CSURFACE PLATFORM · 01 · ARCHITECTURE

Continuous operation · agentless, no installation

Continuous Exposure Management platform delivered as SaaS, with per-tenant isolation and optional integrations for additional depth. Operating model, technical specifications and security posture on a single page.

Datasheet 2026
CSURFACE

OPERATING MODEL

Agentless

discovery and monitoring entirely external · no installation in the customer environment

ISOLATION

Isolated tenant

databases and credential vaults isolated per organization

DATA RETENTION

15 days

operational telemetry window · evidence exportable on demand

Data flow · from the external surface to the alert

01 · COLLECTION
External signals
CT logs, DNS, WHOIS, ASN, headers, runtime scripts, threat feeds

→

02 · CORRELATION
Machine Learning
multiple signals per asset across distinct dimensions for ownership attribution and context

→

03 · ENRICHMENT
Threat Sensor
exploitation probability index, active-exploitation catalog, active exploit, campaign indicators

→

04 · OUTPUT
Prioritized alert
SIEM, ticketing, ChatOps, API, dashboard

There is no agent, instrumentation or traffic originating from the customer environment toward the platform. Collection is entirely external, in line with External Attack Surface Management (EASM) practices.

Technical specifications

DELIVERY MODEL	SaaS · multi-tenant with databases and credential vaults isolated per tenant
HOSTING	AWS · infrastructure controls and certifications inherited from the provider layer
DATA RESIDENCY	Global · region selection by contract
OPERATIONAL RETENTION	15 days · standard telemetry window; evidence exportable on demand for customer retention
ENCRYPTION	At rest and in transit · TLS 1.2+; provider-managed keys
AUTHENTICATION	MFA required for all users · internal vault for integration credentials and tokens
ACCESS CONTROL	Native RBAC · SSO via SAML/OIDC on request
AUDITING	Full audit log · user actions, configuration changes and evidence access are recorded

Security posture · summary. Platform hosted on AWS with databases and credential vaults isolated per tenant. MFA required, full audit log, no sensitive customer data persisted beyond the 15-day window. Infrastructure controls inherit the cloud provider's certifications; application controls are CSURFACE's.

What the platform does *not* do

It does not operate inside the customer perimeter. It does not install an agent on a customer server, endpoint or cloud. It does not require an inbound firewall rule to be opened. It does not exfiltrate application payloads. It does not act as a WAF, EDR or SIEM.

Optional integrations with cloud, CIEM or WAF, when contracted, extend internal visibility without altering the core external operating model.

CSURFACE PLATFORM · 02 · CAPABILITIES

Four capabilities in a single platform · Machine Learning and agentic layer

Contextual discovery through Machine Learning, continuous threat telemetry, credential monitoring with active validation and exploitability validation — operating over the same inventory and the same data model.

Datasheet 2026
CSURFACE

Machine Learning & Agentic Layer

Continuous discovery of the external surface with Machine Learning analyzing multiple signals per asset and an agentic ownership layer. Broad, context-aware coverage with false positives practically eliminated and technically minimal noise.

Threat Sensor

Proprietary emerging-threat telemetry correlating public vulnerability databases, exploitation probability indexes, active-exploitation catalogs, active exploit and campaign indicators. Continuous updating of the prioritization queue — recalculation within 24h of a signal change.

Credential Leak Monitor

Monitoring of breach databases, dark web and public repositories with non-destructive active validation. Confirmation of credential validity before the alert — reducing false positives and unnecessary mobilization.

Exposure Validation

Active exploitability validation where test coverage exists, with passive monitoring across the rest of the surface. It distinguishes "vulnerable" from "exploitable in the observed context" — eliminating the noise of CVEs that are present but not reachable through the published vector.

Asset contextualization · Machine Learning across three dimensions

Each discovered asset is automatically classified across three dimensions — **industry sector**, **owning business area** and **asset category** — so that the platform can decide its importance to the organization without dependence on manual classification. Inference of internal ownership accelerates triage: the alert reaches the right team pre-assigned.

Discovery pipeline · Machine Learning and agentic layer

For each observed asset, the platform analyzes **multiple signals** across distinct dimensions — each dimension contributing clues about the asset's ownership and its relationship to the organization. The output of one layer feeds the next, with progressive refinement: what rises to the following layer is more reliable than what entered.

The result is a **confidence score per asset** that synthesizes all the layers — without dependence on manual classification, with an auditable decision at each stage.

Agentic decision layer. Over the models operates an autonomous agent that evaluates the results, resolves conflicts between converging and diverging signals and decides the final ownership attribution. The agent reduces the required human effort to a minimum — eliminating false positives at the source, with **high precision** across the organizations analyzed. Human feedback, when it occurs, feeds the next round of the model, further reducing the need for future interaction.

Native integrations

- **SIEM** · Splunk, Sentinel, QRadar, Elastic
- **Ticketing** · Jira, ServiceNow
- **ChatOps** · Slack, Microsoft Teams
- **SOAR** · Splunk SOAR, Tines, Cortex XSOAR
- **API** · REST + generic webhook for custom integration

Human feedback model

When the analyst reviews an alert and adjusts attribution, severity or status, that signal enters a **loop** with the ML layer — the model learns with each interaction. The operational consequence is that the need for human intervention **declines progressively** over the period of use, as the agent incorporates the preferences and context specific to that tenant.

See the platform in action — Request an analysis