

FreeNAS - Feature #26131

Add "Auxiliary arguments" field to middleware

10/09/2017 05:30 PM - Joshua Serrine

Status: Done	Estimated time: 0.00 hour
Priority: Important	
Assignee: Vladimir Vinogradenko	
Category: Middleware	
Target version: 11.1-U6	
Severity: Med High	Needs Merging: No
Reason for Closing:	Needs Automation: No
Reason for Blocked:	Support Suite Ticket: MGA-767-84967
Needs QA: No	Hardware Configuration:
Needs Doc: No	

Description

Two customers have asked if s3 sync can be made any faster. One customer I worked with today has a 20TB dataset he is syncing to S3, and it takes more than a week to run through everything. He was hoping to do a "quick" sync after the initial fill, but that doesn't seem to be happening, and it still takes nearly a week.

This is creating a secondary problem as the server gets so busy with iops all the time for the s3 sync that ZFS replication is running at 15MB/sec to a second TrueNAS despite having a 10Gb pipe between the two servers (iperf reported >8Gb/sec between source and destination while the server was in production). The customer I worked with today is currently stuck between a rock and a hard place as they are supposed to have onsite backups to a secondary TrueNAS and offsite backups to S3. Unfortunately the S3 storage is syncing, but taking an extremely long time (almost a week) and the ZFS replication cannot keep up with the workload. The customer is currently evaluating options, but would like to be able to do S3 and replication (although maybe not expect both to be scheduled to run at the exact same time without slowdown). The support ticket is MGA-767-84967.

If there is the option to do things like skip checksums and use date/time stamps, file sizes, or other things that customers may want to consider we should consider making those options available in the WebGUI.

This may be short notice, but I think this should be considered for 11.1 since this issue is likely to get much bigger as time goes on, and I don't want us to be in a situation where this is slated for 11.2 and we have some big ticket customer with a massive server demanding this be added immediately. This issue is clearly for large-scale servers, which is our bread-and-butter.

Thanks.

Related issues:

Related to FreeNAS - Feature #35107: Allow passing auxiliary arguments to rclone	Closed
Related to FreeNAS - Bug #36523: Hide Auxiliary Arguments from Cloud Sync task	Done
Related to FreeNAS - Bug #40968: Fix traceback when trying to add a Cloud Syn...	Done
Related to FreeNAS - Bug #52507: Add Transfers field to Cloud Sync in legacy UI	Ready for Testing

Associated revisions

Revision 0d57a8e1 - 06/17/2018 09:50 AM - Vladimir Vinogradenko

feat(cloud_sync): Allow passing auxiliary arguments to rclone

Ticket: #26131

Revision 1165f6e3 - 06/17/2018 10:05 AM - Vladimir Vinogradenko

feat(cloud_sync): Allow passing auxiliary arguments to rclone

Ticket: #26131

Revision fdcc6764 - 06/22/2018 12:27 PM - Vladimir Vinogradenko

feat(cloud_sync): Allow passing auxiliary arguments to rclone

Ticket: #26131

Revision 5ff265e4 - 06/25/2018 12:04 PM - Dru Lavigne

Doc Auxiliary arguments for cloud sync.

Ticket: #26131

History

#1 - 10/09/2017 08:39 PM - Dru Lavigne

- Status changed from *Untriaged* to *Unscreened*

#2 - 10/13/2017 12:57 PM - Bartosz Prokop

- Assignee changed from *Bartosz Prokop* to *Dru Lavigne*

- Priority changed from *Nice to have* to *Important*

It looks like rclone does not support incremental backups. Maybe it's worth using a different tool for this task?

RC: Please reassingn. This is one of the last tickets without proper 'L3' status.

#3 - 10/14/2017 11:00 AM - Dru Lavigne

- Assignee changed from *Dru Lavigne* to *Marcelo Araujo*

#4 - 10/17/2017 08:32 AM - Marcelo Araujo

- Status changed from *Unscreened* to *Screened*

- Target version set to *TrueNAS 11.1-U2*

#5 - 01/05/2018 01:11 AM - Marcelo Araujo

- Assignee changed from *Marcelo Araujo* to *Bartosz Prokop*

Bartosz do you mind take a look on it? Feel free to send it back to me in case you are busy with other things too.

#6 - 01/05/2018 04:17 AM - Dru Lavigne

- Status changed from *Screened* to *Unscreened*

#7 - 01/05/2018 09:19 AM - Joshua Serrine

How do we go backwards in this process?

#8 - 01/11/2018 12:19 PM - Bartosz Prokop

- Status changed from *Unscreened* to *Screened*

#9 - 01/11/2018 01:35 PM - Warren Block

- Subject changed from *add option to choose between MD5 hashes and date/time tamps for S3 sync* to *add option to choose between MD5 hashes and date/time stamps for S3 sync*

#10 - 01/23/2018 05:48 AM - Ben Gadd

- Due date set to *02/12/2018*

#11 - 01/26/2018 09:11 AM - Dru Lavigne

- Project changed from TrueNAS to FreeNAS
- Category changed from OS to OS
- Status changed from Screened to Not Started
- Target version changed from TrueNAS 11.1-U2 to 11.2-RC2
- Hide from ChangeLog deleted (No)
- Support Department Priority deleted (0)

#12 - 01/26/2018 04:52 PM - Eric Loewenthal

I hate to pile on more work, but this seems to be *begging* for zfs diff.

#13 - 02/08/2018 11:31 AM - Ben Gadd

- Due date deleted (02/12/2018)

#14 - 02/28/2018 06:06 AM - William Grzybowski

- Category changed from OS to Middleware
- Assignee changed from Bartosz Prokop to Vladimir Vinogradenko

Load balancing Bartosz tickets.

#15 - 03/23/2018 02:04 PM - Vladimir Vinogradenko

- Status changed from Not Started to Blocked
- Reason for Blocked set to Waiting for feedback

<https://rclone.org/docs/>

Normally rclone will look at modification time and size of files to see if they are equal.

I've checked this, this is correct:

```
$ rclone --config config -vv --stats 1s sync /mnt/data/aws remote:ixsystems/  
...  
2018/03/23 23:57:33 DEBUG : test: Size and modification time the same (differ by 0s, within tolerance 1µs)  
2018/03/23 23:57:33 DEBUG : test: Unchanged skipping
```

The file on remote is not updated if I update it locally but preserve size and mtime. The reason for slow sync must be something else.

Can you please run verbose rclone by hand to see what is taking it so long? You'll need to create config file

```
[remote]  
type = s3  
env_auth = false  
access_key_id = ...  
secret_access_key = ...  
region =
```

and run it like this

```
rclone --config /path/to/config/file -vv --stats 1s sync /local/path remote:BUCKETNAME/BUCKET/PATH
```

#16 - 03/23/2018 03:30 PM - Joshua Serrine

I'll see if I can reach out to the customer about this one.

Stay tuned.

#17 - 04/16/2018 11:58 AM - Vladimir Vinogradenko

We might try official AWS CLI S3 tool

Install it with

```
# pip install awscli
```

Put your credentials into ~/.aws/credentials:

```
[default]
aws_secret_access_key = ...
aws_access_key_id = ...
```

Run it with

```
# aws s3 sync /local/path s3://bucket-name/path
```

#18 - 05/01/2018 07:38 PM - Vladimir Vinogradenko

- *Severity set to Med High*

#19 - 05/10/2018 09:00 AM - Joshua Serrine

So we tested a directory with 109915 files. No updates were needed in both cases.

```
rclone:
real = 7m46s
user = 2m1s
sys = 0m25s
```

```
aws cli:
real = 3m45s
user = 1m52s
sys = 0m23s
```

So the time was halved.

We then did a directory that had few files, but the files were large (200+GB), total was almost 600GB of files.

```
rclone:
```

We didn't let it finish because the throughput was 8MB/sec and slowly decreasing, with the estimated completion being >6 hours.

```
aws cli:
```

consistently 20MB/sec+, typically 24-29MB/sec.

In my opinion, we should consider switching from rclone to aws cli.

Thanks.

#20 - 05/10/2018 09:28 AM - Vladimir Vinogradenko

That's great!

I think we should add "use awscli" option for S3 targets (which would also prohibit using rclone-specific features like encryption). William, what do you think?

#21 - 05/10/2018 10:10 AM - William Grzybowski

Talked to Vladimir via RC.

We are not going to move away from rclone on the first problem found. It is a fantastic tool and works with dozens of different backends.

We will try to work with maintainers to track down the issue.

If performance is a big issue s3cmd can be used as a workaround for these people in the meantime.

#22 - 05/15/2018 09:43 AM - Joshua Serrine

William Grzybowski wrote:

If performance is a big issue s3cmd can be used as a workaround for these people in the meantime.

I'm trying to schedule a test of a new rclone package to see how it does, but s3cmd is a no-go for my customer.

With a warning like this when you run it:

1. s3cmd

WARNING: !!!!! Support for python3 is currently in a 'Work In Progress' state.

Please don't use s3cmd with python3 on production tasks or with sensitive data as unexpected behaviors could occur !!!!!

The customer isn't even interested in using it for testing purposes.

#23 - 05/15/2018 09:53 AM - William Grzybowski

Joshua Serrine wrote:

William Grzybowski wrote:

If performance is a big issue s3cmd can be used as a workaround for these people in the meantime.

I'm trying to schedule a test of a new rclone package to see how it does, but s3cmd is a no-go for my customer.

With a warning like this when you run it:

1. s3cmd

WARNING: !!!!!!! Support for python3 is currently in a 'Work In Progress' state.

Please don't use s3cmd with python3 on production tasks or with sensitive data as unexpected behaviors could occur !!!!!!!

The customer isn't even interested in using it for testing purposes.

We can add awscli to the image if thats an issue.

#24 - 05/18/2018 09:10 AM - Joshua Serrine

Vlad gave me an install of rclone 1.41. I installed it and the time was 7m15s with no changes on the same dataset I had done above. So faster, but only about 30 seconds faster.

Unfortunately, the dataset with the large files has been running since midnight, so we couldn't run ours since the other was still running. So I was not able to do any raw throughput tests.

Installed 1.38 on the customers box at the end of the call.

#25 - 05/18/2018 10:00 AM - William Grzybowski

Joshua Serrine wrote:

Vlad gave me an install of rclone 1.41. I installed it and the time was 7m15s with no changes on the same dataset I had done above. So faster, but only about 30 seconds faster.

Unfortunately, the dataset with the large files has been running since midnight, so we couldn't run ours since the other was still running. So I was not able to do any raw throughput tests.

Installed 1.38 on the customers box at the end of the call.

Did you use --use-server-modtime flag ?

#26 - 05/23/2018 05:39 AM - William Grzybowski

William Grzybowski wrote:

Joshua Serrine wrote:

Vlad gave me an install of rlcne 1.41. I installed it and the time was 7m15s with no changes on the same dataset I had done above. So faster, but only about 30 seconds faster.

Unfortunately, the dataset with the large files has been running since midnight, so we couldn't run ours since the other was still running. So I was not able to do any raw throughput tests.

Installed 1.38 on the customers box at the end of the call.

Did you use --use-server-modtime flag ?

Ping?

#27 - 05/24/2018 08:18 AM - Joshua Serrine

William Grzybowski wrote:

William Grzybowski wrote:

Joshua Serrine wrote:

Vlad gave me an install of rlcne 1.41. I installed it and the time was 7m15s with no changes on the same dataset I had done above. So faster, but only about 30 seconds faster.

Unfortunately, the dataset with the large files has been running since midnight, so we couldn't run ours since the other was still running. So I was not able to do any raw throughput tests.

Installed 1.38 on the customers box at the end of the call.

Did you use --use-server-modtime flag ?

Ping?

I did not. The message I got in RC made no mention of using that flag.

#28 - 05/24/2018 08:35 AM - William Grzybowski

Joshua Serrine wrote:

William Grzybowski wrote:

William Grzybowski wrote:

Joshua Serrine wrote:

Vlad gave me an install of rlcne 1.41. I installed it and the time was 7m15s with no changes on the same dataset I had done above. So faster, but only about 30 seconds faster.

Unfortunately, the dataset with the large files has been running since midnight, so we couldn't run ours since the other was still running. So I was not able to do any raw throughput tests.

Installed 1.38 on the customers box at the end of the call.

Did you use --use-server-modtime flag ?

Ping?

I did not. The message I got in RC made no mention of using that flag.

I think it would be worth a try.

#29 - 05/24/2018 08:36 AM - Vladimir Vinogradenko

Joshua, please run the same as in <https://redmine.ixsystems.com/issues/26131#note-15> but with --use-server-modtime flag.

Also did this problem

We didn't let it finish because the throughput was 8MB/sec and slowly decreasing, with the estimated completion being >6 hours.

persist after upgrade to 1.41?

#30 - 05/24/2018 10:24 AM - Joshua Sirrine

Vladimir Vinogradenko wrote:

Also did this problem

We didn't let it finish because the throughput was 8MB/sec and slowly decreasing, with the estimated completion being >6 hours.

persist after upgrade to 1.41?

I wasn't able to test it the last time around because the system was doing an update that had been running all night. In a prior call, the time of day was late enough that we were able to test it, but couldn't for the most recent call.

I will reach out to the customer to test this.

#31 - 05/29/2018 10:36 AM - Ben Gadd

- Support Suite Ticket changed from n/a to MGA-767-84967

#32 - 05/29/2018 11:22 AM - Joshua Sirrine

Vladimir,

Is there anything else we want to test? This will be my third call with the customer to test things. If you have more things you'd like tested, I'd prefer to do that rather than a 4th call. You can also choose to be on the call with me if you'd like.

Let me know. I'm still trying to schedule something with the customer.

#33 - 05/29/2018 11:26 AM - Vladimir Vinogradenko

Joshua, no, that's it, after these tests I'll file a bug to rclone.

#34 - 06/06/2018 07:15 AM - Joshua Sirrinc

So the customer and I tried to run rclone 1.41 with the --use-server-modtime and things didn't go well. Files that hadn't been updated in months (and sometimes many years) were resyncing. Screenshot is attached showing the behavior. I informed Vladimir about this issue and the fact that the test had been running for almost 2 hours and wasn't complete and he recommended we try the no-update-modtime parameter instead. Here's the results:

With no-update-modtime I got the following results on 2 runs:

Transferred: 0 Bytes (0 Bytes/s)
Errors: 0
Checks: 135742
Transferred: 0
Elapsed time: 8m17.9s

Transferred: 0 Bytes (0 Bytes/s)
Errors: 0
Checks: 135742
Transferred: 0
Elapsed time: 9m34.6s

So the awscli test we did some weeks ago was still significantly faster.

Thanks.

#35 - 06/06/2018 08:14 AM - Vladimir Vinogradenko

Joshua, it's suspicious to see:

- a) time increase comparing to regular rclone
- b) big time difference in two identical runs

Dataset or network conditions might change since we've ran rclone without these options. Can you please also repeat old test with current dataset and network conditions?

#36 - 06/06/2018 08:24 AM - William Grzybowski

We cant keep bugging our customers to do science experiment all the time.

If think we need to try and replicate the scenario locally. Joshua could you help us set a test scenario?

#37 - 06/08/2018 07:09 AM - Vladimir Vinogradenko

Joshua, have you been running your last tests with -vv? It's my mistake, using -vv affects performance (especially on slow SSH connections). It would

be glad if you would be able to re-run test with just -v. On my synthetic test:

- old rclone: 5m40s
- awscli: 44s
- new rclone: 7s

#38 - 06/12/2018 12:11 PM - Vladimir Vinogradenko

--checkers=128 helped us to speed up syncs with lots of small files

Problem with speed dropping while transferring large files still remains. It's not only --s3-disable-checksum, it's something more.

#39 - 06/13/2018 03:46 AM - Vladimir Vinogradenko

Low speed upload bug fixed. My PR to rclone: <https://github.com/ncw/rclone/pull/2351>

#40 - 06/17/2018 10:10 AM - Vladimir Vinogradenko

- Related to Feature #35107: Allow passing auxiliary arguments to rclone added

#41 - 06/17/2018 10:14 AM - Vladimir Vinogradenko

- Status changed from Blocked to In Progress

New rclone with this fix was released yesterday.

New field will appear in cloud sync: **Auxiliary arguments**

For case with small files user should add: --use-server-modtime --no-update-modtime --checkers=128

For huge files user should add: --s3-upload-concurrency 16 (also --s3-disable-checksum if he likes)

#42 - 06/19/2018 05:36 PM - Dru Lavigne

- Subject changed from add option to choose between MD5 hashes and date/time stamps for S3 sync to Add "Auxiliary arguments" field to Cloud Sync

- Target version changed from 11.2-RC2 to 11.1-U6

- Needs Merging changed from Yes to No

Master PR: <https://github.com/freenas/freenas/pull/1385>

Stable PR: <https://github.com/freenas/freenas/pull/1384>

#43 - 06/20/2018 05:47 AM - Vladimir Vinogradenko

- Status changed from In Progress to Ready for Testing

#44 - 06/25/2018 12:07 PM - Dru Lavigne

- Needs Doc changed from Yes to No

Legacy UI doc commit: <https://github.com/freenas/freenas-docs/commit/5ff265e4a3a2f2bc70b9be83573bc4e82f036a1f>

New UI doc commit will occur in related ticket, once added to the new UI.

#45 - 07/03/2018 12:40 PM - Dru Lavigne

- Related to Bug #36523: Hide Auxiliary Arguments from Cloud Sync task added

#46 - 07/03/2018 01:08 PM - William Grzybowski

- Subject changed from Add "Auxiliary arguments" field to Cloud Sync to Add "Auxiliary arguments" field to middleware

- Status changed from Ready for Testing to Unscreened

Vladimir, as discussed, we will hide this from UI for now.

Thanks!

#47 - 07/05/2018 02:21 PM - Vladimir Vinogradenko

- Status changed from Unscreened to Ready for Testing

PR: <https://github.com/freenas/freenas/pull/1488>

#49 - 08/07/2018 01:08 PM - Bonnie Follweiler

- Status changed from Ready for Testing to Blocked

This test is blocked by <https://redmine.ixsystems.com/issues/40968> since I can't create a Cloud Sync task until that ticket is done

#50 - 08/07/2018 01:08 PM - Bonnie Follweiler

- Related to Bug #40968: Fix traceback when trying to add a Cloud Sync task added

#51 - 08/08/2018 12:36 PM - Bonnie Follweiler

- Reason for Blocked changed from Waiting for feedback to Dependent on a related task to be completed

#52 - 08/13/2018 11:45 AM - Dru Lavigne

- Status changed from Blocked to Ready for Testing

#53 - 08/15/2018 07:24 AM - Bonnie Follweiler

- Status changed from Ready for Testing to Passed Testing

- Needs QA changed from Yes to No

Test Passed in FreeNAS-11.1-U6-INTERNAL4

#54 - 08/15/2018 02:07 PM - Dru Lavigne

- Status changed from Passed Testing to Done

- Reason for Blocked deleted (Dependent on a related task to be completed)

#55 - 10/22/2018 06:08 AM - Vladimir Vinogradenko

- Related to Bug #52507: Add Transfers field to Cloud Sync in legacy UI added