

Request for Proposal: Non-Human and AI Agent Identity

Management Platform

Template from RFPhub.com. Replace everything in square brackets, adjust the Must-have/Nice-to-have priorities, and delete this note and the instructions section before you issue the RFP.

Instructions for the buyer (delete before issuing)

A non-human identity evaluation involves identity, cloud, security operations and engineering teams, and every shortlisted platform will need read access to some of your most sensitive systems. Agree on scope and connector permissions before vendors get involved. The example schedule in this template runs about 12 weeks from kickoff to contract award; stretch it if the proof of value covers many cloud accounts or business units.

1. **Define scope and the sources to connect (weeks 1-2).** List the cloud accounts, SaaS tenants, identity provider, code repositories, CI/CD platforms, Kubernetes clusters, secrets managers and AI agent platforms in the first phase. Record the numbers vendors will price against. If you don't know how many non-human identities you have, give vendors the counts you do know (accounts, tenants, repositories, clusters) and ask them to price on those as well, so a large discovery result doesn't change the price later.
2. **Assemble the buying team.** Include identity and access management, cloud platform, security operations, DevOps or platform engineering, and the teams building AI agents, plus procurement, legal and privacy. Agree early who will act on findings. If remediation will fall to the engineering teams that own the identities, give them a voice in the evaluation.
3. **Decide whether you need an RFI (weeks 2-4).** Products in this category come from different starting points, such as secrets management, privileged access management, identity governance, cloud entitlement management and dedicated non-human identity platforms. They differ in which sources they cover and whether they fix problems or only report them. If you're

unsure which approach fits, whether AI agent governance belongs in this purchase, or your long list has more than five vendors, issue a short RFI first. Ask which sources each vendor covers, what it can remediate, which AI agent platforms it supports today, and how it prices.

4. **Tailor and issue the RFP (weeks 4-7).** Delete requirements that don't apply, adjust priorities, and list your sources and environment details in Section 1. Give vendors three weeks to respond, with a written question-and-answer window in the first week.
5. **Score the responses (weeks 7-8).** Have each evaluator score independently against the rubric before comparing notes, and set aside any vendor that misses a Must-have. Check every 'yes' for whether the feature is generally available today or still in preview or on the roadmap, especially for AI agent features.
6. **Run a proof of value with two or three vendors (weeks 8-11).** Have your security team approve each vendor's connector permissions, then connect the shortlisted platforms read-only to the same set of sources. Compare the inventories side by side: which identities each one found, how accurate the owner suggestions are, and how many findings are real. Test one remediation end to end on a non-production credential, and call references of a similar size.
7. **Negotiate and award (weeks 11-12).** Pin down the license metric and what happens when discovery finds more identities than you estimated. Settle renewal caps, add-on pricing, data deletion and exit assistance before you sign, not after.

Contents

1. Introduction and Background
2. Project Objectives
3. Scope of Work
4. Requirements
5. Vendor Questions
6. Evaluation Criteria and Scoring
7. Pricing Response
8. Submission Guidelines
9. Timeline

1. Introduction and Background

[Company Name] is seeking proposals for a non-human identity management platform that discovers, assigns owners to and secures the service accounts, API keys, tokens, certificates, workload identities and AI agents across our cloud, SaaS, code and CI/CD environments. This RFP describes our requirements, the questions we need answered, and how proposals will be evaluated.

Organization background

- [Describe your organization: industry, size, locations]
- [Regulatory or contractual requirements that apply]
- [Cloud providers in use, with the number of accounts or subscriptions and Kubernetes clusters in scope]
- [SaaS applications, code repositories and CI/CD platforms in scope]
- [Identity provider, secrets managers, privileged access management and identity governance tools in use]
- [AI agent platforms, copilots and automation tools in use, and how agents authenticate today]
- [Known counts of service accounts, API keys and AI agents, if available]

2. Project Objectives

- Build a complete, continuously updated inventory of non-human identities and AI agents
- Assign an accountable owner to every non-human identity, and a human sponsor to every AI agent
- Reduce long-lived static credentials through rotation, expiry and, where possible, federated or short-lived credentials
- Remove unused and excessive permissions from service accounts, workloads and AI agents
- Detect and respond to misuse of non-human credentials
- Produce evidence for audits and access reviews of non-human accounts

3. Scope of Work

The selected vendor will be responsible for:

- Connection of the platform to [cloud accounts, SaaS tenants, identity provider, code repositories, CI/CD platforms, Kubernetes clusters, secrets managers and AI agent platforms]
- Initial discovery, baseline inventory and ownership assignment for in-scope identities
- Configuration of risk scoring, lifecycle policies and remediation workflows
- Integration with our identity provider, cloud IAM, secrets managers, ticketing system, SIEM and identity governance tools
- Training for administrators, security analysts and identity owners
- Ongoing support and service management for the contract term

4. Requirements

For each requirement, enter a response code and a short explanation. Y = available today in the proposed edition; P = partially met or needs configuration/services (explain); R = on the roadmap (give a date); N = not supported. Must-have requirements are pass/fail; Nice-to-have requirements add points.

4.1 Functional

ID	Requirement	Priority	Response	Vendor comments
F1	Continuous discovery of non-human identities across [cloud providers], SaaS applications, our identity provider, Kubernetes clusters, CI/CD platforms and on-premises directories, including service accounts, service principals, managed identities, IAM roles, access keys, API keys, OAuth applications and tokens, and personal access tokens.	Must		
F2	Detection of secrets (keys, tokens, passwords and connection strings) exposed in code repositories, CI/CD pipelines and configuration files, natively or by ingesting findings from our secrets scanning tool, with each secret linked, where possible, to the identity it belongs to.	Must		
F3	Inventory of certificates and keys used for machine authentication, collected directly or imported from our certificate lifecycle management tool or certificate authorities, with expiry dates and owners.	Nice		
F4	Discovery of AI agents built on or running in [AI agent platforms, copilots and automation tools in use], with the identities, API keys and OAuth grants each	Must		

ID	Requirement	Priority	Response	Vendor comments
	agent uses and the permissions they carry.			
F5	Inventory of the tools, APIs and Model Context Protocol (MCP) servers each AI agent is configured to call, and the access each one grants.	Nice		
F6	One record per identity showing its credentials, its permissions, the resources it can reach, the applications and workloads that use it, when it was created, when it was last used and who owns it.	Must		
F7	Ownership assignment for every identity, with suggested owners drawn from creation events, tags, code ownership and usage, and a workflow for owners to confirm or reassign. Every AI agent must have a named human sponsor who is accountable for its purpose and access.	Must		
F8	Detection of orphaned identities when an owner or sponsor leaves or changes role, based on events from our identity provider or HR system, with a reassignment workflow and an option to disable the identity if no one claims it.	Must		
F9	A risk score for each identity that shows the factors behind it, such as privilege level, credential age, exposure in code, missing owner, dormancy and unusual activity.	Must		
F10	Findings for common hygiene problems: credentials past their rotation date or with no expiry, dormant identities, credentials shared across applications or environments, and third-party OAuth applications with broad scopes.	Must		
F11	Remediation workflows that rotate, revoke, disable or delete a credential or identity, either executed through native integrations after approval or assigned to the owner as a ticket, with each finding tracked to closure.	Must		
F12	Automated rotation that also updates every consumer of the credential, or the secrets manager they read from, and confirms that consumers still authenticate after the change.	Nice		
F13	Comparison of granted and used permissions for cloud roles, service accounts and OAuth scopes, flagging unused, excessive and administrative permissions.	Must		

ID	Requirement	Priority	Response	Vendor comments
F14	Right-sized permission policies generated from observed usage, applied through cloud APIs or as infrastructure-as-code pull requests, with a preview of which recent actions the new policy would have blocked.	Nice		
F15	Detection of anomalous use of non-human credentials, such as use from unexpected networks or locations, unusual API calls or volumes, first-time access to sensitive resources and use of a credential that was due to be retired, with the evidence behind each alert.	Must		
F16	Automated response actions (disable an identity, revoke a token or credential, open an incident) triggered by detections, with optional approval steps, natively or through our SOAR platform.	Nice		
F17	Lifecycle policies from creation to decommission: required owner, purpose, environment and expiry at creation; maximum credential lifetime; dormancy thresholds; and a decommission workflow that removes the identity, its credentials and its permissions.	Must		
F18	Periodic reviews in which owners and sponsors confirm that each identity or agent is still needed and its access is still appropriate, run natively or through our identity governance platform, with escalation when owners don't respond.	Must		
F19	For AI agents that act on behalf of a user, a record of the delegating user, the scopes granted and the grant's expiry (for example, through OAuth 2.0 token exchange or on-behalf-of flows), and revocation of those grants when the user leaves.	Nice		
F20	Issuance or brokering of scoped, short-lived credentials for AI agents per task or session, so that agents do not hold standing secrets.	Nice		
F21	Policy enforcement for which tools, APIs and data each AI agent may access, denying unregistered tools and out-of-scope actions by default.	Nice		
F22	An audit trail of AI agent actions that links each action to the agent, its human sponsor or delegating user, the credential used and the resource accessed.	Nice		

4.2 Technical and architecture

ID	Requirement	Priority	Response	Vendor comments
T1	Agentless discovery through each source's APIs, using read-only permissions by default. List every permission the platform needs for each source, identify which are needed only for remediation, and describe how collection stays within each source's API rate limits.	Must		
T2	A list of supported sources (cloud providers, SaaS applications, identity providers, code repositories, CI/CD platforms, Kubernetes, secrets managers and AI agent platforms), stating for each whether the platform discovers identities, shows their activity, analyzes their permissions or can remediate them.	Must		
T3	Self-hosted collectors or a hybrid option for on-premises sources that cannot be reached from the vendor's cloud. State what data the collector sends to the platform.	Must		
T4	Support for [number] identities across [number] cloud accounts, SaaS tenants and repositories. State how often the inventory refreshes and how quickly new identities and activity appear after they are created.	Must		
T5	Retention of inventory history and activity data for at least [number] months, searchable in the platform and exportable.	Must		
T6	Recognition of federated and short-lived workload identities (for example, OIDC federation from CI/CD platforms, cloud workload identity federation and SPIFFE IDs) in the inventory, and reports on static credentials that could be replaced with them.	Nice		

4.3 Integration

ID	Requirement	Priority	Response	Vendor comments
I1	Integration with our identity provider: single sign-on for administrators via SAML 2.0 or OIDC, user and group provisioning via SCIM, and joiner, mover and leaver events used to keep ownership current.	Must		
I2	Integration with cloud IAM in [cloud providers] to read identities, policies and activity logs, and to apply approved remediation through native APIs.	Must		