

OpenClaw for Technical Practitioners

Deploying and operating a practical AI-agent lab

FA-0650 · 1 day

Agentic & Generative AI · DevOps, Cloud & Infrastructure

Build a scoped OpenClaw lab with one agent, an approved model provider, selected tools and a demonstration integration; review permissions, sandboxing and operational checks.

Introduction

This intensive technical introduction follows the OpenClaw architecture from Gateway and workspace configuration to agent instructions, tools, skills and external connectivity. It is intended for practitioners evaluating tool-using AI agents, not introductory computer users.

Participants work through a bounded lab: one agent, one approved model provider, selected tools and one channel or external-service demonstration. Architecture and operational topics use guided exercises or demonstrations where a full implementation would exceed a day. Advanced integrations and production rollout remain follow-on work.

Security and operations are treated as configuration decisions to inspect and validate, not a guarantee of a secure environment. Sandboxing is configurable and is not enabled by default; the Gateway remains on the host. A shared Gateway assumes a trusted operator or team, and separate agent workspaces alone do not isolate mutually untrusted users.

The course covers installation and operational tasks within an authorised training environment. Provider credentials, installation privileges, Docker and any service accounts must be prepared before the course. Skills and their supporting code require review; instructions do not bypass tool policy or grant permissions.

Learning outcomes

The course teaches participants to:

- Explain Gateways, agents, sessions, workspaces, tools, skills, channels and nodes.

- Install or inspect a prepared OpenClaw lab and validate basic Gateway and provider connectivity.
- Configure a scoped agent purpose, instructions, workspace and selected capabilities.
- Understand skill structure, discovery, precedence and verification, and adapt a small lab skill.
- Demonstrate one supported integration with appropriate authentication, routing and access controls.
- Compare tool policies, workspace access and configurable sandbox modes, including their limits.
- Review logs, configuration validation, Doctor and security-audit findings before deciding on repairs.
- Identify unresolved operational risks and produce a follow-on deployment checklist rather than declaring the lab production-ready.

Prerequisites

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

A disposable, authorised training machine or lab environment with the required runtime, Docker capability, approved model-provider access and any selected integration account prepared in advance. Potential charges and organisational permissions are handled before the course.

Training outline

1. 1. Architecture and Installation

Agent architecture

- OpenClaw platform overview
- Agentic AI versus conversational AI

OpenClaw architecture

- Gateway
- Agents
- Workspaces

- Sessions
- Tools
- Skills
- Channels
- Nodes
- Local and remote deployment models: architectural comparison, with one selected lab setup
- OpenClaw configuration structure
- Current OpenClaw ecosystem and operational use cases

Initial configuration and validation

- 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 access and TUI overview where supported

2. 2. Agents, Tools and Skills

Agent behaviour and context

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

Tools, policies and skill development

- OpenClaw tool architecture

Tool permissions

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

OpenClaw skills

- SKILL.md structure
- Workspace, managed and other available skill sources
- Skill discovery and loading
- Skill precedence

ClawHub

- Skill discovery
 - Installation and updates
 - Skill verification, provenance and review of supporting code; scanning is not proof of safety
- A small custom skill example: instructions, supporting resources and review

3. 3. Channels and External Connectivity

- OpenClaw channel architecture
- Channel pairing and authentication
- Supported messaging platforms
- Message routing
- Agent-to-channel configuration
- Multi-channel considerations: configuration and feature differences, not configuration of every channel
- External service integration concepts

4. 4. Security and Operational Controls

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

Docker-backed sandboxing as one supported backend; explicitly enable and inspect the selected lab policy

- Sandbox modes
- Workspace access
- Agent and session sandbox scopes and their limitations
- Tool allowlists and denylists
- Elevated execution: risks, policy controls and when to avoid it
- Multi-agent organisation versus actual security boundaries; separate trust domains for mutually untrusted users
- OpenClaw security auditing

5. 5. Operations and Troubleshooting

- Gateway status and management
- Configuration validation
- Logs and diagnostics
- OpenClaw Doctor: diagnostics, read-only checks and potentially state-changing repair modes
- Security audit workflow
- Skills validation
- Model connectivity troubleshooting
- Channel troubleshooting
- Upgrade considerations: configuration compatibility, backup and review before changes
- Production-readiness gaps: access, secrets, monitoring and operational ownership

6. 6. An Integrated Lab and Follow-On Checklist

- Agent purpose and scope
- Workspace design
- Model configuration
- Required tool selection
- Approved skill selection and a bounded customisation exercise
- Channel integration
- Permission boundaries
- Sandbox strategy
- Security validation
- Operational deployment checklist and unresolved follow-on work

