

Space management in 2.10+

Gerd Behrmann




norden

NordForsk



Nordic e-Infrastructure
Collaboration



```
[piggy] (SrmSpaceManager) admin > ls link groups -h
FLAGS CNT RESVD AVAIL FREE UPDATED NAME
--r-:-o 7 1.4P + 23T = 1.4P 2015-05-17 10:45:19 atlas-disk-group
---c:n- 2 19T + 52T = 71T 2015-05-17 10:45:19 atlas-tape-write-group
--r-:-o 3 25T + 972G = 26T 2015-05-17 10:45:19 atlas-disk-dk-group
```

```
[piggy] (SrmSpaceManager) admin > ls spaces -h
TOKEN LINKGROUP RETENTION LATENCY ALLO USED FREE SIZE DESCRIPTION
95207 atlas-disk-group REPLICATION ONLINE 32G + 2.9P + 1.3P = 4.2P ATLASDATADISK
95216 atlas-tape-write-group CUSTODIAL NEARLINE 10G + 9.5T + 12T = 22T ATLASDATATAPE
95226 atlas-tape-write-group CUSTODIAL NEARLINE 0B + 15T + 6.7T = 22T ATLASMCTAPE
95227 atlas-disk-group REPLICATION ONLINE 0B + 0B + 3.0M = 3.0M ATLASMCDISK
56368439 atlas-disk-dk-group REPLICATION ONLINE 0B + 1.0M + 700G = 700G dcsc
47616811 atlas-disk-group REPLICATION ONLINE 0B + 0B + 29T = 29T ATLASPRODDISK
12101557 atlas-disk-group REPLICATION ONLINE 0B + 45T + 46T = 91T ATLASSCRATCHDISK
45694554 atlas-disk-dk-group REPLICATION ONLINE 0B + 16T + 6.4T = 22T ATLASDKGROUPDISK
19622609 atlas-disk-group REPLICATION ONLINE 0B + 7.8T + 3.2T = 11T ATLASNOLOCALGROUPDISK
30043703 atlas-disk-dk-group REPLICATION ONLINE 0B + 26T + 18T = 44T ATLASDKLOCALGROUPDISK
17040639 atlas-disk-group REPLICATION ONLINE 0B + 0B + 2.2T = 2.2T ATLASHOTDISK
2837931 atlas-disk-group REPLICATION ONLINE 0B + 6.3k + 26T = 26T ATLASGROUPDISK
```

Admin

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > reserve space -al=online -rp=replica -desc=test -lg=sm-group 100m
0 voGroup:null voRole:null retentionPolicy:REPLICA accessLatency:ONLINE linkGroupId:0 size:100000000 created:Sun May 17 10
:55:33 CEST 2015 lifetime:-1ms expiration:NEVER description:test state:RESERVED used:0 allocated:0
```

Client

```
Gerds-MacBook-Pro:dcache-git behrmann$ srm-reserve-space -access_latency=ONLINE -retention_policy=REPLICA
-space_desc=test2 -guaranteed_size=100000000 srm://localhost/
Space token =1
```

Implicitly

```
Gerds-MacBook-Pro:dcache-git behrmann$ srmcp file:///bin/bash srm://localhost/reserved/foo
```

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -a
```

TOKEN	S	LINKGROUP	RETENTION	LATENCY	ALLO	USED	FREE	SIZE	EXPIRES	DESCRIPTION
0	-	sm-group	REPLICA	ONLINE	0 +	0 +	100000000 =	100000000		test
1	-	sm-group	REPLICA	ONLINE	0 +	0 +	100000000 =	100000000	2018-02-10 09:57:54	test2
2	-	sm-group	REPLICA	ONLINE	0 +	628640 +	0 =	628640	2015-05-17 12:00:11	

Admin

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > release space 0
0 voGroup:null voRole:null retentionPolicy:REPLICA accessLatency:ONLINE linkGroupId:0 size:100000000 created:Sun
  May 17 10:55:33 CEST 2015 lifetime:-1ms expiration:NEVER description:test state:RELEASED used:0 allocated:0
```

Client

```
Gerds-MacBook-Pro:dcache-git behrmann$ srm-release-space -space_token=1 srm://localhost
```

Expiration

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -a
```

TOKEN	S	LINKGROUP	RETENTION	LATENCY	ALLO	USED	FREE	SIZE	EXPIRES	DESCRIPTION
0	-	sm-group	REPLICA	ONLINE	0 +	0 +	100000000 =	100000000		test
1	-	sm-group	REPLICA	ONLINE	0 +	0 +	100000000 =	100000000	2018-02-10 09:57:54	test2
2	-	sm-group	REPLICA	ONLINE	0 +	628640 +	0 =	628640	2015-05-17 12:00:11	

Authorization on create

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls link groups -l
FLAGS CNT RESVD      AVAIL      FREE      UPDATED NAME
-orc:no 1      0 + 1073113184 = 1073113184 2015-05-17 11:06:45 sm-group
    behrmann:*
    system-test:*
```

Authorization on release

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -l -e -h
TOKEN S LINKGROUP RETENTION LATENCY OWNER  ALLO  USED  FREE  SIZE      CREATED
    2 - sm-group  REPLICAS  ONLINE  behrmann  0B + 629k +  0B = 629k 2015-05-17 11:00:11
```

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > \sn storageinfoof /reserved
size=512;new=true;stored=false;sClass=test:reserved;cClass=-;hsm=osm;accessLatency=ONLINE;retentionPolicy=REPLICA;
StoreName=test;store=test;group=reserved;bfid=<Unknown>;
```

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > \s PoolManager psu ls linkGroup sm-group -l
sm-group : [ reserved-write-link ]
  AccessLatency:
    onlineAllowed=true
    nearlineAllowed=true
  RetentionPolicy:
    custodialAllowed=true
    outputAllowed=true
    replicaAllowed=true
```

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > \s PoolManager psu ls link reserved-write-link -l
reserved-write-link
  readPref : 10
  cachePref : 0
  writePref : 10
  p2pPref : 0
  section : None
  linkGroup : sm-group
  UGroups :
    any-protocol (links=7;units=1)
    world-net (links=7;units=2)
    reserved-store (links=2;units=1)
```

```
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > \s PoolManager psu ls ugroup reserved-store -l
reserved-store
  unitList :
    test:reserved@osm (type=Store;canonical=test:reserved@osm;uGroups=2)
```

Selecting the link group

- Client must be authorized to the link group
- Link group must allow access latency and retention policy
- Link group must have free space
- Link group must contain link able to accept the file

Link group authorization file

- Explicit reservation by client and implicit reservation on upload require authorization
- Unless you want to support one of the above, **you do not need a link group authorization file!**

```
Gerds-MacBook-Pro:target behrmann$ cat dcache/etc/linkgroup.conf
LinkGroup sm-group
behrmann
system-test
```

—

Client

```
Gerds-MacBook-Pro:target behrmann$ srmcp -space_token=3 file:///bin/bash srm://localhost/reserved/foo2
```

Admin

```
chimera:/# writetag /reserved WriteToken 3  
chimera:/# lstag /reserved  
Total: 3  
OSMTemplate  
WriteToken  
sGroup
```

Implicitly

```
spacemanager.enable.unreserved-uploads-to-linkgroups=true  
Gerds-MacBook-Pro:dcache-git behrmann$ srmcp file:///bin/bash srm://localhost/reserved/foo
```

Client

```
$ srmcp -access_latency=NEARLINE -retention_policy=CUSTODIAL file:///bin/bash srm://localhost/reserved/foo4
[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -e
TOKEN LINKGROUP RETENTION LATENCY ALLO      USED FREE      SIZE          EXPIRES DESCRIPTION
   4 sm-group  CUSTODIAL NEARLINE   0 + 628640 +   0 = 628640 2015-05-17 15:00:46
```

Admin

```
chimera:/# writetag /reserved AccessLatency NEARLINE
chimera:/# writetag /reserved RetentionPolicy CUSTODIAL

$ srmcp file:///bin/bash srm://localhost/reserved/foo5

[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -e
TOKEN LINKGROUP RETENTION LATENCY ALLO      USED FREE      SIZE          EXPIRES DESCRIPTION
   5 sm-group  CUSTODIAL NEARLINE   0 + 628640 +   0 = 628640 2015-05-17 15:04:03
```

Default

```
pnfsmanager.default-retention-policy=REPLICA
pnfsmanager.default-access-latency=ONLINE

$ srmcp file:///bin/bash srm://localhost/reserved/foo6

[Gerds-MacBook-Pro] (SpaceManager@dCacheDomain) admin > ls spaces -e
TOKEN LINKGROUP RETENTION LATENCY ALLO      USED FREE      SIZE          EXPIRES DESCRIPTION
   6 sm-group  REPLICA   ONLINE   0 + 628640 +   0 = 628640 2015-05-17 15:09:13
```

```
# ---- Whether space reservation should be supported.  
#  
# Whether to support space reservations in this dCache instance. This  
# controls whether services expect a spacemanager service to operate. Unless  
# you disable this feature, you must instantiate the spacemanager service by adding it  
# to the layout file on one of the dCache nodes (preferably in the same  
# domain as the poolmanager service).  
#  
(one-of?true|false)dcache.enable.space-reservation=true
```

Check list

- Check ownership of admin managed reservations.
You likely don't want an owner.
- Check whether you really need all link groups.
You likely want to remove those for VOs not using space reservations. Let the links reside outside a group and space manager will fall back to unreserved uploads.
- Check whether you need unreserved uploads to link groups.
You likely don't if you removed all link groups except those for which you have explicit reservations.
Consider creating an admin managed reservation and bind it to the directory tree using the WriteToken tag.
- Check whether you really need a link group authorization file.
If you don't want endusers to create reservations (you likely don't) and don't support unreserved uploads to link groups, then you don't need this authorization file.
- Why? Simpler configuration and better performance.

```
$ srmcp -access_latency-NEARLINE -retention_policy-CUSTODIAL file:///bin/bash srm://localhost/disk/fo
```

```
chimera:/# ls /upload
total 3
drwx--x--x 3  0    0 512 maj 17 14:42 .
dr-xr-xr-x 9  0    0 512 maj 17 11:00 ..
drwxr-xr-x 2 501 1000 512 maj 17 14:42 cc0f4a2b-3ad2-47c6-a888-b1255ef721df
chimera:/# ls /upload/cc0f4a2b-3ad2-47c6-a888-b1255ef721df
total 2
drwxr-xr-x 2 501 1000 512 maj 17 14:42 .
drwx--x--x 3  0    0 512 maj 17 14:42 ..
chimera:/# lstag /upload/cc0f4a2b-3ad2-47c6-a888-b1255ef721df
Total: 6
AccessLatency
ExpectedSize
OSMTemplate
Path
RetentionPolicy
sGroup
chimera:/# readtag /upload/cc0f4a2b-3ad2-47c6-a888-b1255ef721df AccessLatency
NEARLINE
chimera:/# readtag /upload/cc0f4a2b-3ad2-47c6-a888-b1255ef721df RetentionPolicy
CUSTODIAL
chimera:/# readtag /upload/cc0f4a2b-3ad2-47c6-a888-b1255ef721df Path
/disk/foo
```