ETHEON

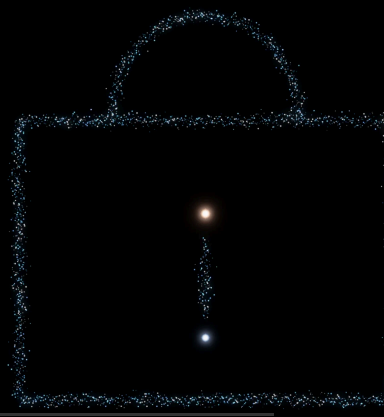

Agentic VM on AWS. A private team of AI agents in your own AWS account.

Deployment
about 8 weeks
$70,000

Run and improve
monthly, 60 hours
from $7,500

How it works. Stage by stage.

Weeks 1 and 2: we agree the team's roles, instructions, tools and limits for your work.

Weeks 3 to 6: we deploy the VM on Amazon EC2, serve the models in your account on Amazon EC2 or Amazon SageMaker AI, and set up each agent.

Week 7: we run the seven release gates and write the evaluation report.

Week 8: we release in stages, each with a rollback, and hand over the runbook.

What you receive. Stage by stage.

The agentic VM in your AWS account, with the agent team set up for your work

Architecture document

Release plan

Evaluation report

Release notes

Runbook

Handover docs

Seven release gates. Passed before every release.

Task quality | Grounding | Safety | Privacy | Latency | Cost | Drift

On AWS. The services we use.

Amazon EC2 for the virtual machine and the models, Amazon SageMaker AI where models are served there, Amazon S3 for files and logs.

Your data. Yours to keep.

We never train models on your data without your opt-in. Your data stays isolated from other clients. The accounts are yours.

Getting started. A reply within one business day.

Request a private offer on AWS Marketplace, or write to hello@etheon.com or call +1 (415) 413-3959. A partner reads every brief, not a form queue.

The work runs in your own AWS account. If you incur additional AWS infrastructure costs, separate from your AWS Marketplace transaction, you are responsible for paying those infrastructure charges.

© 2026 Etheon, Inc.