

Artifactory Optimization with Varnish Enterprise

Improve the Artifactory experience and reduce friction, complexity and costs with industry-leading content delivery software

 80%

Lower Latency

 50%

Reduction in
Licensing Costs

 41%

Faster Dependency
Resolution

 WEEKS

of Developer
Time Saved

The repository manager JFrog Artifactory is found in the DevOps workflows of many enterprises, and its speed, reliability and usability are business critical. Bottlenecks, performance degradation or platform instability can hinder productivity and add costs, while licensing challenges add complexity to site-wide deployments.

Optimizing Artifactory is about improving object delivery and maximizing ROI to get the most out of the investment. Speeding up the retrieval and delivery of artifacts and strengthening system reliability at scale can boost the Artifactory experience while minimizing overheads.

Streamline and Scale Artifactory

Varnish Enterprise content delivery software speeds up digital interactions, handles massive traffic load, and protects infrastructure. Built on top of the feature-rich and robust open-source Varnish Cache and designed for efficiency at enterprise scale, Varnish Enterprise caches artifacts and delivers them at very low latency while handling thousands of requests per second, allowing Artifactory to run at a heavily reduced thread count.

Unlike alternative solutions, Varnish can cache all responses, including previously uncacheable content such as directory listings and 404s, making it the single source of truth for all requests. Varnish also stops upstream performance issues affecting response times, reduces load while revalidating, and ensures secure low-latency access for remote workers. With tight integration into DevOps and CI/CD pipelines, and fast time-to-implementation across cloud, on-premise and hybrid environments, Varnish Enterprise is an efficient pathway towards Artifactory ROI and improved productivity.

Accelerate Workflows

Cache database request traffic and improve dependency resolution time in Gradle and NPM.

Improve Productivity

Ensure cleanups are done on obsolete artifacts that are causing system delays.

Reduce Latency, Enable Scale

Accelerate retrieving and storing of artifacts at large scale and shield database from traffic spikes.

Gain Flexibility

Add additional layers to allow repository access and omissions at a granular level, without cost escalation.

Varnish Enterprise Features for Artifactory Optimization

Persisted disk caching: Extend caching capacity in a cost-effective and seamless way, and persist caches during system restarts.

Origin protect: Combine multiple requests for the same response and send a single request to the origin before satisfying all requests in parallel, for lower latency and improved scalability.

Varnish Configuration Language: Extensive logic for caching personalization allows teams to define and control their Artifactory caching strategy.

Authorization at scale: Authorize users against Artifactory for each requested resource, before serving the response from cache, with the authorization itself also cached if required.

Fast, customizable cache invalidation: Execute thousands of cache purges in milliseconds, ensure cached objects are fresh and out-of-date artifacts removed.

Fine-grained caching controls: Expose the underlying response repository, artifact type and mutability.

Extensive dashboarding: Monitoring and alerts improve operability over other approaches.

Platform and location agnostic: Place caching resources in multiple regions, on any platform, to meet local office needs.

Signed responses: Provide signed 307 redirect responses directly to immutable S3 blob for an artifact, allowing blobs to be cached indefinitely and served to the client, reducing load.

Contact us for more information on how tier 1 enterprises across media, finance and technology verticals are optimizing Artifactory with Varnish Enterprise.

Find out more at <https://www.varnish-software.com/contact-us/>