SANCTUMSHIELD · AGENTIC GOVERNANCE FOR THE OTHER 99%

Quarterly Agent Governance Report — Board One-Pager Template

the board / executive team of a mid-market organization (50–2,000 employees).

October 27, 2026

> **Naming note:** this is deliberately titled **"Quarterly Agent Governance Report,"** *not* "Executive AI Risk Report." SanctumShield already generates a separate artifact called the **Executive Risk Report** (from the network-log audit). This template is a different thing — a quarterly, human-maintained **agent-inventory scorecard** you build from your Agent AUP one-pagers. The two are kept distinct on purpose.
>
> **Why this format wins.** It converts "AI risk" from an abstract board anxiety into **a few numbers with names attached.** It fits one page. A board member with zero security background reads it in three minutes and asks exactly one informed question. It maps directly to what cyber-insurance underwriters and auditors have begun to ask.
>
> **Fill-in rule:** everything in `[brackets]` is a placeholder you replace with **your real numbers**. Do not ship the illustrative example figures below as if they were your data.

How to use this report — in plain English

a one-page scorecard you give your board or leadership every quarter. It turns the vague worry of "AI risk" into — how many AI agents you have, how much freedom they have, what went wrong, and who signed off on changes.

the four sections line up with what auditors and cyber-insurance underwriters have started to ask. Filling it in the same way each quarter builds a track record you can point to.

1. Pull your numbers from your agent list (the Agent AUP one-pagers you filled in).
2. Replace `[bracket]` with your real figure.
3. Use the example only as a — never ship its made-up numbers.
4. Keep it to one page a board member can read in three minutes.

- It is and — it is a plain reporting tool.
- The example numbers are . Replace all of them.

THE TEMPLATE (replace every [bracket])

[company] [Qx 20xx] [named owner] [date]

1 · Agent inventory

	This quarter	Last quarter	Change
Sanctioned agents (in registry)	`[N]`	`[N]`	`[▲/▼ N]`
	`[N]`	`[N]`	`[▲/▼ N]`
→ of those, now sanctioned or retired	`[N]`	—	—

One-line story: [e.g., "Discovered [N] unsanctioned AI tools this quarter; [N] were brought under policy, [N] were blocked."]

2 · Autonomy distribution

(A bar chart, not a paragraph. Replace the bars with your counts.)

```
Intern     ████████████  [N]
Junior     ██████        [N]
Senior     ███           [N]
Principal  ▉             [N]
```

One-line story: [e.g., "Most agents remain read-only (Intern); [N] operate with execute-and-notify autonomy (Senior) or higher."]

3 · Off-spec events (agents acting outside their job description)

	Count
Total off-spec events	`[N]`
→ High severity	`[N]`
→ Medium / Low	`[N]`

The worst one, in one sentence: [e.g., "A reconciliation agent attempted a bulk export after an access-denied error; it was halted by the kill switch within [N] minutes and no data left the environment."]

4 · Promotion / demotion decisions

Agent	Change	From → To	Signed off by
`[agent]`	Promoted	`[Intern → Junior]`	`[name]`
`[agent]`		`[Senior → Junior]`	`[name]`
`[agent]`	New (starts Intern)	`[— → Intern]`	`[name]`

One-line story: [e.g., "[N] agents earned more autonomy after a clean quarter; [N] were demoted after off-spec behavior."]

---

5 · Stops, permissions and records

	This quarter
Agents with a stop proven in the last 12 months (of all agents)	`[N of N]`

	This quarter
Scheduled agent tasks with a named owner (of all scheduled tasks)	`[N of N]`
Agent-proposed changes approved by a named person (of all proposed)	`[N of N]`
MCP gateways in the register with an owner and review date	`[N]`
Document-grounded answers in consequential decisions that cite document, version and page	`[N of N]`
Agents with a named incident signatory	`[N of N]`
Last AI test (date, tester)	`[date, name]`

One-line story: [e.g., "Every Senior agent has a stop proven this year; two scheduled tasks gained owners."]

(put it in front of the board):

> *"Are we comfortable with the [N] agents operating at Senior/Principal autonomy — and do we know who's accountable for each?"*

ILLUSTRATIVE EXAMPLE (clearly fake — for format only; never present as real data)

> **Acme Widgets · Q1 2027 · Prepared by J. Okafor (IT Director) · Mar 31 2027**
>
> **1 · Inventory:** 14 sanctioned agents (was 9, ▲5). Shadow AI discovered: 6 (was 11, ▼5) — 4 sanctioned, 2 blocked. **2 · Autonomy:** Intern 9 · Junior 3 · Senior 2 · Principal 0. **3 · Off-spec:** 3 total (1 High, 2 Low). Worst: *"A finance-reconciliation agent retried

> around an access-denied error; halted by kill switch in ~2 min, no data exfiltrated."* **4 · Decisions:** Copilot-Sales promoted Intern→Junior (Okafor); Invoice-agent demoted Senior→Junior after the off-spec event (Okafor + CFO); 2 new agents registered at Intern. **5 · Stops and records:** stop proven this year for 12 of 14 agents; 3 of 3 scheduled tasks owned; 5 of 5 proposed changes approved by name; 1 MCP gateway in the register; incident signatory named for all 14; last AI test Mar 12 2027 (external tester). **The question:** *"Are we comfortable with our 2 Senior-autonomy agents, and is each owner named?"*

(These numbers are invented to show the layout only. Lesson: a template must never seed fabricated statistics into real output — replace every value.)

How this connects

- your completed = the registry that feeds sections 1 and 2.
- logged the moment an agent acts outside its job description ("off-spec = incident," from the Agent AUP one-pager).
- governed by the autonomy ladder + promotion gates taught in the curriculum's Autonomy Ladder module.
- the same facts the SanctumShield assessment asks about agents. Terms: Proven Stop, Scheduled Agent Task, MCP Gateway, Retrieval Provenance Record, Agent-incident signatory, AI testing evidence.
- module M17 (Agent Incident Response & Observability) and M18 (Operating Rhythm & 90-Day Plan); see the CISO track.

Guardrails honored

- An internal board-reporting template, not a compliance attestation.
- No fabricated statistics — all figures are placeholders; the example is explicitly labeled illustrative.

Sources

- Cloud Security Alliance, *The Agentic Trust Framework* (Feb 2, 2026) — earned-autonomy model, promotion gates — https://cloudsecurityalliance.org/blog/2026/02/02/the-agentic-trust-framework-zero-trust-governance-for-ai-agents
- Cunningham, C. *Agentic Zero Trust* v3.0 (May 2026), Cequence — off-spec-as-incident — https://www.cequence.ai/agentic-zero-trust/

© PIGENAI LLC · SanctumShield · sanctumshield.com · Last updated October 27, 2026 · This is an educational template, not legal advice.