

CONTINUOUS COMPLIANCE FOR THE AI ERA

Audit-ready **every day,** not once a year.

Vendor risk, BAAs, AI governance and audit evidence that collects itself, across your code, cloud, vendors and AI agents. Findings become pull requests. Proof becomes a live Trust Center.

SOC 2

HIPAA

ISO 27001

EU AI ACT

CMMC L2

FEDRAMP

15+ FRAMEWORKS · 196 CONTROLS

THREE TEAMS, ONE RECORD

The same platform, from **whichever side** of the questionnaire you sit.

ENTERPRISE IT & GRC

Run third-party risk continuously

- One standard question set, every vendor, every framework
- Vendor status as an onboarding gate
- BAA · DPA · MSA · NDA register with reminders
- AI inventory and autonomy policy for every agent
- Vendor Watch: breach alerts, SOC 2 verification, Nth-party chain

VENDORS & STARTUPS

Close enterprise deals without the scramble

- Repos, cloud and databases graded against SOC 2 and HIPAA
- Findings fixed as pull requests by Sentinel
- Any questionnaire answered from live posture in a day
- Trust Center instead of a stale PDF
- Vetted auditors and pentesters when you want the real opinion

FEDERAL SUPPLIERS

CMMC Level 2 and FedRAMP readiness

- All 110 NIST 800-171 requirements, crosswalked to evidence you hold
- Live SPRS score, POA&M eligibility, 180-day clock
- SSP, POA&M and shared-responsibility exports, incl. OSCAL
- FedRAMP Rev 5 Low / Moderate / High readiness
- DFARS flow-down and subcontractor gaps on the register

Three artifacts built for a slower world

The annual questionnaire. Sent once, answered once, filed. The vendor changes its stack the following month and nobody asks again for eleven months.

The spreadsheet of dates. BAAs, DPAs, MSAs, NDAs — tracked by whoever inherited the sheet. The expired one surfaces during the deal, or during the investigation.

The evidence hunt. Six weeks before every audit, screenshots and emails are collected by hand to prove things that were true all year — but only provable in that window.

Now add AI. Vendors ship agents that act on your data. Your own teams run copilots on it. The old three artifacts have no row for 'what is this agent allowed to do?'



The cycle got shorter than the review

- **Vendors change monthly, reviews run yearly.**

Model swaps, new sub-processors, new agents — each one a material change to what you approved.

- **Auditors and customers now ask continuously.**

SOC 2 Type II, HITRUST, HIPAA and the EU AI Act all expect evidence over a period, not a snapshot.

- **The boundary is bigger than IT.**

Business units buy AI tools on a card. The first time compliance hears about it is the questionnaire from your customer.

- **Board-level accountability landed.**

NIST AI RMF, ISO 42001 and the EU AI Act name the accountable executive. Annual is no longer defensible.



Six moves — each one replaces a manual cycle

1. Assess every vendor against one standard question set

Mapped to SOC 2, HIPAA, HITRUST, ISO 27001, NIST 800-53 and the EU AI Act — so one answer serves every framework.

2. Gate onboarding until the vendor clears it

Status is a control, not a label: 'Assessing' blocks, 'Approved' unblocks.

3. Keep every agreement current automatically

BAA, DPA, MSA, NDA in one register with 30 / 14 / 7-day reminders.

4. Inventory every AI agent and hold it to a written autonomy policy

Yours and your vendors' — classified by risk, with a kill switch.

5. Let evidence collect itself

Controls monitored continuously; packages assembled from the record, not from screenshots.

6. Prove it on demand

A tamper-evident trail and a Trust Center — for the auditor, the customer and the board.



One standard question set for every vendor

- Every vendor answers the same bank — Operations, Security, Compliance, Best practices, AI governance
- Readiness per section at a glance: Security 8/9 signed off, Compliance 4/10 — you see where to push
- Criticality, PHI handling, SOC 2 expiry and next review on the same card
- Completed questionnaires, SOC 2 reports and pentests attached to the vendor, not an inbox



KollGuard AI

COMPLIANCE INTELLIGENCE

Search menu...

FAVORITES

Compliance Profile

Findings 75

Evidence Library

Reports

Audit Log

Agents 4

OVERVIEW

Dashboard

Daily Digest

CONTINUOUS MONITORING

Scans

Code Scanner

Sensitive Data

AI & AGENTS

AI Governance

Help & docs

Blog

Customize menu

Credentials encrypted in vault

KollGuard / Vendors

Scanned

Third-party register

LegacyLog Tool Observability	Low	Offboarded	High
Microsoft Azure Cloud	High	Approved	Low
MongoDB Atlas Database	High	Approved	Low
Okta Auth	Critical	Approved	Low
OpenAI AI	High	Approved	Medium
Salesforce CRM	Medium	Assessing	Medium
Sentry Observability	Low	Approved	Low
Stripe Payments	High	Approved	Low
Supabase Database	High	Assessing	Medium
Twilio Communication	Medium	Approved	Low

Supabase

Database

CRITICALITY

High

STATUS

Assessing

RISK RATING

Medium

SOC 2

Sep 12, 2026 (expired)

BAA

Signed

NEXT REVIEW

Jul 15, 2026

HANDLES PHI

Yes

CONTACT

—

Review overdue; BAA expiring soon.

Questionnaire readiness

All questionnaires

Supabase — 2026 vendor security assessment

SOC 2 · HIPAA Updated Sep 15, 2026

In Progress

20/33 signed off

4 drafted · 9 open

unlink Open questionnaire →

Operations

6/8 signed off

1 drafted · 1 open

Security

8/9 signed off

1 open

Compliance

4/10 signed off

2 drafted · 4 open

Best Practices

2/6 signed off

1 drafted · 3 open

Link an existing questionnaire...

↕ Link

Or import a new one and pick this vendor in the import form.

Close

Edit vendor


Vendors → Third-party register → Supabase

→ Same questions, every vendor, every year – comparable answers.



Questionnaires in, questionnaires out

- Issue the standard set to vendors — or paste the 300-question one a customer sent you
- KGAi drafts every answer from your real posture: approved policies, BAAs, personnel, dispositioned findings
- Each answer carries a confidence and cites the governing policy; humans approve, then export
- Readiness score reflects answer confidence, not workflow — approving a weak answer doesn't close the gap




Questionnaires → Supabase — 2026 vendor security assessment

→ The same engine answers the questionnaire your customer sent you.



Status is a control, not a label

- 'Assessing' holds onboarding; 'Approved' releases it — the register is the gate
- Criticality and risk rating set the depth of review a vendor gets
- Overdue reviews and expiring agreements surface on the row, before the renewal call
- 17 vendors on this tenant — every one with a status someone is accountable for




Vendors → Third-party register

→ No vendor touches regulated data before it clears the gate.



Every BAA, DPA, MSA and NDA — current, automatically

- One register for all agreement types, with the executed document attached
- 30 / 14 / 7-day expiry reminders to the owner — the expired BAA never surfaces in the deal
- Agreements link to the vendor and to the AI system they cover
- 19 agreements on this tenant; the daily digest lists what expires this month



Governance → Agreements

→ The spreadsheet of dates retires.



Inventory the agents and hold them to a written policy

- Every AI system — vendor-supplied or built in-house — classified by risk under ISO 42001, NIST AI RMF and the EU AI Act
- Per system: human oversight, evaluations, documentation, DPIA status, data it touches
- An autonomy policy in writing: what the agent may do unattended, what needs a human, and the kill switch
- 7 systems on this tenant; patient triage and claims coding flagged high-risk

 **KollGuard AI**
COMPLIANCE INTELLIGENCE

★ FAVORITES

🔍 Compliance Profile

⚠ Findings 75

📁 Evidence Library

📄 Reports

📅 Audit Log

👤 Agents 4

OVERVIEW

🏠 Dashboard

📖 Daily Digest

CONTINUOUS MONITORING

🔍 Scans

🔍 Code Scanner

📄 Sensitive Data

AI & AGENTS

🔍 AI Governance

📄 Help & docs

📖 Blog

⚙ Customize menu

Credentials encrypted in vault

KollGuard / AI Governance

🔍

🔄

🔔

👤 Demo Owner

🔗 AI Governance

Continuous posture for ISO 42001, NIST AI RMF, and the EU AI Act — auto-evidenced from the AI agents you run in Agents and your existing scans.

AI System Register

Every AI system you use, classified by EU AI Act risk tier with ISO 42001 governance evidence.

+ Add AI system

7 systems 3 high-risk 1 high-risk system with evidence gaps

SYSTEM	EU AI ACT	STAGE	DATA	EVIDENCE	
Patient Intake Triage Assistant OpenAI · gpt-4o Triage intake forms and summarize reported symptoms for ...	High-Risk	Production	PHI	✓ Oversight ✓ Eval ✓ Docs ✓ DPIA	✎ 🗑
Patient Support Chatbot OpenAI · gpt-4o-mini Answer non-clinical patient questions (billing, scheduling)	Limited	Production	PII	✓ Oversight · Eval ✓ Docs ✓ DPIA	✎ 🗑
Autonomous Code Review Agent Anthropic · claude-sonnet Review pull requests for security and correctness, open fix ...	Minimal	Production	INTERNAL	✓ Oversight ✓ Eval ✓ Docs · DPIA	✎ 🗑
Claims Coding Copilot Anthropic · claude-sonnet Suggest ICD-10 / CPT codes for billing from clinical notes	High-Risk	Production	PHI	✓ Oversight ✓ Eval ✓ Docs · DPIA	✎ 🗑
Marketing Content Generator OpenAI · gpt-4o Draft blog and marketing copy	Minimal	Production	NONE	✓ Oversight · Eval · Docs · DPIA	✎ 🗑
Meeting Notes Summarizer Anthropic · claude-haiku Summarize internal meetings	Limited	Pilot	INTERNAL	· Oversight · Eval · Docs · DPIA	✎ 🗑

AI Governance


→ 'Do you have an AI inventory?' is now the first question on every questionnaire.

kollguard.com · Northwind Health demo tenant

Evidence that collects itself

- 196 controls across 15+ frameworks — SOC 2, HIPAA, HITRUST, ISO 27001, NIST 800-53, PCI DSS, CIS, EU AI Act
- Controls monitored continuously; readiness per framework is a live number, not a quarterly estimate
- Evidence packages assembled from the record — the auditor gets a period, not a snapshot
- Policies, personnel, agreements and findings feed the same evidence library

→ Six weeks of screenshot hunting becomes an export.




Compliance Profile



Tamper-evident, and public where it should be

- A hash-chained, append-only audit trail — tampering is detected and pinpointed
- A public Trust Center for customers: frameworks, monitored controls, last assessed date
- The questionnaire your customer was going to send you gets shorter — or disappears
- SSO / SAML and SCIM so the whole organization works inside the same record

→ Audit-ready every day looks like this.



kollguard.com/trust/northwind-demo



The whole compliance program — one platform.

COMPLIANCE PROFILE — JURISDICTIONS & FRAMEWORKS

RISK REGISTER — SEEDS FROM REAL FINDINGS

WORKSPACE — EPICS • FEATURES • TICKETS

EVIDENCE LIBRARY — AUDIT-READY PROOF

196 controls · 15+ frameworks

- SOC 2 (TSC 2017)
- HIPAA Security & Privacy Rules
- HITRUST CSF
- ISO/IEC 27001:2022
- NIST SP 800-53 · NIST CSF 2.0
- PCI DSS 4.0
- CIS Controls v8
- GDPR · CCPA
- EU AI Act · ISO/IEC 42001 · NIST AI RMF
- IEC 62443 (physical / OT)



We hold ourselves to the standard we sell.


Server-enforced MFA •
passkeys

AES-256 • external KMS
custody

Per-tenant row-level
security

Hash-chained audit log

1,700+ PRs — every AI-written change goes through the **same governed review loop we sell**. A CISO can review our posture before trusting us with theirs.



WHERE WE HAVE SHOWN IT THIS YEAR

Pitched, exhibited and sponsored, **not just announced.**

JUL 18–19, 2026

AGI Summit 2026

Palace of Fine Arts, San Francisco, CA

JUL 18–19 · STARTUP TRACK ·
CUSTOMER EXCHANGE

AGI Summit 2026

Palace of Fine Arts, San Francisco, CA

- Customer-exchange sessions with founders shipping agents and the buyers vetting them.
- The AI-inventory and autonomy-policy questions we now put to every vendor on the KollGuard register.

JUL 30 – AUG 1, 2026

AI Summit at Stanford

Stanford University, CA

JUL 30 – AUG 1 · EXHIBITOR · MAIN-
STAGE PITCH

AI Summit at Stanford

Stanford University, CA

- Live repo scan on stage: a real codebase graded and mapped to SOC 2 and HIPAA controls in seconds.
- Why AI-written code ships faster than any compliance process was designed for, and how KollGuard closes that gap.

AUG 26–27, 2026

Step San Francisco 2026

AWS Builder Loft + The Midway, San Francisco, CA

AUG 26–27 · SPOTLIGHT STARTUP ·
PITCH COMPETITION

Step San Francisco 2026

AWS Builder Loft + The Midway, San Francisco, CA

- A 3-minute pitch to a founder-and-investor jury: continuous SOC 2 / HIPAA evidence for teams that ship with AI in the loop.
- Kiosk demos of the free security scan, security-questionnaire automation, and AI-agent governance.

SEP 17–18, 2026

ISACA Sacramento Flagship Conference 2026

Wyndham Sacramento, CA

SEP 17–18 · PLATINUM SPONSOR ·
SPEAKER SESSION

ISACA Sacramento Flagship Conference 2026

Wyndham Sacramento, CA

- Continuous compliance for the AI era: vendor risk, BAAs, and audit evidence without the annual scramble.
- Both sides of the vendor security questionnaire: enterprise teams issuing and assessing them, and vendors answering from live posture.

THREE DOORS, NONE OF THEM A SALES CALL

Pick the one that fits **this week.**

START FREE

Scan a repo or URL you own

A graded report in minutes, mapped to SOC 2 and HIPAA controls. Then paste a customer questionnaire and watch KGAI draft the answers from that posture.

kollguard.com/free/security-scan**SEE IT LIVE**

20 minutes on the Northwind tenant

Bring the questionnaire that is eating your week, or one vendor you are unsure about. We run it on a live tenant. No slides.

raj@kollguard.com · 628-333-5116**VERIFY US FIRST**

Read our own Trust Center

Frameworks, monitored controls, sub-processors, last assessed date. Server-enforced MFA, external KMS, per-tenant RLS, hash-chained audit log. BAA available.

kollguard.com/trust/kollguard