Pilots to Production: Becoming an Agentic Enterprise with AWS

Ganesh Hegde

Sr. Worldwide Specialist, GenAI

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. **Amazon Confidential and Trademark.**

The enterprise Agentic AI Journey

01. POCs and experiments

02. Initial pilots and production use cases

03. Multiple agentic use cases in production

04. Enterprise-wide agentic AI adoption

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Why AI pilots stall before production

The gap is not AI capability. It's model readiness and enterprise governance.

Model Performance

Generic models fail at enterprise tasks

- **Unreliable tool-calling**
  Agents pick wrong tools, miss steps
- **Domain hallucination**
  Models don't know your vocabulary, schemas, or business rules
- **Inconsistent reasoning**
  Same prompt, different results

Solved by: Model customization (SageMaker AI)

▶ Covered in next slide

Security & Identity

Agents lack enterprise-grade access control

- **No agent identity**
  Who is this agent? What can it access?
- **No access delegation**
  Agents call APIs, models, and data with no AuthN/AuthZ chain
- **Compliance blockers**
  CISO says no. Regulated industries require auditability before go-live

Solved by: Enterprise identity & governance layer

▶ Ecosystem partners

Governance & Observability

No runtime controls once agents are deployed

- **No prompt governance**
  LLM interactions are ungoverned
- **No visibility into agents**
  Can't trace decisions across multi-agent workflows
- **Cost & usage invisible**
  No quotas, no token tracking, no consumption controls

Solved by: Runtime policy + Bedrock Guardrails

▶ AWS native + ecosystem partners

IMPACT ON AI PRODUCTION READINESS

Slower POC to production • **Limited** AI consumption • **Delayed** production workloads

Solving this requires purpose-built models AND enterprise governance.

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

aws

Agentic capabilities: AWS + Ecosystem Partners

Business Applications: Virtual Assistants | Task Automation | Self-serve Support | Workflows

AI Governance

Responsible AI

Accountability

Security

Data & Privacy

Cost Management

Ecosystem Partners

+ Bedrock Guardrails

Applied across the platform

Agentic AI Platform

Agents

- Orchestrators
- Workflows
- Shared Agents
- Specialized Agents
- Hosting & communications

Tools

- MCP Servers & Gateways
- Shared Tools
- Specialized Tools

Agentic AI Foundations

- Foundation Models
- Evaluation
- Guardrails & Safety
- Observability
- Memory & Prompts
- Agent Registry
- **Agent Identity & Auth**
- KBs, Vector DBs, Data Lakes, APIs

POWERED BY

Bedrock & AgentCore
Runtime, orchestration, tools

SageMaker AI
Model customization & training

Who delivers what

AWS Native
Bedrock, AgentCore

SageMaker AI
Model Customization

Partners
Governance, Identity, Integration

AWS covers the AI platform & runtime.

Partners extend with governance and identity.

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

Build custom agents with Amazon

Open source

Protocols:
MCP, A2A

Frameworks:
CrewAI, LangGraph, Strands Agents

Amazon Bedrock

AgentCore

- Runtime
- Identity
- Gateway
- Code interpreter
- Memory
- Browser tool
- Observability

Model Capabilities

- Model access
- Cost and performance optimization
- Customization
- Guardrails

Knowledge Bases

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Your agents are only as good as the models behind them

Every enterprise has access to the same foundation models. Differentiation comes from customizing them for your domain.

CUSTOMIZE

Amazon SageMaker AI
Model Customization

Serverless Fine-Tuning
SFT, DPO, RLVR, RLAIF across 20+ open-weight models + Amazon Nova

Agent-Guided Workflow
Describe your use case in natural language. AI agent handles data prep through deployment.

SageMaker HyperPod
Continued pre-training at scale across 1000s of GPUs with elastic training

Choice of managed or semi-managed infrastructure

DEPLOY

Amazon Bedrock
Managed Inference + Guardrails

Custom Model Import
Train on SageMaker, deploy to Bedrock for serverless inference

Foundation Model Access
Nova, Claude, Llama, Mistral + your custom fine-tuned models

Bedrock Guardrails
Content safety, hallucination filters, responsible AI at the model layer

Serverless. Scalable. Governed.

ORCHESTRATE

Bedrock AgentCore
Production Runtime for Agents

Runtime | Memory | Identity | Gateway
Observability | Code Interpreter | Browser Tool

Open Frameworks & Protocols
CrewAI, LangGraph, Strands Agents
MCP & A2A protocols

Ecosystem Partner Extensions
API governance, enterprise identity, policy enforcement, integration layers

Extend with partners from AWS Marketplace

SageMaker AI (Fine-tune) → **Model Import** (Deploy to Bedrock) → **AgentCore** (Orchestrate) → **Guardrails** (Content safety) → **Enterprise Governance** (Identity, policy, controls) → **Production-Ready Agents**

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.

aws

The production-ready agentic stack

AWS provides the AI foundation. Ecosystem partners provide complementary enterprise governance.

Enterprise AI Governance Layer

What enterprises need above the AI foundation

Delivered by ecosystem partners

API Governance
Rate limiting, traffic control

Agent Identity
AuthN/AuthZ for agents

Policy Enforcement
Guardrails & compliance

Observability
Tracing & audit

Multi-region, multi-model
Data sovereignty & regional policy

AWS AI Foundation

SageMaker AI
Model Customization & Training

- Serverless fine-tuning (SFT, DPO, RLVR)
- 20+ models + Amazon Nova
- Agent-guided workflow
- HyperPod for large-scale CPT

Managed or semi-managed infrastructure

Amazon Bedrock
Managed Inference & Model Access

- Custom Model Import from SageMaker
- Foundation models (Nova, Claude, Llama)
- Bedrock Guardrails (content safety)
- Knowledge Bases & RAG

Bedrock AgentCore
Production Agentic Runtime

- Runtime, Gateway, Memory
- Agent Identity & Auth
- Observability & tracing
- Open: CrewAI, LangGraph, Strands, MCP

Framework-agnostic. Open protocols.

AWS builds the AI foundation. Ecosystem partners complete the enterprise governance layer.

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Start customizing. Start governing. Start scaling.

01

Customize your models

SageMaker AI serverless fine-tuning gets you from generic to production-grade in days.

aws.amazon.com/sagemaker/ai

02

Build with Bedrock + AgentCore

Deploy custom models to Bedrock. Orchestrate agents with AgentCore. Open frameworks, production runtime.

aws.amazon.com/bedrock

03

Visit the AWS booth

Discuss your agentic AI use cases, the right model customization path, and governance from our partners.

AWS Booth + Partner Ecosystem Lounge

GH

Ganesh Hegde

Sr. Worldwide Specialist, GenAI, AWS

Email
ganeheg@amazon.com

SageMaker AI
aws.amazon.com/sagemaker

Amazon Bedrock
aws.amazon.com/bedrock

© 2026, Amazon Web Services, Inc. or its affiliates. All rights reserved. Amazon Confidential and Trademark.