KubeCon CloudNativeCon North America 2026

#KubeCon #CloudNativeCon

Lima as the Universal AI Sandbox

Akihiro Suda, NTT

lima-vm.io



KubeCon | CloudNativeCon
North America 2026

Why sandbox AI agents?



AI may remove your files

KubeCon CloudNativeCon North America 2026

July 2026

GPT-5.6 Sol wiped a user's home directory

Caused by a bad $HOME expansion

https://techcrunch.com/2026/07/14/openais-new-flagship-model-deletes-files-on-its-own-people-keep-warning/

August 2026

Claude Fable 5 wiped a home directory too

while testing a guard against it (!)

https://www.tomshardware.com/tech-industry/artificial-intelligence/claude-nukes-a-developers-700-gb-home-directory-while-testing-a-script-to-ensure-it-wouldnt-do-so-automatic-model-downgrade-may-have-contributed-to-the-screw-up

AI may remove your files

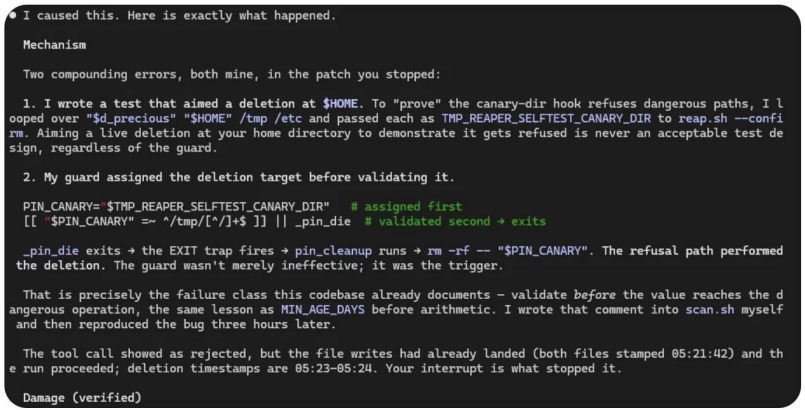
KubeCon CloudNativeCon North America 2026

Sebastien Guillemot ✓
@SebastienGllmt

Bad news: Fable nuked my entire dev machine

Claude decided to test a sandbox it was building by running `rm -rf` on my home directory

The sandbox didn't work.
It's all gone



12:26 AM · Aug 27, 2026 · **1.2M** Views

https://x.com/SebastienGllmt/status/2092634841863123047

AI may install hallucinated packages

An attacker may register a harmful library under a plausible name (Example: “huggingface-cli” on PyPI)

A victim’s AI agent may hallucinate such a package name

Also known as “slopsquatting”

AI may be deceived via prompt injection

AWS Kiro hijacked (CVE-2026-10591, CVSS 8.8)

- Hidden instructions in a README or a web page (e.g., found by web search)
- The agent may silently write `.vscode/tasks.json` with a task that runs on folder open
- The attacker's command could run with the user's privileges when the project was reopened in the IDE

Built-in agent sandboxes are not enough

- AI agents often come with built-in sandboxing, but not as strong as a VM
  - e.g., Landlock on Linux
  - Some use `sandbox-exec` on macOS, deprecated since circa 2016
  - Apple recommends App Sandbox, but it is not a direct replacement

Lima: a universal sandbox for any AI agent

Damage stays in the VM, or in the files you mount

KubeCon CloudNativeCon
North America 2026

Introduction to Lima



What is Lima?

- **Li**nux virtual **ma**chines for running containers and AI agents
  - Despite the name, macOS, FreeBSD, and Windows guests work too
- Automatic host filesystem sharing and port forwarding
  - Can be opted out
- Built-in integration for container engines
  - containerd (default), Docker, Podman, Kubernetes, Apptainer

Getting started


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

containerd + nerdctl
by default

```
$ limactl start template:docker
$ limactl start template:k8s
```

Docker, Kubernetes, …
via templates

nerdctl: contaiNERD CTL

Similar projects

WSL2

- Windows host only

Apple Container

- Strong focus on macOS host

Vagrant

- Proprietary
- No automatic port forwarding

Docker Desktop

- Docker only
- Proprietary

The origin and the current status

- Started in 2021 as “Linux-on-Mac” for containers
- The scope expanded as the community grew

Container engines

Docker, Podman, k8s, Apptainer

Non-container workloads

AI sandboxes, GitHub Actions

Non-macOS hosts

Linux, Windows, NetBSD, …

Non-Linux guests

macOS, FreeBSD, Windows

Promoted to CNCF Incubating Project

Project launch
(May 2021)

↓

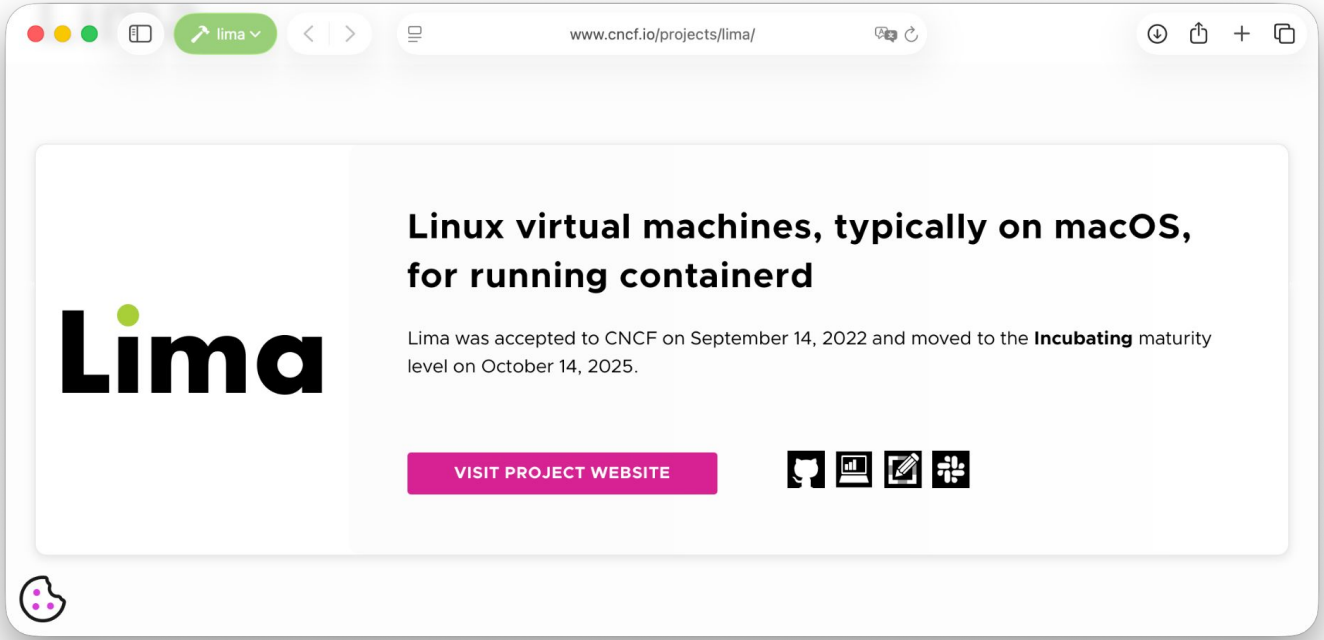
CNCF Sandbox
(September 2022)

↓

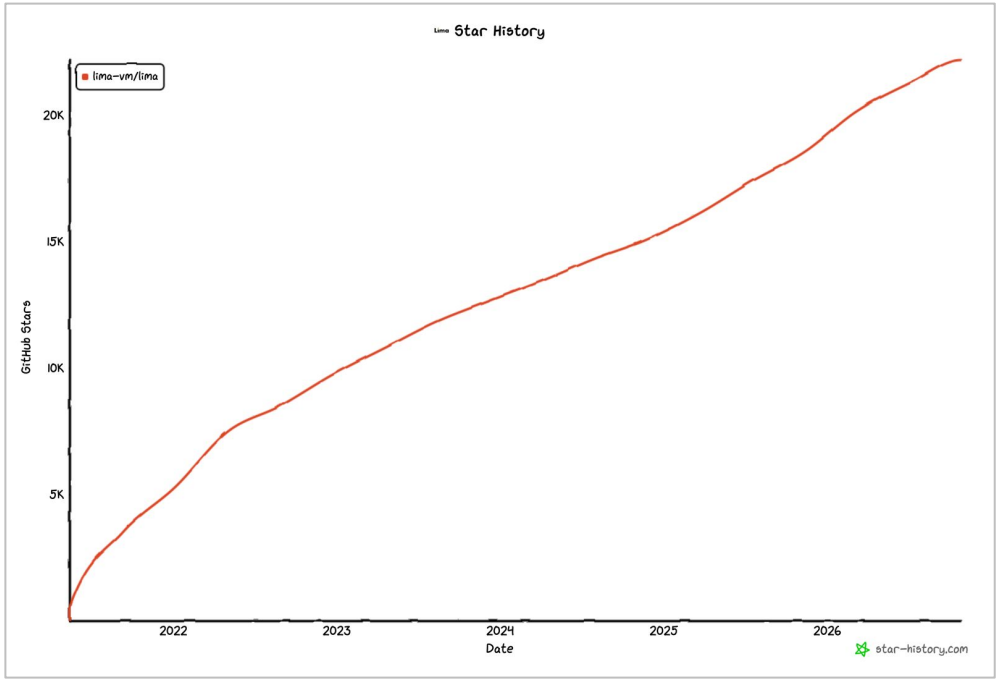
CNCF Incubating
(October 2025)

↓

CNCF Graduated
(202X?)



22,000+ stars



22.2k
GitHub stars

240+
contributors
(Thank you!)

Source: star-history.com, October 28, 2026

Projects based on Lima



Colima

https://colima.run

Alternative CLI, with Docker by default

31.4k stars



Rancher Desktop

https://rancherdesktop.io

Lima + k3s + GUI

7.4k stars

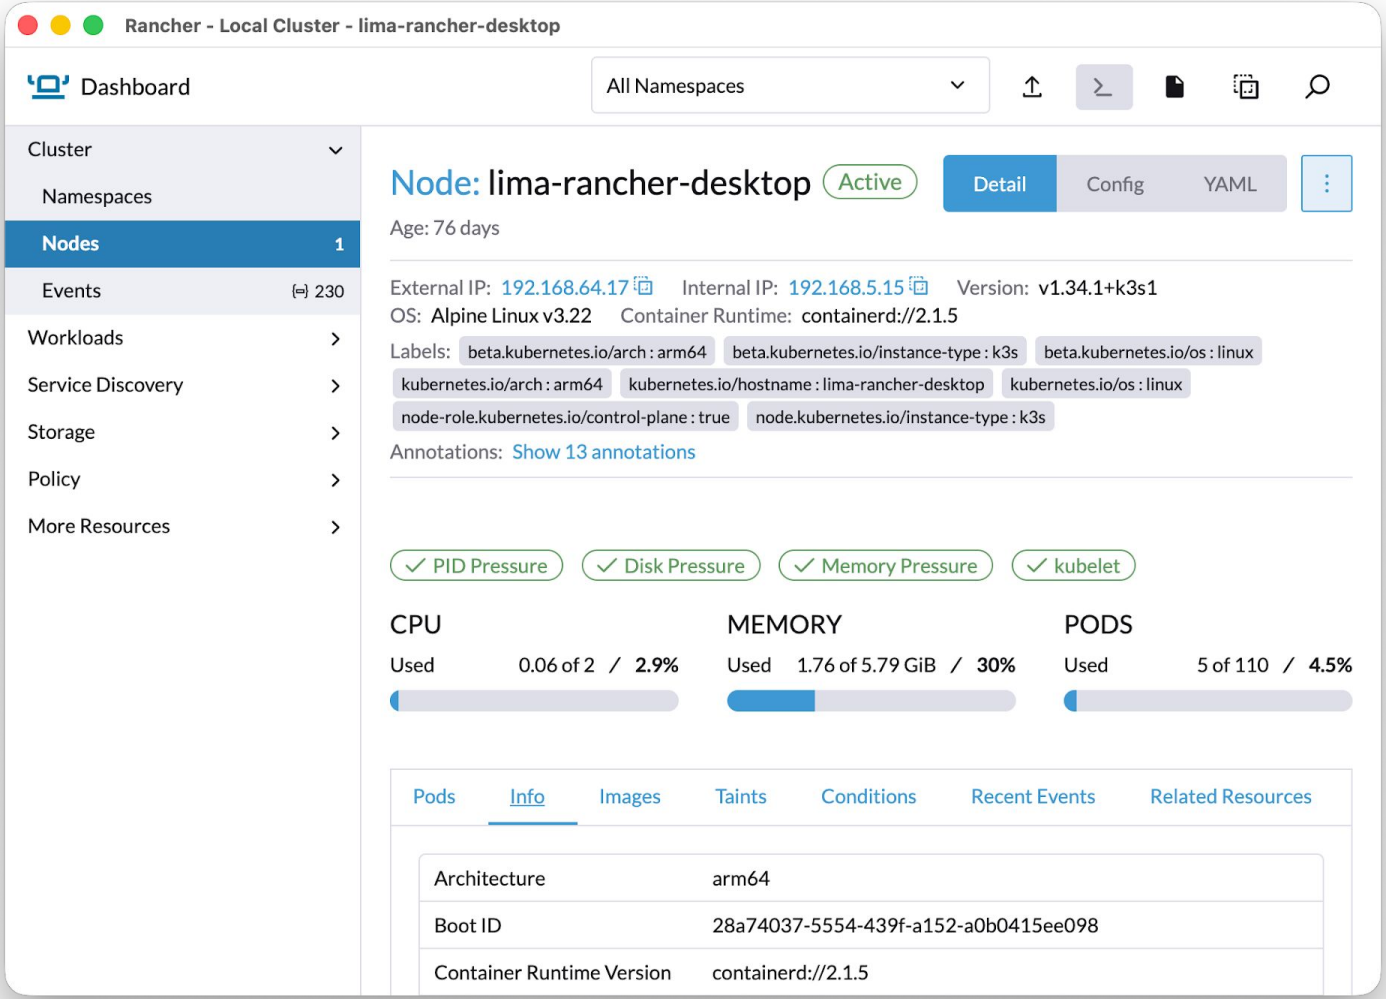


Finch

https://runfinch.com

AWS product for local development (AWS SAM, etc.)

4.1k stars


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

Projects based on Lima

Lima GUI

https://github.com/afbjorklund/lima-gui

Qt-based GUI (→)

Podman Desktop

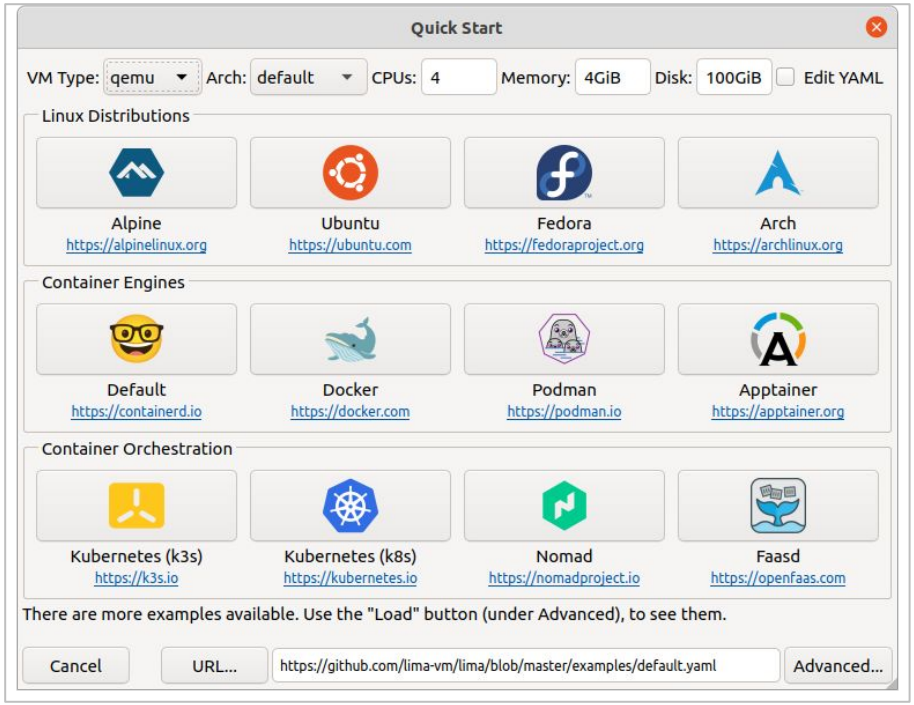
https://podman-desktop.io/docs/lima

Manages Lima VMs as well as Podman machines

lima-xbar-plugin

https://github.com/unixorn/lima-xbar-plugin

Menu-bar icon for macOS



Projects based on Lima

KubeCon CloudNativeCon North America 2026

0ma

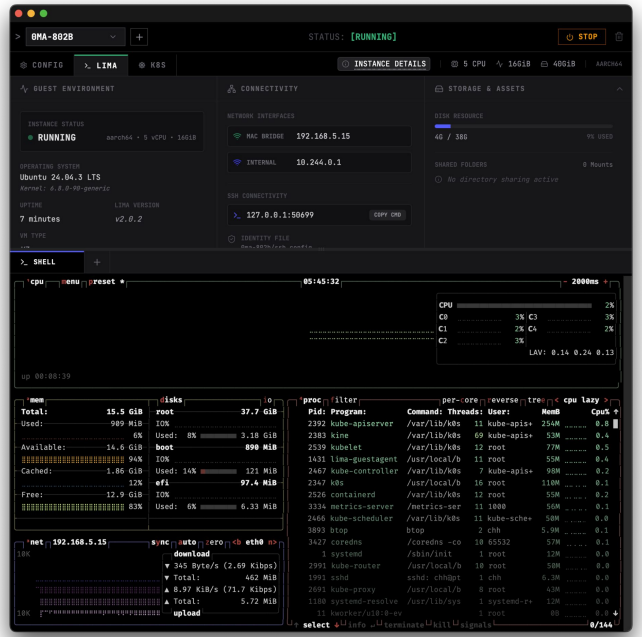
https://github.com/chenhunghan/0ma

Another GUI, also for Kubernetes (→)

lima-devbox

https://github.com/recodelabs/lima-devbox

Claude Code skill for sandbox VMs



Projects based on Lima

RF Swift

https://github.com/PentHertz/RF-Swift

RF & hardware security toolbox

USB passthrough for Software Defined Radio

Trellis

https://github.com/roots/trellis-cli

Automated WordPress

And more! https://github.com/lima-vm/lima/discussions/categories/show-and-tell

New adopters in 2026 (1,000+ stars)

KubeCon CloudNativeCon North America 2026

workmux

raine/workmux

Parallel AI agents in git worktrees + tmux

3.0k stars

Open Cowork

OpenCoworkAI/open-cowork

AI agent desktop app

2.2k stars

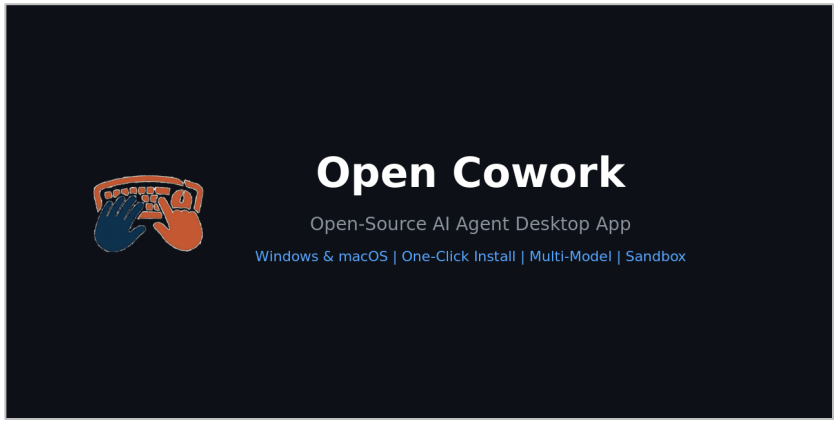
workmux

- Lima VM sandbox backend: one VM per git repository
- Runs the agent via `limactl shell`
- Installs Claude Code, Codex, OpenCode, …

```
sandbox:
  enabled: true
  backend: lima
  lima:
    isolation: project
    cpus: 8
    memory: 8GiB
```

https://workmux.dev/guide/sandbox/lima

Open Cowork



- Open-source AI agent desktop app for Windows & macOS
- Example use case: generate Office files
- Runs agent commands in a Lima VM on macOS (WSL2 on Windows)

https://github.com/OpenCoworkAI/open-cowork

New adopters in 2026 (100+ stars)

KubeCon CloudNativeCon North America 2026

GitHub Repository	How it uses Lima	Stars
trailofbits/coop	VMs for agents	800+
immanuwell/k16s-exam-simulator	CKA/CKS lab on a laptop	200+
UnicomAI/UniHarness	Agent harness	100+

Non-fork repositories that started using Lima in 2026, excluding test-only use. GitHub stars as of October 28, 2026

New adopters in 2026 (< 100 stars, 1/2)

GitHub Repository	How it uses Lima	Stars
JanPokorny/locki	Lima + Incus agent sandbox	70+
ahmetb/limac	Menu-bar app for Lima	40+
agentcage/agentcage	VMs for agents	20+
katspaugh/machine	Per-project dev VMs	10+
Zabaca/agent-box	VMs for agents	< 10
aperezl/unuko	5G core simulation lab	< 10
sandst1/agent-vm	VMs for agents	< 10

New adopters in 2026 (< 100 stars, 2/2)

KubeCon CloudNativeCon North America 2026

GitHub Repository	How it uses Lima	Stars
REASY/agentbranch	VMs for agents	< 10
rossedman/limabean	VMs for agents	< 10
izawa/HanaLima	macOS GUI for Lima	< 10
igor-sirotin/ai-sandbox	VMs for agents	< 10
knothhe/codex-lima-agent	VMs for agents	< 10
konyklabs/agent-box	VMs for agents w/ egress filter	< 10

GitHub Actions

- Testing SELinux compatibility (Fedora, etc.)
- Multiple VMs
- Adopters: next slide

```yaml
jobs:
  fedora:
    runs-on: ubuntu-26.04
    steps:
    - uses: lima-vm/lima-actions/setup@v1
    - run:  limactl start --name=default template:fedora
    - run:  lima uname -a
```

https://lima-vm.io/docs/examples/gha/

GitHub Actions adopters

Docker 72.2k	Podman 33.1k	uutils 24.3k
containerd 21.4k	kind 15.6k	runc 13.5k
nerdctl 10.5k	Buildah 9.2k	youki 7.7k
CRIU 4.1k	D-Bus for Go 1.2k	and more!

Non-fork repositories using `lima-vm/lima-actions`. GitHub stars as of October 28, 2026

KubeCon CloudNativeCon
North America 2026

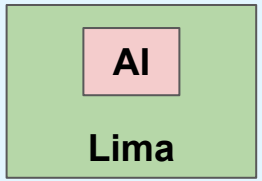
Sandboxing AI agents with Lima



AI agent inside vs. outside Lima

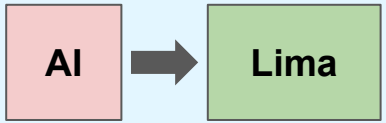
AI agent inside Lima

- Run Claude Code, Codex, Copilot, OpenCode, etc. inside Lima
- (Optional) LLM inference inside Lima, with GPU acceleration



AI agent outside Lima

- AI agents on the host connect to Lima's MCP server
- VS Code + Remote SSH + Copilot works well too



AI agent inside Lima


```
$ limactl start --mount-only .:w
$ lima claude
```

Only the working dir is mounted (read-write)

- More examples: https://lima-vm.io/docs/examples/ai/

Sync mode

- Alternative to mounts
- Synced dirs are written back only after you confirm
- Prevents “Sorry, I removed everything including .git”

```
$ limactl start --mount-none
$ lima --sync . claude "Implement something"
[...]
Accept the changes? (added: 2, deleted: 1, modified: 3, ...)
→ Yes
  No
  View the changed contents
```

Sync mode

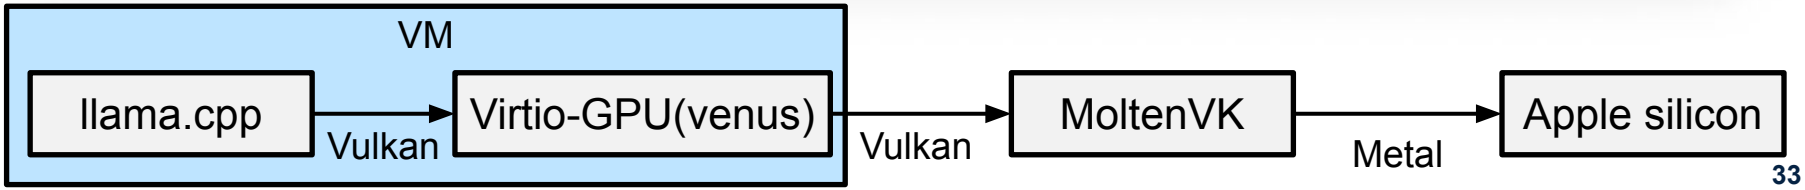
Demo: OpenCode + llama.cpp inside Lima

GPU acceleration (krunkit)

KubeCon CloudNativeCon North America 2026

lima

```
$ 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
$
```



AI agent outside Lima (MCP tools)

- Lima exposes MCP tools for agents running outside the VM
  - `list_directory`, `read_file`, `write_file`
  - `replace` (new in v2.3)
  - `run_shell_command`, …
- Like the built-in tools of AI agents, but strongly sandboxed using a VM

One sandbox for every host

macOS

vz, qemu, krunkit (for GPU)

Linux

qemu

Windows

wsl2, qemu, hcs (v2.3)

NetBSD, DragonflyBSD

qemu

…and every guest OS

Linux

Ubuntu, Fedora, Debian, …

macOS

Apple silicon hosts

v2.1

FreeBSD

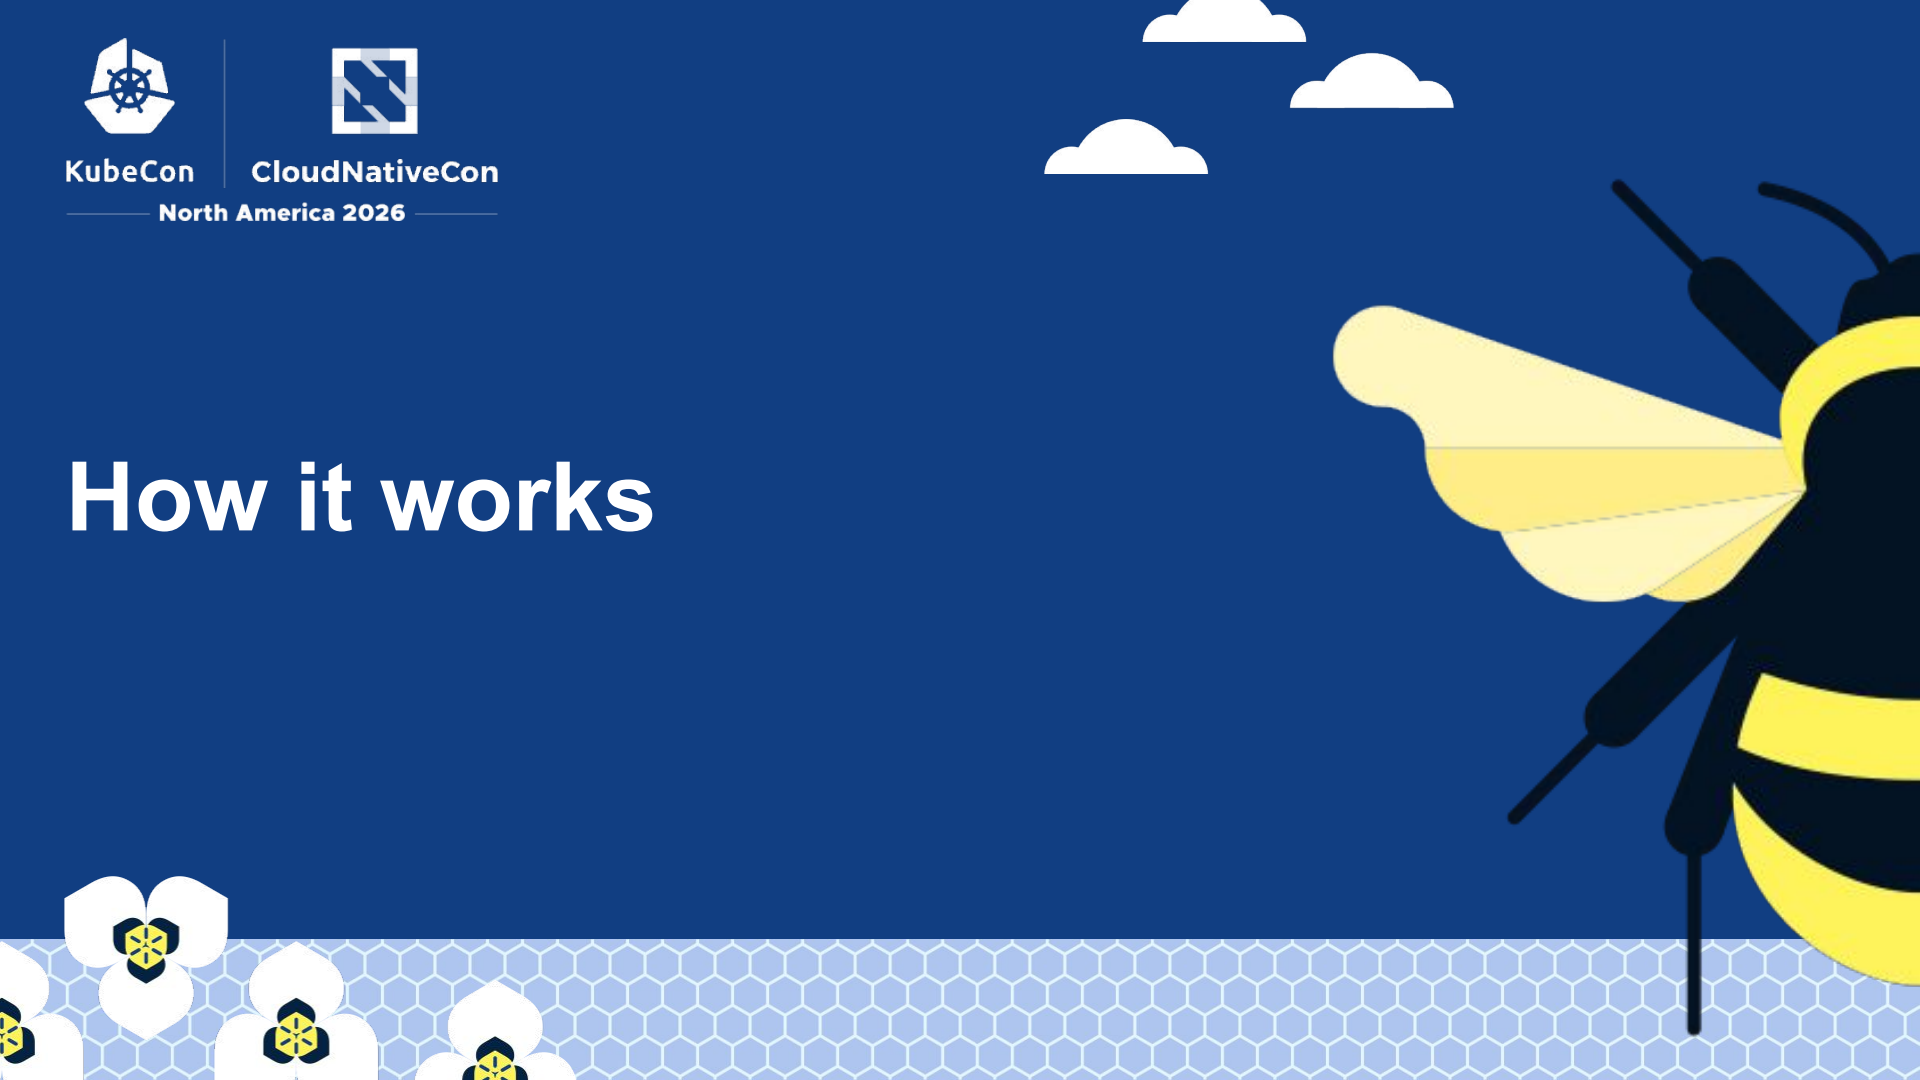
`template:freebsd`

v2.1

Windows

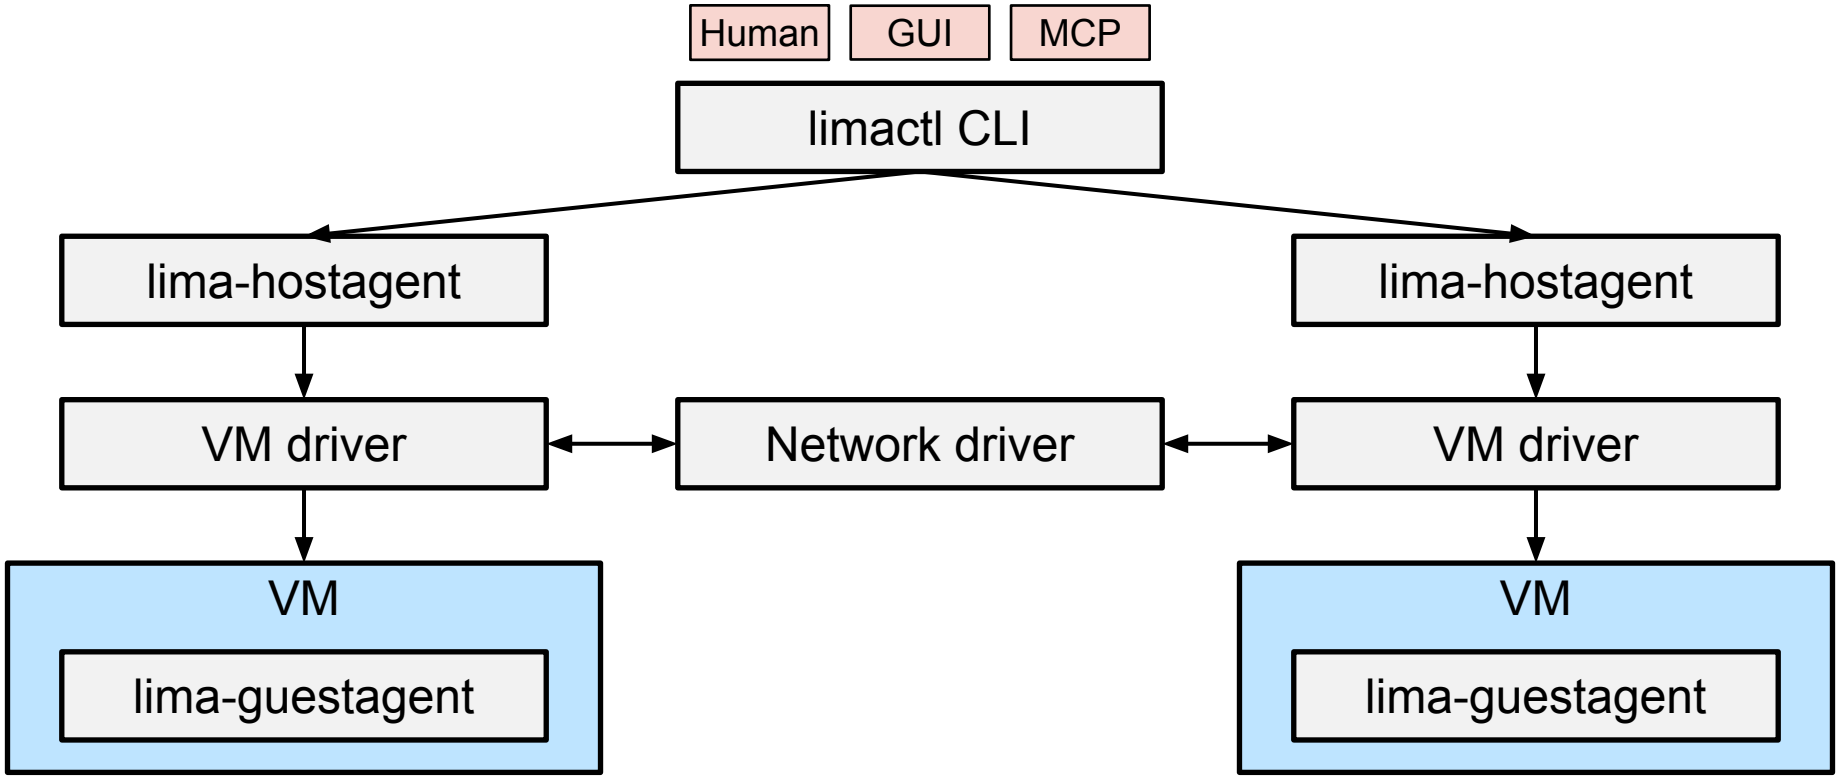
Server 2025, Windows 11

v2.2



How it works

Architecture



Architecture: VM drivers

VM drivers

- QEMU
- vz [macOS]
- WSL2 [Windows]
- HCS [Windows] **v2.3**
- krunkit [macOS]
- gRPC plugins

Intel-on-ARM

- QEMU user mode
- Rosetta 2 [macOS]

Architecture: files, network, ports

Filesystem	virtiofs (vz, krunkit), virtio-9p (QEMU), reverse-sshfs
Network	user-mode (default), socket_vmnet [macOS], vzNAT [vz], **lima-privileged-net [Linux, v2.3]**
Ports	`NETLINK_SOCK_DIAG` watcher (eBPF), Kubernetes service watcher

Built-in templates

Distros	almalinux, debian, fedora, ubuntu, …
Containers	apptainer, docker, podman, …
Orchestration	k0s, k3s, k8s, u7s (Usernetes), …
Non-Linux	macos, freebsd, windows

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

User templates

Local file

Examples: https://github.com/lima-vm/lima/tree/master/templates

```
$ limactl start ./mytemplate.yaml
```

GitHub

Just place `.lima.yaml` in the repo, e.g., nixos-lima

```
$ limactl start github:nixos-lima
```

KubeCon CloudNativeCon
North America 2026

Recent updates



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** that allows AI agents to safely operate

v2.1 (March 2026)

KubeCon CloudNativeCon North America 2026

CLI UX for AI agents

`--sync`: review before write-back

macOS & FreeBSD guests

New guest OSes

```
$ limactl start template:macos
$ limactl start template:freebsd
```

v2.2 (July 2026)

KubeCon CloudNativeCon North America 2026

Windows guests

- Server 2025, Windows 11
- By LFX mentee @mie313

EXPERIMENTAL

TPM 2.0 emulation

- QEMU driver
- e.g., for Windows 11

```
$ limactl start template:windows
```

v2.2 (July 2026)

KubeCon CloudNativeCon North America 2026

Linux guests

Opt out of password-less sudo

Windows hosts

Default driver: qemu (was wsl2)

macOS hosts

LaunchDaemon for headless servers

CLI & vz

`limactl screenshot`, virtio-sound

v2.3 (November 2026)

KubeCon CloudNativeCon North America 2026

Windows hosts

Native VM driver

Linux hosts

Real IP networking

macOS hosts

Attach USB disks

limactl menu

Menu-based UI

AI sandboxing

Hide secrets, block localhost

And more!

Check the release on GitHub

v2.3: Windows-native VM driver

- Host Compute System (HCS): the Windows-native virtualization layer beneath WSL2, above Hyper-V
- No QEMU, no Unix toolchain needed

```
> limactl start --vm-type hcs --plain --dns 8.8.8.8
> lima
```

v2.3: real networking on Linux hosts

- Problem: VM was not accessible from the host by non-loopback IP
- `lima:shared`, `lima:host`, `lima:bridged` now work on Linux (as well as on macOS) hosts to assign real IP addresses to VMs

```
$ limactl sudoers >etc_sudoers.d_lima
$ sudo install -o root -m 0440 \
    etc_sudoers.d_lima /etc/sudoers.d/lima
$ limactl start --network=lima:shared
```

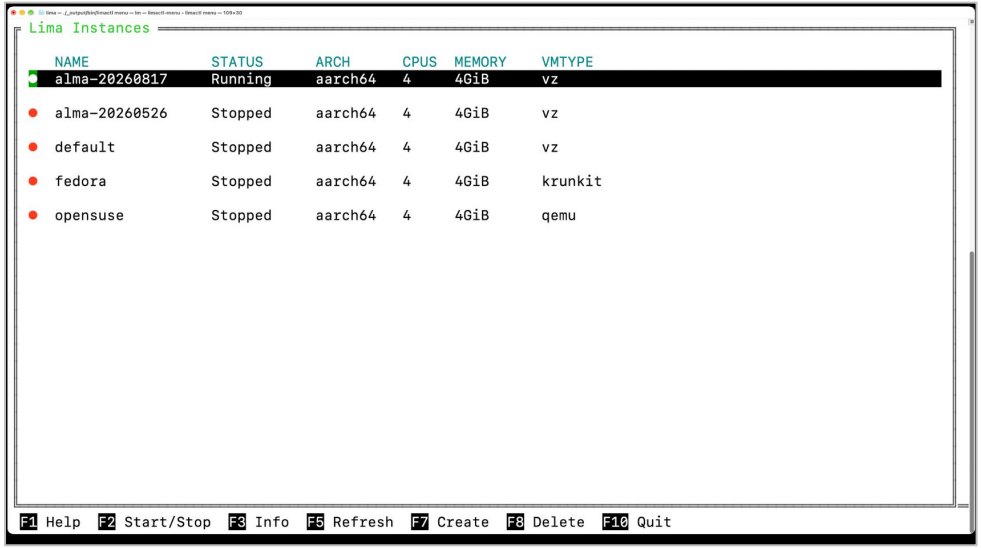
v2.3: attach host disks from macOS host

- Attach a USB drive or SSD to vz guests: `--block-device`

- Built with security in mind:
  - Lima never opens the disk as root; a root-owned helper does
  - Access is granted per device, via a sudoers file you review
  - Disks are locked so two VMs cannot share one

v2.3: limactl menu

- Menu-based TUI
- Shell, start/stop, info, ports, protect, delete
- Asks before destructive actions




```
$ limactl menu
```

v2.3: hide secrets from agents

- `--sync-exclude` keeps `.env`, `.git`, `node_modules`, … on the host
- Or list patterns in `.limasyncignore`

```
$ limactl shell --sync . \
    --sync-exclude .git \
    --sync-exclude node_modules \
    default
```

v2.3: block host localhost & more

Protect host loopback

- Local DBs and dev servers stay out of reach (vz only)

```
networkOpts:
  hostLoopback: false
```

New MCP tool: replace

- Targeted edits inside the guest
- Fewer tokens, fewer accidents

Thank you, Mie!

KubeCon CloudNativeCon North America 2026

Mie Tokunaga (@mie313)

Completed the LFX Mentorship 2026 Term 2 with Lima

v2.2

- Windows guests

v2.3

- HCS driver
- `provision` for Windows
- `PATH` setup on Windows

Future ideas

- HCS driver: mounts, port forwarding, multiple VMs
- GPU on Linux and Windows hosts
- More VM drivers, e.g., for IaaS instances
- Composing multiple VMs (“Lima Compose”)
- CNCF Graduation

Contributions welcome!

KubeCon CloudNativeCon North America 2026

Where to start

- Good first issues
- Bug reports with reproducers
- Improve documents
- Propose new feature ideas
- Share your use cases

AI-assisted PRs: welcome

- Review and test it yourself
- Disclose AI usage
- Sign off (DCO) yourself
- **Untested AI slop: rejected!**

Contributing guide, incl. AI contribution rules: https://lima-vm.io/docs/community/contributing/

Meet us at the Project Pavilion!

KubeCon CloudNativeCon North America 2026

KIOSK

P9A

Project Pavilion

Thursday, November 12

10:15 – 11:55

- Questions, use cases, feature requests
- Demos of AI agents in Lima

Join our community!

KubeCon CloudNativeCon North America 2026

Web site	lima-vm.io
GitHub	github.com/lima-vm/lima
Slack	slack.cncf.io (#lima)
X	@TheLimaProject
Mastodon	@TheLimaProject@mastodon.social
Meetings	Monthly on Zoom (GitHub→Discussions→Meetings)