

SRM Space Manager Configuration

Space Management functions

SRM Version 2.2 introduces a concept of space reservation. Space reservation is a promise by the storage system to make certain amount of storage space of certain type available for usage for a specified period of time. Space reservation is made using `srmReserveSpace` function. In case of successful reservation, a unique name, called space token is assigned to the reservation. Space token can be used during the transfer operations to tell the system to put the files being manipulated or transferred into an associated space reservation. A storage system ensures that the reserved amount of the disk space is indeed available, thus providing a guarantee that a client does not run out of space until all space promised by the reservation has been used. When files are deleted, the space is returned to the space reservation.

A space reservation has a property called retention policy. Possible values of retention policy are `Replica`, `Output` and `Custodial`. The retention policy describes the quality of the storage service that will be provide for the data (files) stored in this space reservation. `Replica` corresponds to the lowest quality of the service, usually associated with storing a single copy of each file on the disk. `Custodial` is the highest quality service, usually interpreted as storage of the data on Tape. Once a file is written into a given space reservation, it inherits the reservation's retention policy.

Another property of the space reservation is called access latency. The two values allowed are `Nearline` and `Online`. `Nearline` means that the data stored in this resevation are allowed to be stored in such a way that retrieving them might requere storage system to perform additional preparatory steps (staging data from tape to a disk cache for example). `Online` means that data is readily available and it will not take long to start reading the date. In case of `dCache Online` means that there will always be a copy of the file on disk, while `Nearline` does not provide such guarantee. As with retention policy, once a file is written into a given space reservation, it inherits the reservation's access latency.

`dCache` however only manages write space, i.e. only space on disk can be reserved and only for write operations. Once files are migrated to tape, they are no longer counted as part of a space, so the space used by these files is returned back into space reservation. When files are read back from tape and cached on disk, they are not counted as part of any space.

SRM Space reservation can be assigned a non-unique description, then the description cab be used in the future to discover all space reservation with a given description.

Properties of the SRM Space Reservation can be discovered using

SrmGetSpaceMetadata, released with srmReleaseSpace. For a complete description of the available functions please see SRM Version 2.2 specification at <http://sdm.lbl.gov/srm-wg/doc/SRM.v2.2.html>

Cache specific concepts

Link Groups

dCache 1.8 PoolManager supports new type of objects called LinkGroups. Each link group corresponds to a number of dCache pools in the following way: LinkGroup is a collection of the Links, each of which is a collection of the PoolGroups associated (Linked, hence a name "Link") with a set of the Pool Selection Units or PSUs. Each link group knows about its available size, which is a sum of all available sizes in all the pools included in this link group. In addition link group has 5 boolean properties called replicaAllowed, outputAllowed, custodialAllowed, onlineAllowed and nearlineAllowed, the values of these properties (true or false) can be configured in PoolManager.conf.

Space Reservations

In dCache 1.8 each SRM Space Reservation is made against the total available disk space of a link group. The total space in dCache that can be reserved is the sum of the available sizes of all Link Groups. If dCache is configured correctly each byte of disk space, that can be reserved, belongs to one and only one Link Group. Therefore it is important to make sure that no pool belongs to more than one pool group, no Pool Group belongs to more than one Link and no Link belongs to more than one LinkGroup.

Files written into a space made within a particular link group will end up on one of the pools referred to by this link group. The difference between the Link Group's available space and the sum of all the current space reservation sizes is the available space in the link group.

Space Manager access control

When SRM Space Reservation request is executed, its parameters, such as reservation size, lifetime, access latency and retention policy as well as user's Virtual Organization (VO) membership information is forwarded to the SRM SpaceManager.

Space Manager uses a special file for listing all the Virtual Organizations (VOs) and all the VO Roles that are permitted to make reservations in the given link group. List of the allowed VOs and VO Roles, together with the total available space and replicaAllowed, outputAllowed, custodialAllowed, onlineAllowed and nearlineAllowed properties of the group is then matched against the information from the user request in order to determine

if a given space reservation can be made in particular link group. Once a space reservation is created, no access control is performed, any user can attempt to store the files in this space reservation, provided he or she knows the exact space token.

Utilization of the space reservation for file storage

SRM Version 2.2 `srmPrepareToPut` and `srmCopy` pull mode transfers allow the user to specify a space reservation token or a retention policy and an access latency. In the protocol, any of these values are optional, and it is up to the implementation to decide what to do, if these properties are not specified. The specification does however require that if a space reservation is given, then any access latency or retention policy specified must match the same properties of the space reservation.

In dCache, if a space reservation is specified, the file will be stored in it (assuming the user has permission to do so in the name space).

If the reservation token is not specified, and implicit space reservation is enabled, then a space reservation will be performed implicitly for each `srm v1.1` and `srm 2.2` `srmPrepareToPut` or `srmCopy` in pull mode. If an Access Latency and a Retention Policy are specified, the user defined retention policy and default access latency. If the user has not specified Access Latency or Retention Policy (or if `srm v1.1` is used), the system will attempt to extract special tags (not surprisingly called “AccessLatency” and “RetentionPolicy”) from PNFS namespace from the directory to which file is being written. If the tags are present, then their values will determine the default Access Latency or Retention Policy that will be used for implicit space reservations. If the tags are not present, then system wide defaults will be used. If no implicit space reservation can be made, the transfer will fail. (Note: some clients also have default values, which are used when not explicitly specified by the user. In this case server side defaults will have no effect.)

If the implicit space reservation is not enabled in dCache 1.8 the pools in the link groups will be excluded from consideration and only the remaining pools will be considered to be the candidates for storing the incoming data, and classical pool selection mechanism will be used. If the space reservation is not used and no LinkGroups are specified, the system behavior will be exactly the same as in dCache 1.7.

SRM Space Manager and LinkGroups

Space Manager is making reservations against space in LinkGroups, LinkGroup is an object created by the PoolManager, that consists of several "Links". The total space

available in the given LinkGroup is a sum of available spaces in all links. An available space in each link is a sum of the available spaces in all pools assigned to a given link. Therefore for the space reservation to work correctly it is essential that each pool belongs to one and only one link, and each link belongs to only one LinkGroup. LinkGroups are assigned several parameters that determine what kind of space the LinkGroup corresponds to and who can make reservation against this space.

Definition of the LinkGroups in the PoolManager.conf

To configure PoolManager to create the new LinkGroup (a new reservable entity in dCache), please use following example (given in the PoolManager). Here we assume that write-link link already exists:

```
psu create linkGroup write-link-group
psu addto linkGroup write-link-group write-link
```

To tell Space Manager if the LinkGroup will be able to store files with given AccessLatency and RetentionPolicy, LinkGroups have 5 attributes: custodialAllowed, outputAllowed, replicaAllowed, onlineAllowed and nearlineAllowed. These attributes can be specified with the following commands:

```
psu set linkGroup custodialAllowed <group name> <true|false>
psu set linkGroup outputAllowed <group name> <true|false>
psu set linkGroup replicaAllowed <group name> <true|false>
psu set linkGroup onlineAllowed <group name> <true|false>
psu set linkGroup nearlineAllowed <group name> <true|false>
```

Please note that these attributes are all set to false by default, so the Srm SpaceManager will not be able to use LinkGroups unless they are defined.

Activating SRM Space Manager

In order to enable the new space reservation: add (uncomment) the following definition in dCacheSetup

```
srmSpaceManagerEnabled=yes
```

dCacheSetup SRM Space Manager Parameters

SpaceManagerDefaultRetentionPolicy: if space reservation request does not specify retention policy we will assign this retention policy by default.

Usage example:

```
SpaceManagerDefaultRetentionPolicy=CUSTODIAL
```

SpaceManagerDefaultAccessLatency: if space reservation request does not specify

access latency we will assign this access latency by default.

Usage example:

```
SpaceManagerDefaultAccessLatency=NEARLINE
```

SpaceManagerReserveSpaceForNonSRMTransfers: if the transfer request come from the door, and there was not prior space reservation made for this file, should we try to reserve space before satisfying the request.

Usage example:

```
SpaceManagerReserveSpaceForNonSRMTransfers=false
```

SpaceManagerLinkGroupAuthiorizationFileName contains the list of FQANs that are allowed to make space reservations in a given link group.

Usage example:

```
SpaceManagerLinkGroupAuthiorizationFileName="/opt/d-cache/etc/LinkGroupAuthiorization.conf"
```

SRM Space Manager Virtual Organization based access control configuration

VO based Authorization

Prerequisites

In order to be able to take advantage of the Virtual Organization (VO) infrastructure and VO based authorization and VO Based Access Control to the Space in dCache, certain things need to be in place:

- User needs to be registered with the VO
- User needs to use voms-proxy-init to create a vo proxy
- dCache needs to use gPlazma and not gPlazma with dcache.kpwd plugin, but other modules that know how to extract VO attributes from the proxy.(see [gPlazma Documentation](#), have a look at gplazmalite-vorole-mapping Plugin).

Only if these 3 conditions are satisfied the VO based authorization of the Space Manager can work.

If a client uses a regular grid proxy, created with grid-proxy-init, and not a Virtual Organization (VO) proxy, which is created with the voms-proxy-init, when he is communicating with SRM server in dCache, then the VO attributes can not be extracted its credential. voms-proxy-init adds a Fully Qualified Attribute Name (FQAN) section(s) to the grid proxy, which contain informaton about user's VO membership, in particular it contain VO Group name and VO Role that the client intends to play at this time. In this

case the name of the user is extracted on basis of the direct Distinguished Name (DN) to use name mapping. For the purposes of the space reservation the name of the user is used as its VO Group name, and the VO Role is left empty.

Access Control configuration

dCache Space Reservation Functionality Access Control is currently performed at the level of the LinkGroups. The access to making reservations in each LinkGroup is controlled by the **SpaceManagerLinkGroupAuthiorizationFile**.

SpaceManagerLinkGroupAuthiorizationFile has following syntax:

LinkGroup Name followed by the list of the Fully Qualified Attribute Names (FQANs), each FQAN on separate line, followed by an empty line, which is used as a record separator, or by the end of file. FQAN is usually a string of the form <VO>/Role=<VORole>. Both <VO> and <VORole> could be set to “*”, in this case all VOs or VO Roles will be allowed to make reservations in this LinkGroup.

Any line that starts with # is a comment and may appear anywhere.

File location is specified by defining

```
SpaceManagerLinkGroupAuthiorizationFileName=<FILENAME>
```

in the dCacheSetup

Example of the **SpaceManagerLinkGroupAuthiorizationFile**:

```
# this is comment and is ignored
```

```
LinkGroup LFSOnly-LinkGroup  
/atlas/Role=/atlas/role1
```

```
LinkGroup CMS-LinkGroup  
/cms/Role=*  
#/dteam/Role=/tester
```

```
LinkGroup default-LinkGroup  
#allow anyone :-)  
*/Role=*  
#/dteam/Role=/tester
```