

GlusterFS bez cap_sys_admin

Ve výchozím stavu používá GlusterFS xattrs, který není dostupný pod kontejnery, bez cap_sys_admin capabilities. Tady je stručný návod k sestavení GlusterFS pod Debianem tak, aby používal uživatelské namespaces:

```
# Nejprve nainstalujeme potřebné závislosti:
apt-get install dpkg-dev fuse
apt-get build-dep glusterfs

# Stáhneme zdrojové kódy glusterfs
apt-get source glusterfs

# Pokud máte problémy s některými GPG klíči, můžete si je přidat:
gpg --keyserver keyserver.ubuntu.com --recv-keys 90CBD8E4

# Upravte zdrojové kódy prostým "najdi a změň"
cd glusterfs-3.5.2/
find . -regex '.*[/^]*\.[ch]' -exec sed -i
's/&quot;trusted\./&quot;user./g' {} \;

# Sestavte balíčky
dpkg-source --commit
dpkg-buildpackage
cd ..

# Nainstalujte je
dpkg -i glusterfs-*.deb
```

Funkční svazek:

```
gluster> volume info

Volume Name: shared
Type: Replicate
Volume ID: 7253c6d6-aec0-4804-9bb5-c4b0150d69bc
Status: Started
Number of Bricks: 1 x 2 = 2
Transport-type: tcp
Bricks:
Brick1: 172.16.0.62:/gluster
Brick2: 172.16.0.1:/gluster
```

Výkonnostní testy na vpsFree.cz:

```
iozone -w -c -e -i 0 -+n -C -r 64k -s 1g -t 8 -F
/mnt/tmp/f{0,1,2,3,4,5,6,7,8}.ioz
      Iozone: Performance Test of File I/O
      Version $Revision: 3.430 $
```

Compiled for 64 bit mode.
Build: linux-AMD64

Contributors:William Norcott, Don Capps, Isom Crawford, Kirby Collins

Al Slater, Scott Rhine, Mike Wisner, Ken Goss
Steve Landherr, Brad Smith, Mark Kelly, Dr. Alain CYR,
Randy Dunlap, Mark Montague, Dan Million, Gavin

Brebner,

Jean-Marc Zucconi, Jeff Blomberg, Benny Halevy, Dave

Boone,

Erik Habbinga, Kris Strecker, Walter Wong, Joshua Root,
Fabrice Bacchella, Zhenghua Xue, Qin Li, Darren Sawyer,
Vangel Bojaxhi, Ben England, Vikentsi Lapa.

Run began: Fri Jul 17 11:48:28 2015

Setting no_unlink
Include close in write timing
Include fsync in write timing
Norettest option selected
Record Size 64 kB

File size set to 1048576 kB

Command line used: ./iozone -w -c -e -i 0 -+n -C -r 64k -s lg -t 8 -
F /mnt/tmp/f0.ioz /mnt/tmp/f1.ioz /mnt/tmp/f2.ioz /mnt/tmp/f3.ioz
/mnt/tmp/f4.ioz /mnt/tmp/f5.ioz /mnt/tmp/f6.ioz /mnt/tmp/f7.ioz
/mnt/tmp/f8.ioz

Output is in kBytes/sec
Time Resolution = 0.000001 seconds.
Processor cache size set to 1024 kBytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Throughput test with 8 processes
Each process writes a 1048576 kByte file in 64 kByte records

Children see throughput for 8 initial writers	=	53489.47 kB/sec
Parent sees throughput for 8 initial writers	=	53204.68 kB/sec
Min throughput per process	=	6546.01 kB/sec
Max throughput per process	=	6947.46 kB/sec
Avg throughput per process	=	6686.18 kB/sec
Min xfer	=	1009216.00 kB
Child[0] xfer count = 1024000.00 kB, Throughput	=	6642.03 kB/sec
Child[1] xfer count = 1045696.00 kB, Throughput	=	6782.18 kB/sec
Child[2] xfer count = 1034880.00 kB, Throughput	=	6712.34 kB/sec
Child[3] xfer count = 1022080.00 kB, Throughput	=	6629.23 kB/sec
Child[4] xfer count = 1011328.00 kB, Throughput	=	6559.43 kB/sec
Child[5] xfer count = 1009216.00 kB, Throughput	=	6546.01 kB/sec
Child[6] xfer count = 1048576.00 kB, Throughput	=	6947.46 kB/sec
Child[7] xfer count = 1028352.00 kB, Throughput	=	6670.79 kB/sec

iozone test complete.

Svazek GlusterFS:

```
iozone -w -c -e -i 0 --n -C -r 64k -s 1g -t 8 -F
/mnt/shared/f{0,1,2,3,4,5,6,7,8}.ioz
  Iozone: Performance Test of File I/O
  Version $Revision: 3.430 $
  Compiled for 64 bit mode.
  Build: linux-AMD64

Contributors:William Norcott, Don Capps, Isom Crawford, Kirby
Collins
              Al Slater, Scott Rhine, Mike Wisner, Ken Goss
              Steve Landherr, Brad Smith, Mark Kelly, Dr. Alain CYR,
              Randy Dunlap, Mark Montague, Dan Million, Gavin
Brebner,
              Jean-Marc Zucconi, Jeff Blomberg, Benny Halevy, Dave
Boone,
              Erik Habbinga, Kris Strecker, Walter Wong, Joshua Root,
              Fabrice Bacchella, Zhenghua Xue, Qin Li, Darren Sawyer,
              Vangel Bojaxhi, Ben England, Vikentsi Lapa.

Run began: Fri Jul 17 12:19:17 2015

Setting no_unlink
Include close in write timing
Include fsync in write timing
No retest option selected
Record Size 64 kB
File size set to 1048576 kB
Command line used: ./iozone -w -c -e -i 0 --n -C -r 64k -s 1g -t 8 -
F /mnt/shared/f0.ioz /mnt/shared/f1.ioz /mnt/shared/f2.ioz
/mnt/shared/f3.ioz /mnt/shared/f4.ioz /mnt/shared/f5.ioz /mnt/shared/f6.ioz
/mnt/shared/f7.ioz /mnt/shared/f8.ioz
  Output is in kBytes/sec
  Time Resolution = 0.000001 seconds.
  Processor cache size set to 1024 kBytes.
  Processor cache line size set to 32 bytes.
  File stride size set to 17 * record size.
  Throughput test with 8 processes
  Each process writes a 1048576 kByte file in 64 kByte records

Children see throughput for 8 initial writers = 23963.82 kB/sec
Parent sees throughput for 8 initial writers = 14037.13 kB/sec
Min throughput per process = 1783.32 kB/sec
Max throughput per process = 5134.14 kB/sec
Avg throughput per process = 2995.48 kB/sec
Min xfer = 364224.00 kB
Child[0] xfer count = 528192.00 kB, Throughput = 2584.27 kB/sec
Child[1] xfer count = 525248.00 kB, Throughput = 2571.77 kB/sec
Child[2] xfer count = 364224.00 kB, Throughput = 1783.32 kB/sec
Child[3] xfer count = 1048576.00 kB, Throughput = 5134.14 kB/sec
```

```

Child[4] xfer count = 526400.00 kB, Throughput = 2577.41 kB/sec
Child[5] xfer count = 1003200.00 kB, Throughput = 4907.09 kB/sec
Child[6] xfer count = 526144.00 kB, Throughput = 2576.12 kB/sec
Child[7] xfer count = 373696.00 kB, Throughput = 1829.70 kB/sec

```

iozone test complete.

Test na ZFS s xattr=sa:

```

iozone -w -c -e -i 0 -+n -C -r 64k -s lg -t 8 -F
/mnt/test/f{0,1,2,3,4,5,6,7,8}.ioz
  Iozone: Performance Test of File I/O
    Version $Revision: 3.430 $
    Compiled for 64 bit mode.
    Build: linux-AMD64

Contributors:William Norcott, Don Capps, Isom Crawford, Kirby
Collins
                Al Slater, Scott Rhine, Mike Wisner, Ken Goss
                Steve Landherr, Brad Smith, Mark Kelly, Dr. Alain CYR,
                Randy Dunlap, Mark Montague, Dan Million, Gavin
Brebner,
                Jean-Marc Zucconi, Jeff Blomberg, Benny Halevy, Dave
Boone,
                Erik Habbinga, Kris Strecker, Walter Wong, Joshua Root,
                Fabrice Bacchella, Zhenghua Xue, Qin Li, Darren Sawyer,
                Vangel Bojaxhi, Ben England, Vikentsi Lapa.

Run began: Wed Jul 22 08:09:08 2015

Setting no_unlink
Include close in write timing
Include fsync in write timing
No retest option selected
Record Size 64 kB
File size set to 1048576 kB
Command line used: ./iozone -w -c -e -i 0 -+n -C -r 64k -s lg -t 8 -
F /mnt/test/f0.ioz /mnt/test/f1.ioz /mnt/test/f2.ioz /mnt/test/f3.ioz
/mnt/test/f4.ioz /mnt/test/f5.ioz /mnt/test/f6.ioz /mnt/test/f7.ioz
/mnt/test/f8.ioz
Output is in kBytes/sec
Time Resolution = 0.000001 seconds.
Processor cache size set to 1024 kBytes.
Processor cache line size set to 32 bytes.
File stride size set to 17 * record size.
Throughput test with 8 processes
Each process writes a 1048576 kByte file in 64 kByte records

Children see throughput for 8 initial writers = 26844.98 kB/sec
Parent sees throughput for 8 initial writers = 26766.06 kB/sec
Min throughput per process = 3354.94 kB/sec

```

```
Max throughput per process           = 3356.57 kB/sec
Avg throughput per process           = 3355.62 kB/sec
Min xfer                             = 1048064.00 kB
Child[0] xfer count = 1048256.00 kB, Throughput = 3355.46 kB/sec
Child[1] xfer count = 1048256.00 kB, Throughput = 3355.55 kB/sec
Child[2] xfer count = 1048384.00 kB, Throughput = 3355.87 kB/sec
Child[3] xfer count = 1048064.00 kB, Throughput = 3354.94 kB/sec
Child[4] xfer count = 1048576.00 kB, Throughput = 3356.57 kB/sec
Child[5] xfer count = 1048256.00 kB, Throughput = 3355.47 kB/sec
Child[6] xfer count = 1048256.00 kB, Throughput = 3355.65 kB/sec
Child[7] xfer count = 1048256.00 kB, Throughput = 3355.47 kB/sec
```

iozone test complete.

From:

<https://kb.vpsfree.cz/> - **Znalostní Báze**

Permanent link:

<https://kb.vpsfree.cz/navody/server/glusterfs>

Last update: **2015/07/30 21:22**