ESIP

Exposure Signal Intelligence Platform

Version: 1.3 · ZenzizenSec · May 2026 · Classification: External

What it is, why it exists, and how it works. A five-minute 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 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?

ESIP provides the global threat signal. You provide the local context.

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 supporting evidence
KEV	Known exploited vulnerabilities	Binary. In the list or not, no gradation or trend
Threat intel	Who is attacking	Not CVE-focused. Hard to operationalise per vulnerability
ESIP	**What is becoming dangerous right now**	**Requires your asset context to complete the picture**

What ESIP Watches

ESIP continuously monitors multiple globally observable data sources, including vulnerability databases, exploit repositories, 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.

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 in that assessment
- What stage of the attack kill-chain the CVE enables

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. Declining means evidence is aging without new activity.
Severity Score	0 – 100	How dangerous this exposure class is right now, based on all available 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. Unclassified means no ATT&CK mapping exists yet for this exposure class.
Confidence	Low · Medium · High · Verified	How reliable is the evidence behind this verdict. Verified means multiple high-trust sources agree. Low means early signal with limited corroboration.

> **Free tier vs commercial tier:** ESIP's free tier surfaces the underlying observable evidence: KEV inclusion, exploit availability, ATT&CK technique mappings, and public detection artifacts. The four-dimension verdict (lifecycle stage, severity score, attack relevance, confidence) is computed by the commercial tier from that evidence and is not present on the free-tier API or GitHub snapshot.
>
> See the ESIP Free Tier Consumer Guide for the free-tier field set. Full commercial documentation is provided under the commercial agreement.

Every signal also includes an evidence payload: what changed, which sources contributed, and when the signal last updated. There are no black boxes: every verdict is explainable.

A Signal in Plain English

Here is what an ESIP signal looks like for a real CVE.

Here is what that looks like in practice:

> **CVE-2021-44228** *(Log4Shell)*
>
> Apache Log4j remote code execution
>
> `Lifecycle Stage` **Confirmed** `Severity Band` **Critical (94)**
> `Attack Relevance` **Entry** `Confidence` **Verified**
>
> **What changed:** *In CISA KEV since Dec 2021. Metasploit module active. EPSS score 0.97. Evidence from 4 independent Verified/High trust sources.*

In plain English: Log4Shell is in CISA's confirmed exploitation catalogue, has a production Metasploit module, and has a 97% EPSS score. Every attacker who wants to exploit it can. If you run Java applications using Log4j and have not patched, this is your highest priority. ESIP does not know if you run Log4j. That is your job. But it tells you this CVE is as dangerous as it gets.

This is the difference between knowing a vulnerability exists and knowing it is your highest priority.

What ESIP Is Not

This is as important as what it is.

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 servers 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.
A CSPM or IAM tool	Configuration and identity risk require access to your environment. ESIP deliberately does not have that.

How Practitioners Use It

ESIP is API-first. Signals are consumed programmatically, not through a dashboard. The three most common integration patterns:

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. Together they produce a prioritised list that reflects current threat reality, not static severity scores.

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 without sending them to read raw intelligence feeds.

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. The MSSP applies their own client asset inventories to filter for relevance per client. Compound risk signals flag CVEs where multiple independent evidence types are reinforcing the same exposure, making them immediate escalation items for any client with the affected technology.

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
- A small percentage of CVEs represent the active risk 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 CVEs that do generate signals are the ones with evidence of real attacker interest. That is a small, high-value list. Exactly what you want for prioritisation.

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. 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.

Instead of asking: *"What vulnerabilities exist?"*
You start asking: ***"What is becoming dangerous right now?"***

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