

TAC and AI Agent Governance

How the TAC AI Agent Governance module — next on the platform — maps to the requirements of NIST AI RMF and ISO/IEC 42001:2023 — specifically for AI agents and non-human identities operating in your environment.

Note: This guide describes how TAC's technical capabilities map to NIST AI RMF and ISO/IEC 42001:2023 controls specifically for AI agents operating against your applications and data. Full conformance to either framework requires a complete AI governance program — including organizational policies, risk assessments, impact analyses, lifecycle management, and process controls — outside the scope of any access control platform. TAC addresses the technical enforcement, monitoring, and accountability layer; not the policy, lifecycle, or assessment layers.

Why AI agents change the compliance picture

AI agents are non-human identities that authenticate, access systems, query data, and take actions inside your enterprise — autonomously, at machine speed, and often with broader permissions than the humans they assist. Both NIST AI RMF and ISO/IEC 42001:2023 treat AI systems, including agents, as requiring identity, access, audit, and accountability controls equivalent to or stronger than those for human users.

Traditional IAM platforms were built for people. They don't extend cleanly to autonomous agents with static API keys, opaque scoping, and no audit trail tied back to a verified identity. The result is a growing population of privileged, ungoverned, non-human identities — the population both frameworks now require organizations to govern.

TAC governs human and non-human identities through one policy engine, producing a unified evidence trail that maps to both frameworks. Your AI governance evidence isn't a separate program — it's an extension of the access governance you already have.

Control Mapping

TAC contributes primarily to the **GOVERN** and **MANAGE** functions — the runtime enforcement, monitoring, and accountability layer for AI agents.

POLICY, ACCOUNTABILITY, AND OVERSIGHT

GOVERN — Partial Coverage

CONTROL	HOW TAC WILL DELIVER
GV-1.4 Policies enforcing accountability for AI systems	Access, authorization, and use policies for AI agents are enforced at runtime by the same policy engine that governs human users. Every agent action is attributable to a verified identity.
GV-1.5 Ongoing monitoring of AI systems	Continuous evaluation throughout every agent session — identity, policy, and request behaviour re-validated in real time. Anomalies surface immediately through SIEM integration.
GV-1.6 Mechanisms for AI inventory and decommissioning	Registry of governed agent identities consumed from your identity sources. One-action revocation immediately stops all further access for a decommissioned or compromised agent across every protected resource.
GV-4.3 Documentation of AI inputs, decisions, and outputs	Immutable audit log records every agent request: identity, source, target resource, parameters, the policy that applied, and the allow/deny decision. Searchable and exportable for review.
GV-6.1 Policies for third-party AI risks	Third-party AI agents accessing your environment are subject to the same identity verification, policy enforcement, and audit requirements as internal agents. No shadow integration paths.

RISK RESPONSE AND OPERATIONAL CONTROL

MANAGE — Primary Coverage

CONTROL	HOW TAC WILL DELIVER
MG-1.3 Risk treatments applied to AI systems	Policy-driven access control treats identified AI risks at runtime: least-privilege enforcement, resource-level restrictions, time-based and condition-based access rules.
MG-2.4 Mechanisms to supersede, disengage, or deactivate AI systems	Real-time session revocation when upstream identity is disabled, policy changes, or anomalous behaviour is detected. Active sessions terminate immediately; no orphaned access remains.

CONTROL	HOW TAC WILL DELIVER
MG-3.1 Third-party AI risks are managed	Same policy engine governs internal and third-party AI agents. Vendor-supplied agents are scoped, monitored, and revocable with the same controls applied to internal workloads.
MG-4.1 Post-deployment AI monitoring	Every request an agent makes is inspected at the proxy — method, URL, headers, parameters, and payload. Continuous evaluation produces a live operational view of agent behaviour.
MG-4.2 Mechanisms to track and respond to errors and incidents	Forensic-grade audit log supports incident investigation: full request-level attribution, immutable record, exportable to SIEM. Compromised agents can be revoked instantly while investigation continues.
MG-4.3 Continuous improvement of AI risk management	Audit log provides longitudinal evidence of agent behaviour to refine policies over time. Policy changes apply immediately to all active and future sessions without redeployment.

OUT OF SCOPE

MAP and MEASURE

The MAP function addresses pre-deployment AI system contextualization, categorization, and impact mapping. The MEASURE function addresses trustworthiness evaluation, fairness metrics, and effectiveness assessment of AI models themselves. These are organizational and methodological activities not addressed by access control platforms. TAC complements them by providing the enforcement and evidence layer once decisions about AI deployment have been made.

Annex A Control Mapping

TAC maps to specific Annex A controls covering AI system operation, monitoring, event logging, responsible use, and third-party AI relationships.

AI SYSTEM LIFE CYCLE — OPERATION AND MONITORING

A.6 — Primary Coverage

CONTROL	HOW TAC WILL DELIVER
A.6.2.6 Operation and monitoring of AI system	Continuous policy evaluation during agent operation. Every request from a deployed agent is authenticated, authorized, inspected, and logged before reaching protected resources.
A.6.2.8 Event logging	Immutable event log records every agent action with full attribution. Supports the operational, security, and audit evidence requirements of ISO 42001 conformance.

RESPONSIBLE USE OF AI SYSTEMS

A.9 — Partial Coverage

CONTROL	HOW TAC WILL DELIVER
A.9.2 Processes for responsible use	Per-agent access policies enforce intended-use boundaries. Resource-level permissions ensure agents can only act within their defined operational scope.
A.9.4 Intended use of the AI system	Out-of-scope agent actions are blocked at the proxy and recorded. Intended-use boundaries are enforced technically, not just documented.

INFORMATION FOR INTERESTED PARTIES

A.8 — Partial Coverage

CONTROL	HOW TAC WILL DELIVER
A.8.4 Communication of incidents	SIEM integration and exportable audit data support incident reporting to interested parties. Full request-level evidence available for regulator, customer, and partner notifications.

THIRD-PARTY AND CUSTOMER RELATIONSHIPS

A.10 — Partial Coverage

CONTROL	HOW TAC WILL DELIVER
A.10.2 Allocation of responsibilities	Identity attribution distinguishes internal agents from third-party agents in the audit log. Accountability for every action is preserved and traceable.
A.10.3 Suppliers	Third-party AI agents from suppliers are scoped, authenticated, monitored, and revocable through the same policy engine that governs internal workloads.

OUT OF SCOPE

A.2, A.3, A.4, A.5, A.7

Controls covering AI policies (A.2), internal organization (A.3), resources for AI systems (A.4), assessing AI impacts (A.5), and data for AI systems (A.7) address AI program governance, organizational structure, model development resources, impact assessment processes, and data handling for model training. These are program-level activities outside the scope of an access control platform. TAC provides the runtime enforcement layer that operates alongside them.

One policy engine. One audit log. Two frameworks.

Both NIST AI RMF and ISO/IEC 42001 expect organizations to demonstrate that AI systems are governed by enforceable policies, monitored in operation, and accountable for their actions. Neither framework dictates how — they expect the organization to provide evidence.

TAC is designed to produce that evidence in normal operation. The same policy engine that grants a clinician access to a patient record can govern an AI agent's access to the same record — with the same evaluation criteria, the same audit trail, and the same revocation mechanism. The evidence supporting your human access governance program and your AI agent governance program is produced by one system and exported in one format.

For organizations preparing for AI governance audits, this matters operationally. You aren't building a parallel evidence pipeline for AI agents. You're extending the one you already have.

PREPARING FOR AN AI GOVERNANCE AUDIT?

Talk to a PortSys specialist

Our team can walk through your AI agent inventory and show how TAC will deliver the technical enforcement and evidence requirements of NIST AI RMF and ISO/IEC 42001:2023.

www.portsys.com/contact