

Container Support - Story #4174

As a user, I can lazy sync a docker repository

11/19/2018 06:30 PM - amacdona@redhat.com

Status:	CLOSED - CURRENTRELEASE	Start date:	
Priority:	Normal	Due date:	
Assignee:	dkliban@redhat.com	% Done:	100%
Category:		Estimated time:	0:00 hour
Sprint/Milestone:	1.0.0	Tags:	
Platform Release:		Sprint:	Sprint 52
Groomed:	Yes	Quarter:	
Sprint Candidate:	Yes		
Description			
A docker repository is composed from tags, manifests and layers.			
We can lazily fetch layers, manifests treated as metadata-like information will be always downloaded to discover the layers.			
Related issues:			
Related to Pulp - Story #4209: As a plugin writer, I can implement lazy sync ...		CLOSED - CURRENTRELEASE	
Related to Pulp - Test #4126: Test sync with different download policies		CLOSED - COMPLETE	
Related to Pulp - Test #4973: Docker - Test sync with different download poli...		NEW	

Associated revisions

Revision 75fc7777 - 04/24/2019 03:14 PM - dkliban@redhat.com

Problem: docker plugin does not support lazy sync

Solution: add support for lazy sync

This patch adds ability to create remotes with policy='on_demand'. When such a remote is used to sync a repository, the Blobs are not downloaded. Only manifests and manifest lists are saved at sync time. Blobs are retrieved on demand from clients.

This patch also includes a set of tests that check that you can pull content from a repo that was synced using a remote that has policy set to 'on_demand'. The tests assert that performing a docker/podman pull results in more artifacts being stored in Pulp than before the image was requested.

closes: #4174 <https://pulp.plan.io/issues/4174>

Revision 75fc7777 - 04/24/2019 03:14 PM - dkliban@redhat.com

Problem: docker plugin does not support lazy sync

Solution: add support for lazy sync

This patch adds ability to create remotes with policy='on_demand'. When such a remote is used to sync a repository, the Blobs are not downloaded. Only manifests and manifest lists are saved at sync time. Blobs are retrieved on demand from clients.

This patch also includes a set of tests that check that you can pull content from a repo that was synced using a remote that has policy set to 'on_demand'. The tests assert that performing a docker/podman pull results in more artifacts being stored in Pulp than before the image was requested.

closes: #4174 <https://pulp.plan.io/issues/4174>

Revision 75fc7777 - 04/24/2019 03:14 PM - dkliban@redhat.com

Problem: docker plugin does not support lazy sync

Solution: add support for lazy sync

This patch adds ability to create remotes with policy='on_demand'. When such a remote is used to sync a repository, the Blobs are not downloaded. Only manifests and manifest lists are saved at sync time. Blobs are retrieved on demand from clients.

This patch also includes a set of tests that check that you can pull content from a repo that was synced using a remote that has policy set to 'on_demand'. The tests assert that performing a docker/podman pull results in more artifacts being stored in Pulp than before the image was requested.

closes: #4174 <https://pulp.plan.io/issues/4174>

Revision 75fc7777 - 04/24/2019 03:14 PM - dkliban@redhat.com

Problem: docker plugin does not support lazy sync

Solution: add support for lazy sync

This patch adds ability to create remotes with policy='on_demand'. When such a remote is used to sync a repository, the Blobs are not downloaded. Only manifests and manifest lists are saved at sync time. Blobs are retrieved on demand from clients.

This patch also includes a set of tests that check that you can pull content from a repo that was synced using a remote that has policy set to 'on_demand'. The tests assert that performing a docker/podman pull results in more artifacts being stored in Pulp than before the image was requested.

closes: #4174 <https://pulp.plan.io/issues/4174>

History

#1 - 11/20/2018 03:06 PM - ipanova@redhat.com

- Description updated

#2 - 11/20/2018 05:00 PM - amacdona@redhat.com

- Description updated

Currently, the pulp3 lazy sync implementation plan is to disable the ArtifactDownload and ArtifactSaver stages. This will not work for Docker since these stages are used to download and save metadata (tags, manifests, and manifest lists) as well as blobs/layers (which would not be downloaded or saved for lazy).

We have a couple options to implement this for docker, and it may also be possible to make a contribution to pulpcore to enable easy implementation for other plugins with multiple metadata files.

Idea 1: Add a field to the DeclarativeArtifact

Using a new field, DeclarativeArtifact.metadata (defaults to false), the ArtifactDownload and ArtifactSaver stages could potentially be updated to skip them when da.metadata=false and lazy=true.

Idea 2: Override ArtifactDownload and ArtifactSaver stages

If we override, we just have to add a little logic that checks the type of the DeclarativeContent.content. If lazy and type is Blob, don't download or save.

#3 - 12/05/2018 04:04 PM - amacdona@redhat.com

- Related to Story #4209: As a plugin writer, I can implement lazy sync for plugins with multiple metadata files added

#4 - 04/18/2019 05:12 PM - ipanova@redhat.com

- Groomed changed from No to Yes

- Sprint Candidate changed from No to Yes

- Sprint set to Sprint 52

#5 - 04/18/2019 09:53 PM - dkliban@redhat.com

- Status changed from NEW to ASSIGNED

- Assignee set to dkliban@redhat.com

#6 - 04/18/2019 10:41 PM - dkliban@redhat.com

- Status changed from ASSIGNED to POST

https://github.com/pulp/pulp_docker/pull/344

#7 - 04/24/2019 06:17 PM - dkliban@redhat.com

- Status changed from POST to MODIFIED

- % Done changed from 0 to 100

Applied in changeset commit:pulp_docker|75fc777756b95e1df5da7bd2b08a32c2945742dc.

#8 - 04/26/2019 10:33 PM - bmbouter

- Tags deleted (Pulp 3)

#9 - 04/29/2019 11:26 PM - kersom

- Related to Test #4126: Test sync with different download policies added

#10 - 06/14/2019 08:19 PM - kersom

- Related to Test #4973: Docker - Test sync with different download policies added

#11 - 11/20/2019 07:21 PM - ipanova@redhat.com

- Project changed from Docker Support to Container Support

#12 - 12/12/2019 11:39 PM - ipanova@redhat.com

- Status changed from MODIFIED to CLOSED - CURRENTRELEASE

#13 - 01/22/2020 08:00 PM - ipanova@redhat.com

- Sprint/Milestone set to 1.0.0