that fail logins at a login endpoint. These rate limiting rules are very customizable, including the timings (# of requests/time). These rules can be applied across the entire tenant, a specific site or group of sites, or a single endpoint. The match criteria also have a very wide range of options such as Response Code, Request Method, Source Country/ASN, and Args.

Site Certificate Management

The ThreatX platform can manage the SLL certificates presented to your site's visitors with Let's Encrypt. The Let's Encrypt integration allows you to offload the overhead and management commonly associated with managing SSL/TLS/TLS certificates while ensuring that an expired certificate is never presented to your site's visitors.

Privacy

ThreatX is committed to privacy and security of our customers' data. The ThreatX platform collects and stores as little corporate data as possible while maintaining the highest level of security and efficacy for the sites we protect. ThreatX has an AICPA certified auditor-issued SOC 2 Type 2 Report covering Security and Availability trust services criteria, including the following:

- Locate and remove or redact specified confidential information as required.
- Regularly and systematically destroy, erase, or make anonymous, confidential information that is no longer required for the purposes identified in its confidentiality commitments or system requirements.

- Erase or destroy records in accordance with the retention policies, regardless of the method of storage.
- Dispose of original archived, backed up, and ad hoc or personal copies of records in accordance with its destruction policies.

The ThreatX platform does not install an agent on servers or workloads, and has no privileged access to origin servers, API endpoints, or any supporting infrastructure related to the web applications the platform protects. The platform sits inline, scrutinizes HTTP and HTTPS requests, and allows or blocks traffic based on attributes inherent in the HTTP request. As such, the platform does not directly interact with customer intellectual property.

Furthermore, the ThreatX Web Application Firewall can be used to satisfy PCI-DSS Requirement 6.6 when deployed within a customer's PCI environment. While the sensors do not store or transmit cardholder data (PANs, CVVs, etc.), maintaining effective security controls is the responsibility of the customer and should be validated by a QSA.

You can find more information about our physical and logical security posture, our controls, and our SOC 2 Type 2 standing on our website. The current report and bridge letter are available to customers who require it for compliance purposes.

Data Gathered

The ThreatX platform gathers the following backend data (summarized):

- Source IP
- User-Agent header
- Request Method (GET/POST/PUT)
- Request Domain (for example, site.com)
- Request Path (/request/path)
- TLS Fingerprint
- ThreatX metadata about security rule matches

The ThreatX sensor does not inspect response data.

Sensitive data is retained only if necessary for business purposes. This includes data required for processing transactions, supporting customers and business functions, and supporting current or historical event analysis. ThreatX requires transaction details to be available in databases and in log format to support customer requests and analysis.

The ThreatX SOC retains the data for 90 days.

Data Redactions

Specific portions of the request are automatically redacted and never sent to the backend, including tokens, credentials, and known patterns such as credit card and social security numbers. This redaction is applied to fields and URL encoded forms.

The remaining sanitized data is reduced to metadata before being sent to the ThreatX platform for analysis, and or visualization to customer security administrators.

Usernames are not automatically redacted, as this data is often critical to security analytics and forensics, for instance in identifying account takeover (ATO) attacks or login rotation.

The ThreatX Soc can help with custom redactions on a case-by-case basis. To scrub specific data, you can contact the Threat SOC.