Faysal Aziz

OpenClaw for Techies

Building Practical AI Agents That Actually Do Things

FA-0001 · 1 day

Agentic & Generative AI

Deploy, connect, secure, and extend a working OpenClaw agent in one intensive day.

Introduction

The shift from conversational AI to agentic AI changes what technical teams can reasonably automate. Instead of limiting an AI system to answering questions, an agent can work with files, execute approved tools, maintain contextual workspaces, interact with messaging platforms, and coordinate actions across systems. OpenClaw sits directly in this emerging operational layer.

OpenClaw is an open-source agent platform built around a Gateway that connects AI agents with communication channels, tools, nodes, sessions, hooks, and external services. Its current architecture supports channels including Slack, Telegram, WhatsApp, Discord, Microsoft Teams, Signal and others, while allowing administrators to choose their own model providers and maintain control of the environment where the agent operates.

This one-day technical course is deliberately focused on deploying and operating OpenClaw rather than studying agent theory. Participants will move through installation, onboarding, model connectivity, Gateway architecture, workspaces, tools, skills, channels and the security controls that matter when an AI agent is allowed to interact with real systems. Current OpenClaw releases provide guided onboarding, configurable workspaces, a skills ecosystem, sandboxed tool execution and built-in security auditing, all important capabilities for technical teams evaluating agents beyond experimentation.

Security receives particular attention because an agent with tools is fundamentally different from an ordinary chatbot. OpenClaw's documentation explicitly defines a single trusted-operator model for each Gateway and recommends separate trust boundaries for mutually untrusted users. It also provides Docker-backed tool sandboxing, per-agent tool policies and configurable workspace access. These are operational considerations that technical practitioners need to understand before granting an agent meaningful access to infrastructure or organizational systems.

The course is designed for technical practitioners, not introductory computer users. The instructor brings over 30 years of industry experience, and the training will emphasize real industry-demanded configurations, operational decisions and deployment practices rather than turning OpenClaw into an academic AI course. Because this is a one-day program, the objective is not to cover every integration or configuration option. The realistic target is to give participants enough architectural understanding and hands-on familiarity to deploy, configure, secure and extend a functional OpenClaw environment.

Learning outcomes

By the end of this course, participants should be able to:

- Explain the OpenClaw architecture and agent operating model.
- Install and perform initial configuration of OpenClaw.
- Configure an OpenClaw workspace and Gateway.
- Connect and manage an AI model provider.
- Understand agents, sessions, tools, skills and channels.
- Configure basic agent behavior and capabilities.
- Understand how OpenClaw skills extend agent functionality.
- Connect OpenClaw to supported communication channels.
- Apply appropriate tool permissions and workspace restrictions.
- Understand sandboxing and agent isolation.
- Perform basic OpenClaw security auditing.
- Troubleshoot common installation, Gateway and configuration problems.
- Identify realistic enterprise and technical use cases for OpenClaw.

Prerequisites

Participants should have:

- Comfortable command-line experience.
- Experience with POSIX / Linux.
- Familiarity with JSON configuration files.
- Good understanding of APIs and API keys.
- Deep understanding of generative AI and LLMs.
- Administrator or appropriate installation privileges on the training machine.
- Internet access during training.
- Access to a supported AI model provider.
- Working Docker knowledge.

Training outline

1. OpenClaw Architecture and Agent Fundamentals

- OpenClaw platform overview
- Agentic AI versus conversational AI

  OpenClaw architecture
  - Gateway
  - Agents
  - Workspaces
  - Sessions
  - Tools
  - Skills
  - Channels
  - Nodes
- Local and remote deployment models
- OpenClaw configuration structure
- Current OpenClaw ecosystem and operational use cases

2. Installation and Initial Configuration

- Platform and runtime requirements
- OpenClaw installation options
- CLI fundamentals
- Guided onboarding
- Model provider configuration
- Authentication and credentials
- Workspace initialization
- Configuration files and directories
- Gateway startup and validation
- Control UI and TUI access

3. Working with Agents

- Agent configuration
- Agent instructions and behavior

- Workspace organization
- Sessions and context
- Model selection
- Agent execution lifecycle
- Multi-agent concepts
- Agent-specific permissions

4. Tools and Skills

- OpenClaw tool architecture

  Tool permissions
  - Allow and deny policies
  - Filesystem access
  - Command execution
  - Browser and external tools

  OpenClaw skills
  - SKILL.md structure
  - Workspace and managed skills
  - Skill discovery and loading
  - Skill precedence

  ClawHub
  - Skill discovery
  - Installation and updates
  - Skill verification
- Custom skill development workflow

5. Channels and External Connectivity

- OpenClaw channel architecture
- Channel pairing and authentication
- Supported messaging platforms
- Message routing
- Agent-to-channel configuration
- Multi-channel considerations
- External service integration concepts

6. Security, Sandboxing and Operational Controls

- OpenClaw trust model
- Gateway security
- Tool execution risks
- Prompt injection considerations
- Secrets and credential exposure

  Docker-backed tool sandboxing
  - Sandbox modes
  - Workspace access
  - Agent and session isolation
- Tool allowlists and denylists
- Elevated tool execution
- Multi-agent security boundaries
- OpenClaw security auditing

7. Operating and Troubleshooting OpenClaw

- Gateway status and management
- Configuration validation
- Logs and diagnostics
- OpenClaw Doctor
- Security audit workflow
- Skills validation
- Model connectivity troubleshooting
- Channel troubleshooting
- Upgrade and configuration considerations
- Production-readiness considerations

8. Building a Practical OpenClaw Environment

- Agent purpose and scope
- Workspace design
- Model configuration
- Required tool selection
- Skill selection and customization
- Channel integration

- Permission boundaries
- Sandbox strategy
- Security validation
- Operational deployment checklist

Disclaimer

This course outline is intended as a structured training guideline rather than an inflexible syllabus. The trainer reserves the right to amend, reorganize, expand, reduce, substitute or omit topics where reasonably necessary to reflect participant competency, available training time, laboratory conditions, software changes, security considerations, current industry practices or developments within the OpenClaw ecosystem. Such adjustments may be made at the trainer's professional discretion without prior notice in order to preserve the relevance, quality and practical value of the training.

© 2026 Faysal Aziz · Kuala Lumpur