Infrastructure - Task #2182

Story # 2166 (Closed): Hazelcast cluster errors need to be isolated

Coordinate d1_indexer locks in the process cluster

2012-01-09 22:26 - Chris Jones

| Status: | Closed | Start date: | 2012-01-09 |
|------------------|--------------------------|-----------------|------------|
| Priority: | Normal | Due date: | |
| Assignee: | Skye Roseboom | % Done: | 100% |
| Category: | d1_indexer | Estimated time: | 0.00 hour |
| Target version: | Sprint-2012.01-Block.1.1 | | |
| Milestone: | CCI-1.0.0 | Story Points: | |
| Product Version: | * | | |
| Description | | | |

Instead of directly locking the pid of an object in the storage cluster to coordinate among processing components, instead create the lock in the Hazelcast process cluster. This should reduce contention on the hzSystemMetadata map calls, but will still allow synchronization, replication, and indexing to coordinate process timing.

History

#1 - 2012-01-09 23:11 - Skye Roseboom

Delay processing index task until the hz lock is released.

May also check replica entry status on the system metadata object - 'queued' or 'requested' means updates coming soon.

Index processing just performs reads from the distributed datastore using event listener on the system metadata map. Indexing does not create locks.

#2 - 2012-01-13 18:43 - Skye Roseboom

- Status changed from New to Closed

Closing...don't think index processing will be effected considering:

Only one update task allowed per PID - so rapid add/update, update/update from locking behavior only generates one task for the most recent event.

Task processor runs on 2 minute schedule - so query for tasks has small chance of occurring between quick add/update, update/update scenarios.