# Infrastructure - Story #2052

## Metacat Create locks SystemMetadata Map as well as Synchronization on the same PID

2011-11-23 18:29 - Robert Waltz

| | | | | |
|---|---|---|---|---|
| **Status:** | Closed | | **Start date:** | 2011-11-23 |
| **Priority:** | Normal | | **Due date:** | |
| **Assignee:** | Ben Leinfelder | | **% Done:** | 100% |
| **Category:** | Metacat | | **Estimated time:** | 0.00 hour |
| **Target version:** | Sprint-2012.01-Block.1.1 | | | |
| **Story Points:** | | | | |

**Description**

Synchronization will lock the SystemMetadata PID before create in Metacat. Metacat in create will lock as well.

Does put on the Hazelcast SystemMetadata Map block on a locked pid?

Are we certain all updates to the SystemMetadata Map will  not cause a deadlock?

## History

#### #1 - 2011-12-06 14:43 - Dave Vieglais

*- Position set to 1*

*- Target version changed from Sprint-2011.48-Block.6 to Sprint-2011.49-Block.6*

#### #2 - 2011-12-19 18:04 - Dave Vieglais

*- Target version changed from Sprint-2011.49-Block.6 to Sprint-2011.51-Block.6*

*- Position deleted (2)*

*- Position set to 1*

#### #3 - 2012-01-04 18:34 - Chris Jones

*- Assignee changed from Chris Jones to Ben Leinfelder*

Assigning to Ben for expediency.

#### #4 - 2012-01-04 18:52 - Matthew Jones

*- Target version changed from Sprint-2011.51-Block.6 to Sprint-2012.01-Block.1.1*

*- Position deleted (14)*

*- Position set to 320*

#### #5 - 2012-01-04 20:50 - Ben Leinfelder

A few things:
1) I don't see a call to "lock" a PID on the Hazelcast SystemMetadata map in d1_synchronization.
2) I'm not sure why d1_synchronization needs to interact directly with the Hazelcast SystemMetadata map. There are methods to get(), create() and registerSystemMetadata() on the CN and it seems reasonable for this d1_synchronization process to use those API methods instead of going through the "backdoor" of the HZ map.

(Perhaps there is a reason for calling the shared map directly and I am missing the point?)

If all interactions with the SystemMetadata map were limited to the CN implementation, I think we would not be at risk for deadlock behavior.

**#6 - 2012-01-13 19:39 - Ben Leinfelder**

*- Status changed from New to Closed*


chris and robert say:
The synchronization process gets a system metadata lock from ILock in the processing cluster rather than the "storeage" cluster (that is shared with Metacat)


**#6 - 2012-01-13 19:39 - Ben Leinfelder**

*- Status changed from New to Closed*


chris and robert say:
The synchronization process gets a system metadata lock from ILock in the processing cluster rather than the "storeage" cluster (that is shared with Metacat)