

Storage Classes

Problematic and Implementation at CCIN2P3

Lionel Schwarz, Jonathan Schaeffer
dcachemaster@in2p3.fr

October 11, 2006



Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

- Improve usage quality of storage resource.
- Having a basic modelisation of the experiments use cases modeling three storage types.
- Speaking a same language with experiments improving comprehension between both sides

Integrated in the SRM 2.2 specification

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

Use Case Tape-Disk

- T1D0 : A file is written to tape and has not to be kept on disk.
- T0D1 : A file is written only on disk.
- T1D1 : A file is kept on disk so it is quickly available, but it has a backup on tape.
It is understood that this is a “temporary” state for a file.

Storage class of a file may vary in the file lifecycle

- Changing a storage class in a directory will result in a mess within the namespace
- storage administrators want to keep it as simple as possible :
Try to keep one unique Storage Class for a given namespace

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

- Integration in SRM v2.2 specification
- Implemented by all SRM servers :
 - Castor
 - dCache
 - dpm

Park description

Purpose

Use Cases

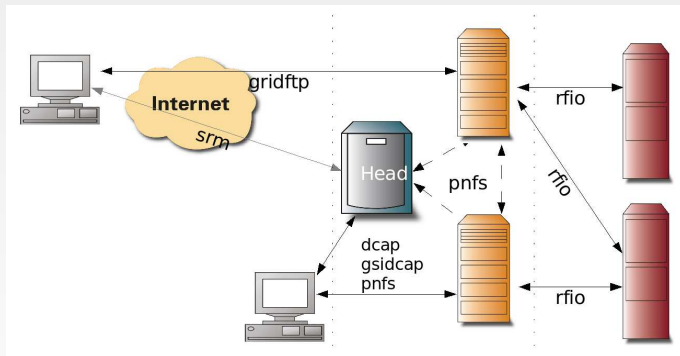
Implementation

dCache approach

dcache@CCIN2P3

Publication

References



- dCache developers do not want to remove features for implementing Storage Classes
- Storage characteristics known as **Storage Groups** of a file is determined by its namespace
- Storage Groups are defined in the PNFS path as tags (hidden files) and are inherited from the parent directory

The physical location of a file is determined by its namespace

- The namespace of a file (and though its Storage Group) can change without affecting its physical location

therefore ...

it is possible to have a reasonable StorageClass support using the existing features of dCache

Purpose

Use Cases

Implementation

dCache approach

dCache@CCIN2P3

Publication

References

- Files are physically stored regarding their PNFS namespace
- HPSS interaction :
 - **COS** (Class Of Service) is not linked with namespace but is associated with :
 - resources (disk type, tape type)
 - policies (filesize, migration, purge ...)
 - **Families** define groups of files which will be stored on a same tape
- **COS** are defined in the PNFS namespace through tags (the same way as for StorageGroups) and used by a migrator script
- **Families** are computed by mapping between HPSS namespace and dCache Storage Class

Purpose

Use Cases

Implementation

dCache approach

dCache@CCIN2P3

Publication

References

Storage Classes Levels

Purpose

Use Cases

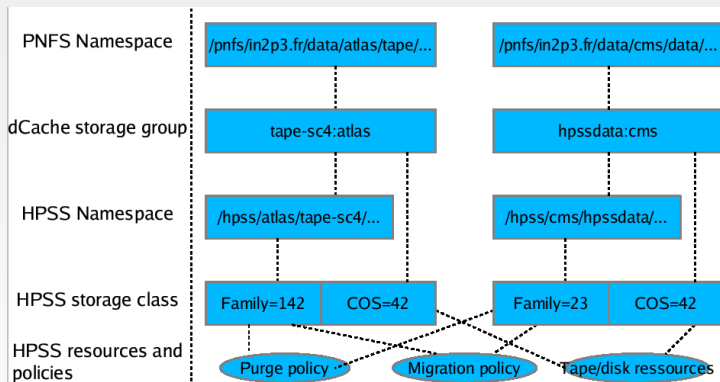
Implementation

dCache approach

dCache@CCIN2P3

Publication

References



Purpose

Use Cases

Implementation

dCache approach

dCache@CCIN2P3

Publication

References

- T1D0 is the default in dCache
- T0D1 corresponds to dCache pools without HPSS backend
- T1D1 still causes problems
 - The guess is to **pin** a file to cache
 - no automatic way to do this
- Transition
 - T1D1 → T1D0 : Set pin duration to 0
 - T0D1 → T1D1 : Should never happen
 - T1D0 → T1D1 : Staging a file

Coping with the Glue description

Actual publication

```
dn: GlueSALocalID=dteam,GlueSEUniqueID=ccsrn.in2p3.fr,  
mds-vo-name=IN2P3-CC,o=grid  
GlueSAAccessControlBaseRule: dteam  
GlueSAPath: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSAPolicyFileLifeTime: Durable  
GlueSAPolicyMaxData: 38553206  
GlueSAPolicyMaxFileSize: 38553206  
GlueSAPolicyMaxPinDuration: 86400  
GlueSAPolicyMinFileSize: 0  
GlueSAPolicyQuota: 53687091  
GlueSARoot: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSASStateAvailableSpace: 38553206  
GlueSASStateUsedSpace: 15133885  
GlueSAType: Durable
```

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

Coping with the Glue description

Tape1 Disk0

```
dn: GlueSALocalID=dteam,GlueSEUniqueID=ccsrn.in2p3.fr,  
mds-vo-name=IN2P3-CC,o=grid  
GlueSAAccessControlBaseRule: dteam  
GlueSAPath: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSAPolicyFileLifeTime: Permanent  
GlueSAPolicyMaxData: 53687091  
GlueSAPolicyMaxFileSize: 53687091  
GlueSAPolicyMaxPinDuration: 0  
GlueSAPolicyMinFileSize: 1048576  
GlueSAPolicyQuota: 53687091  
GlueSARoot: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSASStateAvailableSpace: 53687091  
GlueSASStateUsedSpace: 0  
GlueSAType: Permanent
```

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

Coping with the Glue description

Tape0 Disk1

```
dn: GlueSALocalID=dteam,GlueSEUniqueID=ccsrm.in2p3.fr,  
mds-vo-name=IN2P3-CC,o=grid  
GlueSAAccessControlBaseRule: dteam  
GlueSAPath: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSAPolicyFileLifeTime: Durable  
GlueSAPolicyMaxData: 38553206  
GlueSAPolicyMaxFileSize: 38553206  
GlueSAPolicyMaxPinDuration: ∞  
GlueSAPolicyMinFileSize: 0  
GlueSAPolicyQuota: 53687091  
GlueSARoot: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSASStateAvailableSpace: 38553206  
GlueSASStateUsedSpace: 15133885  
GlueSAType: Durable
```

Coping with the Glue description

Tape1 Disk1

```
dn: GlueSALocalID=dteam,GlueSEUniqueID=ccsrn.in2p3.fr,  
mds-vo-name=IN2P3-CC,o=grid  
GlueSAAccessControlBaseRule: dteam  
GlueSAPath: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSAPolicyFileLifeTime: Permanent  
GlueSAPolicyMaxData: 38553206  
GlueSAPolicyMaxFileSize: 38553206  
GlueSAPolicyMaxPinDuration: 86400  
GlueSAPolicyMinFileSize: 1048576  
GlueSAPolicyQuota: 53687091  
GlueSARoot: dteam:/pnfs/in2p3.fr/data/dteam  
GlueSASTateAvailableSpace: 38553206  
GlueSASTateUsedSpace: 15133885  
GlueSAType: Permanent
```

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

- Pros

- Simple rules to modelize the park
- One Storage Area per Storage Class



One pool per Storage Class

- Cons

- Will drive the metrics people crazy
- Should we publish one entry per namespace ?

- Agenda of the preGDB meeting :
<http://agenda.cern.ch/fullAgenda.php?ida=a058490>
- SRMv2.2 specification :
<http://sdm.lbl.gov/srm-wg/doc/SRM.v2.2.html>
- Glue 1.3 discussion :
<https://twiki.cern.ch/twiki/bin/view/EGEE/GlueSchema13>

Thank you for your attention

Storage Classes

Lionel Schwarz,
Jonathan Schaeffer
dcachemaster@in2p3.fr

Purpose

Use Cases

Implementation

dCache approach
dCache@CCIN2P3

Publication

References

Questions ?