Zero to Production in Neue Haas Grotesk Display

A Field Guide to Deploying AI Agents


Friday night. The laptop closes.
Your competitor-analysis agent worked perfectly all week.
It scraped data, updated the database, and generated insights.
You close your laptop and go home for the weekend.
The agent dies.
Three days of missing data.
A board meeting in 12 hours.

Your agent doesn’t live on your laptop. It lives on a Linux server with no screen, no mouse, no desktop. Just a command line and the truth.

The Window, The Voice, and The Tunnel

The Terminal (The Glass Window)
You can look through it, but you cannot reach inside.

The Shell (The Voice)
The program on the other side that listens, acts, and reports back.

SSH (The Secure Tunnel)
A secure, encrypted phone call to a computer with no screen.

Exploring the Building

Everything starts at / (The Front Door).

`/home` — **The Apartments:** User files and agent code.

`/var` — **The Mailroom:** Where logs and activity accumulate.

`/etc` — **The Manager's Office:** Rules and system configurations.

`/opt` — **The Storage Units:** Optional software installed by hand.

`/tmp` — **The Lobby Whiteboard:** Notes erased every night.

You Don’t Need to Speak Linux. Just Read It.

Every command is just a sentence.

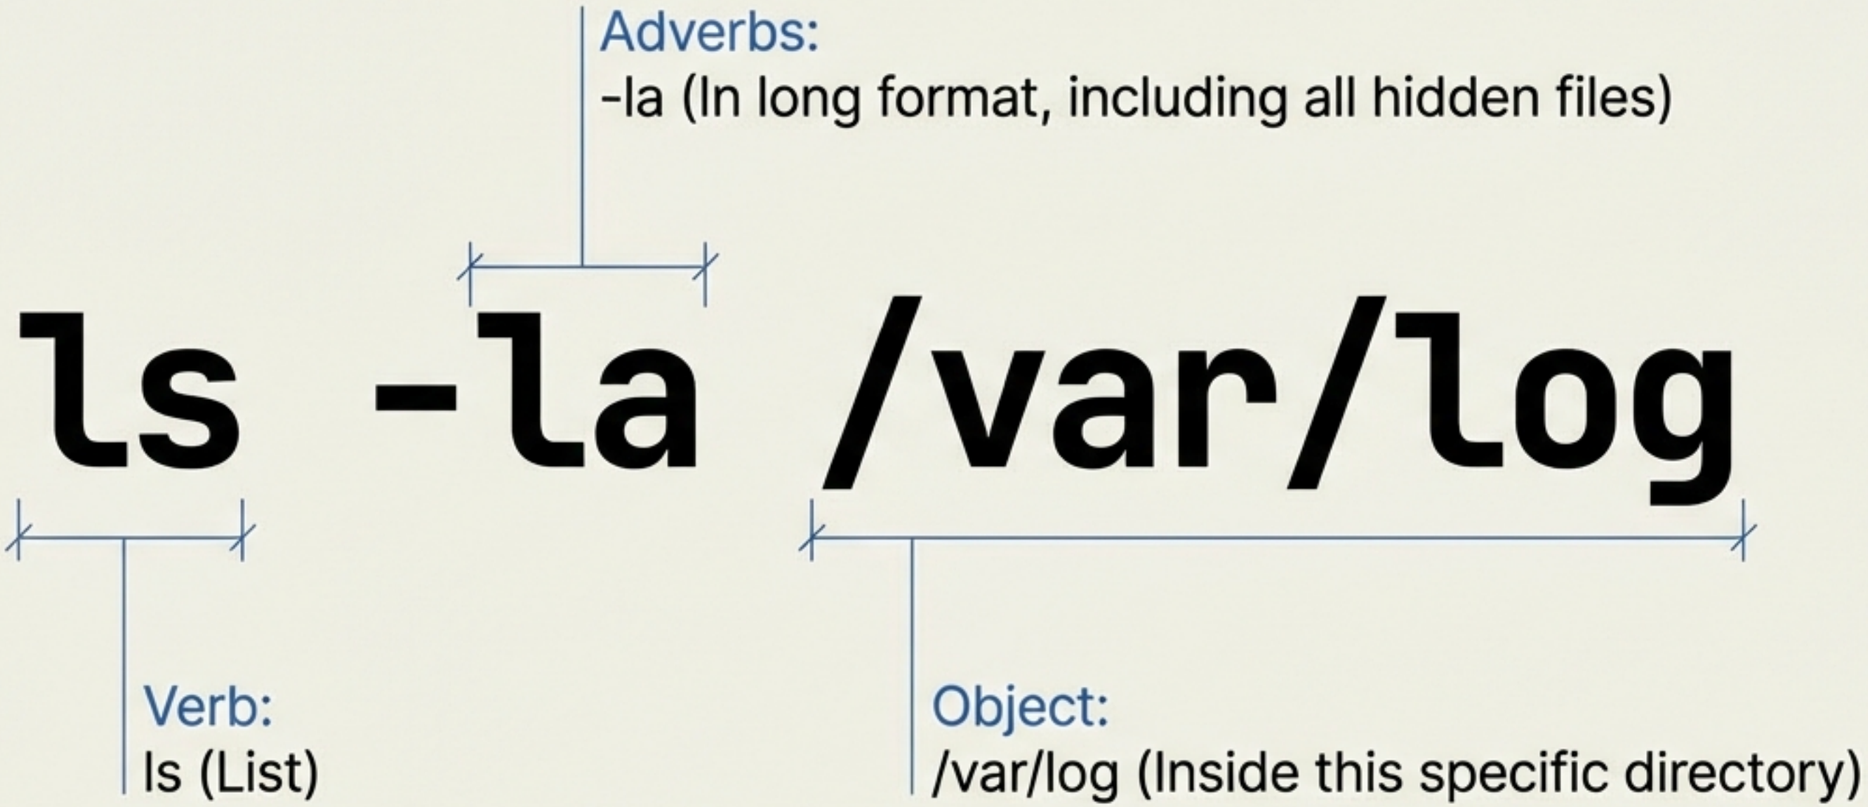

The Keycard System

Decoding `drwxr-xr-x`:

Owner (rwx): The Master Keycard.
Can read, write, and execute.

Group (r-x): The Department Keycard.
Can read and execute, but not modify.

Everyone (r-x): The Visitor Badge.
Can only look.

A Proper Home

Agents dumped in random directories fail because of bad organization.
Professional deployments isolate the components:

/src: Code
/config: Settings
/data: Databases and cache
/logs: Persistent records

Extracting Secrets & Recording History

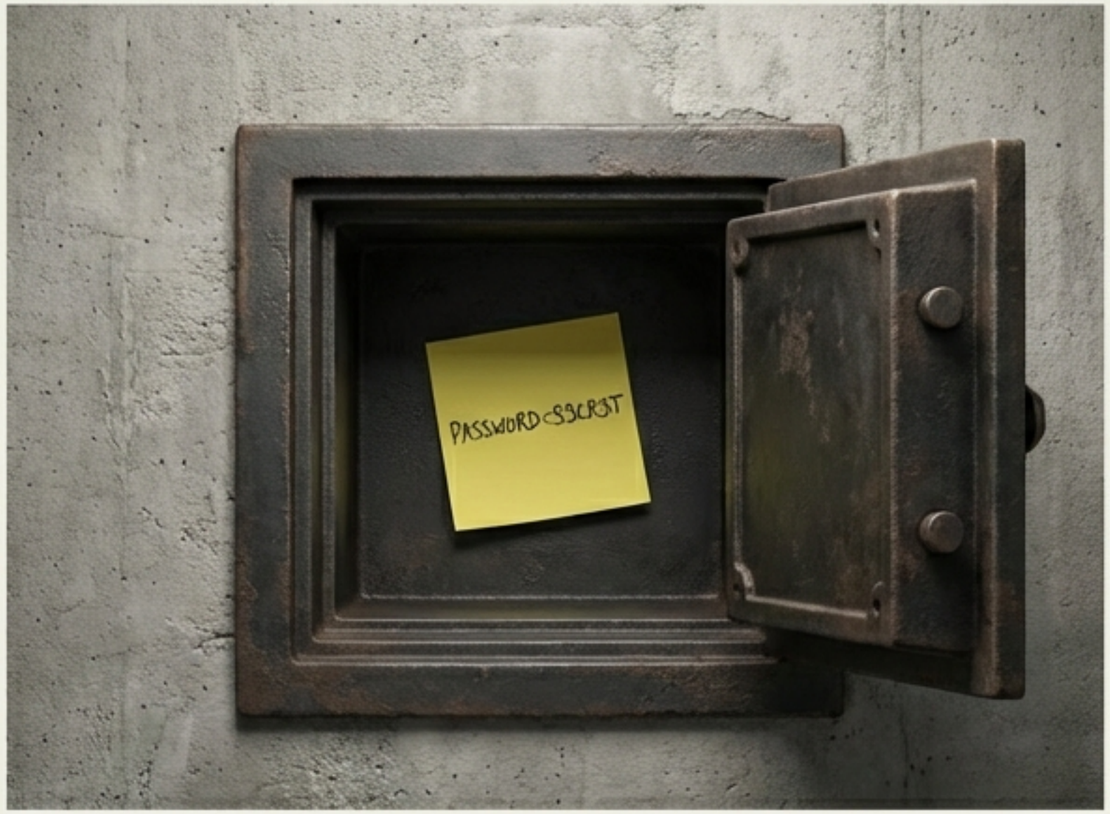

The .env File: Hardcoded passwords are a ticking time bomb. Extract secrets to a .env file and lock it with `chmod 600` (Owner read/write only).

Log Persistence: If a tree falls in a closed terminal, no one hears it. Use redirection and `tee` to send output to a file that survives when you disconnect.

The False Victory

The house is organized. Secrets are safe. Logs are writing.

You start the agent. It works perfectly. You close your laptop to go to bed.

The agent dies. Every. Single. Time.

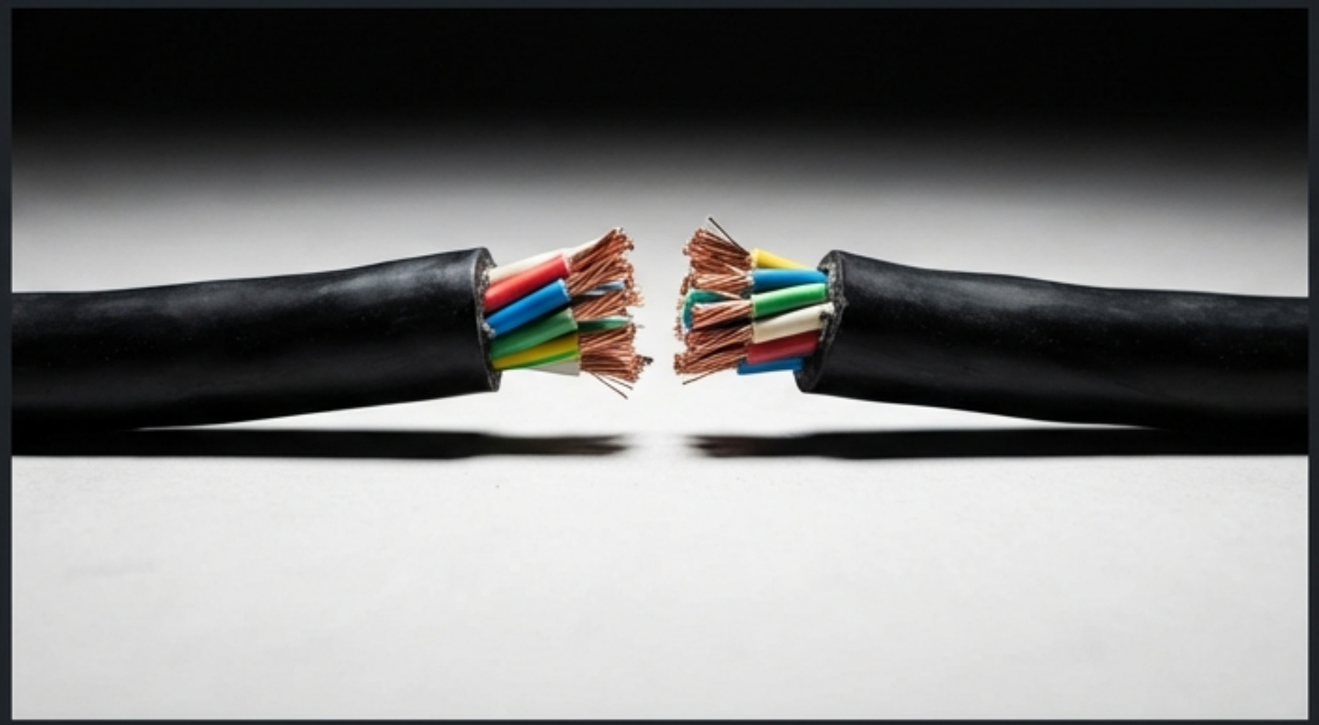

Process vs. Service

A Process is a phone call. When you close your terminal, the session ends, and the server hangs up on your agent.

A Service is a security guard. The building doesn't care if a shift changes. It ensures someone is always at the desk.

Enter systemd: The Building Manager

One file answers five questions:

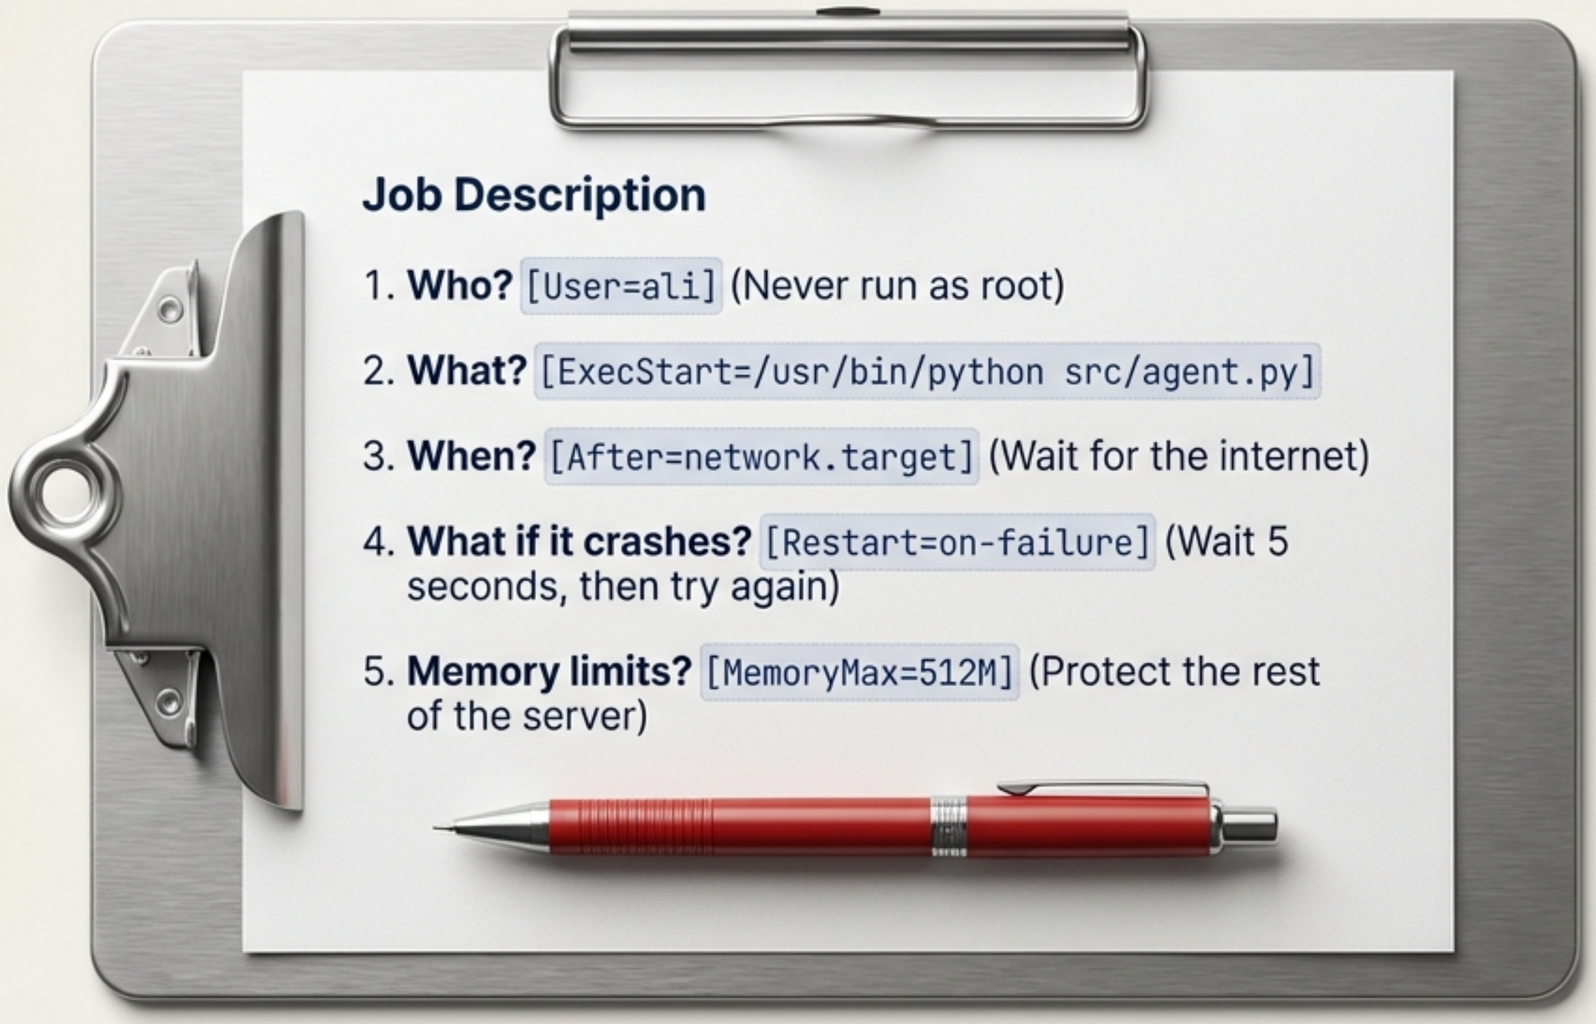

Unkillable.

Close the terminal.
Disconnect from the internet. Reconnect.
The agent is still alive.


47,000
47,000 Reasons to Care
Your agent is unkillable. But Dev checks the authentication logs.
Automated bots have attempted to guess your password 47,000 times in the last three weeks. And your agent is running as root. If one bot guesses correctly, they own the entire server.

Defending the Fortress

- **Badges, not Master Keys:** Create a dedicated agent user. Give the delivery driver the mailbox key, not the front door key.
- **Fingerprints, not Passwords:** Passwords can be brute-forced. SSH Key pairs (id_ed25519) are cryptographic fingerprints. Disable password login entirely.

The 2 AM Crisis

The agent is alive, but the report is completely empty.

The instinct: Restart everything.

The reality: Every bad debugger has one move: restart. Every good debugger has a system. Restarting masks the symptom; it doesn't fix the root cause.

The LNPS Method

1. Logs (journalctl): What did the agent say happened? (No error

2. Network (curl / ping): Can the agent reach the outside world?

3. Process (systemctl): Is it stuck in a loop or sleeping?

4. System (df -h / top): Are we out of memory or disk space?

The Real Culprit

The logs showed the agent successfully queried the database and got zero rows back.

Checking the process revealed the database never started after a reboot.

A blind restart of the agent would have done nothing. Two commands (enable and start on the database) fixed everything.

Turning Knowledge into Execution

The client wants a social media sentiment agent. Will it take another three days of pain?

No. You write a Deployment Spec: The checklist, the executor, and the auditor, all in one file.

The 15-Minute Run

0:00 - Server Access: SSH Key verified.

0:01 - Directory Structure: /opt/agents/ created, .env secured.

0:04 - Application Setup: Dependencies installed.

0:09 - Service Config: systemd unit file applied.

0:11 - Security Checklist: Root disabled, permissions locked.

0:13 - Verification: LNPS checks passed.

You Are In Production

Three days of discovery versus fifteen minutes of execution.

You are no longer guessing. You aren’t just writing code on a laptop anymore. You are **deploying** unkillable, secure agents into the real world.