

Arvados - Story #7995

[Documentation] Document Keep Balance setup in the Install Guide

12/11/2015 04:00 PM - Brett Smith

Status:	Resolved	Start date:	02/17/2017
Priority:	Normal	Due date:	
Assigned To:	Tom Clegg	% Done:	100%
Category:	Documentation	Estimated time:	0.00 hour
Target version:	2017-03-01 sprint		
Description			
It should be as complete as any other page in the install guide. The only caveats are:			
<ul style="list-style-type: none">• It should come with a huge unmissable disclaimer at the top that Keep Balance is still being tested.• There are only two cases where we think it might be safe:<ul style="list-style-type: none">◦ All your Keepstores are backed by their own POSIX filesystem(s)◦ All your Keepstores are backed by shared object storage, one of which has a special service_type, and Data Manager talks to that one alone through its corresponding service_type switch• It should not be linked from the TOC. Enough people want it that we want a single reference to give to interested deployers, but we don't want to generally advertise it.			
Functional requirements:			
<ul style="list-style-type: none">• Document how to do a dry run/log-only run first, then how to switch that to actually deleting blocks once you're satisfied with the result.			
This is how the datamanager token is generated:			
Subtasks:			
Task # 11132: Add page to install guide			Resolved
Task # 11133: Set -enable-trash in keepstore docs			Resolved
Task # 11134: Explain limitations re shared volumes and service_type			Resolved
Task # 11119: Review 7995-keep-balance-docs			Resolved

Associated revisions

Revision 1e6a756a - 02/22/2017 09:36 PM - Tom Clegg

Merge branch '7995-keep-balance-docs'

closes #7995

History

#1 - 12/11/2015 04:03 PM - Brett Smith

- Description updated

- Category set to Documentation

#2 - 12/14/2015 02:57 PM - Brett Smith

- Target version set to Arvados Future Sprints

#3 - 12/15/2015 07:39 PM - Brett Smith

- Description updated

- Story points set to 1.0

#4 - 12/15/2015 07:42 PM - Ward Vandewege

- Description updated

#5 - 12/15/2015 07:46 PM - Brett Smith

- Description updated

#6 - 10/12/2016 08:22 PM - Tom Morris

- Subject changed from [Documentation] Document Data Manager setup in the Install Guide to [Documentation] Document Keep Balance setup in the Install Guide

- Description updated

#7 - 02/15/2017 09:01 PM - Tom Morris

- Target version changed from Arvados Future Sprints to 2017-03-01 sprint

#8 - 02/15/2017 09:06 PM - Tom Clegg

- Assigned To set to Tom Clegg

#9 - 02/17/2017 10:27 PM - Tom Clegg

- Status changed from New to In Progress

#10 - 02/20/2017 10:04 PM - Tom Clegg

7995-keep-balance-docs @ [14304c7af0b0dd7bc9345b6c5aeb61a3bdc1d3b0](https://github.com/14304c7af0b0dd7bc9345b6c5aeb61a3bdc1d3b0)

#11 - 02/21/2017 09:35 PM - Tom Morris

I made a few copy edits and pushed them to the branch. Please review them to make sure that things are still technically correct.

I didn't run linkchecker due to Python dependency issues that I couldn't be bothered to sort out.

In addition, I have the following questions/comments:

- "privileged token" is inconsistent with the name of the script "create_superuser_token"
- Creating the privileged token doesn't seem to include a name or description which can be traced back to its use as a keep-balance token. Is there a way to include some identifying information so that we know which tokens are used for what?
- What is the default setting for delete in keep stores? The implication of the "Enable delete" section is that it's disabled by default, but that's never explicitly mentioned.
- I think we should pick one preferred way of enabling delete and recommend that. Both options (along with there priority ordering for overriding each other) can be documented in the keep-balance reference page (which I can't seem to find)

Bonus semi-related comment:

- Installing keep-store page talks about setting up local file system backed storage and has a separate page for Azure blob, but S3/GCP S3 blob is not documented anywhere that I can find.

#12 - 02/22/2017 04:18 PM - Tom Clegg

Tom Morris wrote:

I made a few copy edits and pushed them to the branch. Please review them to make sure that things are still technically correct.

LGTM thanks

- "privileged token" is inconsistent with the name of the script "create_superuser_token"

Updated (here and in the crunch2 dispatch page I copied it from)

- Creating the privileged token doesn't seem to include a name or description which can be traced back to its use as a keep-balance token. Is there a way to include some identifying information so that we know which tokens are used for what?

We don't have that yet (but it does sound like a good idea)

- What is the default setting for delete in keep stores? The implication of the "Enable delete" section is that it's disabled by default, but that's never explicitly mentioned.

Default is disabled -- added a note to that effect.

- I think we should pick one preferred way of enabling delete and recommend that. Both options (along with there priority ordering for overriding each other) can be documented in the keep-balance reference page (which I can't seem to find)

YAML is the future but the keepstore install page still tells you to use command line flags, so I commented out the YAML option for now.

- Installing keep-store page talks about setting up local file system backed storage and has a separate page for Azure blob, but S3/GCP S3 blob

is not documented anywhere that I can find.

So our related-todo list is

- Name/label option for "create superuser token" script (also, shorthand for scopes, like "keep-balance")
- Update keepstore (and keep-balance) docs to configure keepstore with YAML instead of command line flags
- Document keepstore S3 volumes

#13 - 02/22/2017 04:19 PM - Javier Bértoli

Tom Morris,

- From this text:

Keep-balance can be installed anywhere with network access to Keep services. Typically it runs on the same host as keepproxy.

Keepproxy is optional, as I understand it. If so, can I have more than one keep-balance, installed in 1+ keepstores?

- From this text:

Keep-balance deletes unreferenced and overreplicated blocks from Keep servers, makes additional copies of underreplicated blocks, and moves blocks into optimal locations as needed (e.g. after adding new servers).

I understand that **keep-balance** performs three operations:

1. deletes unreferenced and overreplicated blocks from Keep servers,
2. makes additional copies of underreplicated blocks, and
3. moves blocks into optimal locations as needed

But for this text:

If you are installing keep-balance on an existing system with valuable data, you can run keep-balance in "dry run" mode first and review its logs as a precaution. To do this, use the keepstore `-never-delete=true` flag or remove the `-commit-trash` flag from your keep-balance startup script.

and this snippet:

```
~$ <span class="userinput">printf '#!/bin/sh\nexec keep-balance -commit-pulls -commit-trash 2>&1\n' | sudo tee run</span>
```

I understand that `-never-delete=true` will prevent the **FIRST** of those actions but nothing makes me assume it will prevent the other two. `-commit-trash`, which sounds like a completely different parameter from the runit example, and I suspect I'd need to disable the three parameters independently to have a **REAL** dry run:

- `never-delete=true`
- `commit-trash=false`
- `commit-pulls=false`

Am I right?

Perhaps we need to:

- Go with the established de-facto names for this operation: **--dry-run, -n or -noop** (a new ticket surely?).
- If this is not a priority now, I'd make it extra-clear in the documentation that these parameters will prevent the **THREE** operations or which of them will be **REALLY** affected, or which is the one that will perform a real dry run.

#14 - 02/22/2017 04:38 PM - Tom Clegg

Updated the dry run instructions:

To do this, edit your keep-balance startup script to use the flags `-commit-pulls=false -commit-trash=false`.

#15 - 02/22/2017 04:45 PM - Tom Clegg

Javier Bértoli wrote:

Keepproxy is optional, as I understand it. If so, can I have more than one keep-balance, installed in 1+ keepstores?

Yes, it's possible to run many things (but not the Workbench uploader) without keepproxy.

Added a bold paragraph: **A cluster should have only one keep-balance process running at a time.**

(Does that answer the question?)

#16 - 02/22/2017 05:27 PM - Javier Bértoli

Tom Clegg wrote:

Updated the dry run instructions:

To do this, edit your keep-balance startup script to use the flags `-commit-pulls=false -commit-trash=false`.

I notice you added these two flags and removed `-never-delete=true`. Is that correct, or just missed adding it?

#17 - 02/22/2017 08:05 PM - Tom Clegg

Javier Bértoli wrote:

I notice you added these two flags and removed `-never-delete=true`. Is that correct, or just missed adding it?

That's correct.

`keep-balance -commit-pulls=false -commit-trash=false` means go through the motions but don't tell the keepstore nodes to delete any blocks (or make any additional copies).

`keepstore -never-delete=true` means ignore keep-balance if it tells keepstore to delete any blocks.

#18 - 02/22/2017 08:35 PM - Javier Bértoli

Great, it LGTM, then.

#19 - 02/22/2017 09:45 PM - Tom Clegg

- Status changed from In Progress to Resolved

Applied in changeset `arvados|commit:1e6a756a10a1c0a77aeea5041844ba3a572bdd70`.