

Faster WAN Volume Operations with DPF

Andrew Deason

June 2019

OpenAFS Workshop 2019

Background

- High latency → slow Rx → slow AFS
 - $32window * 1389packetsize / 10msRTT = 4.24MiB/s$
 - $255window * 1389packetsize / 10msRTT = 33.78MiB/s$

- Previous talks:
 - 2012 EAKC “OpenAFS Out-of-Band TCP”
 - 2014 EAKC “OOB Update: Pipelines”

Disposable Protocol Framework (DPF)

- General protocol negotiation
- Rapid development
- TCP, on-the-fly compression (lz4), TLS
- Deployed in production (volserver 1.6)

DPF CLI

```
$ vos release/move/copy -dpf
```

```
$ vos release -no-dpf
```

```
$ vos release -dpf -s2s-dpf-pline $pipeline
```

```
$ vos dump -dpf -dpf-pline $pipeline
```

```
$ vos restore -dpf -dpf-pline $pipeline
```

```
davolserver -dpf -no-dpf -s2s-dpf default-on
```

Pipeline Spec

```
$ eof=net.sinenomine.eof32.duplex
$ fcrypt=net.sinenomine.crypt.fcrypt.clear.client
$ tcp=net.sinenomine.tcp.client
$ vos release -s2s-dpf-pline "$eof:$fcrypt:$tcp"

$ lz4=net.sinenomine.zip.lz4simple.send
$ rx=net.sinenomine.rx.sendrecv
$ vos release -s2s-dpf-pline "$lz4:$rx"
```

Release Progress

```
$ vos status $server  
[...]  
transaction: 310237 created: Wed Jan 30 17:05:24 2019  
lastActiveTime: Wed Jan 30 17:05:24 2019  
attachFlags: offline  
volume: 538137907 partition: /vicepa procedure: Restore  
packetRead: 2 lastReceiveTime: Wed Jan 30 17:05:24 2019  
packetSend: 1 lastSendTime: Wed Jan 30 17:05:24 2019  
█
```

Example Performance Impact

- 1G volume, WAN (across countries)
- Plain Rx:
 - ~20 minutes 16 sec
 - ~0.84 MiB/s
- DPF (tcp, lz4)
 - ~1 minute 12 sec
 - ~13 MiB/s

Specifying Options

- Pipeline spec is cumbersome
- Aliases?
 - `-dpf-pline eof:lz4:ssl:tcp`
- Higher-level flags
 - `-encrypt -compress -tcp`

- Optimization
- Client/fileserver
- Standards?
- 1.9?

Commits

Coming Soon!

Slides

<http://dson.org/talks>

?