Varnish Launches Artifact Firewall, Expands Orca for Kubernetes and AI Workloads

Runtime layer accelerates and governs software and AI artifacts across distributed Kubernetes environments.

AMSTERDAM, NETHERLANDS, March 24, 2026 /EINPresswire.com/ -- Varnish Software today announced the launch of Artifact Firewall, a new runtime enforcement layer for artifact infrastructure that allows organizations to control and govern artifact access at request time. Artifact Firewall can be deployed independently in front of existing artifact repositories or alongside The Varnish Virtual Registry (Orca) to combine policy enforcement with high-performance artifact caching and routing. The company will demonstrate both solutions at KubeCon + CloudNativeCon Europe in Amsterdam.

> “Artifact Firewall adds the ability to enforce security policy at the moment artifacts are requested, which becomes essential when software and AI pipelines are operating at this kind of scale.”
>
> *Adrian J Herrera, GM Data & Storage at Varnish Software*

Varnish is already deployed in production environments to accelerate container pulls, dependency resolution, and large artifact distribution across distributed Kubernetes and GPU-driven workloads. With the addition of Artifact Firewall and S3 optimization capabilities, Varnish now provides a unified runtime layer for artifact performance, policy enforcement, and economic control.

As enterprises scale Kubernetes platforms and AI pipelines, artifact traffic has become larger, more distributed, and more economically volatile. Traditional repository architectures were designed to store artifacts, not to manage how they move across distributed environments. Increasingly, customers are using Orca as a runtime layer that accelerates, routes, and governs artifact traffic before it reaches workloads. This runtime layer operates between artifact repositories and Kubernetes or AI workloads, optimizing artifact delivery without replacing existing infrastructure.

Varnish Orca operates in front of Docker and OCI registries, JFrog Artifactory, Sonatype Nexus, npm, Maven, PyPI, S3, and other artifact repositories. Rather than replacing them, Orca transforms dependency traffic into a resilient caching and high-performance routing layer.
- In large-scale production environments, customers have achieved:
- Up to 80% reduction in artifact latency

- 41–47% faster dependency resolution
- Up to 95% latency reduction for distributed development teams
- 50% reduction in repository licensing footprint
- Significant reductions in NAT gateway, egress, and backend CPU costs

The same architecture is already used in production to accelerate AI training artifacts and model build pipelines, reducing redundant pulls, optimizing cross-region distribution, and improving execution performance in GPU-based environments.

AI models are often hundreds of gigabytes in size. Orca reduces the cost and latency of distributing these models across clusters by caching them closer to the workloads that need them, while Artifact Firewall ensures only approved or verified versions can be pulled into production environments.

Introducing Artifact Firewall for Runtime Governance
Artifact Firewall enforces policy at request time and can operate independently or alongside Varnish Orca. Unlike repository-centric scanners that evaluate artifacts after ingestion, Artifact Firewall evaluates every artifact request before delivery. This ensures vulnerable or non-compliant artifacts can be blocked or hidden consistently across distributed environments.

Key capabilities include:
- Blocking known compromised package versions using package identity and semantic version policies
- Enforcing configurable “package quarantine” delays before new versions become available
- Preventing namespace-based dependency confusion attacks
- Rewriting manifests to hide unsafe versions
- Generating structured audit logs for compliance and investigation
- Managing distributed rulesets via Git

As artifact traffic grows across distributed Kubernetes clusters and AI training environments, enforcing policy solely within centralized repositories becomes increasingly difficult. Artifact Firewall enforces security policy at request time, allowing organizations to evaluate and govern artifacts at the moment they are requested, rather than after they have already propagated through pipelines.

Together, Orca, Artifact Firewall, and S3 optimization capabilities provide runtime control over artifact performance, governance, and infrastructure economics across distributed environments.

“Modern software supply chains now power both application delivery and AI training pipelines,” said Adrian J Herrera, GM Data & Storage at Varnish Software. “Many teams are discovering that centralized artifact architectures struggle to keep up with distributed Kubernetes environments and GPU-driven workloads. Customers are using Orca not just to speed up container pulls, but

to control how artifacts move across their infrastructure. With Artifact Firewall, we're adding the ability to enforce security policy at the moment artifacts are requested, which becomes essential when software and AI pipelines are operating at this kind of scale."

The Varnish Virtual Registry Orca is free to try and available at: https://www.varnish-software.com/orca.

Jenny Lake
Varnish Software
email us here
Visit us on social media:
LinkedIn
Facebook
YouTube
X
