Solutions / AI SOC / Datasheet

CODERED AI SOC COMMAND CENTER (THREAT RESPONDER HUB)

Because your SOC analysts should be threat hunting, not alert chasing.

The Threat Responder Hub is the multi-layer AI, multi-role SOC platform behind PROVINTELLIX 24x7 operation. CodeRed analysts investigate incidents and exposure in the Hub. Customers use their portal to review published findings, track remediation and request incident response.

Arrange a walkthrough | Read the service page

ISO/IEC 27001:2022 certified, with SOC operations aligned to ISO/IEC 27017: 2015 **NACSA licensed** · Managed SOC Monitoring Service · Penetration Testing Service Provider
Global Top 250 MSSP 2024 and 2025

PRODUCT PROFILE

Three access surfaces, one tenant

Internal cases stay inside the SOC. Pending investigations and rejected findings are not visible to customers; publication is what makes a prepared incident visible.

01
SOC workspace
Analysts review active work, affected assets, evidence and recommendations.

02
Customer portal
Customers see published findings, status, remediation details and discussion history.

03
Partner access
Partners review the customer portfolios they are authorised to access.

SCORING REFERENCE

Four scores with different inputs

Each score has its own drivers and its own interpretation. Inspect the components and underlying records before acting. No score establishes attack probability or proves that a recorded control works.

AEAS
Exposure score
Exposure findings assessed through exploitability, impact and threat context, with recorded treatment taken into account.
Higher means more scored exposure

SPS
Security Posture Score
Active incident pressure, evaluated SLA outcomes and the tenant AEAS aggregate.
Lower means less measured pressure

PRS
Predictive Risk Score
Collected external intelligence assessed against the organisation's monitoring profile.
Higher means a stronger threat profile

DPS
Defensive Posture Score
Recorded D3FEND control status relevant to ATT&CK techniques in the tenant's ticket history.
Higher means more recorded coverage

SCORING METHOD

How each score is built

All four run 0 to 100. The weightings are published so that any reading can be reconstructed from the records behind it.

SPS · Security Posture Score

For MXDR tenants. A current-state measure; the dashboard date filter does not change it.

Incident pressure: active incidents weighted by severity, capped at 100 — **40%**

SLA pressure: the missed share of evaluated response and resolution results — **35%**

Exposure pressure: the current tenant AEAS aggregate — **25%**

Limit. Missing AEAS or evaluated SLA data contributes zero pressure. A lower score does not establish that the environment is secure.

AEAS · Exposure score

The tenant aggregate uses open or in-progress scored observations, giving higher severity more weight.

Exploitability: category and CVE context, optional AI refinement — **25%**

Impact severity: category and CVE impact — **20%**

Exploit maturity: public exploit references — **20%**

Threat activity: category, domain and CVE mentions — **20%**

Exposure context: ticket age, severity, resolution SLA status — **15%**

Limit. Severity caps and recorded treatment affect the total; accepted risk can lower it. No scored active observations means no measurement.

PRS · Predictive Risk Score

An assessment against the monitoring profile. The AI result can differ from a weighted sum; the formula fallback applies these weights.

Industry targeting: activity relevant to the monitored industry — **25%**

Actor activity: actors appearing in relevant intelligence — **25%**

Geographic exposure: countries in the monitoring profile — **20%**

Historical pattern: direct domain mentions in historical records — **15%**

Trend momentum: recent activity against the preceding period — **15%**

Limit. Not an attack probability or date forecast. Source coverage, profile settings and data availability change the assessment.

DPS · Defensive Posture Score

Maps ATT&CK techniques from the tenant's tickets over the past year to D3FEND countermeasures, then reads the recorded status of those controls.

RECORDED STATE	COUNTS AS
Active	One toward the relevant control total
Partial	One half toward the relevant control total
Missing	Zero; absent records default to missing
Not applicable	Excluded from the denominator

Limit. Status includes customer declarations; this view does not test controls. An empty set of mapped techniques means insufficient coverage data.

CONTROL MAPPING

From ATT&CK history to D3FEND controls

Techniques recorded in tenant tickets are mapped to related D3FEND countermeasures, with the recorded deployment status by kill-chain stage. Review recorded status and linked tickets before planning work.

01
Locate a gap
Review coverage by kill-chain stage and defensive tactic. Open a cell to see its mapped countermeasures.

02
Check the control
Read the description and MITRE reference. Record whether the control is active, partial, missing or not applicable.

03
Agree the work
Agree implementation work with the control owner, then test its effectiveness.

The broader kill-chain coverage matrix has a different population from the DPS denominator, so the two views will not reconcile exactly.

INVESTIGATION

Incidents, issues and observations

Open an incident to review affected assets, evidence, the recorded assessment, recommendations and status history. Internal cases organise alert bundles; publication makes a prepared customer incident visible through the portal.

Incident
Threat activity
Suspected or confirmed threat activity, with affected assets and investigation or response work.

Issue
Weakness or vulnerability
A security weakness or vulnerability, including relevant CVEs, severity and remediation details.

Observation
Exposure or intelligence
An exposure or intelligence finding that needs validation or follow-up.

Visibility rule: pending investigations stay inside the SOC; rejected findings remain available for review. Release and publication follow the configured operating policy, and customer queries exclude internal cases.

INVESTIGATION ASSISTANCE

CyberLLM, CodeRed GPT and SIMON

Three assistants with three jobs. AI access varies by configuration, and response execution has separate permissions and policy checks. Analysts review and edit AI output before publication.

01
CyberLLM
Structures case evidence and executive synthesis. The four verdicts are real threat, false positive, benign activity and needs investigation.

02
CodeRed GPT
Searches alert and ticket records, retrieves details and builds entity timelines through read-only tools. Natural-language threat hunting.

03
SIMON · SOC chat
Supports ticket questions, remediation guidance, discussion and analyst escalation for customers.

Selected fact checks and an evidence-only fallback support review; they do not guarantee accuracy.

RESPONSE OPERATIONS

Response policies and actions

Response policy determines whether an eligible action stays manual, is recommended, needs approval or can run automatically. The action history records its state.

01
Manual
Keep response decisions with the analyst and the applicable authorisation workflow.

02
Recommend
Generate response recommendations for review without automatic execution.

03
Approval
Queue eligible actions for approval before execution.

04
Full auto
Permit eligible IP blocking when explicitly enabled and policy checks allow.

Check the impact

- Consider tenant authority, protected assets, target impact and expiry before any action.
- Host and account actions always follow approval workflows.
- Trusted-country controls are set in the response-planning workspace.
- Queued actions are not confirmed execution.

Connect supported controls

- Endpoint isolation and release, for example CrowdStrike.
- Identity disable and revoke, for example Microsoft Entra.
- Firewall block and unblock, for example FortiGate.
- Confirm vendor editions, permissions and supported actions during onboarding.

EXTERNAL EXPOSURE AND THREAT INTELLIGENCE

CTEM, asset inventory and advisories

CTEM and CodeRed ASM group external findings by exposure category and attack surface, each observation links the assets and evidence needed for remediation. AVM intelligence and the advisory catalogue give context for tenant-specific investigation.

From finding to asset

- **Find the affected system.** Search by hostname, IP, URL or domain, use tenant and lifecycle filters to narrow the working list.
- **Connect the findings.** Review linked threats and exposure scores beside the asset, then inspect the related ticket evidence.
- **Track treatment.** Record remediation in the ticket, keep fixed, accepted-risk and not-applicable outcomes distinct, with supported revalidation where available.
- Configured probes may collect fresh evidence.

AVM and emerging threat advisories

- ***AVM workflow.*** Search intelligence by organisation, domain, actor, industry or geography, review source and watchlist history. Credential follow-up records resets, MFA and offboarding; recorded outcomes need independent verification.
- **Find.** Search advisories by ID, title, CVE or keyword, filter severity and sort by publication date.
- **Assess.** Review impact and countermeasures, then compare the described behaviour with customer evidence.
- **Discuss.** Ask SIMON contextual questions and bring relevant advisories into a remediation discussion with the SOC.

An advisory listing does not establish customer exposure.

CUSTOMER COORDINATION

Published findings and response service

The SOC investigates in its internal workspace. Customers see published findings, affected assets, recommendations and status in their portal, with a discussion history separate from internal notes.

01
Shared ticket record
Comments, attachments and recommendations stay with the ticket. Configured response and resolution deadlines show on-track, at-risk or breached states. Export a ticket as CSV, Excel or JSON, or create an incident summary PDF.

02
Incident response requests
Where CodeRed IR is included, service requests provide a workflow for specialist assistance. Review request status and tenant-credit information in the Hub, then follow the agreed authorisation and response responsibilities.

03
Service fit
During onboarding, confirm connected sources, supported response actions and who may request or approve them. Set AI access and customer or partner permissions. SOC coverage, assessments, IR credits, service levels and reporting cadence follow the service agreement.