

Deploying AI Agents **Safely** Is Only Half the Job

Every AI agent your team deploys issues credentials the moment it goes live: service accounts, API keys, tokens, all created to get the job done fast. The best practices below will help you deploy safely. But deployment is only the first half of the problem. What happens to that credential once the agent's job is finished is where most organizations lose the thread entirely.

Deployment Best Practices

Inventory every agent before it goes live

Know what's running, what it can access, and who owns it, before you turn it on, not after.

Scope permissions tightly at creation

Grant the agent only what its task requires. Broad access is easy to hand out and hard to claw back later.

Set a rotation and expiry policy up front

Every credential should have a defined shelf life before deployment, not one discovered during incidents.

Assign an explicit owner per credential

"The platform team" is not an owner. A named person or role is.

Log and monitor agent identity activity

Treat every agent like a non-human employee. If it can act, its actions need a trail.

What These Practices Don't Cover

Best practices for deployment protect the moment an agent goes live. They don't protect what happens after. Agents get replaced, projects wind down, teams move on, but the credentials those agents issued don't expire on their own. They sit in your environment with standing access, unowned and unmonitored, until something goes wrong.

What is an orphaned identity?

A credential still holding permissions after the system or agent it was issued to is no longer in use, and no one is responsible for it.

The Scale of the Gap

82:1

For every human identity in an enterprise, there are 82 non-human ones: agents, service accounts, and machine credentials, most without a governance program built for them.

\$4.45M

Average cost of a credential-related breach. None of it is recoverable through cyber insurance if the credential had no assigned owner.