

Clustering classes

- Currently available “full” clustering hint classes:
 - BdbEvsClusteringHint, BdbEvsHdrClusteringHint
 - one-process-per-container strategy (relieves locking problems)
 - BdbCondClusteringHint, BdbCondIndexClusteringHint (not used yet)
 - many-processes-per-container strategy (while one process updates, all others waits, the waiter issues a warning every minute)
 - BdbSpClusteringHint
 - handles registry, dictionary
- Dependency chain:
 - HepODBMS-BdbUtil-BdbApplication-BdbAccess-BdbClustering
- BdbSimpleClusteringHint class

Directory Structure / Db Naming

- Directory Structure
 - baseDir/domain/authLevel/user/comp/comp000000*
 - example: /u1/databases/events/users/becla/raw/raw000000*
- DB naming convension examples:
 - e_u_becla_raw000000*
 - e_s_bb_recHdr000000*
 - c_g_groupA_emcIndex*
 - “.bdb” appended to every database file name

MultiFS Support

- Ability to span multiple File Servers / File Systems
 - requires a daemon running along AMS
 - FS parameters:
 - host name, base directory name, port number
 - FSs can be grouped together
 - *parallel / sequential database distribution across a group*

FS configuration

- ASCII file with a data loaded into a federation (administrator task)
 - “*” supported (all domains or authLevels or users or components)
- see <http://www.slac.stanford.edu/BF/doc/Computing/Databases/Docs/AdminGuide.html>
- every domain/authLevel/userName/component can be assigned:
 - one group of file systems
 - max number of container per db
 - max container size
 - number of databases per directory
 - Defaults
 - FSGroup: catalogFS
 - the rest correspondingly : 10, 200 MB, 256 (configurable per domain)

Other Issues

- Reusing not-full containers
- Protection (currently turned off)
 - db file becomes read-only if db not used, authorization required to change the permission flag
- Problem of simultaneous creation of db with the same name at the same time by more than one process
 - current solution: we allow only one db per federation to be created at the same time
- A pitfall
 - clustering hint may restart a transaction (do not use whilst in persistent context)
 - will be fixed if we move to objy V5 (separate thread)