

FreeNAS - Feature #25182

Add Offline button for faulted devices

07/17/2017 08:58 PM - Stuart Espey

Status:	Resolved	Estimated time:	0.00 hour
Priority:	Important		
Assignee:	William Grzybowski		
Category:	GUI (new)		
Target version:	11.1-BETA1		
Severity:	New	Needs Merging:	Yes
Reason for Closing:		Needs Automation:	No
Reason for Blocked:		Support Suite Ticket:	n/a
Needs QA:	No	Hardware Configuration:	
Needs Doc:	Yes		

Description

Tested with 11.0-U1

In the Storage Volumes tab, if you click on the volume, and then the Volume status button, you can see all the component devices which make up a vdev.

If you click on a component device, then you can click offline.

But if the component device is faulted, then you can not.

You should be able to.

The reason is because if you wish to replace a faulted device with itself, then you have to offline the device, quick wipe it, and then replace the offlined device with the now wiped drive. You might do this if you have run a long test over the drive, and the drive appears okay, and you suspect the issue which caused the faulted device was not due to the device per-se, either way, you want to retry the device.

Without the ability to offline a faulted device in the GUI, you need to offline it via the CLI. The rest of the steps can be done in the UI. Without offlining, you can't wipe the drive without unmounting the pool.

An example of this situation:

```
root@rhea:~ # zpool status tank
pool: tank
state: DEGRADED
status: One or more devices are faulted in response to persistent errors.
       Sufficient replicas exist for the pool to continue functioning in a
       degraded state.
action: Replace the faulted device, or use 'zpool clear' to mark the device
       repaired.
scan: scrub repaired 88K in 5h11m with 0 errors on Sun Jun 18 05:11:13 2017
config:

   NAME                                STATE      READ  WRITE  CKSUM
   tank
   raidz2-0                             DEGRADED   0      0      0
   gptid/e16ecdcb-86f2-11e6-bb77-001cc0071f3f ONLINE     0      0      0
   gptid/e24c858c-86f2-11e6-bb77-001cc0071f3f ONLINE     0      0      0
   gptid/e3250ec0-86f2-11e6-bb77-001cc0071f3f ONLINE     0      0      0
   gptid/e3e200ca-86f2-11e6-bb77-001cc0071f3f FAULTED    0     123    0  too many errors
   gptid/e4b6d326-86f2-11e6-bb77-001cc0071f3f ONLINE     0      0      0

errors: No known data errors
```

because I can't offline the faulted device via UI, I now have to do it via CLI:

```
root@rhea:~ # zpool offline tank gptid/e3e200ca-86f2-11e6-bb77-001cc0071f3f
```

and you can see the device has successfully been offlined.

```
root@rhea:~ # zpool status tank
```

```
pool: tank
state: DEGRADED
status: One or more devices has experienced an unrecoverable error. An
attempt was made to correct the error. Applications are unaffected.
action: Determine if the device needs to be replaced, and clear the errors
using 'zpool clear' or replace the device with 'zpool replace'.
see: http://illumos.org/msg/ZFS-8000-9P
scan: scrub repaired 88K in 5h11m with 0 errors on Sun Jun 18 05:11:13 2017
config:
```

NAME	STATE	READ	WRITE	CKSUM	
tank	DEGRADED	0	0	0	
raidz2-0	DEGRADED	0	0	0	
gptid/e16ecdcb-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
gptid/e24c858c-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
gptid/e3250ec0-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
2223584477133409854	OFFLINE	0	123	0	was /dev/gptid/e3e
200ca-86f2-11e6-bb77-001cc0071f3f					
gptid/e4b6d326-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	

```
errors: No known data errors
```

And now, when I wipe/replace via the GUI, it begins resilvering.

```
pool: tank
```

```
state: ONLINE
status: One or more devices is currently being resilvered. The pool will
continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
scan: resilver in progress since Tue Jul 18 13:44:02 2017
145M scanned out of 8.74T at 7.62M/s, 334h12m to go
27.0M resilvered, 0.00% done
config:
```

NAME	STATE	READ	WRITE	CKSUM	
tank	ONLINE	0	0	0	
raidz2-0	ONLINE	0	0	0	
gptid/e16ecdcb-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
gptid/e24c858c-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
gptid/e3250ec0-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	
gptid/4dfc43ea-6b6b-11e7-bea4-001cc0071f3f	ONLINE	0	0	0	(resilvering)
gptid/e4b6d326-86f2-11e6-bb77-001cc0071f3f	ONLINE	0	0	0	

I'll check the results, and scrub etc, and keep an eye on the drive, and if it faults again, I'll replace it (as its <12 months old), but as the SMART results aren't showing any errors, I can't currently easily RMA, until I prove a drive fault.

I haven't checked if this is supported in the new UI, but if it isn't it should be too.

Related issues:

Associated revisions

Revision 3c4a8058 - 07/18/2017 06:45 AM - William Grzybowski

fix(gui): show Offline for FAULTED disks

Ticket: #25182

Revision f594b7f6 - 07/21/2017 05:44 PM - William Grzybowski

fix(gui): show Offline for FAULTED disks

Ticket: #25182

History

#1 - 07/17/2017 09:05 PM - Stuart Espey

Maybe the correct course of action is to issue a "zpool clear", not sure if that would resolve the faulted state, but either way, you should be able to offline a faulted drive.

#2 - 07/18/2017 03:52 AM - Dru Lavigne

- Assignee changed from Release Council to William Grzybowski

William: thoughts?

#3 - 07/18/2017 04:33 AM - Stuart Espey

Dru Lavigne wrote:

William: thoughts?

So, I'm working through my little disk problem, and I was able to reproduce it again. (FWIW, I believe its a backplane issue)

It appears the right approach to clear the faulted case is to run zpool clear. As soon as that is executed, the FAULTED device is switched to ONLINE and begins resilvering.

So, is there a way in the GUI to clear the status of a pool?

#4 - 07/18/2017 05:28 AM - William Grzybowski

- Status changed from Unscreened to Screened

- Target version set to 11.2-BETA1

While Offline'ing a device in that case is possible your procedure does not make any sense to me.

If such a thing happens the correct way to repair is to scrub the pool, as opposed to offline a disk, wipe it and replace. It has absolutely no advantage.

zpool clear is also simply masquerading the problem, it assumes you have actually replaced something. I suggest you to read more about ZFS to understand how it works, oracle has a nice documentation.

I'll enable Offline button for faulted disks, although thats not a recommended approach.

#5 - 07/18/2017 06:47 AM - William Grzybowski

- Status changed from Screened to Ready For Release
- Priority changed from Expected to Important
- Target version changed from 11.2-BETA1 to 11.1

#6 - 07/18/2017 06:48 AM - William Grzybowski

- Needs QA changed from Yes to No

#7 - 08/31/2017 11:20 AM - Dru Lavigne

- Related to Bug #25737: Clarify replace disk instructions in Guide added

#8 - 09/27/2017 06:40 AM - Dru Lavigne

- Target version changed from 11.1 to 11.1-BETA1

#9 - 10/24/2017 04:47 AM - Dru Lavigne

- Status changed from Ready For Release to Resolved