

Mediaflux[®] + DMF

Zero Data Loss

What About the i-nodes?

- DMF makes multiple copies on tape
- XFS dump used to backup file system i-nodes
 - Typically done once day
 - Sensitive to i-node size (bigger makes slower)
- Issues:
 - If the file-system crashes between dumps, then orphaned files on tape
 - Backup becomes slower as more files in the file system.

Mediaflux + DMF

- Mediaflux now stores DMF BFIDs and keeps track of whether files have been fully committed to tape – optimized storage for large numbers of BFIDs
- Used for two purposes:
 - Reverse mapping of BFID to file and asset path
 - On-the-fly i-node reconstruction
- Add in real-time replication of Mediaflux database and you have zero data loss once a transaction is committed

Examples

```
> asset.create :in file:/Users/jason/stuff/Content/images/brain.jpg :store dmf :namespace medical/images :name brain1.jpg
: id "393"
> asset.get :id max
:asset -id "393" -version "1" -vid "6296"
: type "image/jpeg"
: namespace "/medical/images"
: path "/medical/images/brain1.jpg"
: name "brain1.jpg"
: creator -id "3"
: domain "system"
: user "manager"
: ctime -millisec "1424208620047" "17-Feb-2015 16:30:20"
: mtime -millisec "1424208620047" "17-Feb-2015 16:30:20"
: stime "6296"
: versioned -count "1" "true"

...

:content -id "1" -stime "6296" -versions "1" -total-size "37842"
: ctime -millisec "1424208620047" "17-Feb-2015 16:30:20"
: atime -millisec "1424208620067" "17-Feb-2015 16:30:20"
: type -ext "jpg" "image/jpeg"
: size -h "37.84 KB" "37842"
: csum -base "16" "E25118EC"
: csum -base "10" "3796965612"
: copy-id -stime "6296" "1"
: copy-ctime -millisec "1424208620047" "17-Feb-2015 16:30:20"
: store -oid "20" "dmf"
: url "file:/mnt/cache/stores/dmf/data/0/0/0/0/0/20"
: extended
: dmf
: bfid "53146fe200000000000000048"
: committed-to-tape "false"
```



DMF attributes


Examples – Lookup by BFID

```
> asset.store.dmf.bfid.to.asset :bfid 53146fe200000000000000048
:bfid-to-asset -bfid "53146fe200000000000000048"
:asset -id "393" -version "1"
:path "/medical/images/brain1.jpg"
:store -oid "20" "dmf"
:url "file:/mnt/cache/stores/dmf/data/0/0/0/0/0/20"
```

Examples – DMF Q Status

```
> asset.store.dmf.queue.describe :offset 1
:size "1"
:entry -bfid "53146fe200000000000000048"
  :asset -id "393" -version "1"
    :path "/medical/images/brain1.jpg"
    :store -oid "20" "dmf"
    :url "file:/mnt/cache/stores/dmf/data/0/0/0/0/0/20"
  :request -id "1395" -type "backup"
    :start-time -tz "US/Eastern" -gmt-offset "-5.0" -dst "false" -time "1424190620000" "17-Feb-2015 11:30:20"
    :priority "512"
    :status "active"
    :copy
      :offset "0"
      :length "37842"
      :library-server "lib1"
      :volume-group "libd0v2"
      :chunk
        :chunk-offset "0"
        :file-offset "0"
        :drive-group "libd0"
        :status "wait_zone"
    :copy
      :offset "0"
      :length "37842"
      :library-server "lib1"
      :volume-group "libd0v1"
      :chunk
        :chunk-offset "0"
        :file-offset "0"
        :drive-group "libd0"
        :status "wait_zone"
```

Reverse Mapping



Roadmap

- BFID recording and lookup available now
- On-the-fly i-node reconstruction available towards the end of March, early April 2015
- Real-time Mediaflux database replication available end of March 2015