



# dCache CI/CD migration to Kubernetes

(from our keyboards to your servers)

*Tigran Mkrtchyan for the dCache collaboration,  
DESY-IT GitLab & Kubernetes Teams*



**HELMHOLTZ**

RESEARCH FOR  
GRAND CHALLENGES



Commit 9e9ea4fa authored 6 days ago by Marina Sahakyan [Browse files](#) [Options](#)

**maven-release-plugin] prepare release 9.2.26**

parent [66643955](#)

Branches [Branches containing commit](#)

Tags [9.2.26](#)

No related merge requests found

✓ Pipeline #136231 passed with stages ✓✓✓✓✓✓✓✓ in 52 minutes and 22 seconds

Changes [64](#) Pipelines [1](#)

Status	✓ Passed ⌚ 00:52:22 📅 6 days ago
Pipeline	<a href="#">[maven-release-plugin] prepare release 9.2.26</a> #136231 <a href="#">9.2.26</a> <a href="#">9e9ea4fa</a> <a href="#">latest</a>
Stages	✓✓✓✓✓✓✓✓
Actions	<a href="#">↓</a>



Commit 9e9ea4fa authored 6 days ago by Marina Sahakyan

**[maven-release-plugin] prepare release 9.2.26**

parent 66643955

Branches > Branches containing commit

Tags 9.2.26

No related merge requests found

Pipeline #136231 passed with stages in 59 minutes and 22 seconds

Changes 64 Pipelines 1

Status	Passed	00:52:22	6 days ago
Pipeline	[maven-release-plugin] prepare release 9.2.26	#136231	9.2.26 9e9ea4fa latest
Stages	[Progress indicators]		
Actions	[Download icon]		

## Release 9.2.X

dCache 9.2 is a Golden Release introducing following highlights:

- Performance improvements for the concurrent directory creation and removal

### Incompatibilities

- prior to version 9.2 dCache NFSv4.1 door will publish only `nfs4_1_files` layout. Now on the door publishes all available layout types
- dropped reference count tracking for directory tags
- dropped support of java options `chimera_soft_update` and `chimera_lazy_wcc` in favor of `chimera.attr-consistency` property

dCache v9.2 requires a JVM supporting Java 11 or Java 17.

Download	Rel. Date	md5 hash	Release Notes
<b>dCache 9.2.26 (Debian package)</b>	01.10.2024	544d130787a4b77cb18bb3e417c5d7ee	<b>9.2.26</b>
<b>dCache 9.2.26 (rpm)</b>	01.10.2024	c728784cfb432f38577bc783c963c59e	
<b>dCache 9.2.26 (tgz)</b>	01.10.2024	bc7cd6cb3944ffb584633ac4b4cb867a	
<b>dCache 9.2.25 (Debian package)</b>	11.09.2024	ea99e8f167dcc68ad03c35ea7dc78471	<b>9.2.25</b>
<b>dCache 9.2.25 (rpm)</b>	11.09.2024	2ef6382f7edbe991c8523cac5fcd357	
<b>dCache 9.2.25 (tgz)</b>	11.09.2024	09e79abc408f0572264d5f2f19f33ab8	



Commit 9e9ea4fa authored 6 days ago by Marina Sahakyan

**[maven-release-plugin] prepare release 9.2.26**

parent 66643955

Tags 9.2.26

Pipeline #136231 passed with stages in 59 minutes and 22 seconds

Changes 64 Pipelines 1

Status	Passed	00:52:22
Pipeline	[maven-release-plugin] prepare release 9.2.26	#136231 9.2.26 9e9ea4fa latest
Stages		✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓ ✓
Actions		↓

## Release 9.2.X

dCache 9.2 is a Golden Release introducing following highlights:

- Performance improvements for the concurrent directory creation and removal

### Incompatibilities

- prior to version 9.2 dCache NFSv4.1 door will publish only `nfs4_1_files` layout. Now on the door publishes all available layout types
- dropped reference count tracking for directory tags
- dropped support of java options `chimera_soft_update` and `chimera_lazy_wcc` in favor of `chimera.attr-consistency` property

dCache v9.2 requires a JVM supporting Java 11 or Java 17.

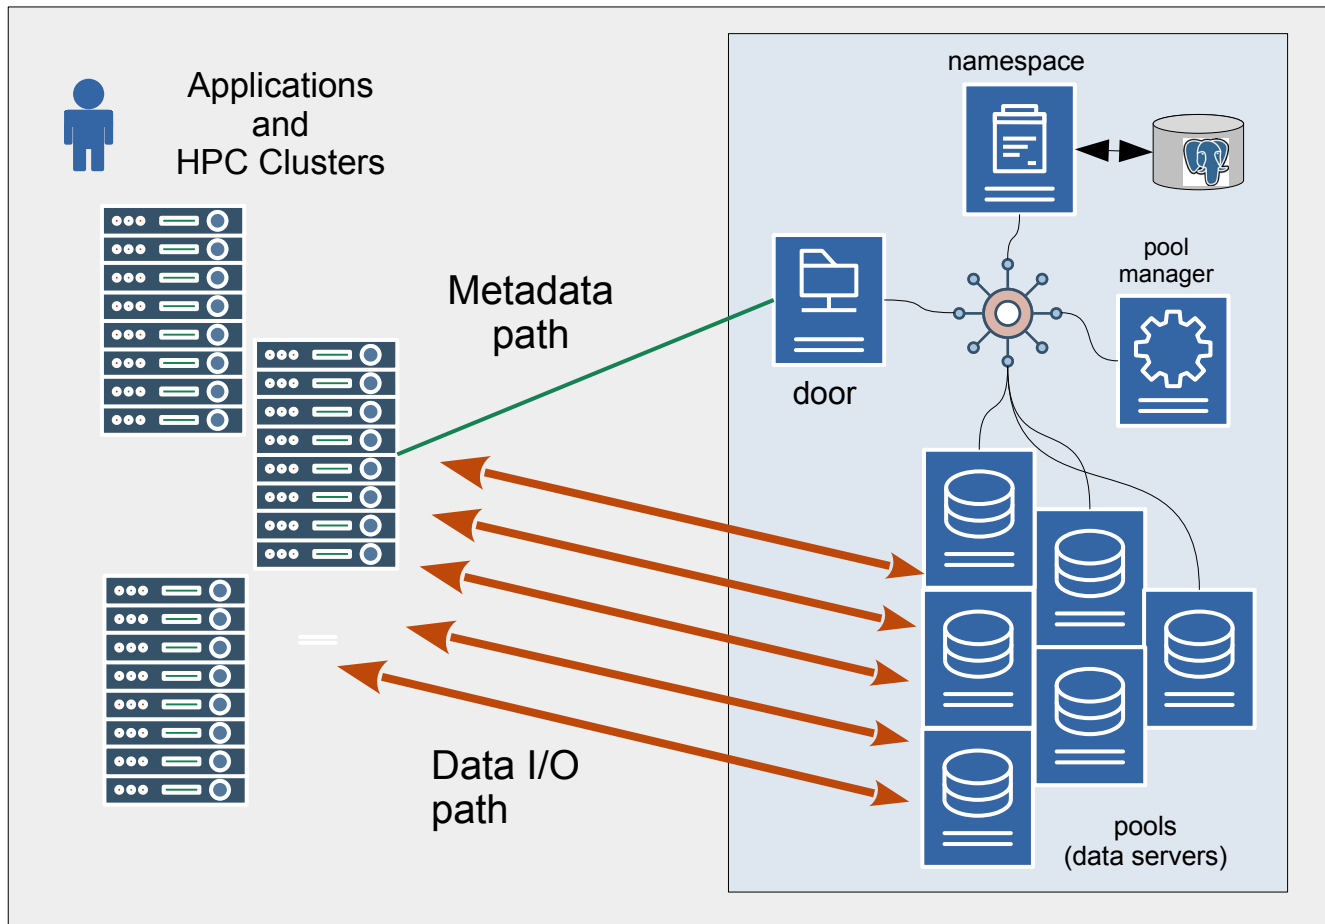
Download	Rel. Date	md5 hash	Release Notes
<b>dCache 9.2.26 (Debian package)</b>	01.10.2024	544d130787a4b77cb18bb3e417c5d7ee	<b>9.2.26</b>
<b>dCache 9.2.26 (rpm)</b>	01.10.2024	c728784cfb432f38577bc783c963c59e	
<b>dCache 9.2.26 (tgz)</b>	01.10.2024	bc7cd6cb3944ffb584633ac4b4cb867a	
<b>dCache 9.2.25 (Debian package)</b>	11.09.2024	ea99e8f167dcc68ad03c35ea7dc78471	<b>9.2.25</b>
<b>dCache 9.2.25 (rpm)</b>	11.09.2024	2ef6382f7edbe991c8523fac5fcd357	
<b>dCache 9.2.25 (tgz)</b>	11.09.2024	09e79abc408f0572264d5f2f19f33ab8	

# Why this Good to Know



- Get to know what we test and what we don't
- Re-use our setup on your testbed
  - Get to know new functionality
- Re-run our test for your custom builds
- Extend our tests with your test case
  - Add your site setup

# Quick into to dCache

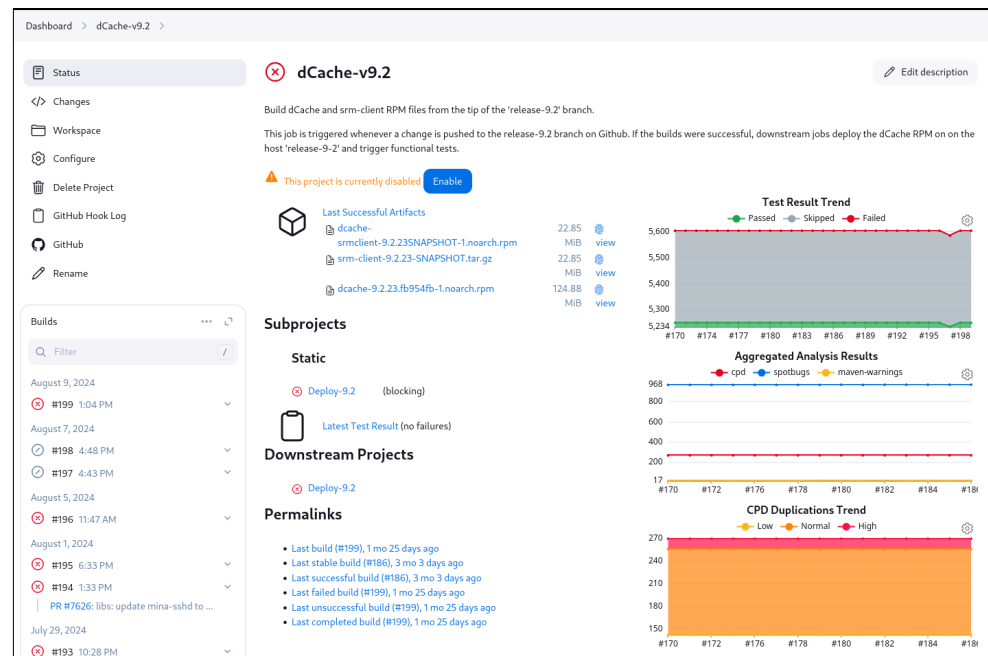


- Distributed Storage systems
  - Client redirection
  - Incoming connection port ranges
- Scales by redirecting clients to desired data server
- Uses 3<sup>rd</sup>-party components
  - ZooKeeper for auto-discovery
  - PostgreSQL Database for metadata server
  - Kafka for event propagation
  - Tape storage for off-line files
- Covers multiple data management use cases
  - Data placement policies
  - End-user analysis
  - Data import/export
  - ...
- Supports multiple AuthN schema
  - Username+password
  - X509 certificates
  - OIDC tokens
  - ...

# Testing Environment (pre-migration)



- Jenkins (since 2007)
- 20 VMs XEN/OpenStack
  - 4 per supported branch
    - Deploy host
    - Grid worker node
    - SRM test suite
    - NFS test suite
- Own X509 CA
- 3<sup>rd</sup> party components in containers




# Testing Environment



openstack®




CI & Agents



User fronting dCache services



dCache internal services



External services used by dCache





# Testing Environment



kubernetes



CI & Runners



User fronting dCache services



dCache internal services



External services used by dCache



# GitHub ⇒ CI (Jenkins) Integration



- Clear roles
  - GitHub: Code repo
  - Jenkins: CI
- Web hook on `push/pull_request`

**Release 9.2.X**

dCache 9.2 is a Golden Release introducing following highlights:

- Performance improvements for the concurrent directory creation and removal

**Incompatibilities**

- prior to version 9.2 dCache NFSv4.1 door will publish only `nf4_1_files` layout. Now on the door publishes all available layout types
- dropped reference count tracking for directory tags
- dropped support of java options `chimera_soft_update` and `chimera_lazy_wcc` in favor of `chimera.attr-consistency` property

dCache v9.2 requires a JVM supporting Java 11 or Java 17.

Download	Rel. Date	md5 hash	Release Notes
<a href="#">dCache 9.2.26 (Debian package)</a>	01.10.2024	544d130787a4b77cb18bb3e417c5d7ee	
<a href="#">dCache 9.2.26 (rpm)</a>	01.10.2024	c728784cfa432f98577bc783c963c59e	<a href="#">9.2.26</a>
<a href="#">dCache 9.2.26 (tgz)</a>	01.10.2024	bc7cd6cb944ffb584633ac4b4cb867a	
<a href="#">dCache 9.2.25 (Debian package)</a>	11.09.2024	ea99e8f167dcc68ad03c35ea7dc78471	
<a href="#">dCache 9.2.25 (rpm)</a>	11.09.2024	2ef63827eedbe991c8523fcac5fcd357	<a href="#">9.2.25</a>
<a href="#">dCache 9.2.25 (tgz)</a>	11.09.2024	09e79abc408f0572264d52f19f33ab8	

# GitHub ⇒ CI (GitLab) Integration



- GitHub → GitLab mirroring
  - GitHub action to synchronize code
- Two public repos
  - GitHub is the ‘official’



**Release 9.2.X**

dCache 9.2 is a Golden Release introducing following highlights:

- Performance improvements for the concurrent directory creation and removal

**Incompatibilities**

- prior to version 9.2 dCache NFSv4.1 door will publish only `nfs4_1_files` layout. Now on the door publishes all available layout types
- dropped reference count tracking for directory tags
- dropped support of java options `chimera_soft_update` and `chimera_lazy_wcc`. In favor of `chimera.attr-consistency` property

dCache v9.2 requires a JVM supporting Java 11 or Java 17.

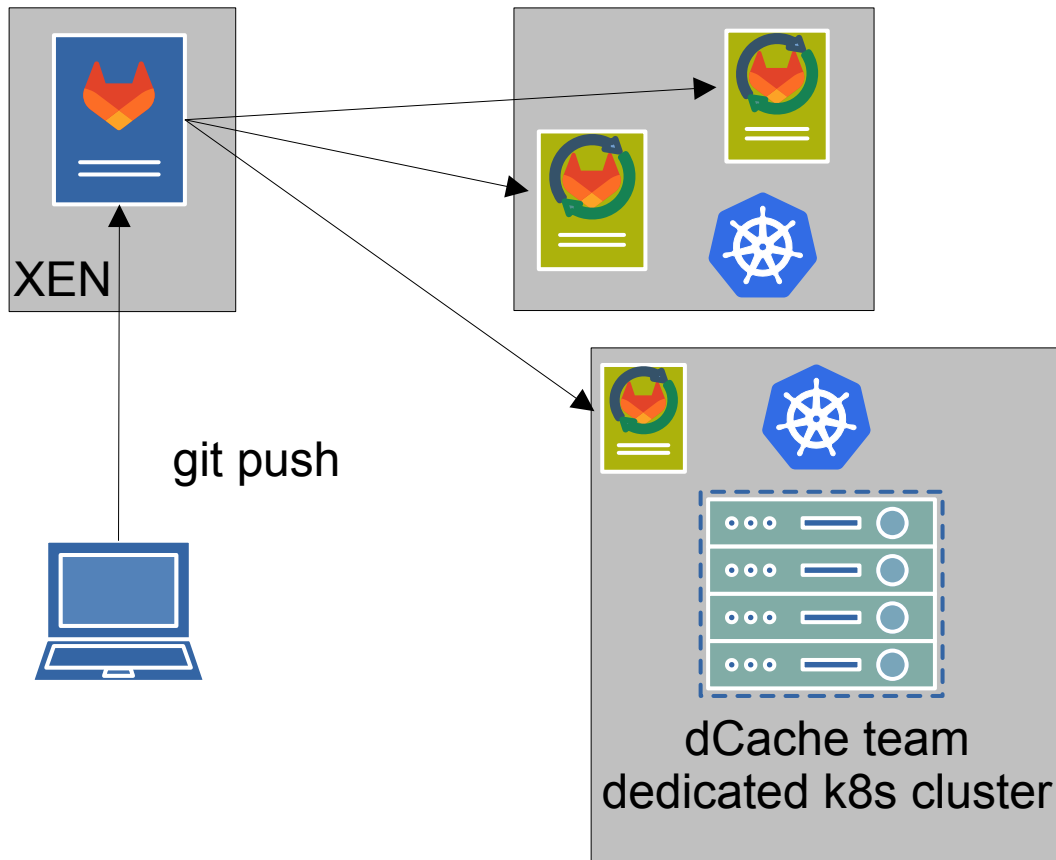
Download	Rel. Date	md5 hash	Release Notes
<a href="#">dCache 9.2.26 (Debian package)</a>	01.10.2024	544d130787a4b77cb18bb3e417c5d7ee	<a href="#">9.2.26</a>
<a href="#">dCache 9.2.26 (rpm)</a>	01.10.2024	c728784cb432f38577bc783c963c59e	<a href="#">9.2.26</a>
<a href="#">dCache 9.2.26 (tgz)</a>	01.10.2024	bc7cfdcb3944ff584633ac4b4cb867a	<a href="#">9.2.26</a>
<a href="#">dCache 9.2.25 (Debian package)</a>	11.09.2024	ea99e8f167dccb68ad03c35ea7dc78471	<a href="#">9.2.25</a>
<a href="#">dCache 9.2.25 (rpm)</a>	11.09.2024	2ef6382f7edbe991c8523fca5fcd357	<a href="#">9.2.25</a>
<a href="#">dCache 9.2.25 (tgz)</a>	11.09.2024	09679abc408f0572264d5f2f19f33ab8	<a href="#">9.2.25</a>

# GitLab Infrastructure at DESY



GitLab central service

Shared runners



- DESY-IT provided GitLab service
  - In production since 2020
- Shared runners for non k8s jobs
- Special runner for k8s-aware jobs
  - dCache-team dedicated cluster
  - Network isolation
  - Access privileges

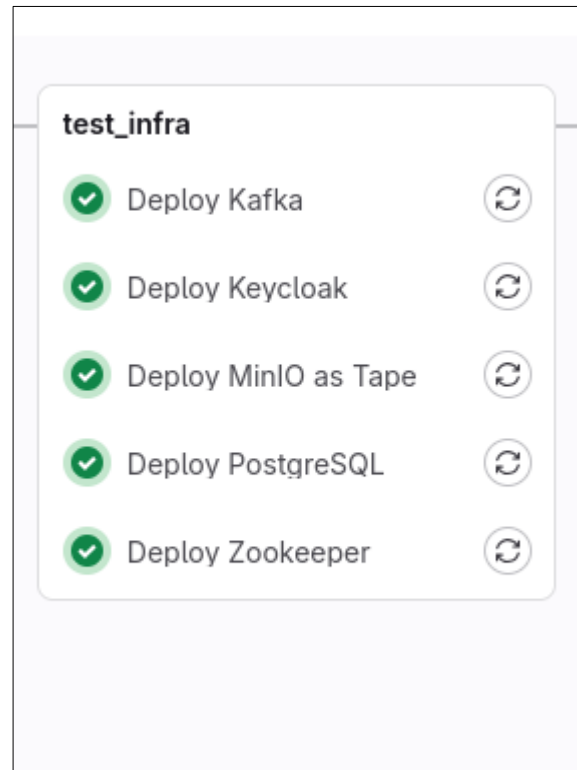


stages:

- build *# build rpm, tar, deb, oci*
- sign *# PGP sign rpms*
- testenv\_pre *# prepare k8s env*
- test\_infra *# deploy pg, zk, ...*
- test\_deploy *# deploy dcache helm*
- testing *# run all tests*
- testenv\_post *# collect logs*
- upload *# publish packages*



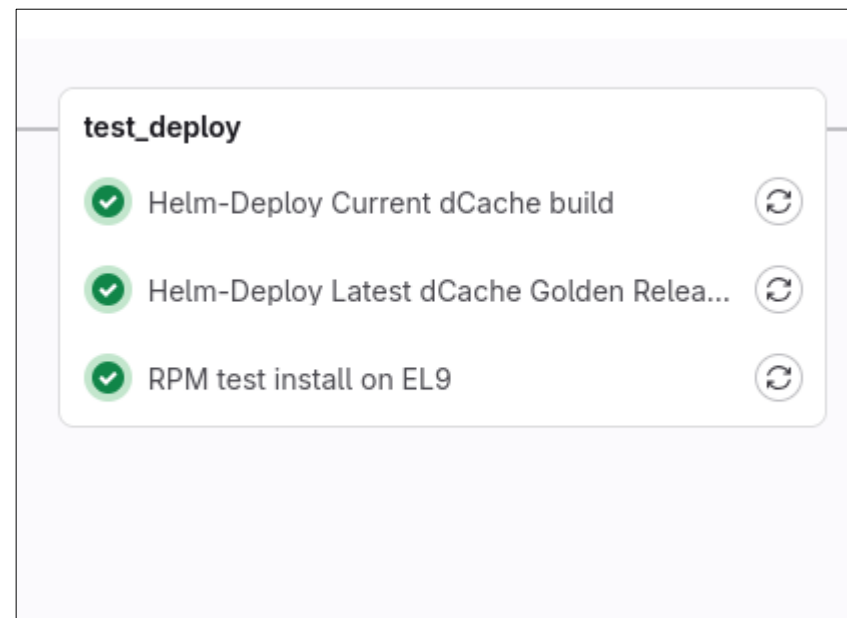
- Mimic a small site
  - Tape library
    - MinIO
  - IdP
    - Keycloak
  - X509-CA
  - PostgreSQL
  - ZooKeeper
  - Kafka
  - Worker nodes



# Test Deployment



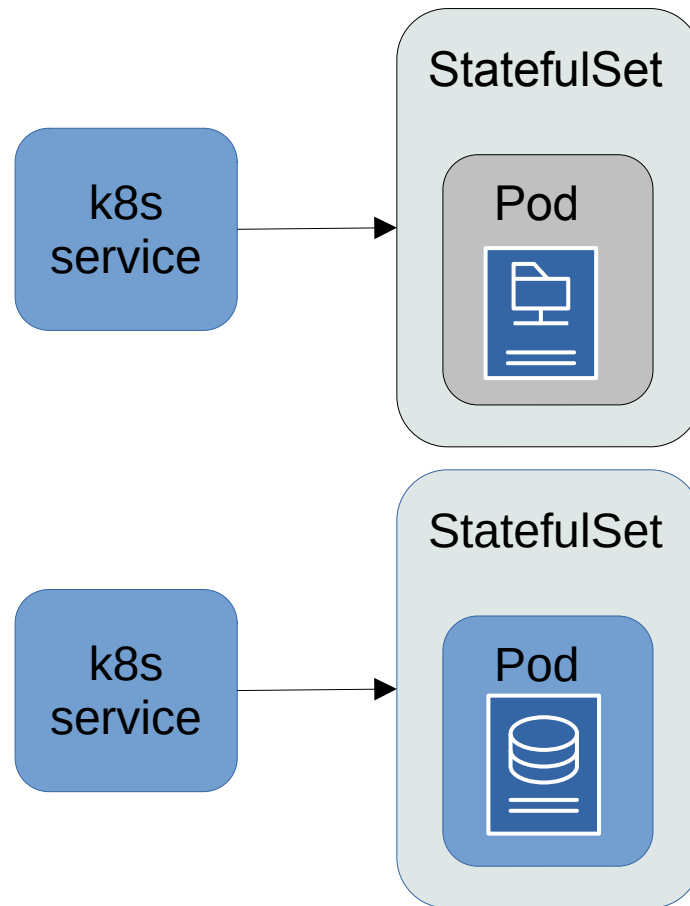
- Full dCache deployment for a typical site
- Version-mix for backward compatibility test



# dCache K8S services



- Each door published as a service
- Each protocol on a pool is published as a service
- Port range exposed as a service per port
- Pools expose themselves by service name





# Happy YAML Coding (pool)!



```
apiVersion: v1
kind: Service
metadata:
  name: my-tier-2-pool-a-svc
spec:
  ports:
    - name: nfs-mover
      port: 32049
      targetPort: 32049
    - name: xroot-mover
      port: 31094
      targetPort: 31094
    - name: http-mover
      port: 38080
      targetPort: 38080
```

```
[my-tier-2-pool-a-svc]

[my-tier-2-pool-a-svc/pool]
localaddresses=my-tier-2-pool-a-svc
pool.name=pool-a
pool.path=/pool
pool.mover.nfs.port.min=32049
pool.mover.nfs.port.max=32049
pool.mover.xrootd.port.min=31094
pool.mover.xrootd.port.max=31094
pool.mover.http.port.min=38080
pool.mover.http.port.max=38080
pool.mover.https.port.min=38083
pool.mover.https.port.max=38083
```

# Happy YAML Coding (pool, door)!



```
- name: wan-port-0
  port: 28000
  targetPort: 28000
- name: wan-port-1
  port: 28001
  targetPort: 28001
- name: wan-port-2
  port: 28002
  targetPort: 28002
- name: wan-port-3
  port: 28003
  targetPort: 28003
- name: wan-port-4
  port: 28004
  targetPort: 28004
```

TCP pots for gridftp can't be assigned dynamically, and require an advance mapping.

This must be done on ftp door and all pools to support various gridftp transfer modes.

# dCache Helm Chart










- Uses freshly built container (identified by commit hash or tag)
- Pre-configured, full functional setup
  - doors: dcap(g), frontend, ftp, http(s), nfs, srm, xrootd(g)
  - pools: a, b, c
  - core: PnfsManager, PoolManager,  
billing, cleaner-disk  
PinManager, SpaceManger,  
gPlazma, TransferService
- Not intended for proportional use (as of now)!





- Grid toolkit with almalinux-9
  - dccp, gridftp, srm, gfal-xxx,
  - 3<sup>rd</sup>-party copy
- SRM spec compatibility tests
  - test suite since srm-2.0 deployment
- xroot-gsi test
- Simple WebDAV with x509
- Basic Token-based authN
- REST-API
  - Bulk/Tape/QoS
  - Migration
- NFS protocol compatibility
  - No kernel client tests!

**testing**

- ✓ Frontend test suite 
- ✓ Grid EL9 WN tests 
- ✓ NFS4.x protocol compliance tests 
- ✓ Run OIDC test 
- ✓ SRM S2 test suite 
- ✓ gsi\_xroot\_tests 
- ✓ webdav\_with\_x509\_tests 

# Testing!



- Spawn WN
- Copy test script
- Run test
- Collect results

# Example: SRM test suite (ciao Flavia)



SRM S2 test suite:

stage: testing

extends: .kubernetes\_image

script:

- kubectl -n \$K8S\_NAMESPACE apply -f .ci/s2-runner.yaml
- kubectl -n \$K8S\_NAMESPACE cp .ci/init-el9-ui.sh s2-tester:/init-el9-ui.sh
- kubectl -n \$K8S\_NAMESPACE cp .ci/run-s2.sh s2-tester:/run-s2.sh
- kubectl -n \$K8S\_NAMESPACE exec s2-tester -- /bin/sh /run-s2.sh
- kubectl -n \$K8S\_NAMESPACE cp s2-tester:/TEST-basic.xml TEST-basic.xml
- kubectl -n \$K8S\_NAMESPACE cp s2-tester:/TEST-avail.xml TEST-avail.xml
- kubectl -n \$K8S\_NAMESPACE cp s2-tester:/TEST-usecase.xml TEST-usecase.xml

artifacts:

reports:

junit:

- "TEST\*.xml"

# Still Tested Manually



- Kernel NFS I/O
  - fio, mdtest, xfs-tests
- DB schema migration
- HA, Fail-over
- ...



# The Full 'Thing'



## [maven-release-plugin] prepare release 9.2.26

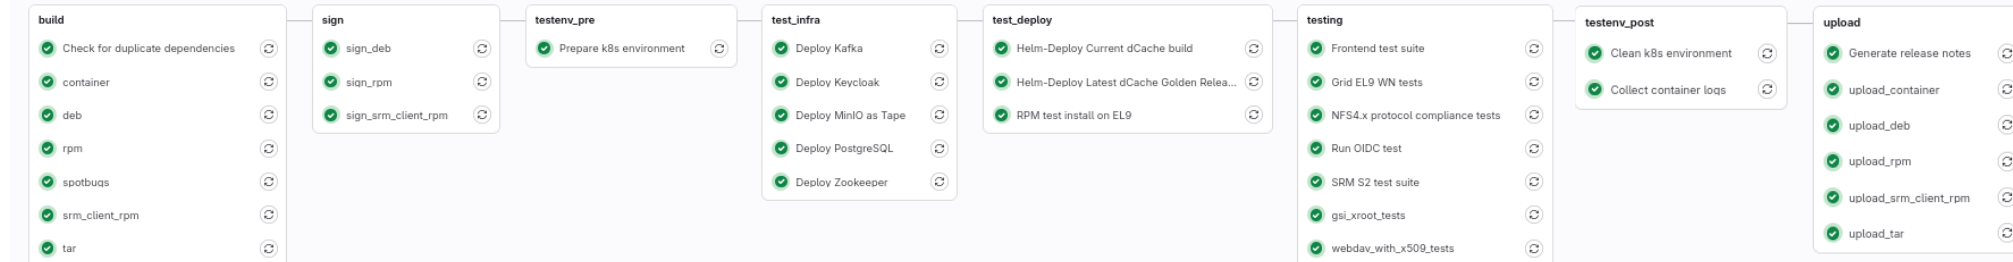
Passed guthub-sync created pipeline for commit 9e9ea4fa 6 days ago, finished 6 days ago

For 9.2.26

latest 26 jobs 52 minutes 22 seconds, queued for 7 seconds

Pipeline Needs Jobs 26 Tests 5579

Group jobs by Stage Job dependencies





# The Full 'Thing'



**[maven-release-plugin] prepare release 9.2.26**  
Passed guthub-sync created pipeline for commit 9e9ea4fa 6 days ago, finished 6 days ago  
For 9.2.26  
latest 26 jobs 52 minutes 22 seconds, queued for 7 seconds

Pipeline Needs Jobs 26 Tests 5579

Group jobs by Stage Job dependencies

Stage	Jobs
build	Check for duplicate dependencies, container, deb, rpm, spotbuas, srm_client_rpm, tar
sign	sign_deb, sign_rpm, sign_srm_client_rpm
testenv_pre	Prepare k8s environment
test_infra	Deploy Kafka, Deploy Keycloak, Deploy MinIO as Tape, Deploy PostgreSQL, Deploy Zookeeper
test_deploy	Helm-Deploy Current dCache build, Helm-Deploy Latest dCache Golden Relea..., RPM test install on EL9
testing	Frontend test suite, Grid EL9 WN tests, NFS4.x protocol compliance tests, Run OIDC test, SRM S2 test suite, gsi_xroot_tests, webdav_with_x509_tests
testenv_post	Clean k8s environment, Collect container logs
upload	Generate release notes, upload_container, upload_deb, upload_rpm, upload_srm_client_rpm, upload_tar

Can we build?

# The Full 'Thing'



**[maven-release-plugin] prepare release 9.2.26**  
Passed guthub-sync created pipeline for commit 9e9ea4fa 6 days ago, finished 6 days ago  
For 9.2.26  
latest 26 jobs 52 minutes 22 seconds, queued for 7 seconds

Pipeline Needs Jobs 26 Tests 5579

Group jobs by Stage Job dependencies

```
graph LR; build --> sign; sign --> testenv_pre; testenv_pre --> test_infra; test_infra --> test_deploy; test_deploy --> testing; testing --> testenv_post; testenv_post --> upload;
```

**build**

- Check for duplicate dependencies
- container
- deb
- rpm
- spotbuas
- srm\_client\_rpm
- tar

**sign**

- sign\_deb
- sign\_rpm
- sign\_srm\_client\_rpm

**testenv\_pre**

- Prepare k8s environment

**test\_infra**

- Deploy Kafka
- Deploy Keycloak
- Deploy MinIO as Tape
- Deploy PostgreSQL
- Deploy Zookeeper

**test\_deploy**

- Helm-Deploy Current dCache build
- Helm-Deploy Latest dCache Golden Relea...
- RPM test install on EL9

**testing**

- Frontend test suite
- Grid EL9 WN tests
- NFS4.x protocol compliance tests
- Run OIDC test
- SRM S2 test suite
- gsi\_xroot\_tests
- webdav\_with\_x509\_tests

**testenv\_post**

- Clean k8s environment
- Collect container logs

**upload**

- Generate release notes
- upload\_container
- upload\_deb
- upload\_rpm
- upload\_srm\_client\_rpm
- upload\_tar

Can we  
deploy?

# The Full 'Thing'



**[maven-release-plugin] prepare release 9.2.26**  
Passed guthub-sync created pipeline for commit 9e9ea4fa 6 days ago, finished 6 days ago  
For 9.2.26  
latest 26 jobs 52 minutes 22 seconds, queued for 7 seconds

Pipeline Needs Jobs 26 Tests 5579

Group jobs by Stage Job dependencies

```
graph LR; build --> sign; sign --> testenv_pre; testenv_pre --> test_infra; test_infra --> test_deploy; test_deploy --> testing; testing --> testenv_post; testenv_post --> upload;
```

Stage	Jobs
build	Check for duplicate dependencies, container, deb, rpm, spotbuas, srm_client_rpm, tar
sign	sign_deb, sign_rpm, sign_srm_client_rpm
testenv_pre	Prepare k8s environment
test_infra	Deploy Kafka, Deploy Keycloak, Deploy MinIO as Tape, Deploy PostgreSQL, Deploy Zookeeper
test_deploy	Helm-Deploy Current dCache build, Helm-Deploy Latest dCache Golden Relea..., RPM test install on EL9
testing	Frontend test suite, Grid EL9 WN tests, NFS4.x protocol compliance tests, Run OIDC test, SRM S2 test suite, gsi_xroot_tests, webdav_with_x509_tests
testenv_post	Clean k8s environment, Collect container logs
upload	Generate release notes, upload_container, upload_deb, upload_rpm, upload_srm_client_rpm, upload_tar

Did we brake it for users?

# Congratulations!



**You are  
a  
certified dCache  
CI-pipeline expert!**



# Full dCache Deployment in k8s



```
$ helm install chimera bitnami/postgresql
```

```
$ helm install cells bitnami/zookeeper
```

```
$ helm install billing bitnami/kafka
```

```
$ helm install --set image.tag=10.2.0 \  
my-store dcache/dcache
```





- Can't use ingress
  - dCache redirects clients to pools
  - **Productive sites use NodePort**
- Too-long hostnames
  - POSIX limits Full Qualified Host Names to 64 chars  
*old-store-pool-d-0.old-store-pool-d-svc.dcache-build-127946.svc.cluster.local*
- Yaml-hell
  - Charts quickly become too complex and fragile



# Work in Progress/Wish list



- GitHub pull-request validation
- Full dCache configuration with helm
- External helm charts as internal dependencies
- Cloud-native
- Mastering rolling updates

# Get Involved



- Use our container in your testing
- Help us to make helm charts production ready
- Help us with documentation
- Add your test scenario
- Share your experience and knowledge
- Share your needs



"Rosie the Riveter", National Museum of American History, Miller, J. Howard  
[https://americanhistory.si.edu/collections/search/object/nmah\\_538122](https://americanhistory.si.edu/collections/search/object/nmah_538122)



Thank you!



# Questions?



- dCache Helm Chart

<https://github.com/dCache/dcachelhelm>

- dCache CI-pipeline

<https://github.com/dCache/dcachelblob/master/.gitlab-ci.yml>

- dCache Containers

<https://hub.docker.com/r/dcache/dcacheltags>