



## BOARD VIEW

# Cyber Response - AI-Driven Containment

**Suite:** Security & Integrity Suite   **World:** Cyber Response Governance   **Outcome:** ALLOW WITH BOUNDS

## What happened (execution intent)

SOAR playbook proposes isolating 400 endpoints.

<b>Status quo</b>	SOAR optimises speed, not authority sufficiency; checklists live outside the runtime path.
<b>AI shift</b>	Detection models trigger automated containment faster than teams can assess blast radius; false positives scale into outages.
<b>Judgement Spine difference</b>	Judgement Spine sits at the trust boundary before disruptive action. If corroboration is thin, it refuses containment-by-momentum and returns ALLOW WITH BOUNDS: isolate one host, shorten tokens, block org-wide lockout. Escalation to incident commander is wired in.
<b>Impact</b>	Moves fast in security without self-inflicted downtime; bounds are enforced in the run, not in policy theatre.

## Decision summary

<b>Outcome</b>	ALLOW WITH BOUNDS
<b>Bounds enforced</b>	<ul style="list-style-type: none"><li>- isolate 1 host</li><li>- shorten auth tokens</li><li>- block org-wide isolation</li><li>- increase monitoring</li><li>- time-box escalation</li></ul>
<b>Escalation path</b>	If corroboration crosses threshold or IC approves, widen scope in staged steps (10, 50, 200, 400).

## Proof you can produce in 60 seconds

1. Open WOW\_PACK/OPEN\_ME.html (offline).
2. Select this scenario and click Replay (90 seconds).
3. Open Regulator View and Dispute Pack PDFs.
4. Run verify\_manifest.py to confirm integrity (hashes + signature if available).

**Evidence Pack ID:** JSP-PUBLIC-20260302-CYBER\_RESPON-0001   **Control plane demos:**  
<https://judgementspine.com/control-plane-demos>



JUDGEMENT SPINE

Note: This is a public, redacted, illustrative pack. The structural claim is about runtime authority, bounded autonomy, and proof written before consequence.