

# Enterprise AI Governance & *Data Security Briefing*

For IT, SecOps and legal review. How The Alliance deploys custom AI systems without operational or regulatory risk.

## AT A GLANCE

### Four commitments on every implementation

- **Zero Data Retention (ZDR).** Enterprise API protocols are configured for zero retention by model providers. Client data is never used to train foundation models.
- **100% client ownership.** You own 100% of the IP, code repositories and prompt pipelines. Everything is transferred at hand-off.
- **Human control.** Human-in-the-Loop (HITL) approval checkpoints, Role-Based Access Control (RBAC) and safety guardrails are built in from Day 1.
- **Framework alignment.** SOC 2 Type II and ISO 27001 alignment, HIPAA and GDPR support, with automated audit logging.

## Core pillars of security and architecture

Guardrails are built into your existing infrastructure from Day 1. Your data stays in your tenant, your team keeps full control, and you own every line of code.

### Zero Data Retention (ZDR): your data stays in your tenant

All custom agents run inside your company-owned tenant (Azure, AWS, M365, Google Workspace or OpenAI Enterprise). Enterprise API protocols are configured for zero retention by model providers, so your proprietary data is not used to train third-party models. Audit logs of agent activity stay in your environment.

### Role-Based Access Control (RBAC) and Active Directory

Agents inherit your existing directory privileges, SSO and permission rules. An agent can only read or process files that the employee using it is authorized to view, which reduces the risk of internal data leakage.

**Human-in-the-Loop (HITL) approval checkpoints**

High-risk workflows (financial movements, legal outreach, external communications) require mandatory human sign-off. Your team retains full kill-switch authority and audit oversight.

**100% IP ownership and zero vendor lock-in**

You own 100% of the IP. All custom code, workflows, prompt pipelines and UI documentation are transferred into your code repositories at hand-off.

**Compliance and framework compatibility**

| STANDARD / FRAMEWORK                  | HOW WE DELIVER IT          | WHAT THAT MEANS                                                                                                      |
|---------------------------------------|----------------------------|----------------------------------------------------------------------------------------------------------------------|
| SOC 2 Type II and ISO 27001 alignment | Tenant-isolated processing | Agents run inside your own tenant with tenant-isolated API processing, automated audit logging and token encryption. |
| HIPAA and GDPR support                | Zero-retention APIs        | Built on BAA-eligible infrastructure, with no persistent data logging on public models.                              |
| WCAG 2.1 AA                           | Accessible interfaces      | Custom agent dashboards and co-building interfaces are built to meet WCAG 2.1 AA.                                    |
| OAuth 2.0, SAML and SSO               | Directory integration      | Works with Okta, Microsoft Entra ID and Google Cloud Identity.                                                       |

**Governance levels and required guardrail stacks**

Match the execution tier to your use case to see the guardrail stack it requires.

LEVEL 1 · LOW RISK

**Internal Productivity**

Summaries, drafting, internal search.

**Required guardrail stack:** Standard RBAC, read-only access, and usage logging.

LEVEL 2 · MEDIUM RISK

**Workflow Automation**

CRM updates, report generation, routine customer comms.

**Required guardrail stack:** RBAC, human-in-the-loop review, audit trails, and rate limiting.

LEVEL 3 · HIGH RISK

**Autonomous Agent Execution**

Direct database writes, payment triggers, and regulatory filings.

**Required guardrail stack:** Full RBAC, mandatory human sign-off, kill-switch authority, real-time audit logging, and tenant isolation.

---

## Technical deliverables included with every implementation

- Data flow and architecture diagrams, mapped for IT and SecOps review.
- System prompt and guardrail audits that enforce brand and compliance boundaries.
- API log and audit trails with full transparency into token utilization and user actions.
- Incident response and kill-switch protocols.

---

## Questions from your legal and security teams

### Will our proprietary business data be used to train LLMs?

No. Every agent runs through zero-data-retention API connections inside your own tenant. Your proprietary data, prompts, and outputs are never logged, stored, or used to train third-party foundation models.

### How do custom agents handle sensitive PII (Personally Identifiable Information)?

PII is processed only within the access boundaries already defined by your directory permissions and RBAC policies. Sensitive fields can be masked, redacted, or excluded entirely from agent context depending on the workflow's risk tier.

### What happens if an external LLM API experiences downtime?

Workflows are built with fallback logic and human-in-the-loop checkpoints so they degrade gracefully instead of failing silently. High-risk autonomous workflows are configured to automatically pause and alert your team rather than take action while an API is unavailable.

### Can we run custom workflows completely on-premise or within a private cloud (AWS/Azure)?

Yes. Custom agents can be deployed entirely within your private AWS or Azure environment, including private endpoints and VPC-isolated model access, depending on your existing infrastructure and compliance requirements.

---

**Next step:** [book a 30-minute governance review](#). Tell us your stack and your first use case and we will map the guardrails it needs.

© 2026 AI Workforce Alliance, Inc. This briefing describes how The Alliance configures and delivers implementations. It is not a certification or legal advice.