AI Agent Governance Guide

vTECH io AI Agent Governance Guide: Policies and Best Practices for Secure Digital Employees

February 2026

This guide provides actionable frameworks to govern AI agents responsibly, mitigating risks like data leakage while maximizing value. Ideal for IT/security leaders in regulated industries.

Section 1: Core Principles

- Accountability: Assign human owners for every agent.
- Transparency: Document agent capabilities, decisions, and logs.
- Least Privilege: Grant only necessary access.
- Continuous Monitoring: Audit trails and anomaly detection.
- Compliance Alignment: Map to NIST, CJIS, HIPAA, etc.

Section 2: Governance Framework Template

- Policy Statement — "AI agents are treated as non-human identities with defined roles, requiring approval, monitoring, and revocation processes."
- Roles & Responsibilities — CISO owns security; IT manages deployment; Business units request use cases.

- Onboarding Checklist — Define purpose, risk assessment, RBAC setup, testing.
- Access Control Policy — Use RBAC/ABAC; no hardcoded secrets; zero-trust enforcement.

Section 3: Risk Assessment Template
- Identify: Autonomy level, data access, tools integrated.
- Analyze: Prompt injection, leakage potential.
- Mitigate: Guardrails, monitoring.
- Review: Quarterly or post-incident.

Section 4: Best Practices Quick List
- Unique identities per agent.
- Dynamic/context-aware permissions.
- Output filtering & redaction.
- Behavioral monitoring tools.
- Employee awareness training.
- Decommissioning protocol for unused agents.

Section 5: Resources & Next Steps
- Links to external reports.
- vTECH io Contact: Schedule a consultation for customized implementation.