

Infrastructure - Decision #6392

CN getChecksum REST API design review

2014-09-18 16:46 - Skye Roseboom

Status:	New	Start date:	2014-09-18
Priority:	Normal	Due date:	
Assignee:	Dave Vieglais	% Done:	0%
Category:	Requirement	Estimated time:	0.00 hour
Target version:			
Milestone:	None		
Description			
There is a difference in REST API design between MN and CN getChecksum service.			
MN.getChecksum requires the MN to calculate or produce a checksum that is 'known to be correct' - rather than returning the checksum recorded in system metadata for the associated document. Conversely the CN.getChecksum service returns the checksum as recorded in the system metadata record (not the checksum of the actual file on the CN).			
This seems to be somewhat inconsistent. MN which house large data files are expected to calculate/produce accurate checksum values while the CN which houses science metadata and resource maps does not. This makes auditing of CN replicas for correctness to be implemented differently than MN auditing. (requires auditing process to download documents from each CN to calculate checksum on the actual bytes).			
Suggestion is to move the CN.getChecksum service to be consistent with the MN.getChecksum service. Creates a more uniform REST API design and decreases need to d/l files across the CN cluster for auditing.			

History

#1 - 2014-09-18 17:03 - Matthew Jones

This design was intentional, in that the MN is reporting the 'actual' checksum for its holding, whereas the CN is reporting the putative checksum against which the MN checksums can be compared. This was because the CNs were authoritative for SystemMetadata, and held the canonical copy of that information. If you change the semantics of CN.getChecksum() to return the actual checksum, then: 1) there will be no mechanism to get the recorded checksum except by retrieving the SystemMetadata itself (probably ok), and 2) they will no longer be able to respond to getChecksum() for all objects as it does now. Some clients may be relying on the CN for the current behavior, so from an API stability perspective the safer thing to do would be to add a new method to return the actual checksum. That said, it's unlikely that clients are calling getChecksum() at all, because that info is in SystemMetadata. So... I suspect that it would not be very disruptive to change it. Maybe a quick log analysis on the CN would tell us how often it's being called from outside the CN.

Regarding the download of files across the CN cluster for auditing, it's unclear why that needs to happen. Each CN already has a copy of every file that it houses sitting on the filesystem. So a download should be unnecessary.

#2 - 2014-09-18 17:09 - Skye Roseboom

There is a single replica verified date per replica. If the CN are performing independent auditing, it will be unclear from a single date - which CN has audited itself. This is why a single CN runs an audit on a pid and audits the document on all 3 CN.

The REST API already contains services like /object which are not applicable for every document on the CN although they all appear in the object listing.