AEGYS

CONFIDENTIAL

Security & AI Posture Assessment

Prepared for Acme Corp

Where your security program stands, how ready you are to adopt AI, and what to fix first.

Acme Corp (~400 employees, B2B software) · May 2026 · Confidential
aegys.io · secure@aegys.io

Executive summary

Acme runs a credible, growing security function with real strengths in identity and backups. What it lacks is connective tissue: clear ownership of security decisions, a response plan that has actually been exercised, and any guardrails for the AI tools already in daily use across the company. The fastest risk reduction here is organizational, not technical, and almost none of the top priorities require new spend.

Posture at a glance

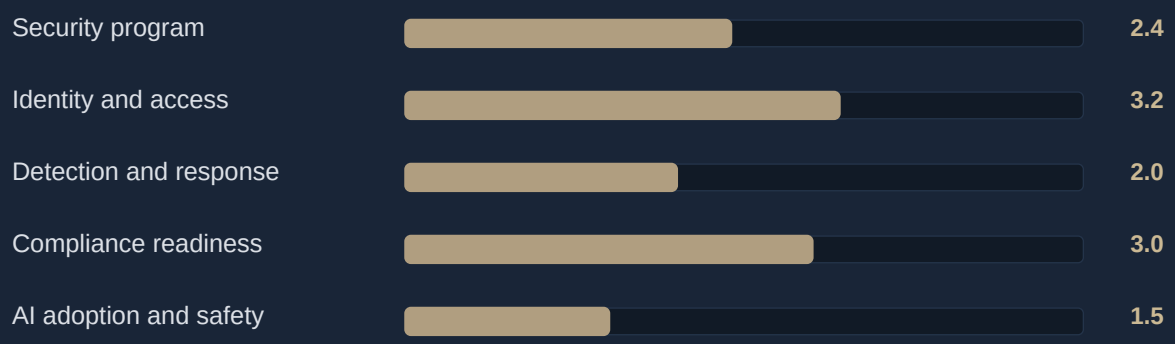

Maturity by domain on a 0 to 5 scale, weighted to what is verified in practice.

Where to start

- **Name an owner.** Assign one accountable owner for security and AI decisions, with a monthly cadence. Most gaps below trace back to this.
- **Publish the one-page AI guardrail** and turn off training-on-your-data in the two tools that default to it. This closes the largest data-exposure risk in a week.
- **Close five CIS IG1 gaps** (endpoint logging, account cleanup, recovery testing). Low cost, high coverage against the attacks you are most likely to see.

Scope and approach

This readout reflects a three-week review of Acme's security program, AI usage, and compliance readiness: interviews with IT and engineering leads, a no-blame shadow-AI survey across teams, a review of identity, logging, and backup configurations, and a mapping of findings to the frameworks below.

What we assess	Security program maturity, AI adoption and shadow-AI exposure, and compliance readiness.
Frameworks	NIST CSF 2.0 and NIST AI RMF 1.0; MITRE ATLAS and the OWASP LLM Top 10 for AI threats; CIS Controls v8 (IG1) and SOC 2 for baseline and compliance.
Maturity scale	1 Initial, 2 Developing, 3 Defined, 4 Managed, 5 Optimized. Ratings reflect what is real and verified, not what is documented.

Security program maturity (NIST CSF 2.0)

Maturity across the six NIST CSF 2.0 functions. The headline is consistent fundamentals undercut by missing ownership and unexercised response.

Function	Maturity	Key finding and priority action
Govern	Developing (2/5)	Security decisions are made informally with no named owner for risk acceptance. Stand up a lightweight monthly governance cadence; this is the keystone fix.
Identify	Defined (3/5)	Asset and data inventories exist but are partially stale, and the vendor inventory is incomplete. Refresh and assign an owner.
Protect	Defined (3/5)	Multi-factor authentication covers most external apps; a few legacy systems and service accounts remain exceptions, and patching is reactive. Close the exceptions.
Detect	Developing (2/5)	Centralized logging covers cloud but not endpoints, and alert ownership is unclear. Extend logging to endpoints and name an owner.
Respond	Developing (2/5)	An incident response plan exists on paper but has never been exercised. Run a tabletop in the next quarter.
Recover	Defined (3/5)	Backups run and are mostly tested; recovery-time objectives are undocumented. Set and test RTOs for the top systems.

AI adoption and shadow-AI exposure

AI is already in daily use at Acme well ahead of any policy. Exposure is rated against NIST AI RMF, with technique context from MITRE ATLAS and the OWASP LLM Top 10.

Area	Exposure	What we found, and the next step
Shadow AI	High	The survey found nine AI tools in active use across teams; only one is on a business tier with training opt-out. Consolidate to two sanctioned tools.
Data controls	High	No standard governs what data may enter AI tools, and at least two default to training on input (OWASP LLM02). Turn training off and publish the guardrail.
AI in security and IT	Early	Assistants are used ad hoc for drafting, with no sanctioned workflow. Pick one high-volume, low-stakes task and ship a reviewed win.
Guardrail and oversight	Absent	There is no one-page acceptable-use guidance and no named person to ask. This is the single fastest thing to fix.

The pattern we see most

Acme's AI risk is not exotic. It is ordinary data exposure: capable people pasting company information into tools nobody vetted, because the safe path was never made the easy one. The fix is a guardrail and two sanctioned tools, not a ban.

Compliance posture

Compliance is read as a byproduct of real controls, not a separate program. Acme is in reasonable shape with clear, finite gaps.

SOC 2	Targeting a Type II within 12 months. Roughly 60 percent of common controls are evidenced informally; the gap is documentation and consistency, not capability.
CIS IG1	Thirteen of eighteen IG1 safeguards are substantially in place. Gaps cluster in centralized logging, account management, and data-recovery testing.
PCI DSS	Currently out of scope: Acme stores no cardholder data. Confirm with a quick data-flow review before any payments feature ships.

Prioritized roadmap

A sequence, not a backlog. Each phase is achievable with the team Acme already has.

Now (0 to 30 days)	• Name one accountable owner for security and AI decisions. • Publish the one-page Green / Yellow / Red AI guardrail. • Turn off training-on-your-data in the two flagged tools. • Enable MFA on the remaining legacy exceptions.
Next (30 to 90 days)	• Consolidate to two sanctioned AI tools with single sign-on. • Extend centralized logging to endpoints; assign alert ownership. • Run a tabletop incident-response exercise. • Close the five open CIS IG1 safeguards.
Later (90+ days)	• Formalize the SOC 2 program around controls that now exist. • Set and test recovery-time objectives for top systems. • Establish a quarterly AI and vendor-risk review.

How we would run this with you

A Security and AI Posture Assessment takes three to four weeks: a short discovery, a shadow-AI survey, a controls review, and this readout. Reach us at secure@aegys.io · aegys.io

Method and references

Ratings are AEGYS's assessment judgment against the frameworks below, on a 1 to 5 maturity scale, weighted toward what is verified in practice rather than what is written in a policy. This document is a point-in-time readout, not an audit or a certification.

National Institute of Standards and Technology (2024). *Cybersecurity Framework (CSF) 2.0.* NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST AI 100-1. https://doi.org/10.6028/NIST.AI.100-1

MITRE. *ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems.* https://atlas.mitre.org

OWASP. *Top 10 for LLM Applications (2025).* https://genai.owasp.org/llm-top-10/

Center for Internet Security. *CIS Critical Security Controls v8, Implementation Group 1.* https://www.cisecurity.org/controls

AICPA. *SOC 2, Trust Services Criteria.* https://www.aicpa-cima.com