

Arvados Epics - Story #16107

Storage classes

01/30/2020 09:31 PM - Ward Vandewege

Status:	Resolved	Start date:	03/01/2021
Priority:	Normal	Due date:	09/30/2021
Assigned To:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
Requirements and high level design			
Each collection has one or more desired storage classes. The Keep blocks making up the collection inherit the storage classes from the collection.			
<ul style="list-style-type: none">• Each keep volume is assigned one or more storage classes• The system asynchronously copies/moves blocks between Keep volumes to fulfill the desired storage classes for each block.			
When uploading a block, the Python and Go SDKs support passing a list of storage classes to influence placement of the data block. Each data block will be stored on volumes that fulfill the desired storage classes.			
<ul style="list-style-type: none">• A block that has two or more storage classes may be fulfilled by a single volume that fulfills all storage classes, or multiple volumes for each storage class.• #17349 The "replicas_desired" field expresses the number of replicas per storage class.<ul style="list-style-type: none">◦ A single volume can be configured to count as multiple replicas (existing behavior)◦ #17350 for keep-balance] A block assigned to two storage classes with N=2 replicas could have 2, 3 or 4 actual copies. For example, a block requesting storage classes A and B could be written to two volumes with storage classes "A and B", or two "A" volumes and two "B" volumes, or "A", "B" and "A,B".◦ If sufficient replicas cannot be written for each storage class during upload, it is a fatal error (existing behavior)• Storage classes are advisory and do not guarantee that a block will only be stored on a certain class or will never be stored on a certain class. Some circumstances where blocks may not be stored on only the requested storage class:<ul style="list-style-type: none">◦ User has specified an impossible situation, such as changing a collection to a storage class that isn't fulfilled by any volume◦ The same block may be referenced by multiple collections with different storage classes◦ A storage class for a collection may have been changed but the system has not caught up to it yet◦ In these cases the block will remain on its original storage volume◦ There will be a reporting tool (keep-balance or other) that reports when there is a mismatch between the desired and actual storage classes for a block. The tool will also provide a way to report which blocks (associated collections) on a certain storage class also exist on a different storage class.• When writing a block, the desired storage classes are passed in the keep service request• #17389 Keepproxy will understand storage classes and forwards blocks to the appropriate keepstore service.• When reading a collection, the Python and Go SDK support passing a list of storage classes to inform block lookup.• There is always a "default" storage class. It is an error if there is not at least one volume with the "default" class.<ul style="list-style-type: none">◦ If an upload doesn't specify a storage class it will use the 'default' storage class• Container and container request records gain a field specifying the storage class of the output collection.• #17351#17390 The following Arvados component swill gain support for specifying storage classes for data upload: arv-put, arvados-cwl-runner, crunch-run<ul style="list-style-type: none">◦ When writing a block, the desired storage classes are passed in the keep service request◦ Keepproxy will understand storage classes and forwards blocks to the appropriate keepstore service.◦ When reading a collection, the Python and Go SDK support passing a list of storage classes to inform block lookup.◦ There is always a "default" storage class. It is an error if there is not at least one volume with the "default" class.◦ If an upload doesn't specify a storage class it will use the 'default' storage class◦ Container and container request records gain a field specifying the storage class of the output collection.• Workbench 2 collection view will display the storage class of the collection.• Work will be tested with unit, functional and integration testing using Amazon S3			
Related issues:			
Related to Arvados - Feature #17349: [Keep API] Revisit "replicas_desired" a...		Rejected	
Related to Arvados - Feature #17350: [keep-balance] Expected behaviour with d...		Rejected	
Related to Arvados - Feature #17351: [arv-put] Storage classes		Resolved	06/03/2021

Related to Arvados - Feature #17388: [arv-copy] Storage classes revisit	Resolved	07/23/2021
Related to Arvados - Feature #17389: Storage classes support in keepproxy	Resolved	06/18/2021
Related to Arvados - Bug #17390: Set storage classes for intermediates and fi...	Resolved	08/12/2021
Related to Arvados - Feature #17391: [keepstore] expose the volume storage cl...	Closed	
Related to Arvados - Feature #17392: Support writing blocks to correct storag...	Resolved	04/12/2021
Related to Arvados - Feature #13382: [keepstore] Write new blocks to appropri...	Resolved	04/02/2021
Related to Arvados - Feature #17394: Go SDK CollectionFS writes files to corr...	Resolved	07/15/2021
Related to Arvados - Feature #17395: Control storage class of container / con...	Resolved	06/30/2021
Related to Arvados - Feature #17393: Go and Python SDK propagate correct stor...	Duplicate	
Related to Arvados - Support #17447: Scoping/grooming storage classes	Resolved	
Related to Arvados - Story #17465: Support writing blocks to correct storage ...	Resolved	06/01/2021
Related to Arvados - Feature #11184: [Keep] Support multiple storage classes	Resolved	
Related to Arvados - Feature #17572: arv-mount understands storage classes	Resolved	06/21/2021
Related to Arvados - Story #17573: User interface for exposing / changing sto...	Resolved	07/15/2021
Related to Arvados - Feature #17574: keep-balance updates API server with cor...	Resolved	07/29/2021
Related to Arvados - Feature #17696: Exported config has default storage clas...	Resolved	08/19/2021
Related to Arvados - Story #17697: Design for reporting tools to determine wh...	Resolved	
Related to Arvados - Feature #17698: Parallelize writes inside keepstore when...	Resolved	08/13/2021
Related to Arvados - Feature #17967: Prioritize reads from different storage ...	Resolved	08/05/2021
Blocked by Arvados - Feature #17994: [api] storage class fields should be sup...	Resolved	08/27/2021
Blocked by Arvados - Feature #17995: [api] add method to get collections wher...	Resolved	08/27/2021

History

#1 - 01/30/2020 09:35 PM - Peter Amstutz

- Start date changed from 08/01/2020 to 10/01/2020

- Due date changed from 11/30/2020 to 12/31/2020

#2 - 04/01/2020 05:26 PM - Peter Amstutz

- Start date changed from 10/01/2020 to 11/01/2020

- Due date changed from 12/31/2020 to 01/31/2021

#3 - 09/09/2020 05:30 PM - Peter Amstutz

- Start date changed from 11/01/2020 to 02/01/2021

- Due date changed from 01/31/2021 to 04/30/2021

#4 - 09/21/2020 03:07 PM - Peter Amstutz

- Subject changed from Storage classes to Storage classes / glacier

#5 - 09/21/2020 03:08 PM - Peter Amstutz

- Start date changed from 02/01/2021 to 01/01/2021

- Due date changed from 04/30/2021 to 03/31/2021

#6 - 11/16/2020 04:09 PM - Peter Amstutz

- Due date changed from 03/31/2021 to 06/30/2021

#7 - 12/09/2020 04:13 PM - Peter Amstutz

- Subject changed from Storage classes / glacier to Storage classes

#8 - 12/09/2020 04:23 PM - Peter Amstutz

- Start date changed from 01/01/2021 to 03/01/2021

#9 - 02/05/2021 08:53 PM - Nico César

- Description updated

#10 - 02/05/2021 08:53 PM - Nico César

- Related to Feature #17349: [Keep API] Revisit "replicas_desired" and "storage_classes_desired". added

#11 - 02/05/2021 09:37 PM - Nico César

- Related to Feature #17350: [keep-balance] Expected behaviour with different "replicas_desired" and "storage_classes_desired" values added

#12 - 02/05/2021 09:47 PM - Nico César

- Description updated

#13 - 02/08/2021 08:12 PM - Nico César

- Related to Feature #17351: [arv-put] Storage classes added

#14 - 02/15/2021 07:01 PM - Nico César

- Related to Feature #17388: [arv-copy] Storage classes revisit added

#15 - 02/15/2021 07:33 PM - Nico César

- Related to Feature #17389: Storage classes support in keepproxy added

#16 - 02/15/2021 07:44 PM - Nico César

- Description updated

#17 - 02/15/2021 07:49 PM - Nico César

- Related to Bug #17390: Set storage classes for intermediates and final outputs added

#18 - 02/15/2021 07:52 PM - Nico César

- Description updated

#19 - 02/15/2021 07:53 PM - Nico César

- Description updated

#20 - 02/16/2021 06:25 PM - Nico César

- Related to Feature #17391: [keepstore] expose the volume storage classes added

#21 - 02/16/2021 06:29 PM - Nico César

- Related to Feature #17392: Support writing blocks to correct storage classes in Go SDK added

#22 - 02/16/2021 06:53 PM - Nico César

- Related to Feature #13382: [keepstore] Write new blocks to appropriate storage class added

#23 - 02/16/2021 07:06 PM - Peter Amstutz

- Related to Feature #17394: Go SDK CollectionFS writes files to correct storage class added

#24 - 02/16/2021 07:22 PM - Peter Amstutz

- Related to Feature #17395: Control storage class of container / container_request output added

#25 - 02/16/2021 07:41 PM - Peter Amstutz

- Related to Feature #17393: Go and Python SDK propagate correct storage class to keepstore added

#26 - 03/03/2021 07:15 PM - Peter Amstutz

- Related to Support #17447: Scoping/grooming storage classes added

#27 - 03/10/2021 05:40 PM - Peter Amstutz

- Related to Story #17465: Support writing blocks to correct storage classes in Python SDK added

#28 - 04/12/2021 03:29 PM - Tom Clegg

- Related to Feature #11184: [Keep] Support multiple storage classes added

#29 - 04/21/2021 04:12 PM - Peter Amstutz

- Related to Feature #17572: arv-mount understands storage classes added

#30 - 04/21/2021 04:49 PM - Peter Amstutz

- Related to Story #17573: User interface for exposing / changing storage classes on a collection added

#31 - 04/21/2021 04:51 PM - Peter Amstutz

- Related to Feature #17574: keep-balance updates API server with correct storage_classes_confirmed added

#32 - 05/10/2021 02:29 PM - Peter Amstutz

Customer request: can specify desired default storage classes, which can be more than one.

#33 - 05/10/2021 02:31 PM - Peter Amstutz

Reporting tools to determine what data is on multiple storage classes.

#34 - 05/10/2021 02:32 PM - Peter Amstutz

Way to prioritize reading data from volumes with faster storage classes.

#35 - 05/10/2021 02:41 PM - Peter Amstutz

Parallelize writes inside keepstore when there are writes that request multiple storage classes.

#36 - 05/10/2021 02:48 PM - Peter Amstutz

Has use case crunch-run starting its own keepstore server.

#37 - 05/19/2021 03:50 PM - Peter Amstutz

- Related to Feature #17696: Exported config has default storage class(es), SDKs use the configured default storage class if not overridden added

#38 - 05/19/2021 03:51 PM - Peter Amstutz

- Related to Story #17697: Design for reporting tools to determine what data is on multiple storage classes. added

#39 - 05/19/2021 04:00 PM - Peter Amstutz

- Related to Feature #17698: Parallelize writes inside keepstore when there are writes that request multiple storage classes. added

#40 - 06/30/2021 03:14 PM - Peter Amstutz

- Due date changed from 06/30/2021 to 07/31/2021

#41 - 06/30/2021 08:04 PM - Peter Amstutz

- Due date changed from 07/31/2021 to 08/31/2021

#42 - 08/03/2021 07:36 PM - Peter Amstutz

- Related to Feature #17967: Prioritize reads from different storage classes added

#43 - 08/17/2021 08:07 PM - Ward Vandewege

- Blocked by Feature #17993: [deduplication-report] supports storage classes added

#44 - 08/17/2021 08:07 PM - Ward Vandewege

- Blocked by Feature #17994: [api] storage class fields should be supported in filters added

#45 - 08/17/2021 08:07 PM - Ward Vandewege

- Blocked by Feature #17995: [api] add method to get collections where replication_confirmed < replication_desired added

#46 - 08/23/2021 02:04 PM - Peter Amstutz

- Due date changed from 08/31/2021 to 09/30/2021

#47 - 10/13/2021 06:26 PM - Peter Amstutz

- Status changed from New to In Progress

#48 - 10/27/2021 04:57 PM - Peter Amstutz

- Blocked by deleted (Feature #17993: [deduplication-report] supports storage classes)

#49 - 10/27/2021 04:57 PM - Peter Amstutz

- Status changed from In Progress to Resolved