Securing AI Agents on AWS

Rowan Udell · AWS Security Hero

The Problem

Your AI agent needs permissions to be useful.

But most teams give it **way too many**.

Architecture Overview

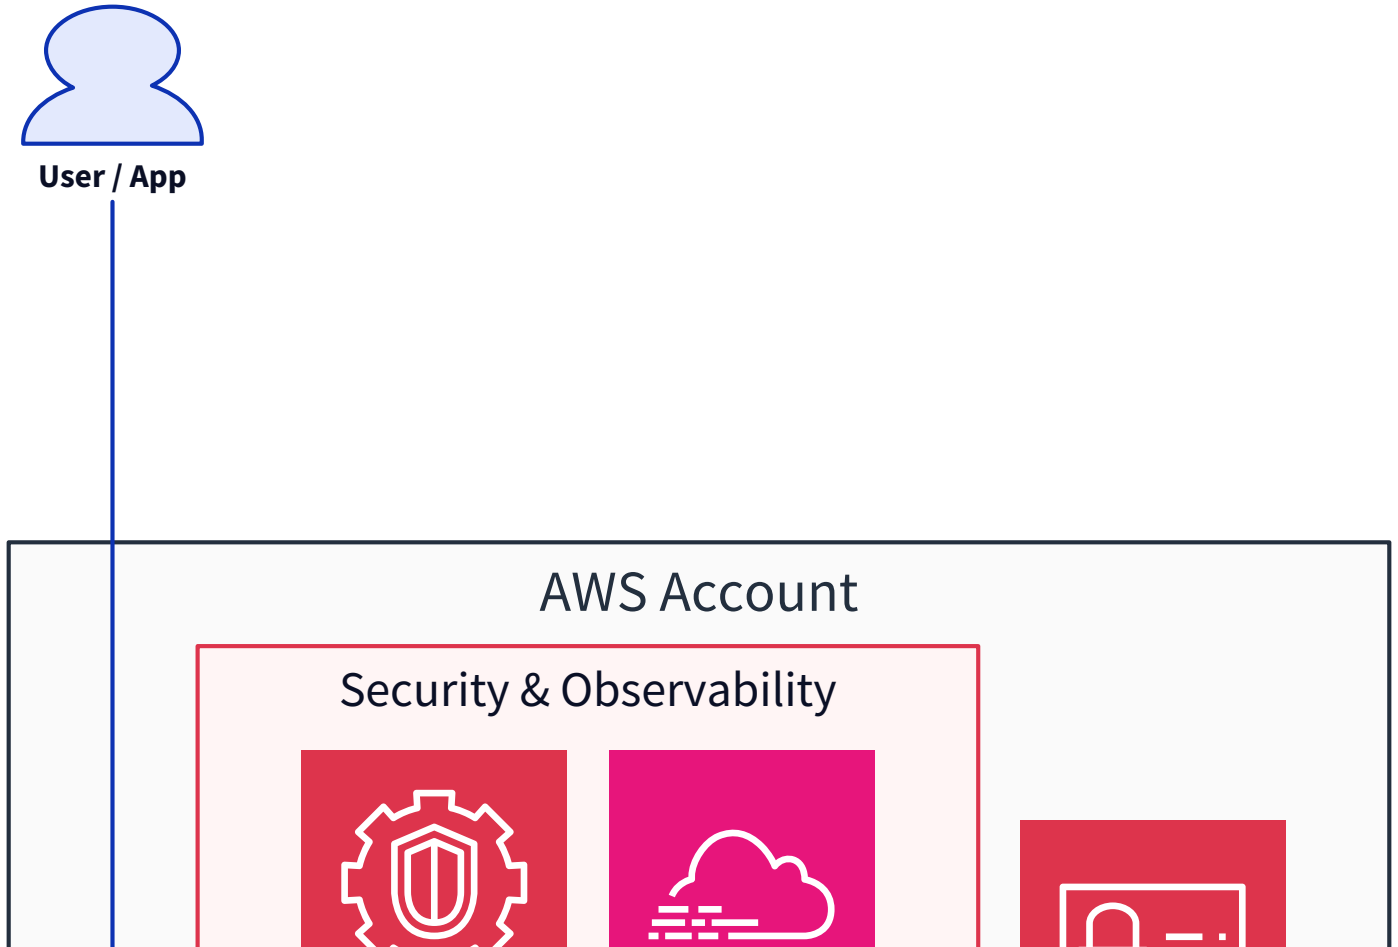

The Lethal Trifecta

Three risks that compound each other:

1. **Overprivileged IAM role** — agent can do more than it should
2. **No prompt injection defence** — attacker controls agent behaviour
3. **No observability** — you won't know when it goes wrong

IAM: Least Privilege for Agents

```
# What most teams do
aws iam attach-role-policy \
  --role-name agent-role \
  --policy-arn arn:aws:iam::aws:policy/AdministratorAccess   # 🚨

# What you should do: scope to exact actions + resources
aws iam create-policy --policy-document file://agent-policy.json
```

Key Takeaways

- Treat agent IAM roles like **production service accounts** — not developer sandboxes
- Log every `InvokeModel` call via CloudTrail
- Apply Bedrock Guardrails before the prompt reaches the model

Want the full checklist?

DM me if you're building agents on AWS and want an architecture review.

auditready.cloud