THE OPENCLAW FIELD MANUAL

25+ Use Cases, Architecture Deep Dive & Security Runbook

INSIDE THIS GUIDE:

- → Accurate history: Clawdbot to OpenClaw
- → 25+ real-world use cases
- → Hardened security runbook
- → Gateway, Channels & Skills architecture
- → Skill Risk Matrix (T1/T2/T3)
- → NemoClaw & what comes next

2

The Agent That Actually Does Things

OpenClaw is the most-starred AI agent project of early 2026 — and one of the most misunderstood. Unlike standard AI assistants that wait to be asked, **OpenClaw executes tasks autonomously**, retains memory across sessions, and connects to the tools you already use — all processed locally on your own hardware.

> **THE TMINUSAI AXIOM**
>
> "Reliability is built on localized infrastructure, not cloud-based logic. Own your agent stack."

Verified Timeline (Public Record)

Date	Event	Name at the Time
Nov 2025	Peter Steinberger releases the first version	Clawdbot
Jan 27, 2026	Renamed after trademark complaint from Anthropic (re: "Clawd")	Moltbot
Jan 30, 2026	Final rename: "Moltbot never quite rolled off the tongue"	OpenClaw
Jan 2026	25,000 GitHub stars in a single day. Moltbook social network launches.	OpenClaw
Feb 14, 2026	Steinberger joins OpenAI; project moves to open-source foundation	OpenClaw
Mar 2, 2026	247,000 GitHub stars, 47,700 forks	OpenClaw
Mar 2026	Nvidia announces NemoClaw — enterprise-grade fork in development	OpenClaw / NemoClaw

SECTION 1 — Architecture Deep Dive

How OpenClaw Actually Works

OpenClaw's power comes from three interlocking components. Understanding the architecture is mandatory before you deploy any skill that touches production systems.

Component 1: The Gateway

The Gateway is a Node.js process that runs continuously in the background. It is the central control plane that:

- Normalizes incoming messages from any connected channel (WhatsApp, Telegram, etc.)
- Loads the session memory and installed Skills for the current user
- Passes the full context to the configured LLM (Claude, GPT-5, DeepSeek, or local)
- Executes the LLM's determined actions and returns results to the chat thread
- Persists session history and memory to disk — survives restarts

Component 2: Channels

Channels are the messaging integrations that serve as the human interface. OpenClaw supports 20+ channels including: WhatsApp, Telegram, Discord, Slack, Signal, iMessage. The agent is accessible anywhere you already are.

Component 3: Skills

Skills are modular add-ons stored as directories containing a `SKILL.md` file with metadata and instructions. The community-maintained registry (ClawHub) hosts 100+ preconfigured skills. Agents can even write their own new skills dynamically.

> **IMPORTANT: MODEL AGNOSTIC BY DESIGN**
>
> OpenClaw works with any LLM via API key. Supported out of the box: Anthropic Claude, OpenAI GPT-5 / GPT-5.2, Google Gemini, DeepSeek, and any local model hosted via Ollama. Switching models requires zero reconfiguration of the rest of the stack.

SECTION 2 — Skill Risk Matrix

The TminusAI Skill Risk Matrix (SRM)

Before installing any skill from ClawHub, evaluate it against the SRM. A January 2026 audit by Kaspersky identified 512 vulnerabilities in OpenClaw — 8 classified as critical. Vetting is not optional.

Tier	Permission Level	Risk	TminusAI Deployment Rule
T1: Read-Only	Reads files/data without modification	Low	Safe for laptop or home server. No isolation required.
T2: Transactional	Specific, reversible actions (e.g., send Slack message)	Medium	Requires explicit pairing (DM code confirmation). Log all actions.
T3: Executive	Shell access, file modification, API write operations	HIGH	MUST isolate in Docker or dedicated VPS. Zero-trust policy.

SECTION 3 — 25+ Real-World Use Cases

What People Are Actually Doing With OpenClaw

Category A: Personal Productivity

- **CEO Morning Briefing:** Agent scans Gmail, Outlook, Slack for "High Priority" at 6am; synthesizes into a Telegram summary by 6:30am.
- **Inbox Zero:** Agent categorizes, unsubscribes, drafts responses, and archives — clearing thousands of emails autonomously.
- **Calendar Defense:** Agent blocks focus time, reschedules low-priority meetings, sends polite decline templates.
- **Weekly Meal Planning:** Agent reads dietary preferences from Notion, queries recipes, creates grocery list, syncs to Apple Reminders.
- **Daily Standup Prep:** Agent reads Jira tickets updated yesterday, drafts 3-line standup summary to Slack by 9am.

Category B: Developer & Technical Workflows

- **Automated PR Review:** Webhook triggers agent on every GitHub PR; model checks diff, posts inline comments, summarizes risk.
- **Overnight Coding Agent:** Agent runs Codex tasks (refactor, feature build) while you sleep; commits results to staging branch.
- **Smart Home Automation:** Agent reads biomarker data (WHOOP, Apple Health) and adjusts air quality, lighting, temperature accordingly.
- **AWS CloudFormation Deployments:** Agent provisions infrastructure from a Telegram command — with confirmation step before execution.
- **Bug Triage Agent:** Agent monitors error logs, classifies severity, creates Jira tickets, and pages on-call for P0 incidents.

Category C: Content & Research

- **Research Aggregator:** Agent browses 10+ sources on a topic, extracts key claims, synthesizes into a Notion page with citations.
- **Social Listening Monitor:** Agent scans Reddit, Twitter, LinkedIn for brand mentions; sends daily digest with sentiment analysis.
- **Competitor Intelligence:** Agent monitors competitor websites for product updates, pricing changes, and new blog posts.
- **Newsletter Draft Generator:** Agent collects the week's key stories, drafts a newsletter, posts draft to Notion for your review.

Category D: Business Operations

- **Lead Research Assistant:** Agent enriches CRM leads — finds LinkedIn profiles, company info, recent news, and email addresses.
- **Contract Monitoring:** Agent reads vendor contracts, flags renewal dates, extracts key terms, and logs to Notion 60 days ahead.
- **Customer Ticket Triage:** Agent reads support tickets, classifies by product area and urgency, routes to the correct Slack channel.
- **Expense Report Generator:** Agent reads bank/card statements, categorizes transactions, drafts expense report in Google Sheets.
- **Invoice Processing:** Agent reads invoice PDFs, extracts vendor name, amount, due date — logs to accounting system.

Category E: Advanced / Power Users

- **Moltbook Agent:** Agent maintains a public profile on Moltbook (the agent-native social network) and interacts with other agents.
- **Personal Second Brain:** Agent integrates Obsidian notes, builds a knowledge graph, answers queries from your own notes.
- **Renting Humans:** For tasks that require real-world action, OpenClaw agents can post requests to Taskrabbit, Fiverr, or similar.
- **Insurance Dispute Agent:** Agent composes and sends dispute emails to insurers (real reported case: successfully triggered reinvestigation).
- **Raspberry Pi Home Server:** Agent runs 24/7 on a low-power Raspberry Pi; accessible via WhatsApp from anywhere in the world.

SECTION 4 — Hardened Security Runbook

Zero-Trust Deployment Protocol

OpenClaw requires broad system access to function: email credentials, API keys, calendar tokens, browser cookies, filesystem access, and terminal permissions. The following runbook is non-negotiable for production deployments.

#	Security Control	Why It Matters	Implementation
1	Bind Gateway to loopback only	Prevents remote exploitation of open port	Use 127.0.0.1, not 0.0.0.0. Use Tailscale for remote access.
2	Store secrets in a vault	~/.openclaw/ is a primary infostealers target	Use environment variables or 1Password/Bitwarden vault.
3	Human-in-the-loop for T3 skills	Prevents autonomous deletion, financial actions	Require "Confirm" button or 2FA approval for destructive ops.
4	Isolate T3 skills in Docker	Limits blast radius of compromised skill	Run shell/file-access skills in isolated Docker container.
5	Audit prompt injection surface	Emails/pages can carry malicious instructions to the agent	Sanitize untrusted inputs; limit which content sources feed the LLM.
6	Log all agent actions	Essential for debugging and incident response	Enable verbose logging; rotate logs to a secure location.

MAINTAINER WARNING (VERBATIM)

"If you cannot understand how to run a command line, this is far too dangerous of a project for you to use safely." — OpenClaw maintainer "Shadow," February 2026 (Discord). Heed this advice seriously.

WHAT IS NEXT — NemoClaw & The Enterprise Layer

The Ecosystem Is Growing Fast

As of March 2026, Nvidia is developing **NemoClaw** — an enterprise-grade fork of OpenClaw with hardened security, governance controls, and native integrations with Salesforce, Cisco, Google, Adobe, and CrowdStrike. OpenAI (who hired OpenClaw's creator) is expected to influence the project's roadmap significantly.

NEXT STEPS

Continue with Guide 3: Multimodal AI Mastery to learn how to feed images, audio, and video into your agent workflows — the critical missing layer most OpenClaw operators overlook. Available at tminusai.com.