AI Agent Inventory

Bounding Agency: Identity, Authority, and Trust in the Age of AI Agents

The AI-IAM Practitioner Series · Book 1 · Rahul Bhardwaj · Arvik Press

Tool, version and date	AI Agent Inventory v1.0, September 2026
Purpose	Answers the first question: what is acting in the environment, what can it reach and do, who owns it, and under what authority?
Source in the book	Chapter 4 (Table 4.2, worked records); minimum fields in Chapter 17, Table 17.2; authoritative sources in Chapter 12, Table 12.2.
When to use it	Always: it is the input to every other tool. Its scope follows the IAM threshold of Chapter 1, so it includes assistants acting with a user's delegated access, not only autonomous agents.
What it produces	One record per agent or use case, with evidence and confidence.
Scope	Fill in: which agents, use cases, environments and business units this record covers.
Completed by / reviewed by	Fill in: who completed it, who reviewed it, and when.
How to use this workbook	Cells shaded pale yellow are for you to fill in; drop-down lists appear where the book defines the options. Example rows are shaded blue-grey and written in italics: replace or delete them. Reference sheets reproduce the book's tables.
Evidence conventions	For every evaluative entry record the evidence source and date. Evidence strength, strongest first: tested control or log; configuration; documented policy; interview. Use N/A only with a reason. Record Unknown as a finding, never as zero.
Design note	There is deliberately no 'delegation depth' field. How far redelegation should go is an open design question (Chapter 7).
Already have a CMDB, IGA, GRC or AI inventory?	Add the missing fields there rather than running a parallel spreadsheet.
Status	**A working diagnostic. Not a certification, audit opinion or compliance attestation.**
Licence	© 2026 Rahul Bhardwaj. Published by Arvik Press. You may copy and adapt the tools for use within your own organization. They may not be resold or republished.
Updates	Check arvikpress.com/bounding-agency for newer versions and corrections.
Examples	Inventory rows 3–4 are CrystalSoft examples from the book (Table 4.2). CrystalSoft Energy and its people, agents and systems are fictional, created to illustrate the book. Values marked with an asterisk (*) are illustrative.

	IDENTITY				OWNERSHIP			CAPABILITY				AUTHORITY		CREDENTIALS AND LIFECYCLE			EVIDENCE		
Agent or use case	**Principal and identity type**	**Platform and runtime**	**Environment**	**Version and provenance**	**Business sponsor**	**Technical owner or operator**	**Degree of agency**	**Tools and systems**	**Data sensitivity**	**Read, write and action capability**	**Reversibility**	**Authority mode and source**	**Delegation capability (which agents it may invoke or delegate to)**	**Human-approval mode**	**Credential and authentication pattern**	**Lifecycle state**	**Risk notes**	**Evidence source and date**	**Confirmed or inferred**	**Confidence**
Billing Resolution Agent: resolve customer billing complaints	*Workload identity (its own principal)*	*CrystalSoft's cloud agent platform*	*Production*	*Version and deployment record kept*; model lineage not tracked*	*Head of Customer Operations*	*Platform engineering*	*High: interprets the goal and selects steps*	*CRM, third-party billing platform, ticketing, knowledge base, email drafting*	*Customer PII, billing data*	*Reads billing history; writes CRM records and tickets; requests credits*	*Credits reversible with effort; sent emails are not*	*Both: own identity for platform operations; on Priya's behalf for customer-specific actions*	*Can request credits from the Payments Agent only*	*Priya approves customer emails; credits above $500 need human approval*	*Short-lived, task-scoped**	*Active*	*High (customer data, financial action); review quarterly**	*Platform registry + sponsor attestation*	*Confirmed**	*High*
Payments Agent: post approved credits and refunds	*Its own workload identity (separate principal)*	*Finance platform**	*Production*	*Version record kept**	*Finance**	*Finance platform team*	*Lower: executes credit requests within policy**	*Billing platform payment functions**	*Payment data**	*Posts credits and refunds*	*Credits reversible with effort*	*Own authority, acting on requests from approved agents**	*None**	*Credits above $500 need human approval*	*Short-lived**	*Active*	*High (financial action); review quarterly**	*Finance platform registry*	*Inferred**	*Medium**

IDENTITY				OWNERSHIP			CAPABILITY				AUTHORITY			CREDENTIALS AND LIFECYCLE			EVIDENCE		
Agent or use case	Principal and identity type	Platform and runtime	Environment	Version and provenance	Business sponsor	Technical owner or operator	Degree of agency	Tools and systems	Data sensitivity	Read, write and action capability	Reversibility	Authority mode and source	Delegation capability (which agents it may invoke or delegate to)	Human-approval mode	Credential and authentication pattern	Lifecycle state	Risk notes	Evidence source and date	Confirmed or inferred	Confidence

Table 17.2 — AI Agent Inventory (v1.0): Minimum Fields

Group	Fields
Identity	Agent or use case; principal and identity type; platform and runtime; environment; version and provenance
Ownership	Business sponsor; technical owner or operator
Capability	Degree of agency; tools and systems; data sensitivity; read, write and action capability; reversibility
Authority	Authority mode and source; delegation capability (which agents it may invoke or delegate to); human-approval mode
Credentials and lifecycle	Credential and authentication pattern; lifecycle state
Evidence	Risk notes; evidence source and date; confirmed or inferred; confidence

Table 12.2 — Chapter 4 Inventory Fields and Their Authoritative Sources

Inventory field	Likely authoritative source	Owner
Principal / identity type	Identity provider or workload identity platform	IAM
Name / purpose	Use-case approval record	Business sponsor, via AI governance
Sponsor / operator	IGA record (sponsor); platform registry (operator)	Business; platform
Environment / status	Agent platform and deployment pipeline	Platform
Degree of agency	Design record and use-case approval	AI engineering, confirmed by AI governance
Authority mode	Authority approval and grant configuration	IAM, approved by the sponsor
Delegation capability	Grant and policy configuration	IAM; owner of the receiving authority
Tools / data	Agent configuration; data catalog	AI engineering; data governance
Read / write	Entitlements and grant configuration	IAM
Reversibility	Business process assessment	Business sponsor
Human approval	Policy configuration	IAM implements; sponsor and risk decide
Credential pattern	Token service, secrets manager, PAM	IAM
Provenance / version	Deployment pipeline and model registry	Platform; AI engineering
Risk / review	GRC or AI governance record	Risk; AI governance

Chapter 12: build a federated inventory, one record per agent assembled from fields whose authoritative source is explicit.