PRACTITIONER TOOLKIT

# Practitioner Questions

*Bounding Agency: Identity, Authority, and Trust in the Age of AI Agents — every chapter's questions in one place*

Each chapter of the book ends with questions to take back into your organization. They are collected here so you can use them in a workshop, an assessment kickoff or a design review. Answer them with evidence, and record what you do not know: an Unknown is a finding.

## Chapter 1 · From Human Identity to Agent Identity

1. Which AI systems in your environment can do more than generate information? Can they read, invoke or change something in an enterprise system?
2. When they act, do they use their own identity, a workload identity or a user's delegated access?
3. Who sponsors or owns each of them, and would you know if that person changed role or left?
4. Could your team tell a conventional automation from an agent that selects its own actions toward a goal?

## Chapter 2 · Identity, Agency and Authority

1. When an AI-enabled system acts in your environment, can you identify the actual actor, not only the user who initiated the request?
2. How much discretion does that actor have to choose or initiate intermediate actions?
3. Are you treating its technical permissions as if they automatically define what it is legitimately authorized to do?
4. When the system acts on behalf of someone else, do your records preserve both the actor and the subject?

## Chapter 3 · Why Today's IAM Is Necessary—but Not Sufficient

1. Which IAM capabilities in your environment already apply to agents today?
2. Which of your controls assume a stable identity, static context, simple delegation or a discrete review point?
3. Where are you using AI to improve IAM, and where are you governing AI itself as an enterprise actor?

4. Which of your perceived gaps are genuinely new, and which are adaptations of controls you already have?

## Chapter 4 · Establishing Agent Identity

1. Can you name the principal that represents each production agent?
2. Can you distinguish the enduring agent or deployment identity from the short-lived credentials it uses?
3. Who is the accountable sponsor for each agent, and is that different from the technical operator?
4. Do you know whether each agent can write or change state, whether its actions are reversible and when human approval is required?
5. How confident are you that your inventory sees embedded, SaaS-provided and platform-created agents?

## Chapter 5 · Authenticating the Agentic Enterprise

1. How do your production agents authenticate today: shared secret, certificate, workload identity, token or delegated user session?
2. Are their credentials short-lived and bound to the intended workload or client where your architecture supports it?
3. When one agent calls another, can each side verify the actual caller rather than only a shared platform?
4. Which security signals can cause trust in a principal or session to be re-evaluated, and how fast do they reach relying systems?
5. Can your teams explain what authentication proves, and what it does not prove about authority?

## Chapter 6 · Authorization in a World of Autonomous Action

1. Which of your agents' actions are consequential enough to need action-level rather than resource-level authorization?
2. What context can your policy engine trust and enforce today, and which inputs come from the agent itself?
3. Where would you use short-lived or just-in-time grants, step-up or human approval?
4. If authority must be withdrawn, how long could existing tokens or downstream grants remain usable?
5. Where can one agent both plan and execute a high-impact transaction without an independent checkpoint?

## Chapter 7 · Delegated Authority and the Chain of Trust

1. Who originally granted the authority this agent is exercising?
2. Can each downstream actor receive less authority than its upstream delegator?
3. Can the resource distinguish the actual actor from the principal on whose behalf it acts?
4. How many times may authority be redelegated, and who decides?
5. If the original grant is revoked, which downstream credentials remain usable, and what evidence survives to reconstruct the chain?

## Chapter 8 · Privilege, Lifecycle and Access Governance

1. Who sponsors and who operates each agent today, and what happens when either changes?
2. Which events trigger an access review besides a calendar date?
3. Can you find dormant or ownerless agents and stale credentials, including integration tokens?
4. Which agents can obtain privileged access, and is that access time- or task-bounded?
5. When an agent is retired, can you prove its third-party credentials and integrations went with it?

## Chapter 9 · Architecture and the Technology Ecosystem

1. Which existing identity, governance, credential, policy and observability components can your agent program reuse today?
2. Where must rich agent context be translated before it reaches legacy systems, and what will the translator log?
3. Which protocols in your design provide interoperability, and which actually provide identity or security controls?
4. Where is agent registration authoritative, and where is discovery only advertised metadata?
5. Which emerging components are you treating as production dependencies, and what is your plan if their specifications change?

## Chapter 10 · Identity Risk and the Agentic Threat Surface

1. If this agent identity were compromised, what could it do, not merely access?
2. What authority could it pass downstream, and would that authority stay bounded?
3. Which of its actions are irreversible or high-consequence, and which controls sit outside the agent layer?
4. Which signals would tell you that behavior is outside the approved task or authority?
5. What can you revoke immediately, and what might remain valid downstream?
6. Can you distinguish an unapproved agent from an approved agent behaving unexpectedly?

## Chapter 11 · Governance, Regulation and Accountability

1. Who is accountable for this agent use case, and who operates it technically?
2. Who approves the authority the agent can exercise, and material changes to that authority?
3. Can you reconstruct actor, subject, authority source, policy decision, action and outcome, including across third-party boundaries?
4. Which actions need additional oversight because of consequence, reversibility or regulation, and is that oversight working?
5. Which obligations actually apply to this use case and jurisdiction, and who confirmed it?
6. What would leadership need to see to understand exposure without a misleading aggregate score?

## Chapter 12 · IAM Cannot Do This Alone

1. Who decides whether an agent use case may enter production?
2. Which source is authoritative for identity, owner, authority, model and runtime, tools, data and risk fields?
3. Who approves material changes to an agent's authority or delegation, including delegation to another team's agent?
4. Who receives runtime signals, and who can actually contain access?
5. Who decides an agent is retired, and who proves the decommissioning?
6. Where are there handoffs today with no explicit decision owner?

## Chapter 13 · The Agentic IAM Operating Model

1. Can every agent and use case be in all six capabilities?
2. Where is agent identity authoritative, and who keeps it current?
3. Where are authentication and access decisions enforced for agents?
4. Where is delegated authority traced and bounded, and where does that stop?
5. Which runtime identity and authority signals can trigger an intervention today?
6. Who governs authority, evidence and accountability, and which enterprise interfaces must work for the model to operate?

## Chapter 14 · The Agentic IAM 360° Assessment

1. What evidence supports each assessment conclusion, and how strong is it?
2. Where is the state truly unknown rather than weak?
3. Which capabilities are adequate for the specific use cases in scope?

4. Where does an existing control need adaptation rather than replacement?
5. Which BUILD items are genuine capability gaps rather than immature implementation?
6. Which critical gaps are hidden by otherwise strong capability coverage?

## Chapter 15 · From Assessment to Roadmap

1. Which gaps materially increase risk or block a priority use case?
2. What must exist before this initiative can succeed?
3. Which current IAM controls should be adapted rather than replaced?
4. Which high-risk use cases need THEN capabilities sooner, or should wait until they exist?
5. Which emerging technologies should be piloted rather than standardized enterprise-wide?
6. What decisions does leadership need to make now, and what evidence will show each initiative worked?

## Chapter 16 · Your First 90 Days

1. What can I confirm in the first week, and what remains unknown?
2. Which agents or use cases combine the most authority, data sensitivity, tool access and consequence?
3. Who owns each use case, and who can approve or change its authority?
4. Which existing IAM controls can I reuse immediately?
5. Where is the organization blind because evidence or ownership is missing?
6. What decisions must leadership make before the next 30 days, and what artifacts will exist by Day 90 that did not exist on Day 1?

## Chapter 17 · The Agentic IAM Practitioner Toolkit

1. Which tool answers the question I have right now?
2. What evidence do I need before I complete this field or rating?
3. Is this item Unknown, N/A, adequate, adaptable or genuinely missing?
4. Which outputs should feed another tool rather than be re-entered by hand?
5. Which tools must be repeated when an agent's authority or use case changes?
6. What should I version and retain as program or audit evidence?

*A working diagnostic. Not a certification, audit opinion or compliance attestation.*

*© 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.*