



CYBER COMMAND CENTER

# The Agentic SOC layer

AI does the volume. People keep the judgement. Built AI-native — one multi-tenant platform for investigation, response and automation, not a layer bolted onto a legacy SOAR.

SIEM AGNOSTIC

PER-TENANT SOAR

UNIFIED THREAT INTEL

SOVEREIGN DEPLOYMENT

ARABIC NATIVE

**1,000+**

SIEM ALERTS PER SHIFT

**~20%**

MEANINGFULLY INVESTIGATED

**4.8M**

WORKFORCE GAP, UNHIREABLE

## 01 THE SOC PROBLEM, AND WHAT C3 DOES ABOUT IT

### Attacks move at machine speed

Breakout is measured in seconds. Agents investigate on arrival, verdict and ATT&CK mapping attached.

### Mundane operational tasks consume the shift

Agentic automation and playbooks absorb the repetitive work end to end.

### Residency rules block the deal

Your data centre or your own cloud tenancy.

### Most alerts never investigated

Every alert is investigated to the same depth, whoever is on shift.

### Senior skills are scarce and uneven

Agents carry the investigative depth; your analysts supervise and decide.

### Audit asks how the verdict was reached

Findings, evidence and a full audit log on every investigation.

## 02 AGENTIC TRIAGE, INVESTIGATION, AUTOMATION

### Detection to approved response

#### AI triage and prioritisation

Alerts enriched, scored and queued by real priority.

#### Investigation and recommendation

Agents gather evidence in a sandbox, then draft the containment step.

#### Approval-gated execution

Disruptive actions need runbook or ticket approval. All logged.

## 03 AGENT LIBRARY

80+ INTEGRATIONS · 40+ AGENTS

**Autonomous alert triage**

Only the alerts needing human judgement surface.

**Incident response**

IOCs correlated, ATT&CK mapped and actions drafted in minutes.

**Playbook automation**

AI executes the workflow, you hold the guardrails.

**Compliance automation**

10+ frameworks: NESA, CBUAE, ISO 27001, PCI-DSS v4.0, SOC 2.

**Threat hunting**

Continuous sweeps for identity abuse, persistence and C2.

**User access review**

Entitlements reviewed across identity providers, alert tuning proposed.

## 04 HUMAN IN COMMAND

**Approval gates on every action**

Runbook or ticket approval, four-eyes on core systems, kill switch on all agents.

**Supervise, override, correct**

Analysts overrule any verdict. Corrections carry a reason code, reviewed weekly.

**Evidence for every decision**

Findings, evidence and a full audit log per investigation. Agents pause per tenant.

## 05 SOVEREIGNTY BY ARCHITECTURE

**Residency by design**

YOUR DATA CENTRE

YOUR CLOUD TENANCY

Cases, evidence and AI logs stay in region, and models sit behind the gateway — no prompt with incident data leaves as an API call.

## 06 CUSTOMEROPS, FOR MSSPS

**The whole client lifecycle, in one portal**

PROPOSAL

ONBOARDING

DELIVERY &amp; SLA

REPORTS

Proposals convert straight into tenant onboarding, renewals automated. Every engagement is a tracked work item with SLA clocks, and clients watch delivery as it happens.

## MEASURED LIVE

MTTA · MTTR · escalation rate · false-positive rate · playbook success · agent runs · connector health

## ALREADY IN PRODUCTION

**10x**FASTER  
RESPONSE**75%**FEWER ALERTS  
REACH A  
HUMAN**<1 min**MEAN TIME TO  
ACKNOWLEDGE

Live today in a multi-tenant deployment.

## DIFFERENTIATORS

- Complete data sovereignty
- SIEM agnostic
- Policy gates to cut AI cost per alert
- Per-tenant agentic SOAR
- CustomerOps client portal (MSSP)