vNode

vNode Security Assessment by Cure53 Confirms Strong Isolation Against Container Breakouts

Table of Contents

Intro: What is vNode 2

How vNode Compares 2

Why Security Matters 3

Cure53 Independent Audit 4

Key Findings 4

Highlights from the Report 4

Summary of Results 5

Why This Matters 5

See vNode in Action 6

©2025 vCluster

Intro: What is vNode

vNode is a next-generation container runtime from vCluster Labs that provides strong, lightweight isolation at the node layer, even on bare metal.

Instead of relying on hypervisors or virtual machines, vNode uses an additional layer of Linux kernel namespaces and seccomp-based syscall filtering to isolate workloads running on the same Kubernetes node.

This design enables secure multi-tenancy and container breakout prevention with near-native performance, making it ideal for GPU-accelerated and bare-metal environments where VM-based approaches add too much overhead.

How vNode Compares

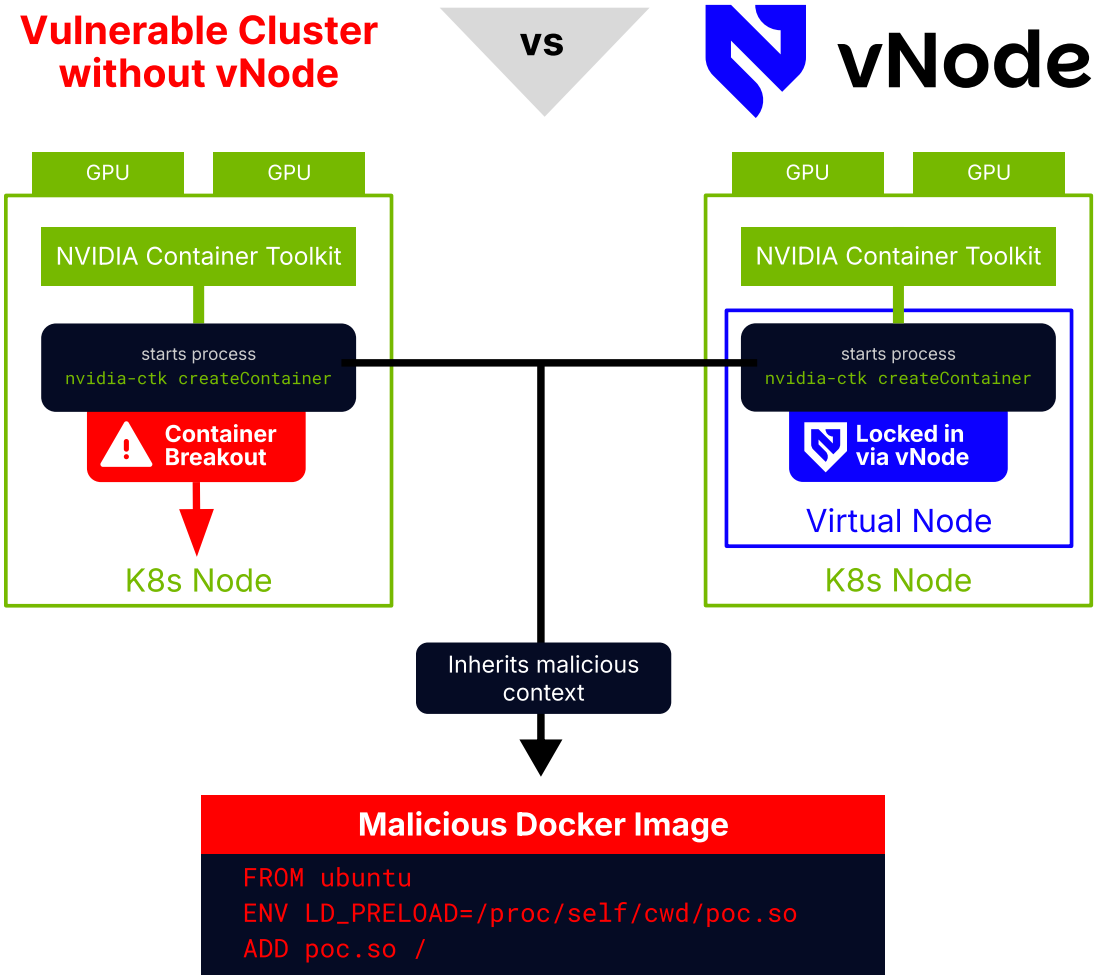
Unlike VM-based or micro-VM isolation technologies, vNode introduces a nested container runtime model combining VM-like isolation with container-level performance. The table below highlights how vNode differs from common runtime isolation tools used in Kubernetes environments:

	vNode	Kata Containers	gVisor	KubeVirt
Technology Type	Nested Container Runtime	Container Runtime	Container Runtime	VM-based Platform
Isolation Model	Containers run inside another container's user namspace	Isolates workloads in lightweight Micro-VMs	Isolates via syscall filtering (Seccomp)	Runs workloads as full virtual machines on separate nodes
Low Overhead	✓	✗	✗	✗
Fast Startup Time	✓	✗	✓	✗
Low Performance Impact	✓	✗	◐	✗
High Tenant Autonomy	✓	✓	✗	✓
High Security Strength	✓	✓	✓	✓
High Networking & Storage Isolation	✓	✓	✗	✓
Low Failure Blast Radius	✓	✗	◐	✓
Compatibility with Cloud Providers	✓	✗	◐	◐
Kubernetes Native	✓	✓	✓	◐
Ease of Use	✓	✗	✗	✗
Commercial Support & Maintenance	✓	✗	✗	✓

Why Security Matters

Container breakouts, where a compromised container gains access to the host or other workloads, represent one of the most serious risks in multi-tenant Kubernetes environments.

Recent high-profile exploits like NVIDIAScape have shown how vulnerable traditional container isolation can be.



vNode was built to close this gap, providing true node-level isolation for containerized and GPU workloads on bare metal, without the overhead of hypervisors.

Read blog →

Cure53 Independent Audit

In September 2025, vCluster Labs engaged Cure53, a globally recognized cybersecurity firm, to perform a white-box pentest and source code audit of vNode's core system and integrations with runc and sysbox.

The audit included over 30 days of testing by seven senior security researchers, covering vNode's runtime architecture, isolation boundaries, syscall filtering, and container escape resistance.

Key Findings

Cure53's report confirms that vNode's isolation design is effective and resilient, with no container escapes identified during testing.

Proven Protection Against Container Breakouts

"Cure53 was unable to identify any container escapes during the assessment, therefore the security posture of vNode can be described as impressive."

Engineered for Security and Stability

"Overall, Cure53 found the modifications of *runc* to be minimal and excellently engineered."

Defense in Depth at the Syscall Layer

"Seccomp filters for allowed syscalls were analyzed" and "showed that harmful syscalls are filtered out, including syscalls which allow loading kernel modules or accessing the hard disk directly."

Strong Runtime Isolation Between Containers

"The accessibility of this TTRPC endpoint was tested, but all attempts to read/write to the socket from other containers failed."

Secure by Design — Authenticated and Least Privilege

"Communication between the vNode container and the Kubernetes control plane"..."verified that this communication is properly authenticated" and "analysis of the associated service account showed that it is provisioned with minimal permissions and has no capacity for modifying states."



Highlights from the Report

Together, these findings confirm that vNode delivers strong, lightweight isolation on bare metal and GPU infrastructure, preventing container breakouts without the need for VMs or hypervisors.

©2025 vCluster

Summary of Results

- No container escapes found during Cure53's extensive testing
- Effective namespace and syscall isolation verified
- Strong authentication and limited permissions between vNode and Kubernetes control plane
- Minimal code modifications, audited as "excellently engineered"
- Impressive overall security posture confirmed by independent experts

Why This Matters

This independent audit confirms that vNode delivers what modern Kubernetes workloads need most:



"Security-strengthened isolation without the drag of virtualization."

vNode helps prevent breakout scenarios like NVIDIAScape-class exploits before they reach the host, giving platform teams confidence to run high-stakes, GPU-accelerated, and multi-tenant workloads securely on bare metal.

See vNode in Action

Proven breakout prevention.
Lightweight isolation.
Built for Kubernetes and GPUs.

Sign up for the vNode Beta →



vNode