Exposure Signal Intelligence Platform

Public-Facing Output Contract

Version 1.5 · June 2026 · ZenzizenSec: External

Audience: Customers, integrators, partners, platform vendors, auditors

Classification: ZenzizenSec: External

This document answers exactly three questions:

- What will I receive?
- What does it mean?
- How stable is it?

It is not an API reference. It is not an engineering specification. It contains no database tables, ingestion logic, replay procedures, or internal architecture. Those documents are for engineers building ESIP: this document is for organizations integrating with it.

1. What an Exposure Signal Is

An Exposure Signal is a deterministic, evidence-based assessment of a CVE's current threat state. It is derived from globally observable public data: what is actually happening in the wild with a given vulnerability, right now.

It is not a CVSS score. CVSS measures theoretical severity at time of disclosure and does not change.

It is not a scanner finding. ESIP has no visibility into your environment.

It is not a vendor advisory. Advisories describe what a vendor knows. ESIP describes what attackers are doing.

> **The distinction that matters** A static EPSS score of 0.82 is data. An EPSS that jumped from 0.10 to 0.82 overnight while a weaponized exploit appeared in a public tool is a signal. ESIP detects, classifies, and delivers the latter. With schema v1.4, ESIP also tells you whether that signal is getting worse, better, or holding steady: and why.

1.1 What ESIP Covers

ESIP produces signals for CVEs where globally observable evidence exists: active exploitation, weaponized exploit tools, significant exploit probability movement, or confirmed patch availability. ESIP currently covers tens of thousands of active signals across more than 250,000 published CVEs. CVEs with no signal have no observed attacker activity: absence of signal is accurate, not a gap.

> **Absence of a signal is not clearance** It means ESIP has not observed globally meaningful evidence. Whether a CVE affects your environment is a question only you can answer with your asset inventory.

1.2 What ESIP Does Not Provide

- Remediation instructions or patch guidance
- Exploit code or proof-of-concept details
- Asset inventory or exposure assessment
- Organization-specific risk scoring
- Vulnerability scanning or detection
- Predictions about future attacker behavior or time-to-breach estimates

2. Delivery Model

2.1 Access

Signals are delivered via a REST API over HTTPS. Authentication uses an API key passed in the X-API-Key request header. All requests require a valid key.

2.2 Endpoints

Endpoint	Purpose
GET /v1/signals/{cve_id}	*Current signal for one CVE*
GET /v1/signals	*All active signals, paginated. Filter by severity_band, lifecycle_stage, confidence.*
GET /v1/signals/changes?updated_since=	*Signals changed since a timestamp: the recommended polling pattern. Responses ordered by signal_version ascending, then exposure_class_id ascending for deterministic pagination. Note: temporal context changes that do not trigger a signal_version change will not appear in this feed: poll GET /v1/signals/{cve_id} directly to detect context-only changes.*
GET /v1/compound-risk	*CVEs with active compound risk conditions*
GET /v1/schema	*Current schema version and field definitions*
GET /v1/changelog	*Schema version history*

2.3 Data Format

All responses are JSON. All timestamps are UTC in ISO 8601 format (2026-03-21T14:22:00Z). Null field values are explicit: a null means no value is available, not that the field is missing. Treat nulls explicitly in your integration.

2.4 Rate Limits

API responses include X-RateLimit-Limit, X-RateLimit-Remaining, and X-RateLimit-Reset headers. When a limit is exceeded the API returns HTTP 429 with a Retry-After header specifying the wait in seconds. Implement exponential backoff on 429 responses.

3. The Signal Object

Every signal response has the same structure. The six top-level objects are described below. temporal_context may be null if context has not yet been computed for a CVE. compound_risk is null when no compound condition is active.

3.1 Top-Level Fields

Field	Type	Required	Meaning
exposure_class_id	string	Always	*The CVE identifier this signal describes. Example: CVE-2026-41940. The permanent identity anchor: never changes.*
exposure_type	string	Always	*Always cve at current release.*
signal_id	UUID	Always	*Globally unique identifier for this signal record. Never reused. Stable across verdict updates. Use as your correlation key in SIEM, ticketing, and audit systems.*
generated_at	ISO 8601 UTC	Always	*When this specific API response was generated.*
signal_version	integer	Always	*How many times this signal's verdict has changed. Starts at 1. Increments on every verdict update. Never decreases. Compare your stored signal_version against the current value to detect changes efficiently.*
tier1_signals	array	Always	*Active evidence signals contributing to this verdict. See Section 3.2.*
verdict	object	Always	*The scored verdict. See Section 3.2.*
evidence	object	Always	*Human-readable change explanation. See Section 3.2.*
temporal_context	object / null	Always	*How the threat is evolving over time. Null until context is first computed for this CVE. See Section 3.3.*
compound_risk	object / null	Always	*Multi-source amplification. Null means no compound condition is active. See Section 3.2.*
cve_metadata	object	Always	*NVD metadata: title, CVSS score, published date.*
_meta	object	Always	*Response metadata: schema version, source health, request ID. See Section 3.2.*

3.2 Field Definitions

tier1_signals[]: Evidence Signals

Each entry represents one active Tier 1 signal contributing to the verdict. Only evidence signals are included: internal scheduler events are filtered from this array.

Field	Type	Meaning
signal_type	string	*What this signal represents. See Section 4 for the full catalog.*
fired_at	ISO 8601 UTC	*When this signal was first detected.*
source_id	string	*Which data source triggered this signal (e.g. cisa_kev, metasploit, exploitdb, first_epss, enisa_euvd_exploited, github_poc).*
source_trust_tier	string	*Confidence tier of the source: low, medium, high, or verified. Affects the confidence calculation.*
decay_expires_at	ISO 8601 UTC	*When this signal will fully decay if no new evidence arrives.*

> **weight_decay_factor removed in v1.3** The decay weight for each signal type is no longer included in this array. To access decay weights and decay recomputation history, use GET /v1/signals/{id}/evidence which returns decay_state[] and decay_recomputation_events[] arrays.

verdict: Scored Verdict

Field	Type	Meaning
lifecycle_stage	enum	*Current threat stage. Six values: see Section 5.*
entered_stage_at	ISO 8601 UTC	*When the current stage was entered. Resets on each transition.*
severity_score	integer 0–100	*Computed score reflecting weighted evidence adjusted for age. 0 = Dormant. Use severity_band for alerting, not raw score.*
severity_band	enum / null	*Named band: Critical (70–100), High (50–69), Medium (20–49), Low (5–19), Informational (1–4). Null when Dormant. Build alert rules on bands: they are contract-stable.*
severity_band_range	object / null	*Machine-readable band boundaries: {"min": 70, "max": 100}. Null when Dormant.*

Field	Type	Meaning
attack_relevance.technique_id	string	*ATT&CK technique ID (e.g. T1078). Use for control mapping and detection prioritisation.*
attack_relevance.technique_name	string	*ATT&CK technique name (e.g. Valid Accounts).*
attack_relevance.tactic	string	*ATT&CK tactic (e.g. defense-evasion).*
attack_relevance.mapping_confidence	enum	*High / Medium / Low: confidence of the ATT&CK mapping.*
confidence	enum	*Low, Medium, High, or Verified. Derived from source trust tiers and corroboration count. Verified = real-world exploitation confirmed by a government catalog.*
remediation_available	boolean	*True if a patch or mitigation exists globally. Independent of lifecycle stage.*

evidence: Human-Readable Explanation

Field	Type	Meaning
what_changed	string / null	*Plain-English description of the most recent change. Read this first when signal_version increments.*
supporting_signals	array	*Active Tier 1 signal type names contributing to this verdict.*
source_count	integer	*Number of independent data sources currently contributing. Higher values indicate broader corroboration.*
highest_trust_source	string	*Trust tier of the highest-trust contributing source: low, medium, high, or verified.*

compound_risk: Multi-Source Amplification

Null means no compound condition is currently active. When compound risk resolves, signal_version increments.

Field	Type	Meaning
output_class	*string*	*amplified_confidence at current release. Means multiple independent sources confirm the same CVE.*
strength_tier	*enum*	*Weak (2 signals, same source), Moderate (2 signals, independent sources), Strong (3+ signals, 2+ independent sources).*
contributing_signal_types	*array*	*Signal types contributing to this compound condition.*

Field	Type	Meaning
contributing_source_ids	*array*	*Source IDs contributing. Independence across different organizations is what makes compound risk meaningful.*
detected_at	*ISO 8601 UTC*	*When this compound condition was first detected. Preserved across re-evaluations.*

_meta: Response Metadata

Field	Type	Meaning
signal_schema_version	*string*	*Schema version that produced this response. Currently 1.5. Read on every response: when it changes, review /v1/changelog before assuming behavior is identical.*
generated_at	*ISO 8601 UTC*	*When this response was generated.*
degraded_sources	*integer*	*Count of data sources currently delayed or degraded. Non-zero means some signal inputs may be stale.*
api_request_id	*UUID*	*Unique identifier for this request. Include in support tickets.*

3.3 temporal_context: How the Threat Is Evolving

The temporal_context object is new in schema v1.4. It answers questions a static score cannot: is this threat getting worse or better, how fast is evidence accumulating, and when did something significant last happen.

> **temporal_context can be null** Null means the Temporal Context engine has not yet processed this CVE since its deployment. It will self-populate on the next engine run after new evidence arrives. For most CVEs with recent signal activity, temporal_context is populated.

Field	Type	Meaning
current_trend	enum	*Escalating / Stable / Declining. Derived from evidence activity: not from score movement. A falling score can coexist with an Escalating trend when new evidence arrives while a constraint holds the score floor.*
trend_drivers[]	array / []	*Why current_trend has the value it does. Non-empty when current_trend is Escalating or Declining. Empty array when Stable. Each entry: driver (enum name), optional source_id, optional days_ago.*
latest_material_change	string / null	*The signal_type of the most recent meaningful evidence event. A material change is a new exploitation confirmation, weaponized exploit, significant EPSS movement, or stage transition.*
latest_material_change_source	string / null	*The source_id that produced the most recent material change.*
latest_material_change_at	ISO 8601 UTC / null	*Timestamp of the most recent material change.*
time_since_last_material_change_days	integer	*Days since the most recent material change. Computed at response time. 0 = changed today.*
signal_velocity	enum	*Low / Medium / High. Measures how fast evidence is accumulating over the last 14 days. This is NOT a severity indicator: a High-velocity CVE can have a low severity score. It measures rate of change, not danger level.*
threat_age_days	integer	*Days since the first ever signal fired for this CVE. Computed at response time.*
weaponization_age_days	integer / null	*Days since the first weaponized_exploitation_available signal fired.*

Field	Type	Meaning
		Null if no weaponized evidence has ever been observed.
exploitation_phase	string / null	*Response Position phase. Always null in v1.4: will be populated in a future release.*
context_version	integer	*Monotonically incrementing per CVE. Increments when current_trend, trend_drivers, signal_velocity, or latest_material_change changes. Does not increment on daily elapsed-time changes alone.*

> **Reading the story from temporal_context** Four fields tell the story: current_trend (Escalating), trend_drivers (new_weaponized_source from exploitdb), time_since_last_material_change_days (4), signal_velocity (High). Together: a new weaponized exploit appeared 4 days ago from three corroborating sources, evidence is accumulating rapidly. The score may appear flat due to a floor constraint: the temporal context explains the truth.

Trend Drivers

trend_drivers lists the specific evidence events that explain why current_trend has its value. Common driver values: new_weaponized_source (a new weaponized exploit appeared), new_verified_exploitation (government-confirmed exploitation), epss_increasing (exploit probability rising sharply), multiple_sources_converging (three or more independent sources active), evidence_base_weakening (no new activity and existing signals aging), key_signal_expiring (a high-trust signal approaching expiry). Consumer integrations should check for specific driver values by name, not by array position.

4. Signal Catalog

These are the Tier 1 signals ESIP produces. Signal names describe what the evidence means: not which data source produced it.

Signal Type	Meaning	Typical Confidence	Consumer Action
exploitation_confirmed_in_wild	*CVE is in CISA KEV or ENISA EUVD. Real-world exploitation confirmed by a government agency.*	*Verified*	*Immediate. No waiting for further corroboration.*
weaponized_exploitation_available	*A production-ready exploit tool exists. Accessible to non-specialist attackers.*	*High*	*Urgent. Weaponized availability materially increases actual exploitation risk.*
exploit_capability_emerging	*Public exploit code or reference detected. An attacker now has capability.*	*Medium*	*Elevated priority. Watch for escalation to weaponized.*
exploit_likelihood_increasing	*EPSS score moved meaningfully upward. Risk trajectory changing.*	*Low–Medium*	*Early warning. Monitor for escalation.*
remediation_path_available	*A vendor patch or mitigation globally confirmed.*	*Low*	*Informational. Apply asset context to determine urgency.*
compound_risk_detected	*Multiple independent signals reinforce the same CVE.*	*Medium–High*	*Elevated priority relative to severity score alone.*

5. Lifecycle Stages

A CVE moves through stages as evidence accumulates and ages. Stage transitions are deterministic: they follow defined rules, not editorial judgment.

Stage	Meaning	Severity Score Range
Emerging	*Exploit probability increasing but no exploit code or confirmed exploitation. Watch this CVE.*	*Typically 10–40*
Active	*Exploit code exists or exploitation is beginning. Prioritize based on your asset inventory.*	*Typically 40–70*
Confirmed	*Real-world exploitation confirmed: KEV or EUVD inclusion. Highest urgency. Score always ≥70.*	*Always ≥ 70*
Declining	*Evidence aging. No new signals arriving. Severity decreasing. Monitor for reactivation.*	*Variable, trending down*
Dormant	*All evidence fully decayed. No active signals. Severity score = 0. History preserved.*	*0*
Remediation_Available	*A vendor patch or mitigation exists. Coexists with any other stage.*	*Unchanged by this flag*

6. Severity Bands

Use bands for alerting and routing. Use the raw score only for fine-grained sorting within a band. Band boundaries are contract-stable: changing them is a Major version event requiring 90 days notice.

Band	Score Range	Recommended Response
Critical	*70–100*	*Immediate. Includes all KEV and EUVD-confirmed CVEs. Do not wait for Verified confidence.*
High	*50–69*	*Same-day response. Escalate when confidence ≥ Medium.*
Medium	*20–49*	*Next remediation cycle. Include in daily or weekly digest.*
Low	*5–19*	*Monitor. Alert only on band escalation.*
Informational	*1–4*	*Patch management awareness. No security alerting required.*
null	*0 (Dormant)*	*No active signals. No action required.*

7. Stability Guarantees

7.1 Determinism

Given the same evidence, the same scoring rules, and the same point in time, ESIP always produces the same signal verdict. signal_version increments exactly once per verdict change. Reprocessing identical evidence produces identical output without incrementing signal_version.

7.2 Field Stability Classes

Stability Class	Fields	Behavior
Immutable	*exposure_class_id, signal_id, exposure_type*	*Never change. Safe as permanent correlation anchors.*
Monotonic	*signal_version, context_version*	*Only increases. Never resets.*
State machine	*lifecycle_stage, severity_band, current_trend*	*Follow defined transition rules. Names are contract-stable.*
Dynamic	*severity_score, confidence, tier1_signals[], evidence.*, temporal_context.**	*Updated regularly. Poll for updates.*
Computed at response time	*temporal_context.threat_age_days, weaponization_age_days, time_since_last_material_change_days*	*Derived from anchor dates at response time. Not stored as integers.*

8. Change Management

The signal schema is a contract. Every API response carries _meta.signal_schema_version: read this field on every response.

Change Type	Example	Notice Period
Patch	*New optional _meta informational field*	*None: additive, non-breaking*
Minor	*New signal type; existing field deprecated*	*30 days before taking effect*
Major	*New verdict dimension; lifecycle stage renamed; score range changed*	*90 days + migration guide*

> **Forward compatibility rule** Ignore unknown fields rather than rejecting them. New optional fields may be added in Patch and Minor versions without notice. Integrations that fail on unknown fields will break on routine schema updates.

9. Consumer Responsibilities

- Store signal_version per CVE and compare on each poll: this is the most efficient change detection method.
- Read _meta.signal_schema_version on every response. When it changes, review the schema changelog before assuming behavior is identical.
- Treat null fields as explicit states: null compound_risk means no active compound condition, null severity_band means Dormant, null temporal_context means context not yet computed.
- Apply your own asset context. ESIP signals describe global class-level threat state. Whether a Critical CVE is dangerous in your environment depends on whether you run the affected software. ESIP does not know your asset inventory.
- Use severity_band for alert routing, not raw severity_score. Band boundaries are contract-stable; score values within a band fluctuate continuously.
- For temporal_context changes that do not trigger a signal_version change: poll GET /v1/signals/{id} and compare context_version. The change feed does not surface context-only changes.
- Handle API errors gracefully. Retry transient errors with exponential backoff. Respect Retry-After on 429. Do not treat a 404 as permanent.

11. Data Sensitivity Statement

ESIP signals contain globally observable threat intelligence derived from public data sources. They describe CVE-class threat state at a global level: not findings from any consumer's environment.

ESIP signals do not contain: customer-specific data, organizational asset information, personally identifiable information, confidential vulnerability research, or exploit code.

This is relevant for data classification, procurement review, and regulatory compliance. ESIP output is global threat intelligence, not customer data.

ESIP Public-Facing Output Contract v1.5 · ZenzizenSec · June 2026
This document is the authoritative external signal contract. Where engineering specifications and this document conflict, this document governs consumer-facing commitments.