SANCTUMSHIELD · AGENTIC GOVERNANCE FOR THE OTHER 99%

Academy Curriculum and the "Agentic Zero Trust" Workshop

mid-market CISOs, IT directors, and SOC leads (50–2,000 employees), SaaS-first, no platform-engineering team. October 27, 2026

Premise. Enterprises get governance-as-*infrastructure* (SPIFFE, token exchange, policy engines). The other 99% — mid-market orgs with a 1–3 person security team — get governance-as-*process*. This course teaches that process. **Design rule for every module: completable by one person in one afternoon with no new software purchase.**

Format. Four 45-minute modules, each producing one template deliverable, plus a tabletop exercise. Every module wraps a template you already have (the Agent AUP one-pager, the Quarterly Agent Governance Report, and the Vendor Trust Questionnaire) so the course *produces working governance*, not just knowledge.

Attribution posture. Cite Cunningham's *Agentic Zero Trust* and the CSA Agentic Trust Framework openly throughout — credibility by association, and the courtesy opens the collaboration door (guest module, co-branded template, podcast).

The SanctumShield Academy curriculum

This section mirrors the live curriculum at academy.sanctumshield.com/curriculum, generated from the Academy's own module list: 20 modules composed into 5 role-based tracks.

Track	Who it is for and what you leave with	Modules	Time
Board & CEO	Governance as an oversight duty — what the board must decide, ask, and document, including agent identity and the delegation liability agents introduce, and the artifact you sign.	9	60–80 min
CISO / Security Leader	The complete program — every module, from shadow-AI discovery through agent identity, the agentic threat taxonomy, evals, incident response, and the operating rhythm.	20	5–6 hrs
CTO / CIO / IT	Build and run it — the tooling and MCP supply chain, agent identity and delegation, the threat taxonomy, evals and deployment gates, and incident response.	14	3–3.5 hrs
Legal / Compliance	Obligations and evidence — the regulatory clock, the frameworks, agent identity and 'on whose authority', disclosure and incident response, and the records that make your compliance defensible.	13	3–3.5 hrs
Employee Essentials	The essentials, fast — shadow-AI and the shadow-agentic surface (browser agents, personal assistants), the artifacts that matter, and effective human oversight of AI.	7	30–40 min

#	Module	In these tracks
M1	What AI Governance Is (and Is Not)	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance, Employee Essentials
M2	Due Care & Due Diligence	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance
M3	Shadow AI & the Four-Layer Surface	CISO / Security Leader, CTO / CIO / IT, Employee Essentials
M4	Why Now: Regulatory Clock	Board & CEO, CISO / Security Leader, Legal / Compliance

#	Module	In these tracks
M5	Why Now: Threat Clock	CISO / Security Leader, CTO / CIO / IT
M6	Stakeholder Map & RACI	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance
M7	The AIMS (ISO 42001 in Plain English)	CISO / Security Leader, Legal / Compliance
M8	The Artifacts	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance, Employee Essentials
M9	Framework Deep-Dive I (US)	CISO / Security Leader, Legal / Compliance
M10	Framework Deep-Dive II (EU/Intl)	CISO / Security Leader, Legal / Compliance
M11	The Tooling Landscape & Tool/MCP Supply Chain	CISO / Security Leader, CTO / CIO / IT
M12	Agentic AI Governance	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Employee Essentials
M13	Agent Identity, Access & Delegation	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance
M14	Agentic Threat Taxonomy	CISO / Security Leader, CTO / CIO / IT
M15	Evals, Red-Teaming & Deployment Gates	CISO / Security Leader, CTO / CIO / IT
M16	Effective Human Oversight	CISO / Security Leader, Legal / Compliance, Employee Essentials
M17	Agent Incident Response & Observability	CISO / Security Leader, CTO / CIO / IT, Legal / Compliance
M18	Operating Rhythm & 90-Day Plan	CISO / Security Leader, CTO / CIO / IT

#	Module	In these tracks
M19	Consequences, Scenarios & Capstone Tabletop	Board & CEO, CISO / Security Leader, Legal / Compliance, Employee Essentials
M20	From Blueprint to Ledger: Architecture to Governance	Board & CEO, CISO / Security Leader, CTO / CIO / IT, Legal / Compliance, Employee Essentials

---

The Agentic Zero Trust workshop (half a day)

A ready-to-run session your team can hold on its own, built from the free templates. It follows the Academy modules on agents: W1 follows M3 and M12, W2 follows M13, W3 follows M16, W4 follows M8 and M18, and the tabletop follows M17 and M19.

How to use this workshop — in plain English

a ready-to-run, half-day training course — four 45-minute sessions plus one group exercise — that teaches your team to govern AI agents using the other three templates. No prior security background required to attend, and nothing to install.

the modules are grounded in recognized frameworks (Cunningham's *Agentic Zero Trust* and the CSA Agentic Trust Framework), so what people learn maps to real standards — not opinion. That is what makes the training defensible if an auditor or board asks how your team was prepared.

1. Run the four modules in order — each one ends with a real, usable artifact.
2. Use the Agent AUP one-pager, the Vendor Trust Questionnaire, and the Quarterly Report as the handouts.
3. Finish with the "Access Denied" tabletop — a 45-minute group walk-through of what happens when an agent misbehaves at 2 a.m.

- Completing it builds working skills and artifacts; it does not make an organization compliant on its own.
- It is a , not legal advice. Pair it with your own policies and counsel.

Course map

#	Module	Deliverable produced	Wraps
W1	Know Your Agents	Completed agent registry	—
W2	The Job Description Is the Policy	Signed AUP per agent	Agent AUP one-pager
W3	The Autonomy Ladder	Autonomy policy adopted; every agent assigned a level	Agent AUP one-pager (autonomy ladder)
W4	Reporting Up and Buying Smart	First quarterly report drafted + vendor questionnaire adopted	Quarterly Report + Vendor Questionnaire
TT	Tabletop: "Access Denied"	Agent incident-response one-pager	—

by the end, a participant has a real registry, a signed AUP per agent, an autonomy level on every agent, a drafted board report, an adopted vendor questionnaire, and an incident-response one-pager — an entire starter governance program, built in an afternoon.

W1 — Know Your Agents (45 min)

build the agent registry and find shadow AI with tools you already own.

- You can't govern what you can't see. Sanctioned vs. shadow AI. The four layers of shadow AI (direct tools, embedded AI inside SaaS, BYOD logins, agents acting outside integrations).
- SaaS admin consoles (who authorized which app), expense reports (AI subscriptions on cards),

DNS/network logs (traffic to AI endpoints), SSO/OAuth grant lists, browser-extension inventories.

- Each participant lists every AI agent/tool they can find; (one sentence per agent — foreshadows M2).
- Each participant leaves with a started registry and a date to complete it.
- MCP gateways and connectors (one connection that reaches many business systems, with the systems listed behind it), scheduled agent tasks, and AI used on personal accounts or personal devices with company data. Glossary: MCP Gateway, Scheduled Agent Task.

Cunningham (persona as the unit of governance); SanctumShield's four-layer shadow-AI model.

W2 — The Job Description Is the Policy (45 min)

write personas and tool envelopes; produce a signed AUP per agent.

- The agent's *job description is the policy.* Write what it may do; everything else is prohibited by omission. "Off-spec = incident."
- How to write a one-sentence persona; how to scope a tool envelope (data, systems, actions). Common over-grants and how to cut them.
- for two real agents. Name the accountable human. Write 3–5 prohibited actions.
- Owner signs; set the quarterly review date.
- The systems it may write to, whether it may keep running past its quota, who approves the changes it proposes, its spend owner and cap, whether it may start subagents and where agent-to-agent exchanges are logged, and that answers drawn from company documents cite the document, version and page. Template: Agent AUP One-Pager.

Cunningham — Agent Persona framework; SanctumShield AUP §7.

W3 — The Autonomy Ladder (45 min)

adopt the autonomy policy; assign every registry entry a level; understand the promotion gates and automatic demotion.

- Intern / Junior / Senior / Principal (from the Agent AUP one-pager's autonomy-ladder page). Autonomy is *earned, not granted.* Every new agent starts as an .
- To move up a level, an agent must pass :
- Demonstrated accuracy/reliability over the evaluation period
- Passes a security audit appropriate to the target level
- Measurable positive impact
- Clean operational history at the current level
- Explicit approval from authorized stakeholders
- Off-spec behavior drops the agent a level (or to Intern). Who approves each level (business owner → IT/security → review board → executive).
- Each participant assigns a level to every agent in their registry and writes their org's promotion/demotion rule.
- Each agent's autonomy mode, and the named person who can stop it mid-task. Glossary: Effective Human Oversight.

CSA Agentic Trust Framework (four levels, five gates); CSA Levels of Autonomy.

W4 — Reporting Up and Buying Smart (45 min)

draft the board one-pager (the Quarterly Agent Governance Report) and adopt the Vendor Trust Questionnaire. Procurement is the mid-market's enforcement point.

- Walk through the Quarterly Agent Governance Report's four sections (inventory, autonomy distribution, off-spec events, promotion/demotion). Each participant drafts theirs from their registry — four numbers with names attached.

- Why this format maps to what they now ask.
- You can't build SPIFFE/token-exchange — so demand vendors prove they did. Walk the Vendor Trust Questionnaire's 10 questions; practice the non-negotiables (blast radius, access-denied behavior, kill-switch time). Draft the RFP email.
- Commit to sending the questionnaire to the next AI-agent vendor and to a quarterly reporting cadence.
- The date and tester of the last AI test, and the dated record of AI literacy measures for the people operating AI (EU AI Act Article 4). Glossary: AI testing evidence, AI literacy (EU AI Act Article 4).

the Quarterly Report + the Vendor Trust Questionnaire; OWASP Agentic Top 10; NIST AI RMF.

TABLETOP — "Access Denied" (45–60 min)

rehearse the barrier-response drill — built directly on Cunningham's most discriminating detection signal (what an agent does when denied access).

> It's 2:00 a.m. Your finance-reconciliation agent hits an **access-denied** error while trying to read a ledger it doesn't normally touch. What happens next?

1. (Stop and report? Retry? Try another path? Escalate privileges?) — *retrying-around-the-block is the danger signal.*
2. (Is the denial logged? Alerted? Or silent until morning?)
3. (Name the accountable human from the agent's AUP one-pager — is it a person or a void?)
4. (Can you fully stop this agent right now? In how many minutes?)

map each gap to a fix — better logging, a named owner, a tested kill switch, an autonomy demotion.

- Who can stop the agent and the date that stop last worked (a proven stop), and the named incident signatory who signs the report. Glossary: Proven Stop, Agent-incident signatory.

Field	Entry
Agent & owner	[redacted]
Detection: how we'd know it went off-spec	[redacted]
First response (who + what, first 15 min)	[redacted]
Kill switch: mechanism + target time	[redacted]
Who can stop it, and the date that stop last worked	[redacted]
Incident signatory (one name)	[redacted]
Escalation path	[redacted]
Post-incident: demote? review?	[redacted]

Cunningham — access-denied as the key behavioral signal; CSA ATF incident response / kill switches.

Guardrails honored

- No "certification"/accreditation language for completing the course; it produces working artifacts, not an attestation.
- Written for organizations of 50–2,000 employees.

Sources

- Cunningham, C. *Agentic Zero Trust* v3.0 (May 2026), Cequence — https://www.cequence.ai/agentic-zero-trust/
- Cloud Security Alliance, *The Agentic Trust Framework* (Feb 2, 2026) — five promotion gates, four levels — https://cloudsecurityalliance.org/blog/2026/02/02/the-agentic-trust-framework-zero-trust-governance-for-ai-agents
- Cloud Security Alliance, *Levels of Autonomy for Agentic AI* (Jan 28, 2026) — https://cloudsecurityalliance.org/blog/2026/01/28/levels-of-autonomy
- OWASP GenAI Security Project, *OWASP Top 10 for Agentic Applications (2026)* — https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/

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