

Blackboard's Docstore Merge Tool

Introduction

In previous releases of Blackboard Content System, the five top-level directories, `/users`, `/courses`, `/orgs`, `/institution` and `/library` (with the two special second-level directories `/library/Library Content` and `/library/eReservers`), each had their own database schema. These schemas are known as Xythos Document Store (docstore) schemas. The CMS meta-schema is used to store access information for these docstores (in the `XY_FILE_SYSTEMS` table), so one knows in which docstore to look for files contained in either of these top-level directories (`XY_ROOT_DIRECTORIES`).

In release 9.1, this design is abandoned in favor of a single docstore holding all five traditional top-level directories. This new docstore (named `BBLEARN_CMS_DOC` in fresh 9.1 installs, and `CMS_DOC` in upgraded installs) also holds the new `/internal` top-level directory (TLD), which is not managed by Xythos rules but by the Blackboard application. However the five existing TLDs (`/users`, `/courses`, `/orgs`, `/institution`, `/library`) will only be there in fresh 9.1+ installs or after merging the docstores, but not in upgraded installs whose docstores were not (yet) merged.

A command line script is being provided to migrate the five traditional docstores one by one to the new docstore, creating a single merged docstore. Note that this is a major upgrade operation and needs to be executed with the Blackboard application stopped on all appservers.

Merging the docstores does not rename the other database schemas, so upgraded installs will still use a distinct naming system which is not compatible with fresh 9.1 installs. The docstore merge does not alleviate this problem, although it removes one major distinction between the two types of DB setups. In detail:

- upgraded (classic) installs contain these DB schemas: `BB_BB60`, `BBADMIN`, `BB_BB60_STATS`, `CMS`, `CMS_DOC`, `CMS_FILES_USERS`, `CMS_FILES_COURSES`, `CMS_FILES_ORGS`, `CMS_FILES_INST`, `CMS_FILES_LIBRARY`
- fresh 9.1 installs with the default BBUID of `BBLEARN` contain these DB schemas: `BBLEARN`, `BBLEARN_ADMIN`, `BBLEARN_STATS`, `BBLEARN_CMS`, `BBLEARN_CMS_DOC`
- fresh 9.1 installs with the custom BBUID of `BB9` (example) contain these DB schemas: `BB9`, `BB9_ADMIN`, `BB9_STATS`, `BB9_CMS`, `BB9_CMS_DOC`
- fresh 9.1 installs with the custom BBUID of `BB_BB60` would contain these DB schemas: `BB_BB60`, `BB_BB60_ADMIN`, `BB_BB60_STATS`, `BB_BB60_CMS`, `BB_BB60_CMS_DOC`. The application gets confused in this scenario, so don't do this!

Docstore Merging

There are a number of reasons to merge the docstores:

- Simplified database management. Once all docstores are merged into one, there are fewer DB schemas to manage and fewer places to look for information.
- Disk space savings. Certain tables such as `XYF_PRINCIPAL_LOOKUPS` no longer need to be replicated in each schema.
- Far fewer database connections will be needed (this is a big concern for clients with many appservers due to memory demands on the DB server). Every docstore needs its own pool of database connections. Determining the optimal connection pool sizing for each docstore can be tricky. With only one docstore, only this one pool of database connections is of concern.

However, the max connections should be set high for it, considerably higher than the shipping default of 25.

- Application code will be able to re-use DB connections for queries on different top-level-directories, whereas currently this requires multiple connections from multiple connection pools. It is possible that the traditional Content System performance issues with DB connections to multiple docstores that can lead to application-level deadlock can only truly be resolved by merging docstores.
- All QA testing from release 9.1 onward will assume that there is a single docstore. To remain in a "standard" configuration, clients should ideally use this expected setup too. That does not mean they are not supported if they did not merge yet, but rather that any issues they might face are likely to remain undetected by us.

The Docstore Merge Tool

It is not recommended that clients run the docstore merge tool until after they have upgraded to at least 9.1 SP5. There are significant issues with the tool in the prior versions. However the tool is fully functional beginning in 9.1 SP5.

The following is the command line for the tool:

```
blackboard/tools/admin/MergeDocstores.(bat | sh)
```

Be careful to only run one one instance of this script, running multiple copies at the same time will cause problems.