

OU XROOTD SITE REPORT

HORST SEVERINI
XROOTD WORKSHOP 2019
JUNE 2019

Outline

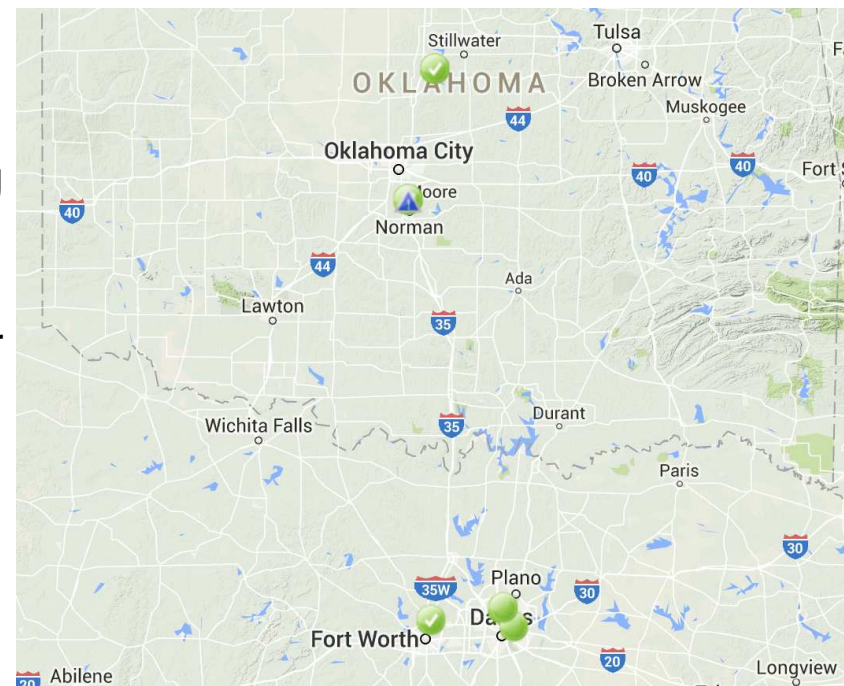
- Introduction
- Computing and Storage Hardware
- Network
- XRootD Configuration
- Third Party Copy Testing

US ATLAS SWT2 Center

- University of Oklahoma
 - Oklahoma Center for High Energy Physics (OCHEP)
 - OU Supercomputing Center for Education and Research (OSCER)

- University of Texas Arlington
 - Chemistry and Physics Building (CPB)
 - Arlington Regional Data Center (ARDC) in Fort Worth

- Langston University
 - Lucille



OU_OSCER_ATLAS Tier 2 Hardware

- 52 Nodes (2400 Slots) – 2 GB RAM per Slot
- 10 Support Nodes (1 GK, 1 DTN/Proxy, 1 redirector, 7 storage)
- 700 TB of usable xrootd storage (7 T630s with 16 8 TB drives, RAID6, xfs)
- SALT (CentOS 7.3), SLURM 17.11, OSG 3.4, XRootD 4.9.1
- Part of generic OSCER HPC cluster
- Rest of OSCER Schooner Hardware
 - 600 Nodes (about 12k Slots) – 2-4 GB RAM per Slot
 - Opportunistically available for ATLAS production



Network

- OU connected at 100 Gbps to I2 and ESnet via OneNet
- OSCER connected at dual 10 Gbps to OneNet
- Internally, cluster already ipv6 ready; working on external connectivity
- Working on 40 Gbps DMZ for OSCER cluster
 - Migrate DTN from central OU network to DMZ
 - Put that subnet onto LHCONE
 - perfSonar node already on that subnet



XRootD Configuration

- Proxy config file, `xroot-se.config`

```
all.adminpath /var/spool/xrootd
all.pidpath /var/run/xrootd
all.export /xrd
pss.origin dms.oscer.ou.edu:1094
# enable this to run as a proxy server
ofs.osslib libXrdPss.so
# enable Third Party Copy
ofs.tpc oids fcreds ?gsi =X509_USER_PROXY autorm pgm \
/etc/xrootd/xrdcp-tpc.sh
# ENABLE_SECURITY_BEGIN
xrootd.seclib /usr/lib64/libXrdSec.so
sec.protparm gsi -vomsfun:libXrdSecgsiVOMS.so -vomsfunparms:\
certfmt=pem|vos=atlas,dteam|grps=/atlas,/dteam|grpopt=10
sec.protocol gsi -dlgpxy:1 -exppxy:=creds -ca:1 -crl:3
acc.authdb /etc/xrootd/auth_file
acc.authrefresh 60
ofs.authorize
```



```
# ENABLE_SECURITY_END
xrootd.chksum Adler32 /etc/xrootd/xrdadler32.sh
```

- **auth_file**

```
g /atlas /xrd a
g /dteam /xrd/dteam a
```

- **xrdadler32.sh**

```
#!/bin/sh
/usr/bin/xrdadler32 root://$XRDXROOTD_PROXY/$1 | awk '{print $1}'
```

- **New auth_file; Doesn't currently work because it breaks OCHEP Stage-Outs**

```
= usatlas1 o: atlas g: /atlas r: production
x usatlas1 /xrd a
= usatlas2 o: atlas g: /atlas r: pilot
x usatlas2 /xrd a
g /atlas /xrd rl
g /dteam /xrd/dteam a
```



Third Party Copy

- `xrdcp-tpc.sh`

```
#!/bin/sh
src=$1
dst=$2
DSTHOST=dms.oscer.ou.edu:1094
echo $dst | grep -q -e ^\xrd\atlas
if [ $? -eq 0 ]; then
    tgtsp=`echo $dst | awk -F\ / '{print toupper($3)}'`
    echo $tgtsp | grep -q -e ^ATLAS.*DISK$
    [ $? -eq 0 ] || tgtsp="public"

    /usr/bin/xrdcp --server -f $src root://$DSTHOST/${dst}?\
oss.cgroup=$tgtsp
else
    /usr/bin/xrdcp --server -f $src root://$DSTHOST/${dst}
fi
```

TPC Testing

- Local testing:

```
[hs@osgitb1 ~]$ xrdcp -d 1 -f --tpc only root://se1.oscer.ou.edu:1094\
//xrd/srm-test/tpctest/data11_7TeV.00191217.physics_Muons.merge.AOD.f413
_m1019._1b0158._0001.1 root://se1.oscer.ou.edu:1094//xrd/atlasscratchdis
[2019-06-06 11:56:36.440530 -0500][Info  ][Utility  ]
We are NOT using delegation
[213.7MB/3.889GB] [ 5%] [==>                               ] [35.6
[960MB/3.889GB] [ 24%] [=====>                           ] [120MB/
[1.625GB/3.889GB] [ 41%] [=====                           ] [166.
[2.165GB/3.889GB] [ 55%] [=====                           ] [184.
[2.875GB/3.889GB] [ 73%] [=====>                           ] [210.
[3.618GB/3.889GB] [ 93%] [=====>                           ] [231.
[3.75GB/3.889GB] [ 96%] [=====>                           ] [213.3
[3.889GB/3.889GB] [100%] [=====] [209.6MB/s]
```

- Does spacetoken/group assignment correctly



TPC Testing continued

- Remote testing:

```
[hs@osgitb1 ~]$ xrdcp -d 1 -f --tpc delegate only root://griddev04.\
slac.stanford.edu:2094//xrootd/atlas/tpctest/file12.junk \
root://se1.oscer.ou.edu:1094//xrd/dteam/file12.junk
[976.6kB/976.6kB] [100%] [=====] [488.
[976.6kB/976.6kB] [100%] [=====] [488.3kB/s]
[hs@osgitb1 ~]$
[hs@osgitb1 ~]$ xrdcp -d 1 -f --tpc delegate only root://se1.oscer.\
ou.edu:1094//xrd/dteam/file12.junk root://griddev04.slac.\
stanford.edu:2094//xrootd/atlas/tpctest/junk.horst
[976.6kB/976.6kB] [100%] [=====] [325.
[976.6kB/976.6kB] [100%] [=====] [325.5kB/s]
```

- Does delegation correctly



Summary and Conclusions

- OU XRootD storage very stable and performant
- TPC working well, recently put into production
- Able to transfer 800+ MB/s, which is probably close to current available OCSEER core network limit
- Need to fix auth_file, not sure if possible before OCHEP retirement
- Next steps: pre-assigned space groups, enable HTTP for XRootD
- Looking forward to network improvement once on DMZ / LHCONE

