

# Container Breakout: What it is, how it happens, and how to prevent host takeover

Gridinsoft Help Center

## What it is

A container breakout is when an attacker escapes a container (e.g., a Docker/Kubernetes pod) and reaches the host OS. From there, they can read sensitive files, tamper with other containers, or run malware with host-level access. (Escaping a virtual machine is similar but is usually called a VM escape.)

## How it happens

- Kernel bugs: containers share the host kernel; a vuln can break isolation.
- Over-privileged containers: --privileged, extra capabilities, or unsafe syscalls.
- Dangerous mounts: mounting /var/run/docker.sock, host dirs, or device files.
- Weak pod policies: no admission controls, no seccomp/apparmor, no read-only FS.

## What you might notice

- Unexpected host processes spawned by a container
- New mounts, modified binaries, or odd activity in /var/run, /etc, /proc
- Lateral movement between pods/nodes; secrets accessed from other workloads

## If you suspect a breakout

- Isolate the node from the network; cordon/drain in Kubernetes if possible.
- Capture evidence (logs, container and kernel events) before stopping containers.
- Rotate secrets (K8s secrets, cloud creds, registry tokens) from a clean admin host.
- Rebuild compromised nodes from clean images; don't just restart containers.

## Prevent it

- Least privilege: drop capabilities, avoid --privileged, use read-only FS and non-root users.
- Harden the host: keep kernel up to date; enable seccomp/AppArmor/SELinux.
- Safe mounts: never mount the Docker socket; restrict hostPath volumes.
- Policies & scanning: apply Pod Security Standards/OPA; scan images and block risky ones.
- Network & secrets: segment workloads; use minimal, scoped tokens and short-lived creds.