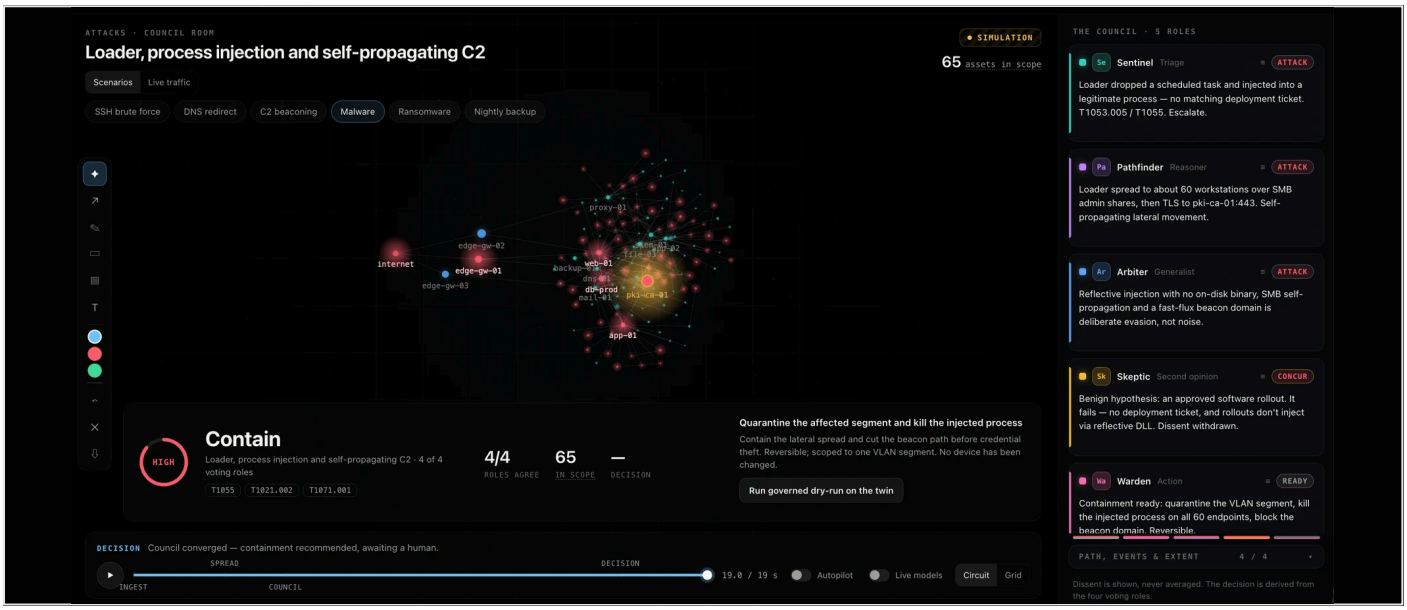
BREAKWATER ASOC / AGENTIC AI SECURITY OPERATIONS CENTER

Agentic AI investigation. Human-controlled response.

Breakwater Response brings the Council's multiple AI reasoning roles into the ASOC investigation. Challenge the attack story, rehearse response options and prepare an evidence-backed proposal for human authorization.



Response / Scripted Council investigation and response proposal.

Investigate together

Inspect AI answers with their tools and source references.

Challenge the story

Compare explanations, dissent and missing evidence.

Prepare the response

Review target, impact, permissions and recovery.

From AI investigation to an action you can govern

Connect grounded investigation and Council reasoning to a safety case: the evidence, target, proposed action and remaining risk. Review the intervention and its service impact before granting scoped authority.

The capture illustrates a Council scenario. Live model inference and operational execution are separate deployment capabilities.

RESPONSE / AGENTIC INVESTIGATION AND GOVERNED ACTION

Investigate with context. Act within defined authority.

Grounded AI prepares the response. Your team governs its scope, approval, verification and recovery.

THE GOVERNED PATH / SCOPE AND EVIDENCE TRAVEL WITH THE ACTION

01	02	03	04
Investigate	**Rehearse**	**Authorize**	**Verify / recover**
AI examines the case	Check the intervention	Review the safety case	Check service behavior

01 Ground the investigation / Challenge the explanation

AI hunting and Chat expose available tools and references. Sentinel, Pathfinder, Arbiter, Skeptic and Warden add distinct reasoning perspectives to challenge the attack story.

02 Rehearse the decision / Preserve the service

Use a digital twin or supported lab exercise to examine the intervention, required service behavior and recovery. Carry the result and its assumptions into the safety case.

03 Bind authority / Recheck before execution

The governed path checks approval expiry and matching evidence, actions and targets. Impact classes determine required reviewers. Supported executors apply the action, check health and compensate when required.

Confirm model backend, evidence connectors, approval classes and executor scope. Operational drafts, scripted scenarios, lab actions and live site execution are distinct paths.

See Breakwater in action.

Request a demonstration focused on your priorities.

Explore Response **Request a demo**