BMS CYBER DEFENCE
Expert Cyber Security for SMEs

bmscyberdefence.co.uk
admin@bmscyberdefence.co.uk

OpenClaw AI Assistant

Security Best-Practice Guide

Authored by **Jason Lewis** | BMS Cyber Defence | Version 1.0 | February 2026

VPS | AWS | Oracle Cloud | Bare Metal | Free Public Guide

AI is transforming the way we work — and agentic AI tools like OpenClaw are leading that charge. This guide gives you everything you need to deploy OpenClaw securely, stay on the right side of UK law, and protect your business. Best of all, it's completely free.

Prepared by BMS Cyber Defence — specialists in making enterprise-grade security accessible for SMEs across the UK.

LIVE SECURITY ALERT — Update OpenClaw to v2026.1.29 Immediately

CVE-2026-25253 (CVSS 8.8) was disclosed in January 2026. All versions prior to v2026.1.29 are vulnerable to 1-click remote code execution. Upgrade before proceeding with this guide, then rotate your gateway token and all API keys.

About BMS Cyber Defence

BMS Cyber Defence was founded with a clear mission: to give small and medium-sized businesses access to the same quality of cyber security expertise previously available only to large enterprises — at a price that makes sense. Based in Telford, Shropshire, we serve SMEs across the UK with a full spectrum of cyber security services, delivered by accredited and certified specialists.

In today's threat landscape, AI tools like OpenClaw, autonomous agents, and cloud services are becoming standard business tools — and with them come new risks that traditional IT support simply isn't equipped to address. BMS Cyber Defence has made AI security and UK regulatory compliance a core specialism, helping businesses deploy these powerful technologies confidently and safely.

We don't just tell you what the risks are — we help you fix them. From Cyber Essentials certification (which can be a prerequisite for winning public sector contracts) to full AI risk assessments, supply chain security reviews, and incident response planning, our team gives your business practical, actionable protection.

What We Do	What That Means For You
AI Security & Regulatory Compliance	Assess your AI tools against NCSC, DSIT, and ICO requirements. Stay compliant before regulators come calling.
Cyber Essentials & Cyber Essentials Plus	Government-backed certification — increasingly required to win contracts with MSPs, NHS, and public sector clients.
Supply Chain Security Reviews	Ensure your security posture meets the demands of regulated clients under the new Cyber Security and Resilience Bill.
Vulnerability Assessments & Pen Testing	Find your weaknesses before attackers do. Professional testing by certified ethical hackers.
Incident Response Planning	Be ready when it matters most. We help you build and test your response procedure before an incident occurs.
Security Awareness Training	Your people are your biggest risk — and your best defence. Practical training that actually changes behaviour.
All-in-One Cyber Packages incl. Insurance	Comprehensive protection bundled for SMEs — security controls and cyber insurance in one straightforward package.

Ready to talk? **Visit bmscyberdefence.co.uk** or email admin@bmscyberdefence.co.uk

1. What Is OpenClaw?

OpenClaw (formerly Moltbot/Clawdbot) is an open-source, self-hosted personal AI assistant. Rather than relying on a third-party cloud service, it runs on your own infrastructure and connects to messaging channels — Telegram, WhatsApp, Slack, Discord, and 50+ others. You instruct it in plain English; it executes complex, multi-step tasks autonomously on your behalf.

Unlike workflow builders such as n8n or Make — where every automation must be manually wired together — OpenClaw understands natural-language instructions and acts independently: browsing the web, managing calendars, running commands, publishing content, and more. It gained 145,000 GitHub stars within weeks of launch in late 2025. With that popularity came serious security scrutiny, which is exactly why this guide exists.

Architecture in Plain English

The OpenClaw Gateway runs on your server and holds all your data and configuration. You connect from your phone or laptop via SSH or Tailscale. The AI 'thinking' happens on cloud model providers (OpenAI, Anthropic, etc.) — your server just coordinates. This means you do NOT need expensive hardware. A $5/month VPS is genuinely sufficient.

Official docs: docs.openclaw.ai | GitHub: github.com/openclaw/openclaw

2. Why Self-Host?

Self-hosting puts you in control: your data stays on infrastructure you own, your agent runs 24/7 without your laptop needing to be on, and you're not dependent on a third-party platform's pricing or availability decisions.

⚠ Do NOT Run OpenClaw on a Personal Machine

OpenClaw can execute terminal commands and manage files. Running it on your everyday laptop is a serious security risk. An isolated VPS or cloud instance protects your personal files if something goes wrong. Think of it as digital hygiene — not paranoia.

3. Choosing Your Infrastructure

OpenClaw is lightweight — 2 vCPUs and 4 GB RAM is plenty. Here are your main options:

Platform	Cost	Best For	Notes
VPS (Hetzner, DigitalOcean, Vultr)	~£4–10/mo	Most users — easiest to manage	One-click Docker templates available
AWS EC2 / Lightsail	Free tier or ~£3–8/mo	Teams already using AWS	Lightsail is simplest; EC2 for full control
Oracle Cloud Always Free	£0/mo — permanently	Zero-budget evaluation	4 ARM vCPUs, 24 GB RAM, never expires
Hetzner / Fly.io / Railway	£4–10/mo	European data residency	Official guides at docs.openclaw.ai/vps
Bare Metal / Mini PC	Hardware cost only	Air-gapped or high-trust on-premises	Needs static IP or VPN — see Section 10

VPS hosting overview: docs.openclaw.ai/vps

✔ SME Tip from BMS

For most SMEs trying OpenClaw for the first time, a £4/month Hetzner VPS or DigitalOcean Droplet is the ideal starting point. You can always scale up. Oracle's Always Free tier is excellent for evaluation with zero financial commitment.

4. Initial Server Hardening

Complete all steps in this section before installing OpenClaw. These apply to any Ubuntu/Debian host — VPS, AWS EC2, or bare metal.

4.1 — Create a Non-Root User

```
# Always run as root immediately after first login
adduser openclaw
usermod -aG sudo openclaw
su - openclaw
```

4.2 — Harden SSH Access

SSH is the most common attack vector on internet-facing servers. Disable password login immediately.

```
# On your LOCAL machine — copy your SSH key to the server
ssh-copy-id openclaw@YOUR_SERVER_IP

# On the SERVER — edit the SSH configuration
sudo nano /etc/ssh/sshd_config

# Set these values:
PermitRootLogin no
PasswordAuthentication no
PubkeyAuthentication yes
MaxAuthTries 3

sudo systemctl restart sshd
```

⚠ Before You Restart SSH

Test your key-based login in a SECOND terminal window before closing your current session. Disabling password auth without a working key will lock you out completely.

4.3 — Firewall

```
sudo apt install ufw -y
sudo ufw default deny incoming
```

```
sudo ufw default allow outgoing
sudo ufw allow 22/tcp      # SSH only
# NEVER open port 3000 (OpenClaw Gateway) publicly
sudo ufw enable
sudo ufw status verbose
```

4.4 — Node.js Version (Patches Two Known CVEs)

Two resolved CVEs require Node.js 22.12.0 LTS or later. Many long-lived VPS instances run older versions — check yours now.

```
node --version     # Must be v22.12.0 or later

# Upgrade if needed:
curl -fsSL https://deb.nodesource.com/setup_22.x | sudo bash -
sudo apt install -y nodejs
node --version     # Confirm
```

4.5 — Automatic Security Updates

```
sudo apt install unattended-upgrades -y
sudo dpkg-reconfigure --priority=low unattended-upgrades
```

5. Docker — Secure Container Setup

Docker containers isolate OpenClaw from the rest of your server. If something goes wrong inside the agent, it cannot easily affect your host OS or other services.

5.1 — Install Docker

```
sudo apt update && sudo apt install -y ca-certificates curl gnupg
curl -fsSL https://get.docker.com | sudo sh
sudo usermod -aG docker $USER
newgrp docker
docker --version
```

5.2 — Hardened Docker Compose File

Never Bind to 0.0.0.0

42,000+ OpenClaw instances were found publicly exposed in January 2026 because operators used 0.0.0.0 instead of 127.0.0.1. Always use 127.0.0.1 in port bindings. The ClawShield audit tool specifically checks for this.

```
# docker-compose.yml — secure production configuration
version: '3.8'
services:
  openclaw:
    image: openclaw/openclaw:latest
    restart: unless-stopped
    read_only: true
    cap_drop:
      - ALL
    security_opt:
      - no-new-privileges:true
    ports:
      - '127.0.0.1:3000:3000'   # LOOPBACK ONLY — never 0.0.0.0
    volumes:
      - openclaw-data:/app/data
    environment:
      - OPENCLAW_GATEWAY_TOKEN=${OPENCLAW_GATEWAY_TOKEN}
    logging:
      driver: 'json-file'
      options:
        max-size: '10m'
        max-file: '5'
volumes:
  openclaw-data:
```

6. Installing OpenClaw

Option A — CLI (Any Linux Server)

```
curl -fsSL https://openclaw.ai/install.sh | bash -s -- \
  --install-method git --no-prompt --no-onboard

openclaw --version    # Confirm: must be 2026.1.29 or later
openclaw onboard      # Mandatory — do not skip
```

Option B — One-Click Marketplace Templates

Hostinger, DigitalOcean, Hetzner, and Railway all offer pre-built OpenClaw Docker templates. The DigitalOcean 1-click droplet is particularly strong — it auto-generates a gateway token, enforces non-root execution, applies rate-limiting firewall rules, and configures DM pairing automatically.

- **Hostinger guide:** hostinger.com/tutorials/how-to-set-up-openclawVisit guide
- **DigitalOcean hardened deployment:** digitalocean.com/community/tutorials/how-to-run-openclawVisit guide

Option C — AWS EC2 / Lightsail

1. **Provision:** Launch EC2 (t3.small, Ubuntu 22.04) or Lightsail. Assign an Elastic/static IP.
2. **Security Group:** Allow port 22 inbound only. Do NOT open port 3000 publicly under any circumstances.
3. **Node.js:** Install Node.js 22.12.0 LTS (Section 4.4) and verify version.
4. **Docker & OpenClaw:** Follow Sections 5 and 6A on the instance.
5. **Access:** SSH tunnel or Tailscale only — see Section 7.2.

7. Securing the Gateway Token

The Gateway Token is the single credential controlling your entire OpenClaw instance. Anyone who holds it can run system commands on your server. Treat it like a root password combined with unrestricted API access.

> ⚠ **Protect Your Gateway Token**
>
> Never share it. Never commit it to a repository. Never transmit it over an unencrypted channel. Store it in a password manager — not a text file, not your shell history.

7.1 — Finding and Rotating the Token

```
# Find the current token
cat ~/.openclaw/config.json | grep -i token

# Or inside Docker
docker exec openclaw_gateway env | grep OPENCLAW_GATEWAY_TOKEN

# Rotate immediately if you suspect compromise
openclaw gateway token rotate
```

7.2 — Secure Remote Access

SSH Tunnel — Most Secure

```
# Run this on your LOCAL machine
ssh -L 3000:localhost:3000 openclaw@YOUR_SERVER_IP

# Then open in your local browser: http://localhost:3000
# The Gateway never touches the public internet
```

Tailscale — Best for Teams

Tailscale creates an encrypted WireGuard mesh between your devices. The Gateway remains on loopback; only your authorised devices can reach it. Free for personal/small team use.

```
curl -fsSL https://tailscale.com/install.sh | sh
sudo tailscale up
tailscale serve --bg 3000
```

Tailscale: tailscale.com

8. API Key Management

- **Dedicated keys:** Create a separate API key for OpenClaw on each provider. Never share keys between applications.
- **Spend limits:** Set hard monthly limits on every provider dashboard. An unexpected spike is often the first sign of compromise.
- **No hardcoding:** Store keys as environment variables. Never put them in config files, shell history, or source code.
- **Rotation:** Rotate all API keys every 90 days, or immediately after any suspected exposure.

```
export OPENAI_API_KEY='sk-...'
export ANTHROPIC_API_KEY='sk-ant-...'

# Scan for accidentally committed credentials
pip install detect-secrets==1.5.0
detect-secrets scan --baseline .secrets.baseline
```

9. Messaging Channel Security

- **DM pairing:** Enable private DM pairing so only authorised accounts can send commands to your agent.
- **Telegram:** In BotFather, set /setprivacy to disabled so the bot ignores public group messages.
- **No public channels:** Never add your OpenClaw bot to public groups — untrusted parties could issue commands or craft prompt injection attacks.
- **Webhook rotation:** Rotate messaging platform webhook secrets on the same 90-day schedule as API keys.
- **Prompt injection vigilance:** Treat all content from external sources — emails, websites, documents — as potentially adversarial. Configure HITL (Human-in-the-Loop) approval for actions triggered by external content.

10. On-Premises & Bare Metal

- **Network isolation:** Static local IP; router-level firewall blocking all inbound except from trusted subnets.
- **Remote access:** Tailscale or self-hosted WireGuard. Never open a router port for the OpenClaw Gateway.

- **Disk encryption:** Enable LUKS full-disk encryption to protect data if hardware is physically stolen.
- **UPS:** Uninterruptible power supply to prevent data corruption from sudden power loss.
- **Log management:** Configure logrotate to prevent disk exhaustion over time.

11. Ongoing Maintenance

Keeping OpenClaw Updated

```
cd ~/openclaw
docker compose pull
docker compose up -d
docker compose exec openclaw openclaw --version
```

Backups

```
tar -czf openclaw-backup-$(date +%Y%m%d).tar.gz ~/.openclaw
aws s3 cp openclaw-backup-*.tar.gz s3://your-bucket/backups/
```

The Big Red Button

```
# Immediately stop the agent
docker compose down openclaw

# Monitor logs for anomalies
docker compose logs openclaw | grep -iE
'error|unauthorized|fail|exfil|curl|wget'
```

PART 2: SECURITY IN DEPTH

CVEs, UK Regulation, Advanced Hardening & Tools

12. Known CVEs & Active Security Alerts

OpenClaw exploded in popularity in January 2026 — attracting attackers just as fast. Scanning for exposed instances began the same day the project hit Hacker News. Security researchers confirmed over 42,000 instances were publicly exposed at peak. Monitor github.com/openclaw/openclaw/security for new advisories.

CVE ID	CVSS	Status	Component	Summary
CVE-2026-25253	8.8	PATCHED	Control UI / WebSocket	1-click RCE via cross-site WebSocket hijacking. The Control UI accepted a gateway URL without validation, auto-connecting and sending the auth token to attackers. Fixed in v2026.1.29.
CVE-2026-21636	TBC	PATCHED	Node.js Permission Model	Permission model bypass in Node.js older than v22.12.0. Upgrade to Node.js 22.12.0 LTS immediately.
CVE-2025-59466	TBC	PATCHED	Node.js async_hooks	Denial-of-service vulnerability in async_hooks. Resolved by Node.js 22.12.0 LTS.
Moltbook DB Leak	N/A	ACTIVE	Adjacent Platform	1.5M API tokens exposed due to missing Row-Level Security on Moltbook. Rotate any credentials used there.
ClawHub Malware	N/A	ACTIVE	Third-party Skills	Malicious skills distributing macOS infostealer malware on the ClawHub repository. Audit all third-party skills (see Section 17).
Infostealer	N/A	ACTIVE	Host / Config Files	Active infostealer campaigns targeting OpenClaw config files, gateway tokens and API keys on compromised hosts (February 2026).

- **CVE-2026-25253 full technical write-up:** DepthFirst researcher analysisdepthfirst.com
- **Belgium CCB emergency advisory:** Government advisory — highest priority patchingccb.belgium.be
- **Adversa AI deep-dive:** Most comprehensive independent security analysisadversa.ai

13. UK Regulatory Framework for Agentic AI

OpenClaw is an agentic AI system — one that takes autonomous actions with real-world consequences. Deploying it, particularly where personal data is involved, triggers obligations under UK law. The good news: compliance with these frameworks and secure deployment go hand in hand. Do it right once, and you're covered on both fronts.

13.1 — NCSC: Guidelines for Secure AI System Development

Jointly published by the UK National Cyber Security Centre, US CISA, and 21 international agencies, these guidelines are the closest thing to a global standard for AI system security. They insist that security must be designed in from the start — not bolted on after a breach.

- **Secure design:** Minimise attack surface; give the agent only the permissions it needs for defined tasks
- **Secure development:** Vet supply chain for skills and plugins; scan for credential leakage
- **Secure deployment:** Harden runtime; enforce network egress; bind to loopback only; use Docker
- **Secure operation:** Maintain logs, rotate credentials, patch promptly, retire cleanly

NCSC guidelines: ncsc.gov.uk/collection/guidelines-secure-ai-system-development

13.2 — DSIT: AI Cyber Security Code of Practice (January 2025)

Published by the Department for Science, Innovation and Technology and now adopted globally as ETSI Technical Standard TS 104 223, this voluntary Code establishes 13 security principles for AI systems. It explicitly references ICO data protection guidance — so if your OpenClaw instance handles personal data, both sets of obligations apply.

- DSIT Codegov.uk/government/publications/ai-cyber-security-code-of-practice

- Implementation Guide (PDF)assets.publishing.service.gov.uk — Implementation Guide AI Cyber Security Code

13.3 — ICO: Tech Futures Report on Agentic AI (January 2026)

The ICO's January 2026 report is the most significant regulatory signal on agentic AI from a UK data protection perspective. Five points every OpenClaw operator needs to understand:

6. **You remain responsible:** Agent autonomy does not transfer data protection liability. You are the data controller.
7. **Least privilege is a legal requirement:** UK GDPR Article 5(1)(c) requires limiting agent access to only what is strictly necessary.
8. **Define your purposes narrowly:** Broad 'do anything' agent permissions create purpose-creep compliance risk.
9. **Significant decisions need human override:** Financial transactions, legal notices — any action with legal effect needs a human check.
10. **Proactive controls are expected:** The ICO expects controls to be in place before deployment, not in response to incidents.

ICO Agentic AI report: ico.org.uk

13.4 — Cyber Security and Resilience Bill

This forthcoming UK legislation (currently progressing through Parliament) will significantly expand the NIS Regulations 2018. While it primarily targets Managed Service Providers and critical national infrastructure, its reach extends directly to SMEs through supply chain obligations.

If you supply services to a regulated organisation — an NHS trust, an energy company, a major MSP — expect your cyber resilience to become a contractual requirement. MSPs facing fines of up to £17 million or 4% of worldwide turnover will contractually push security requirements down their supply chains. Cyber Essentials certification will increasingly be the minimum bar.

✔ Get Ahead of the Curve

The Cyber Security and Resilience Bill is an opportunity, not just a burden. SMEs that achieve Cyber Essentials or Cyber Essentials Plus now are positioning themselves to win and retain contracts from regulated clients who will soon be required to ensure their suppliers are secure.

UK Official Regulatory References	
NCSC AI Guidelines	Secure AI System Development — joint international standardhttps://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development
DSIT Code of Practice	AI Cyber Security Code of Practice — now ETSI TS 104

	223https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice
ICO Tech Futures	Agentic AI & Data Protection, January 2026https://ico.org.uk
CS&R Bill	Cyber Security and Resilience Bill — GOV.UKhttps://www.gov.uk/government/collections/cyber-security-and-resilience-bill
Cyber Essentials	Government-backed certification for SMEshttps://www.ncsc.gov.uk/cyberessentials/overview

14. Secure by Default: The Applied Framework

The following framework distils NCSC, DSIT, and ICO requirements into concrete OpenClaw architecture decisions. This is the minimum viable security posture for any deployment involving personal data, organisational systems, or external API access.

14.1 — Least Privilege (The Most Important Principle)

Treat OpenClaw as a privileged insider that could be socially engineered via prompt injection. Never grant it more access than is required for a specific, defined task.

- **Identity:** Run under a dedicated, non-privileged system user. Never root.
- **Tokens:** Fine-grained, scoped API keys (e.g., GitHub read-only for one repo) not master credentials.
- **Isolation:** Ephemeral Docker container. Consider wiping workspace after each session for sensitive use cases.
- **Persistence:** System directories mounted read-only. Write access only to a specific /output/ path.

14.2 — The Positive Security Model

Define only what the agent is permitted to do — don't try to block bad things. This is harder to bypass than blocklists and aligns with the NCSC's secure-by-design principle.

- **Command allow-list:** Define permitted binaries (git, python, ls). Block rm, chmod, sudo at OS level.
- **Network egress control:** Firewall restricts outbound to known providers only (api.openai.com etc.). All other egress blocked.
- **Human-in-the-Loop:** Manual approval required for destructive actions: deleting files, sending messages, financial transactions.

14.3 — UK Compliance Summary

Security Layer	Requirement	OpenClaw Implementation
Environment	Sandboxed, isolated from host OS	Docker: --read-only, --cap-drop=ALL, non-root user
Data Access	No access to PII or credentials without scope	Fine-grained keys; no .env files in agent workspace
Audit Logs	Immutable, stored externally to agent	/var/log/openclaw (root-owned, agent cannot write)
Ethics / ICO	Legal-effect decisions need human override	HITL configured for destructive/significant actions
Supply Chain	Third-party components vetted before use	Cisco Skill Scanner; cryptographic signing in production
Incident Response	Events detectable, logged, reportable	Anomaly monitoring; documented IR procedure

15. Immutable Logging & Monitoring

DSIT Code and NCSC guidelines both require audit logs to be immutable and stored externally to the agent. The OpenClaw process must not be able to modify or delete its own logs.

```
# Protected log directory — root-owned, agent cannot write
sudo mkdir -p /var/log/openclaw
sudo chown root:root /var/log/openclaw
sudo chmod 755 /var/log/openclaw

# Pipe container logs to this protected path
docker compose logs --follow 2>&1 | sudo tee /var/log/openclaw/agent.log
```

```
# /etc/logrotate.d/openclaw
/var/log/openclaw/*.log {
    daily
    rotate 30
    compress
    missingok
}
```

16. Security Products & Tooling

A growing ecosystem of security tools has emerged specifically for OpenClaw. Use these alongside the hardening steps above — they are not a replacement for secure configuration.

16.1 — SecureClaw (Adversa AI)

Open-source OWASP-aligned security suite. Acts as a gateway plugin enforcing hardening rules, and as a behavioural skill giving the agent self-awareness of security risks including prompt injection. Adversa AI also published the most thorough independent OpenClaw security analysis available.

adversa.ai

16.2 — ClawShield

Security preflight tool that audits your clawdbot.json for risky configurations before deployment. Specifically checks for 0.0.0.0 binding, public Gateway exposure, and overly permissive skill settings.

16.3 — ClawdStrike (Backbay Labs)

Runtime enforcement library with 7 built-in guards at near-zero latency: path access controls, secrets detection, and network egress enforcement. Ideal for production deployments.

16.4 — Lasso Security: Intent Deputy

Real-time intent monitoring that flags when the agent diverges from its stated purpose — a strong indicator of prompt injection in progress. Commercial-grade, enterprise-ready.

lasso.security

16.5 — VirusTotal Integration

Integrate VirusTotal as an OpenClaw skill via MCP for pre-execution file scanning (70+ engines), URL reputation checking before browsing, and automated triage of suspicious files.

virustotal.com

17. Skill Security & Supply Chain

OpenClaw Skills are local code packages — the primary vector for supply-chain attacks. The ClawHub community repository has already distributed macOS infostealer malware. Apply this checklist for every third-party skill.

17.1 — Cisco Skill Scanner

Cisco's AI Defense team released Skill Scanner (GitHub) for auditing skill packages. It checks metadata for hidden instructions, flags curl/wget/chmod in skill logic, and detects silent exfiltration calls.

github.com/cisco

17.2 — Manual Verification Checklist

11. **SOUL.md integrity:** Any skill that overwrites SOUL.md (the agent's core safety file) is reprogramming the agent's safety rules. Treat this as a critical security incident.
12. **Cryptographic signing:** In production, require signed skills. Reject unsigned packages from unknown authors.
13. **Network isolation test:** Load the skill in a Docker container with no internet access. If it fails without internet but shouldn't need it — reject it.
14. **Read the source:** Obfuscated logic, base64-encoded payloads, or unexplained network calls are immediate rejection criteria.

18. Incident Response

If you suspect a compromise — unexpected config changes, API spend spikes, unknown connections in logs, or suspicious commands — act immediately.

15. **Isolate:** docker compose down openclaw — stop the agent first, investigate second.
16. **Preserve evidence:** Copy all logs from /var/log/openclaw before making any system changes.
17. **Rotate everything:** Gateway token, all API keys, all messaging platform tokens.
18. **Audit:** Review 48 hours of logs. Identify the attack vector.

19. **Patch and rebuild:** Apply all outstanding patches. Consider a clean rebuild if the host OS may be compromised.
20. **Report if required:** Personal data involved? UK GDPR requires ICO notification within 72 hours.ico.org.uk/make-a-report

19. Pre-Live Security Checklist

Tick every item before putting any OpenClaw instance into use. Section references link back to this guide.

☐	OpenClaw version is 2026.1.29 or later (CVE-2026-25253 patched) → *§12*
☐	Node.js version is 22.12.0 LTS or later (patches CVE-2025-59466 and CVE-2026-21636) → *§4.4*
☐	Non-root system user created for running OpenClaw → *§4.1*
☐	Root SSH login disabled; password authentication disabled; key-only access → *§4.2*
☐	Firewall configured: port 22 inbound only; Gateway port 3000 never public → *§4.3*
☐	Automatic security updates enabled → *§4.5*
☐	Docker running with --read-only and --cap-drop=ALL → *§5.2*
☐	Gateway bound to 127.0.0.1 (loopback) — NOT 0.0.0.0 — confirmed in compose file → *§5.3*
☐	Remote access via SSH tunnel or Tailscale only — no public reverse proxy → *§7.2*
☐	Gateway Token in a password manager — not in shell history or a text file → *§7.1*
☐	All API keys stored as environment variables; detect-secrets baseline created → *§8*
☐	Spend limits set on every provider dashboard (OpenAI, Anthropic, Google, etc.) → *§8*
☐	Messaging channel DM pairing enabled; bot not in any public channels or groups → *§9*
☐	Immutable log directory created (/var/log/openclaw); logrotate configured → *§15*
☐	All third-party skills scanned with Cisco Skill Scanner before installation → *§17.1*
☐	SOUL.md integrity verified; no skill permitted to overwrite it → *§17.2*
☐	'Big Red Button' kill command tested: docker compose down openclaw → *§11*
☐	Backup created, transferred offsite, and restore-tested → *§11*
☐	Incident response procedure documented → *§18*
☐	ICO data protection obligations assessed for any personal data processed → *§13.3*
☐	DSIT Code of Practice reviewed; applicable principles noted → *§13.2*
☐	Cyber Security and Resilience Bill supply chain implications assessed → *§13.4*

20. References & Official Resources

OpenClaw Official

- **Official Documentation:** https://docs.openclaw.ai
- **VPS Hosting Overview:** https://docs.openclaw.ai/vps
- **Security Advisories:** https://github.com/openclaw/openclaw/security
- **GitHub Repository:** https://github.com/openclaw/openclaw

CVE References

- **CVE-2026-25253 — NVD:** https://nvd.nist.gov/vuln/detail/CVE-2026-25253 — CVSS 8.8, 1-click RCE via WebSocket hijacking
- **DepthFirst Technical Analysis:** https://depthfirst.com/post/1-click-rce-to-steal-your-moltbot-data-and-keys
- **Belgium CCB Advisory:** https://ccb.belgium.be/advisories/warning-critical-vulnerability-openclaw-allows-1-click-remote-code-execution-when
- **Adversa AI Security Deep-Dive:** https://adversa.ai/blog/openclaw-security-101-vulnerabilities-hardening-2026/

UK Official Guidance

- **NCSC: Guidelines for Secure AI System Development:** https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development
- **DSIT: AI Cyber Security Code of Practice:** https://www.gov.uk/government/publications/ai-cyber-security-code-of-practice
- **DSIT: Implementation Guide (PDF):** https://assets.publishing.service.gov.uk/media/679cae441d14e76535afb630/Implementation_Guide_for_the_AI_Cyber_Security_Code_of_Practice.pdf
- **ICO: Tech Futures Agentic AI Report:** https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2026/01/ai-ll-get-that/
- **Cyber Security and Resilience Bill:** https://www.gov.uk/government/collections/cyber-security-and-resilience-bill
- **Cyber Essentials — NCSC:** https://www.ncsc.gov.uk/cyberessentials/overview
- **ICO: Report an incident:** https://ico.org.uk/make-a-report/

Deployment Guides

- **Hostinger: How to Set Up OpenClaw:** https://www.hostinger.com/tutorials/how-to-set-up-openclaw
- **DigitalOcean: How to Run OpenClaw:** https://www.digitalocean.com/community/tutorials/how-to-run-openclaw
- **Cognio Labs: Oracle Cloud Always Free:** https://cognio.so/clawdbot/self-hosting

Security Tooling

- **Tailscale — Secure Private Networking:** https://tailscale.com
- **VirusTotal:** https://www.virustotal.com
- **Adversa AI / SecureClaw:** https://adversa.ai
- **Lasso Security: Intent Deputy:** https://www.lasso.security

BMS CYBER DEFENCE

Expert Cyber Security for Small & Medium Businesses

AI is moving fast. UK regulation is catching up. Your security should be ahead of both.

Whether you're exploring OpenClaw for the first time, need help getting compliant with NCSC guidance and the incoming Cyber Security and Resilience Bill, or want Cyber Essentials certification to unlock new contracts — BMS Cyber Defence is ready to help. We work specifically with SMEs because we believe great security shouldn't require an enterprise budget.

bmscyberdefence.co.uk

admin@bmscyberdefence.co.uk

Follow us on LinkedIn, Facebook & Twitter/X | @BMSCyberDefence

AI Assistance Disclosure

This document was generated with the assistance of Anthropic Claude. Jason Lewis of BMS Cyber Defence provided the conceptual framework, iterative prompts, and final editorial oversight to ensure accuracy, technical rigour, and stylistic alignment with BMS Cyber Defence standards. Jason Lewis and BMS Cyber Defence remain solely responsible for the integrity of the content presented in this guide.

This guide is provided free of charge as a public resource. Share freely. | Authored by Jason Lewis, BMS Cyber Defence | Version 1.0 | February 2026