AGENTIC AI SECURITY OPERATIONS CENTER

Agentic security for the systems your business depends on.

Connect assets, software and cryptographic evidence across IT, IoT and OT. Investigate with AI, rehearse response options and govern operational change.

Protect a critical service. Preserve operational availability.

01	02	03	04	05
Discover	**Investigate**	**Rehearse**	**Respond**	**Verify**
Establish context	AI reasons over evidence	Compare options	Authorize change	Assess the outcome

Evidence / Ownership / Permissions / Decision record

Operational context

Relate exposure to the systems and services it could affect.

Agentic investigation

Use grounded AI and multiple reasoning roles to examine the case.

Human control

Review the proposal, its scope and the recovery plan.

ASOC operating model. Supported inputs, integrations and execution paths are established for the selected deployment.

WHY BREAKWATER ASOC

Connect the evidence. Govern the response.

Examine the threat and its operational consequences together. Bring connected evidence, AI investigation and rehearsal into the decision about what to change.

01 Connect the evidence

OPERATIONAL CONTEXT

Assets & connections
Relationships / Exposure

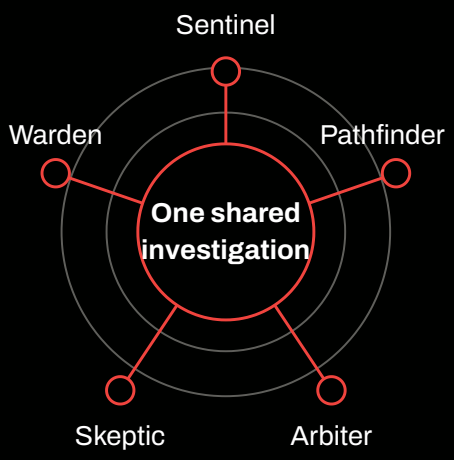
Software & cryptography
Source / SBOM / CBOM

Security observations
Signals / Findings

Sources and gaps for investigation →

02 Investigate with agentic AI

THE COUNCIL / REASONING ROLES



Challenge explanations
Preserve dissent
Rehearse options

03 Govern the response

PROPOSAL TO DECISION

Response proposal

Target & scope

Impact & recovery

Permissions

Human authorization

Verify the outcome
Retain the decision

ASOC operating model / Evidence informs the investigation. Human authority governs operational change.

Understand the path to impact.

Integrated attack paths connect candidate exposures to operational assets. Inspect the relationships and source evidence behind the route.

Bind authority to the action.

Response's governed path checks targets, evidence and approval expiry. Reassess a changed proposal and define verification and recovery.

Rehearsal and execution use the supported model, lab or deployment path. Confirm inputs, handoffs and approval policy for the selected environment.

THE COUNCIL / AGENTIC AI REASONING ROLES

Agentic AI investigation. Human-controlled response.

Each Council role examines a different part of the same case. Consider unusual workstation activity after a rollout: an intrusion, or authorized maintenance?

Sentinel	Triage the signal	01	Classifies the activity and identifies a relevant attack technique. **Does the workstation activity warrant an intrusion investigation?**
Pathfinder	Reason about the attack path	02	Examines a plausible attack sequence and the attacker's possible next step. **Where could this access lead, and which relationship needs checking?**
Arbiter	Interpret the wider pattern	03	Assesses whether the activity suggests human-operated intrusion or automated noise. **Does the sequence support a coordinated intrusion?**
Skeptic	Test a benign explanation	04	Proposes an alternative explanation and examines whether it survives the evidence. **Could the authorized rollout explain the same observations?**
Warden	Prepare the response option	05	Proposes a specific containment action for review, rehearsal and authorization. **Which host, session or key should the team consider restricting?**

Make disagreement useful.

If maintenance could explain the signal, inspect the rollout record before accepting the attack hypothesis. Retain unresolved questions in the case.

Carry reasoning into action.

Review the proposed containment against service dependencies. The next step is rehearsal and human authorization, followed by outcome verification.

The Council routes reasoning tasks to the configured model backend. Roles may use distinct models or share a model, depending on deployment. Evaluate the actual outputs and evidence behind each conclusion.

THE PRODUCTS BEHIND THE AGENTIC AI SOC

Three products.
AI with operational context.

Breakwater ASOC connects asset intelligence, software evidence and agentic response. Each product contributes a distinct mechanism for turning a security question into an informed decision.

01 **Discover / Ground AI investigation in operational context**

Connect assets and exposures through integrated attack paths. Use cited prompts and authorized MCP assistant access to examine routes, evidence gaps and intervention questions.

02 **Provenance / Give AI assistance a traceable source**

Combine source-linked findings, replayable validation and a configured in-product model helper. Connect software dependencies and cryptographic uses to the evidence behind the review.

03 **Response / Orchestrate the agentic investigation**

Use the Council's multiple AI reasoning roles to challenge explanations, compare modeled interventions and prepare response proposals for human authorization.

A shared operating context

Asset identity, source revision, observation time and ownership matter at each handoff. Establish the supported interfaces and explicit mappings between network assets and software evidence in your evaluation.

Asset context

Identity, connections and candidate routes.

Source context

Findings, dependencies and cryptographic uses.

Response context

Reasoning, proposal, authorization and outcome.

ENTERPRISE FIT

Fit the workflow to your operations.

A focused evaluation should establish operational usefulness, integration scope and control over data and change.

01 **Work with existing controls**

Define the evidence sources, imports and supported integrations. Keep your systems of record and change-management responsibilities explicit.

02 **Set the data boundary**

Review local processing, retention, identity, model dependencies and update procedures. Validate the chosen connected, restricted or disconnected configuration.

03 **Test the decision as well as the investigation**

Use a known case to assess AI source support, evidence gaps and proposal quality. Test changed targets, expired approval and failed verification on an agreed supported path.

Agree what success looks like.

Select one environment and a critical service. Name the inputs, responsible owners and acceptance criteria. Review the results together before expanding scope.

Demonstrations use defined data and scenarios. Customer inputs and live actions require separately agreed readiness, permissions and supported execution arrangements.

See Breakwater in action.

Request a demonstration focused on your priorities.

Explore the ASOC platform **Request a demo**