Identity-Aware AI Security

An Enterprise Workbook for Governing AI Through the Lens of Identity

Version 6.0.1 | August 2026 | startmakingsense.org

Version 6.0.1 Change Summary

Version 6.0.1 establishes a five-pillar enterprise architecture for identity-aware AI security, restores the detailed pillar and appendix material required to operate that architecture, and organizes the workbook around an iterative portfolio approach to authorization-risk reduction.

The version:

- Uses the five pillars of Authorization, Data, Services, SecOps, and Governance.
- Distinguishes the enterprise portfolio unit—an **AI application**—from the fixed **IAAI security interoperability use cases** used in the canonical catalog and public Interoperability Graph.
- Uses a portfolio approach that includes homegrown and vendor-introduced AI applications, including capabilities discovered but not explicitly enabled or authorized.
- Distinguishes standing entitlements, vaulted credentials, and real-time policy decisions as the three canonical preventative authorization-control primitives.
- Preserves a layered Data/Services model for data retrieval, abstraction, transformation, and consequential action.
- Adds a unified lifecycle for implementing IAAI: Discover, Assess, Implement, and Govern.
- Retains the three-pass adoption path, detailed seam-testing patterns, assessment instruments, representative technology/standards examples, end notes, index, and version history.

Table of Contents

- Executive Summary
- Workbook At a Glance
- 1. Why Identity-Aware AI Security
- 2. How to Use This Workbook
- 3. IAAI Architecture Overview
- 4. Getting Started: Your First 90 Days of Portfolio Authorization-Risk Reduction
- 5. Pillar A — Authorization: Identity-Aware Authorization Policy Management
- 6. Pillar B — Data: Identity-Aware Data Operations
- 7. Pillar C — Services: Identity-Aware Service Operations
- 8. Pillar D — SecOps: Post-AI Security Operations
- 9. Pillar E — Governance: Enterprise AI Governance
- 10. Cross-Pillar Interoperability and Control Seams
- 11. Implementing IAAI: Discovery, Assessment, Implementation, and Governance Lifecycle
- 12. Operating and Improving IAAI as an Enterprise System
- Appendix A — Operational Glossary and Cross-Reference Guide
- Appendix B — Selected Interoperability Patterns
- Appendix C — Representative Technologies, Standards, and Interoperability Examples
- Appendix D — Portfolio and AI Application Assessment Instruments, Evidence Templates, and Roadmap Worksheets
- End Notes and Sources
- Index
- Version History

List of Figures and Tables

- **Figure 1.** Identity-Aware AI Security Five-Pillar Architecture
- **Table 1.** The five pillars and their architectural responsibilities
- **Table 2.** Portfolio attention tiers
- **Table 3.** Control maturity by domain
- **Table 4.** Vendor-introduced AI-application lifecycle states
- **Table 5.** IAAI product capability category anchors

Executive Summary

AI is becoming a new enterprise access, interpretation, and action layer. It can retrieve, correlate, transform, and explain information across data estates;

invoke tools and APIs; alter business records; initiate workflows; communicate externally; and operate infrastructure. These capabilities create value, but they also concentrate authority in systems that may act faster, across more boundaries, and with less visible context than traditional human interfaces.

The central governance question is not only whether an AI model is accurate, safe, or compliant. It is also: **who or what is acting, on whose behalf, under what authority, against which data and services, with what constraints, evidence, and accountability?**

Identity-Aware AI Security (IAAI) is an enterprise architecture and operating workbook for answering that question. It organizes responsibilities into five connected pillars.

Pillar	Full responsibility	Primary role
A — Authorization	Identity-Aware Authorization Policy Management	Establishes identity context, lifecycle, delegation, entitlement and credential administration, authorization policy, policy-decision capability, and authorization evidence.
B — Data	Identity-Aware Data Operations	Governs data access, retrieval, filtering, classification, metadata, abstraction, transformation, lineage, derived artifacts, and data policy enforcement points.
C — Services	Identity-Aware Service Operations	Governs tools, APIs, MCP operations, workflows, applications, communications, transactions, administrative actions, infrastructure control planes, and service policy enforcement points.
D — SecOps	Post-AI Security Operations	Provides compensating controls, guardrails, DLP, telemetry, posture analysis, detection, investigation, response, resilience, and security feedback.
E — Governance	Enterprise AI Governance	Sets purpose, ownership, risk appetite, portfolio attention, decision rights, lifecycle conditions, exception treatment, assurance, evidence expectations, and accountability.

The pillars are connected responsibilities—not sequential project phases, mutually exclusive product categories, required organizational silos, or a demand to

purchase five separate technology stacks. A technology, standard, team, process, or product can contribute to several pillars. Conversely, no single product should be assumed to implement the entire architecture.

The workbook manages a portfolio of **AI applications**. An AI application is a bounded enterprise deployment or use of AI for a business purpose, including its people, agents, models, runtime, connected data, tools/services, workflows, external effects, ownership, lifecycle, and governance context. This differs from an **IAAI security interoperability use case**, which is a fixed canonical assertion about a recurring cross-pillar security interaction and is used in the public Interoperability Graph and Appendix A.

The operating objective is iterative authorization-risk reduction. The enterprise discovers AI applications, preserves awareness even where information is incomplete, assesses where authority and consequence concentrate, implements preventative and compensating/governance controls proportionately, validates evidence, makes residual risk accountable, learns from operation, and reprioritizes the portfolio.

Five ideas must remain distinct:

- **Portfolio attention tier** determines how much review, evidence, testing, governance, and specialist analysis an AI application needs now.
- **Control maturity** describes the current and target strength of a particular control domain.
- **Finding severity or urgency** describes the impact and time sensitivity of a specific gap, incident, or control failure.
- **Residual risk** is the risk remaining after applicable controls are considered.
- **Roadmap priority** determines where the next investment can reduce the most material portfolio risk.

The workbook also distinguishes preventative from compensating/governance controls. Preventative controls constrain excessive, unsafe, or out-of-scope authority before a consequential data or service operation occurs. Compensating/governance controls reduce residual risk, constrain deployment, inspect behavior, detect and respond to misuse, make remaining risk accountable, and improve the architecture over time.

The companion **Enterprise AI Authorization Risk Calculus** remains a separate work. This workbook provides the portfolio, architecture, evidence, and governance method that can trigger and consume the Companion Authorization-Risk Exercise for selected high-consequence authority paths.

Workbook At a Glance

This workbook is designed for different readers and different points of entry. Executives and governance participants can use the summary, governance, implementation, and assurance material to establish direction and accountability. Architects and control owners can use the architecture, pillar chapters, seam patterns, technology reference, and worksheets to design, assess, test, and improve controls. Readers do not need to proceed linearly, but Sections 1–4 establish the shared concepts needed to use the detailed material consistently.

If you need to...	**Start here**	**Then use**
Understand why IAAI is needed	Executive Summary; Section 1	Sections 2–3
Establish a practical enterprise starting point	Section 4	Section 11; Appendix D
Inventory AI applications, including vendor-introduced capabilities	Sections 2 and 4	Worksheets D.2–D.3
Understand the architecture	Section 3	Sections 5–10
Map identity, delegation, policy, and lifecycle	Section 5	Pattern B.2.3; Worksheets D.5.1–D.5.4
Map data operations and Data PEPs	Section 6	Pattern B.2.1; Worksheets D.5.2 and D.5.4
Map services, tools, APIs, MCP, and Service PEPs	Section 7	Pattern B.2.2; Worksheets D.5.3 and D.5.4
Plan guardrails, DLP, telemetry, detection, and response	Section 8	Pattern B.2.5–B.2.7; Worksheet D.5.5
Establish governance, exceptions, assurance, and portfolio review	Section 9	Patterns B.2.8–B.2.9; Worksheets D.5.6–D.7
Test cross-pillar interactions	Section 10	Appendix B; Worksheet D.5.4
Implement the full lifecycle	Section 11	Appendix D
Build long-term shared-control leverage and learning	Section 12	Sections 8–11
Select standards or technology capabilities	Appendix C	End Notes and Sources

1. Why Identity-Aware AI Security

1.1 The enterprise AI authorization problem

AI is becoming a new access, interpretation, and action layer over the enterprise. Copilots, chatbots, retrieval-augmented generation systems, agents, agent frameworks, models, and AI-enabled workflows can reach data across mailboxes, repositories, data platforms, SaaS applications, APIs, observability systems, and business processes. They can synthesize information across organizational and technical silos. They can also act: invoke tools, trigger workflows, create or modify records, communicate externally, change configurations, and initiate transactions.

The intended value of AI is inseparable from a new form of risk. A mis-scoped identity, credential, entitlement, authorization policy, tool permission, data path, or service path can permit not only a local access event but a chain of retrievals, transformations, recommendations, messages, transactions, infrastructure changes, or other consequential outcomes. AI compresses the time between access and action, reuses grants across repeated tasks, and can combine otherwise ordinary permissions into effective authority that was not fully anticipated when each permission was granted.

1.2 Why identity-aware authorization is the primary control surface

Identity-aware authorization asks a preventative question before a governed operation proceeds: which authenticated identity is acting; in what delegated, workload, service, or agent context; against which data or service; for which operation; under what constraints; and with what accountable owner?

This is a more precise primary control surface than relying only on after-the-fact pattern recognition. It can constrain authority before sensitive data is retrieved, a tool is invoked, a record is changed, a message is sent, or an infrastructure action is performed. It does not make security inspection, DLP, guardrails, detection, incident response, privacy analysis, legal judgment, or governance unnecessary. It gives those controls a stronger identity, policy, and evidence substrate.

Without reliable, governable authorization, an enterprise can create AI access paths that are less controlled than the pre-AI systems the AI can reach.

1.3 AI applications, interoperability use cases, and authority paths

An **AI application** is the primary portfolio unit. It is a bounded application or deployment of AI for a business purpose, including users, agents, models, runtime, connected data, tools, services, workflows, external effects, ownership, lifecycle, and governance context. An AI application can be homegrown,

vendor-introduced, embedded in an enterprise platform, or composed from several shared services and agents.

An **IAAI security interoperability use case** is different. It is a fixed canonical assertion about a recurring security interaction among one or more pillars. For example, a canonical use case can describe an access-request workflow that updates a standing entitlement for a data operation with Authorization, Data, and Governance participation. These fixed assertions support the public IAAI Interoperability Graph and Appendix A; they are not the items the enterprise inventories as portfolio applications.

An **authority path** is a selected identity-to-data/service-action chain within an AI application. It is the appropriate unit for deep evaluation when effective authority, delegation, reach, consequence, or reversibility warrants more detailed analysis.

1.4 Scope and non-goals

This workbook is an architecture, governance, and assessment guide. It is not a vendor ranking, product-selection guide, legal opinion, complete protocol specification, single-product reference architecture, or substitute for business, safety, security, privacy, operational, or governance judgment.

The workbook assumes heterogeneous enterprise portfolios. Its purpose is to make responsibility boundaries, control interactions, evidence, and improvement paths clear—not to imply that every product belongs exclusively to one pillar or that every enterprise must centralize all decisions in one technology.

---

2. How to Use This Workbook

2.1 Purpose and reader paths

This workbook helps enterprises make Identity-Aware AI Security actionable. It provides architecture, design prompts, interoperability patterns, assessment methods, technology/standards references, evidence templates, and implementation guidance for AI applications that access enterprise data, invoke services and tools, communicate externally, alter business records, or otherwise create meaningful operational effects.

It supports two complementary reader paths.

Executive and governance path. Executives, board members, business sponsors, risk leaders, governance participants, legal/compliance leaders, and transformation leaders can begin with the Executive Summary, Sections 1–4, Section 9, Section 11, Section 12, and the portfolio instruments in Appendix D. This

path emphasizes purpose, accountability, attention, evidence, risk treatment, implementation cadence, and shared-control investment.

Architect and control-owner path. Enterprise architects, IAM/IGA/PAM leaders, data/platform teams, API/workflow owners, AI-platform teams, SecOps leaders, data-security leaders, and governance implementation owners should use Sections 1–3, then Sections 5–10, Appendices A–C, and the detailed Appendix D instruments. This path emphasizes identity context, policy decision and enforcement points, data/service boundaries, telemetry, patterns, tests, evidence, and technology options.

2.2 The portfolio principle

The enterprise should manage authorization risk as a portfolio of AI applications, not as isolated one-time approval exercises. The objective is not to apply every prompt or advanced control to every AI application on day one. The objective is to build comprehensive awareness, identify where authority and consequence concentrate, apply progressively stronger preventative and compensating/governance controls, validate evidence, learn from findings, and reprioritize investments.

Every discovered AI application remains in the portfolio while it exists. Unknown information is recorded as **Unknown** rather than becoming a reason to exclude the application from inventory. This applies especially to vendor-introduced AI capabilities that may arrive through SaaS, collaboration, CRM, ERP, data, security, or other enterprise platforms.

2.3 Vendor-introduced AI applications

A vendor-provided capability does not ride for free into the enterprise architecture. If it can access enterprise data, act through enterprise services, influence decisions, communicate externally, or expand agentically across a vendor platform, it is an AI application for portfolio purposes.

Lifecycle state	Meaning
Discovered — not explicitly enabled/authorized	Known capability available through an existing vendor relationship, service default, evolving feature, or tenant configuration, without a confirmed enterprise authorization or activation decision.
Available — pending evaluation	Potentially activatable capability awaiting business ownership, value/risk, technical-reach, vendor-evidence, and control evaluation.
Pilot or limited evaluation	Bounded testing under defined conditions.
Enabled or operational	Activated for internal or external use.

Lifecycle state	Meaning
Restricted, suspended, or disabled	Limited, paused, or withdrawn because of governance, control, risk, lifecycle, or business reasons.
Retired	No longer in use; associated identities, access, connections, evidence, and records require retirement treatment.

Each vendor-introduced AI application needs an enterprise business owner for the value sought and risk appetite tolerated. Assessment considers vendor architecture/security documentation, available configuration, data/service reach, identity model, contractual terms, change notices, vendor evidence, and any required vendor engagement.

2.4 Portfolio, AI application, and authority path

Level	Unit of work	Core question	Primary purpose
Portfolio	All known AI applications, shared dependencies, and common controls	Where should the enterprise invest next to reduce the most material authorization risk?	Prioritization, common-control planning, capacity allocation, governance, and trend measurement
AI application	One bounded business-purpose deployment or use of AI	What identity, data, service, SecOps, and governance controls are needed here?	Proportionate design, assessment, evidence, remediation, and lifecycle management
Authority path	One material identity-to-data/service-action chain within an AI application	What effective authority and consequence can arise along this path?	Selected deep analysis of toxic combinations, escalation, bypass, and high-consequence risk

2.5 Portfolio attention tiers

Portfolio attention tier determines the depth of review, evidence, testing, governance, and specialist analysis an AI application requires now. It is not a maturity score, residual-risk rating, or permanent label.

Tier	Name	Typical characteristics	Typical treatment
Tier 1	Managed baseline	Currently limited/reasonably bounded reach or incomplete evidence requiring monitored baseline attention	Portfolio visibility, ownership, lifecycle state, preliminary scope, proportionate baseline/routing, and reassessment triggers
Tier 2	Enhanced review	Broader internal data, changing integrations, meaningful workflow effect, or unresolved control/evidence questions	Fast diagnostic, targeted pillar/seam assessment, added evidence, and defined remediation/routing
Tier 3	Material authority	Material data, business action, autonomy, cross-boundary reach, external effect, or consequential decision influence	Full relevant pillar assessment, seam testing, governance review, roadmap, and documented authority bounds
Tier 4	High-consequence authority	Financial, legal, safety, regulated, privileged, irreversible, external, or infrastructure consequence	Deep evidence/testing, explicit authority bounds, senior risk treatment, and documented companion-exercise routing decision

Tiering guides attention. It does not decide whether an AI application is included in the portfolio. A newly discovered item may remain provisionally unclassified or temporarily Tier 1 while the enterprise obtains enough information to reassess it.

2.6 Control maturity and preventative-control depth

Control maturity is assessed by domain, not as one blended score for an AI application.

Maturity	Description
Basic	Documented ownership and minimum safeguards; standing entitlements, manual controls, or fragmented evidence may predominate.

Maturity	Description
Intermediate	Managed and repeatable controls; controlled credential patterns, systematic review, narrower scope, defined workflows, and stronger evidence may be present.
Advanced	Context-aware, real-time or near-real-time authorization and enforcement where appropriate; tested integration, high-quality telemetry, automated lifecycle controls, and verified feedback loops.

The canonical preventative authorization-control primitives provide a useful baseline view of depth:

Primitive	Baseline depth	Qualification
Standing entitlement	Basic	May gain depth credit through enhancements such as approval workflows, time bounds, review/certification, narrower scope, segregation of duties, evidence, and compensating controls.
Vaulted credential	Intermediate	Protects a retained privileged credential or secret; it does not become an advanced real-time policy decision merely because vault/PAM controls are strong.
Real-time policy decision	Advanced	Evaluates a requested operation in current context and can be enforced at a relevant PEP.

This primitive/depth relationship is a major input to authorization-risk analysis for a particular element of an AI application. It is not a universal replacement sequence, a complete risk rating, or a whole-application maturity score. Enterprises commonly begin with extensive standing entitlements and some privileged shared credentials in vaults; they selectively deepen preventative authorization based on risk while using compensating controls and governance to reduce risk during the transition.

2.7 Preventative and compensating/governance controls

Preventative controls constrain excessive, unsafe, or out-of-scope authority before a consequential data or service operation occurs. Compensating/governance controls reduce residual risk, constrain deployment, inspect behavior, detect and respond to misuse, make remaining risk accountable, and improve the architecture over time.

2.8 Three adoption passes

Pass	Objective	Typical output
Pass 1 — Portfolio visibility	Know what exists, including vendor-introduced capabilities, and where authority/consequence may concentrate	AI-application inventory, owner, lifecycle state, preliminary attention, dependencies, unknowns
Pass 2 — Baseline and routing	Confirm minimum information and controls; decide which applications need deeper work	Fast diagnostic, current/target control view, routing decision, gaps, reassessment trigger
Pass 3 — Material AI-application deep dive	Design, test, evidence, and remediate material identity/data/service paths	Authority map, pillar/PEP/seam assessment, findings, roadmap, companion-exercise routing where warranted

These passes help an enterprise enter the lifecycle. They do not replace the comprehensive Discover, Assess, Implement, and Govern method in Section 11.

2.9 Four question families

For each material AI application, ask:

1. **Authorization:** Who or what is acting, on whose behalf, with what identity, lifecycle, delegation, entitlement, credential, and policy context?
2. **Data:** What data operation is permitted, at what scope, classification, freshness, representation, transformation, lineage, and derived-artifact boundary?
3. **Services:** What tool, API, MCP operation, workflow, message, transaction, configuration change, or other consequential action is permitted under what constraints?
4. **SecOps and Governance:** What compensating controls, evidence, decision rights, exceptions, lifecycle conditions, and accountability apply?

---

3. IAAI Architecture Overview

3.1 Architecture at a glance

The IAAI architecture differentiates the Authorization capability that establishes identity-aware policy leverage from the Data and Services contexts that

enforce policy (preventative controls); the SecOps capability that provides compensating controls; and the Governance function that determines acceptable risk and directs evolution.

Figure 1. Identity-Aware AI Security Five-Pillar Architecture

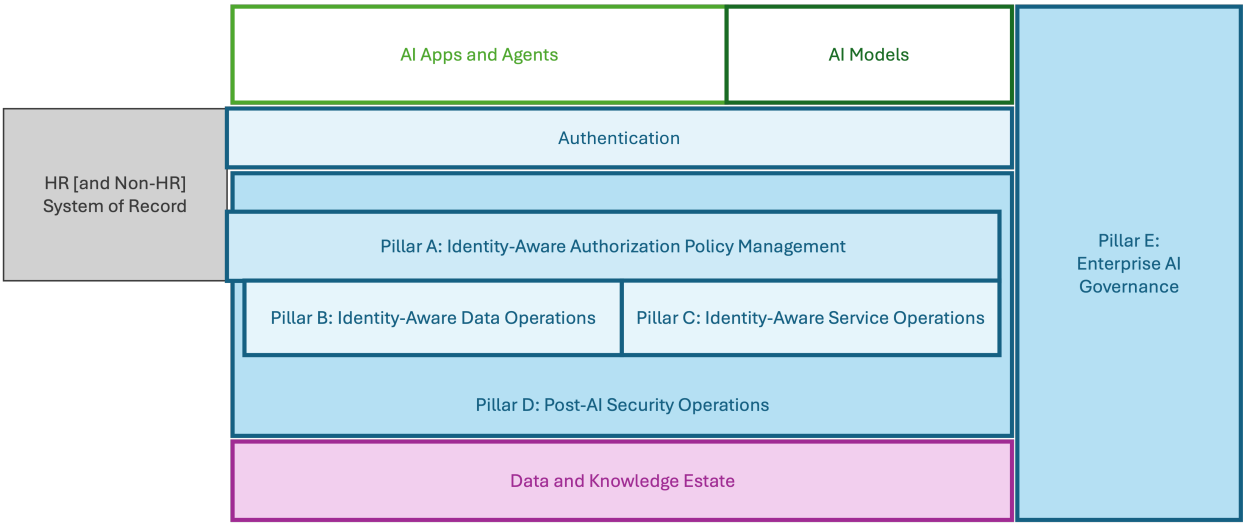

Figure 1: Identity-Aware AI Security Five-Pillar Architecture

MCP, APIs, service mesh, and workflow interfaces are interoperability contexts, not additional pillars. Pillar E is cross-cutting rather than a layer below the other pillars. Pillar D spans the operating plane and broader enterprise security estate; it is not merely an output-inspection filter.

3.2 The five pillars

Pillar	Architectural responsibility	Boundary and interaction
Authorization	Establishes identity context; administers identity lifecycle, standing entitlements, vaulted credentials, delegation, policy, PDP capability, and evidence.	Supplies governed identity/policy context to Data and Services boundaries; provides evidence to SecOps and Governance.

Pillar	Architectural responsibility	Boundary and interaction
Data	Governs identity-aware data operations: retrieval, query, filtering, platform-configured abstraction, transformation scope, lineage, metadata, and derived artifacts.	Hosts Data PEPs; consults Authorization where required; supplies data-operation evidence.
Services	Governs identity-aware service operations: tools, APIs, MCP operations, workflows, communications, transactions, administration, and infrastructure actions.	Hosts Service PEPs; may invoke Data operations; supplies service-operation evidence.
SecOps	Applies compensating controls, inspection, posture, detection, investigation, response, and security feedback.	Consumes evidence from A/B/C; provides findings and corrective feedback to A/B/C/E.
Governance	Governs purpose, ownership, portfolio, risk appetite, decision rights, exceptions, lifecycle, assurance, and accountability.	Directs requirements across A/B/C/D and receives evidence/findings for risk treatment and prioritization.

3.3 PDPs, PEPs, and a coordinated portfolio

A **policy decision point (PDP)** evaluates policy using available identity, subject/actor, workload, purpose, resource, operation, data classification, approval, risk, and contextual inputs. A **policy enforcement point (PEP)** applies a decision at an operating boundary.

Pillar A provides enterprise policy administration, lifecycle, and PDP leverage. Pillar B hosts Data PEPs. Pillar C hosts Service PEPs. Native/local PDP functions may remain necessary for local resource semantics, performance, and platform-specific controls. The enterprise should use a coordinated enterprise/native PDP portfolio: document policy boundaries, align policy semantics, make synchronization and exceptions explicit, and retain auditability.

3.4 Common context chain

```
Subject → actor/agent → delegation/workload context → purpose → data/service resource
→ operation → constraints → decision → enforcement → evidence → outcome → feedback
```

Loss of relevant context at any point can create overbroad access, untraceable action, weak investigation, or ungoverned exceptions.

3.5 Core architectural concepts

Concept	Definition
Human identity	A specific authenticated person with governed entitlements, lifecycle, and accountable manager or owner.
Personal agent identity	A non-human identity acting for a particular person; it should have authority derived from and not exceeding the sponsoring person's approved scope.
Enterprise agent identity	A non-human identity acting for an organizational purpose or process; it has named owner, defined purpose, lifecycle, and constrained function-specific authority.
Workload identity	Cryptographically verifiable runtime identity for a workload, service, or agent instance.
Subject	The person, organization, or principal on whose behalf action is taken.
Actor	The agent, workload, service, or person performing the action.
Standing entitlement	Persistent authorization object such as a role, group, ACL, grant, cloud permission, or relationship.
Vaulted credential	Retained privileged credential or secret protected, rotated, brokered, and released under controlled conditions.
Real-time access token	Time-bounded runtime artifact carrying identity/delegation context; not a canonical preventative-control primitive.
Real-time policy decision	Current-context determination that a requested operation is permitted or denied.
Data PEP	Enforcement point at a data-operation boundary.
Service PEP	Enforcement point at a service-operation boundary.
Governance decision	Human deliberation or subjective determination about risk, approval, constraint, exception, or direction.
Governance event	Processed record or operating action implementing or documenting governance-driven change.

4. Getting Started: Your First 90 Days of Portfolio Authorization-Risk Reduction

4.1 Start with awareness, not perfection

Start small, make the enterprise AI-application portfolio visible, apply a proportionate baseline, choose high-value deep dives, strengthen shared controls, and repeat. Do not wait for complete information, universal advanced authorization, or perfect metrics before reducing real authorization risk.

The first rule is awareness: enter every discovered AI application in the portfolio, including applications that are vendor-introduced, embedded in an existing solution, or discovered without an explicit activation decision. Record unknowns. A known but poorly understood AI application is more governable than an unseen one.

4.2 The three passes

Pass	What to do	Who to involve	What you produce
Portfolio visibility	Discover AI applications, agents, data connections, tools/services, lifecycle, ownership, vendor-introduced capabilities, and initial authority/ consequence indicators.	Executive sponsor, business owners, AI/platform lead, governance lead, representatives from IAM, Data, Services, and SecOps.	Portfolio inventory, owner, lifecycle state, preliminary attention, unknowns, dependencies.
Baseline and routing	Run a fast diagnostic; confirm minimum ownership, scope, identity, data/service reach, evidence, and governance conditions; select deeper work.	AI-application owner plus relevant control owners.	Baseline status, current/target view, routing decision, gaps, reassessment triggers.

Pass	What to do	Who to involve	What you produce
Material deep dive	Map authority, data, services, PEPs, seams, evidence, findings, remediation, exception, and governance needs.	Cross-functional assessment team.	Detailed profile, evidence, findings, roadmap, and companion-exercise routing where warranted.

4.3 Days 1–15 — Establish the cadence

1. Name an executive sponsor and small cross-functional working group.
2. Establish scope: production AI applications, important pilots, vendor-introduced capabilities that are available or discovered, and embedded AI with enterprise data or service reach.
3. Adopt common terms: AI application, IAAI security interoperability use case, authority path, attention tier, control maturity, preventative/compensating controls, identity types, PDPs/PEPs, standing entitlement, vaulted credential, and real-time policy decision.
4. Establish an operational review cadence for material applications and open findings, plus a periodic portfolio review cadence for priority and shared-control decisions.
5. Create one portfolio register and one integrated remediation roadmap.

Continue with: Section 11.2 for Discover and Section 11.5 for governance cadence; Worksheet D.2.1.

4.4 Days 16–30 — Build a usable portfolio

For each known AI application, record:

- AI-application ID, business purpose, and value sought.
- Sponsor, business owner, technical owner, vendor/supplier where applicable.
- Lifecycle state, including discovered—not explicitly enabled/authorized where applicable.
- Agents, models, runtime, platforms, and shared dependencies.
- Main data sources, highest known classification, tools, APIs, MCP contexts, workflows, and external effects.
- Identity/authority pattern and preliminary attention tier.
- Known concerns, unknowns, vendor evidence, and shared-control dependencies.

Do not postpone inventory because a field is unknown. Unknown is a valid initial value and a future assessment task.

Continue with: Section 11.2; Worksheets D.2.1–D.2.3.

4.5 Days 31–45 — Apply baseline and route work

Run the Fast Diagnostic for active production applications, high-priority pilots, and materially available vendor-introduced applications. Confirm ownership, lifecycle, data/service reach, identity pattern, baseline evidence, exceptions, and reassessment triggers. Select a small number of Tier 3/4 deep dives, recurring shared-control gaps, and a representative Tier 1/2 sample.

Continue with: Section 11.3; Worksheets D.3.1–D.3.3.

4.6 Days 46–75 — Deepen material AI applications

For selected applications, complete the identity/authority map, Data and Services operation maps, relevant pillar instruments, PEP/seam coverage map, SecOps/evidence profile, Governance profile, and roadmap. Define preventative requirements, compensating/governance controls, evidence, temporary-exception conditions, and closure criteria. Route high-consequence authority paths to the Companion Authorization-Risk Exercise when appropriate.

Continue with: Section 11.4; Worksheets D.4–D.6.

4.7 Days 76–90 — Create leverage and repeat

Identify repeated patterns and shared dependencies. Improve one or two common controls such as agent identity lifecycle, policy administration, metadata quality, retrieval filtering, API/service enforcement, telemetry correlation, or exception workflow. Update reference architectures, tests, and templates. Review open exceptions and reprioritize the next cycle.

Continue with: Sections 11.5–11.6 and Section 12; Worksheets D.7–D.8.

4.8 What you should have at day 90

By day 90, an enterprise should have a usable portfolio register; named owners; visible vendor-introduced capabilities; preliminary attention and routing; baseline status for priority applications; selected material deep dives; an integrated roadmap; defined review cadence; initial shared-control priorities; and a repeatable method for the next cycle.

---

5. Pillar A — Authorization

Identity-Aware Authorization Policy Management

The enterprise capability that establishes, administers, and makes available the identity-aware authorization context and policy decisions needed to govern AI activity.

5.1 The problem

Every identity-aware control in Data and Services depends on one fundamental condition: the enterprise can consistently identify who or what is acting and determine what that identity may do, under what conditions, for what purpose, and with what evidence. Without this capability, Data and Services teams must define and manage policy locally, creating inconsistent semantics, fragmented lifecycle practices, unreviewable privilege, weak delegation evidence, and unscalable governance.

Pillar A does not perform all enforcement itself. Its role is to provide enterprise-wide identity-aware authorization-management capability: identity, lifecycle, entitlement, credential, delegation, policy administration, PDP leverage, and evidence that Data PEPs, Service PEPs, SecOps, and Governance can use. It is a shared capability, not necessarily a monolithic system. Most enterprises compose it from multiple technologies covering different environments, identity types, resources, and policy domains.

The architectural test is not whether every decision originates from one product. It is whether distributed authorization decisions are understood, governed, auditable, aligned with the enterprise policy model, and enforceable at relevant Data and Services boundaries.

5.2 Core principle

Pillar A provides shared, authoritative, callable identity-aware authorization-management capability. A system making a material identity-aware decision—whether for data retrieval, data-platform abstraction, tool/API use, MCP operation, workflow action, delegated agent activity, workload/service communication, or conditional access—should use an explicit Pillar A policy/identity context rather than silently maintain unmanaged local policy.

The enterprise should maximize common policy leverage while preserving native/local PDP capabilities where resource semantics, performance, or platform-specific control require them. Local decisions are legitimate when their policy boundaries, alignment, synchronization, exception handling, evidence, and governance are explicit.

5.3 What Pillar A manages

Managed capability	What it includes	Why it matters for AI applications
Human identity lifecycle	Joiner/mover/leaver events, authentication, provisioning, deprovisioning, access reviews, and organizational context.	A human's role, employment, and responsibilities affect the authority that can be delegated to personal agents and accessed through AI applications.
Personal-agent identity lifecycle	Registration, sponsor, purpose, derived scope, certification, suspension, and retirement.	A personal agent must not silently accumulate authority beyond its sponsor's approved scope.
Enterprise-agent identity lifecycle	Purpose, owner, function, entitlement, credential, review, suspension, and retirement.	Enterprise agents act for organizational purposes, not merely as proxies for an individual.
Workload, service, and application identity	Service authentication, workload attestation, rotation, audience, and trust context.	AI agents and services require verifiable runtime identities at service boundaries.
Standing entitlement administration	Roles, groups, relationships, ACLs, cloud permissions, data grants, and other persistent authorization objects.	Most enterprise access still relies on standing entitlements; their scope, ownership, review, and drift are central risk inputs.
Vaulted credential administration	Retained privileged credentials/secrets, ownership, rotation, checkout/check-in, session control, release, and audit.	Privileged service or administrative access requires controlled use of retained credentials.
Policy definition and versioning	Policy-as-code, policy review, testing, versioning, distribution, exception handling, and rollback.	AI data/service policy must be testable and change-controlled rather than embedded invisibly in application code.
PDP capability	Enterprise-wide and native/local policy-decision services.	Decisions need policy, identity, resource, operation, and context at runtime or relevant evaluation time.
Access requests and certification	Requests, approvals, time bounds, reviews, recertification, revocation, and evidence.	An enhanced standing-entitlement workflow may deepen control beyond a basic unmanaged grant.

Managed capability	What it includes	Why it matters for AI applications
AI platform, model, tool, and MCP access	Who may use an AI platform, model, agent, tool, data connector, or MCP capability, under what conditions.	Tool/platform access is itself an authorization concern before downstream data/service operations occur.
Authorization posture and evidence	Entitlement, credential, policy, decision, lifecycle, exception, and coverage evidence.	Evidence supports SecOps analysis and Governance treatment of effective authority, drift, and risk.

5.4 Identity types, ownership, and lifecycle

Identity-aware AI security begins with identity types that are explicit enough to support lifecycle, ownership, policy, evidence, and accountability.

Identity type	Definition	Ownership and lifecycle expectation
Human identity	Authenticated employee, contractor, partner, or other person with enterprise relationship and entitlements.	Managed through established identity lifecycle and certification processes.
Personal agent identity	AI agent acting on behalf of a particular human for that person's work.	Sponsoring human and accountable manager/owner; authority derived from and constrained by the human's approved scope; reviewed and revoked with the relationship or purpose.
Enterprise agent identity	AI agent acting for an enterprise function, team, workflow, or process.	Named enterprise owner and purpose; function-scoped authority; periodic review; suspension/deprovisioning when purpose, control, or governance conditions change.
Workload identity	Verifiable runtime identity for a workload, agent instance, or service.	Managed by technical owner with rotation, attestation, trust-domain, and deployment lifecycle controls.

Identity type	Definition	Ownership and lifecycle expectation
Service/ application identity	Identity associated with an application/service boundary.	Named owner; scoped authority; no unowned shared credential or broad persistent privilege.
Delegated identity context	Subject/actor relationship, audience, scope, purpose, and chain of action.	Preserved through relevant calls; evidence must permit reconstruction of who requested and who acted.

An AI application has an enterprise-governed lifecycle under Pillar E. The agents, workloads, services, credentials, and entitlements within it have component lifecycles managed by their technical owners, especially within Pillar A. Material component changes—new agent capability, altered delegation, broader credential, new tool access, changed workload identity, or changed policy—can trigger AI-application reassessment.

Pillar E may require restriction, suspension, or termination of an agent lifecycle when active governance policy, risk conditions, ownership, evidence, or purpose requirements are not met.

5.5 Identity governance, administration, and governance

Identity governance and administration (IGA) is a principal Pillar A capability. It manages identity lifecycle, roles, standing entitlements, access requests, provisioning, certification, deprovisioning, and related evidence. Routine technical administration belongs in Pillar A.

Enterprise AI Governance is different. Pillar E decides whether an AI application, authority pattern, exception, effective-access posture, or residual risk is acceptable and what direction or constraint is required. Pillar E can require an access request to receive elevated review, restrict an agent purpose, impose a condition, or terminate use; Pillar A translates those decisions into technical administration, policy, lifecycle, and evidence activities.

A governance decision is human deliberation. A governance event is the recorded approval, restriction, exception, review completion, or technical-change request that implements or documents the decision.

5.6 Authorization-control primitives

Primitive	Purpose	Baseline depth	Important limitations
Standing entitlement	Persistent permission through role, group, ACL, relationship, grant, or cloud permission.	Basic	Can become stale, inherited, broad, unreviewed, or reused beyond original intent. Enhanced request/review/time-bound/constraint/evidence patterns can deepen assurance.
Vaulted credential	Controlled protection, rotation, brokering, and release of retained privileged credential/secret.	Intermediate	Protects credential use but is not a current-context policy decision and does not become advanced merely through vault controls.
Real-time policy decision	Current-context evaluation of a requested operation.	Advanced	Has preventative value only when a relevant PEP enforces it and required context/policy availability are reliable.

A real-time access token can carry subject, actor, audience, scope, and other runtime context. It is useful for authentication, federation, delegation, and authorization context, but is not a fourth canonical preventative primitive.

5.7 Authorization enforcement contexts

Context	Where policy is evaluated/enforced	Questions that matter
Data retrieval and query	Repository, data platform, data API, search, vector, graph, semantic layer, or transformation boundary.	Who may retrieve which source/representation under which metadata, purpose, and scope constraints?
Data-platform abstraction	Data platform automatically masks, filters, suppresses, or returns a policy-configured representation in response to an authorized read.	What representation may this identity retrieve under existing policy?

Context	Where policy is evaluated/enforced	Questions that matter
Service/tool/ API invocation	API gateway, service gateway, MCP authorization component, workflow, application, service mesh, or agent runtime.	Which method, parameter, object, recipient, threshold, state, rate, and external effect are permitted?
Delegation and agent-to-agent activity	Orchestration, agent runtime, token exchange, workload identity, and service boundary.	Who is subject, who is actor, what scope/audience/purpose may be passed, and what chain is evidenced?
AI platform/ model/tool access	IdP, SSO, conditional access, enterprise browser, agent platform, or tool registry.	Which identities may access which AI capability, under what device/session/ location/purpose conditions?
MCP operations	MCP client/server, authorization server, gateway, Data/Service PEP, and event stream.	Does an MCP interaction expose data, a service action, or both; what identity, policy, audit, and lifecycle context applies?
Privileged administration	PAM/vault, cloud control plane, administrative API, service account, or workload access boundary.	Who may use retained privileged credentials, for what approved purpose, with what session/audit controls?

5.8 Delegation, federation, and workload identity

Delegation must preserve enough information to establish the subject, actor, purpose, audience, scope, and relevant chain of action. OAuth 2.0 and OpenID Connect can establish delegated authentication/authorization context; OAuth 2.0 Token Exchange (RFC 8693) can support subject/actor propagation across service and agent chains; DPoP can reduce bearer-token replay risk by sender-constraining tokens; and Rich Authorization Requests can express transaction-specific authorization details. See Appendix C.2.1.

Workload identity establishes cryptographically verifiable runtime identity for services and agents. SPIFFE/SPIRE, cloud workload identity, certificate systems, and service-mesh patterns can establish mTLS, rotation, and trust context. They authenticate workloads; policy still determines whether an authenticated workload may perform the requested data or service operation.

5.9 Policy administration, PDPs, and policy-as-code

Authorization policies should be versioned, testable, reviewable, deployable, observable, and reversible. They can be expressed in policy-as-code engines, relationship-authorization services, native platform policies, or other PDP capabilities. Policies should not be hidden as unreviewed application logic where the enterprise cannot test, audit, govern, or compare their behavior.

Relevant inputs may include authenticated identity, subject/actor relationship, role/relationship, resource, operation, classification, purpose, location, device/session, time, transaction state, approval state, risk signal, threshold, rate, and policy version.

A coordinated enterprise/native PDP portfolio should document which decisions are enterprise-wide, which are local, why local semantics are required, how policies align, how changes propagate, how exceptions are represented, how availability/fallback works, and how decisions are evidenced.

5.10 Privileged access, vaults, and retained credentials

Vault/PAM/secrets platforms protect retained privileged credentials and secrets. They can enforce named checkout/check-in, ownership, rotation, session control, controlled release, recording, and audit. These controls are essential when a retained privileged credential is required, such as a sensitive administrator account, privileged service credential, root-level cloud credential, or sensitive API secret.

Short-lived delegated tokens, workload identity documents, and runtime artifacts are discussed under authentication, delegation, credential issuance, and real-time authorization context as appropriate. A product may provide both vault/PAM functions and dynamic credential/token functions; the workbook distinguishes these roles rather than collapsing them.

5.11 Authorization posture, CIEM, and evidence

Pillar A produces evidence about identity lifecycle, ownership, purpose, entitlements, certification, credential use, policy version, decisions, exceptions, and policy coverage. This evidence supports authorization posture.

Authorization posture spans three connected responsibilities:

- Pillar A administers authorization objects, provides evidence, and corrects identity/entitlement/credential/delegation/policy conditions.
- Pillar D performs or consumes security-oriented CIEM, cloud-permission, identity-security, effective-access, drift, anomaly, and posture analysis.
- Pillar E evaluates risk treatment, accountability, exceptions, portfolio priority, and compliance with segregation-of-duties policies governing cloud-infrastructure entitlements.

5.12 Implementation patterns

Pattern A1: IGA as authoritative entitlement source

IGA should be the authoritative provisioning and review source for AI-application access where applicable—not merely for conventional applications. Identity and entitlement changes should propagate to AI platforms and connected systems through controlled interfaces such as SCIM, native connectors, or governed APIs. Human, personal-agent, and enterprise-agent identities require distinct lifecycle, owner, purpose, and certification treatments.

Pattern A2: Policy-as-code

Manage data/service authorization policy as versioned, testable artifacts. Review policy changes before production promotion; run tests in delivery pipelines; record the policy version used in material decisions; and preserve rollback/exception controls.

Pattern A3: Differentiated agent identity governance

Personal agents receive constrained authority derived from the sponsor's approved scope. Enterprise agents receive purpose-bound authority from an accountable enterprise owner. Neither should accumulate privileges simply because an agent becomes more capable or more connected.

Pattern A4: Zero-trust application-layer authorization

Authenticate and authorize AI service calls at the identity and application layers, not merely through network location, firewall, or IP allowlist. Service-to-service calls should carry verifiable identity and delegation context; Data and Service boundaries should enforce policy appropriate to their resource semantics.

Pattern A5: Workload identity for agents and services

Use workload identity to reduce reliance on long-lived shared secrets for distributed agent and service environments. Bind the workload identity into authorization decisions where relevant; rotate credentials; and retain deployment/attestation evidence.

5.13 Mini-case: Vendor collaboration copilot

AI application context. A collaboration-platform vendor introduces a copilot capability that can summarize meetings, search files, draft messages, and invoke approved workflow connectors. The capability is discovered in the enterprise tenant but has not been explicitly enabled.

Why Authorization matters. The enterprise must identify the business owner, tenant activation state, user/group eligibility, personal-agent or

enterprise-agent identity model, delegated scope, connector entitlements, and any default permissions inherited from the collaboration platform.

Representative control/technology context. An IdP/conditional-access platform, IGA access-review process, platform-native permissions, and policy/PDP integration may work together. OAuth/OIDC delegation and vendor documentation can clarify subject/actor scope.

Relevant interoperability references. See Appendix A canonical use cases for identity context, standing entitlement, and lifecycle; see Appendix B Pattern B.2.3.

Evidence/test. Verify that disabled/unauthorized tenant state prevents access; test user-group activation, connector scope, delegated identity evidence, and removal of access when an employee or agent lifecycle changes.

Worksheet connection. Use Worksheets D.2.1, D.3.1, D.5.1, and D.5.4.

5.14 Workbook prompts

- A1. What human, personal-agent, enterprise-agent, workload, service, and delegated identities participate?
- A2. Who owns each non-human identity, and what business purpose and lifecycle conditions apply?
- A3. What standing entitlements, vaulted credentials, and real-time policy decisions govern material operations?
- A4. Where can a standing entitlement receive increased assurance through request, approval, time-bound, review, segregation-of-duties, or evidence controls?
- A5. How do subject/actor, purpose, audience, and scope propagate through material delegation chains?
- A6. Which PDPs make decisions, what policy boundaries exist, and how are local/native decisions aligned?
- A7. Which Data and Services PEPs enforce relevant decisions?
- A8. How are lifecycle, certification, revocation, exception, suspension, and retirement handled?
- A9. What authorization posture, CIEM, cloud-permission, and segregation-of-duties evidence exists?

---

6. Pillar B — Data

Identity-Aware Data Operations

The enterprise capability that governs data operations and enforces policy at the points where AI applications retrieve, query, filter, abstract, and use governed data.

6.1 The problem

AI applications can reach data at machine speed across document repositories, SaaS applications, data platforms, warehouses, lakes, lakehouses, semantic layers, APIs, search services, vector indexes, graphs, and knowledge systems. A retrieval path that lacks identity-aware authorization, metadata filtering, and evidence can place data into model or agent context before later controls can meaningfully reduce exposure.

Pillar B governs the data operation: what data may be reached, in what scope, with what classification, purpose, owner, freshness, metadata, filtering, representation, abstraction, transformation, lineage, derived-artifact treatment, and evidence.

RAG is an important pattern, but it is not the whole Data problem. Identity-aware data operations include retrieval, query, filtering, transformation, inference, aggregation, correlation, abstraction, derived data, and governed delivery of data artifacts.

6.2 Core principle

Apply identity-aware authorization and required metadata/classification filters before data reaches the model or agent context. The enterprise should make Data PEPs explicit, testable, evidenced, and aligned with Authorization policy/identity context.

Missing, stale, disputed, or incomplete metadata is a risk condition. It can require narrower scope, remediation, compensating controls, elevated governance review, or a decision not to expose data. It is not a reason to bypass policy.

6.3 Scope of data operations

Data-operation class	**Examples**
Retrieval and access	Document/object access; table, row, column, field, API, graph, or repository retrieval; search and query.
Filtering and representation	ACL/entitlement filtering, classification filtering, masking, redaction, suppression, tokenization, pseudonymization, de-identification, semantic abstraction, controlled views.
Semantic and graph retrieval	Vector similarity search, semantic retrieval, RAG context assembly, graph traversal, node/edge/path access, knowledge graph inference.
Transformation and derivation	Aggregation, correlation, join, classification, scoring, analysis, summarization, feature/metric creation, derived artifact generation.

Data-operation class	Examples
Delivery and export	Governed data-product delivery, download, export, downstream transfer, retention, and reuse.

6.4 Data PEPs

A Data PEP applies authorization and related constraints at a data-operation boundary. It can exist in a repository, content-management system, warehouse, lakehouse, data product, semantic layer, data API, search/indexing service, vector retrieval service, graph system, data mesh gateway, transformation pipeline, or governed abstraction service.

A Data PEP should receive sufficient context to evaluate relevant identity, subject/actor, resource, operation, classification, purpose, owner, metadata, policy, and constraints. It should not rely solely on an upstream application claim that filtering occurred.

6.5 Retrieval, search, vector, graph, and RAG controls

Retrieval systems—including enterprise search, semantic/vector search, graph retrieval, and RAG pipelines—must apply authorization and metadata filtering before results are returned to an AI application's context window, memory, prompt, or agent state.

Key design questions include:

- Are source ACLs, entitlements, sensitivity labels, classifications, ownership, and retention conditions correctly represented in the index or query path?
- Do indexing and embedding pipelines preserve required metadata and remove data when source access or lifecycle changes?
- Does query-time filtering use verified identity/delegation context rather than only application-supplied attributes?
- Are graph nodes, edges, paths, relationships, and inferences governed rather than assuming graph traversal is harmless?
- Can the enterprise test stale ACL propagation, stale embeddings, missing metadata, index lag, cross-tenant boundary failures, and overbroad semantic recall?

6.6 Data abstraction, representation, and transformation boundary

Data abstraction remains a major IAAI capability even though it is not a separate pillar.

When an AI application reads data and the data platform automatically applies an existing identity-aware policy—such as row/column filtering, masking, suppression, de-identification, controlled semantic view, or permitted representation—the result is a Pillar B Data response to an identity-authorized request. The AI application is receiving the representation that the data platform policy already permits.

When an agent requests a net-new abstraction, transformation, aggregation, correlation, derivation, write, or other data change, the request is a Pillar C Service operation. It requires Service PEP treatment, purpose/parameter/approval/autonomy constraints, and evidence. The resulting operation can still involve Pillar B because it reads or writes governed data, creates derived artifacts, or requires Data PEP enforcement.

The practical heuristic is:

- Read/retrieval operations can often be served by Pillar B without complication from Pillar C.
- Write/transformation operations ordinarily require Pillar C identity-aware service operations, even when their output is data.

6.7 Derived artifacts, lineage, and downstream use

A permissible source retrieval does not automatically authorize every transformation or downstream artifact. Governed data operations should consider whether aggregation, correlation, join, inference, summarization, scoring, feature creation, semantic abstraction, export, retention, reuse, or delivery creates a representation with a different sensitivity, audience, purpose, or risk profile.

Maintain appropriate lineage from source to derived artifact. Record classifications, owner, policy, transformation, retention, downstream recipients, and evidence sufficient to investigate or reassess the artifact.

6.8 Evidence and feedback

Data-operation evidence should include, as appropriate, the AI application, subject/actor/workload context, source/product/index, classification/metadata state, query/retrieval scope, filters applied, policy/decision context, transformation, derived artifact, delivery/export, exception, and outcome. This evidence supports Pillar D correlation and Pillar E assurance.

6.9 Implementation patterns

Pattern B1: Source-authorized retrieval

Apply source authorization at or before the retrieval boundary. Do not assume that a broad AI application service account can safely retrieve all data merely because an agent later promises to filter it.

Pattern B2: Metadata-aware retrieval policy

Use classification, purpose, owner, retention, geography, freshness, lineage, and other metadata as inputs to Data PEP decisions. Build remediation and safe restriction paths for absent or stale metadata.

Pattern B3: Entitlement-aware semantic retrieval

Filter vector, semantic, search, and RAG results by verified identity/entitlement and policy metadata before returning content to an AI application. Test index synchronization and alternate retrieval paths.

Pattern B4: Governed data products and abstraction services

Use curated data products, semantic layers, controlled views, aggregates, and de-identified representations to provide AI applications with an appropriate data representation rather than unrestricted raw records.

Pattern B5: Derived-artifact governance

Treat material derived data as a governed artifact with owner, lineage, classification, retention, policy, and downstream-use conditions.

6.10 Mini-case: HR knowledge assistant

AI application context. An internal HR assistant retrieves policy documents, benefits information, job architecture, and employee-related guidance from a document repository, enterprise search index, and data platform.

Why Data matters. The assistant must not load restricted employee records, legal investigations, compensation data, or stale policy documents into context merely because a user can access the assistant.

Representative control/technology context. Source-system permissions, sensitivity labels, enterprise search ACL propagation, vector metadata filters, policy-aware data APIs, and governed semantic views can provide layered controls.

Relevant interoperability references. See Appendix A data-operation use cases; Appendix B Patterns B.2.1, B.2.4, and B.2.6.

Evidence/test. Test a user changing roles, stale index permissions, missing classification, denied document retrieval, and retrieval of a permitted policy summary versus prohibited raw employee records.

Worksheet connection. Use Worksheets D.5.2, D.5.4, and D.5.5.

6.11 Workbook prompts

- B1. What sources, repositories, indexes, graphs, APIs, data products, and derived artifacts are reachable?
- B2. What classification, owner, purpose, retention, lineage, freshness, and quality metadata is required?
- B3. Where is the Data PEP for each material operation?
- B4. How are entitlement and metadata filters applied before semantic/vector/graph/RAG results reach AI context?
- B5. Which representations are automatically provided by existing data-platform policy, and which requested transformations become Services operations?
- B6. What transformations, derivations, exports, and downstream uses are permitted, prohibited, or constrained?
- B7. What happens when policy, metadata, source availability, identity, or index state is incomplete or stale?
- B8. What evidence supports investigation, assurance, and lifecycle reassessment?

---

7. Pillar C — Services

Identity-Aware Service Operations

The enterprise capability that governs tools, APIs, MCP operations, workflows, communications, transactions, administrative functions, and other consequential actions performed by AI applications.

7.1 The problem

AI applications can invoke tools, APIs, MCP capabilities, workflows, applications, communications, transactions, and infrastructure control planes. These actions can create business, legal, financial, operational, security, or external consequences. A permission to retrieve data does not automatically authorize an agent to create a record, send a message, submit a purchase request, change a customer status, deploy infrastructure, or initiate a transaction.

Pillar C governs the service operation: whether an actor may invoke a capability, against which object or recipient, with which method, parameter, threshold, workflow state, rate, time, destination, external-effect constraint, approval condition, and autonomy limit.

7.2 Core principle

Place Service PEPs at material tool, API, workflow, application, communication, transaction, and infrastructure boundaries. Preserve relevant subject, ac-

tor, delegation, purpose, scope, approval, and policy context through chained operations. Constrain autonomy through purpose, action scope, approval state, confidence, reversibility, rate, thresholds, escalation conditions, and stop/pause mechanisms.

7.3 Service-operation scope

Service-operation class	Examples
Tools and APIs	Invoking a tool, API method, application function, or service endpoint.
MCP interactions	Using MCP tools, resources, prompts, or server-provided capabilities that expose data operations, service operations, or both.
Workflows and business processes	Starting, approving, denying, routing, scheduling, escalating, or completing workflow steps.
Communications and delivery	Sending email, chat, notifications, publishing content, delivering documents, or communicating with employees, customers, partners, suppliers, regulators, or other external parties.
Transactions and commitments	Procurement, payment, order, contract, approval, booking, account, pricing, or other financial/legal commitment.
Administrative and infrastructure actions	Provisioning, configuration, deployment, cloud, Kubernetes, infrastructure-as-code, identity, or other privileged control-plane action.
Agent-to-agent/service actions	Delegation, task assignment, service invocation, handoff, and autonomous action chaining.

7.4 Service PEPs

A Service PEP applies authorization and constraints at the service-operation boundary. It can exist in an API/service gateway, MCP authorization component, workflow engine, enterprise application, SaaS platform, cloud control plane, service mesh, agent runtime, orchestration layer, or other application boundary.

A Service PEP should evaluate more than identity alone where relevant: method, parameter, object, recipient, destination, threshold, workflow state, rate, time, location, purpose, data classification, delegated scope, approval condition, expected external effect, and reversibility.

7.5 MCP as interoperability context and control surface

MCP is an interoperability context and potential control surface, not a pillar or a substitute for authorization. It can expose a data resource, a service/tool operation, or both.

- An MCP interaction that retrieves governed data requires Pillar B Data analysis and Data PEP enforcement.
- An MCP interaction that invokes a tool, workflow, or consequential action requires Pillar C Service analysis and Service PEP enforcement.
- Enterprise-managed MCP authorization can connect the interaction to enterprise identity, scoped token, PDP, audit, lifecycle, and policy controls.
- MCP servers and related components can emit events to Pillar D for correlation, detection, investigation, and response.

7.6 Data and Services layering

A service operation often retrieves, changes, or delivers governed data. The architecture preserves both questions:

- Pillar B governs data operation: permitted source/scope, filtering, metadata, existing platform abstraction, transformation, lineage, and derived artifact.
- Pillar C governs service operation: invocation, method, parameter, workflow, recipient, destination, transaction, approval, external effect, and autonomy.

For example, an agent can be permitted to retrieve a supplier profile through a Data PEP but prohibited from sending that profile to an external supplier, changing a vendor record, or submitting a purchase request without Service PEP constraints and approval.

7.7 Bounded autonomy and tool chaining

Tool chaining and autonomy are architecture and governance decisions, not merely implementation conveniences. Define bounded autonomy explicitly:

- Permitted task classes and business purpose.
- Reachable tools, APIs, MCP capabilities, services, and data products.
- Permitted methods, parameters, objects, recipients, destinations, and transaction limits.
- Approval gates, confidence thresholds, separation-of-duties conditions, and workflow-state constraints.
- Rate limits, time windows, budget limits, reversibility, and rollback paths.
- Escalation triggers, pause/kill capability, and human review conditions.
- Evidence and audit requirements.

7.8 High-consequence paths

Communications, financial/legal commitments, customer-facing changes, privileged administration, and infrastructure control-plane operations commonly merit Tier 3 or Tier 4 attention. They may require deeper authority-path analysis, explicit constraints, enhanced testing, documented governance decisions, and Companion Authorization-Risk Exercise routing.

7.9 Evidence and feedback

Service-operation evidence should capture actor, subject/delegation context, service/tool, method, parameter or parameter class, affected object, recipient/destination, workflow/approval state, decision, enforcement, resulting effect, exception, and outcome. This supports Pillar D investigation and Pillar E assurance.

7.10 Implementation patterns

Pattern C1: Method- and parameter-aware API authorization

Do not treat an API token or connection permission as blanket authorization. Constrain specific methods, objects, parameter ranges, thresholds, destinations, and rates.

Pattern C2: Workflow-state and approval-aware action

Use workflow state, segregation of duties, approval, time, and threshold conditions to constrain whether an AI application may initiate, advance, or complete a business process.

Pattern C3: Controlled external communications

Treat recipient, destination, attachment, content type, publishing state, and approval as service-operation constraints. Pair with Data controls where governed data is included.

Pattern C4: Bounded infrastructure administration

Use workload identity, PAM/vault controls, Service PEPs, policy, approval, blast-radius constraints, telemetry, and SecOps monitoring for administrative or infrastructure actions.

Pattern C5: MCP capability governance

Treat an MCP server's data/tool resources as governed capabilities. Bind them to enterprise identity, policy, audit, lifecycle, and appropriate Data/Service PEPs.

7.11 Mini-case: Procurement support agent

AI application context. A procurement agent retrieves supplier and contract information, drafts purchase requests, routes approvals, and can submit a request to a procurement system under defined conditions.

Why Services matters. Retrieval of supplier data does not authorize submission of a purchase request, selection of a vendor, change of threshold, communication with a supplier, or commitment of enterprise funds.

Representative control/technology context. A Service PEP at the procurement API or workflow platform can enforce method, cost center, supplier, amount, category, approval state, delegation, and threshold constraints. Data PEPs govern supplier/contract retrieval.

Relevant interoperability references. See Appendix A service-operation use cases; Appendix B Patterns B.2.2, B.2.4, B.2.5, and B.2.8.

Evidence/test. Attempt threshold override, unauthorized supplier change, duplicate submission, bypass of workflow state, external communication to an unapproved recipient, and policy-service unavailability.

Worksheet connection. Use Worksheets D.5.2–D.5.6 and D.6.1.

7.12 Workbook prompts

- C1. What services, tools, APIs, MCP contexts, workflows, communications, transactions, and control planes are reachable?
- C2. Where is the Service PEP for each material action?
- C3. Which method, parameter, object, recipient, destination, rate, threshold, time, workflow, approval, and external-effect constraints apply?
- C4. Where does a service operation also invoke Data operations, and are both PEPs visible?
- C5. What autonomy, escalation, approval, reversibility, and stop conditions are defined?
- C6. Which actions create financial, legal, customer, privileged, regulatory, or infrastructure consequences?
- C7. What failure, fallback, bypass, or policy-unavailability behavior applies?
- C8. What service-operation evidence supports investigation and assurance?

---

8. Pillar D — SecOps

Post-AI Security Operations

The enterprise security capability that applies compensating controls, visibility, detection, investigation, response, resilience, and technical feedback to AI applications and the broader post-AI enterprise.

8.1 The problem

Pillar D is not merely an AI "safety net." It represents the broader enterprise security estate adapting to AI contexts: prompts, model interactions, retrieved context, agents, tools, APIs, workflows, data flows, services, workloads, endpoints, browsers, infrastructure, suppliers, and business processes.

AI changes the character of familiar risks. It can accelerate data movement, expand service reach, add agent/tool interfaces, introduce new model and connector dependencies, and enable prompt injection, unsafe behavior, model abuse, agent compromise, and novel interaction paths.

AI also accelerates the attacker. Adversaries can use AI to scale reconnaissance, craft tailored social engineering, generate or adapt malicious code, probe weak configurations, manipulate AI interactions, automate exploitation, and iterate rapidly against defensive controls. Defenders must protect a large, changing estate of identities, devices, workloads, applications, data paths, and business processes; an attacker may need only one exploitable weak link.

Identity-aware AI security reduces this asymmetry by reducing practical reach and improving the fidelity of legitimate-activity evidence. When the enterprise can identify actor, delegation, permitted data/service operation, policy, expected scope, and result, compensating controls and behavioral analytics can make more precise decisions.

8.2 Core principle

Pillar D is compensating, detective, investigative, responsive, posture-improving, and feedback-producing capability. It does not replace Pillar A authorization or Pillar B/C PEP enforcement. It can block or constrain risky activity when policy coverage, integration, precision, or threat semantics are incomplete, but it should not be treated as the primary entitlement engine or PDP.

Pillar D must also use AI responsibly to accelerate defense. AI-assisted security operations can correlate high-volume signals, identify patterns, prioritize alerts, summarize incidents, recommend containment, generate investigation hypotheses, assist detection engineering, and automate bounded response actions. These uses are themselves AI applications and must be governed through the same five-pillar architecture.

8.3 SecOps capability domains

Domain	Role
AI guardrails and safety frameworks	Inspect and constrain prompts, context, outputs, tool selection/calls, agent behavior, and action sequences against safety/guardrail policy. Compensating controls, not substitutes for authorization.
AI-aware DLP	Inspect sensitive data in prompts, context, outputs, files, API payloads, browser interactions, and transfers; block, redact, watermark, quarantine, alert, and log where appropriate.
AI threat, model, and runtime security	Address prompt injection, jailbreaks, poisoning, model theft, adversarial inputs, malicious connectors/tools, runtime compromise, and unsafe behavior.
AI red teaming and adversarial testing	Test prompts, retrieval, tools, agents, identities, integrations, models, pipelines, and service actions; feed findings into architecture improvement.
Activity observability	Capture agent traces, governed prompts/context where appropriate, retrievals, tool calls, policy decisions, workflow actions, and outcomes.
SIEM and SOAR	Collect, correlate, analyze, detect, investigate, report, contain, and coordinate bounded response.
DSPM	Discover, classify, assess, and prioritize AI-reachable data exposure, metadata gaps, index scope, and misconfiguration.
Identity/cloud permission posture	Analyze effective permissions, excessive privilege, drift, toxic combinations, workload/service-account risk, and anomalies.
Endpoint, browser, API, network, workload, cloud, and runtime security	Protect AI activity across managed devices, browsers, APIs, integrations, services, containers, infrastructure, and deployments.
Model/pipeline/artifact/ supply-chain security	Protect models, packages, dependencies, containers, agent definitions, prompts/templates, connectors, indexes, pipelines, and artifacts.

8.4 Telemetry boundary and end-to-end correlation

Authentication, authorization, Data-operation, and Service-operation systems must emit usable evidence. Emitting telemetry does not make the originating technology a Pillar D technology. Pillar D begins where capabilities collect, inspect, correlate, analyze, detect, investigate, block, quarantine, remediate, or otherwise act on security-relevant signals.

SecOps should correlate relevant context across identity, delegation, policy decision/version, data operation, service operation, MCP event where applicable, guardrail/DLP event, workload activity, and outcome. Correlation IDs, common timestamps, actor/subject context, policy identifiers, resource identifiers, and governed retention/access to evidence are important design requirements.

8.5 Compensating controls and preventative controls

Guardrails and DLP can block or constrain activity that is otherwise authorized. They are valuable where a relevant AI application is not covered by a PDP, policy coverage is incomplete, the threat is behavioral/semantic, an integration is immature, or residual risk requires additional protection.

Their presence does not prove that preventative authorization is adequate. A guardrail/DLP event should create evidence and, where appropriate, technical feedback to Authorization, Data, Services, and Governance.

8.6 AI security inspection and detection patterns

Pattern D1: Prompt, context, and output inspection

Inspect governed prompts, retrieved context, outputs, and policy-relevant behavior for sensitive-data movement, unsafe instruction, prohibited content, prompt injection, policy violation, and anomalous activity. Retain enough evidence to investigate without indiscriminately retaining sensitive content.

Pattern D2: DLP and data-movement controls

Inspect and constrain sensitive data movement across prompts, responses, files, APIs, browser interactions, and external destinations. Pair with Data and Services controls rather than assuming DLP establishes original authorization.

Pattern D3: Identity-aware anomaly detection

Use subject, actor, workload, policy, entitlement, data/service operation, expected scope, and outcome context to improve behavioral baselines and anomaly detection.

Pattern D4: Posture and exposure management

Use DSPM, CIEM, cloud posture, identity security, runtime posture, and configuration analysis to identify exposed data, stale metadata, unsafe connectors, excessive authority, toxic combinations, weak boundaries, and drift.

Pattern D5: AI-specific SIEM/SOAR detections

Detect suspicious delegation, abnormal retrieval, anomalous tool use, policy-denial bursts, guardrail/DLP violations, unusual external actions, workload identity misuse, and correlated control failures. Use SOAR for bounded containment, evidence preservation, ticketing, and remediation coordination.

Pattern D5a: AI-assisted security operations

Use AI to accelerate defensive work where the AI application is governed, evidence-based, and proportionate to risk. Appropriate examples include alert triage, event correlation, investigation summarization, detection-content assistance, threat-hunting hypotheses, case prioritization, remediation recommendation, and bounded SOAR actions.

A SOC copilot or agent must not silently receive broader data or service authority than the analyst, responder, or approved enterprise purpose allows. It needs defined identity, data-operation scope, service-operation scope, autonomy limit, approval conditions, evidence trail, and pause/review capability. High-consequence containment actions should remain bounded by policy, confidence thresholds, reversibility, and human escalation where appropriate.

Pattern D6: Red teaming and resilience testing

Test identity, delegation, policy, retrieval, tool use, autonomy, service action, guardrail/DLP coverage, telemetry, and response under misuse, failure, bypass, stale-context, and adversarial conditions.

8.7 Feedback and corrective response

Material SecOps findings should result in owned, verifiable correction:

- Authorization: identity, entitlement, credential, delegation, lifecycle, or policy correction.
- Data: source, classification, metadata, filtering, lineage, index, or Data PEP correction.
- Services: tool/API/workflow, Service PEP, method/parameter, autonomy, approval, or destination correction.
- Governance: risk treatment, exception, lifecycle, portfolio priority, assurance, or accountability record.

8.8 Supply chain and resilience

Secure models, packages, dependencies, containers, agent definitions, prompts/templates, connectors, indexes, pipelines, and artifacts. Test policy-service availability, telemetry loss, degraded identity source, stale metadata, fallback behavior, manual override, exception expiry, and recovery paths. Explicitly decide safe failure posture for each material boundary rather than allowing accidental fail-open behavior.

8.9 Standards, methods, and interoperability considerations

A heterogeneous enterprise relies on standards, protocols, policy languages, open interfaces, architecture patterns, frameworks, and vendor integrations to carry identity, lifecycle, delegation, workload trust, authorization context, data/service context, telemetry, and governance evidence.

These are not all the same kind of thing. Formal standards, emerging specifications, policy languages, open-source projects, vendor APIs, architecture patterns, and governance frameworks should be evaluated according to role, maturity, implementation support, operational fit, and security properties.

Area	Illustrative foundations	Role
Identity and lifecycle	OAuth 2.0, OpenID Connect, OAuth Token Exchange, SCIM, SPIFFE/SPIRE	Authentication, federation, delegation, provisioning, workload identity
Authorization	PDP/PEP, policy-as-code, Rego/OPA, Cedar, relationship authorization, AuthZEN	Policy definition, decision, enforcement interoperability
Agent/tool/ service interoperability	MCP, enterprise-managed MCP authorization, OpenAPI, API gateways, service mesh	Controlled connectivity and enforcement context
Telemetry	OpenTelemetry, cloud audit logging, SIEM interfaces	Evidence correlation, detection, response, assurance
Governance	NIST AI RMF, NIST SP 800-207, ISO/IEC 42001, applicable obligations	Risk, control, accountability, and assurance framing

See Appendix C for distinctions, maturity, representative examples, and implementation cautions.

8.10 Mini-case: SOC investigation copilot

AI application context. A SOC copilot summarizes correlated alerts, retrieves security telemetry, drafts investigation steps, recommends containment, and can initiate a bounded SOAR playbook after approval.

Why SecOps matters. The copilot may improve response speed but can also access sensitive logs, expose investigation content, recommend harmful actions, or gain excessive containment authority.

Representative control/technology context. SIEM/SOAR, identity-aware data access, Service PEPs for response actions, AI guardrails, case-management approval, bounded automation, and evidence logging work together.

Relevant interoperability references. See Appendix B Patterns B.2.5–B.2.7 and B.2.9.

Evidence/test. Test whether the copilot can access only permitted cases/logs; verify approval before high-consequence action; simulate prompt injection and telemetry loss; trace a finding through correction and Governance closure.

Worksheet connection. Use Worksheets D.5.1–D.5.6, D.6.1, and D.7.1.

8.11 Workbook prompts

- D1. Which preventative controls are incomplete, and which compensating controls address residual risk?
- D2. What guardrail, DLP, browser, API, runtime, endpoint, and infrastructure controls apply?
- D3. Which identity, policy, Data, Services, MCP, and outcome events can be correlated end-to-end?
- D4. What detections identify abnormal delegation, retrieval, tool use, external effect, or policy failure?
- D5. How do posture findings become technical correction and Governance records?
- D6. What response playbooks, containment powers, approvals, and evidence rules apply?
- D7. What red-team, adversarial, bypass, stale-context, and resilience tests are required?
- D8. How are models, connectors, indexes, packages, and pipelines secured?
- D9. What telemetry-loss and policy-unavailability scenarios have been tested?
- D10. How are AI-generated security recommendations reviewed before material action?
- D11. What feedback loops return findings to Authorization, Data, Services, and Governance?
- D12. Which shared SecOps improvements reduce risk across multiple AI applications?

- D13. What residual risks require Governance treatment?
- D14. Where could governed AI assistance materially reduce security-response time or improve detection quality? For each proposed use, what identity, data-operation scope, service-operation scope, autonomy limit, approval condition, and evidence trail are required?

---

9. Pillar E — Governance

Enterprise AI Governance

The peer enterprise function that establishes purpose, decision rights, risk appetite, accountability, lifecycle direction, assurance, and portfolio learning for AI applications.

9.1 The problem

Governance is not an implementation-only technical control plane. It establishes the conditions under which the enterprise uses AI: purpose, business ownership, value sought, risk appetite, authority boundaries, lifecycle, attention tier, approval, exception treatment, assurance, evidence, and accountability.

Governance should not be a one-time committee review of every AI application. It is an iterative portfolio practice that directs technical and operational controls, receives evidence and findings, makes residual risk explicit, strengthens shared controls, and reprioritizes investment over time.

9.2 Core principle

Pillar E governs aggregate AI risk and the governance component of identity-aware authorization. It decides or directs what the enterprise is willing to permit, constrain, defer, prohibit, accept, test, evidence, and improve. It does not replace technical control owners; it ensures their work is connected to purpose, risk appetite, accountability, and assurance.

A manual committee is a valid starting point. Over time, decisions should become tracked technical and operational requirements with owners, evidence, implementation work, review dates, and closure conditions.

9.3 Governance of the AI-application portfolio

Every discovered AI application remains visible in the portfolio. This includes homegrown applications, pilots, enterprise agents, embedded vendor capabilities, vendor-introduced AI applications discovered but not explicitly enabled/authorized, and applications that are only partly understood.

Governance maintains or directs:

- Business purpose, value sought, sponsor, business owner, technical owner, vendor/supplier, and lifecycle state.
- Data/service reach, authority/autonomy pattern, dependencies, and external effects.
- Portfolio attention tier, reassessment triggers, review depth, and evidence expectations.
- Risk appetite, prohibited uses, required approvals, policy direction, and control conditions.
- Exception, residual-risk, risk-acceptance, and remediation treatment.
- Assurance, compliance, audit, board reporting, and portfolio trend analysis.

9.4 Governance decision and governance event

A **governance decision** is a human deliberation or subjective determination: approve, constrain, defer, prohibit, accept risk, require evidence, impose a condition, or change policy direction.

A **governance event** is a processed record or action that implements or documents governance-driven change: an approval record, attention-tier assignment, exception issuance, review completion, policy-change request, implementation ticket, evidence update, or closure verification.

The distinction matters because a technical workflow can record or execute a governance event, but it should not be mistaken for the underlying decision when judgment or accountable risk acceptance is required.

9.5 Governance and control interaction

Governance direction must become verified operating controls:

```
Governance purpose, risk appetite, decision, or condition
→ Authorization / Data / Services / SecOps requirements
→ implementation, test, evidence, and operating record
→ Governance assurance, risk treatment, learning, and reprioritization
```

The interaction label is neutral even though the explanation often describes how governance direction becomes control requirements. Evidence, incidents, posture findings, and exceptions also inform Governance in return.

9.6 Effective authority, authorization posture, and CIEM governance

Governance evaluates effective authority—not only intended roles or policy statements. It considers effective entitlements, cloud permissions, service accounts, workload identities, local policy blind spots, excessive privilege, toxic combinations, separation-of-duties concerns, drift, exceptions, and the combined data/service reach of AI applications.

Authorization administers access objects and evidence. SecOps performs or consumes security-oriented CIEM, cloud-permission, effective-access, posture, and anomaly analysis. Governance evaluates risk treatment, accountability, portfolio priority, and compliance with segregation-of-duties policies governing cloud-infrastructure entitlements.

9.7 Exceptions and risk treatment

Exceptions must be explicit, owned, time-bound where feasible, evidenced, monitored, reviewed, and connected to remediation or renewal decisions. Each exception should identify:

- The affected AI application, authority path, control, and pillar(s).
- The business rationale and relevant value.
- The authority/consequence affected.
- Compensating controls and monitoring.
- Residual risk and risk-acceptance authority.
- Start, expiry, review, remediation, and closure conditions.

Risk acceptance is not a missing decision, stale ticket, unowned workaround, or an indefinite "temporary" condition.

9.8 Lifecycle and assurance

Governance directs discovery, inventory, evaluation, pilot, activation, material-change review, deployment, periodic reassessment, incident/posture-triggered reassessment, restriction, suspension, and retirement. It can require agent/component restriction or termination when active governance conditions are no longer satisfied.

Assurance tests both design and operation: whether governance direction became technical change, whether controls operated as intended, whether evidence is complete, whether exceptions remain valid, and whether residual risk remains within accepted appetite.

9.9 Implementation patterns

Pattern E1: Portfolio intake and lifecycle governance

Use a common portfolio record for AI applications, including vendor-introduced capabilities. Preserve unknowns; assign owners; capture lifecycle and activation state; and establish reassessment triggers.

Pattern E2: Proportionate attention and routing

Use attention tiers to route review, evidence, testing, governance, and Companion Authorization-Risk Exercise decisions. Do not use tiers as a reason to omit applications from visibility.

Pattern E3: Governance requirements as control work

Translate decisions into specific Authorization, Data, Services, and SecOps requirements with accountable owners, milestones, test/evidence expectations, exception treatment, and closure criteria.

Pattern E4: Effective-authority risk treatment

Use authorization posture, CIEM, and SecOps analysis to identify excessive privilege, toxic combinations, cloud-infrastructure separation-of-duties concerns, drift, and exception risk. Decide treatment, track remediation, and verify closure.

9.10 Mini-case: Vendor CRM agentic sales feature

AI application context. A CRM vendor introduces an agentic feature that can summarize accounts, draft customer messages, recommend next actions, update records, and eventually invoke workflow automations. The feature is discovered in the tenant but not explicitly enabled.

Why Governance matters. The enterprise must identify value sought, business owner, allowed lifecycle state, data/service reach, vendor evidence, configuration choices, external customer impact, sales-policy conditions, and the decision to evaluate, pilot, enable, restrict, or decline the feature.

Representative control/technology context. AI governance workflow, GRC/exception tracking, CRM permissions, Data/Service PEPs, vendor documentation, approval conditions, and SecOps evidence may all participate.

Relevant interoperability references. See Appendix B Patterns B.2.4, B.2.8, and B.2.9.

Evidence/test. Verify discovery/activation state, owner assignment, customer-data scope, outbound communication restrictions, record-update constraints, vendor audit evidence, and reassessment after vendor capability change.

Worksheet connection. Use Worksheets D.2.1–D.3.3, D.5.3, D.5.6, and D.6.2.

9.11 Workbook prompts

- E1. What business purpose, value, sponsor, owner, vendor/supplier, and lifecycle state apply?
- E2. Is the application homegrown, vendor-introduced, embedded, shared, or composite, and is it explicitly enabled/authorized?
- E3. What attention tier applies now, what is unknown, and what triggers reassessment?
- E4. Which decisions require business, security, data, privacy, legal, risk, or executive authority?

- E5. How are governance decisions translated into Authorization, Data, Services, and SecOps requirements?
- E6. What evidence returns to validate implementation, assurance, and residual risk?
- E7. What effective-access, CIEM, toxic-combination, drift, or cloud-infrastructure segregation-of-duties concerns exist?
- E8. What exceptions exist, who accepts risk, what compensating controls apply, and when are they reviewed or closed?
- E9. Which shared-control investments improve multiple AI applications?
- E10. Which authority paths require Companion Authorization-Risk Exercise routing?

---

10. Cross-Pillar Interoperability and Control Seams

10.1 Why interaction seams matter

A strong control can fail when required context or evidence is lost between responsibilities. Identity may not survive delegation. A Data PEP may not receive policy or metadata. A Service PEP may not constrain a parameter or recipient. A vector index may retain stale access metadata. A guardrail may detect a violation without producing correction. A governance decision may not become a technical control. Telemetry may not permit reconstruction of the resulting action.

This section provides the cross-pillar model. Appendix B provides detailed fixed patterns, questions, tests, evidence ideas, failure conditions, and representative examples. Appendix D provides assessment and roadmap instruments.

10.2 Principal interaction contexts

Interaction context	Participating pillars	Core question	Detailed reference
Authorization and Data-Operation Enforcement	Authorization; Data; Governance where applicable	Does the Data PEP receive usable identity, policy, metadata, purpose, and scope context before governed data reaches AI context?	Pattern B.2.1

Interaction context	Participating pillars	Core question	Detailed reference
Authorization and Service-Operation Enforcement	Authorization; Services; Governance where applicable	Does the Service PEP constrain actual method, parameter, object, recipient, destination, state, threshold, and effect?	Pattern B.2.2
Identity Propagation, Delegation, and Workload Identity	Authorization; Data; Services; SecOps; Governance	Can the enterprise reconstruct subject, actor, workload, chain, purpose, scope, and outcome?	Pattern B.2.3
Data and Service-Operation Layering	Data; Services; Authorization; Governance	Are data scope and consequential service context both governed when both are material?	Pattern B.2.4
Compensating Controls at Data and Service Boundaries	Data; Services; SecOps; Authorization; Governance	Do guardrails/DLP complement preventative enforcement and create feedback rather than obscure missing controls?	Pattern B.2.5
Operational Evidence into SecOps	Authorization; Data; Services; SecOps; Governance	Can security operations correlate decision, data operation, service action, event, and outcome?	Pattern B.2.6
Technical Security Feedback and Corrective Response	SecOps; Authorization; Data; Services; Governance	Do findings become owned, verifiable technical correction and accountable governance treatment?	Pattern B.2.7
Governance/ Control Interaction Patterns	Governance; Authorization; Data; Services; SecOps	Do purpose, risk, approval, conditions, and exceptions become specific controls and evidence?	Pattern B.2.8
Authorization Posture, CIEM, and Governance Risk Treatment	Authorization; SecOps; Governance; Data/Services where applicable	Does the enterprise govern effective authority, drift, toxic combinations, and exceptions?	Pattern B.2.9

10.3 Interaction-testing principles

Test normal operation and failure conditions:

- Policy unavailability or degraded PDP response.
- Stale or missing identity, delegation, entitlement, metadata, classification, index, or lifecycle state.
- Bypass or alternate path.
- Fallback and manual override.
- Partial enforcement, policy drift, or inconsistent native/local PDP behavior.
- Telemetry loss, schema mismatch, missing correlation ID, access limitation, or retention failure.
- Vendor feature activation/change without expected governance review.
- Exception expiry, control rollback, or incomplete remediation.

A useful test record identifies scenario, AI application/authority path, participating pillars, expected decision, enforcement boundary, expected evidence, failure posture, compensating response, accountable owner, and closure condition.

10.4 How the pillars work together

Authorization establishes managed identity, policy, lifecycle, and decision capability. Data and Services enforce relevant constraints where operations occur. SecOps observes, protects, detects, investigates, responds, and returns technical findings. Governance directs purpose, risk, lifecycle, exceptions, assurance, and portfolio investment. The architecture works when these responsibilities exchange enough context and evidence to constrain authority, explain decisions, and improve the system over time.

---

11. Implementing IAAI: Discovery, Assessment, Implementation, and Governance Lifecycle

11.1 Implementation objective

Implement IAAI as a recurring enterprise lifecycle, not a one-time architecture project. The aim is to increase awareness of the full AI-application estate, improve preventative authorization depth where risk warrants, use compensating/governance controls while gaps remain, validate operation with evidence, and continuously reprioritize shared and application-specific improvement.

The lifecycle has four phases: Discover, Assess, Implement, and Govern. The three initial adoption passes in Section 4 help enterprises enter the lifecycle without waiting for comprehensive maturity.

11.2 Discover

Discover includes identifying AI applications, inventorying them, linking shared components/dependencies, and conducting early segmentation.

Portfolio discovery activities

- Identify homegrown AI applications, embedded copilots, agents, workflows, model integrations, data-retrieval paths, tools/APIs/MCP contexts, and external effects.
- Identify vendor-introduced AI applications through tenant configuration, product documentation, vendor notices, procurement, architecture review, security discovery, data/service logs, browser activity, and business-unit inquiry.
- Record applications that are discovered but not explicitly enabled/authorized.
- Assign a durable application ID; sponsor; business/technical owner; vendor/supplier; lifecycle state; and initial business purpose/value.
- Link components: agents, models, runtime, data sources, connectors, tools, APIs, MCP servers, workflows, shared controls, and suppliers.
- Record unknowns rather than suppressing visibility.

Discover outputs

- Authoritative AI-application portfolio.
- Linked component/dependency inventory.
- Initial lifecycle/activation states.
- Ownership gaps and discovery backlog.
- Early authority/consequence indicators.

11.3 Assess

Assess includes remaining segmentation, attention tier, baseline/rebaseline, control-maturity view, evidence review, and routing.

Assess activities

- Estimate data sensitivity/reach, service authority, autonomy, external effect, financial/legal/safety consequence, privileged/infrastructure reach, cross-boundary exposure, change velocity, control maturity, and evidence quality.
- Assign provisional or evidence-supported attention tier.
- Run the Fast Diagnostic and record current maturity → target maturity by domain.
- Identify material Authority paths, PEP locations, interaction contexts, failures, dependencies, and vendor-control boundaries.

- Determine whether the AI application remains at managed baseline, requires targeted review, needs a material deep dive, must be restricted/deferred, or should route an authority path to the Companion Authorization-Risk Exercise.
- Rebaseline after material component change, vendor capability change, incident, posture finding, exception, lifecycle change, or new evidence.

Assess outputs

- Attention tier and rationale.
- Current/target control profile by Authorization, Data, Services, SecOps, and Governance domain.
- Gaps, unknowns, evidence needs, and reassessment triggers.
- Routing decision and priority.

11.4 Implement

Implement includes deepening analysis, designing controls, configuring/integrating them, testing, validating evidence, remediating findings, and creating reusable patterns.

Implement activities

- Map identity, subject/actor, delegation, workload, entitlement, credential, PDP, and policy boundaries.
- Map Data operations, metadata, Data PEPs, representations, transformations, derived artifacts, and evidence.
- Map Services operations, Service PEPs, methods, parameters, workflow states, recipients, thresholds, autonomy, and effects.
- Map SecOps controls, telemetry, correlation, detection, response, posture, red-team, resilience, and feedback.
- Translate Governance decisions into technical requirements, owners, milestones, evidence, exception conditions, and closure criteria.
- Test normal, denied, failure, bypass, fallback, stale-context, policy-unavailability, telemetry-loss, exception, and vendor-change scenarios.
- Verify operation through evidence; record findings and corrective actions.
- Convert successful controls, policy profiles, test cases, patterns, and templates into shared enterprise assets.

Implement outputs

- Architecture and authority-path maps.
- PDP/PEP and interaction coverage maps.
- Control design/configuration records.
- Tests, evidence, findings, exceptions, and verified closure.
- Reusable reference architectures and shared-control candidates.

11.5 Govern

Govern includes review, prioritization/reprioritization, risk/exception treatment, lifecycle decisions, assurance, accountability, and investment allocation.

Govern activities

- Review portfolio changes, Tier 3/4 applications, material authority paths, vendor-introduced capabilities, open findings, exceptions, incidents, posture trends, and shared-control gaps.
- Decide whether to authorize, constrain, defer, prohibit, restrict, suspend, retire, or reassess AI applications and components.
- Evaluate effective authority, CIEM, cloud permission, toxic-combination, drift, and segregation-of-duties evidence.
- Accept or reject residual risk through accountable decision rights.
- Prioritize shared-control and application-specific roadmap work.
- Verify closure and trigger reassessment as conditions change.

Govern outputs

- Decisions and governance events.
- Approved conditions, exceptions, risk treatment, and lifecycle actions.
- Prioritized integrated roadmap.
- Assurance record, portfolio metrics, and board/executive reporting inputs.

11.6 Integrated evidence, exceptions, and roadmap

Maintain one integrated roadmap. Do not separate "security roadmap," "AI roadmap," "identity roadmap," and "governance roadmap" so completely that cross-pillar work cannot be prioritized together.

Each roadmap item should identify affected AI application(s) and shared control, pillars, attention tier, severity/urgency, residual risk, target maturity, owner, dependency, due date, evidence of completion, and reassessment trigger.

Exceptions should identify affected authority, control gap, business rationale, compensating controls, residual risk, acceptance authority, monitoring, expiry/review, remediation, and closure conditions.

11.7 Role and cadence model

Role	Core contribution
Executive sponsor	Sets mandate, resolves cross-functional barriers, and sponsors portfolio investment.
Business owner	Owns value sought, use conditions, and business accountability for an AI application.

Role	Core contribution
Technical owner	Owns application architecture, integration, operation, and component change management.
IAM/IGA/PAM owner	Owns identity, entitlement, credential, lifecycle, and policy-management capability.
Data owner	Owns data classification, purpose, quality, lifecycle, and Data-operation requirements.
Service/API/workflow owner	Owns Service PEP, action constraints, workflow/approval, and effect boundaries.
SecOps owner	Owns detection, response, posture, telemetry, and security feedback.
Governance authority	Owns decision rights, risk appetite, exceptions, assurance, and portfolio direction.
Risk acceptance authority	Accepts or rejects documented residual risk within delegated authority.
Enterprise architecture	Maintains cross-pillar coherence, patterns, shared-control roadmap, and technology alignment.

Use an operational cadence for material applications and open findings, plus a portfolio cadence for prioritization, shared-control investment, vendor changes, and assurance.

11.8 Companion Authorization-Risk Exercise interface

The Companion Authorization-Risk Exercise is recommended for Tier 3/4 or other high-consequence authority paths. A documented routing decision should state whether it is used and why.

Inputs can include AI-application inventory, business purpose, identity/authority map, Data/Services operation map, PDP/PEP coverage, evidence, exceptions, findings, control maturity, and residual-risk conditions. The exercise returns authority-path findings, assumptions, recommended constraints, and risk-treatment inputs to the same Governance and roadmap process.

11.9 Appendix D integration

Appendix D provides the instruments that operationalize the lifecycle:

- Discover: Worksheets D.2.x.
- Assess: Worksheets D.3.x.
- Implement: Worksheets D.4.x–D.6.x.
- Govern and improve: Worksheets D.7.x–D.8.x.

12. Operating and Improving IAAI as an Enterprise System

12.1 Shared-control leverage

The enterprise gains leverage when a shared improvement reduces risk across many AI applications. Typical shared-control investments include AI-application inventory, agent registry, identity lifecycle, workload identity, policy administration, Data PEP patterns, metadata quality, retrieval filtering, Service PEP patterns, API/MCP governance, telemetry correlation, exception workflow, and evidence/assurance automation.

Prioritize shared controls when they reduce concentrated authority, repeated findings, broad dependency risk, or recurring implementation burden.

12.2 Operating-model alignment

IAAI requires business, technology, data, identity, security, service, governance, legal/privacy, procurement, and risk functions to work through explicit decision rights and handoffs. The goal is not to centralize every task. It is to ensure that purpose, authority, control, evidence, and accountability remain connected.

12.3 Learning from operation

Use incidents, posture findings, access reviews, red-team results, exceptions, vendor changes, control failures, successful patterns, architecture changes, and lifecycle events as learning inputs. Convert recurring lessons into policy profiles, reference architectures, reusable tests, templates, implementation accelerators, vendor-assessment questions, and roadmap priorities.

12.4 Maturity as a continuous, domain-specific practice

Maturity is not a one-time score. An enterprise may have advanced real-time policy decisions in one high-risk Service boundary, intermediate controls in a data platform, basic standing-entitlement patterns in older applications, and strong compensating SecOps controls while preventative coverage is being improved. The portfolio approach allows risk-prioritized improvement without implying that every AI application must achieve the same maturity at the same time.

12.5 Adapting to change

Reassess when AI applications change, vendors introduce capabilities, models or agents gain new functions, new data/service connections appear, iden-

tity/permission patterns drift, incidents occur, standards mature, regulatory obligations change, or evidence shows that prior assumptions were incomplete.

The intended outcome is not perfect centralization, universal advanced maturity, or the elimination of all risk. It is a more governable enterprise: one that can identify authority, constrain it at meaningful boundaries, observe and investigate its use, make residual risk accountable, and improve its portfolio through evidence and learning.

Appendix A — Operational Glossary and Cross-Reference Guide

A.1 Purpose and terminology distinction

This appendix provides an operational glossary, lifecycle map, task/cross-reference guide, and the **Canonical IAAI Security Interoperability Use-Case Catalog**.

An AI application is the primary portfolio unit. An IAAI security interoperability use case is a fixed canonical assertion about recurring interactions among pillars. The catalog helps enterprises recognize and assess common interoperability conditions within AI applications; it is not itself the AI-application portfolio.

A.2 Operational glossary

Term	Definition	Do not confuse with
AI application	Bounded enterprise deployment/use of AI for a business purpose, including relevant agents, models, data, services, lifecycle, and governance.	A fixed canonical interoperability use case.
Authority path	Material identity-to-data/service-action chain inside an AI application.	A complete AI application.
IAAI security interoperability use case	Fixed canonical assertion about a recurring cross-pillar security interaction.	Portfolio record or product category.
Standing entitlement	Persistent role, group, ACL, grant, relationship, or cloud permission.	Real-time policy decision.

Term	Definition	Do not confuse with
Vaulted credential	Retained privileged credential/secret controlled through vault/PAM/secrets practices.	Runtime token or PDP decision.
Real-time policy decision	Current-context permit/deny determination.	Possession of a token or credential.
PDP	Policy decision point.	PEP, which applies a decision at a boundary.
Data PEP	Enforcement point for a data operation.	Service PEP.
Service PEP	Enforcement point for a service operation.	Data PEP.
Governance decision	Human deliberation or accountable subjective determination.	Governance event.
Governance event	Processed record/action that implements or records a governance-driven change.	The decision itself.
Portfolio attention tier	Current depth of review/evidence/testing/governance needed.	Control maturity or residual risk.
Control maturity	Current/target strength of a control domain.	Attention tier.
Effective authority	Practical authority from combined identities, policies, data, services, delegation, and paths.	One permission considered alone.

A.3 Lifecycle map

`Discover → Assess → Implement → Govern → Discover`

Discover includes inventory and initial segmentation. Assess includes attention, baseline/rebaseline, evidence, and routing. Implement includes deepening, design, control implementation, testing, and validation. Govern includes review, prioritization/reprioritization, exceptions, risk treatment, assurance, and lifecycle decisions. Learning informs every phase.

A.4 Task and cross-reference guide

Task	Start here	Then use
Discover/inventory AI applications	Sections 2, 4, 11	Worksheets D.2.1–D.2.3
Assess vendor-introduced capability	Sections 2.3, 9.3, 11	Worksheets D.2.1, D.3.1, D.5.6
Assign attention tier	Section 2.5	Worksheets D.3.1–D.3.3
Map identity/delegation	Section 5	Pattern B.2.3; Worksheet D.5.1
Identify Data PEP	Section 6	Pattern B.2.1; Worksheet D.5.2
Identify Service PEP	Section 7	Pattern B.2.2; Worksheet D.5.3
Conduct interaction/seam testing	Section 10	Appendix B; Worksheet D.5.4
Establish SecOps evidence	Section 8	Pattern B.2.6; Worksheet D.5.5
Treat authorization posture/CIEM risk	Sections 5.11, 9.6	Pattern B.2.9; Worksheets D.5.6, D.6.1
Manage exception	Sections 9.7, 11.6	Worksheet D.6.2
Route Companion Exercise	Sections 11.3, 11.8	Worksheet D.3.3; Worksheet D.8.1

A.5 Canonical IAAI Security Interoperability Use-Case Catalog

The catalog identifies common fixed interoperability assertions. Control depth is described as basic, intermediate, or advanced preventative control where appropriate, or compensating/governance control. The three authorization primitives provide baseline depth: standing entitlement/basic, vaulted credential/intermediate, and real-time policy decision/advanced. Enhanced standing-entitlement patterns may receive increased depth credit; vaulted credentials do not become advanced merely through vault treatment.

A.5.1 Authorization and identity management

ID	Interoperability use case	Participating pillars	Control classification
UC-01	Human identity is authenticated and receives governed AI-application access through enterprise identity context.	A, E	Basic Preventative Control

ID	Interoperability use case	Participating pillars	Control classification
UC-02	Personal-agent identity acts for a sponsoring human within constrained derived authority.	A, E	Basic Preventative Control
UC-03	Enterprise-agent identity receives purpose-bound ownership, lifecycle, and authority.	A, E	Intermediate Preventative Control
UC-04	Workload/service identity authenticates a non-human runtime at a governed boundary.	A, C, D	Intermediate Preventative Control
UC-05	Subject/actor delegation context is preserved via an authorization enforcement mechanism.	A, B, C, D	Advanced Preventative Control
UC-06	Standing entitlement authorizes a governed data or service operation via an authorization enforcement mechanism.	A, B or C	Basic Preventative Control
UC-07	Access request workflow updates standing entitlement for data operation.	A, B, E	Intermediate Preventative Control
UC-08	Vaulted credential is controlled, rotated, released, and audited for privileged operation.	A, C, D, E	Intermediate Preventative Control
UC-09	Real-time policy decision permits or denies a requested governed operation at a relevant enforcement point.	A, B or C	Advanced Preventative Control
UC-10	Identity/entitlement/ lifecycle evidence supports authorization posture, SecOps analysis, and Governance treatment.	A, D, E	Compensating/ Governance Control

A.5.2 Data operations

ID	Interoperability use case	Participating pillars	Control classification
UC-11	Data operation applies source authorization before governed data reaches AI application context.	A, B	Basic to Advanced Preventative Control
UC-12	Classification, purpose, ownership, and metadata constrain a data operation.	A, B, E	Intermediate Preventative Control
UC-13	Semantic, vector, graph, search, or RAG retrieval applies entitlement and metadata filters before AI context assembly.	A, B	Advanced Preventative Control
UC-14	Data platform applies existing policy-based masking, filtering, suppression, or abstraction in response to an authorized read.	A, B	Intermediate Preventative Control
UC-15	Derived artifact is governed for lineage, classification, retention, and downstream use.	B, C, E	Intermediate Preventative Control
UC-16	Missing/stale metadata triggers restriction, remediation, compensating treatment, or governance decision.	B, D, E	Compensating/ Governance Control
UC-17	Data-operation evidence supports SecOps correlation and Governance assurance.	B, D, E	Compensating/ Governance Control

A.5.3 Services operations

ID	Interoperability use case	Participating pillars	Control classification
UC-18	Service operation authorizes tool/API method invocation through a Service PEP.	A, C	Advanced Preventative Control

ID	Interoperability use case	Participating pillars	Control classification
UC-19	Service operation constrains method, parameter, object, recipient, destination, threshold, workflow state, or rate.	A, C	Advanced Preventative Control
UC-20	MCP interaction applies enterprise identity, policy, audit, and lifecycle controls to a governed capability.	A, B and/or C, D	Advanced Preventative Control
UC-21	Agent-requested data transformation, aggregation, correlation, derivation, or write is authorized as a service operation.	A, B, C	Advanced Preventative Control
UC-22	Workflow or autonomous action applies approval, segregation-of-duties, threshold, state, and escalation controls.	A, C, E	Advanced Preventative Control
UC-23	Communication, transaction, commitment, or external effect is constrained by identity-aware service policy.	A, C, E	Advanced Preventative Control
UC-24	Administrative or infrastructure action applies workload identity, privileged-access, service-policy, and evidence controls.	A, C, D, E	Advanced Preventative Control
UC-25	Service-operation evidence supports SecOps investigation and Governance assurance.	C, D, E	Compensating/ Governance Control

A.5.4 SecOps and governance

ID	Interoperability use case	Participating pillars	Control classification
UC-26	AI guardrail constrains an otherwise-authorized risky data or service operation.	B and/or C, D	Compensating Control

ID	Interoperability use case	Participating pillars	Control classification
UC-27	AI-aware DLP detects or constrains sensitive data movement in AI interaction.	B, C, D	Compensating Control
UC-28	Operational evidence correlates identity, policy, data, service, MCP, and outcome context.	A, B, C, D	Compensating/ Governance Control
UC-29	SIEM/SOAR detects, investigates, and coordinates bounded response to AI-related risk.	A, B, C, D, E	Compensating/ Governance Control
UC-30	Posture analysis identifies AI-reachable data, effective permission, configuration, drift, or exposure risk.	A, B, C, D, E	Compensating/ Governance Control
UC-31	Red teaming/adversarial testing produces findings for architecture improvement.	A, B, C, D, E	Compensating/ Governance Control
UC-32	Governance/control interaction translates purpose, conditions, and risk treatment into operating controls and evidence.	A, B, C, D, E	Compensating/ Governance Control
UC-33	Authorization posture and CIEM analysis inform effective-authority and segregation-of-duties risk treatment.	A, D, E	Compensating/ Governance Control
UC-34	Vendor-introduced AI application is discovered, lifecycle-governed, and routed for enterprise authorization-risk management.	A, B, C, D, E	Compensating/ Governance Control
UC-35	AI-assisted security operation is governed as an AI application with bounded authority and evidence.	A, B, C, D, E	Compensating/ Governance Control

Appendix B — Selected Interoperability Patterns

B.1 How to Read These Patterns

B.1.1 Participating Pillars, Interaction Context, and Scope

Each pattern describes a fixed interoperability context involving one or more pillars. Pattern titles and participating-pillar descriptions are neutral: they identify responsibilities involved in a recurring assertion without encoding a one-way flow label. The narrative may explain common operating sequences, dependencies, evidence pathways, and feedback where useful.

Use these patterns for design, assessment, testing, material change, incident/posture response, and assurance. They do not replace detailed pillar chapters or become a second complete assessment method.

Each pattern includes participating pillars, primary concern, questions, tests/evidence ideas, failure/fallback/bypass considerations, and representative implementation notes.

B.2 Interoperability Patterns

B.2.1 Authorization and Data-Operation Enforcement

Participating pillars

- Pillar A — Authorization: identity, entitlement, delegation, policy, PDP, lifecycle, and policy-decision context.
- Pillar B — Data: Data PEP enforcement at repository, query, search, vector, graph, API, abstraction, or transformation boundary.
- Pillar E — Governance: ownership, policy direction, exception, and evidence expectations where applicable.

Primary concern. Material data operations should receive usable identity, policy, metadata, purpose, resource, operation, and scope context before governed data enters AI application context.

Questions. Where is the Data PEP? Which source permissions, metadata, classifications, and policies govern it? How does identity/delegation arrive? Does platform-configured abstraction occur automatically for a read, or is a net-new transformation requested as a Service operation?

Tests and evidence. Compare source permissions with index/vector/graph permissions; test user/agent lifecycle change; test stale ACL propagation, missing metadata, denied retrieval, policy unavailability, alternate retrieval path, and evidence of filters/policy version.

Failure/fallback considerations. Do not silently broaden retrieval when policy/metadata/index state is unavailable. Define safe restriction, remediation, compensating controls, and Governance treatment.

B.2.2 Authorization and Service-Operation Enforcement

Participating pillars

- Pillar A — Authorization: authentication, delegation, policy administration, PDP context, and decision.
- Pillar C — Services: Service PEP enforcement for tools, APIs, MCP operations, workflows, communications, transactions, and infrastructure actions.
- Pillar E — Governance: authority boundary, approval, exception, risk treatment, and accountability where applicable.

Primary concern. Service authorization must constrain actual action, not merely connection. Relevant constraints can include method, parameter, object, recipient, destination, threshold, workflow state, rate, time, purpose, approval, and external effect.

Tests and evidence. Attempt disallowed parameter, recipient, threshold, workflow, rate, or destination changes; test policy unavailability and manual override; verify decision, enforcement, approval, and resulting-effect evidence.

B.2.3 Identity Propagation, Delegation, and Workload Identity

Participating pillars

- Pillar A — Authorization: subject/actor propagation, token exchange, workload identity, credential context, lifecycle, and PDP inputs.
- Pillar B — Data: Data PEP consumption of delegation/workload context.
- Pillar C — Services: Service PEP consumption of delegation/workload context.
- Pillar D — SecOps: correlation and anomalous-delegation detection.
- Pillar E — Governance: owner, purpose, lifecycle, and authority accountability.

Primary concern. The enterprise should reconstruct who requested, who acted, which workload/service executed, what authority was delegated, what scope/audience applied, and what outcome occurred.

Tests and evidence. Test actor loss, impersonation, audience mismatch, stale token, delegation beyond sponsor scope, cross-boundary path, workload replacement, token replay controls, and policy/evidence correlation.

B.2.4 Data and Service-Operation Layering

Participating pillars

- Pillar B — Data: source, scope, filtering, platform-configured abstraction, transformation, lineage, and derived artifact.

- Pillar C — Services: requested transformation, delivery, tool invocation, recipient/destination, workflow, communication, transaction, and other consequence.
- Pillar A — Authorization: identity, policy, delegation, and decision context.
- Pillar E — Governance: risk treatment where data and service consequence combine.

Primary concern. Preserve both questions when both are material: what data operation is permitted, and what service operation is permitted?

Tests and evidence. Test permitted retrieval with prohibited external delivery; permitted data-platform abstraction with prohibited agent-requested transformation; permitted summary with prohibited raw export; and evidence from both relevant PEPs.

B.2.5 Compensating Controls at Data and Service Boundaries

Participating pillars

- Pillar B — Data: Data PEP and governed data context.
- Pillar C — Services: Service PEP and consequential action context.
- Pillar D — SecOps: guardrails, DLP, runtime inspection, detection, and response.
- Pillar A — Authorization: identity and policy context that refines compensating controls.
- Pillar E — Governance: residual-risk, exception, deployment, and accountability conditions.

Primary concern. Guardrails and DLP should protect against incomplete or insufficient preventative coverage without obscuring missing authorization/PEP controls.

Tests and evidence. Simulate otherwise-authorized but unsafe prompt, export, destination, or tool action; verify block/alert, evidence, owner, corrective feedback, exception, and follow-up remediation.

B.2.6 Operational Evidence into SecOps

Participating pillars

- Pillar A — Authorization: authentication and authorization evidence source.
- Pillar B — Data: data-operation evidence source.
- Pillar C — Services: service-operation and MCP evidence source where applicable.
- Pillar D — SecOps: collection, correlation, detection, investigation, and response.

- Pillar E — Governance: consumer of material risk analytics and escalated evidence where applicable.

Primary concern. SecOps should correlate authenticated actor, delegation context, policy decision/version, data operation, service action, MCP event where relevant, guardrail/DLP event, and outcome.

Tests and evidence. Exercise missing correlation ID, telemetry loss, clock skew, schema mismatch, sampling, restricted-log access, event retention, and incident reconstruction.

B.2.7 Technical Security Feedback and Corrective Response

Participating pillars

- Pillar D — SecOps: finding, investigation, containment, response, and technical recommendation.
- Pillar A — Authorization: identity, entitlement, credential, delegation, lifecycle, or policy correction.
- Pillar B — Data: metadata, classification, source, filtering, index, lineage, or Data PEP correction.
- Pillar C — Services: tool, API, workflow, autonomy, Service PEP, parameter, recipient, or destination correction.
- Pillar E — Governance: risk treatment, exception, portfolio priority, assurance, and closure accountability.

Primary concern. A material finding should become owned, verifiable correction and accountable governance treatment.

Tests and evidence. Trace a finding from detection through containment, remediation, retest, evidence, Governance event, exception/acceptance where needed, and closure; test recurrence and exception expiry.

B.2.8 Governance/Control Interaction Patterns

Participating pillars

- Pillar E — Governance: purpose, decision rights, risk appetite, attention, conditions, exceptions, and assurance.
- Pillar A — Authorization: identity, lifecycle, entitlement, credential, and policy requirements.
- Pillar B — Data: classification, source, Data PEP, transformation, and retention requirements.
- Pillar C — Services: authority, approval, autonomy, destination, threshold, workflow, and external-effect requirements.
- Pillar D — SecOps: compensating-control, telemetry, detection, response, and feedback requirements.

Primary concern. Demonstrate that governance conditions become specific control requirements, implementation work, evidence, and closure criteria—and that evidence/finding returns inform ongoing governance.

Tests and evidence. Sample a governance decision and trace requirements, tickets/configuration, tests, evidence, exception record, review, and outcome. Test vendor feature activation/change for required governance interaction.

B.2.9 Authorization Posture, CIEM, and Governance Risk Treatment

Participating pillars

- Pillar A — Authorization: entitlement, credential, policy, lifecycle, and decision evidence.
- Pillar D — SecOps: CIEM, identity-security, cloud-permission, posture, and anomaly analysis.
- Pillar E — Governance: effective-authority risk treatment, segregation-of-duties compliance, exception, prioritization, and accountability.
- Pillar B — Data: high-risk data reach and metadata context where applicable.
- Pillar C — Services: high-consequence service reach and action context where applicable.

Primary concern. Govern effective authority, excessive privilege, drift, toxic combinations, separation-of-duties concerns, local-policy blind spots, and exception accumulation—not merely intended roles.

Tests and evidence. Analyze dormant identity, inherited privilege, stale service account, agent ownership gap, excessive cloud permission, high-risk data/service combination, separation-of-duties conflict, unresolved exception, and remediation closure.

B.3 Applying the Patterns

Apply patterns during architecture/design, material change, implementation, test, incident/posture response, assurance, and lifecycle review. Use Appendix D to record the relevant AI application, authority path, evidence, tests, findings, exception, owner, roadmap item, and reassessment trigger.

Appendix C — Representative Technologies, Standards, and Interoperability Examples

C.1 Architecture before products

This appendix is architecture-before-products and category-first. It identifies representative standards, methods, open-source projects, vendor APIs, commer-

cial products, and frameworks that can contribute to IAAI. It is not a market ranking, procurement shortlist, endorsement, or assertion that a named technology fully implements a pillar.

Validate actual current capability, deployment model, policy semantics, integration depth, data/service coverage, operational fit, security properties, support model, and contractual commitments before architectural dependency.

C.2 Source-type distinctions

Type	Meaning
Formal protocol or standard	Defined interoperability behavior maintained by recognized standards body or open community.
Emerging specification or extension	Important but maturity/implementation support requires validation.
Policy language or authorization method	Expresses/evaluates policy; not necessarily an interchange standard.
Open-source project or implementation	Implements a protocol, language, or capability.
Vendor API or implementation	Provider-specific interface/product capability.
Architecture pattern	Reusable design method implemented through multiple technologies.
Governance/control framework or regulation	Directs risk, evidence, and accountability; not an interoperability protocol.

C.3 Capability-category anchors

The IAAI Product Capability Category Taxonomy provides a durable capability structure for this appendix and the related Interoperability Graph. Group anchors do not imply exclusive pillar assignment; a product may contribute to several categories and pillars.

C.3.1 Authorization categories

- Identity Governance and Administration Platforms
- Enterprise Agent Identity Management
- Identity Provider and Conditional Access Platforms
- Privileged Access Management, Vault, and Secrets Platforms
- Workload Identity and Service Authentication Platforms
- Policy-as-Code Authorization Engines
- Relationship Authorization and Authorization Decision Platforms
- Authorization Administration and Policy Observability Platforms
- Zero-Trust Access Proxies for AI Services

C.3.2 Data categories

- Enterprise Data Platforms, Warehouses, Lakes, and Lakehouses
- Governed Data Products, Semantic Layers, and Data Abstraction Services
- Data Classification and Sensitivity Labeling Platforms
- Document Repositories and Enterprise Content Management Systems
- Enterprise Search and Content Indexing Services
- Vector Search and Semantic Retrieval Services
- Data APIs, Data Mesh Gateways, and Data-Service Access Controls
- Graph Data, Knowledge Graph, and Relationship Traversal Platforms
- Data Lineage, Catalog, and Metadata Management Platforms
- Data Transformation, Derived-Artifact, and AI Data Pipeline Controls

C.3.3 Services categories

- AI Orchestration Frameworks with Policy Hooks
- AI Agent Runtimes and Service-Operation Platforms
- MCP Tool and Data Server Implementations
- Enterprise-Managed MCP Authorization Components
- API Gateways, Service Gateways, and API Management Platforms
- Service Meshes and Workload Identity Access Controls
- Workflow, Business Process, and Autonomous Action Platforms
- Enterprise Application and SaaS Service-Control Platforms
- Communication, Delivery, and External-Effect Platforms
- Transaction, Procurement, Payment, and Commitment Platforms
- Cloud, Infrastructure, and Administrative Control Planes
- Enterprise Browsers with AI-Aware Controls

C.3.4 SecOps categories

- AI Guardrail and Safety Frameworks
- AI-Aware Data Loss Prevention Platforms
- AI Threat, Model, and Agent Runtime Security Platforms
- AI Red-Teaming and Adversarial Testing Tools/Services
- AI Agent, Tool, and Service Activity Observability Platforms
- SIEM Platforms
- SOAR Platforms
- Cloud Audit Logging and Infrastructure Telemetry Platforms
- Data Security Posture Management Platforms
- Identity Security, Cloud Permission, and CIEM Analysis Platforms
- Endpoint, XDR, and Telemetry Platforms with AI Activity Visibility
- API, Application, Integration, and Network Security Platforms
- Cloud, Workload, Container, and Runtime Security Platforms
- Model, Pipeline, Artifact, and AI Supply-Chain Security Platforms

C.3.5 Governance categories

- Enterprise AI Governance and Risk Management Platforms
- GRC, Enterprise Risk, and Policy Management Platforms with AI Modules
- Identity Governance, Authorization Posture, and CIEM Governance Capabilities
- AI Regulatory, Privacy, Compliance, and Assurance Platforms
- Responsible AI, Model Risk, and Impact Assessment Toolkits
- AI Inventory, Registry, and Lifecycle Governance Capabilities

C.4 Identity, authentication, delegation, and lifecycle

Standard or method	Type and steward	Maturity	IAAI role	Implementation consideration
OAuth 2.0	Formal IETF framework	Established	Delegated authorization framework and scoped access context	Define audiences, scopes, token lifecycle, resource-server enforcement, and client security.
OpenID Connect	OpenID Foundation specification	Established	Federated authentication and identity assertions	Define issuer trust, claims, assurance, session, and relying-party validation.
OAuth 2.0 Token Exchange, RFC 8693	Formal IETF standard	Established	Subject/actor propagation across agents and services	Preserve audience, scope, actor/subject semantics, and audit correlation.
DPoP, RFC 9449	Formal IETF standard	Established but implementation-dependent	Sender-constrains OAuth access tokens to reduce bearer-token replay risk	Requires compatible authorization/resource servers, client key management, nonce/replay handling, and deliberate token-boundary design.

Standard or method	Type and steward	Maturity	IAAI role	Implementation consideration
RAR, RFC 9396	Formal IETF standard	Established but implementation-dependent	Carries transaction-specific authorization details	Value depends on consistent semantics and resource-server enforcement.
AuthZEN	Emerging OpenID Foundation specification family	Emerging/ watch	Authorization interaction and decision API patterns that may aid PDP/PEP interoperability	Validate profiles, semantics, latency, governance, and production support.
OpenID4VP	OpenID Foundation specification	Emerging/ use-case-specific	Verifiable presentations for identity, role, qualification, organization, or compliance claims	Define issuer trust, revocation, privacy, lifecycle, and justification for complexity.
SD-JWT	IETF emerging standards-track method	Emerging/ use-case-specific	Selective disclosure of identity attributes	Define issuer/verifier trust, disclosure minimization, claim semantics, holder binding where needed, and status/revocation.
SCIM 2.0	Formal IETF standard	Established	Provisioning/ deprovisioning and lifecycle propagation	Validate schemas, connector coverage, deprovisioning semantics, and agent extensions where applicable.
SPIFFE/ SPIRE	Open standard and OSS implementation	Established/ use-case-dependent	Workload identity and service authentication	Integrate attestation, trust domains, rotation, registration, and policy inputs.

C.5 Authorization, Data, Services, SecOps, and Governance examples

Authorization

Representative capabilities include IGA, IdP/conditional access, PAM/vault/secrets, workload identity, policy-as-code, relationship authorization, authorization observability, and zero-trust access proxies. Examples may include commercial IGA/identity providers, cloud IAM services, policy engines such as OPA/Rego, Cedar, Cerbos, relationship-authorization systems, SPIRE, and enterprise PAM/vault offerings. Named examples should be assessed for actual deployment and integration coverage.[1][2][3]

Data

Representative capabilities include data platforms, governed data products, semantic layers, classification/labeling, repositories, search/indexing, vector retrieval, graph systems, data APIs/gateways, lineage/catalog, and transformation/derived-artifact controls. Data platforms may supply native Data PEPs, but enterprises should evaluate alignment with Authorization policy and whether semantic/vector/index paths preserve required metadata and entitlements.[4][5]

Services

Representative capabilities include agent orchestration frameworks, agent runtimes, MCP servers and authorization components, API gateways, service meshes, workflow engines, enterprise applications, communication platforms, transaction systems, cloud control planes, and enterprise browsers. The relevant question is not whether the product is "an AI product," but whether it exposes or constrains a material service-operation boundary.[6][7]

SecOps

Representative capabilities include guardrails, DLP, AI threat/runtime security, red teaming, activity observability, SIEM/SOAR, cloud audit logging, DSPM, CIEM/identity security, endpoint/XDR, API/application security, workload/runtime security, and supply-chain security. These capabilities are complementary; guardrail/DLP coverage should not be mistaken for full preventative authorization.[8][9]

Governance

Representative capabilities include AI governance/risk-management platforms, GRC/policy systems, authorization-posture/CIEM governance capabilities, privacy/compliance/assurance systems, responsible-AI/model-risk methods, and AI inventory/registry/lifecycle systems. Governance may be implemented

across multiple systems and operating processes rather than a single product category.[10]

C.6 Interoperability foundations

Area	Examples	Role
Policy decision and enforcement	PDP/PEP, OPA/Rego, Cedar, XACML, relationship authorization	Express and apply policy at Data and Services boundaries.
Agent/tool/ service interoperability	MCP, enterprise-managed MCP authorization, OpenAPI, API gateways, service mesh	Expose and govern controlled capabilities; MCP can involve Data, Services, and SecOps evidence.
Data operations	Native data controls, semantic layers, vector filters, graph policy, data APIs	Data PEP and metadata-aware enforcement.
Telemetry	OpenTelemetry, cloud audit, SIEM interfaces	Cross-pillar correlation and assurance.
Governance	NIST AI RMF, NIST SP 800-207, ISO/IEC 42001, applicable obligations	Risk, control, assurance, and accountability framing.

Appendix D — Portfolio and AI Application Assessment Instruments, Evidence Templates, and Roadmap Worksheets

D.1 Start Here: Choose Your Path

If you need to...	Use this path	Primary outputs
Establish visibility	Pass 1 — Portfolio visibility	AI-application portfolio, owners, lifecycle/activation states, unknowns, preliminary attention
Evaluate baseline and route work	Pass 2 — Baseline and routing	Fast diagnostic, current/target profile, routing, gaps, reassessment triggers

If you need to...	Use this path	Primary outputs
Assess material authority	Pass 3 — Material AI-application deep dive	Authority/data/service/PEP/ seam/evidence/governance profile, findings, roadmap
Analyze high-consequence authority path	Companion path	Documented routing decision and Companion Authorization-Risk Exercise inputs/outputs

D.2 Pass 1 — Portfolio Visibility

D.2.1 AI-application inventory record

Field	Record
AI-application ID and name	
Business purpose and value sought	
Sponsor; business owner; technical owner	
Vendor/supplier; homegrown/ vendor-introduced/ embedded/composite	
Lifecycle/activation state	Discovered—not explicitly enabled/authorized / Available—pending evaluation / Pilot / Enabled or operational / Restricted-suspended-disabled / Retired
Agents, models, runtime, and platforms	
Main data sources and highest known classification	
Tools, APIs, MCP contexts, workflows, and external effects	
Identity/authority pattern	
Preliminary attention tier or provisional/unclassified status	
Known concerns, gaps, and unknowns	
Vendor evidence, configuration, contract/change considerations	

Field	Record
Shared-control dependencies	
Next review date/trigger	

D.2.2 Linked component and dependency record

Component/dependency	Type	Owner	Lifecycle/change state	Relationship to AI application	Relevant pillar(s)
	Agent/model/data source/tool/API/MCP/workflow/shared control/vendor				

D.2.3 Discovery and ownership gap log

Unknown or gap	AI application/component	Risk/attention implication	Owner	Next evidence/action	Due/review date

D.3 Pass 2 — Baseline and Routing

D.3.1 Quick Reference: Attention, maturity, severity, risk, exception, priority

Term	Decision rule
Attention tier	Determines review/test/governance depth now.
Control maturity	Current/target strength by specific domain.
Finding severity/urgency	Impact and time sensitivity of a gap/failure.
Residual risk	Risk remaining after controls.
Exception status	Approved temporary departure from target control with owner/conditions.
Roadmap priority	Relative investment sequence and shared-control leverage.

D.3.2 Fast Diagnostic

Question	Yes	Partial	No	Unknown	Evidence/ next step
Is the AI application in the portfolio with business/technical ownership?					
Is lifecycle/activation state known, including vendor-introduced state where applicable?					
Are participating identities/agents/workloads owned and lifecycle-managed?					
Are data sources, classifications, and Data PEPs known for material operations?					
Are tools/services/external effects and Service PEPs known for material operations?					
Are Delegation, entitlement, credential, and policy boundaries identified?					
Are SecOps controls/evidence sources and correlation needs known?					
Are Governance conditions, exception status, and reassessment triggers known?					
Has vendor evidence/ configuration/change exposure been assessed where applicable?					

Why this matters: a fast diagnostic does not prove safety. It identifies whether an AI application can remain at monitored baseline, requires targeted work, should receive a material deep dive, or needs restriction/deferment pending evidence.

Next step: Use D.3.3 for maturity and routing; then use D.4–D.6 for material deep dive.

D.3.3 Current maturity → target maturity and routing

Domain	Current: basic/ intermediate/ advanced	Target	Gap	Evidence/ test	Owner
Authorization					
Data					
Services					
SecOps					
Governance					

Routing decision	Record
Attention tier/provisional status and rationale	
Continue monitored baseline	
Conduct targeted pillar/interaction assessment	
Conduct material AI-application deep dive	
Restrict, defer, suspend, or retire	
Route Companion Authorization-Risk Exercise; rationale whether used or not	
Reassessment trigger/date	

D.4 Pass 3 — Identity, Data, and Services Deep Dive

D.4.1 Identity and authority map

Element	Description/evidence
Subject(s)	
Actor(s), agent(s), workload(s), service(s)	
Purpose, owner, lifecycle, certification, retirement	
Delegation, token, credential, and trust path	
Standing entitlements and enhanced controls	

Element	Description/evidence
Vaulted credentials	
Real-time policy decisions	
PDPs, policy boundaries, versions, exceptions, availability/fallback	

D.4.2 Data-operation map

Data source/ product/ index	Classification/ metadata	Read/ automatic abstraction/ transform	Data PEP	Filters/ lineage/ derived artifact	Evidence

D.4.3 Service-operation map

Service/ tool/API/ MCP/ workflow	Operation/ effect	Service PEP	Method/ parameter/ recipient/ threshold constraints	Approval/ autonomy	Evidence

D.4.4 PEP and interaction coverage map

Authority path	Participating pillars	PDP	Data PEP	Service PEP	Evidence chain	Failure/ bypass test	Owner

D.5 Pass 3 — SecOps, Governance, and Evidence Deep Dive

D.5.1 SecOps and evidence profile

Control/ evidence domain	Current coverage	Gap	Test/ evidence	Feedback destination
Guardrails				
DLP				

Control/ evidence domain	Current coverage	Gap	Test/ evidence	Feedback destination
Telemetry/ correlation				
Detection/SIEM				
SOAR/response				
Posture/DSPM/ CIEM				
Red team/resilience				
Vendor evidence/change monitoring				

D.5.2 Governance, lifecycle, and assurance profile

Topic	Record
Business purpose/value and ownership	
Vendor-introduced/activation/lifecycle state	
Attention tier and rationale	
Required governance decisions and conditions	
Governance events/evidence	
Effective authority/CIEM/segregation-of-duties concerns	
Residual risk	
Exception status and expiry	
Assurance/review cadence	
Material-change and reassessment triggers	
Suspension/retirement conditions	

D.5.3 Test-depth and failure-scenario record

Scenario	Expected behavior	Participating pillars	Evidence	Failure/ fallback/ compensating response	Owner/ closure
Policy unavailable					
Stale identity/ delegation					
Stale/missing metadata/ index					
Bypass/ alternate path					
Telemetry loss					
Vendor capability/ configuration change					
Exception expiry					

D.6 Findings and Exceptions

D.6.1 Finding record

Field	Record
Finding ID and date	
Affected AI application(s)/shared control	
Pillar(s)/interaction context(s)	
Severity/urgency	
Authority/consequence affected	
Evidence	
Root cause/hypothesis	
Immediate containment	
Corrective action	
Owner and due date	
Verification/closure evidence	

D.6.2 Exception record

Field	Record
Exception ID and affected control	
AI application/authority path	
Business rationale	
Authority/consequence affected	
Compensating controls	
Residual risk	
Risk-acceptance authority	
Start, expiry, and review dates	
Remediation plan	
Monitoring/evidence	
Closure criteria	

D.7 Roadmap and Portfolio Learning

D.7.1 Integrated roadmap

Roadmap item	Application-specific/ shared	Pillar(s)	Attention urgency	Severity/	Target ma-tu-rity	Owner	Dependency	Decision date	Completion evi-dence

D.7.2 Portfolio learning log

Pattern, incident, finding, or vendor change	Reusable lesson	Control/ template/ test update	Affected portfolio segment	Owner	Adoption evidence

D.8 Worked Journey: Procurement Agent

Pass 1 — Portfolio visibility

A procurement AI application retrieves supplier and contract information, drafts purchase requests, routes approvals, and can submit requests to a procurement system. The initial record captures supplier data, contract records, procurement APIs, workflow, external supplier communication, employee subject, agent actor, workload identity, vendor dependencies, and potential financial commitment. It is provisionally routed to Tier 3 because it can create material business action.

Pass 2 — Baseline and routing

The Fast Diagnostic finds named ownership and data classification but identifies broad shared API authority, incomplete vector-index filtering for supplier artifacts, and incomplete evidence about an available vendor-provided autonomous workflow feature. The application is routed to a material deep dive. A shared portfolio gap is identified: delegation-aware Service PEP patterns for procurement APIs.

Pass 3 — Material deep dive

The team maps employee subject, agent actor, workload identity, supplier/contract Data PEPs, procurement Service PEP, approval thresholds, recipient restrictions, DLP/guardrails, telemetry, exception treatment, and failure scenarios. The roadmap adds query-time filtering, parameter-aware procurement API authorization, human approval above threshold, correlation IDs, vendor-feature lifecycle review, and policy-unavailability tests.

Portfolio learning

The enterprise converts the resulting approach into a reusable procurement-agent reference architecture, delegated-authority policy profile, Service PEP test suite, vendor-introduced-feature evaluation checklist, and shared roadmap item.

End Notes and Sources

1. IETF, “OAuth 2.0 Authorization Framework”.
2. OpenID Foundation, “OpenID Connect Core 1.0”.
3. IETF, “RFC 8693: OAuth 2.0 Token Exchange”.
4. IETF, “RFC 9449: OAuth 2.0 Demonstrating Proof of Possession (DPoP)”.
5. IETF, “RFC 9396: OAuth 2.0 Rich Authorization Requests”.
6. Model Context Protocol, “Model Context Protocol Specification”.
7. OpenTelemetry, “OpenTelemetry Documentation”.
8. National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0)”, 2023.
9. National Institute of Standards and Technology, “Zero Trust Architecture (SP 800-207)”, 2020.

10. International Organization for Standardization, “ISO/IEC 42001: Artificial intelligence — Management system”.
11. IETF, “RFC 7644: System for Cross-domain Identity Management: Protocol”, 2015.
12. SPIFFE, “SPIFFE Overview”.
13. OpenID Foundation, “AuthZEN Working Group”.
14. OpenID Foundation, “OpenID for Verifiable Presentations”.
15. IETF, “Selective Disclosure for JWTs (SD-JWT)”.
16. IAAI Product Capability Category Taxonomy, “Version 6 Reference Taxonomy for the Enterprise Workbook and Interoperability Graph,” approved working taxonomy, August 20, 2026.

Index

Concepts

- AI application — Executive Summary; Sections 1.3, 2.2–2.4, 9.3, 11.2; Worksheet D.2.1
- Authority path — Sections 1.3, 2.4, 11.3, 11.8; Worksheet D.4.4
- Authorization posture — Sections 5.11, 9.6; Pattern B.2.9
- Companion Authorization-Risk Exercise — Sections 2.4, 4.6, 11.8; Worksheet D.3.3
- Control maturity — Sections 2.6, 11.3; Worksheet D.3.3
- Data PEP — Sections 3.5, 6.4; Pattern B.2.1; Worksheet D.4.2
- Delegation — Sections 3.5, 5.8; Pattern B.2.3; Worksheet D.4.1
- Effective authority — Sections 2.4, 9.6; Pattern B.2.9
- Governance decision — Sections 3.5, 5.5, 9.4
- Governance event — Sections 3.5, 5.5, 9.4
- IAAI security interoperability use case — Sections 1.3; Appendix A
- MCP — Sections 3.1, 5.7, 7.5, 8.9; Appendix C
- Portfolio attention tier — Section 2.5; Worksheet D.3.3
- Real-time policy decision — Sections 2.6, 3.5, 5.6; Appendix A
- Service PEP — Sections 3.5, 7.4; Pattern B.2.2; Worksheet D.4.3
- Standing entitlement — Sections 2.6, 3.5, 5.6; Appendix A
- Vendor-introduced AI application — Sections 2.3, 4, 9.3, 11.2; Worksheet D.2.1
- Vaulted credential — Sections 2.6, 3.5, 5.6; Appendix A
- Workload identity — Sections 3.5, 5.8; Pattern B.2.3

Tasks

- Assess vendor-introduced capability — Sections 2.3, 9.3, 11.2; Worksheets D.2.1, D.3.2
- Assign attention tier — Section 2.5; Worksheet D.3.3
- Conduct interaction testing — Section 10; Appendix B; Worksheet D.5.3
- Create AI-application inventory — Sections 4.4, 11.2; Worksheet D.2.1
- Identify Data PEP — Section 6.4; Worksheet D.4.2
- Identify Service PEP — Section 7.4; Worksheet D.4.3
- Initiate companion exercise — Sections 11.3, 11.8; Worksheet D.3.3
- Map delegation — Sections 5.8, 11.4; Worksheet D.4.1
- Reassess after change — Sections 9.8, 11.3; Worksheet D.5.2
- Verify closure — Sections 11.4–11.6; Worksheet D.6.1

Roles

- Business owner — Sections 2, 9, 11.7; Worksheet D.2.1
- Data owner — Sections 6, 11.7; Worksheet D.4.2
- Governance authority — Section 9; Worksheet D.5.2
- IAM/IGA/PAM owner — Sections 5, 11.7; Worksheet D.4.1
- Risk acceptance authority — Sections 9.7, 11.5; Worksheet D.6.2
- SecOps owner — Sections 8, 11.7; Worksheet D.5.1
- Service/API/workflow owner — Sections 7, 11.7; Worksheet D.4.3
- Technical owner — Sections 2, 11.7; Worksheet D.2.1

Standards and methods

- AuthZEN — Appendix C.4
- DPoP — Appendix C.4
- MCP — Sections 7.5, 8.9; Appendix C.6
- OAuth 2.0 — Appendix C.4
- OAuth Token Exchange — Appendix C.4
- OpenID Connect — Appendix C.4
- OpenID4VP — Appendix C.4
- OpenTelemetry — Sections 8.4, 8.9; Appendix C.6
- RAR — Appendix C.4
- SCIM — Sections 5.3, 5.12; Appendix C.4
- SD-JWT — Appendix C.4
- SPIFFE/SPIRE — Sections 5.8, 5.12; Appendix C.4

Version History

Version	Date	Summary
6.0.1 Enriched Integrated Manuscript	August 2026	Restored detailed five-pillar architecture, portfolio-first lifecycle, AI-application terminology, vendor-introduced AI scope, rich authorization/data/services/SecOps/ governance detail, interaction patterns, capability taxonomy, standards treatment, assessment instruments, mini-cases, end notes, and index.
6.0.1 Initial Integrated Manuscript	August 2026	Established complete initial structural integration of front matter, main body, appendices, and back matter.