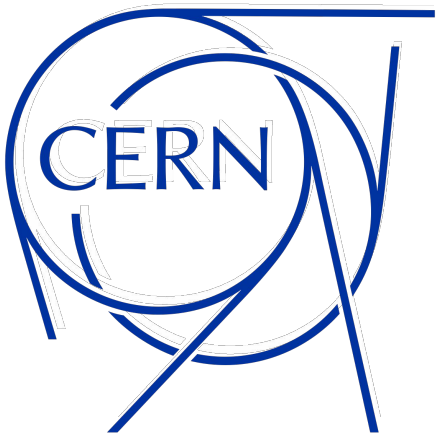

Managing your images with scans, immutability and tag retention

Ankur Kothiwal

10th October, 2024

Harbor registry

Harbor is an open source registry that secures artifacts with policies and role-based access control, ensures images are scanned and free from vulnerabilities, and signs images as trusted.

Harbor is a Cloud Native Computing Foundation Graduated project

Why are we managing our own Registry?

- **Control over full software supply chain**
- **Control over security policies**
  - Vulnerability scanning and security policies
  - Access control
  - Images Signing
- **Custom policy management**
  - Tag retention policy
  - Image immutability
- **Eliminates pull rate limits**
- **Garbage collection for better storage management**

Harbor features

- **Policy**
  - Image scanning
  - Immutability
  - Tag retention
  - Quotas
  - Sign Artifacts
- **Multitenancy**
  - RBAC and Project Isolation
- **Garbage Collection**
- **SBOM (will be fully supported in the next release)**

Image Scanning

- **Vulnerability Scanning:**

  Harbor provides static analysis of vulnerabilities in images through open source project Trivy (default)

- **Configuring Scan Policies:**

  Harbor allows administrators to create vulnerability policies that prevents images from being deployed or pushed if they do not meet specific criteria (e.g., no "Critical" or "High" severity vulnerabilities)

- **Notifications and Webhooks:**

  Harbor supports **webhooks** and notifications for image scanning events. This enables teams to be alerted whenever vulnerabilities are found, allowing for quicker remediation.

Immutability

By default, users can repeatedly push an image with the same tag to repositories in Harbor. This causes the previous image to effectively be overwritten with each push, in that the tag now points to a different image and the image that previously used the tag now becomes tagless.

Tag immutability ensures that once an image is created and tagged, it cannot be overwritten or modified. This prevents accidental or malicious overwrites of images.

Tag Retention

Tag retention in Harbor is a feature designed to manage storage and prevent clutter in the registry by automatically deleting old, unused, or irrelevant image tags based on defined policies.

As a Harbor system administrator, you can define rules on repositories that govern how many artifacts of a given repository to retain, or for how long to retain certain artifacts.

For eg.

- Keep the latest N images
- Delete images older than a certain date
- Retain images with specific tags (such as with **stable-***)

Sign Artifacts

Harbor allows developers to sign the artifacts that they push to the registry thus ensuring authenticity and integrity.

Harbor supports Notary and **Cosign** services for it.

Garbage Collection

Garbage collection in Harbor is designed to clean up and reclaim storage by removing unreferenced or unused layers of container images.

Harbor allows administrators to **schedule** garbage collection tasks.

Notification and Webhook

Webhook notifications allow to trigger external services when specific events occur in the Harbor registry.

These notifications are useful for integrating Harbor with other tools or workflows, such as CI/CD pipelines, monitoring systems, or alerting mechanisms.

For example: Send notification on Mattermost if when project is about to exceed its quota

SBOM (coming in the next release)

Software Bill of Materials (SBOM) acts as an inventory list, documenting all components used in a software project. It provides transparency by listing dependencies, their versions, and associated licenses present it the software or container image.

Since version 2.11 Harbor supports now automatic generation of SBOMs in combination with its default scanner - Trivy.

Addition to that, users can also click the `GENERATE SBOM` button to manually generate an SBOM of a given artifact.

Best Practices

- **Fix critical vulnerabilities**: Act on vulnerabilities by patching the base images or updating dependencies.
- **Enforce policies**: Consider setting policies that prevent the deployment of images with critical vulnerabilities.
- **Enable image immutability** on critical projects to prevent overwriting tags.
- Set **tag retention policies** to automatically delete old or unnecessary image tags and free up storage.
- Sign Artifacts
- Always pull and run **trusted, signed images** to minimize security risks, particularly in production environments
- **Assign roles based on necessity**: Harbor offers roles such as project admin, developer, guest, and more. Assign users only the permissions they need to minimize risks and adhere to the principle of **least privilege**.

For more info please check: https://kubernetes.docs.cern.ch/docs/registry/

Thank you

home.cern

Backup

Garbage Collection

Over time, as images are pushed and deleted, some layers may no longer be used but remain in the storage as they might be referenced by other images. This leads to inefficient use of resources. GC helps to reclaim this space by removing these orphaned objects.