SOPHIZO

INDUSTRY PLAYBOOK · TECHNOLOGY & SOFTWARE

If you cannot govern your own AI, you cannot sell it.

Agentic AI for incident response, code velocity, and internal governance you can defend in front of an enterprise buyer.

PUBLISHED 2026 · **SOPHIZO RESEARCH**

Fractional AI & RevOps Leadership
sophizo.net

Tech companies are first adopters and first cautionary tales.

Tool sprawl is the silent tax of every engineering org. Tickets pile up, security alerts pile up, code reviews pile up. And the answer until now has been to hire faster. Agentic AI changes the math, but the same agents that accelerate your roadmap can become the audit liability that costs you your next enterprise deal.

Why this matters now

Building agents you would be willing to ship to a regulated customer is the same project as building agents you would be willing to ship to yourself. An AI Officer's job is to make those two projects identical.

> *You sell AI to enterprise buyers. Could you pass your own AI vendor questionnaire?*

Software vendors are the supply chain. Your customers are downloading your AI risk along with your product.

Beyond the six global regimes, technology & software carries the overlays below. Each one has its own enforcement model and its own evidence expectation.

EU CRA EU Cyber Resilience Act + AI Act for software

EUROPEAN UNION HIGH

Applies to. Any product with digital elements placed on the EU market. Includes AI-enabled software.

Key obligation. Security by design, vulnerability handling for the product lifetime or 5 years, CE marking, SBOM disclosure.

Evidence. Technical documentation, SBOM, vulnerability disclosure procedures, conformity assessment.

SOC 2 AI AICPA SOC 2 + AI Trust Services Criteria

UNITED STATES, AICPA HIGH

Applies to. SaaS vendors with AI features being evaluated by enterprise procurement.

Key obligation. Extend Trust Services Criteria to cover AI-specific controls. Increasingly demanded in vendor security questionnaires.

Evidence. SOC 2 Type II report with AI control narrative, AI governance policy, control test results.

US State AI US State AI Laws (CO SB 205, CA AB 2013, TX TRAIGA)

UNITED STATES, STATE-LEVEL HIGH

Applies to. Software using AI in consequential decisions for residents of regulated states.

Key obligation. Pre-deployment impact assessments, transparency to consumers, opt-out where applicable. Colorado AI Act effective Feb 2026.

Evidence. Impact assessments, consumer disclosures, opt-out records, model documentation.

Open Source AI Open Source AI Licensing + EU AI Act FOSS Carve-Outs

GLOBAL MODERATE

Applies to. Companies distributing or using open-source AI models in commercial products.

Key obligation. OSAID compliance, model card publication, dataset documentation, EU AI Act FOSS exemption boundaries.

Evidence. Model cards, dataset documentation, license compliance review.

Global emphasis for this sector. EU AI Act · ISO/IEC 42001 · NIST AI RMF.

Four capability domains. One operating layer.

01 Incident Response and SecOps

- Triage, containment, and case-prep agents
- Adversary-aware playbook execution
- SOAR augmentation, not replacement
- Identity-aware agent permissions

02 Code Generation, Test and Deploy

- PR drafting with style and security checks
- Test generation and flake remediation
- Deploy agents with rollback policy
- Review-queue triage

> *Copilot saved your engineers 14% of their time. Where did that 14% land in the P&L?*

Capability domains, continued.

03 Internal Workflow Automation

- Eng/IT helpdesk autonomous resolution
- Procurement and access-request agents
- Runbook execution and documentation
- Onboarding and offboarding orchestration

04 Customer-Facing AI Governance

- Model and prompt registries customers can audit
- Red-team-as-a-service for shipped features
- Data-residency and consent orchestration
- Disclosures and policy artifacts

What production deployments look like at scale.

50%+
FASTER INCIDENT RESOLUTION

200%+
PRODUCTION ROI IN MATURE DEPLOYMENTS

6 to 12 mo
FASTEST PAYBACK IN ANY INDUSTRY

Production-stage benchmarks compiled from GitHub Octoverse, Stack Overflow Developer Survey, and Linux Foundation Research and CNCF benchmarks (2024 to 2025). Your spread depends on test coverage baseline, deploy-cadence discipline, and how mature your model registry already is.

The AI Officer Mandate.

Three responsibilities a Fractional AI Officer owns from day one in technology & software.

01
Internal governance as the role model. Your customers will demand what you practice.

02
Secure-by-design agent platforms with identity, secrets, and blast-radius controls.

03
Red-team and assurance that scales with shipping velocity.

06 · A 90-DAY PATH

How a Sophizo engagement starts in Technology & Software.

DAYS 1 TO 30

Diagnose

MAP THE OPERATING REALITY

- Agent inventory: shadow pilots, vendor agents, internal prototypes
- Customer AI vendor questionnaire dry-run against your own posture
- Kill-switch and identity audit on every agent in production
- Model registry baseline and red-team cadence assessment

DAYS 31 TO 60

Architect

DESIGN THE AUTONOMY BOUNDARY

- Agent permissions and escalation policy
- Evidence file and audit trail design
- First production pilot scoped with rollback plan
- Cross-functional governance committee charter

DAYS 61 TO 90

Operate

SHIP AND INSTRUMENT

- First agent in production with HITL controls
- Operator coaching and policy refinement
- P&L instrumentation by use case
- Quarterly review cadence established

What we will not do.

We do not run your eng org, set your roadmap, or replace your Head of AI. We do not ship customer-facing agents without a model registry, a kill switch, and a published red-team cadence. Even if it means slipping your launch. We pass on companies where the CEO sees AI as marketing, because the customers who matter are about to ask hard questions and the company will not be ready to answer them.

07 · WHAT THE BOARD NEEDS TO HEAR

Five things the board needs to hear about AI in software companies.

Five cited insights for the next risk-committee meeting. Each one is sourced. Each one is what an experienced AI Officer would put in front of the board if they walked in tomorrow.

01 · DEVELOPER LEVERAGE

Agent-assisted developers ship measurably faster. The variance is real.

GitHub's 2024 enterprise productivity study reports a 26 percent increase in pull requests per developer. BCG's September 2024 GenAI study found agent-assisted developers complete coding tasks 55 percent faster. The variance depends entirely on task type and codebase maturity. Treat the benchmarks as ranges, not promises.

Source. GitHub Octoverse 2024 Enterprise Productivity Study; BCG GenAI Software Development Study, September 2024.

02 · RFP GATE

Enterprise AI procurement now demands attestation.

Gartner's 2024 RFP analysis indicates roughly 60 percent of Fortune 500 RFPs for AI-enabled software now require evidence of NIST AI RMF alignment or ISO 42001 readiness. Vendors without an attestation pack lose deals at the security review, not at the product demo.

Source. Gartner 2024 AI RFP and Procurement Benchmark; ISO/IEC 42001:2023 AI Management Systems Standard.

03 · SBOM EXPANSION

Model and prompt registries are becoming AI's SBOM.

Federal procurement under CISA's SBOM guidance is extending to model provenance. Executive Order 14110 is in active review, but the underlying registry expectations have not retreated. Software companies that already maintain a model registry will move through federal security reviews materially faster.

Source. White House Executive Order 14110, October 2023; CISA Software Bill of Materials guidance, 2024.

Two more, then the framework.

04 · CLOUD BENCHMARK

Agent gross margin is the new SaaS unit economic.

Bessemer's 2024 State of the Cloud and Battery's OpenCloud Index both track agent gross margin as a leading indicator of valuation premium. The companies winning Series C and later rounds in 2025 are the ones who can show per-invocation cost of goods and segmented agent margin by cohort.

Source. Bessemer Venture Partners 2024 State of the Cloud Report; Battery Ventures 2024 OpenCloud Index.

05 · BUYER BEHAVIOR

Customer agent governance is now a sales-enablement artifact.

Sapphire Ventures' 2024 B2B SaaS buyer survey found 73 percent of buyers require model risk attestation as a pre-purchase requirement. The model registry, the red-team report, and the NIST AI RMF crosswalk now belong in the sales deck, not the appendix of the security questionnaire.

Source. Sapphire Ventures 2024 B2B SaaS Buyer Survey; Vendr 2024 SaaS Buying Patterns Report.

The Agent Unit Economics Stack.

Every customer-facing agent has a P&L. Track it like a SKU, govern it like a regulated component, sell it like a feature with attached attestation. Anything less is a financial story your investors will eventually ask about.

LINE 1

Cost of goods

Per-invocation compute, third-party calls, and retrieval cost, fully loaded.

LINE 2

Margin by class

Per agent decision class, segmented by customer cohort.

LINE 3

Governance file

Model card, registry entry, validation report, and red-team artifact.

LINE 4

Customer attestation

NIST AI RMF crosswalk and SOC mapping, ready for buyer security review.

LINE 5

Production controls

Kill switch tested quarterly. Drift monitoring on every shipped feature.

From John Utley.

> *Software companies are first adopters and first cautionary tales for a reason. Ship an agent to a Fortune 500 customer without a model registry and a red-team artifact and you are not closing the deal. You are extending the security review by a quarter. Instrument before you launch.*

John Utley

FOUNDER, SOPHIZO · SEATTLE, WA

John Utley founded Sophizo to give growth-stage companies the AI and revenue architecture work historically reserved for the Fortune 500. He writes and advises on agentic AI governance, predictive forecasting, and operating-model design for boards and operators across technology & software and adjacent sectors.

Test your operating picture against these.

1

You sell AI to enterprise buyers. Could you pass your own AI vendor questionnaire?

2

Copilot saved your engineers 14% of their time. Where did that 14% land in the P&L?

3

How many of your shipped features have a model registry entry an enterprise buyer can audit?

Frequently asked questions.

We already use Copilot. Why bring in an AI Officer?

Copilot accelerates individuals. An AI Officer makes the company itself faster. By removing handoffs, governing the sprawl of pilots, and setting up the assurance posture you will need when an enterprise buyer asks for it. Tools alone do not change throughput at the org level.

How do you handle agent security?

Every agent runs with its own identity, scoped permissions, and a kill switch. We treat agents the same way we treat any other privileged service. And we red-team them on the same cadence.

Will agentic AI eat our developer headcount?

It changes what developers do, not how many you have. The orgs winning right now are using the freed capacity to ship more product, not to ship the same product with fewer people.

If this maps to your operating reality, we should talk.

The Diagnostic Sprint is two weeks. Board-ready output. Tailored to technology & software.

ENGAGE
sophizo.net/checkout/diagnostic-sprint

INDUSTRY PAGE
sophizo.net/industries/technology-software

EMAIL
john@s

Primary research behind this brief.

Every claim, statistic, and citation in this playbook traces back to one of the primary sources below. Pressure-test any of them with your team. We have done the same.

01. GitHub.
Octoverse 2024 Enterprise Productivity Study.

02. BCG.
GenAI Software Development Study, September 2024.

03. Gartner.
2024 AI RFP and Procurement Benchmark.

04. ISO/IEC.
42001:2023 AI Management Systems Standard.

05. White House.
Executive Order 14110 on Safe, Secure, and Trustworthy AI, October 2023.

06. CISA.
Software Bill of Materials guidance, 2024.

07. Bessemer Venture Partners.
2024 State of the Cloud Report.

08. Battery Ventures.
2024 OpenCloud Index.

09. Sapphire Ventures.
2024 B2B SaaS Buyer Survey.

Editorial note. This brief is a field reference compiled by Sophizo Research. It is not legal, accounting, or clinical advice. Cite the primary regulator guidance for binding interpretation. Where statistics are quoted, the most recent published figure as of early 2026 is used.