

Porting GlusterFS to FreeBSD

Mike Ma (mikema@)

Mentored by Sean Bruno (sbruno@)

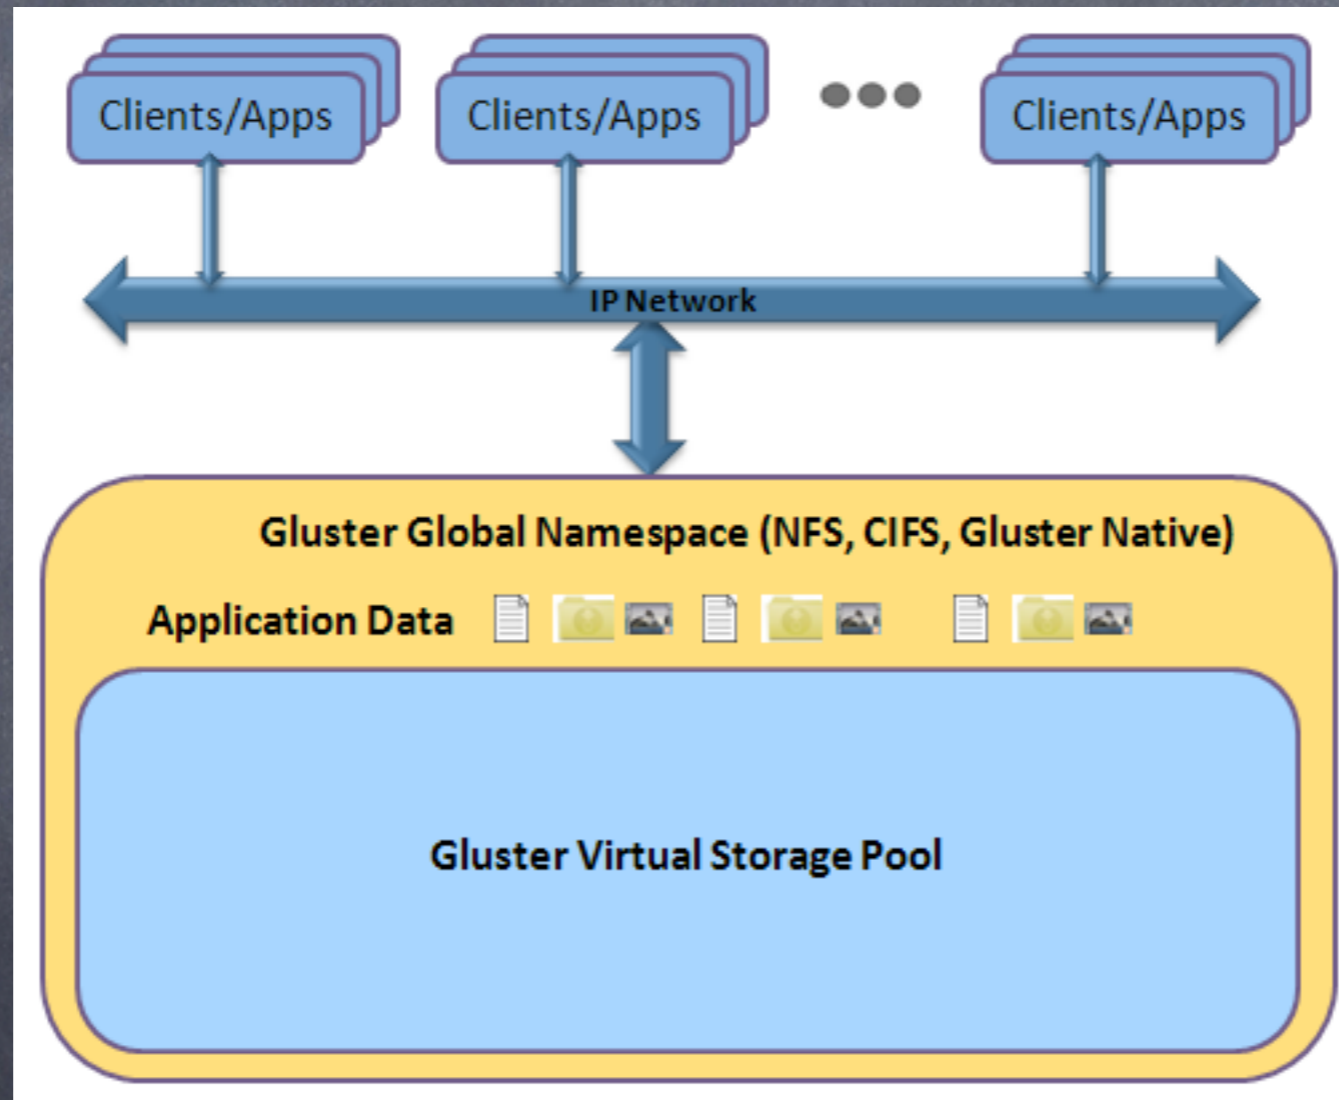
Outline

- Background
- Setup Steps
- Experimental Result
- Known Issues
- Conclusion

Motivation

- GlusterFS
 - open source, distributed file system
 - cluster machines into storage cloud
 - userspace filesystem
 - written for Linux initially, later ported to NetBSD
 - FUSE is now available in FreeBSD 10

Architecture



Outline

- Background
- Setup Steps
- Experimental Result
- Know Issues
- Conclusion

Setup Steps

- Download Source code:

- `git clone https://github.com/cosql/glusterfs`

- Build and install

- `./autogen.sh && ./configure && make && make install`

- Start **glusterd** first

- `glusterd`

Setup steps

• Initialize disk partition

- `newfs /dev/da2s1a`
- `mount /dev/da2s1a /mnt`

• Create a gluster volume

- `gluster volume create gv0 replica 2 $server_ip:/mnt/exp1
$server_ip:/mnt/exp2`
- `gluster volume gv0 start`

• Mount

- `glusterfs --log-file=log --log-level=INFO --volfile-id=gv0 --
volfile-server=$server_ip <mount_point>`

Outline

- Background
- Setup Steps
- **Experimental Result**
- Known Issues
- Conclusion

Experiment Settings

- VirtualBox

- 8 core cpu
- 8GB RAM
- 40GB VDI, ufs and zfs

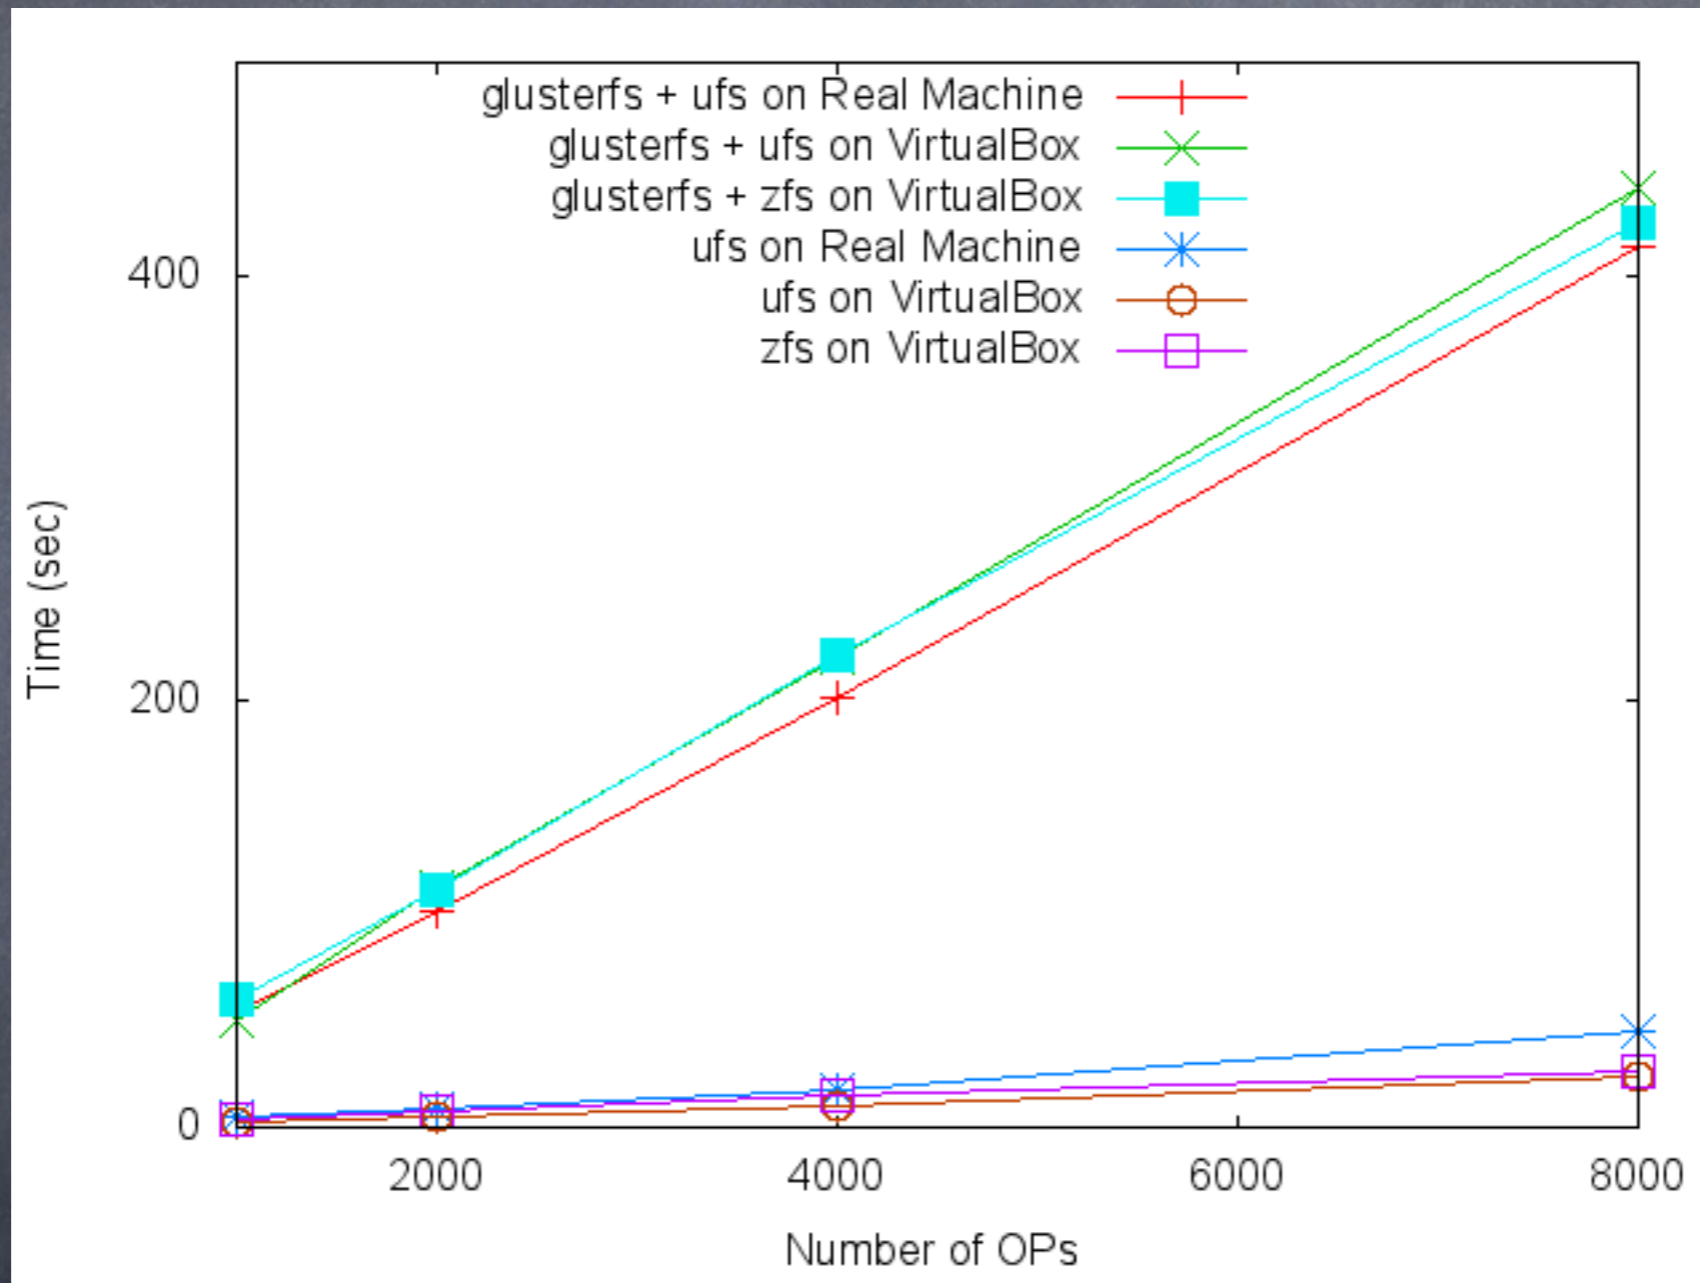
- Real Machine

- Intel(R) Core(TM)2 Duo CPU E6850 @3.00GHz
- 2GB RAM
- 2GB usb disk, formatted as ufs

- FreeBSD-10-Current

- 2 replicas

FSX



• Run with $-N \{1000, 2000, 4000, 8000\}$

Blogbench

Gluster		Underlying ufs		Gluster		Underlying zfs	
Rd	Wr	Rd	Wr	Rd	Wr	Rd	Wr
2	5490	30	66577	1	159	22	13180
1	121	28	54191	2	174	17	7884
2	1697	29	48407	1	148	10	7125
1	128	29	45987	1	285	9	6262
1	230	26	46340	2	245	10	5166

👁 On VirtualBox

Outline

- Background
- Setup Steps
- Experimental Result
- **Known Issues**
- Conclusion

Known Issues

- "mount -t fuse.glusterfs" not working
- tests under tests/ directory not ported
- libgfapi not supported
- build with clang could fail

Outline

- Background
- Setup Steps
- Experimental Result
- Known Issues
- **Conclusion**

Conclusion

- Pre-alpha state
- Need more tests on physical machines
- Plan to submit the changes to gluster.org after GSoC is over

The End

🌀 Thank You!