



WHITEPAPER

Threat Intelligence Platform

How aiTIP360 collects, enriches, scores and delivers intelligence on IP addresses, domains and URLs.

Seceon aiTIP360

Ver-3.0

ABOUT THIS PAPER

A blocklist tells you someone reported an address. It doesn't tell you whether to block it, who owns it, what else lives next to it, or why it showed up today.

This paper walks through how aiTIP360 builds that picture: where the data comes from, what the platform adds to it, how every score is put together, and how the result reaches the tools your team already uses. It's written for security engineers, SOC leads and architects evaluating a threat intelligence platform, and it assumes no prior knowledge of aiTIP360.

If you only have five minutes, read the executive summary on the next page.

CONTENTS

Executive summary.....	3
<i>The whole paper in half a page.</i>	
1. Why lists aren't enough.....	3
<i>Three problems with feeds that buying another feed won't fix.</i>	
2. What's in the database.....	4
<i>Known-bad, known-good, and the context in between, all in one local store.</i>	
3. Web / Domain reputation.....	4
<i>What comes back when an analyst looks up a domain or URL.</i>	
4. IP reputation.....	6
<i>What comes back when an analyst looks up an IP address.</i>	
5. Neighbours and relationships.....	7
<i>The company an indicator keeps: shared hosts, certificates and nameservers.</i>	
6. How the score is built.....	7

Named rules, five verdict bands, and how scores fade.

7. Categories and threat context..... 8

Threat types, malware families, actors, campaigns and ATT&CK.

8. Keeping false positives down..... 8

The mechanisms that hold scores down when they shouldn't rise.

9. Getting it into your other tools..... 9

TAXII and the REST API, in both directions.

10. Air-gapped and on-premises deployment..... 9

Delivered as a Docker package on a dedicated server, including air-gapped sites.

Executive summary

Most security teams don't lack threat intelligence. A typical SOC already runs a commercial subscription, a sector ISAC, national CERT advisories and a handful of community lists. What's missing is a way to turn all of it into a decision someone can stand behind. IP and web / domain reputation is where that decision usually starts, and it's the focus of this paper.

aiTIP360 keeps one local database of what's known about IPs, domains, URLs and file hashes: what's been reported as bad, what's known to be good, and the context in between. Every indicator gets a score from 0 to 100, one of five verdicts, and a stored record of every rule that contributed. Any verdict can be opened and explained to an analyst deciding whether to escalate, an auditor testing a control, or a customer asking why their traffic was blocked.

Agreement, not volume. The score rewards independent sources that agree, weighted by each source's reliability rating. In the Prism composite, one loud report can't push an indicator into the blocking band on its own.

Context the feeds don't carry. Registration data, location, network ownership, DNS and passive DNS history, certificates, server behaviour and neighbouring infrastructure are attached by the platform itself, whatever a feed claims.

Noise control built in. Popularity caps, warning lists, conflict penalties, rate-limited score drops and per-type decay are part of the scoring model, not filters bolted on afterwards.

Your data centre, your data. Seceon delivers it as a package that runs as Docker containers on a dedicated server in your data centre, including in air-gapped networks. It answers lookups from its own database, and shares indicators over TAXII (2.1 and 1.1) and a REST API.

1. Why lists aren't enough

The trouble usually starts when two feeds disagree, or when one of them is simply old. None of the problems below is fixed by buying another feed.

Single-source risk. When one vendor reports an address and nothing else does, the sighting looks authoritative because it's the only thing on the screen. It might be right. It might also be a misclassified scanner, a shared hosting address that briefly carried a bad site, or a home broadband address a carrier recycled three weeks ago.

Data that ages faster than it's used. Phishing domains and command-and-control URLs can be gone in a few days, but they sit in feed archives for months. Static lists pile up. They rarely retire anything.

Verdicts nobody can defend. A bare number is no help in the conversations that matter. It's the 2 a.m. question: would you be happy blocking this with nobody awake to check? You need to know how many sources agree, how fresh the sighting is, who registered the domain, where the IP lives and what the server says when you knock. And afterwards you'll want to show your working.

The same goes for the wider picture of a domain. Experienced analysts weigh its DNS intelligence as a whole, including the MX, SPF, DKIM and TXT records that show how its mail is configured; the content classification of what it serves; whether it is a lookalike or similar domain imitating a brand it has no connection to; and whether it belongs to APT infrastructure tied to a named actor or campaign rather than commodity crimeware. Reputation is the sum of all of that, not a single list entry.

2. What's in the database

aiTIP360 holds three kinds of knowledge side by side, in one database on the aiTIP360 server in your data centre.

What's known to be bad. Indicators from threat feeds, TAXII collections, MISP instances and file imports, deduplicated and linked to every source that reported them, so an address reported by six sources is one record with six sightings. More than a hundred feeds are catalogued. Each gets a licence review, and about half are cleared and switched on out of the box.

What's known to be good. The 100,000 most popular domains, refreshed weekly from the Tranco ranking, plus warning lists of benign IPs, domains and URLs that your team keeps by hand or imports in MISP format.

What's known about everything in between. Ownership, location, DNS history, certificates and server behaviour, gathered by enrichment workers and kept with timestamps so you can see how an indicator changed over time.

It handles IPv4 and IPv6 addresses and CIDR blocks, domains, URLs, and MD5, SHA-1, SHA-256 and SHA-512 hashes. A lookup for a single address also checks the netblocks that cover it, so an IP inside a listed hijacked range is caught even if nobody reported that exact address. Feeds refresh on their own schedules, which you can set per feed down to every 5 minutes.

1 INGEST	2 NORMALIZE	3 ENRICH	4 SCORE	5 DELIVER
Feeds, TAXII, MISP, file import	Deduplicate, link sources	Context workers in parallel	Named rules, evidence kept	TAXII 2.1 / 1.1, REST API

Enrichment runs in parallel queues, so a burst from one busy feed doesn't hold up lookups, and it re-runs on its own as evidence gets older. Each kind of lookup is a self-contained connector that declares which indicator types it handles and feeds its findings into the scoring rules, so new signals are added as connectors without touching the scoring engine, and each one can be switched on or off per deployment.

3. Web / Domain reputation

Web indicators turn over faster than IP addresses, so web / domain reputation leans on signals that expose newly built infrastructure, not only on reporting history. Say an analyst pastes a domain from a suspicious email into aiTIP360. Here's what comes back on the intel card, and where each piece comes from.

SIGNAL	WHAT THE ANALYST LEARNS
Who reported it	Every feed, TAXII collection and analyst submission that listed the domain or URL, and how much each source is trusted.
Registration	Registrar, registrant organization, abuse contact, status and expiry from RDAP, with WHOIS history.
DNS intelligence	Current A, AAAA, MX and NS records, plus passive DNS history with first-seen and last-seen dates and observation counts. A missing MX record counts against the domain.

SIGNAL	WHAT THE ANALYST LEARNS
Certificates	Every certificate from transparency logs, names that share a certificate, and how recently the newest one was issued.
Related hosts and domains	Other domains that have resolved to the same IP, domains on the same certificate or nameserver, and hosts reached by redirect.
Server response	Status, final URL, content type, Server header, page title and the presented TLS certificate. Login-lure titles, name mismatches and executable downloads raise the score.
Redirects	Shortened and redirecting links unwound, with a change of host flagged as evasion.
Popularity	Whether the domain is in the top-100,000 list, which caps the score.
Warning-list status	Whether your team or an imported MISP warning list has marked it benign, which holds the score down and keeps it out of exports.
Classification	Each domain and URL is classified by threat category, such as phishing, malware delivery or command-and-control, with malware family and campaign tags where a source provides them. The live probe also records the page title and content type.
ATT&CK context	MITRE ATT&CK techniques linked to the indicator through tags and STIX relationships.
Sightings	First and last seen, with sightings broken down by feeds, SIEM, EDR, IDS and analysts.

Four of these deserve a closer look.

Who registered it

Registration data comes from RDAP, the modern replacement for WHOIS: registrar, registrant organization where it's published, abuse contact, domain status and expiry date. Every lookup is stored, so the card shows WHOIS history too.

What it resolves to, now and before

We query current A, AAAA, MX and NS records. A domain with no MX record is rarely a real business, and that counts against it. Each resolution is also written to a passive DNS store with first-seen and last-seen dates and an observation count, so the card shows the addresses and nameservers a domain has used over time, not only today.

Its certificates

Certificate transparency logs list every certificate issued for the domain, with issuer, validity dates, serial number, how many exist, and other names that share a certificate. A certificate issued in the last 7 days counts for more than one issued in the last 30, because fresh certificates on fresh infrastructure are how phishing kits usually start.

What the server says

We make one careful HTTP request and follow up to five redirects. The card records the status code, final URL, content type, Server header and page title, plus the TLS certificate the server actually presented: subject, issuer, expiry and SHA-256 fingerprint. A certificate whose name doesn't match the host, a page title that reads like a login lure, or an executable served where a web page should be will each raise the score. Shortened and redirecting links are unwound separately, and a change of host along the chain is treated as an evasion signal.

How popular it is

If the domain is in the top-100,000 list, the card says so and the score is capped. One stray report shouldn't be enough to make a major site a block candidate.

4. IP reputation

An IP address tells a different story, so the card looks different. The context matters most when the sources are thin or disagree.

SIGNAL	WHAT THE ANALYST LEARNS
Location	Country, city and coordinates from a local MaxMind GeoLite2 database. No network call is needed.
Network and ownership	ASN and organization, plus RDAP registry data: network name, owner, the CIDR it belongs to and when that range was allocated. Ranges allocated in the last 90 days get extra weight.
Kind of network	Large hosting providers and CDNs are recognized, and known internet scanners are classified as scanners rather than attackers.
Reverse DNS	The PTR name, checked against the naming patterns of legitimate cloud and CDN providers. Generic or throwaway names stand out.
Who reported it	Every source that listed the address and how much each is trusted, and whether it's on a high-confidence list such as a spam-source or hijacked-netblock list or only a secondary one.
Anonymizers	Tor exits are flagged as context, not as proof of anything.

Put together, these separate the two cases that matter. An address several sources agree on, in a freshly allocated range at a bulk host with a generic PTR name, is a different thing from an address reported once inside a long-established enterprise netblock. In the second case the report is more likely to be the anomaly. If you hold licences for AbuseIPDB, Shodan or VirusTotal, optional connectors add their report categories, open ports and hash verdicts to the same card.

5. Neighbours and relationships

Bad infrastructure rarely lives alone, so when the platform scores an indicator it also looks at the company it keeps.

- Other indicators on the same ASN, the same ISP and the same /24 network
- Domains that have resolved to the same IP, from passive DNS history
- Related domains that share a certificate or a nameserver, and hosts linked by a redirect
- Malware families, campaigns and threat actors that a source has tied to the indicator

The relationship graph expands two hops out and counts how many neighbours are themselves malicious. Analysts can pivot from any node to related IPs, domains, URLs and hashes without writing a query.

6. How the score is built

The score is a running total of named rules. We don't hide a model behind it. Corroboration across sources counts most, then recency, then each source's reliability rating. Smaller adjustments come from severity, repeat reports, sightings in the last 24 hours, and bonuses for government-sourced reports and for indicators seen in more than one source category. Enrichment findings from sections 3 to 5 come after that, and every point is stored with the rule, the source and the evidence text.

IP addresses have one difference. If your deployment holds an AbuseIPDB licence or a Cyber Threat Alliance (CTA) membership, the headline score for an IP follows the higher of those two vendor scores, and the Prism composite with its full rule breakdown sits alongside it on the card. Without those sources, IPs are scored by Prism like everything else.

SCORE	VERDICT	WHAT MOST TEAMS DO WITH IT
90 and above	Very malicious	Block automatically, if your policy allows
65 and above	Malicious	Shared with SIEM and enforcement tools
25 and above	Suspicious	Analyst queue, SIEM correlation
5 and above	Unusual	Watch it
Below 5	No evidence	Retired, or suppressed as a false positive

How fast scores fade

Each indicator type has its own clock, which restarts every time the indicator is seen again. Score history is kept, so you can see when something peaked and how it cooled.

TYPE	SCORE REACHES ZERO AFTER
DGA domain	3 days
URL, domain	7 days
IP address	30 days
File hash	180 days

7. Categories and threat context

A score tells you how worried to be. The category tells you what you're dealing with. Indicators carry a threat type drawn from over twenty that the platform recognizes, among them botnet C2, phishing, ransomware, banking trojans, malware infrastructure, spam sources, scanners, brute-force sources, anonymizers, compromised hosts and web attacks. Severity and tags travel with them.

Where a source provides it, you also get the malware family, whether a C2 server is online or offline, the port it was seen on, and links to named threat actors and campaigns. Indicators tied to a named actor or campaign are marked on the card, so targeted-attack infrastructure stands out from commodity crimeware. MITRE ATT&CK technique IDs are picked up from tags and STIX relationships, shown on the card, and can be exported as an ATT&CK Navigator layer covering the whole corpus.

8. Keeping false positives down

A reputation service that cries wolf gets switched off. That's why a good part of the engine is about holding scores back, not pushing them up.

Popularity caps. Domains on the top-site list are capped, as described above.

Conflict penalty. When a blocklist and the popularity list disagree about the same domain, the score goes down instead of siding with the louder source.

Warning lists. Anything on your warning lists is forced to a low score and left out of exports.

Rate-limited drops. A score can't collapse in one step when an indicator is seen again, so verdicts don't jump around between refreshes. That matters when firewalls consume them automatically.

Automatic retirement. Stale indicators retire on their own instead of piling up.

In the Prism composite, a single feed can't push an indicator into the 90+ band. That verdict needs independent enrichment or graph evidence on top of what the feeds say.

9. Getting it into your other tools

Reputation is only useful where decisions get made, so aiTIP360 is built to hand its data on. Everything goes through two standard interfaces.

ROUTE	WHAT IT'S FOR
TAXII 2.1 and TAXII 1.1	A TAXII server publishes a collection of scored indicators that any TAXII client, such as a SIEM, SOAR or firewall manager, can poll on its own schedule. A TAXII client pulls from your enterprise, sector and government sources. TAXII 2.1 with STIX 2.1 is the default. TAXII 1.1 with STIX 1.x is available for older peers and is enabled per deployment.
REST API	Single lookups for any IP, domain, URL or hash, bulk enrichment of up to 1,000 indicators per call for SOAR playbooks, and indicator import.

Data comes in the same way, over TAXII or the REST API, alongside the built-in feeds and file import. Your own commercial, sector or government feeds sit in the same database and carry their own reliability weight, so trusted intelligence can count for more than a general public list.

10. Air-gapped and on-premises deployment

Seceon delivers aiTIP360 as a package that runs as Docker containers on a dedicated server in your data centre. We specify the server for your indicator volume and retention. The database, scoring engine, relationship graph, passive DNS history and geolocation are all local. If the internet link drops, lookups and scoring carry on.

Feeds don't need a direct internet connection either. They can come from a TAXII server or MISP instance inside your network, or as files carried in through your normal transfer process and imported.

High-availability layouts are planned the same way, as part of the deployment design.

aiTIP360 supports deployment in air-gapped networks. Every air-gapped environment has its own transfer process and security rules, so we plan these deployments with each customer. For details, please reach out to Seceon to discuss your environment.

See it on your own data

Send us 500 or so indicators you already know the answer for, malicious and benign. We'll run them through aiTIP360 and show you every verdict and the evidence behind it, so you can judge it against cases your team has already worked.

DOCUMENT HISTORY

RELEASE	DATE	SUMMARY
—	March 2025	Initial release
Ver-3.0	October 2026	Current release

ABOUT SECEON

Seceon's Open Threat Management (OTM) platform brings detection, investigation and response together in a single data pipeline, for enterprises and for the managed service providers that serve them. It's used by more than 9,800 customers and 850 partners worldwide, across financial services, government, healthcare, manufacturing and critical infrastructure.



Seceon, Inc. • 238 Littleton Rd, Suite #200, Westford, MA 01886 • +1-978-923-0040 • www.seceon.com