Advanced Generative AI Development on AWS

Category: AI, Data & Analytics | **Vendor:** Amazon Web Services

Duration: 24.00 hours (3 days) **19.5 CPD Hours** **Rating:** ⭐ 4.6 (5,878 reviews)

Course Information

Delivery Format: Instructor Led - Online

Course Overview

The Advanced Generative AI Development on AWS is designed for developers seeking to master the implementation of production-ready generative AI solutions on AWS. The course addresses the needs of organizations embarking on their generative AI journey and how to build comprehensive generative AI strategies that align with broader business objectives. This advanced 3-day instructor-led training builds expertise across the entire generative AI stack - from foundation models to enterprise integration patterns. In addition, you will learn about advanced data processing techniques, vector database implementation and retrieval augmentation, sophisticated prompt engineering and governance, agentic AI systems and tool integration, AI safety and security measures, performance optimization and cost management strategies, comprehensive monitoring and observability solutions, testing and validation frameworks. The course structure follows AWS's proven model for generative AI adoption, progressing from experimentation to production-ready implementations.

About This Course

The Advanced Generative AI Development on AWS is designed for developers seeking to master the implementation of production-ready generative AI solutions on AWS. The course addresses the needs of organizations embarking on their generative AI journey and how to build comprehensive generative AI strategies that align with broader business objectives. This advanced 3-day instructor-led training builds expertise across the entire generative AI stack - from foundation models to enterprise integration patterns. In addition, you will learn about advanced data processing techniques, vector database implementation and retrieval augmentation, sophisticated prompt engineering and governance, agentic AI systems and tool integration, AI safety and security measures, performance optimization and cost management strategies, comprehensive monitoring and observability solutions, testing and validation frameworks. The course structure follows AWS's proven model for generative AI adoption, progressing from experimentation to production-ready implementations.

Who Should Attend

- Software developers
- Technical Professionals

Prerequisites & Entry Requirements

General Prerequisites:

- AWS Technical Essentials
- Generative AI Essentials on AWS
- 2 or more years of experience building production grade applications on AWS or with open-source technologies, general AI/ML or data engineering experience
- 1 year of hands-on experience implementing generative AI solutions

Learning Outcomes

Upon successful completion of this course, participants will be able to:

- Develop production-ready generative AI solutions using AWS services that meet enterprise requirements for security, scalability, and reliability
- Evaluate and select appropriate foundation models for specific business use cases, including benchmarking performance and implementing dynamic model selection architectures
- Design and implement resilient foundation model systems with circuit breakers, cross-region deployment, and graceful degradation strategies
- Build comprehensive data processing pipelines for multi-modal inputs, including validation workflows and optimization techniques
- Implement sophisticated vector database solutions using Amazon Bedrock Knowledge Bases, OpenSearch, and hybrid approaches for effective retrieval augmentation
- Create and manage advanced prompt engineering frameworks, including chain-of-thought reasoning and enterprise-wide prompt governance systems
- Develop autonomous AI agents using Amazon Bedrock Agents, implementing complex reasoning patterns and tool integration capabilities
- Implement comprehensive AI safety and security controls, including content filtering, privacy preservation, and adversarial testing mechanisms
- Optimize performance and manage costs through token efficiency strategies, batching implementations, and intelligent caching systems
- Design and implement comprehensive monitoring and observability solutions for foundation model applications
- Create systematic testing and validation frameworks for continuous quality assurance of AI applications
- Integrate generative AI solutions within enterprise environments using secure, compliant, and scalable architectural patterns

Detailed Course Outline

Day 1

Module 1: Foundation Model Selection and Configuration

- Enterprise foundation model evaluation framework
- Dynamic model selection architecture patterns
- Resilient foundation model system designs
- Cost optimization and economic modeling

Module 2: Advanced Data Processing for Foundation Models

- Comprehensive data validation and quality assurance
- Multi-modal data processing pipelines
- Input optimization and performance enhancement

Module 3: Vector Databases and Retrieval Augmentation

- Enterprise vector database architecture
- Advanced document processing and chunking strategies
- Sophisticated retrieval system implementation
- Hands-on Lab: Develop Retrieval Augmented Generation (RAG) Applications with Amazon
- Bedrock Knowledge Bases

Day 2

Module 4: Prompt Engineering and Governance

- Advanced prompt engineering frameworks

- Complex prompt orchestration systems

- Enterprise prompt governance and management

- Hands-on Lab: Develop conversation pattern with Amazon Bedrock APIs

Module 5: Implementing Agentic AI Frameworks with Amazon Bedrock AgentCore

- Agentic AI Frameworks

- Amazon Bedrock AgentCore

Module 6: AI Safety and Security

- Comprehensive content safety implementation

- Privacy-preserving AI architecture

- AI governance and compliance frameworks

Day 3

Module 7: Performance Optimization and Cost Management

- Token efficiency and cost optimization

- High-performance system architecture

- Intelligent caching systems implementation

- Hands-on Lab: Building Secure and Responsible Gen AI with Guardrails for Amazon Bedrock

Module 8: Monitoring and Observability for Generative AI

- Foundation model monitoring systems

- Business impact and value management

- AI-specific troubleshooting and diagnostics

Module 9: Testing, Validation, and Continuous Improvement

- Comprehensive AI evaluation frameworks

- Quality assurance and continuous improvement

- RAG system evaluation and optimization

Module 10: Enterprise Integration Patterns

- Enterprise connectivity and integration architecture

- Secure access and identity management

- Cross-environment and hybrid deployments

Module 11: Course wrap-up

- Next steps and additional resources

- Course summary, Foundation Model Selection and Configuration

Enterprise foundation model evaluation framework

Dynamic model selection architecture patterns

Resilient foundation model system designs

Cost optimization and economic modeling

Advanced Data Processing for Foundation Models

Comprehensive data validation and quality assurance

Multi-modal data processing pipelines

Input optimization and performance enhancement

Vector Databases and Retrieval Augmentation

Enterprise vector database architecture

Advanced document processing and chunking strategies

Sophisticated retrieval system implementation

Hands-on Lab: Develop Retrieval Augmented Generation (RAG) Applications with Amazon

Bedrock Knowledge Bases

Prompt Engineering and Governance

Advanced prompt engineering frameworks

Complex prompt orchestration systems

Enterprise prompt governance and management

Hands-on Lab: Develop conversation pattern with Amazon Bedrock APIs

Agentic AI and Tool Integration

Agentic AI architecture and evolution

Amazon Bedrock Agents implementation

AWS Agentic AI service ecosystem

Tool integration and production observability

AI Safety and Security

Comprehensive content safety implementation

Privacy-preserving AI architecture

AI governance and compliance frameworks

Performance Optimization and Cost Management

Token efficiency and cost optimization

High-performance system architecture

Intelligent caching systems implementation

Hands-on Lab: Building Secure and Responsible Gen AI with Guardrails for Amazon Bedrock

Monitoring and Observability for Generative AI

Foundation model monitoring systems

Business impact and value management

AI-specific troubleshooting and diagnostics

Testing, and Continuous Improvement

Comprehensive AI evaluation frameworks

Quality assurance and continuous improvement

RAG system evaluation and optimization

Enterprise Integration Patterns

Enterprise connectivity and integration architecture

Secure access and identity management

Cross-environment and hybrid deployments

Course wrap-up

Next steps and additional resources

Course summary

Skills You Will Learn:

- Develop production-ready generative AI solutions using AWS services that meet enterprise requirements for security, scalability, and reliability
- Evaluate and select appropriate foundation models for specific business use cases, including benchmarking performance and implementing dynamic model selection architectures
- Design and implement resilient foundation model systems with circuit breakers, cross-region deployment, and graceful degradation strategies
- Build comprehensive data processing pipelines for multi-modal inputs, including validation workflows and optimization techniques
- Implement sophisticated vector database solutions using Amazon Bedrock Knowledge Bases, OpenSearch, and hybrid approaches for effective retrieval augmentation
- Create and manage advanced prompt engineering frameworks, including chain-of-thought reasoning and enterprise-wide prompt governance systems
- Develop autonomous AI agents using Amazon Bedrock Agents, implementing complex reasoning patterns and tool integration capabilities
- Implement comprehensive AI safety and security controls, including content filtering, privacy preservation, and adversarial testing mechanisms
- Optimize performance and manage costs through token efficiency strategies, batching implementations, and intelligent caching systems
- Design and implement comprehensive monitoring and observability solutions for foundation model applications
- Create systematic testing and validation frameworks for continuous quality assurance of AI applications
- Integrate generative AI solutions within enterprise environments using secure, compliant, and scalable architectural patterns

Frequently Asked Questions

Q: What delivery options are available for Advanced Generative AI Development on AWS?

We offer multiple delivery formats:

- Live Instructor-Led Classroom Online (Virtual/Live Online)
- Traditional Instructor-Led Classroom Training (ILT)
- On-site delivery at your offices anywhere in United Kingdom
- Private dedicated courses customized for your team

Q: How many CPD hours does this course provide?

The 3-day Advanced Generative AI Development on AWS course provides up to 19.5 CPD hours of structured learning. CPD certificates can be provided upon request.

Q: What is the duration of the Advanced Generative AI Development on AWS training?

The training takes place over 3 day(s), with each day lasting approximately 24.00 hours including breaks for lunch and refreshments.

Q: Do you provide corporate training for Advanced Generative AI Development on AWS?

Yes, we provide corporate training, dedicated training, and closed classes for Advanced Generative AI Development on AWS. Training can take place anywhere in United Kingdom including London, Manchester, Birmingham, Edinburgh, or live online allowing teams from across United Kingdom or internationally to attend.

Q: What related terms do people search for?

Popular related searches include: Advanced Generative AI Development on AWS, Machine Learning, AGAID, Amazon Web Services, vILT

Q: Why choose Nexus Human for Advanced Generative AI Development on AWS?

Nexus Human is recognized as one of the leading training providers. Our trainers have won multiple awards including:

- Small Firms Best Trainer Award
- National Training Partner of the Year (Ireland) - Multiple Years
- Global Top 30 Instructor Awards (2012, 2019, 2021)
- Tech Excellence Award Nominations
- Learning Performance Institute (LPI) External Training Provider Sponsor 2024

Q: Are there any discount codes available?

Yes! Use discount code **PENPAL5** when booking your Advanced Generative AI Development on AWS training. Please note that only one discount code can be used per booking and cannot be combined with other special offers.

Nexus Human

Professional Training & Development

Email: info@nexushuman.com

Website: www.nexushuman.com

Phone: +353 1 889 8700 (Ireland) | +44 20 8142 8961 (UK)

© 2026 Nexus Human. All rights reserved. This brochure was generated on 14/11/2026.