PRODUCT DATASHEET · 2026

CodeRed AI Log Manager

Universal Log Acquisition & SIEM Forwarding Appliance

Every customer runs a different mix of log sources: cloud platforms, identity providers, databases, network estates, and the one appliance no vendor ships a connector for. CodeRed AI Log Manager collects from all of them, normalises the output, and forwards it to your SIEM.

When a source falls outside the built-in catalogue, you describe it in plain language and the Log Manager generates a working connector for you. You skip the professional-services engagement and the wait for a vendor release.

⬡ AI-Generated Connectors · Any Source

Built for continuous operation

Scheduled collection with retry, timeout and checkpointing. A source outage creates a delay, never a silent gap in the evidence trail.

Built for the SOC floor

Live status over WebSocket, keyboard-first navigation, and a complete audit trail so shift handover never loses context.

PLATFORM OVERVIEW

The Integration Layer, Shipped as a Product

Log onboarding is where SOC deployments lose weeks. Each source needs authentication, scheduling, paging, checkpointing, failure handling and format conversion before one event reaches the SIEM. CodeRed AI Log Manager ships that layer as a maintained product, so your team stops rebuilding it for every customer.

System Architecture

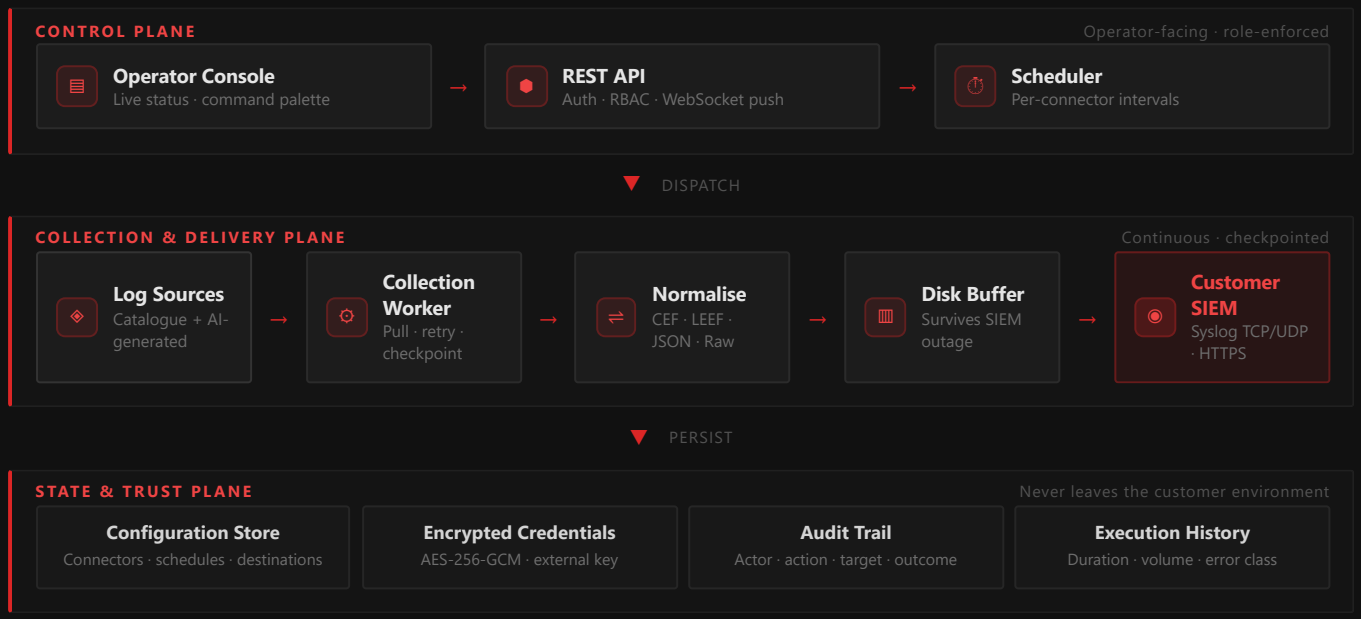

Why This Matters

A new customer signs. Their estate has cloud audit logs, an identity provider, three databases, an EDR, and a firewall appliance from a vendor nobody on the team has integrated before.

Traditionally that last one turns into a change request, a scoping call and a six-week wait, and the customer pays for monitoring that cannot see a third of their estate.

What Changes Here

The catalogue covers the common sources on day one. For anything else, an operator describes the source, reviews the generated connector, and tests it before it goes live. Onboarding takes an afternoon instead of a quarter.

Deployment Model

ONE APPLIANCE PER ENVIRONMENT

Each customer environment runs its own appliance, so tenant separation is physical. Its credentials are encrypted under a key that exists only there, and a compromise of one deployment cannot reach another.

AIR-GAPPED OPERATION

Collection and forwarding run with no internet access and fetch nothing at runtime. Only connector generation needs outbound connectivity, and even that is optional.

NO INBOUND REQUIREMENT

The appliance opens no inbound path from the internet. Outbound access is needed only to the sources being collected and the destination SIEM.

DATA ACQUISITION

Any Source. Including the Ones Nobody Supports.

The catalogue covers the sources most estates run. The generator covers everything else. That combination is the difference between a monitoring service that sees most of a customer's estate and one that sees all of it.

Built-In Connector Catalogue

CATEGORY	SUPPORTED SOURCES
Cloud	AWS CloudTrail (API), AWS CloudTrail (S3), AWS CloudWatch Logs, AWS GuardDuty, AWS WAF, Azure Event Hub, Google Cloud Audit Logs, Huawei Cloud WAF
Identity	Active Directory, Entra ID (Azure AD), Microsoft Graph Security
Database	Microsoft SQL Server, MySQL / MariaDB, PostgreSQL
EDR	CrowdStrike Falcon, Microsoft Defender for Endpoint
SaaS	Google Workspace, Office 365 Management Activity
XDR	Trend Micro Vision One
Email Security	Trend Micro Email Security
Network	Cato Networks SASE
Infrastructure	Docker Host

Connectors load as plugins. If one fails to load, the Log Manager isolates and reports it, and collection continues for the rest of the estate.

AI Connector Generation

For sources outside the catalogue

Describe the source the way you would explain it to a colleague: what it is, how it authenticates, how it pages. The Log Manager writes a working connector, checks it against a static security gate, repairs it if the gate rejects it, and presents both the script and its declared network access for operator approval.

EXAMPLE BRIEF

"Pull firewall traffic and event logs from a FortiAnalyzer over its REST API. It authenticates with an API token and supports paging by timestamp."

What the Operator Reviews

- The generated script in full, before it is saved
- The exact hosts the connector is permitted to reach
- The security gate verdict and any repair rounds taken
- A live connection test, run before activation

Generated connectors save inactive by default. Nothing runs on a schedule until an operator has tested it and switched it on.

Generation Pipeline

Describe → Generate → Security Gate → Repair → Review → Test → Activate

Generated Code Is Never Trusted

A connector written by a machine gets less privilege than one written by the vendor. Two independent layers enforce that, and neither depends on the other being correct.

Layer 1 · Static Gate

Whitelist analysis runs before the script executes. The analyser rejects anything it cannot account for, because a gate that fails open is no gate at all. It refuses shell access, raw sockets, dynamic imports and code-from-data constructs.

Layer 2 · Runtime Confinement

Static analysis cannot see through every construct, so the sandbox confines execution anyway: capped memory, CPU and wall-clock time, a hard ceiling on processes and open files, and egress limited to the hosts that connector declared.

— *Two layers, on purpose*

The static gate filters and the sandbox contains. Security evaluators ask what happens if an attacker bypasses the gate, and the answer is that the sandbox never depended on it.

COLLECTION & DELIVERY

Continuous Collection, Verifiable Delivery

Execution Control

Per connector, per source

Every connector carries its own schedule and failure policy. A source that rate-limits hard and one that publishes in bursts do not share a configuration.

Schedule interval	Configurable per connector
Execution timeout	Per connector; 300 s default
Retry count & interval	Per connector
Checkpointing	Advances only on a successful pull
Overlap protection	A slow run cannot collide with its own next schedule
Stale detection	No logs beyond threshold; 24 h default
Credential expiry warning	14 days ahead by default
Run history	Duration, volume and categorised error per execution

A failed pull re-reads the window

Checkpoints advance only when a run succeeds. When a pull fails, the appliance re-reads that window and may deliver duplicates. It leaves no gap in the evidence trail, the one outcome a SOC cannot audit its way out of.

Health States

Healthy: logs arriving | Error: last run failed | Stale: nothing beyond threshold | Inactive: disabled

Watch the stale state. A connector reporting success while its source has stopped publishing stays invisible to any dashboard that tracks only errors.

Normalisation & Forwarding

Into the SIEM you already own

Events are buffered to disk on collection, then forwarded. Buffering decouples the two, so a SIEM maintenance window does not discard events already pulled from the source.

Forwarding methods	Syslog TCP, Syslog UDP, HTTPS POST
Output formats	CEF, LEEF, JSON, Raw
Destination	Global default, overridable per connector
Delivery buffer	Per connector, consumed on successful forward

Monitoring & Alerting

The console pushes state over a live socket instead of polling, so an operator sees a failure the moment it happens.

- Estate health: healthy, error, stale and inactive at a glance
- Aggregate throughput and total collected volume
- Pipeline view of sources converging on the collector and relaying to the SIEM
- Flow animates only on legs actually carrying data
- Per-connector execution history with categorised failures
- Command palette for keyboard-first navigation during an incident

Alert Conditions

Status change | Credential expiry | SIEM unreachable | Agent disconnected

NOTIFICATION CHANNELS

Email | Webhook

Webhook delivery routes into an existing ticketing or paging system so alerts reach the roster on shift instead of an unattended mailbox.

SECURITY ARCHITECTURE

An Appliance That Holds Every Customer Credential

This platform stores the keys to a customer's entire cloud and identity estate, which makes it a target in its own right. It is built for that threat model and assumes hostile network conditions from the start.

Authentication & Access

Password storage	bcrypt, with over-length inputs rejected instead of truncated
Session tokens	Signed JWT; 60 min access, 7 day refresh
Roles	Administrator, Operator, Viewer, enforced server-side on every endpoint
Brute-force throttle	Counted per username *and* per source address
Lockout escalation	Doubles per repeat offence to a one-hour ceiling
Strike memory	Survives a successful login, so guessing right once does not reset the penalty
User enumeration	Identical timing and error text for unknown and known accounts

— *Two axes, both counted*

Counting per username alone lets an attacker spray many accounts from one host. Counting per address alone lets them rotate addresses against one account. CodeRed counts both, because either measure on its own is easy to bypass.

Credential Protection

Encryption at rest	AES-256-GCM, authenticated
Key custody	Supplied by the environment; never written to the database, never logged
Browser exposure	Secrets are never returned to the console
Reveal path	Separate, explicitly audited endpoint
Audit redaction	Secret values stripped from stored audit payloads

Untrusted Script Confinement

Applies to every operator-supplied and AI-generated connector for its whole lifetime, not only on first run.

Import control	Whitelist; shell, socket, dynamic-import and serialisation modules denied
Call control	Code-from-data constructs rejected
Memory	512 MB address-space cap
CPU & wall clock	Both capped; process group killed on expiry
Processes & files	32 processes, 128 open files, core dumps disabled
Network egress	Only hosts declared by that connector are reachable

Application Hardening

Content-Security-Policy	Self-origin only; framing, plugins and inline script denied
Clickjacking	Framing denied outright
MIME sniffing	Disabled
Referrer leakage	Suppressed
Device APIs	Camera, microphone, geolocation and payment disabled
Cross-origin isolation	Same-origin opener policy enforced

Audit Trail

Every privileged action records actor, action, target and outcome. The appliance logs a failed authentication as an unattributed event, since no user row exists for an attempt against an account that is not real, which is what a defender needs to see.

- Create / update / delete
- Activate / deactivate
- Clone
- Test connection
- Manual run
- Login / logout
- Password change
- Credential reveal
- Lockout events
- Acknowledge / respond

TECHNICAL SPECIFICATIONS

Platform Specifications

Platform Architecture

Architecture	Containerised services with REST API and live socket layer
Console	Single-page web application
Database	Enterprise relational database with managed migrations
Task execution	Distributed worker queue with interval scheduler
Real-time	WebSocket status push
Encryption	AES-256-GCM at rest
Deployment	On-premises or customer cloud; no vendor lock-in
Air-gapped	Supported for collection and forwarding

Data Acquisition

Built-in catalogue	Cloud, identity, database, EDR, SaaS, XDR, email, network, infrastructure
Custom connectors	Unlimited, AI-generated or operator-supplied
Connector loading	Plugin discovery; failures isolated and reported
Scheduling	Per-connector interval, timeout, retry
Resumption	Per-connector checkpointing
Health states	Healthy, Error, Stale, Inactive

Output & Delivery

Forwarding	Syslog TCP, Syslog UDP, HTTPS POST
Formats	CEF, LEEF, JSON, Raw
Destination control	Global default with per-connector override
Buffering	On-disk per connector, consumed on delivery

Security & Access Control

Authentication	Signed session tokens with refresh
Password hashing	bcrypt
Roles	Administrator, Operator, Viewer
Brute-force defence	Dual-axis throttle with escalating lockout
Credential storage	AES-256-GCM, key held outside the database
Script confinement	Static gate plus runtime resource and egress limits
Web hardening	Strict CSP, anti-framing, MIME and referrer controls
Audit trail	Actor, action, target and outcome per event

Operations

Estate dashboard	Live health, throughput and volume
Pipeline view	Source-to-SIEM topology with live flow
Alert conditions	Status change, credential expiry, SIEM unreachable, agent disconnected
Notification channels	Email, webhook
Execution history	Per-run duration, volume and error category
Operator navigation	Command palette, keyboard-first
Accessibility	WCAG 2.1 AA contrast, visible keyboard focus, reduced-motion support

Console Branding

Custom logo	Operator-uploaded, validated by file signature
Product naming	Configurable for partner deployments

SIZING

Host resources and buffer retention are sized against the source mix and retention window of each deployment. Our solutions team will size the appliance against your estate as part of onboarding.

Stop Scoping Connectors. Start Collecting.

Bring us the source your current platform cannot integrate. We will build the connector for it on the call.

Get In Touch

Email
business@coderedcyber.ai

Website
https://coderedcyber.ai

Provintell Technologies Sdn. Bhd.
Powered by CodeRed Cyber

See It Running

A working session against a live appliance. Bring the source you have been unable to onboard and watch a connector get written, gated, tested and activated inside the session.

BOOK A TECHNICAL WALKTHROUGH

https://coderedcyber.ai

Deploys Alongside

The Log Manager is one component of the CodeRed platform. It feeds any SIEM, and integrates directly with the rest of the CodeRed estate.

CodeRed Server Agent

Endpoint telemetry across Linux, macOS and Windows. Covered in its own datasheet.

CodeRed Threat Responder Hub

Detection, response, intelligence, orchestration and reporting for the SOC floor. Covered in its own datasheet.

Why Teams Choose It

- ✓ Any source is in scope
- ✓ Onboarding measured in hours
- ✓ Works with the SIEM you already own
- ✓ Credentials stay in the customer environment
- ✓ Runs air-gapped