

Infrastructure - Task #2814

Story # 2720 (Closed): KNB, LTER, PISCO, ESA, SANPARKS roll out

Email all existing KNB replica partners re: 2.0.0 upgrade

2012-05-24 22:09 - Ben Leinfelder

Status:	Closed	Start date:	2012-05-24
Priority:	Normal	Due date:	
Assignee:	Ben Leinfelder	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:	Sprint-2012.35-Block.5.2	Story Points:	
Milestone:	CCI-1.0.0		
Product Version:	*		

Description

Existing replica partners will need to disable replication with the KNB when we upgrade to 2.0.0 and either reconfigure with client certificates so that traditional replication can continue or plan to become a distinct MN in the future.

In particular:

Brazil

iEcolab (Spain)

History

#1 - 2012-05-31 22:19 - Ben Leinfelder

- Status changed from New to In Progress

Here is my first draft:

Hello, fellow Metacat administrators.

We will be releasing a new version of Metacat soon and upgrading the KNB to become a DataONE Member Node. Because you control a KNB replication partner, this change will affect you.

Existing content

When the KNB becomes a DataONE Member Node, content that you have shared with the KNB will also be shared throughout the DataONE network. Any existing access control policies for your metadata and data will remain fully enforced and no private information will be accessible to users who should not have access to it. If you do not wish to have content that your Metacat instance has shared with the KNB also shared with the DataONE network, please let us know as soon as possible.

Replication changes

When the KNB upgrade is complete you will need to upgrade your Metacat instance to 2.0.0 in order to continue replicating content from your server to the KNB server. If you are running an older version at that time (1.9.5 and below) you will no longer be able to replicate documents to the KNB. If you take no action and leave your Metacat instance as is, you will see warning messages in your server log files indicating that replication to the KNB has failed, but there will be no ill-effects to either the KNB or to your server. We, of course, strongly encourage you to continue participating as a contributing member of the KNB by upgrading your Metacat deployment when convenient.

The Metacat documentation has been updated and includes instructions on how to configure Metacat replication between servers running version 2.0.0. A working draft of that documentation is available for preview here:

<http://dev-testing.dataone.org:8080/hudson/job/Metacat/javadoc/replication.html>.

Becoming a DataONE Member Node

Many of our replication partners (LTER, PISCO, ESA, SANParks) have chosen to become DataONE Member Nodes instead of continuing with the legacy KNB/Metacat replication approach. Becoming a DataONE Member Node offers many of the same benefits that KNB replication provides (redundant backup, expanded audience) and represents the future direction of our federated data sharing approach. We recognize that you may not be in a position to make this switch at this time, but we do hope that you will consider it in your future data management plans. We would be happy to discuss with you in more detail the process of becoming a DataONE Member Node if this option interests you.

Please let us know if you have any questions or concerns about the impending upgrade or how it will impact your current data management solution. Thank you,

-The Metacat development team

#2 - 2012-06-01 04:10 - Matthew Jones

Revisions from Matt Jones:

Hello fellow Metacat administrators,

We will be releasing a major new version of Metacat (2.0.0) soon with many significant new features and bug fixes. At NCEAS, we have been working with partners to expand the networked-replication approach that we've used in the KNB to include not just Metacat repositories, but also data repositories running a variety of software systems. This is happening through the DataONE project, which represents the future direction that we will be supporting for federating data repositories. Upon release of Metacat 2.0.0, we will be upgrading the KNB and several other Metacat installations to become DataONE Member Nodes. Because you control a KNB replication partner, this new release and related changes will affect you.

Existing content

When the KNB becomes a DataONE Member Node, we will be switching to a new replication approach that replaces KNB replication services and that provides greater control over replication (on a per-object basis) and broader replication of content. Content that you have shared with the KNB will now be shared via the DataONE replication services. Any existing access control policies for your metadata and data will remain fully enforced and no private information will be accessible to users who should not have access to it. This represents a great improvement to the replication services available in Metacat, but if for some reason you do not wish to have content that your Metacat instance has shared with the KNB replicated via the DataONE network, please let us know as soon as possible so that we can discuss the issues.

Replication changes

When the KNB upgrade is complete you will need to upgrade your Metacat instance to 2.0.0 in order to continue replicating content from your server to the KNB server. If you are running an older version at that time (1.9.5 and below) you will no longer be able to replicate documents to the KNB.

If you take no action and leave your Metacat instance as is, you will see warning messages in your server log files indicating that replication to the KNB has failed, but there will be no ill-effects to either the KNB or to your server. We, of course, strongly encourage you to continue participating as a contributing member of the KNB by upgrading your Metacat deployment when convenient. There are also significant bugs that were fixed in 2.0.0. The Metacat documentation has been updated and includes instructions on how to configure Metacat replication between servers running version 2.0.0. A working draft of that documentation is available for preview here: <http://dev-testing.dataone.org:8080/hudson/job/Metacat/javadoc/replication.html>.

Becoming a DataONE Member Node

Many of our replication partners (e.g., LTER, PISCO, SANParks) have chosen to become DataONE Member Nodes instead of continuing with the legacy KNB/Metacat replication approach. Becoming a DataONE Member Node offers many of the same benefits that KNB replication provides (redundant backup, expanded audience), as well as some new benefits (finer-grained replication policies, replication to non-Metacat nodes, ability to use DOIs as identifiers, etc.) and represents the future direction of our federated data sharing approach. Metacat 2.0 fully supports operation as a DataONE Member Node, and so it is fairly easy to configure the system to participate. However, we recognize that you may not be in a position to make this switch at this time, but we do hope that you will consider it in your future data management plans. We would be happy to discuss with you in more detail the process of becoming a DataONE Member Node if this option interests you.

Deprecation of existing features

Due to our changes in supporting DataONE services, we will be deprecating the existing Metacat and EcoGrid programming interfaces that some of you may have used to access Metacat via your systems. We will keep these services available in Metacat for backwards compatibility for at least one year, but we plan to remove them at some time after that to reduce the maintenance overhead associated with having all three services. In addition, Metacat 2.0.0 will be supporting the DataONE Authentication design, which relies upon public key cryptography via SSL certificates for authenticating users and systems. This is a significant change from our existing LDAP user directory system, which we will also continue to support for at least a year and maybe longer, depending on your needs. If you rely upon the Metacat or EcoGrid API for your installation, or have your own LDAP-system needs, please be sure to contact us in order to discuss your needs for migration. We would like this transition to be as seamless as possible by helping to minimize disruption associated with the transition while still realizing the major benefits from creating a larger network via DataONE.

Please let us know if you have any questions or concerns about the impending upgrade or how it will impact your current data management solution.

Thank you,

Ben Leinfelder, Chris Jones, and Matt Jones
National Center for Ecological Analysis and Synthesis

#3 - 2012-06-04 18:44 - Ben Leinfelder

- Status changed from In Progress to Closed

Sent to Brazil and iEcolab