March 2026

Why Kubernetes Security Comes Down to Linux Security

Marina Moore, Edera

whoami

- Head of Edera Research
- Co-chair CNCF TAG Security and Compliance
- TUF maintainer
- Live in New York City

Agenda

1. Container escapes
2. History of containers
3. Scope of cgroups & namespaces
4. Demo
5. What this means for you

Containers

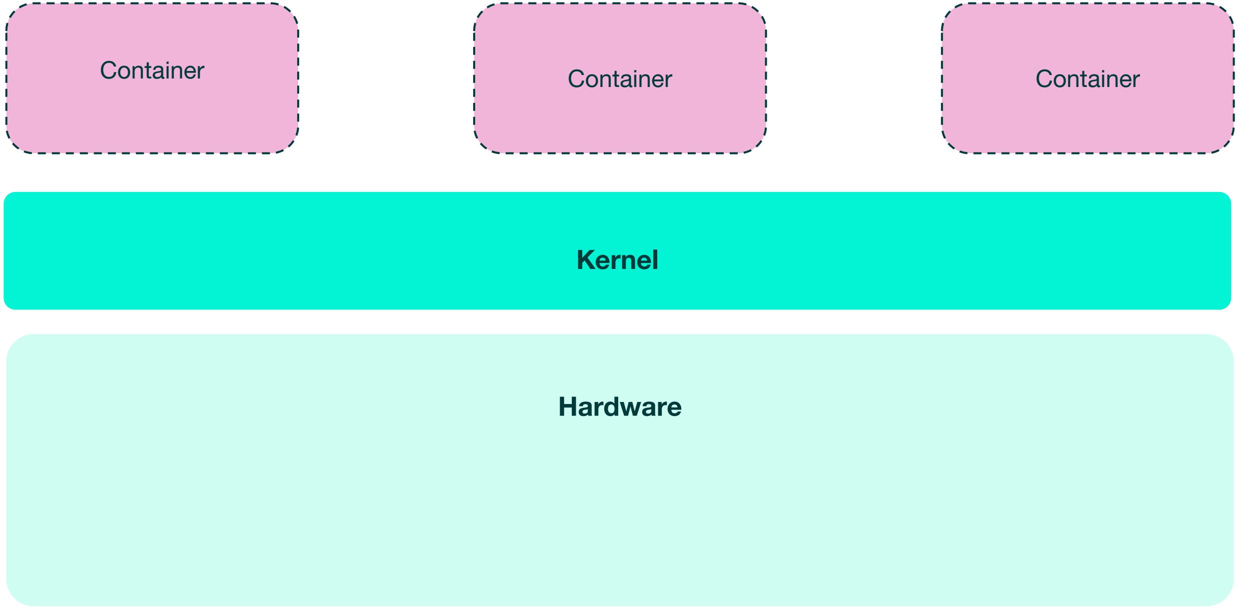

Container Escapes

- Violation of the container boundary
- Enables unauthorized operations on the host system
    - Code execution
    - Escalation of privileges
    - Host filesystem access
    - Information disclosure
    - Data tampering

Container Escapes

- CVE-2024-1753: Bug in Buildah and Podman Build allows containers to mount host files to the build system, allowing container escape at build time
- CVE-2024-0132: NVIDIA container toolkit TOCTOU vulnerability gives container access to host filesystem, leading to container escape
- CVE-2024-21626 (Leaky Vessels): container escape in runc and buildkit gives access to the host filesystem
- CVE-2025-9074: Docker Desktop vulnerability allows for privilege escalation and host filesystem access
- CVE-2025-23266 (NVIDIAScape): NVIDIA container toolkit flaw that leads to privilege escalation

Containers are not a security boundary

They are a mechanism to control resource usage.

History of Containers: pre-history

- Virtual machines
    - To isolate a process, bring a whole operating system
- Container = OS-level virtualization
    - Shared kernel with process isolation!

History of Containers

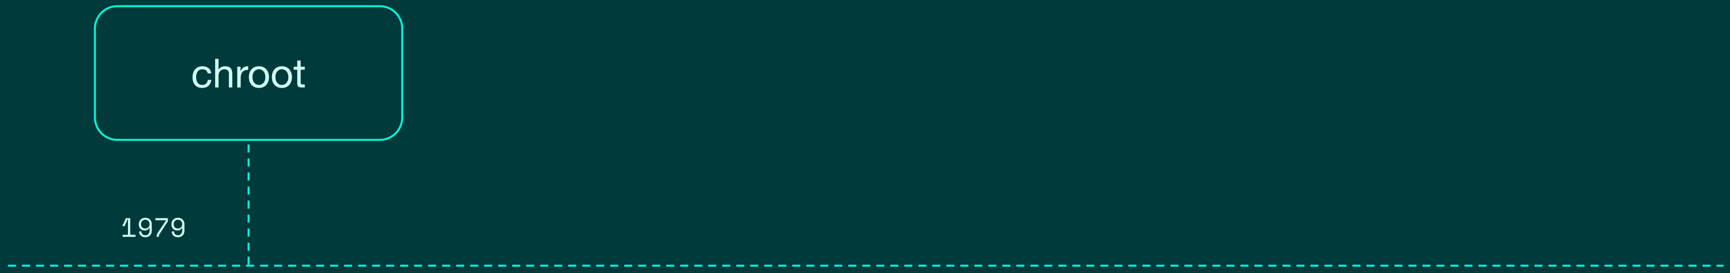

chroot

1979

History of Containers

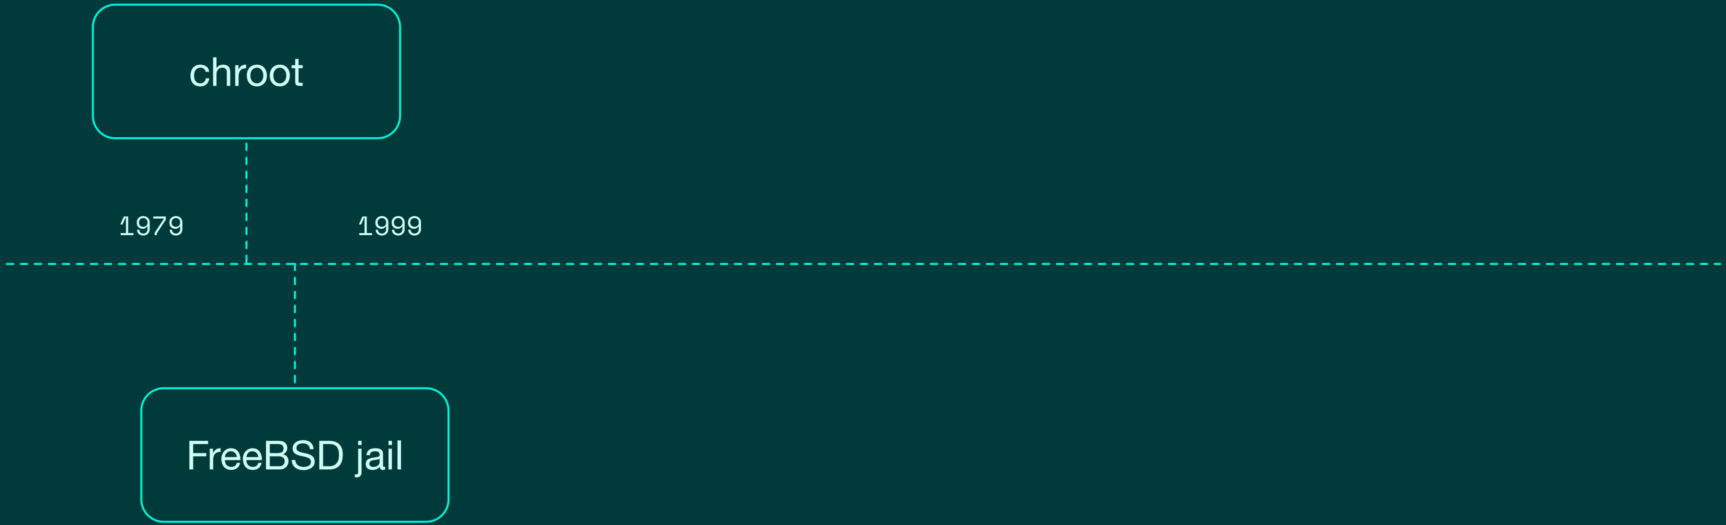

chroot

1979

1999

FreeBSD jail

History of Containers

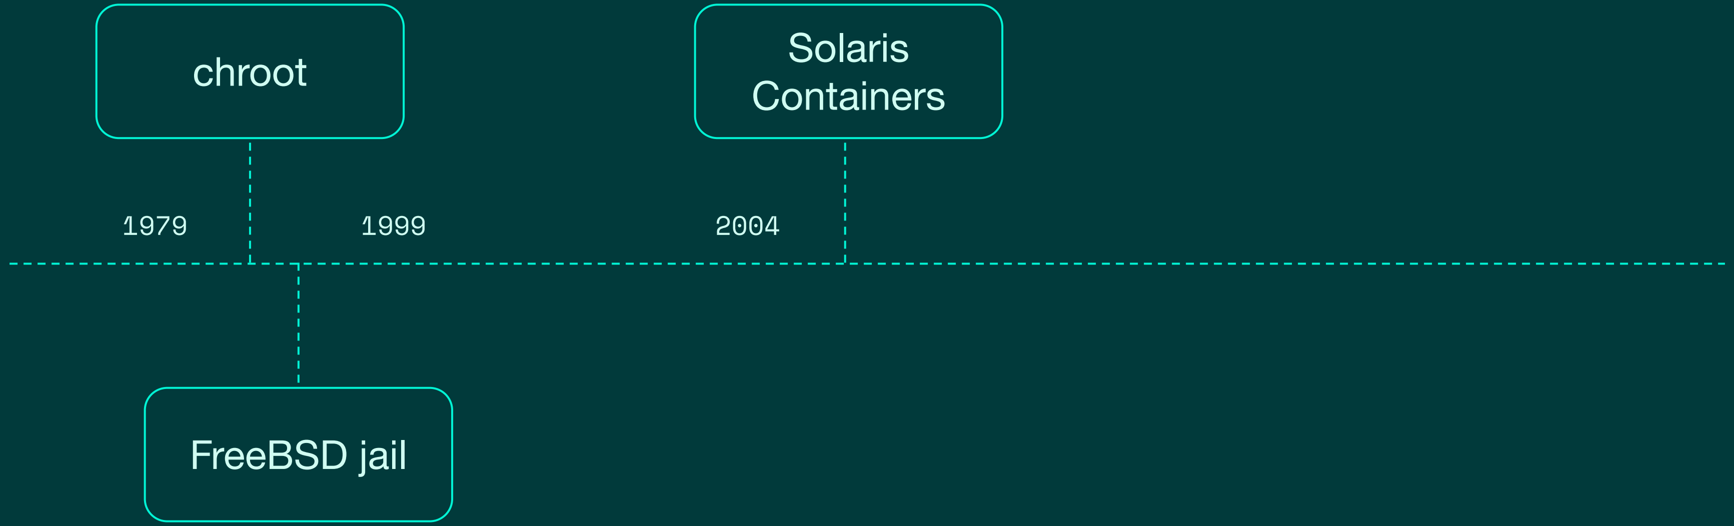

chroot

Solaris Containers

1979

1999

2004

FreeBSD jail

History of Containers

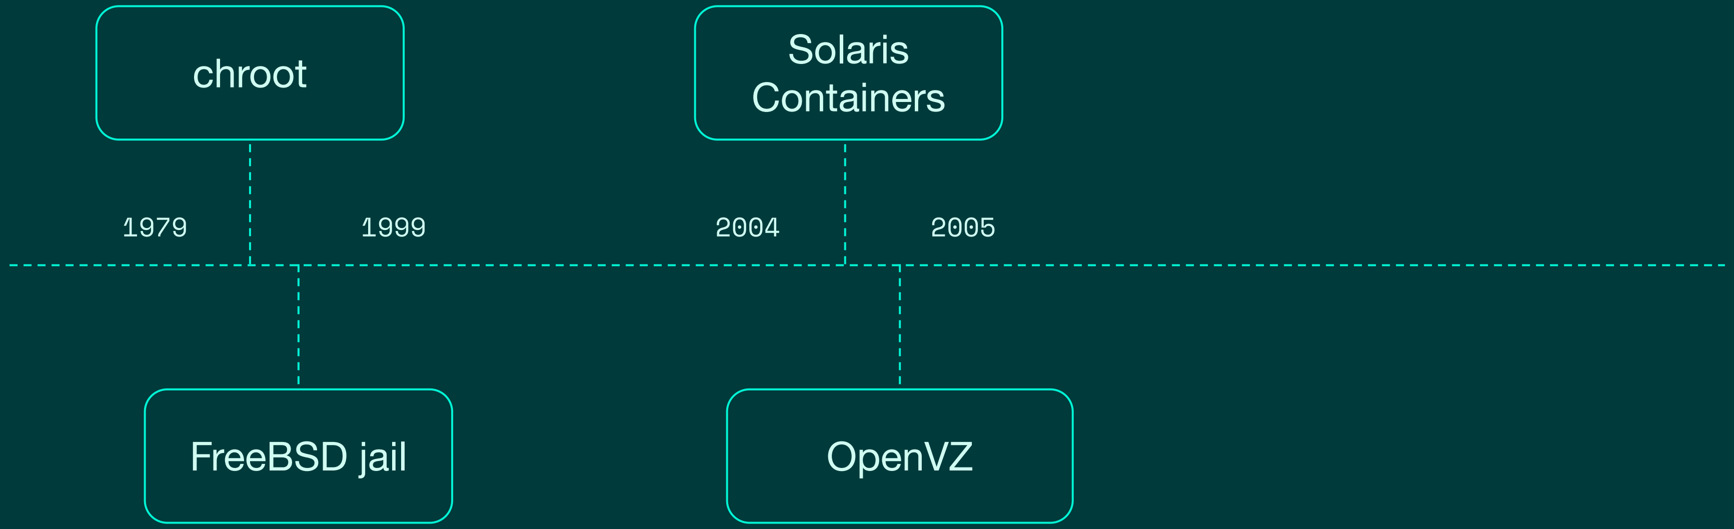

History of Containers

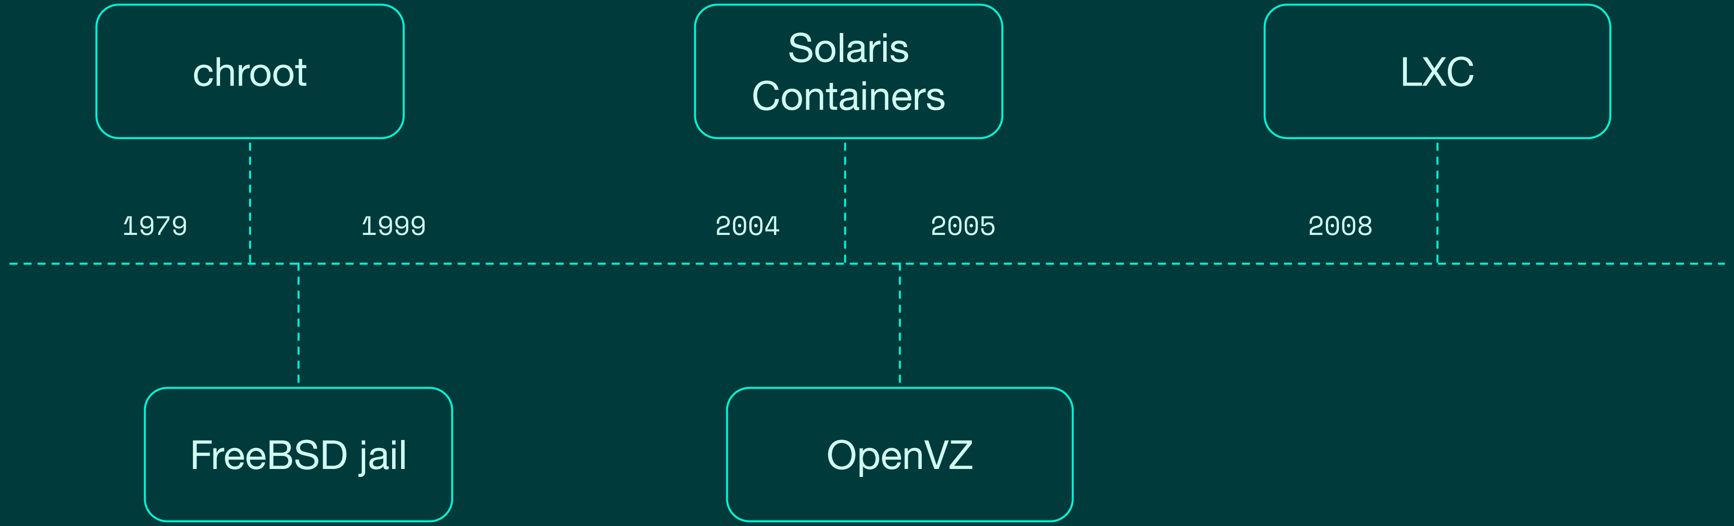

History of Containers

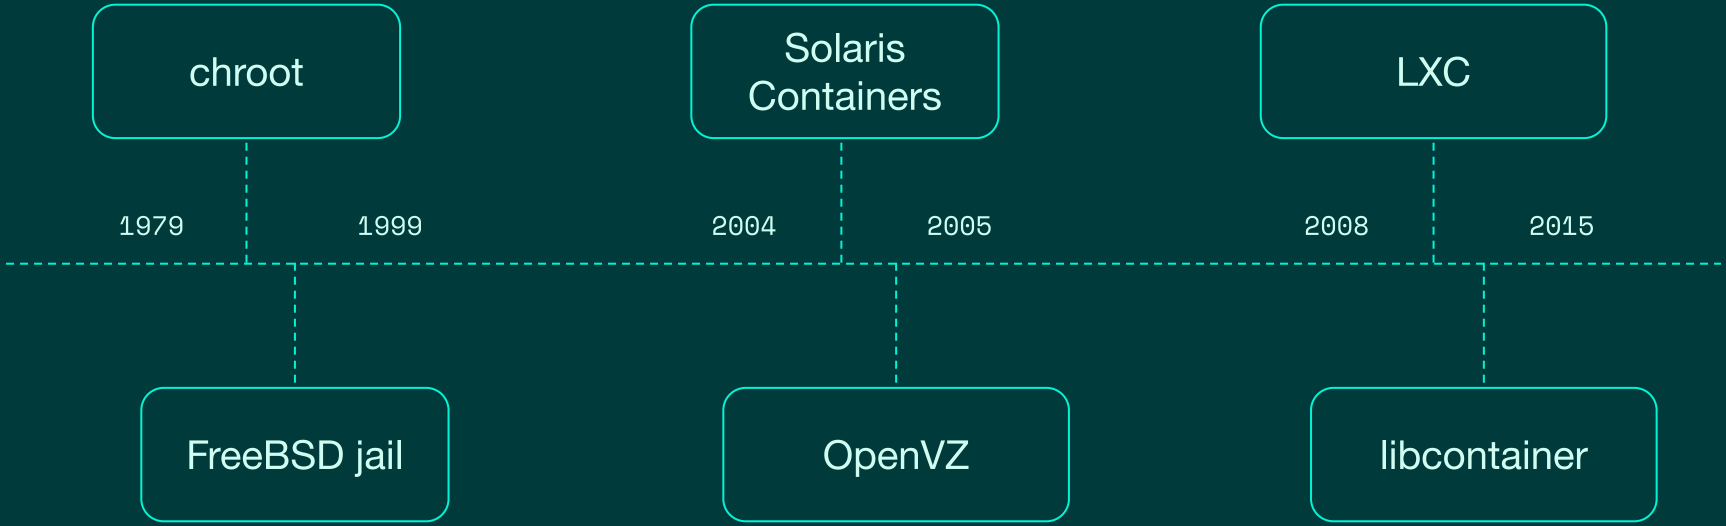

chroot

1979

FreeBSD jail

1999

Solaris Containers

2004

OpenVZ

2005

LXC

2008

libcontainer

2015

Containers: where we are now

- OS virtualization, but with:
    - Namespaces
    - Cgroups
    - Capabilities
    - security profiles
    - network interfaces
    - …

Container Security = Linux Security

- All these virtualization mechanisms are in the Linux kernel
- Container escapes stem from
  - File system access
  - Violations of namespaces, cgroups, etc
  - Network access
  - Privileged commands
- All of this is Linux security

Scope of namespaces and cgroups

- In scope: Userspace
- Out of scope:
  - Networking
  - File descriptors

Demo time

History of Containers: post-history

- Multi-kernel
- Sysbox
- gVisor
- User namespaces
- microVMs
- WASM

What this means for you

- To understand containers, learn about Linux!
- Understand which resources your container can see, what this means
- OS virtualization has been built up over time, we can keep making it better

AI Agent Sandboxing

- Automonous AI agents
  - “With great power comes great responsibility”
- What can an agent do inside a container?

Thank You

Marina Moore
Research Scientist
marina@edera.dev