
Subject: Can migration continue after IO errors?

Posted by [wsap](#) on Mon, 04 Oct 2021 17:05:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

Ran vzmigrate on a container that resulted in this in the phaul log:

2021-10-02 20:06:11.382: 1023556: Preliminary FS migration

2021-10-02 20:47:21.046: 1023556: Error in send_image_block (ploop-copy.c:665): Error from pread() size=1048576 pos=352539639808: Input/output error

Traceback (most recent call last):

File "/usr/libexec/phaul/p.haul", line 11, in <module>

load_entry_point('phaul==0.1', 'console_scripts', 'p.haul')()

File "/usr/lib/python2.7/site-packages/phaul/shell/phaul_client.py", line 49, in main

worker.start_migration()

File "/usr/lib/python2.7/site-packages/phaul/iters.py", line 175, in start_migration

self.__start_restart_migration()

File "/usr/lib/python2.7/site-packages/phaul/iters.py", line 295, in __start_restart_migration

fsstats = self.fs.start_migration()

File "/usr/lib/python2.7/site-packages/phaul/fs_haul_ploop.py", line 127, in start_migration

total_xferred += ploopcopy.copy_start()

File "/usr/lib64/python2.7/site-packages/libploop/__init__.py", line 16, in copy_start

return libploopapi.copy_start(self.h)

RuntimeError: Error in send_image_block (ploop-copy.c:665): Error from pread() size=1048576 pos=352539639808: Input/output error

Kernel: 3.10.0-1160.21.1.vz7.174.13

OS: Virtuozzo Linux release 7.9

When phaul encountered that error, vzmigrate was hung with no data transferring, the destination ploop file not increasing, and the source vzmigrate still running, but not actually doing anything.

Is there anything I can pass to vzmigrate or an environment variable that can be set prior to running that will ensure that when it hits IO errors, it reports on them, but continues on to further files/blocks?

I have reason to believe this is because of bad sectors on the disk, which is exactly why I'm trying to migrate containers away from this node, however if it halts the migration upon *any* error like this, I can't actually get the containers away from the node with failing disks... bit of a catch 22.

Any suggestions?
