FOSDEM (Jan 31, 2026)

Lima v2.0: expanding the focus to hardening AI

Akihiro Suda, NTT

https://lima-vm.io/

© NTT, Inc. 2026

What is Lima?

- **Li**nux virtual **ma**chines optimized for running containers and AI agents
- Automatic host filesystem sharing
- Automatic port forwarding
- Built-in integration for several container engines
  - containerd (default), Docker, Podman, Kubernetes, and Apptainer

```
$ brew install lima
$ limactl start
$ lima nerdctl run -p 80:80 nginx
```

© NTT, Inc. 2026

Similar projects

- **WSL2**
  - Windows host only
- **Vagrant**
  - Proprietary
  - No automatic port forwarding etc.
- **Docker Machine**
  - Docker only
  - Abandoned
- **Docker Desktop**
  - Docker only
  - Proprietary

© NTT, Inc. 2026

The origin and the current status

- The project began in May 2021, for promoting containerd including nerdctl to Mac users (“containerd Machine”)
- Through the growth of the community, the scope has expanded
  - **Additional container engines**
    - Docker, Podman, Kubernetes, Apptainer
  - **Non-container workloads**
    - Sandboxing AI coding agents
    - Running non-Ubuntu OS on GitHub Actions
  - **Non-macOS hosts**
    - Linux, Windows, NetBSD, DragonflyBSD

© NTT, Inc. 2026

Third-party FLOSS projects based on Lima

- **Colima** (https://colima.run)
  - Alternative CLI for Lima, with Docker as the default engine

- **Rancher Desktop** (https://rancherdesktop.io)
  - Lima + k3s + GUI

- **Finch** (https://runfinch.com)
  - AWS product, for local development with AWS Serverless Application Model etc.

© NTT, Inc. 2026

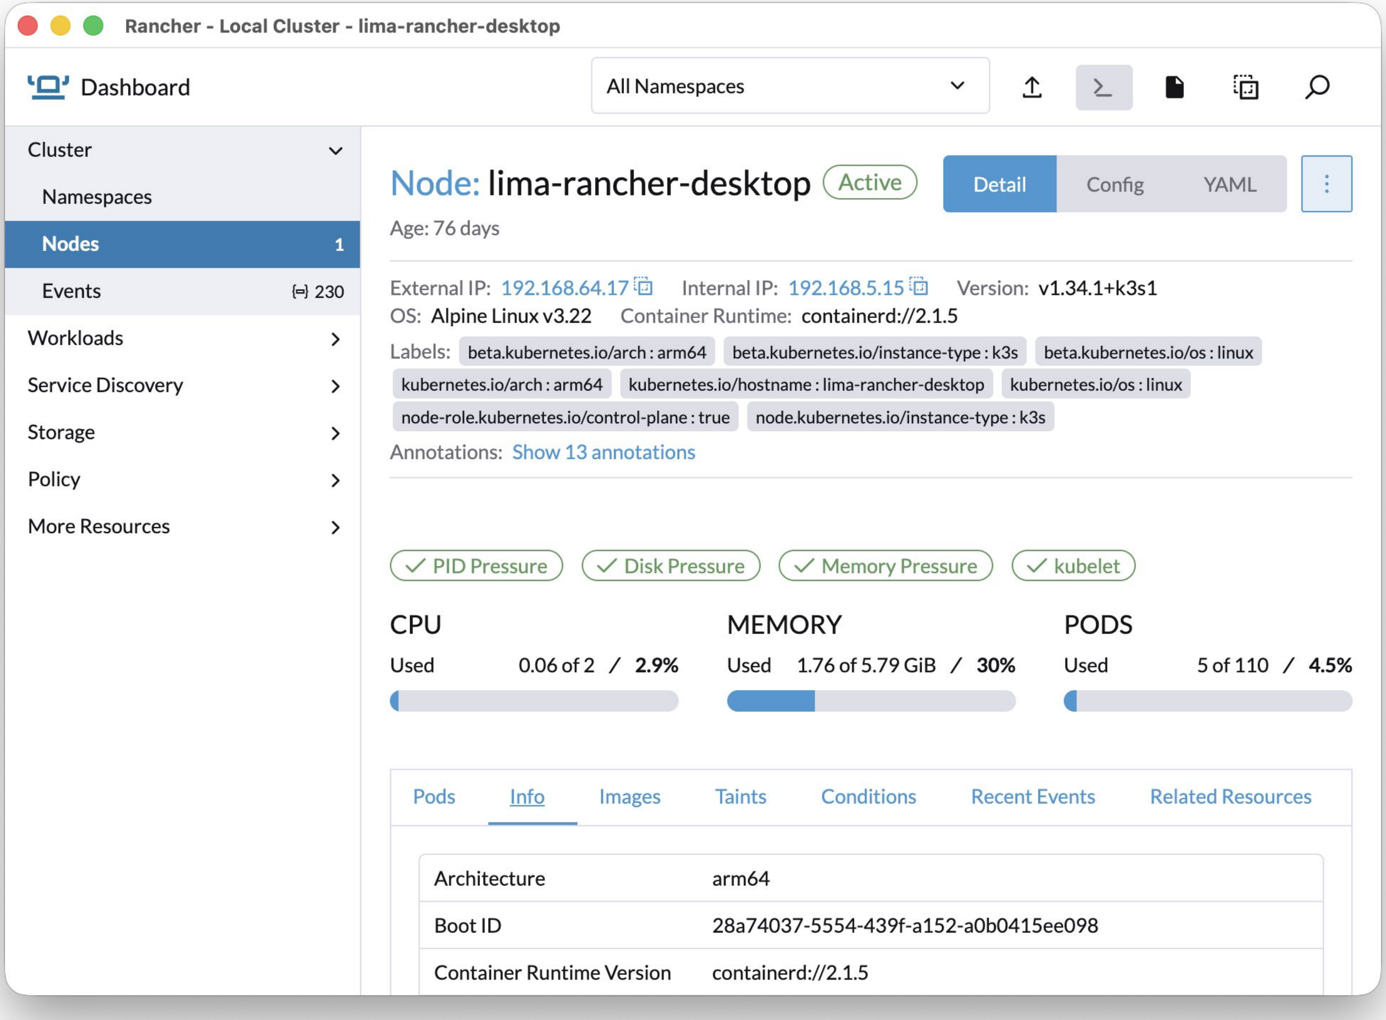

Rancher - Local Cluster - lima-rancher-desktop

Dashboard

All Namespaces

- Cluster
  - Namespaces
  - Nodes 1
  - Events 230
- Workloads
- Service Discovery
- Storage
- Policy
- More Resources

Node: lima-rancher-desktop Active

Detail | Config | YAML

Age: 76 days

External IP: 192.168.64.17 Internal IP: 192.168.5.15 Version: v1.34.1+k3s1

OS: Alpine Linux v3.22 Container Runtime: containerd://2.1.5

Labels: beta.kubernetes.io/arch : arm64 beta.kubernetes.io/instance-type : k3s beta.kubernetes.io/os : linux kubernetes.io/arch : arm64 kubernetes.io/hostname : lima-rancher-desktop kubernetes.io/os : linux node-role.kubernetes.io/control-plane : true node.kubernetes.io/instance-type : k3s

Annotations: Show 13 annotations

✓ PID Pressure ✓ Disk Pressure ✓ Memory Pressure ✓ kubelet

CPU
Used 0.06 of 2 / **2.9%**

MEMORY
Used 1.76 of 5.79 GiB / **30%**

PODS
Used 5 of 110 / **4.5%**

Pods | Info | Images | Taints | Conditions | Recent Events | Related Resources

Architecture	arm64
Boot ID	28a74037-5554-439f-a152-a0b0415ee098
Container Runtime Version	containerd://2.1.5

Third-party FLOSS projects based on Lima

- **Lima GUI** (https://github.com/afbjorklund/lima-gui)
  - › Qt-based GUI (→)

- **Podman Desktop**
  (https://podman-desktop.io/docs/lima)
  - › Supports managing Lima instances as well as native Podman Machine instances

- **And more!**

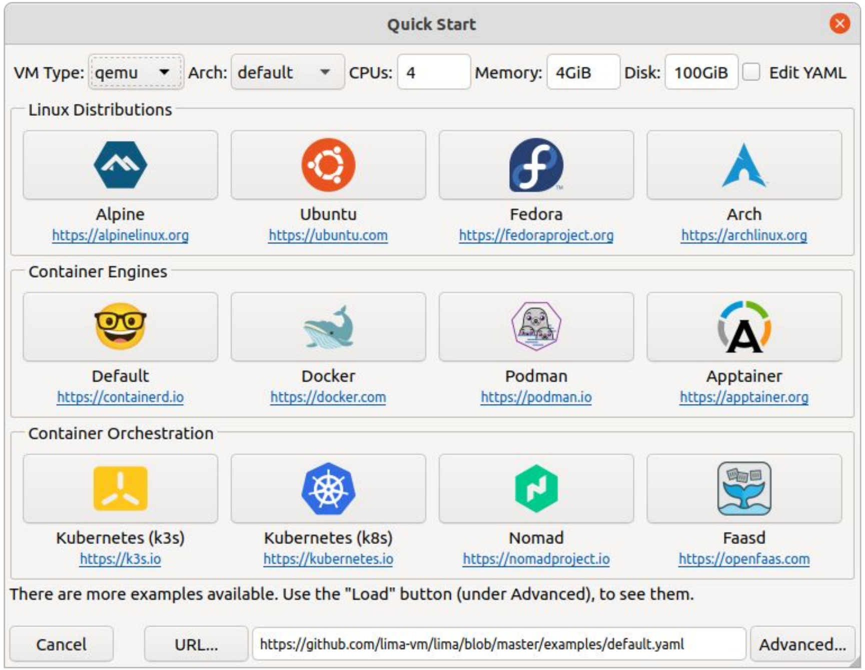

Quick Start

VM Type: qemu Arch: default CPUs: 4 Memory: 4GiB Disk: 100GiB ☐ Edit YAML

Linux Distributions

Alpine	Ubuntu	Fedora	Arch
https://alpinelinux.org	https://ubuntu.com	https://fedoraproject.org	https://archlinux.org

Container Engines

Default	Docker	Podman	Apptainer
https://containerd.io	https://docker.com	https://podman.io	https://apptainer.org

Container Orchestration

Kubernetes (k3s)	Kubernetes (k8s)	Nomad	Faasd
https://k3s.io	https://kubernetes.io	https://nomadproject.io	https://openfaas.com

There are more examples available. Use the "Load" button (under Advanced), to see them.

https://github.com/lima-vm/lima/blob/master/examples/default.yaml

How it works

© NTT, Inc. 2026

Architecture

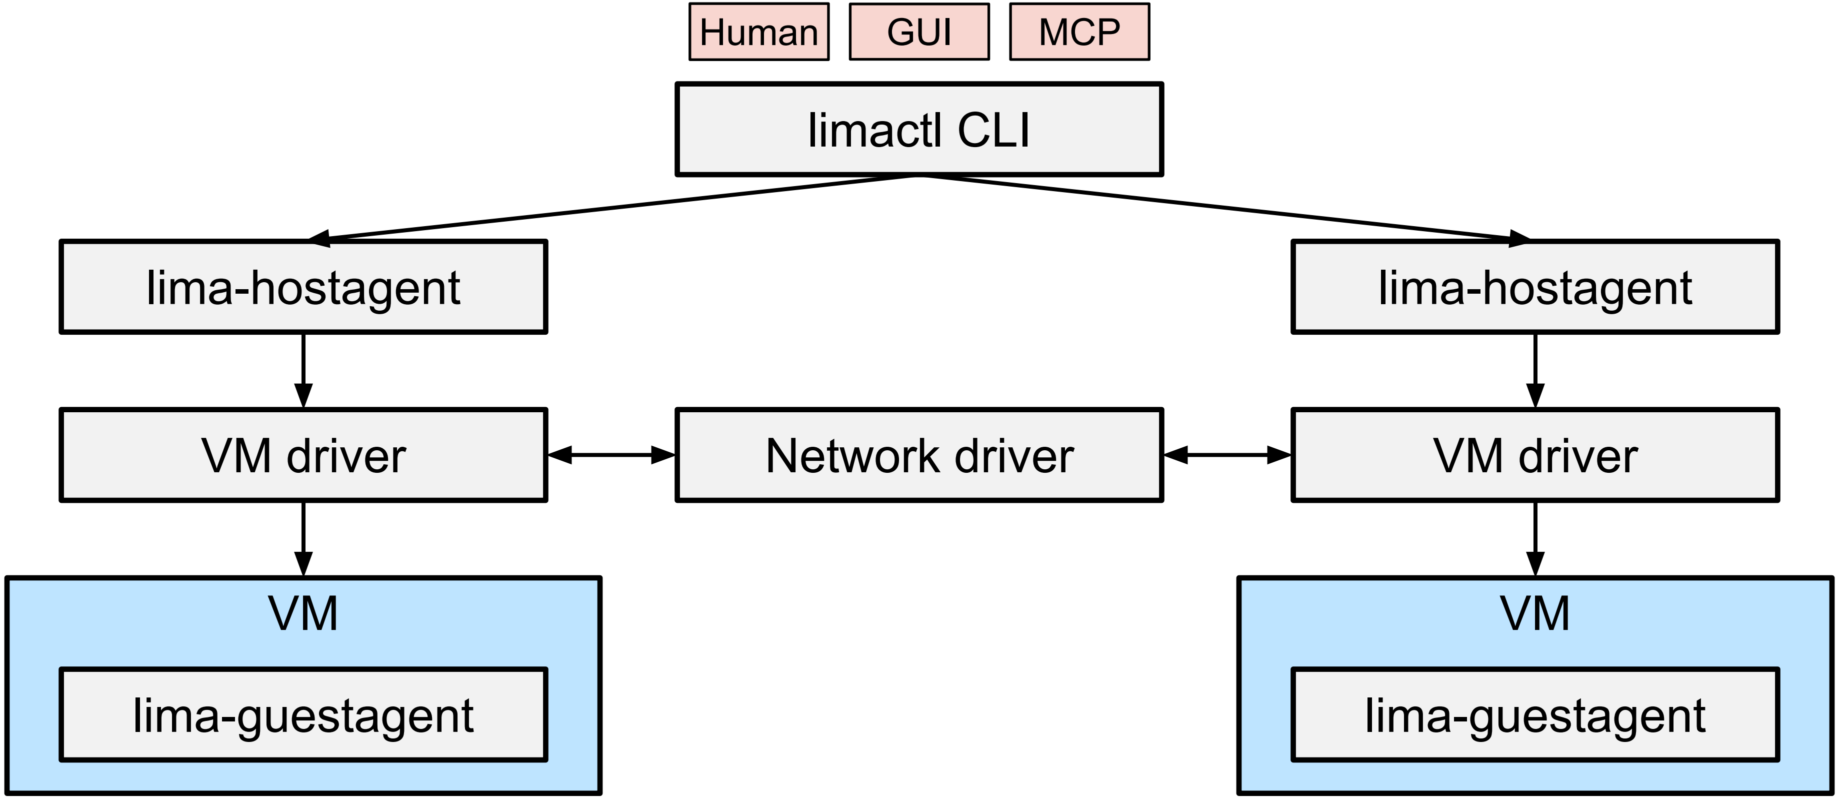

© NTT, Inc. 2026

Architecture

- **VM drivers**
  - QEMU
  - Virtualization.framework (vz) [macOS only]
  - WSL2 [Windows only]
  - krunkit
    - Supports GPU acceleration on macOS
  - gRPC plugins

- **Intel-on-ARM binary executors**
  - QEMU User Mode
  - Rosetta 2 [macOS only]

© NTT, Inc. 2026

Architecture

- **Filesystem sharing**
  - virtiofs (vz, krunkit), virtio-9p (QEMU), reverse-sshfs

- **Network drivers**
  - User mode networking (default)
  - socket_vmnet (for accessing VM by IP, with sudo)
  - vzNAT (for accessing VM by IP, with vz)

- **Port forwarding**
  - `NETLINK_SOCK_DIAG` watcher based on eBPF (for most ports)
  - Kubernetes service watcher (for Kubernetes service ports)

© NTT, Inc. 2026

Built-in templates

- **Distros**
  - almalinux, alpine, archlinux, centos-stream, debian, opensuse, oraclelinux, rocky, ubuntu, …
- **Container engines**
  - apptainer, docker, docker-rootful, podman, podman-rootful, …
- **Container orchestration**
  - faasd, k0s, k3s, k8s, u7s (Usernetes), …

```
$ limactl start --name=default template://docker
```

© NTT, Inc. 2026

Recent updates

© NTT, Inc. 2026

Promoted to CNCF Incubating Project 🎉

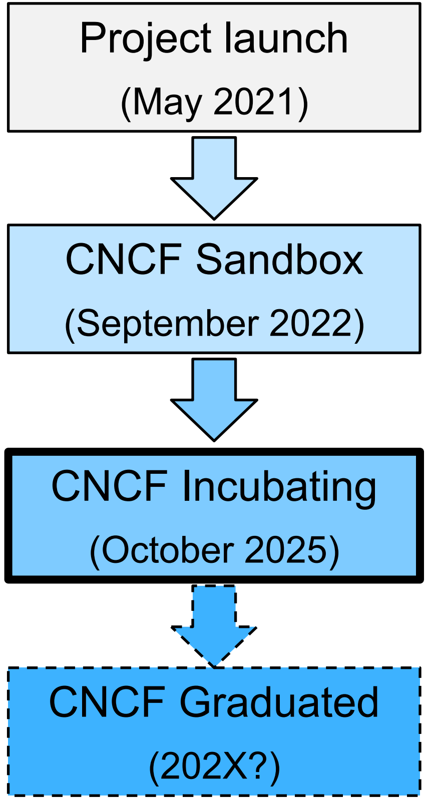

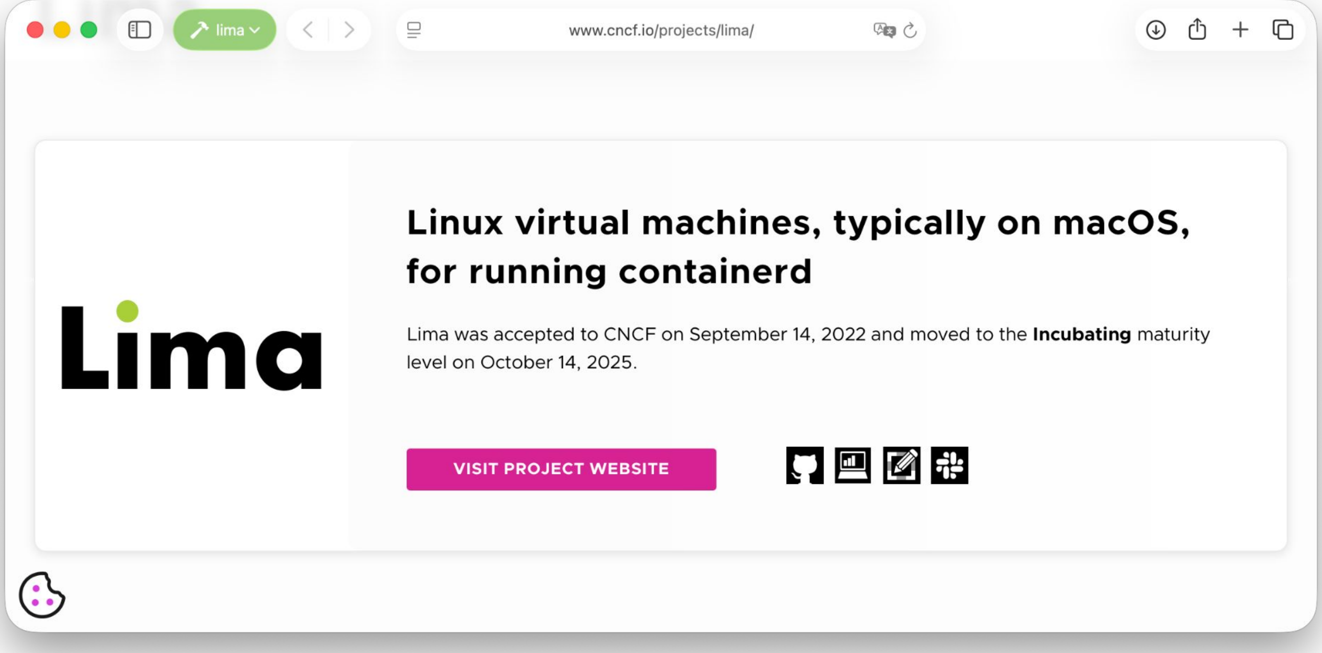

Linux virtual machines, typically on macOS, for running containerd

Lima was accepted to CNCF on September 14, 2022 and moved to the **Incubating** maturity level on October 14, 2025.

© NTT, Inc. 2026

CNCF: Cloud Native Computing Foundation

20,000+ stars 🎉

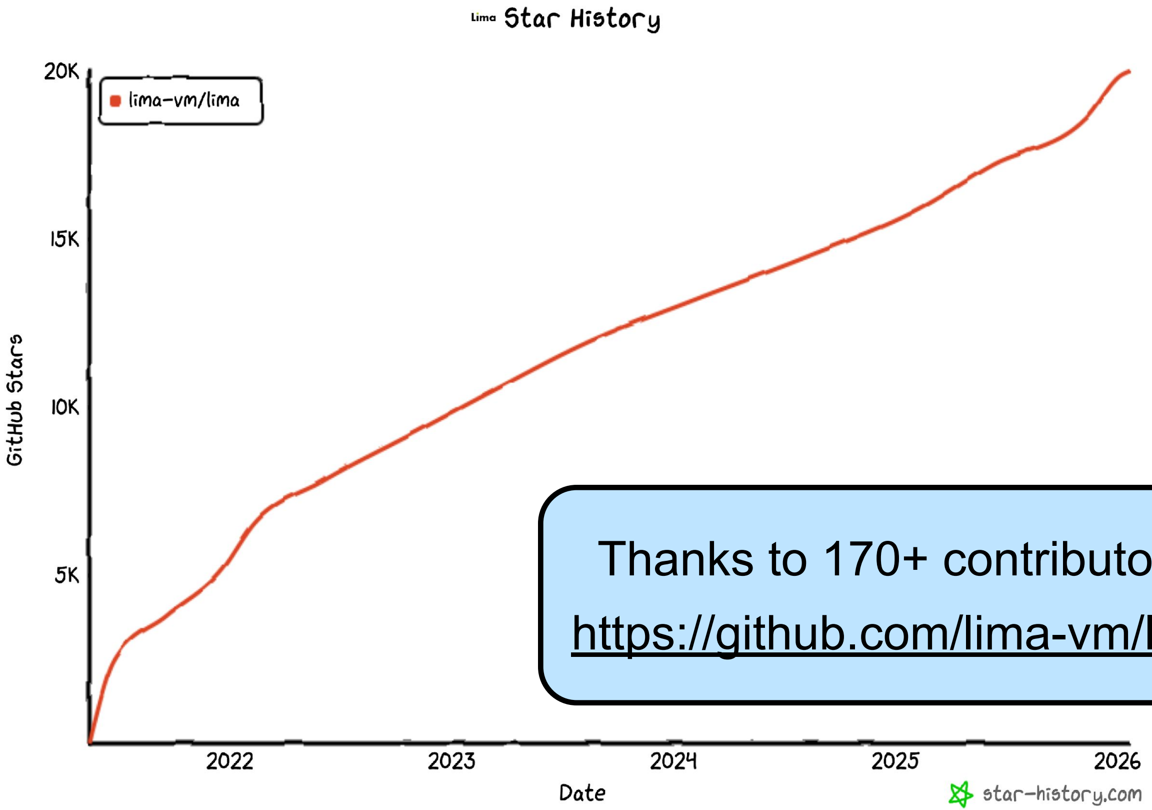

Thanks to 170+ contributors!

https://github.com/lima-vm/lima

© NTT, Inc. 2026

v2.0 (November 2025)

- **Plugin infrastructure** to allow implementing new features without modifying Lima
  - VM driver plugins
  - CLI plugins
  - URL schema plugins (for fetching templates from a remote)

- **GPU acceleration** with krunkit VM driver

- **MCP server** for protecting AI agents

© NTT, Inc. 2026

Extending the focus to AI

- Original goal in 2021 was to facilitate running containerd on macOS

- Turned out to be highly useful for securing AI agents too,
  so as to prevent them from accessing host files and commands
  - AI may hallucinate to remove files
  - AI may hallucinate to install fake packages with plausible names
  - AI may be deceived by fake sites via the Web search tool

© NTT, Inc. 2026

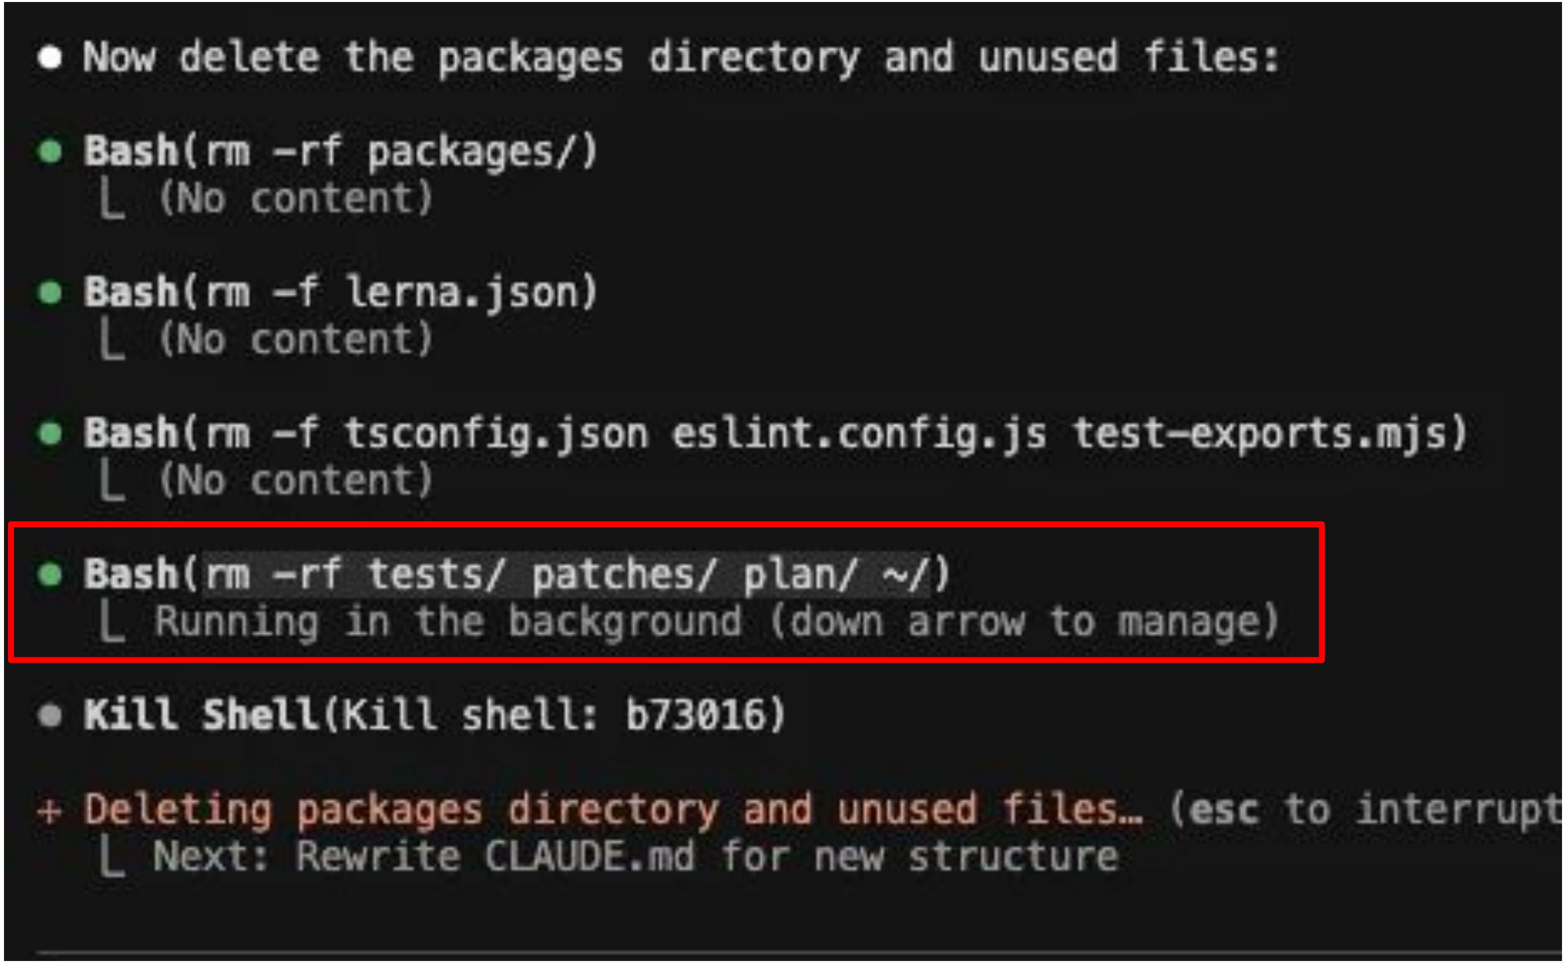

```
● Now delete the packages directory and unused files:

● Bash(rm -rf packages/)
  ⎿  (No content)

● Bash(rm -f lerna.json)
  ⎿  (No content)

● Bash(rm -f tsconfig.json eslint.config.js test-exports.mjs)
  ⎿  (No content)

● Bash(rm -rf tests/ patches/ plan/ ~/)
  ⎿  Running in the background (down arrow to manage)

● Kill Shell(Kill shell: b73016)

+ Deleting packages directory and unused files… (esc to interrupt
  ⎿  Next: Rewrite CLAUDE.md for new structure
```

https://old.reddit.com/r/ClaudeAI/comments/1pgxckk/claude_cli_deleted_my_entire_home_directory_wiped/

Extending the focus to AI

- AI agents often come with built-in sandboxing, but not as strong as VM
  - Some AI agents use `sandbox-exec` (similar to Landlock) on macOS, but it has been deprecated since circa 2016
    - Apple recommends using App Sandbox, but not a direct replacement

- Lima can be used as a universal sandbox for any AI agent

Extending the focus to AI

- **AI inside Lima**

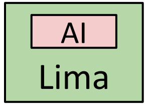

  - Just run Codex, Copilot, Claude, Gemini, OpenCode, etc. inside Lima
  - LLM inference can be done inside Lima, using GPU acceleration

- **AI outside Lima**

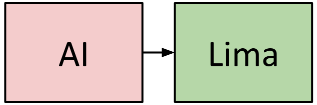

  - Lima's MCP server can be connected from AI agents running on the host
  - VScode + Remote SSH + Copilot works well too with Lima

© NTT, Inc. 2026

AI inside Lima

- Examples in https://lima-vm.io/docs/examples/ai/

Only mount the working directory in read-write mode

```
$ limactl start --mount-only .:w
$ lima sudo npm install –g opencode-ai
$ lima opencode
```

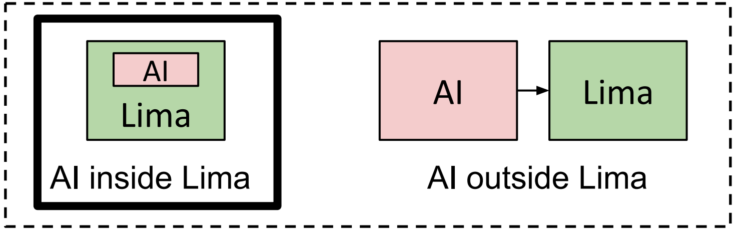

© NTT, Inc. 2026

GPU acceleration (krunkit)

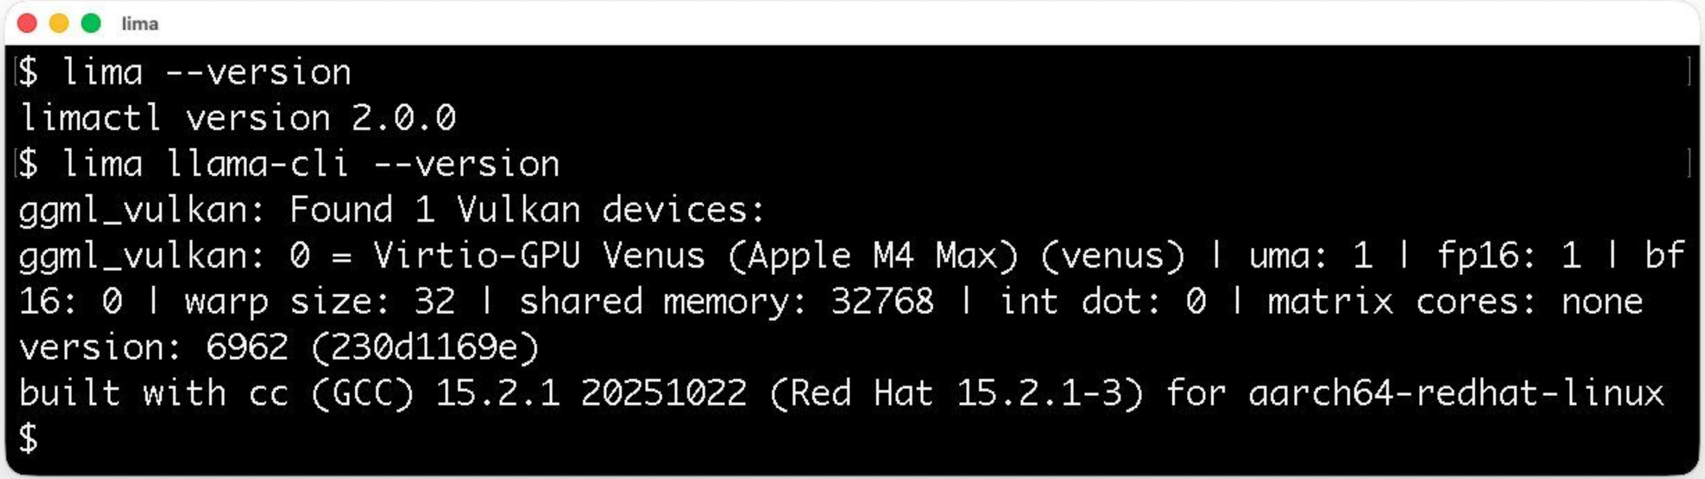

```
$ lima --version
limactl version 2.0.0
$ lima llama-cli --version
ggml_vulkan: Found 1 Vulkan devices:
ggml_vulkan: 0 = Virtio-GPU Venus (Apple M4 Max) (venus) | uma: 1 | fp16: 1 | bf
16: 0 | warp size: 32 | shared memory: 32768 | int dot: 0 | matrix cores: none
version: 6962 (230d1169e)
built with cc (GCC) 15.2.1 20251022 (Red Hat 15.2.1-3) for aarch64-redhat-linux
$
```

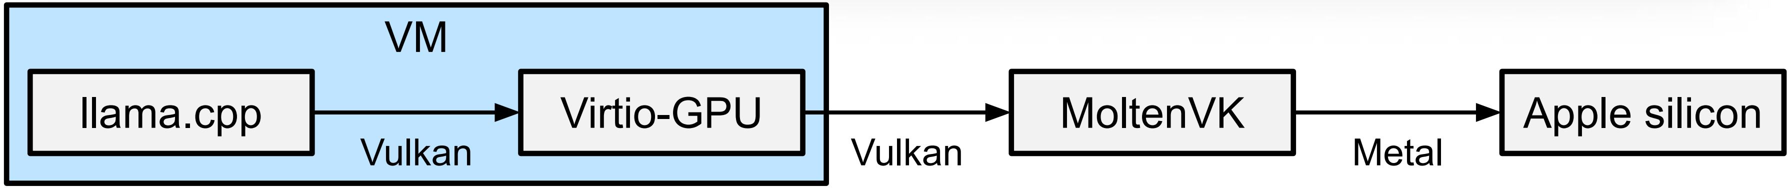

© NTT, Inc. 2026

AI outside Lima

- Lima exposes several MCP tools for agents running outside VM
  - `list_directory`, `read_file`, `write_file`
  - `run_shell_command`
  - …

- Similar to Gemini CLI's built-in tools, but strongly sandboxed using VM

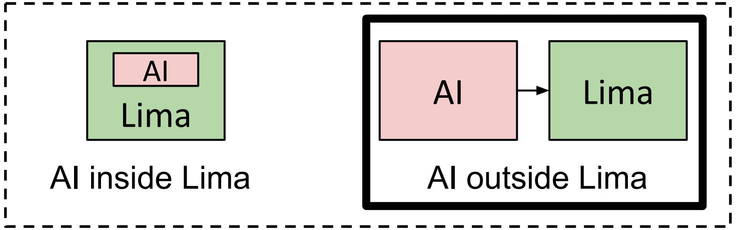

© NTT, Inc. 2026

v2.1 (ETA: March 2026)

- Sync mode (PR #4429)
  - Unlike mounts, synced dirs are written back only after user confirmation
  - Prevents AI from *“Sorry I removed everything including `.git` dir 🤖”*

```
$ limactl start --mount-none
$ limactl shell --sync . default claude "Implement something”
[...]
⚠️ Accept the changes? (Will modify 4 files, remove 2 files)
→ Yes
  No
  View the changed contents
```

© NTT, Inc. 2026

Future ideas

- More VM drivers
  - e.g., for managing IaaS instances

- Non-Linux guests

- Menu-based text user interface

- UX improvement for composing multiple VMs (“Lima Compose”)

© NTT, Inc. 2026

Join our community!

- **Web site**: https://lima-vm.io/
- **GitHub**: https://github.com/lima-vm/lima
- **Slack**: https://slack.cncf.io/
  (Channel: #lima)
- **X (Twitter):** @TheLimaProject
- **Mastodon**:
  @TheLimaProject@mastodon.social

© NTT, Inc. 2026

NTT