A Day in the Life

Deploying *AI Agents*

Last time, Alex was dealing with secrets sprawl in GitLab pipelines. Now the company is deploying AI agents across the business — and the credential problem just got a lot more complicated.

WITHOUT AEMBIT ⚠ | **WITH AEMBIT** ✓

01 DEPLOYMENT DAY

Three sets of agents are live — with long-lived credentials.

Three AI agents are running across the business – handling research, generating reports, querying internal data. Because the devs used OAuth to let the user give the agent access, each one gets the full identity and access of the user that created them.

> ***To Alex, any action an agent takes looks the same as an action taken by a user.***

Agents go live with blended identities and managed access.

The same three agents – Claude, ChatGPT, the custom LLM – received blended identities from Aembit that are tied to the user and the agent itself. Access policies are defined before deployment. Credentials are issued at runtime, scoped to each specific task, and expire automatically when the session ends.

> ***"Set the policy once. Aembit handles the credential at runtime." Short-lived. Scoped. Gone when the job's done.***

02 THE AUDIT

Agents are accessing resources – but which agents? Which resources?

Compliance wants to know which systems the AI agents accessed over the last 90 days. Alex starts digging through server logs. The service accounts don't distinguish between agents and users. With shared credentials, there's no clear attribution.

> ***Alex's answer: "I can pull the server logs... it'll take a few days." The auditor notes the gap. The CISO is looped in.***

The compliance question with a clean answer.

Same question from compliance. Alex opens Aembit. Every access event is there — which agent, which resource, which user's session it was tied to, what policy authorized it, when it happened.

> ***"I have that right here." Alex doesn't have to piece activity together from shared credentials and logs.***

03 THE NEXT DEPLOYMENT

The business wants more. Security's scrambling to roll it out safely.

Another business unit wants to deploy a new agent. The CISO puts a hold on it until the current access model is documented and justified. That documentation doesn't exist yet. Alex is now doing retroactive paperwork for deployments that already happened.

> ***Approval timeline: 6+ weeks. Workflows are delayed and stakeholders are frustrated.***

The business wants more agents. Security says yes.

The new agent deployment comes up for approval. The CISO sees that no agents will be given persistent access, reviews the policy configuration, and approves in the same meeting.

> ***The rollout no longer stalls over credential and access questions. Alex moves on to the next thing***

THE DIFFERENCE

Identity and access management built for AI agents.

Aembit is the identity control plane for every agent — Claude, ChatGPT, Gemini, and custom LLMs.

- Just-in-time credentials
- Least-privilege access
- Runtime policy enforcement
- Full audit trail
- Works with all major agents
- Compliance-ready reporting

See how Aembit secures AI agent access at your organization.

Works with

Claude | ChatGPT | Gemini | Custom agents

Request a demo at **aembit.io**

aembit