ESIP

Exposure Signal Intelligence Platform

Version: 1.6 · ZenzizenSec · June 2026 · Classification: External

What it is, why it exists, and how it works. A quick read for security practitioners, product teams, and anyone evaluating ESIP. Technical details live in the full specification suite.

The Problem

Security teams are not failing to find vulnerabilities. They are failing to decide which ones matter.

Every tool produces more CVEs. Every feed produces more data. But none of them clearly answer the question practitioners actually need:

> ***What is becoming dangerous right now?***

Most prioritisation tools answer: what vulnerabilities exist? They score CVEs, rank them, and produce a list. Those lists are useful, but they are static snapshots. They do not tell you what changed, what attackers are actively doing, or what became dangerous between Monday and Tuesday.

That gap between static scores and dynamic threat reality is what ESIP is designed to close.

What ESIP Is

ESIP answers one question: what is becoming dangerous right now? It does not describe vulnerabilities. It describes how they are changing.

ESIP is a global exposure signal feed. It ingests publicly available threat intelligence data, detects meaningful changes in the exploit landscape, and produces structured signals per CVE class.

> ***An EPSS score of 0.82 is data. An EPSS score that jumped from 0.10 to 0.82 overnight while a weaponized exploit appeared in a public tool is a signal.***

ESIP is not a scanner. It does not connect to your environment, scan your assets, or know anything about your infrastructure. It watches the global threat landscape and tells you what is moving.

You bring your asset inventory. You apply ESIP signals to your environment. Together, those two things answer the question practitioners actually need answered: which of our exposures matters most right now?

> **Zero footprint in your environment** ESIP requires no agent installation, no network credentials, no cloud connector, and no access to your infrastructure. All signals are derived from globally observable public data. ESIP does not collect, process, or store any data about your assets, vulnerabilities, or environment. There is nothing to harden and no infrastructure attack surface to manage. The privacy review surface is limited to API key handling and outbound HTTPS: no asset, vulnerability, or environmental data leaves your network. Authentication credentials and usage metadata are retained for security audit and rate-limit enforcement.

How ESIP Is Different

Most tools describe vulnerabilities. ESIP describes how they are changing.

Tool	What it tells you	Limitation
CVSS	*How severe a vulnerability could be*	*Static. Does not change with the threat landscape.*
EPSS	*Likelihood of exploitation in 30 days*	*Probabilistic. Does not indicate active exploitation or evidence.*

Tool	What it tells you	Limitation
KEV	*Known exploited vulnerabilities*	*Binary. In the list or not; no gradation, no trend.*
Threat intel	*Who is attacking*	*Not CVE-focused. Hard to operationalise per vulnerability.*
ESIP	*What is becoming dangerous right now, and whether it is getting worse*	*Requires your asset context to complete the picture.*

What ESIP Watches

ESIP continuously monitors 13 globally observable data sources, including vulnerability databases, exploit repositories, government exploitation catalogs, and public threat intelligence feeds. All publicly available. No customer data involved.

No single source is trusted above all others. ESIP detects meaningful changes across multiple sources and requires corroboration before producing a signal. The strength of a signal reflects both the evidence present and the trust tier of the sources that contributed it.

Source type	Example sources	What it contributes
Government exploitation catalogs	*CISA KEV, ENISA EUVD*	*Confirmed real-world exploitation. The highest-trust signal in the system.*
Exploit probability	*FIRST EPSS*	*Daily probability score. ESIP watches the delta: is this rising fast?*
Weaponized exploit tools	*Metasploit*	*Production-ready exploit modules. Accessible to any attacker, not just skilled ones.*
Exploit references	*ExploitDB*	*Published exploit code. A new entry means someone built a working exploit.*
Proof-of-concept code	*GitHub*	*Early-stage capability. Lower trust; filtered for noise before contributing to signals.*
CVE metadata	*NVD*	*Reference layer: affected products, severity scores, patch availability.*

Signal Freshness

New evidence is reflected in signals within hours of source publication, not days. For procurement teams asking how current the data is:

Source	Target latency from source publication
CISA KEV (exploitation confirmed)	*Target: within 4 hours*
ENISA EUVD (exploitation confirmed)	*Target: within 36 hours*
FIRST EPSS (exploit probability)	*Target: within 26 hours of daily publication*

Source	Target latency from source publication
Metasploit and ExploitDB (weaponized exploits)	*Target: within 12 hours*
GitHub PoC (proof-of-concept)	*Target: within 18 hours*
NVD (CVE metadata, patch status)	*Target: within 36 hours*

When any source is delayed or degraded, every API response carries a degraded_sources count and a cache_serving_stale flag so consumers know exactly which signals may be behind.

What ESIP Produces

ESIP does not return raw data. Each signal is a structured verdict that tells you how dangerous a CVE is right now, how that risk is evolving, how confident the system is, and what attack stage the CVE enables.

Signal Verdict

Dimension	Values	What It Means
Lifecycle Stage	*Emerging · Active · Confirmed · Declining · Dormant · Remediation Available*	*Where this CVE sits in its threat lifecycle. Confirmed means real-world exploitation is happening. Emerging means early warning indicators are moving.*
Severity Score	*0 to 100*	*Current evidence-derived threat severity for this exposure class. Reflects all globally observable evidence. Not a static CVSS score. Recomputed as evidence changes.*
Attack Relevance	*Entry · Execution · Escalation · Movement · Impact · Unclassified*	*What stage of the attack kill-chain this CVE enables. Drives control-mapping decisions.*
Confidence	*Low · Medium · High · Verified*	*How reliable is the evidence behind this verdict. Verified means government-confirmed real-world exploitation.*

> **Every verdict is explainable** ESIP produces no black-box scores. Every signal carries a full evidence payload: which sources contributed, what each source observed, when each observation was made, and which scoring constraints were applied. A security team can justify every prioritisation decision to a CISO or auditor without saying 'the tool said so.' For organizations subject to regulatory scrutiny, explainability is not optional.

> **Deterministic by design** Every ESIP verdict is deterministic, explainable, versioned, and evidence-backed. The same evidence always produces the same output. Scores are replayable: any past verdict can be reconstructed from its source observations. This is not a property of the UI or the API: it is an architectural guarantee of the signal engine.

Temporal Context

Schema v1.4 adds Temporal Context as a first-class output layer. It answers the questions a static score cannot.

> **The score appears flat but the threat is intensifying** A CVE held at Critical by a government exploitation floor can simultaneously show new weaponized exploit activity, multiple converging sources, and high evidence velocity. The score does not move. The threat is getting worse. Temporal Context is the layer that surfaces this.

Field	Values	What It Means
current_trend	*Escalating · Stable · Declining*	*Is the threat picture getting worse, holding, or fading? Derived from evidence activity, not from score movement. The most important operational signal in the temporal layer.*
trend_drivers	*List of specific driver events*	*Why current_trend has the value it does. Examples: new_weaponized_source, new_verified_exploitation, multiple_sources_converging, evidence_base_weakening. Typically populated when Escalating or Declining.*
signal_velocity	*Low · Medium · High*	*How fast evidence is accumulating. High means three or more material changes in the last 14 days. This is NOT a severity indicator: a High-velocity CVE can have a low severity score.*
latest_material_change	*Signal type and date*	*When did something meaningful last happen? The most recent exploitation*

Field	Values	What It Means
		confirmation, weaponized exploit, or significant EPSS movement.
threat_age_days	*Integer*	*How long has this CVE been active in the global threat landscape? Useful for distinguishing new emergent threats from long-running confirmed ones.*

Temporal Context is computed server-side from evidence history. It is deterministic, versioned, and auditable: the same evidence always produces the same trend classification. It is not generated by AI or editorial judgment.

Compound Risk

Compound Risk fires when multiple independent Tier 1 signals are simultaneously active for the same CVE. It is not a simple count of signals; it measures independent corroboration across distinct sources and organizations.

Strength Tier	Condition	What It Means
Weak	*2 signals from the same source*	*Same organization detecting the same CVE twice. Corroboration is limited.*
Moderate	*2 signals from independent sources*	*Two separate organizations have detected exploitation evidence. Confidence is materially elevated.*
Strong	*3+ signals from 2+ independent sources*	*Multiple independent organizations corroborating the same CVE simultaneously. Treat as highest-confidence evidence.*

> **Why Compound Risk matters** A CVE with a Metasploit module AND a CISA KEV entry AND an ExploitDB weaponized entry is fundamentally different from a CVE with only one of those. Compound Risk surfaces that difference explicitly. No manual cross-referencing required.

Free Tier vs Commercial Tier

Free Tier	Commercial Tier
Observable evidence: KEV inclusion, exploit availability, ATT&CK technique mappings, detection artifacts	*Full signal verdict: lifecycle stage, severity score, attack relevance, confidence*
Static snapshot, 24-hour lag	*Real-time signal freshness (within hours of source updates)*
Active and Confirmed lifecycle stages only	*All six lifecycle stages*
No compound risk signals	*Compound risk with strength tier*
No temporal context	*Full Temporal Context: current_trend, trend_drivers, signal_velocity, threat age*
GitHub snapshot + free public API	*Commercial REST API with SLA*

The free tier tells you what is happening. The commercial tier tells you what matters most and how the threat is changing.

A Signal in Plain English

Here is what a complete ESIP commercial signal looks like for two real CVEs.

CVE-2021-44228 (Log4Shell): Long-running confirmed threat

Field	Value
Lifecycle Stage	*Confirmed*
Severity Band	*Critical (score: 94)*
Attack Relevance	*Entry (T1190, Initial Access)*
Confidence	*Verified*
current_trend	*Stable*
trend_drivers	*[ ] (empty: no new evidence in the last 14 days)*
signal_velocity	*Low*
threat_age_days	*1,640+*
What changed	*In CISA KEV since Dec 2021. Metasploit module active. EPSS 0.97. Evidence from 4 independent Verified/High trust sources.*

Reading: Log4Shell is as dangerous as it gets and has been for years. No new evidence activity. The score is held at Critical by the KEV floor. This is a steady-state remediation item, not an escalation event.

CVE-2026-41940 (cPanel/WHM auth bypass): Actively escalating

Field	Value
Lifecycle Stage	*Confirmed*
Severity Band	*Critical (score: 70, at KEV floor)*
Attack Relevance	*Entry (T1078, Valid Accounts)*
Confidence	*Verified*
current_trend	*Escalating*
trend_drivers	*new_weaponized_source (ExploitDB, 4 days ago); new_weaponized_source (Metasploit, 13 days ago). Note: values shown are illustrative of the shape; query /v1/signals/CVE-2026-41940 for current values.*

Field	Value
signal_velocity	*High*
threat_age_days	*31*
What changed	*Weaponized exploit detected (ExploitDB). Evidence from 7 independent sources.*

Reading: Both CVEs are Confirmed Critical. But CVE-2026-41940 has a new weaponized exploit that appeared 4 days ago, three corroborating sources, and rapidly accumulating evidence. This is where your remediation urgency goes first.

> **This is the difference** Without temporal context, both CVEs look identical: Confirmed Critical Verified. With temporal context, one is a stable long-term threat and the other is actively getting worse. Same score. Completely different situation.

What ESIP Is Not

ESIP is not...	Why not
A vulnerability scanner	*ESIP never touches your environment. No agents, no credentials, no cloud connections.*
An asset manager	*ESIP signals at the CVE class level. It does not know which of your systems are affected.*
A CVSS replacement	*CVSS scores base severity. ESIP measures current threat activity. Both matter. They answer different questions.*
A threat intelligence platform	*Threat intel tells you who is attacking. ESIP tells you what is becoming dangerous. Related but different.*
A prioritisation engine	*ESIP produces the global intelligence half of a prioritisation decision. You supply the asset context half.*
An AI narrative generator	*Temporal Context is computed deterministically from evidence. It is not generated by AI, not editorial, and not probabilistic. Same evidence always produces the same verdict.*

How Practitioners Use It

ESIP is API-first. Signals are consumed programmatically, not through a dashboard.

Vulnerability Management Platform Integration

A VM platform pulls ESIP signals via API and applies them to its CVE-asset mapping. CVEs with Active or Confirmed lifecycle stage get elevated priority in the remediation queue. The VM platform knows which assets are affected; ESIP knows how dangerous the CVE is right now and whether it is getting worse.

With temporal context: the platform can surface current_trend = Escalating as a separate urgency flag, independently of severity band. A Medium-band CVE that is Escalating with High velocity may warrant faster action than a Critical-band CVE that is Stable with Low velocity.

CTEM Programme Input

A Continuous Threat Exposure Management programme runs a weekly cycle. The analyst pulls all Confirmed and Active signals from ESIP, filters by attack_relevance = Entry to focus on initial access exposures, then cross-references against the asset inventory to identify what is locally relevant.

The what_changed evidence field explains why each CVE is elevated this week, useful for briefing stakeholders. The trend_drivers field explains whether the situation is improving or deteriorating.

MSSP / Multi-Client Deployment

An MSSP ingests the ESIP feed once and applies it across all client environments. ESIP signals are class-level and client-agnostic. The same signal applies to any organisation running the affected technology.

Compound risk signals flag CVEs where multiple independent evidence types are reinforcing the same exposure. Escalating trend with High velocity signals flag CVEs where the threat picture is actively worsening. Both become immediate escalation items for any client with the affected technology in their inventory.

What Changes When You Use ESIP

- Decisions are driven by current threat activity, not fixed CVSS ratings.
- Emerging signals surface before exploitation is confirmed, giving your team time to act.
- current_trend surfaces whether a threat is getting worse, holding, or fading: independently of the severity score.
- Only a small percentage of CVEs represent the active threat set. That is the list that matters.
- Every prioritisation decision is backed by explainable evidence: useful for briefings and audit trails.

Signal Coverage: Setting Expectations

ESIP signals a small fraction of the total CVE universe. This is intentional.

> ***A platform that signals everything signals nothing.***

At any given time, approximately 3 to 10 percent of published CVEs will have an active ESIP signal. The remainder have not crossed any qualifying threshold: no confirmed exploitation, no significant EPSS movement, no published exploit code.

Absence of a signal is not a safety guarantee. It means that CVE has not triggered threshold conditions that warrant a signal. Your base vulnerability management process handles the rest.

The Signal Lifecycle

Signals are not static. They follow a lifecycle as evidence appears, builds, and ages.

Stage	What It Means
Emerging	*Early warning. EPSS is moving upward significantly. No exploit code yet. Watch this one.*
Active	*Exploit evidence present. A PoC or Metasploit module exists. An attacker has a capability.*
Confirmed	*In CISA KEV or ENISA EUVD. Real-world exploitation confirmed. Highest urgency; treat as immediate.*
Declining	*Evidence is aging. No new observations. Signal still present but losing weight.*
Dormant	*All evidence expired. Score drops to zero. Signal preserved for history. No current threat activity.*
Remediation Available	*A vendor fix exists globally. Additive; sits alongside any other stage. ESIP cannot tell you whether you have applied the fix.*

The Shift

ESIP does not replace your vulnerability management programme. It changes how you make decisions within it.

Before ESIP	With ESIP
What vulnerabilities exist?	*What is becoming dangerous right now?*
Ranked by static CVSS score	*Prioritised by current evidence and threat activity*
Same CVE looks the same every week	*CVE status reflects this week's evidence reality*
Cannot explain why a CVE is high priority	*Every verdict backed by explainable, auditable evidence*
No signal on whether threats are worsening	*current_trend and trend_drivers surface the direction of change*

> ***That is a small change in the question. It is a large change in what you do with the answer.***

How to Get Started

ESIP is available in two tiers. The free tier is the recommended evaluation path before committing to a commercial agreement.

Path	What you get
Free tier (no commitment)	*Access to observable evidence data via the free public API and GitHub daily snapshot. Includes KEV inclusion, exploit availability, ATT&CK technique mappings, and detection artifacts. No verdict scoring, no temporal context. Use this to evaluate data quality and integration fit before purchasing.*
Commercial tier	*Full signal verdicts, Temporal Context, Compound Risk, real-time freshness, and API SLA. Delivered under a commercial agreement with ZenzizenSec.*

> **Evaluation path** Start with the free tier. Register at zenzizensec.com to obtain a free API key. Integrate against the free public API at public.exposuresignal.io to validate the data model and assess fit. When ready to evaluate commercial capabilities, contact ZenzizenSec to arrange a commercial trial against api.exposuresignal.io.

ESIP Overview v1.6 · ZenzizenSec · June 2026
Full specification suite and commercial documentation available. Contact ZenzizenSec for access.