BREAKWATER ASOC / AGENTIC AI SECURITY OPERATIONS CENTER

Investigate the threat. Keep control of the response.

Breakwater Response brings agentic AI reasoning, digital rehearsal and governed action into ASOC. Follow a suspicious-activity investigation through a proposal your team can authorize and verify.

ILLUSTRATIVE SCENARIO / SUSPICIOUS ACTIVITY NEAR OPERATIONS

Containment could help. It could also interrupt the service.

An engineering workstation shows unexpected activity after a software rollout. Security suspects lateral movement; operations needs the workstation available. The team must distinguish an attack from legitimate change and review a response with service impact in view.

01 **Observe**
Records & scope

02 **Investigate**
Council reasoning

03 **Rehearse**
Action & effects

04 **Authorize**
Bounded proposal

The decision

What evidence justifies intervention, and which response preserves the service requirements?

Investigate competing explanations.

Compare the attack hypothesis with maintenance activity, source evidence and missing context.

Prepare a bounded proposal.

Name the action, target, expected effect, service checks and recovery owner before authorization.

RESPONSE'S ROLE IN ASOC

Turn agentic reasoning into a governed decision.

The Council brings multiple AI reasoning roles to the case. Rehearsal informs the intervention; human authority governs consequential change. Supported executors check the scope and verify the result.

01 / THE COUNCIL

Challenge the attack story. Preserve the reasoning.

The Council gives AI models distinct reasoning tasks on a shared case: classify the activity, trace a possible attack, test alternatives and prepare an action for review.

01 Connect the evidence

OPERATIONAL CONTEXT

Assets & connections
Relationships / Exposure

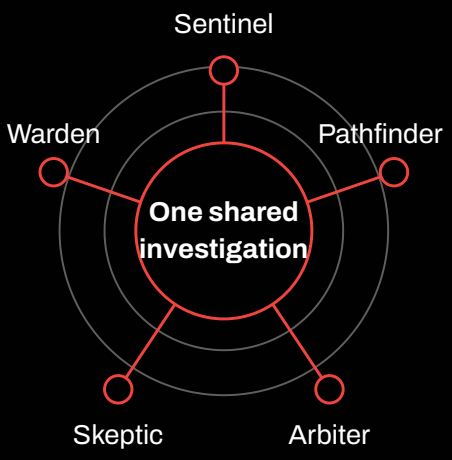
Software & cryptography
Source / SBOM / CBOM

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.

Challenge the explanation.

For the workstation scenario, compare the attack hypothesis with the rollout record and service context. Keep missing evidence and dissent visible.

Prepare the human decision.

Carry the target, scope, service impact and recovery into the proposal. The Council informs the decision; the responsible team grants authority.

Explore the product workflow

Explore the Council workflow →

02 / AGENTIC AI REASONING ROLES

Agentic AI investigation. Human-controlled response.

Each Council role examines a different part of the investigation. Compare its rationale with the evidence and the other roles before selecting a response.

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.

03 / DIGITAL REHEARSAL & SERVICE IMPACT

Rehearse the action you intend to authorize.

A useful rehearsal examines the proposed operation against the intended target. Security and operations review the expected effect, required service behavior and recovery together.

FROM PROPOSAL TO RELEVANT REHEARSAL

01	02	03	04
Propose	**Rehearse**	**Compare**	**Review**
Operation & target	Supported twin / lab	Match intended work	Effect & omissions

Execution gate

The governed executor checks dry-run parity and stops a failed or mismatched rehearsal.

Compare the intervention with its operational cost.

Review lens	Question for the team	Evidence to retain
Threat effect	Does the restriction address the supported attack explanation?	Target, operation and modeled effect.
Service continuity	Which legitimate connections or tasks must remain available?	Owner-defined service checks and dependencies.
Recovery	What restores the intended state if the change or health check fails?	Compensation steps, owner and recovery checks.

Give the owner a precise proposal.

Compare a scoped connection restriction with broader isolation. Agree which options the adapter can exercise.

Carry assumptions into approval.

Record omitted effects: timing, device behavior and undocumented dependencies. Review material gaps with the service owner.

Confirm the rehearsal environment, supported adapter and service tests before granting site authority.

04 / AUTHORITY, VERIFICATION & RECOVERY

Bind permission to the plan. Verify what happened.

The approved target, action and evidence must still match at execution time. A successful command is only one part of the result; the service must behave as required.

GOVERNED EXECUTION / CONTROLS TO EVALUATE

Keep authority specific and time-bound.

Scope	Review the action, targets, evidence and residual risk as one safety case.
Reviewers	Apply the impact class: human signoff, independent reviewers and process-engineering review as required.
Validity	Check approval expiry and the fingerprint of evidence, actions and targets.
Verification	Run the configured health checks and inspect the service behavior agreed with its owner.

AFTER THE ACTION / RETAIN THE ACTUAL OUTCOME

Verified

Retain the execution and service-check results.

Recovery attempted

On supported failures, attempt the recorded compensation.

Recovery incomplete

Stop further work and require manual intervention.

ONE AGENTIC AI SOC / THREE CONNECTED PRODUCTS

Agentic investigation. Accountable action.

Discover supplies asset context; Provenance supplies software evidence. Response connects agentic investigation to scoped authority and outcome verification.

Agree supported inputs, interfaces and evidence mappings for the selected deployment.

Bring a response decision worth rehearsing.

Request a Response demo →

Explore Response