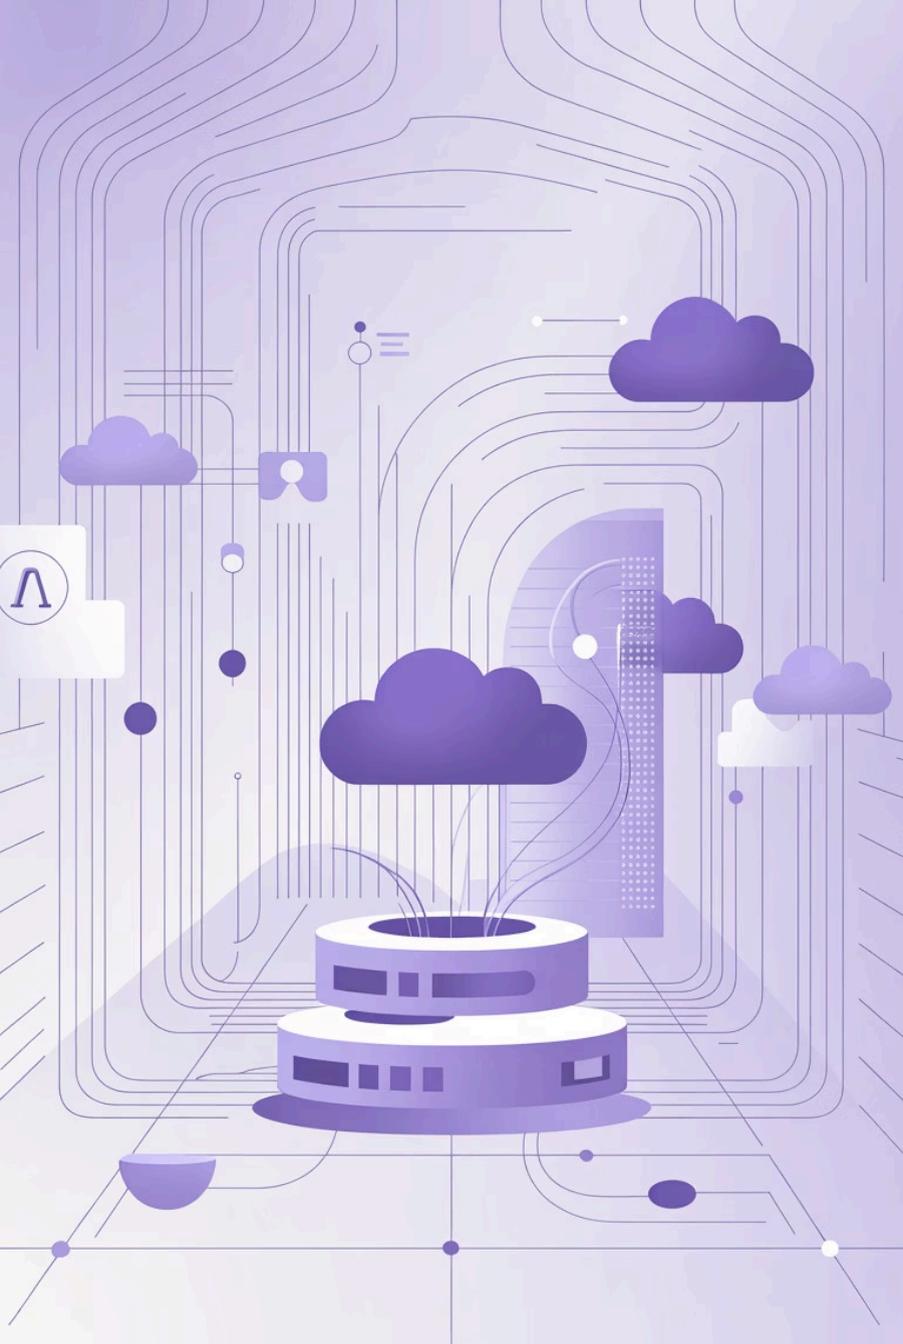


AWS AI - BUILDING AGENTIC AI APPLICATIONS ON AWS

TUTORIAL SERIES:
EXPERT LEARNING GUIDES





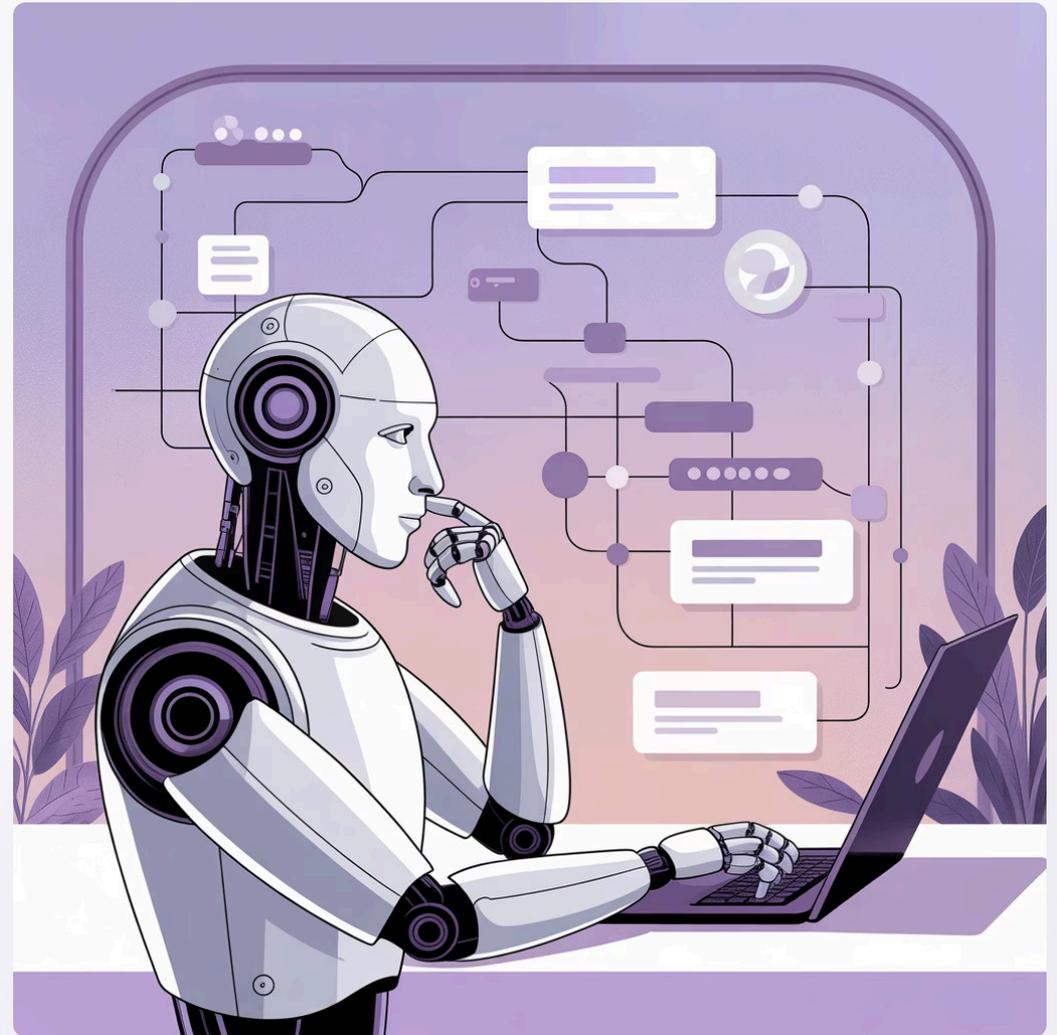
Best Practices for Building Agentic AI Applications on AWS

A comprehensive guide to designing, deploying, and scaling autonomous AI agents on Amazon Web Services

What is Agentic AI?

Autonomous AI agents represent a fundamental shift in artificial intelligence capabilities. These sophisticated systems perceive their environment, reason through complex scenarios, and take decisive action to achieve multifaceted goals with minimal human oversight.

Unlike traditional reactive assistants that simply respond to queries, agentic AI systems autonomously plan strategies, execute intricate multi-step workflows, and adapt dynamically to changing conditions and new information.



- ❏ **AWS Perspective:** Agentic AI represents the evolution from simple chatbots to digital teammates that meaningfully impact enterprise outcomes and transform business operations.



Why Build Agentic AI on AWS?



Scalable Infrastructure

AWS provides enterprise-grade, elastic infrastructure specifically optimised for demanding agentic AI workloads, ensuring performance at any scale.



Cutting-Edge Tools

Access Amazon Bedrock AgentCore, AWS Lambda, and Amazon ECS for serverless and containerised deployments that accelerate development.



Enterprise Ready

Built-in observability, security controls, and compliance frameworks ensure production readiness and regulatory adherence from day one.

Key AWS Tools & Frameworks for Agentic AI

Amazon Bedrock AgentCore

Build, deploy, and operate AI agents at scale without managing complex infrastructure. Fully managed service enabling rapid agent development with enterprise-grade reliability.

AWS Strands Agents SDK

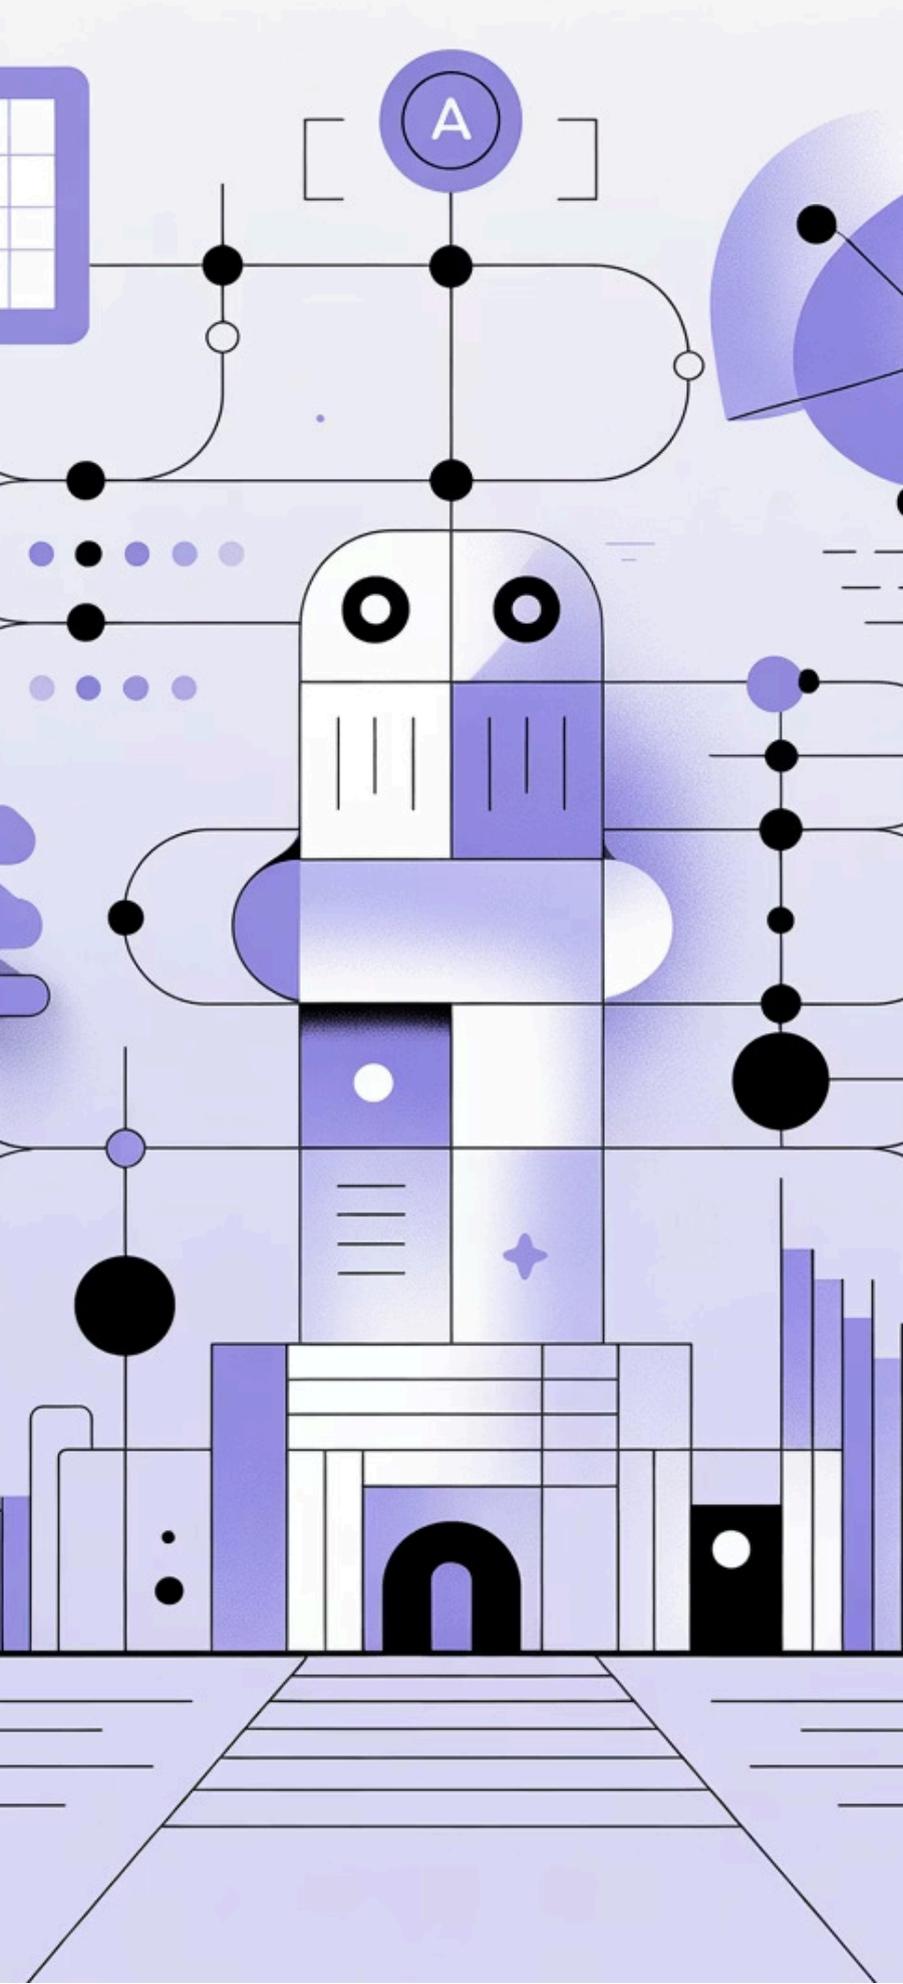
Open-source, model-driven framework designed for building sophisticated multi-agent collaboration workflows with flexibility and control.

LangGraph

Powerful open-source framework supporting advanced multi-agent orchestration with Model Context Protocol (MCP) for seamless integration.

Amazon ElastiCache & Neptune Analytics

Enable persistent, scalable memory architectures for agents using Valkey for vector storage and Neptune for entity relationships.



Designing Agentic AI Architectures: Patterns & Workflows

01

Reusable Agent Patterns

Deploy specialized patterns: reasoning agents for complex logic, retrieval-augmented agents for knowledge integration, coding agents for development, and workflow orchestrators for process management.

02

Implement Agent Loops

Create iterative cycles where LLMs reason through problems, make informed decisions, and invoke appropriate tools until objectives are successfully achieved.

03

Multi-Agent Orchestration

Build event-driven workflows with supervisor agents delegating tasks to specialized sub-agents, enabling sophisticated collaboration for complex requirements.

04

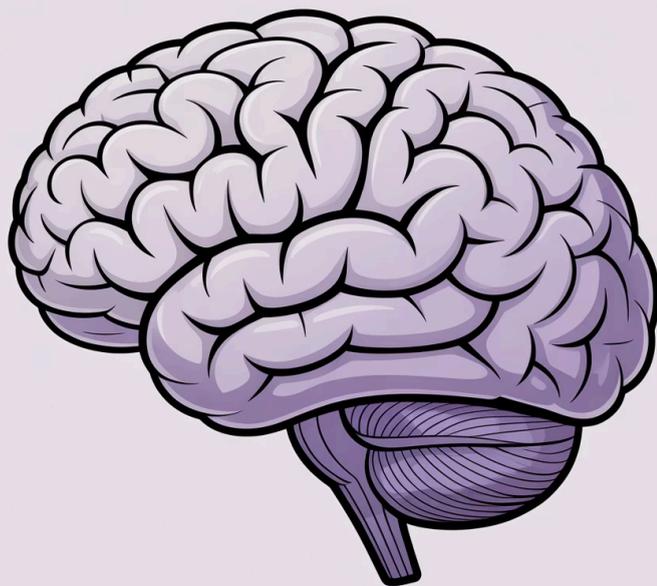
Observability & Control

Implement comprehensive monitoring to track agent decisions, ensure alignment with business goals, and maintain transparent, auditable operations.

Building Persistent Memory for Agentic AI

The Memory Challenge

Stateless agents severely limit the ability to maintain context across multiple sessions and deliver truly personalised experiences that users expect from intelligent systems.



The AWS Solution

Integrate Mem0 Open Source with Amazon ElastiCache for Valkey (vector storage) and Neptune Analytics (entity relationships) to create robust memory architectures.

- **Short-term memory:** Session-specific context and immediate conversation history
- **Long-term memory:** Cross-session learning, user preferences, and adaptive behaviours
- **AgentCore Runtime:** Support for multiple LLM providers including Amazon Bedrock, Anthropic, and OpenAI

Step-by-Step: Building Agentic AI on AWS

Example Use Case: Intelligent Mortgage Assistant



Start with Single-Agent Workflows

Deploy Amazon Bedrock agents for focused tasks like mortgage preapproval processing, document verification, and initial customer interactions.



Expand to Multi-Agent Systems

Introduce supervisor agents that intelligently route requests to specialised agents handling credit checks, property valuation, and risk assessment.



Deploy Scalable Compute

Leverage AWS Lambda for serverless execution or Amazon ECS for containerised agents, ensuring automatic scaling and cost optimisation.



Add Persistent Memory & Real-Time Data

Incorporate memory systems and live data integration to enable dynamic, context-aware interactions that learn from each customer engagement.

Best Practices Summary



Choose the Right Framework

Evaluate your use case requirements: select managed services like Bedrock for rapid deployment, or DIY frameworks like Strands and LangGraph for maximum customisation and control.



Design Modular Architectures

Build event-driven, loosely coupled systems that promote scalability, maintainability, and seamless integration with existing enterprise infrastructure.



Implement Robust Memory Management

Enable personalised, multi-session experiences through sophisticated memory systems that allow agents to learn, adapt, and improve over time.



Prioritise Security & Compliance

Leverage AWS native security controls, implement responsible AI guardrails, and ensure privacy protection throughout your agent lifecycle.



Monitor & Iterate Continuously

Deploy comprehensive observability tools to track agent performance, validate decision-making, and ensure ongoing alignment with business objectives.

Real-World Impact & Business Value

Automate Complex Workflows

Transform DevOps automation, financial services operations, and supply chain optimisation with intelligent, autonomous decision-making.

Accelerate Development

Deploy AI copilots like Kiro autonomous agent to dramatically boost developer productivity and code quality.

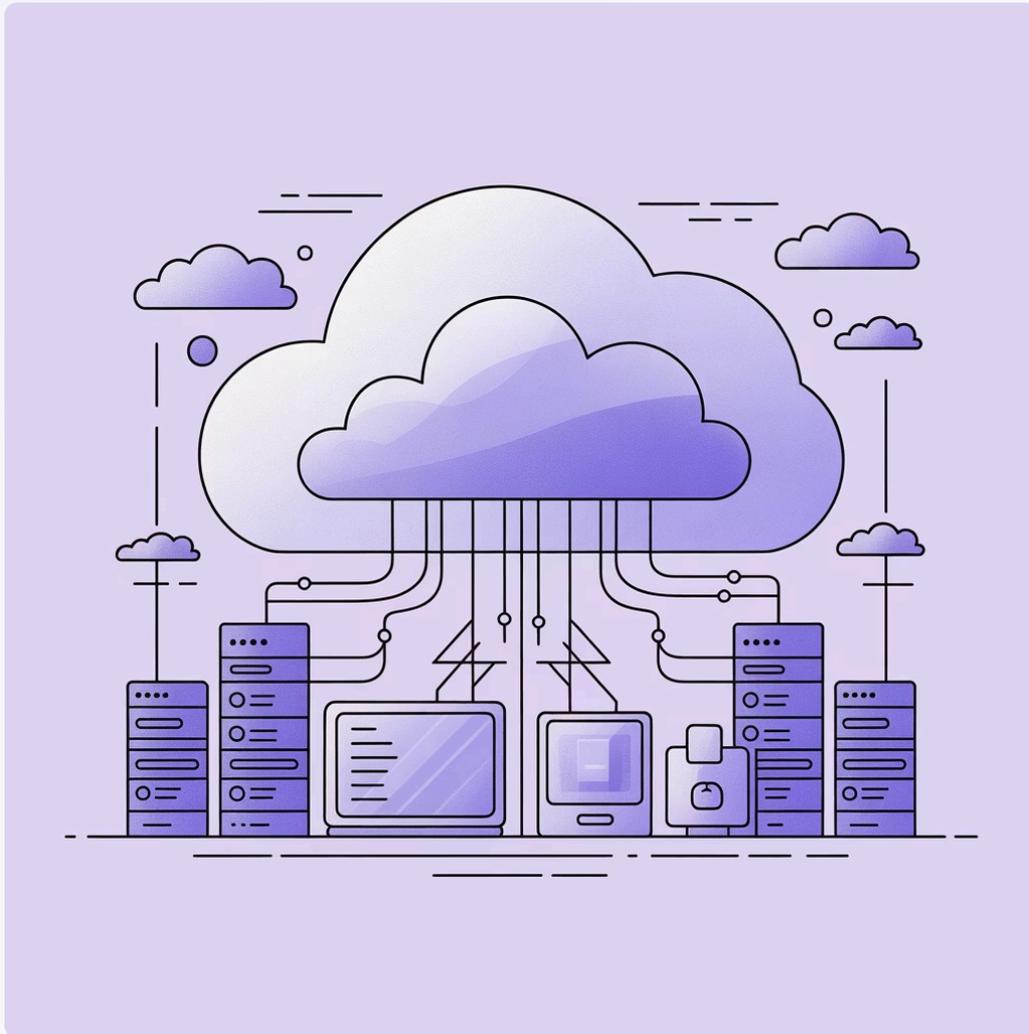
Reduce Operational Costs

Automate decision flows that previously required constant human oversight, freeing teams for strategic initiatives.

Drive Innovation

Build adaptive, goal-driven AI agents that transform legacy systems and unlock new business capabilities.

Ready to Build Your Agentic AI Application on AWS?



Start Your Journey Today



Explore AWS Resources

Visit the [AWS Agentic AI Hub](#) for comprehensive documentation, case studies, and architectural guidance.



Get Hands-On Experience

Access practical workshops on GitHub: [Building Agentic AI on AWS](#)



Connect with Experts

Engage with AWS AI specialists to accelerate your implementation and overcome technical challenges.

Embrace the future: Build autonomous, intelligent agents that transform your business and unlock unprecedented operational efficiency.