

ceph-volume - Bug #24993

ceph-volume fails to create OSD with Python 3

07/19/2018 08:24 AM - Michael Jones

Status:	Resolved	Start date:	07/19/2018
Priority:	High	Due date:	
Assignee:	Alfredo Deza	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Reviewed:	
Source:	Community (user)	Affected Versions:	v13.2.0
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No		
Severity:	2 - major		

Description

ceph version 13.2.0 (79a10589f1f80dfe21e8f9794365ed98143071c4) mimic (stable)
Python 3.5.5
OS: Gentoo

```
fenrir ~ # ceph-volume lvm create --bluestore --data /dev/sda
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/
bootstrap-osd/ceph.keyring -i - osd new alff1660-7007-4d52-9ec9-bbb7bfc63fd8
--> TypeError: memoryview: a bytes-like object is required, not 'str'
```

The log gives no indication of WHY ceph-volume was unable to prepare the device.

Note, this was a disk that was the bluestore storage of an OSD created by ceph-disk, so this isn't a hardware issue, or an issue of a non-working cluster. I'm just trying to migrate my OSD to ceph-volume.

/var/log/ceph/ceph-volume.log:

```
[2018-07-19 03:03:31,064][ceph_volume.main][INFO ] Running command: ceph-volume lvm create --blu
estore --data /dev/sda
[2018-07-19 03:03:31,074][ceph_volume.process][INFO ] Running command: /usr/bin/ceph-authtool --g
en-print-key
[2018-07-19 03:03:31,143][ceph_volume.process][INFO ] stdout AQDTRVBblYwpCBAA5BCd44IVFlUJhyrm2ox9
kQ==
[2018-07-19 03:03:31,146][ceph_volume.process][INFO ] Running command: /usr/bin/ceph --cluster ce
ph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new alf
f1660-7007-4d52-9ec9-bbb7bfc63fd8
[2018-07-19 03:03:31,152][ceph_volume.devices.lvm.prepare][ERROR ] lvm prepare was unable to compl
ete
[2018-07-19 03:03:31,152][ceph_volume.devices.lvm.prepare][INFO ] will rollback OSD ID creation
[2018-07-19 03:03:31,153][ceph_volume][ERROR ] exception caught by decorator
Traceback (most recent call last):
  File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
    return f(*a, **kw)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
    terminal.dispatch(self.mapper, subcommand_args)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
    terminal.dispatch(self.mapper, self.argv)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
```

```

instance.main()
File "/usr/lib/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
    self.create(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
    prepare_step.safe_prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_
prepare
    self.prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 245, in prepa
re
    self.osd_id = prepare_utils.create_id(osd_fsid, json.dumps(secrets), osd_id=args.osd_id)
File "/usr/lib64/python3.5/site-packages/ceph_volume/util/prepare.py", line 72, in create_id
    show_command=True
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 200, in call
    stdout_stream, stderr_stream = process.communicate(stdin)
File "/usr/lib64/python3.5/subprocess.py", line 803, in communicate
    stdout, stderr = self._communicate(input, endtime, timeout)
File "/usr/lib64/python3.5/subprocess.py", line 1441, in _communicate
    input_view = memoryview(self._input)
TypeError: memoryview: a bytes-like object is required, not 'str'

```

Creating an lvm volume by hand doesn't help matters either.

```

pvcreate /dev/sda ; vgcreate VGroup1 /dev/sda ; lvcreate -l100%VG VGroup1 ; ceph-volume lvm create
--bluestore --data VGroup1/lvol0
Physical volume "/dev/sda" successfully created.
Volume group "VGroup1" successfully created
Logical volume "lvol0" created.
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/
bootstrap-osd/ceph.keyring -i - osd new 5a225676-2d28-49bb-a538-f9419f510ffe
--> TypeError: memoryview: a bytes-like object is required, not 'str'

```

/var/log/ceph/ceph-volume.log:

```

[2018-07-19 03:11:29,138][ceph_volume.main][INFO ] Running command: ceph-volume lvm create --blu
estore --data VGroup1/lvol0
[2018-07-19 03:11:29,149][ceph_volume.process][INFO ] Running command: /usr/bin/ceph-authtool --g
en-print-key
[2018-07-19 03:11:29,227][ceph_volume.process][INFO ] stdout AQCxR1BbSDP7DBAAjwyHEhhoyoUnF41tq0If
XQ==
[2018-07-19 03:11:29,230][ceph_volume.process][INFO ] Running command: /usr/bin/ceph --cluster ce
ph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 5a2
25676-2d28-49bb-a538-f9419f510ffe
[2018-07-19 03:11:29,235][ceph_volume.devices.lvm.prepare][ERROR ] lvm prepare was unable to compl
ete
[2018-07-19 03:11:29,236][ceph_volume.devices.lvm.prepare][INFO ] will rollback OSD ID creation
[2018-07-19 03:11:29,237][ceph_volume][ERROR ] exception caught by decorator
Traceback (most recent call last):
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
    return f(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
    terminal.dispatch(self.mapper, subcommand_args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main

```

```

terminal.dispatch(self.mapper, self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
self.create(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
prepare_step.safe_prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_
prepare
self.prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 245, in prepa
re
self.osd_id = prepare_utils.create_id(osd_fsid, json.dumps(secrets), osd_id=args.osd_id)
File "/usr/lib64/python3.5/site-packages/ceph_volume/util/prepare.py", line 72, in create_id
show_command=True
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 200, in call
stdout_stream, stderr_stream = process.communicate(stdin)
File "/usr/lib64/python3.5/subprocess.py", line 803, in communicate
stdout, stderr = self._communicate(input, endtime, timeout)
File "/usr/lib64/python3.5/subprocess.py", line 1441, in _communicate
input_view = memoryview(self._input)
TypeError: memoryview: a bytes-like object is required, not 'str'

```

```

fenrir ~ # pvdisplay
--- Physical volume ---
PV Name           /dev/sda
VG Name           VGroup1
PV Size           931.51 GiB / not usable 1.71 MiB
Allocatable       yes (but full)
PE Size           4.00 MiB
Total PE          238467
Free PE           0
Allocated PE      238467
PV UUID           hiLPNH-Tr4L-ZGna-s4QW-yPqu-iPgC-Vcquyh

```

```

fenrir ~ # vgdisplay
--- Volume group ---
VG Name           VGroup1
System ID
Format            lvm2
Metadata Areas    1
Metadata Sequence No 2
VG Access         read/write
VG Status         resizable
MAX LV            0
Cur LV           1
Open LV           0
Max PV            0
Cur PV           1
Act PV            1
VG Size           931.51 GiB
PE Size           4.00 MiB
Total PE          238467
Alloc PE / Size   238467 / 931.51 GiB
Free PE / Size    0 / 0
VG UUID           XD8JW4-CmeL-mYf2-Idp7-JBt4-pJ2d-JNqJ7t

```

```

fenrir ~ # lvdisplay
--- Logical volume ---
LV Path           /dev/VGroup1/lvol0

```

```
LV Name          lv010
VG Name          VGroup1
LV UUID          MUBIEC-wxRu-Fje3-iTl6-dw6u-kKz8-vdeGPY
LV Write Access  read/write
LV Creation host, time fenrir, 2018-07-19 03:11:27 -0500
LV Status        available
# open          0
LV Size          931.51 GiB
Current LE       238467
Segments        1
Allocation       inherit
Read ahead sectors auto
- currently set to 256
Block device     254:0
```

So, I have no idea how to proceed here.

```
TypeError: memoryview: a bytes-like object is required, not 'str'
```

Isn't a meaningful error message. Heck, combine "ceph-volume" with the python error message on Google and you get 2 results total, and neither of them has anything to do with either ceph or python.

```
lvm prepare was unable to complete
will rollback OSD ID creation
```

Doesn't give me any direction in how to trouble shoot... since I can (and did) create the logical volume by hand, I'm at a complete loss on what to investigate.

And, it's also a liar. It totally did not roll back the OSD id creation. Had to purge that by hand (ceph osd purge # --yes-i-really-mean-it).

History

#1 - 07/19/2018 08:37 AM - Michael Jones

Creating a partition doesn't help any either.

```
fenrir ~ # gdisk /dev/sda
GPT fdisk (gdisk) version 1.0.3
```

```
Partition table scan:
```

```
  MBR: not present
  BSD: not present
  APM: not present
  GPT: not present
```

```
Creating new GPT entries.
```

```
Command (? for help): n
Partition number (1-128, default 1):
First sector (34-1953525134, default = 2048) or {+}size{KMGTP}:
Last sector (2048-1953525134, default = 1953525134) or {+}size{KMGTP}:
Current type is 'Linux filesystem'
Hex code or GUID (L to show codes, Enter = 8300):
Changed type of partition to 'Linux filesystem'
```

```
Command (? for help):
```

```
Command (? for help): w
```

```
Final checks complete. About to write GPT data. THIS WILL OVERWRITE EXISTING
PARTITIONS!!
```

```
Do you want to proceed? (Y/N): y
OK; writing new GUID partition table (GPT) to /dev/sda.
The operation has completed successfully.
fenrir ~ # ceph-volume lvm --log-level 5 create --bluestore --data /dev/sdal
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 0e5db055-f7ce-4a3b-b5fc-af707e6fa64c
--> TypeError: memoryview: a bytes-like object is required, not 'str'
```

```
[2018-07-19 03:36:51,553][ceph_volume.main][INFO ] Running command: ceph-volume lvm --log-level 5 create --bluestore --data /dev/sdal
[2018-07-19 03:36:51,561][ceph_volume.process][INFO ] Running command: /usr/bin/ceph-authtool --gen-print-key
[2018-07-19 03:36:51,657][ceph_volume.process][INFO ] stdout AQCjTVBb/UKoJhAAhP6yGqhIgxBCPqB21Xf8Bg==
[2018-07-19 03:36:51,660][ceph_volume.process][INFO ] Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new 0e5db055-f7ce-4a3b-b5fc-af707e6fa64c
[2018-07-19 03:36:51,666][ceph_volume.devices.lvm.prepare][ERROR ] lvm prepare was unable to complete
[2018-07-19 03:36:51,666][ceph_volume.devices.lvm.prepare][INFO ] will rollback OSD ID creation
[2018-07-19 03:36:51,667][ceph_volume][ERROR ] exception caught by decorator
Traceback (most recent call last):
  File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
    return f(*a, **kw)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
    terminal.dispatch(self.mapper, subcommand_args)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
    terminal.dispatch(self.mapper, self.argv)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
    self.create(args)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
    prepare_step.safe_prepare(args)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_prepare
    self.prepare(args)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 245, in prepare
    self.osd_id = prepare_utils.create_id(osd_fsid, json.dumps(secrets), osd_id=args.osd_id)
  File "/usr/lib64/python3.5/site-packages/ceph_volume/util/prepare.py", line 72, in create_id
    show_command=True
  File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 200, in call
    stdout_stream, stderr_stream = process.communicate(stdin)
  File "/usr/lib64/python3.5/subprocess.py", line 803, in communicate
    stdout, stderr = self._communicate(input, endtime, timeout)
  File "/usr/lib64/python3.5/subprocess.py", line 1441, in _communicate
    input_view = memoryview(self._input)
TypeError: memoryview: a bytes-like object is required, not 'str'
```

#2 - 07/19/2018 10:58 AM - Alfredo Deza

This looks like Python 3 compatibility problem. Can you confirm that is the case in your system?

#3 - 07/19/2018 11:00 AM - Alfredo Deza

If it is Python 3 as I suspect, would it be possible for you to try this patch?

```
diff --git a/src/ceph-volume/ceph_volume/process.py b/src/ceph-volume/ceph_volume/process.py
index 872bd09304..908f0f578e 100644
--- a/src/ceph-volume/ceph_volume/process.py
+++ b/src/ceph-volume/ceph_volume/process.py
@@ -194,6 +194,7 @@ def call(command, **kw):
     stderr=subprocess.PIPE,
     stdin=subprocess.PIPE,
     close_fds=True,
+    encoding='utf8',
     **kw
 )
     if stdin:
```

#4 - 07/19/2018 11:53 AM - Alfredo Deza

- Status changed from New to Verified
- Assignee set to Alfredo Deza
- Priority changed from Normal to High

I verified this is happening on Python 3 instances. subprocess wants to receive stdin in bytes not 'str' so we must set the encoding, except Python 2 doesn't take that argument.

#5 - 07/19/2018 11:12 PM - Michael Jones

After applying the suggested patch and rebuilding:

```
fenrir ~ # ceph-volume lvm zap /dev/sda
--> TypeError: init() got an unexpected keyword argument 'encoding'
```

#6 - 07/20/2018 12:39 PM - Alfredo Deza

Yeah that would happen on a Python 2 environment. I am working on a patch to make it work for both. You didn't confirm if the error happened in a Python 3 server

#7 - 07/20/2018 12:40 PM - Alfredo Deza

- Subject changed from *ceph-volume fails to create OSD on raw sata hdd* to *ceph-volume fails to create OSD with Python 3*
- Status changed from Verified to In Progress

#8 - 07/20/2018 04:47 PM - Michael Jones

I use python 3.5.5, as indicated on line 2 of my original response.

The shebang at the top of the file is for python 3.5, and running the ceph-volume script directly with python 3.5 all spit out the error I indicated.

#9 - 07/20/2018 04:51 PM - Michael Jones

```
fenrir ~ # eselect python list
Available Python interpreters, in order of preference:
[1] python3.5
[2] python2.7
fenrir ~ # python3.5 /usr/sbin/ceph-volume lvm zap /dev/sda
--> TypeError: init() got an unexpected keyword argument 'encoding'
```

#10 - 07/20/2018 05:01 PM - Alfredo Deza

PR <https://github.com/ceph/ceph/pull/23141>

Feel free to try the change, tested against Python 3.5

#11 - 07/21/2018 06:48 AM - Michael Jones

Patch at <https://github.com/ceph/ceph/pull/23141/commits/728c225258ac5cd3274f2d5ef2528d87d38423fc.patch>

does not resolve the issue:

```
fenrir ~ # ceph-volume lvm zap /dev/sda
--> Zapping: /dev/sda
Running command: /sbin/cryptsetup status /dev/mapper/
stdout: /dev/mapper/ is inactive.
Running command: /sbin/wipefs --all /dev/sda
Running command: /bin/dd if=/dev/zero of=/dev/sda bs=1M count=10
--> TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

#12 - 07/21/2018 11:45 AM - Alfredo Deza

That sounds like a new (different) issue. Can you get us the log output when you ran that command **or** run it with the debug env var:

```
CEPH_VOLUME_DEBUG=1 ceph-volume lvm zap /dev/sda
```

Since you are seeing this for 'zap' and no longer for 'create', I am going to move ahead with the fix for this bug and close it once the PR merges.

#13 - 07/21/2018 12:49 PM - Alfredo Deza

Would you mind sharing what distro are you using? Perhaps it would be far easier on our end to try and get functional testing on that distro vs. attempting to find these on a case-by-case basis

#14 - 07/21/2018 05:18 PM - Michael Jones

It's the same issue

```
fenrir ~ # ceph-volume lvm create --bluestore --data /dev/sda
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new
3620d19e-a27d-4c36-88b3-2fe04567b31b
Running command: /sbin/vgcreate --force --yes ceph-76f7b0d0-a023-4dff-95a3-2a4d8e9c23c0 /dev/sda
--> Was unable to complete a new OSD, will rollback changes
--> OSD will be fully purged from the cluster, because the ID was generated
Running command: /usr/bin/ceph osd purge osd.10 --yes-i-really-mean-it
--> TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

#15 - 07/21/2018 05:19 PM - Michael Jones

The distro that I'm using is in the third line of my initial report.

#16 - 07/21/2018 05:19 PM - Michael Jones

Portage 2.3.40 (python 3.5.5-final-0, default/linux/amd64/17.0, gcc-7.3.0, glibc-2.26-r7, 4.14.52-gentoo x86_64)

```
=====  
System Settings =====  
System uname: Linux-4.14.52-gentoo-x86_64-AMD_E-350D_APU_with_Radeon-tm-_HD_Graphics-with-gentoo-2.4.1  
KiB Mem: 16134312 total, 14082204 free  
KiB Swap: 10485756 total, 10485756 free  
Timestamp of repository gentoo: Wed, 18 Jul 2018 15:24:42 +0000  
Head commit of repository gentoo: f9675b0067f57e7c11b60de9a5f8b8e3b7305286
```

Head commit of repository jonesmz-public-overlay: c83a2294f2ef3f6c1b4f8bc1086fc9eec1aec37d

Head commit of repository steam-overlay: 1237b523da636a247376b25cd4ec59c16d5b0104

```
sh bash 4.4_p12  
ld GNU ld (Gentoo 2.30 p2) 2.30.0  
distcc 3.2rc1 x86_64-pc-linux-gnu [disabled]  
app-shells/bash: 4.4_p12::gentoo  
dev-lang/perl: 5.24.3-r1::gentoo  
dev-lang/python: 2.7.14-r1::gentoo, 3.5.5::gentoo  
dev-util/cmake: 3.9.6::gentoo  
dev-util/pkgconfig: 0.29.2::gentoo  
sys-apps/baselayout: 2.4.1-r2::gentoo  
sys-apps/sandbox: 2.13::gentoo  
sys-devel/autoconf: 2.69-r4::gentoo  
sys-devel/automake: 1.11.6-r3::gentoo, 1.15.1-r2::gentoo  
sys-devel/binutils: 2.30-r2::gentoo  
sys-devel/gcc: 7.3.0-r3::gentoo  
sys-devel/gcc-config: 1.8-r1::gentoo  
sys-devel/libtool: 2.4.6-r3::gentoo  
sys-devel/make: 4.2.1::gentoo  
sys-kernel/linux-headers: 4.13::gentoo (virtual/os-headers)  
sys-libs/glibc: 2.26-r7::gentoo  
Repositories:
```

```
gentoo  
location: /usr/portage  
sync-type: git  
sync-uri: git://anongit.gentoo.org/repo/sync/gentoo.git
```


priority: -1000

jonesmz-public-overlay

location: /usr/portage-overlays/jonesmz-public-overlay

sync-type: git

sync-uri: <https://github.com/jonesmz/gentoo-overlay.git>

masters: gentoo

steam-overlay

location: /usr/portage-overlays/steam-overlay

sync-type: git

sync-uri: <https://github.com/anyc/steam-overlay.git>

masters: gentoo

priority: 50

Installed sets: @archive, @pc-base-system, @portage, @vcs

ACCEPT_KEYWORDS="amd64"

ACCEPT_LICENSE="* -@EULA"

CBUILD="x86_64-pc-linux-gnu"

CFLAGS="-O2 -pipe -march=x86-64 -mtune=generic -O2 -pipe"

CHOST="x86_64-pc-linux-gnu"

CONFIG_PROTECT="/etc /usr/share/gnupg/qualified.txt"

CONFIG_PROTECT_MASK="/etc/ca-certificates.conf /etc/dconf /etc/env.d /etc/gconf /etc/gentoo-release /etc/revdep-rebuild /etc/sandbox.d /etc/terminfo"

CXXFLAGS="-O2 -pipe -O2 -pipe -march=x86-64 -mtune=generic -O2 -pipe"

DISTDIR="/usr/portage-distfiles"

EMERGE_DEFAULT_OPTS="--jobs --keep-going --newuse --deep --backtrack=3000 --complete-graph --with-bdeps=y"

ENV_UNSET="DBUS_SESSION_BUS_ADDRESS DISPLAY PERL5LIB PERL5OPT PERLPREFIX PERL_CORE PERL_MB_OPT PERL_MM_OPT

XAUTHORITY XDG_CACHE_HOME XDG_CONFIG_HOME XDG_DATA_HOME XDG_RUNTIME_DIR"

FCFLAGS="-O2 -pipe"

FEATURES="assume-digests binpkg-logs buildpkg clean-logs compress-build-logs compressdebug config-protect-if-modified distlocks ebuild-locks

fixlafiles installsources merge-sync multilib-strict news nostrip parallel-fetch parallel-install preserve-libs protect-owned sandbox sperms split-elog

split-log strict unknown-features-warn unmerge-logs unmerge-orphans userfetch userpriv usersandbox usersync xattr"

FFLAGS="-O2 -pipe"

GENTOO_MIRRORS="http://distfiles.gentoo.org"

LANG="en_US.utf8"

LDFLAGS="-Wl,-O1 -Wl,--as-needed"

LINGUAS="en en_US"

MAKEOPTS="-j3"

PKGDIR="/usr/portage-packages"

PORTAGE_COMPRESS="xz"

PORTAGE_CONFIGROOT="/"

PORTAGE_RSYNC_OPTS="--recursive --links --safe-links --perms --times --omit-dir-times --compress --force --whole-file --delete --stats

--human-readable --timeout=180 --exclude=/distfiles --exclude=/local --exclude=/packages --exclude=/.git"

PORTAGE_TMPDIR="/var/tmp"

USE="acl amd64 avahi btfs bzip2 clang crypt cxx dbus gd gudev hardened iconv ipv6 libtirpc lm_sensors multilib ncurses nls nptl openmp pam pcre pie python readline samba seccomp ssl ssp systemd threads udev udisks unicode v4l xattr xtpax zeroconf zlib" ABI_X86="64"

ALSA_CARDS="ali5451 als4000 atiixp atiixp-modem bt87x ca0106 cmipci emu10k1x ens1370 ens1371 es1938 es1968 fm801 hda-intel intel8x0

intel8x0m maestro3 trident usb-audio via82xx via82xx-modem ymfpci" APACHE2_MODULES="authn_core authz_core socache_shmcb unixd actions

alias auth_basic authn_alias authn_anon authn_dbm authn_default authn_file authz_dbm authz_default authz_groupfile authz_host authz_owner

authz_user autoindex cache cgi cgid dav dav_fs dav_lock deflate dir disk_cache env expires ext_filter file_cache filter headers include info log_config

logio mem_cache mime mime_magic negotiation rewrite setenvif speling status unique_id userdir usertrack vhost_alias"

CALLIGRA_FEATURES="karbon plan sheets stage words" COLLECTD_PLUGINS="df interface irq load memory rrdtool swap syslog"

CPU_FLAGS_X86="mmx sse sse2 mmxext" ELIBC="glibc" GPSD_PROTOCOLS="ashtech aivdm earthmate evermore fv18 garmin garminxt

gpsclock isync itrax mtk3301 nmea ntrip navcom oceanserver oldstyle oncore rtc104v2 rtc104v3 sirf skytraq superstar2 timing tsip tripmate tnt

ublox ubx" GRUB_PLATFORMS="coreboot efi-64 emu qemu pc" INPUT_DEVICES="libinput" KERNEL="linux" L10N="en en-US"

LCD_DEVICES="bayrad cfontz cfontz633 glk hd44780 lb216 lcdm001 mtخورب ncurses text" LIBREOFFICE_EXTENSIONS="presenter-console

presenter-minimizer" OFFICE_IMPLEMENTATION="libreoffice" PHP_TARGETS="php5-6 php7-0" POSTGRES_TARGETS="postgres9_5

postgres10" PYTHON_SINGLE_TARGET="python3_5" PYTHON_TARGETS="python2_7 python3_5" QEMU_SOFTMMU_TARGETS="arm aarch64

x86_64" QEMU_USER_TARGETS="arm aarch64 x86_64" RUBY_TARGETS="ruby23" USERLAND="GNU" VIDEO_CARDS="vesa"

XTABLES_ADDONS="quota2 psd pknock lscan length2 ipv4options ipset ipp2p iface geoip fuzzy condition tee tarpit sysrq steal rawnat logmark

ipmark dhcpmac delude chaos account"

Unset: CC, CPPFLAGS, CTARGET, CXX, INSTALL_MASK, LC_ALL, PORTAGE_BINHOST, PORTAGE_BUNZIP2_COMMAND,

PORTAGE_COMPRESS_FLAGS, PORTAGE_RSYNC_EXTRA_OPTS

=====
Package Settings

sys-cluster/ceph-13.2.0:gentoo was built with the following:

USE="cephfs fuse mgr radosgw ssl systemd tcmalloc -babeltrace -dpdk -jemalloc -ldap -ltnng -mgr-frontend (-static-libs) (-system-boost) -test -xfs

-zfs" ABI_X86="(64)" CPU_FLAGS_X86="sse sse2 -sse3 -sse4_1 -sse4_2 -ssse3" PYTHON_TARGETS="python2_7 python3_5 -python3_4

-python3_6"

#17 - 07/21/2018 05:21 PM - Michael Jones

```
fenrir ~ # CEPH_VOLUME_DEBUG=1 ceph-volume lvm zap /dev/sda
--> Zapping: /dev/sda
Running command: /sbin/cryptsetup status /dev/mapper/
stdout: /dev/mapper/ is inactive.
Running command: /sbin/wipefs --all /dev/sda
Traceback (most recent call last):
File "/usr/sbin/ceph-volume", line 6, in <module>
main.Volume()
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 37, in init
self.main(self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
return f(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
terminal.dispatch(self.mapper, subcommand_args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
terminal.dispatch(self.mapper, self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/zap.py", line 169, in main
self.zap(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/zap.py", line 102, in zap
wipefs(path)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/zap.py", line 21, in wipefs
path
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 103, in stdout
return _Write(prefix=blue(' stdout: ')).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

```
fenrir ~ # CEPH_VOLUME_DEBUG=1 ceph-volume lvm create --bluestore --data /dev/sda
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new
3c3ce324-87c4-4312-a526-1bbd1ca5dc97
Running command: /sbin/vgcreate --force --yes ceph-41b5a381-f255-4bb9-95fb-7804492a31e1 /dev/sda
--> Was unable to complete a new OSD, will rollback changes
--> OSD will be fully purged from the cluster, because the ID was generated
Running command: /usr/bin/ceph osd purge osd.10 --yes-i-really-mean-it
Traceback (most recent call last):
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_prepare
self.prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 283, in prepare
block_lv = self.prepare_device(args.data, 'block', cluster_fsid, osd_fsid)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 196, in prepare_device
api.create_vg(vg_name, arg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/api/lvm.py", line 305, in create_vg
name] + list(devices)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 107, in stderr
return _Write(prefix=yellow(' stderr: ')).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

During handling of the above exception, another exception occurred:

```
Traceback (most recent call last):
File "/usr/sbin/ceph-volume", line 6, in <module>
main.Volume()
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 37, in init
self.main(self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
return f(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
terminal.dispatch(self.mapper, subcommand_args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
terminal.dispatch(self.mapper, self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
self.create(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
prepare_step.safe_prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 220, in safe_prepare
rollback_osd(args, self.osd_id)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/common.py", line 31, in rollback_osd
'--yes-i-really-mean-it'])
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 107, in stderr
return _Write(prefix=yellow(' stderr: ')).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

#18 - 07/23/2018 02:52 PM - Alfredo Deza

After the last patch you've applied I see:

Since you are using 13.2.0 it is missing this patch <https://github.com/ceph/ceph/pull/22640>

Would be super useful if you could apply that and hopefully get some log info on why that is failing (the last TypeError is still a bug though)

#19 - 07/23/2018 03:52 PM - Michael Jones

Applied patch by hand.

```
fenrir /etc/portage/patches/sys-cluster/ceph # CEPH_VOLUME_DEBUG=1 ceph-volume lvm create --bluestore --data /dev/sda
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new
d07cd461-512c-414c-9e76-c8bad03a7aea
Running command: /sbin/vgcreate --force --yes ceph-8f9720f6-a3a1-4646-a130-b70c37641a07 /dev/sda
--> Was unable to complete a new OSD, will rollback changes
--> OSD will be fully purged from the cluster, because the ID was generated
Running command: /usr/bin/ceph osd purge osd.7 --yes-i-really-mean-it
Traceback (most recent call last):
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_prepare
self.prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 283, in prepare
block_lv = self.prepare_device(args.data, 'block', cluster_fsid, osd_fsid)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 196, in prepare_device
api.create_vg(vg_name, arg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/api/lvm.py", line 305, in create_vg
name] + list(devices)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 107, in stderr
return _Write(prefix=yellow(" stderr: ")).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith("\n"):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

During handling of the above exception, another exception occurred:

```
Traceback (most recent call last):
File "/usr/sbin/ceph-volume", line 6, in <module>
main.Volume()
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 37, in init
self.main(self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
return f(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
terminal.dispatch(self.mapper, subcommand_args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
terminal.dispatch(self.mapper, self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
self.create(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
prepare_step.safe_prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 220, in safe_prepare
```

```
rollback_osd(args, self.osd_id)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/common.py", line 31, in rollback_osd
'--yes-i-really-mean-it'])
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 107, in stderr
return _Write(prefix=yellow(' stderr: ')).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

#20 - 07/23/2018 03:55 PM - Michael Jones

It's getting a little farther though.

Managing to create lvm stuff, where as it wasn't before, and managing to purge the OSD that it failed to create.

Still needs to clean up it's LVM stuff if it fails though. I doubt many admins want to go around hunting for volume groups for failed OSDs.

```
fenrir /etc/portage/patches/sys-cluster/ceph # lvdisk ; vgdisplay ; pvdisplay
fenrir /etc/portage/patches/sys-cluster/ceph # CEPH_VOLUME_DEBUG=1 ceph-volume lvm create --bluestore --data /dev/sda
Running command: /usr/bin/ceph-authtool --gen-print-key
Running command: /usr/bin/ceph --cluster ceph --name client.bootstrap-osd --keyring /var/lib/ceph/bootstrap-osd/ceph.keyring -i - osd new
16bcab21-f380-48f7-8ce4-a3d21486bf78
Running command: /sbin/vgcreate --force --yes ceph-49f6041d-5b02-4ad0-b562-17d5a0bb50d8 /dev/sda
--> Was unable to complete a new OSD, will rollback changes
--> OSD will be fully purged from the cluster, because the ID was generated
Running command: /usr/bin/ceph osd purge osd.7 --yes-i-really-mean-it
Traceback (most recent call last):
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 216, in safe_prepare
self.prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 283, in prepare
block_lv = self.prepare_device(args.data, 'block', cluster_fsid, osd_fsid)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 196, in prepare_device
api.create_vg(vg_name, arg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/api/lvm.py", line 305, in create_vg
name] + list(devices)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 103, in stdout
return _Write(prefix=blue(' stdout: ')).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
```

During handling of the above exception, another exception occurred:

```
Traceback (most recent call last):
File "/usr/sbin/ceph-volume", line 6, in <module>
main.Volume()
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 37, in init
self.main(self.argv)
```

```

File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 59, in newfunc
return f(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/main.py", line 153, in main
terminal.dispatch(self.mapper, subcommand_args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/main.py", line 38, in main
terminal.dispatch(self.mapper, self.argv)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 182, in dispatch
instance.main()
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 69, in main
self.create(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/decorators.py", line 16, in is_root
return func(*a, **kw)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/create.py", line 26, in create
prepare_step.safe_prepare(args)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/prepare.py", line 220, in safe_prepare
rollback_osd(args, self.osd_id)
File "/usr/lib64/python3.5/site-packages/ceph_volume/devices/lvm/common.py", line 31, in rollback_osd
'--yes-i-really-mean-it!'])
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 133, in run
log_descriptors(reads, process, terminal_logging)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 55, in log_descriptors
log_output(descriptor_name, read(descriptor, 1024), terminal_logging, True)
File "/usr/lib64/python3.5/site-packages/ceph_volume/process.py", line 31, in log_output
getattr(terminal, descriptor)(message)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 107, in stderr
return _Write(prefix=yellow(" stderr: ")).raw(msg)
File "/usr/lib64/python3.5/site-packages/ceph_volume/terminal.py", line 92, in raw
if not string.endswith('\n'):
TypeError: endswith first arg must be bytes or a tuple of bytes, not str
fenrir /etc/portage/patches/sys-cluster/ceph # lvdisplay ; vgdisplay ; pvdisplay
--- Volume group ---
VG Name          ceph-49f6041d-5b02-4ad0-b562-17d5a0bb50d8
System ID
Format           lvm2
Metadata Areas   1
Metadata Sequence No 1
VG Access        read/write
VG Status        resizable
MAX LV           0
Cur LV          0
Open LV          0
Max PV           0
Cur PV          1
Act PV           1
VG Size          931.51 GiB
PE Size          4.00 MiB
Total PE         238467
Alloc PE / Size  0 / 0
Free PE / Size   238467 / 931.51 GiB
VG UUID          y2kAlm-41rF-be9F-RtVu-L1uP-ZLIP-CZSkTc

--- Physical volume ---
PV Name          /dev/sda
VG Name          ceph-49f6041d-5b02-4ad0-b562-17d5a0bb50d8
PV Size          931.51 GiB / not usable 1.71 MiB
Allocatable     yes
PE Size          4.00 MiB
Total PE         238467
Free PE         238467
Allocated PE     0
PV UUID          wzcV3C-hWuT-J196-MnJi-z8Ws-UNJI-zvD0j0

```

```
fenrir /etc/portage/patches/sys-cluster/ceph #
```

#21 - 07/25/2018 12:41 PM - Alfredo Deza

Was able to narrow down this one, and I have a test and a fix.

The fix is here: <https://github.com/ceph/ceph/pull/23141/commits/2133087e9016684c55a8770dd2d8bef5557bd2fd>

I believe that should get us to a working OSD. Let me know

#22 - 07/25/2018 04:29 PM - Michael Jones

- File *ceph-osd.7.log* added

- File *ceph-volume.log* added

Patch allows ceph-volume lvm create /dev/sda to work, apparently successfully.

The command returns success, and the log messages indicate that it thinks everything worked properly.

The OSD process that gets spawned after ceph-volume finishes core-dumps immediately, and now my harddrive is unresponsive to any program except direct writes like dd, even after a reboot (which apparently requires me to press the power button. Software reboot hangs for apparently forever).

Reproduced the new problem on two different drives.

I'm very frustrated that the ceph project deprecated ceph-disk, which was working without issues for me, in favor of a new tool that

- 1) Requires new kernel modules, and new userland utilities, with a whole new set of headaches for management.
- 2) Is obviously not fully baked yet.

```
fenrir /var/log/ceph # ceph-volume lvm zap /dev/sdb
--> Zapping: /dev/sdb
--> Unmounting /var/lib/ceph/osd/ceph-7
Running command: /bin/umount -v /var/lib/ceph/osd/ceph-7
stderr: umount:
stderr: /var/lib/ceph/osd/ceph-7 unmounted
stderr:
Running command: /sbin/wipefs --all /dev/sdb
stderr: wipefs: error: /dev/sdb: probing initialization failed: Device or resource busy
--> RuntimeError: command returned non-zero exit status: 1
```

#23 - 07/25/2018 04:31 PM - Michael Jones

- File *ceph.audit.log* added

- File *ceph.log* added

Redmine refused to upload all my log files, apparently.

#24 - 07/25/2018 04:39 PM - Alfredo Deza

Since these patches are now working for the problem you reported I will close this once the functional tests are passing on the PR.

If you would still prefer ceph-disk, you can do so up until Mimic. It is no longer present for the next (nautilus) version of Ceph.

#25 - 07/25/2018 05:01 PM - Michael Jones

You can close this issue.

Please open a new one for the segmentation fault (or tell me that I need to, so I know to do it.)

If you would still prefer ceph-disk, you can do so up until Mimic. It is no longer present for the next (nautilus) version of Ceph.

That's... the problem.

Don't mark something as deprecated until the replacement **actually** works.

#26 - 07/25/2018 05:34 PM - Alfredo Deza

Sorry that you feel that way about using the new tooling. We do test ceph-volume and it does work in our environments without problems, we acknowledge that Python 3 compatibility is an ongoing process and we are trying to get more functional testing around that to prevent issues like the ones you found.

ceph-disk has a long tail of problems that were recently documented here: <http://docs.ceph.com/docs/master/ceph-volume/intro/#ceph-disk-replaced>

ceph-volume does have some issues but it works well and does not display any of the impossible-to-fix issues that ceph-disk had. I am happy to accept more issues as you find them, they are super useful! I appreciate your patience while we get these fixes sorted out.

For core dumps I would suggest opening a ticket against the rados project: <http://tracker.ceph.com/projects/rados/issues/new>

#27 - 07/25/2018 06:10 PM - Michael Jones

Don't deprecate a feature until the replacement has been deployed and in use by the community for a full release.

You simultaneously deprecated a working feature (even if there are problems with it), and introduced the new tool at the same time.

You should have announced that ceph-disk was deprecated on the nautilus release, a full release after ceph-volume became available. Or announced it's deprecation for the mimic release (as you did), with it's removal the release **after** the nautilus release, which is not what you're doing.

Unless both tools live side by side for a full release (without one of them being deprecated), you're just begging for something to go wrong.

But whatever, this isn't the place to discuss this kind of thing, and I doubt you personally care about my opinion on the matter.

Thank you for working with me to repair the ceph-volume issues.

ceph-volume does have some issues but it works well and does not display any of the impossible-to-fix issues that ceph-disk had.

I mean, I disagree that it works well... but I have a different perspective on the matter than you do.

I'd be very interested in a blog post describing all of these issues that ceph-disk has, and what steps the Ceph developers took to get udev to fix it's problems. Udev is at the heart of the majority of linux machines out there, and I don't disagree with you that it's... not fun, so I'm curious what led the Ceph project to abandon it in favor of LVM (which is **not** as widely available as udev) and whether the udev developers were simply unwilling, or unable, to address any of the problems.

Created a new issue for rados.

<http://tracker.ceph.com/issues/25106>

#28 - 07/25/2018 07:02 PM - Alfredo Deza

- Status changed from In Progress to Closed

#29 - 07/25/2018 07:09 PM - Alfredo Deza

Don't deprecate a feature until the replacement has been deployed and in use by the community for a full release.

ceph-disk was initially deprecated in Luminous, we backtracked and it is now deprecated in Mimic. Luminous users will not see a deprecation warning. The deprecation banner is in Mimic only.

You simultaneously deprecated a working feature (even if there are problems with it), and introduced the new tool at the same time.

This was discussed in the ceph-devel mailing list, a consensus was reached and the deprecation was set. For someone installing **today**, it will not see a deprecation warning in Luminous.

You should have announced that ceph-disk was deprecated on the nautilus release, a full release after ceph-volume became available. Or announced it's deprecation for the mimic release (as you did), with it's removal the release **after** the nautilus release, which is not what you're doing.

These opinions would've been valuable for sure, when the discussion happened. The community agreed that waiting another full release was fine. We can't really revert this to accommodate your opinion.

Unless both tools live side by side for a full release (without one of them being deprecated), you're just begging for something to go wrong.

Both tools are living side by side for a full release (Luminous) without one of them being deprecated.

But whatever, this isn't the place to discuss this kind of thing, and I doubt you personally care about my opinion on the matter.

I would welcome your ideas and discuss them further in the ceph-devel mailing list (or the ceph-users mailing list) if you prefer.

#30 - 07/27/2018 12:30 PM - Alfredo Deza

- Status changed from Closed to In Progress

Reopening because we saw some failures with strings not fully decoded in stdin:

```
Traceback (most recent call last):
  File "/usr/sbin/ceph-volume", line 6, in <module>
    main.Volume()
  File "/usr/lib/python2.7/dist-packages/ceph_volume/main.py", line 37, in __init__
    self.main(self.argv)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/decorators.py", line 59, in newfunc
    return f(*a, **kw)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/main.py", line 153, in main
    terminal.dispatch(self.mapper, subcommand_args)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
  File "/usr/lib/python2.7/dist-packages/ceph_volume/devices/simple/main.py", line 33, in main
    terminal.dispatch(self.mapper, self.argv)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/terminal.py", line 182, in dispatch
    instance.main()
  File "/usr/lib/python2.7/dist-packages/ceph_volume/devices/simple/activate.py", line 235, in main
    self.activate(args)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/decorators.py", line 16, in is_root
    return func(*a, **kw)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/devices/simple/activate.py", line 124, in activate
    data_device = self.get_device(data_uuid)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/devices/simple/activate.py", line 80, in get_device
    encryption_utils.luks_open(self.dmccrypt_secret, device, uuid)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/util/encryption.py", line 90, in luks_open
    process.call(command, stdin=key, terminal_verbose=True, show_command=True)
  File "/usr/lib/python2.7/dist-packages/ceph_volume/process.py", line 205, in call
    stdin.encode(encoding='utf-8', errors='ignore')
UnicodeDecodeError: 'ascii' codec can't decode byte 0xa3 in position 3: ordinal not in range(128)
```

Previous fixes are in master only, backports started failing.

#31 - 07/27/2018 08:06 PM - Alfredo Deza

- Status changed from In Progress to Resolved

Master PR <https://github.com/ceph/ceph/pull/23289>

Files

ceph-osd.7.log	635 KB	07/25/2018	Michael Jones
ceph-volume.log	30.6 KB	07/25/2018	Michael Jones
ceph.audit.log	9.12 KB	07/25/2018	Michael Jones
ceph.log	8.21 KB	07/25/2018	Michael Jones