

Infrastructure - Story #3729

Member Nodes should be the authoritative source of System Metadata

2013-04-26 15:38 - Chris Jones

Status:	Closed	Start date:	2013-05-07
Priority:	High	Due date:	
Assignee:	Dave Vieglais	% Done:	100%
Category:	Documentation	Estimated time:	0.00 hour
Target version:	CCI-2.0.0		
Story Points:			
Description			
<p>After gaining some experience with CN-based authority for system metadata, we've realized that there are many use cases that require system metadata to be managed by the MNs authoritatively, and by the CNs secondarily as a cached version. The main use case involves access control. When an ITK client creates an object through MN.create(), control of the system metadata is transferred to the CN once synchronization happens. After that point, the ITK client (and scientist) has to make CN.setAccessPolicy() calls to make any changes. If the MN is set to sync once a week, this is problematic.</p> <p>Ultimately, the CN stores and manages system metadata in order to track replicas of an object, and to perform auditing on those replicas. This change would really just require that the CN remains the authority of the ReplicaList, whereas the MN would become the authority of all other fields. By doing so, ITK clients will be able to interact with their MN without delay, and the CNs will work off of cached versions of the system metadata (in Hazelcast and persisted in Metacat).</p> <p>To effect these changes, specific sub tasks include:</p> <ol style="list-style-type: none">1) Design of the system - sequence diagram and use case changes2) Update architecture and guide docs to express that the MN is the authoritative source for system metadata3) Ensure MN API calls reflect that the MN is responsible for incrementing serial version<ol style="list-style-type: none">3.1) Serial version will be used to track all fields except the ReplicaList3.2) Consider adding an optional serialVersion attribute to a ReplicaList, and use it to track versions of the list on the CN4) Include a push notification of system metadata change (CN.systemMetadataChanged())5) Change the CN stack to accommodate the MN-based authority<ol style="list-style-type: none">5.1) Replication5.2) Synchronization5.3) DAO layer for replication6) Change all MN stacks to implement the new features (Metacat, GMN, Mercury, EDAC, Dryad, Merritt)7) Add MN.updateSystemMetadata() interfaces (d1_common_{java python}, d1_libclient_{java python})8) Refactor CN sysmeta methods to delegate to MN.			
Subtasks:			
Task # 3743: Change system metadata design documents to reflect MN authority			Closed
Task # 3744: Add MN API calls to update system metadata			Closed
Task # 3745: Add Replica version attribute			Rejected
Task # 3746: Add CNCORE.systemMetadataChanged() call			Rejected
Task # 7175: update operations documentation			Closed
Task # 3787: CN should detect changed sysmeta and resync objects when needed			Closed
Task # 3788: Add CN.updateSystemMetadata method			Closed
Task # 7179: add API documentation for CN.updateSystemMetadata			Closed
Task # 7180: add d1_libclient_java implementation			Closed
Task # 7225: Add cn.updateSystemMetadata implementation in Metacat			Closed

History

#1 - 2013-04-26 15:53 - Chris Jones

- Description updated

#2 - 2014-10-01 16:45 - Robert Waltz

- Target version set to CCI-2.0.0

- Start date set to 2014-12-01

- Due date set to 2014-12-01

#3 - 2016-01-05 18:50 - Chris Jones

- Assignee changed from Chris Jones to Dave Vieglais

Outstanding tasks in this story can be moved to other stories since the V2.0.0 release is complete.

#4 - 2016-03-25 01:14 - Dave Vieglais

- Status changed from New to Closed

- % Done changed from 0 to 100