dCache Installation and Security Tips

OSG Site Administrators workshop
Indianapolis
August 6-7 2009

Fermilab

Tanya Levshina
tlevshin@fnal.gov
Fermilab

Talk Outlines

- Part I:
  - dCache Status
  - Navigating through dCache releases
- Part II:
  - Changes vdt-dCache installation
  - Community toolkit
- Part III
  - dCache security

dCache Status - production release

- Reliable Deletion Registration
  - Eliminates leak of pool space, false “File Not Found” errors
  - Reliable Deletion Registration
- XACML gPlazma plugin
- Logging Context Propagation
  - dCache moved to Log4j for logging
  - New log4j management commands
- Info Service

dCache status – new features

- New Pool Code, Pool Migration Module
  - Improved Pool Space Accounting, reduce memory usage
  - Support for Berkley DB Pool Repository
  - Migration Module for copying files between pools
- Chimera
  - Better performance then pnfs, in production
  - Prerequisite for NFS4.1 and some ACL features
- Access Control Lists
  - Subset of the NFS version 4 ACLs
  - Used in conjunction with POSIX file permission mask
  - ACLs are stored in the database on PnfsManager
  - Work with Chimera (Full support) and PNFS (Only Subset of functions), SRM ACL is not yet implemented

dCache release policy

- Two dimensional release system:
  - A feature release is produced once in 2 months
  - Feature releases have version numbers 1.9.x-1, where x is a positive integer.
  - Each feature release is followed up by a number of maintenance releases containing bug fixes
- There is ONLY one recommended production release at any given time

Current releases

- Versions lower then 1.9.2 are no longer supported
- 1.9.2-x
  - gPlazma XACML authorization mechanism
  - storage of the authorization in database by SRM
  - PnfsManager has been updated to work with the latest version of Chimera
- 1.9.3-x
  - ACL
  - NFS4.1 support but no security yet, not connected to gPlazma
  - asynchronous srmLs support
  - refactoring of the pool code, first steps to refactor the pool manager
- 1.9.4-x
  - AC for staging from tape
  - scalable xrootd redirector
  - performance improvements in Chimera.
- 1.9.5 TBD – Golden release, for the duration of first round of LHC run

Changes in vdt-dcache installation

- New configuration file:

```
# ++++++++++++++++++++++++++++++++++++++++++++
# SECTION 1 - Some General Stuff
# You must fill in all values in this section
# ++++++++++++++++++++++++++++++++++++++++++++
# File System Domain
# Example: fnal.gov
MY_DOMAIN="fnal.gov"
# Would you like java to be installed by the dcache install script?
# Options: yes or no
INSTALL_JDK="yes"
# Would you like to install/use Gratia dCache storage and transfer probes?
# Options: yes or no
INSTALL_DCACHE_GRATIA_PROBES="yes"
…
```

- Run configuration script, that will install
  - Postgres
  - Pnfs
  - dCache-server
  - Srm-watch
  - gratia storage and trasfer probes
  - Java
  - Jython
  - Community toolkit
- For upgrade: run convergence script to transfer the old format to a new one

OSG Community Toolkit

- Core dCache operation:
  - Collects disk usage of all pools
  - Verifies checksum of all files. For each file, it compares the checksum as computed by the individual pool, and the checksum as stored in PNFS.
  - Verifies the discrepancy in disk usage of all pools, by collecting information from individual pools. If there is a difference in total disk as compared to used disk and free disk, discrepancy is reported.
  - Provides all the current transfer rates in dCache
  - Cleans out broken transfers for a site which does not have a tape backend. It can be used to trigger resets of all stuck dCache files.
  - finds and cleansup files still resident on disks of pools, but which have been otherwise removed from PNFS
  - Finds files on a specific pool
  - Calculates number of replicas for all files
  - Finds path for a specific pnfsid
  - Finds pnfsid for a specific path
  - Finds pools with a specific file
- dCache Chronicle
  - Sends an email with overall system summary, based on all active PoolGroups in the storage system at the site.
  - For each PoolGroup found to have active pools, total is computed to measure overall disk capacity.
  - Total is also computed for space which is free, or precious, or cached and removable.
  - Report includes total number of active pools and disk utilization percentage.

dCache Security Architecture

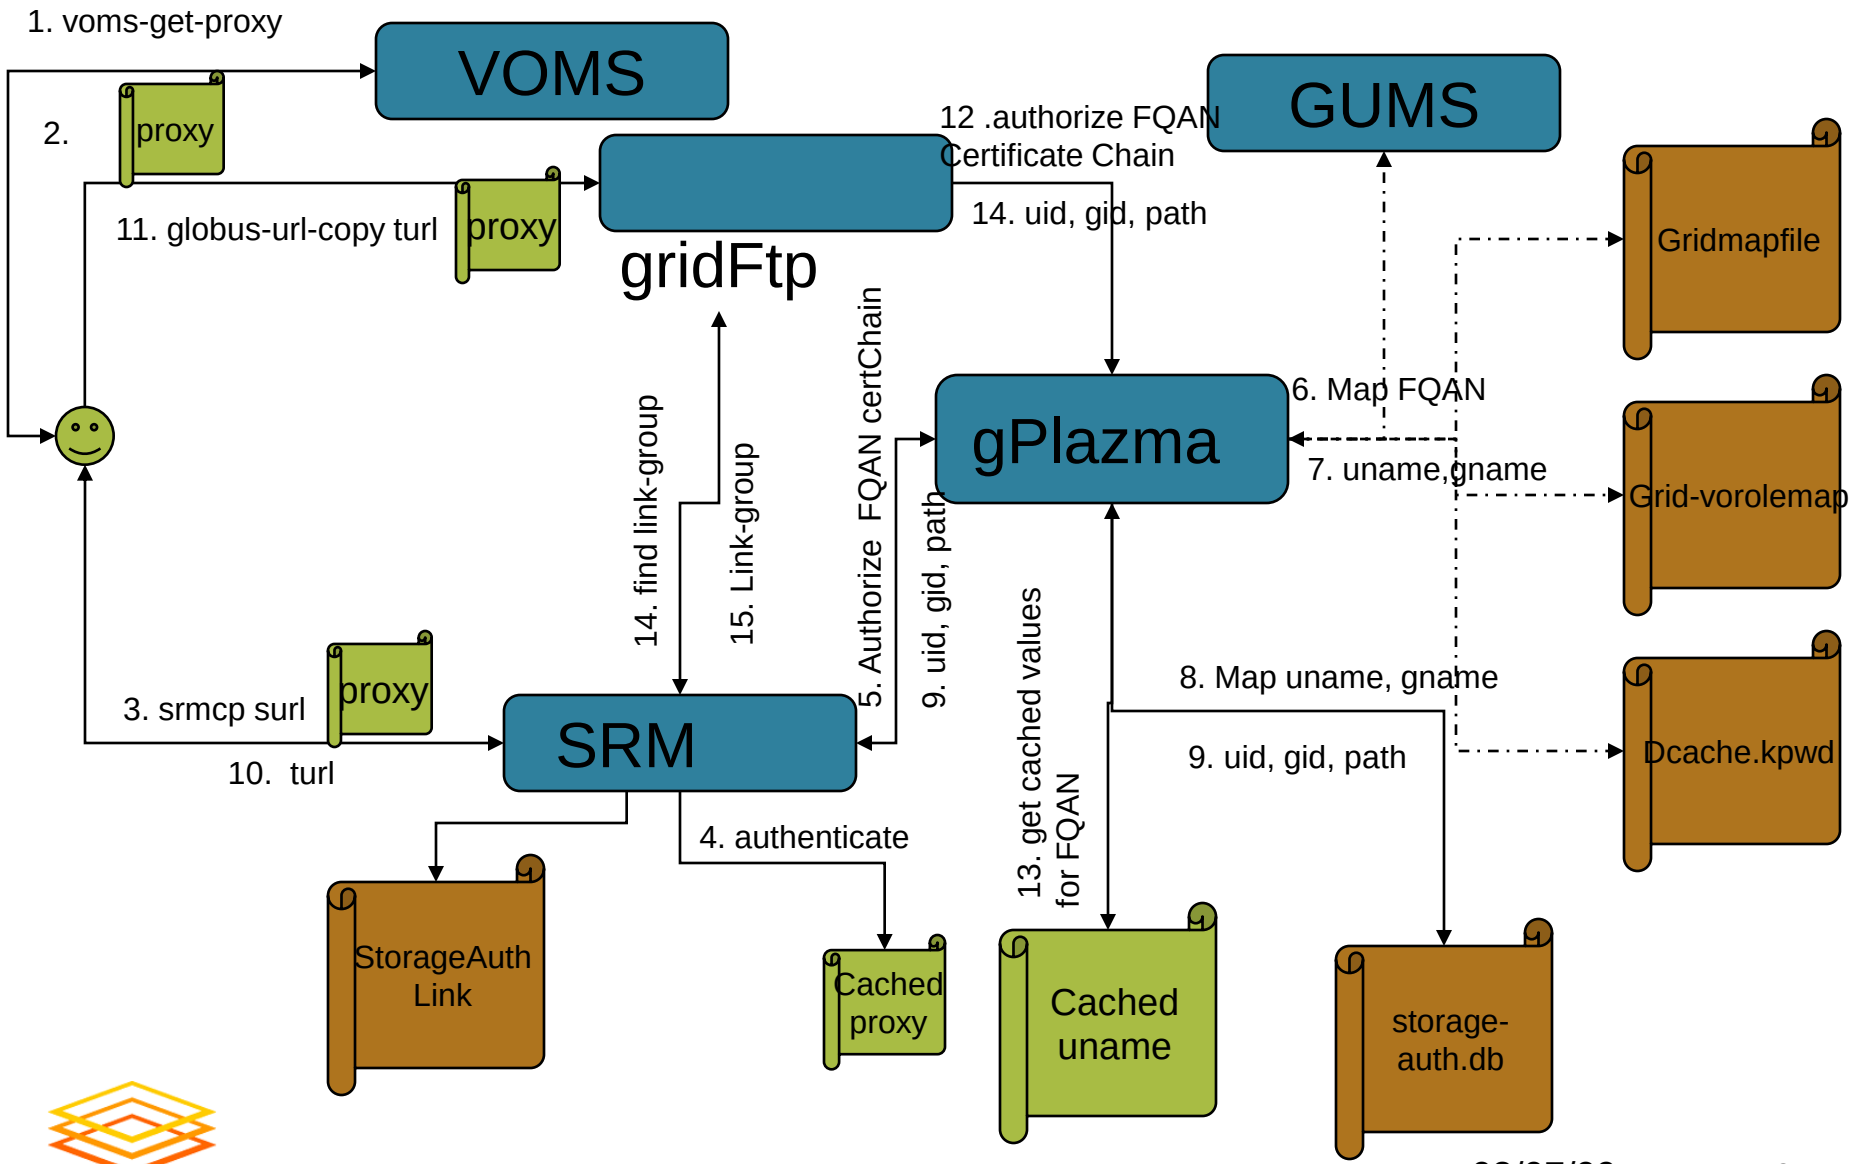

dCache Security - SRM

- SRM Web Services run on top of httpg. Provides authentication and delegation of X509 based credential.
- Stores one copy of the user's credential in memory
  - substitute stored credential with the new one with longer lifetime for all requests by the same user.
  - Provides a persistent storage of the user's credential in order to support the execution of the requests after the restart
- dCache Space Reservation Access Control is currently performed at the level of the LinkGroups.
  - The access to each LinkGroup is controlled by the setting in LinkGroupAuthorization.conf
  - It specifies the list of FQANs that are allowed to make space reservations in a given link group.

dCache Security – gPlazma (I)

- Receives user credential information and returns the authorization decision and site-specific user information such as uid, gid, and rootpath in return.
- Currently supports 5 plugins which implement various authorization methods:
  - kpwd : The dcache.kpwd file maps a user's DN to a local username, and the same file is used in a second mapping of the username to the uid, gid, and rootpath.
  - grid-mapfile : From the gridmapfile, the user's DN is mapped to a username. A second file, storage-authzdb, is used for the mapping of the username to the uid, gid, and rootpath.
  - gplazmalite-vorole-mapping : From gplazmalite-vorole-mapping file maps the user's FQAN The mapping of username to uid, gid, and rootpath is through the storage-authzdb file.
  - saml-vo-mapping : The DN and Role are mapped to a username via a callout to a GUMS server. The GUMS service may run an extension which returns the uid, gid, and rootpath as well. Otherwise, the mapping of username to uid, gid, and rootpath is through the storage-authzdb file.
  - xacml-vo-mapping : similar to saml-vo-mapping but allows gPlazma to acquire authorization mappings from any service which supports the obligation profile for grid interoperability. Servers presently supporting XACMLmapping are the latest releases of GUMS and SCAS.

dCache Security – gPlazma (II)

- storage-authzdb contains mapping of user name to uid, gid and root path
- The policy is defined in dcachesrm-gplazma.policy
  - On/off
  - Priority
  - timeouts
- gPlazma could be configured to perform VOMS attribute validation.
  - Verifies the signature on an attribute against the CA certificate that signed the voms server certificate.
  - Requires
    - "*.lsc" files in /etc/grid-security/vomsdir for each authorized voms server
    - the following attributre in dcachesrm-gplazma.policy set : vomsValidation="false"

dCache Security – GridFTP, GSIdCap

- gPlazma cell can be called from GridFTP and GSIdCap doors
- The configuration files must exist on the gPlazma node
- You can control whether the GridFTP door or SRM will authorize locally, or use the gPlazma for authorization. The default is to use the gPlazma cell for authorization:

  # useGPlazmaAuthorizationModule=false
  # useGPlazmaAuthorizationCell=true

- If both values are set to false, the GridFTP door or SRM will use the dcache.kpwd lookup method. A dcache.kpwd file must be present on the GridFTP door or SRM node in that case.
- It is possible to use gPlazma methods on the GridFTP door or SRM without calling the gPlazma cell by setting

  useGPlazmaAuthorizationModule=true

Useful links

- dCache Book
- dCache downloads
- Srm Design Document
- OSG community toolkit
- VDT dCache Home Page