

Shoveler testing and validation

Katy Ellis, XRootD workshop, 12th Sept 2024



Science and
Technology
Facilities Council

Caveats

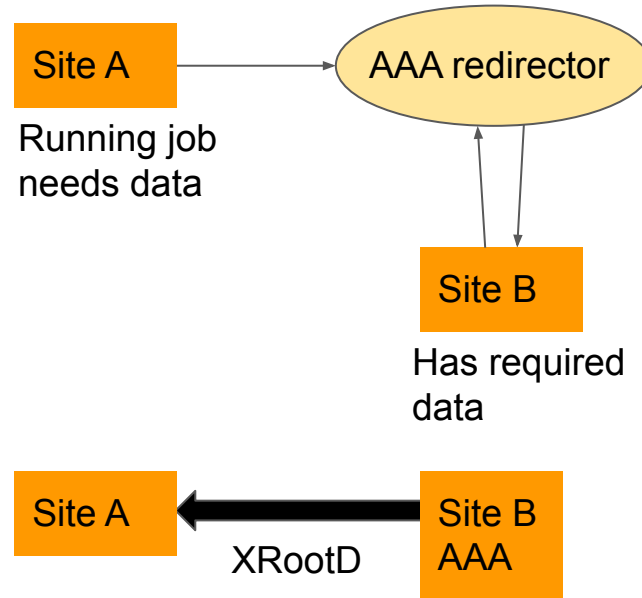



- This is very much a work in progress!
- Speaking from a CMS perspective

- Note - I'm planning to give an updated version of this talk at CHEP in a few weeks

Introduction

- The XRootD transfer protocol is used extensively by **CMS** jobs
 - Streaming data from local storage or remote storage via AAA
 - Writing output data to storage



XRootD monitoring for WLCG


- Until recently you could view XRootD transfer rates in the WLCG dashboard, which used the 'GLED' software - data sent to 'COLLECTOR'
- GLED monitoring has long been considered incomplete and unreliable
 - Packets lost when transmitted to the collector?
 - <https://zenodo.org/records/4688624>
- This monitoring was turned off in June due to:
 - Shoveler having been developed (although not fully deployed)
 - Not wanting to upgrade GLED for a new OS
- The new XRootD monitoring is called Shoveler
 - Originated from the WLCG Monitoring Taskforce
 - Developed by Derek Weitzel (UNL) and run by Borja Garrido Bear (CERN)

Why is it so important for CMS?

- ATLAS mostly bring entire data files to their jobs
 - CMS data is streamed - only parts needed by the job
- CMS make extensive use of remote reads (AAA) as well as local
 - Are jobs failing because either network or slow storage?
 - Is job efficiency low at certain sites due to slow reads?
 - In an era of finite network bandwidth, how much WAN capacity is CMS using?
 - As I showed in the data challenge, XRootD ~ FTS (see next slide) according to GLED
 - But now I'm not sure if this includes local reads
- To make improvements, we need accurate and informative monitoring

Old XRootD monitoring - CMS

Transfer Throughput



How Shoveler works - passes UDP to message bus

1. Install Shoveler software on real or virtual machine in your data centre [GitHub - opensciencegrid/xrootd-monitoring-shoveler](#)

-Install the RPM from the latest release.

-Start the systemd service with:

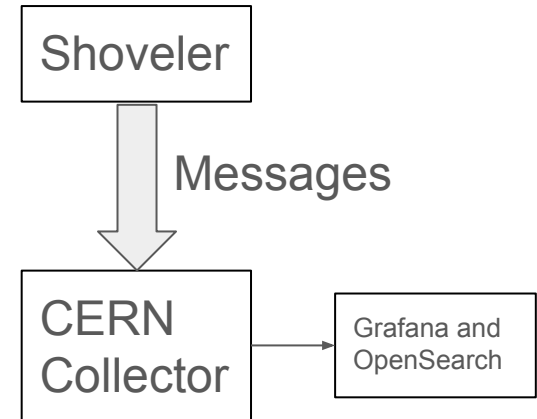
```
systemctl start xrootd-monitoring-shoveler.service
```



Shoveler

How Shoveler works

1. Install Shoveler software on real or virtual machine in your data centre [GitHub - opensciencegrid/xrootd-monitoring-shoveler](#)
2. (Shoveler needs permissions to send data to the CERN collector)



How Shoveler works

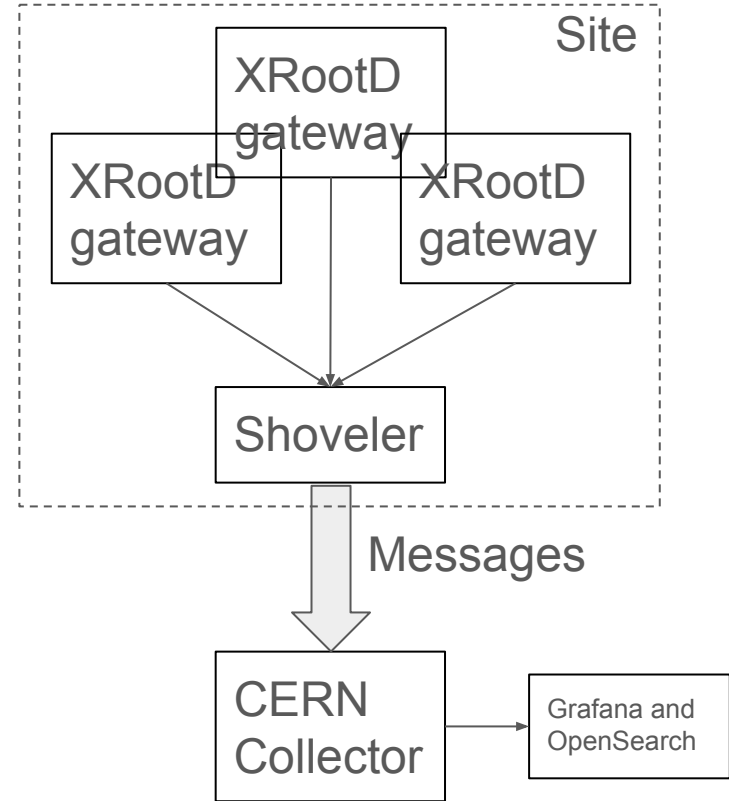
1. Install Shoveler software on real or virtual machine in your data centre [GitHub - opensciencegrid/xrootd-monitoring-shoveler](#)
2. (Shoveler needs permissions to send data to the CERN collector)
3. Add config to xrootd access points

```
xrootd.monitor all auth fstat 10s ops lfn xfr 1 ident 1m dest fstat info user redir 172.16.105.115:9931  
dest fstat info user redir shoveler-monit01.gridpp.rl.ac.uk:9993
```

**RAL example
only - do not use!**

New Shoveler
machine at RAL
(now using IP addr)

Existing RAL-based
Kibana monitoring



How Shoveler works

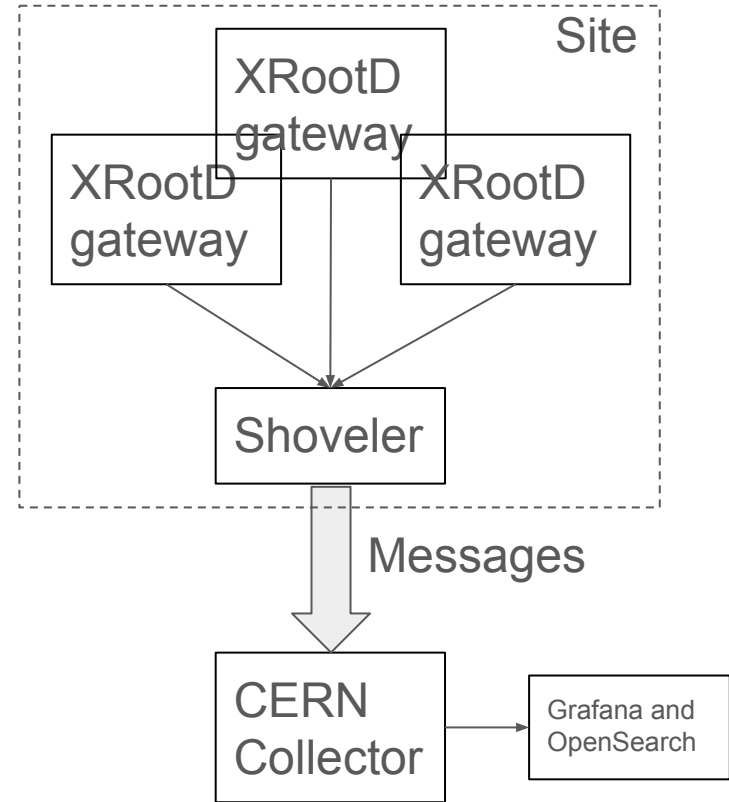
1. Install Shoveler software on real or virtual machine in your data centre [GitHub - opensciencegrid/xrootd-monitoring-shoveler](#)
2. (Shoveler needs permissions to send data to the CERN collector)
3. Add config to xrootd access points

```
xrootd.monitor all auth fstat 10s ops 1m xm-1 ident 1m dest fstat info user redir 172.16.105.115:9931  
dest fstat info user redir shoveler-monit01.gridpp.rl.ac.uk:9993
```

**RAL example
only - do not use!**

New Shoveler
machine at RAL
(now using IP addr)

Existing RAL-based
Kibana monitoring



N.B. If the connection between Site and CERN is lost, data is stored on the Shoveler machine in /tmp/

Shoveler from its own docs

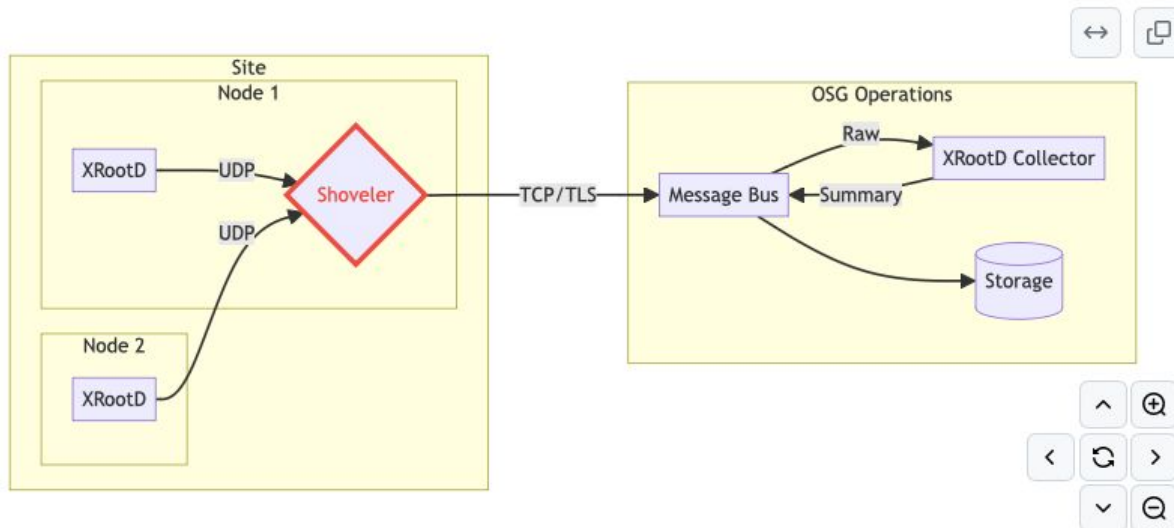
<https://github.com/opensciencegrid/xrootd-monitoring-shoveler>

XRootD Monitoring Shoveler

This shoveler gathers UDP monitoring messages from XRootD servers and sends them to a reliable message bus.

Unit Testing passing Go Linting passing CodeQL Static Analysis passing [reference](#) [license](#) [Apache-2.0](#)

[Documentation](#) · [Report Bug](#) · [Request Feature](#)



Long-term testing of Shoveler at RAL

- RAL has been running Shoveler at RAL for ~2 years
- Only monitoring the AAA proxy gateways until recently
 - Observe what data other sites read from RAL disk storage
 - Possibility to compare with internal monitoring
- Issues fixed:
 - Shoveler process dying silently (although appears still to be running)
 - Ability to identify transfer by VO
 - Lack of rate monitoring

The dCache exception

- For some reason, dCache endpoints cannot use Shoveler
 - Some of CMS' largest sites run dCache, e.g. FNAL
 - A lot of missing traffic
- A script was written to try to produce a similar output as Shoveler
 - Evidence that some sites are using this
- My current focus is on validating Shoveler, but clearly both are important

What is being monitored so far?

- (I do not have a comprehensive list)
- CERN
- RAL (some parts)
- DESY and KIT (test)
- US sites?
- Definitely not all sites - rollout campaign is on hold

How can I see the monitoring?

- In the WLCG > XRootD transfers dashboard
<https://monit-grafana.cern.ch/d/000000444/xrootd-transfers?orgId=20>

Remember, this is still incomplete, and not yet validated!

WLCG

Home > Dashboards > Transfers > XRootD Transfers

Group By: vo, Bin: 1d, VO: All, Source Country: All, Dest Country: All

Remote Access: true, IPv6: All, Filters +

> Next Generation (4 panels) ← This is Shoveler


> dCache (3 panels) ← This is the dCache 'script'

> GLED (4 panels) ← This is the old GLED monitoring

Observation:
It would be nice to have these combined!

Shoveler monitoring in Grafana - all VOs


Transfer Throughput



Shoveler monitoring in OpenSearch

Use index:

monit_prod_xrootdng_enr*



Time	data.vo	data.file_lfn
> Aug 20, 2024 @ 16:31:28.000	cms	atlas:datadisk/rucio/data17_13TeV/6c/31/DAOD_PHYSLITE.39756495._000479.pool.root.1
> Aug 20, 2024 @ 16:31:28.000	cms	atlas:datadisk/rucio/data17_13TeV/6c/31/DAOD_PHYSLITE.39756495._000479.pool.root.1
> Aug 20, 2024 @ 16:31:28.000	cms	atlas:datadisk/rucio/data17_13TeV/21/b2/DAOD_PHYSLITE.39756495._000480.pool.root.1
> Aug 20, 2024 @ 16:31:26.000	cms	atlas:scratchdisk/rucio/panda/48/cb/panda.0820162507.986438.lib._40917469.39700958953.lib.tgz
> Aug 20, 2024 @ 16:31:26.000	alice	/eos/alice/grid/11/49035/86aa4bf4-e749-11ee-8010-0242a36ce054
> Aug 20, 2024 @ 16:31:26.000	alice	/eos/alice/cond/10/35749/e6cc59de-d22d-11ec-8dd8-08f1eaf024ee
> Aug 20, 2024 @ 16:31:26.000	cms	/lhcb:buffer/lhcb/LHCb/Collision24/BNOC.DST/00235369/0110/00235369_01106375_1.bnoc.dst

?!
?!
?!
?!
?!

?!

Shoveler monitoring in OpenSearch


```
> Sep 11, 2024 @ 10:49:56.000 cms /lhcb:prod/lhcb/LHCb/Collision17/BHADRON.MDST/00206514/0000/00206514_00001317_1.bhadron.mdst
> Sep 11, 2024 @ 10:49:56.000 cms missing directory
> Sep 11, 2024 @ 10:49:56.000 cms /lhcb:prod/lhcb/LHCb/Collision18/BHADRON.MDST/00210361/0002/00210361_00024568_1.bhadron.mdst
> Sep 11, 2024 @ 10:49:56.000 cms missing directory
> Sep 11, 2024 @ 10:49:55.000 cms /lhcb:buffer/lhcb/MC/2011/SIM/00239604/0000/00239604_00000848_1.sim

> Sep 11, 2024 @ 10:49:51.000 /store/mc/RunIISummer20UL18NanoAODv9/ZZTo4L_TuneCP5_13TeV_powheg_pythia8/NANOADSIM/106X_upgrade2018_realistic_v16_L1v1-v2/40000/240589F7-3596-934C-8B80-7EC2F7DB9FEE.root DE-DESY-ATLAS-T2
> Sep 11, 2024 @ 10:49:51.000 /store/mc/RunIISummer20UL16NanoAODAPVv9/WWto4Q_4f_TuneCP5_13TeV-amcatnLoFXFX-pythia8/NANOADSIM/106X_mcRun2_asymptotic_preVFP_v11-v3/70000/6B22927A-6D17-9446-9A6A-E8002ACF96AC.root DE-DESY-ATLAS-T2
```


Shoveler monitoring in Grafana

▼ Next Generation


Transfer Throughput




Average Operation Time



Transferred Volume



Transfers Finished



● Similar for dCache

Failure failure


- XRootD (and hence Shoveler) has no obvious way to monitor xrootd transfer failures
- We can count number of accesses, the volume of data, length of connection and calculate transfer rates...but we cannot monitor failures as we do for FTS transfers
- Our jobs (CMSSW) can tell if a job has failed with a local or remote read...so why can't Shoveler at least give a success/failure count?
- (Even better would be to categorise error messages)



Many of the usual group-by options in Grafana

~ Next Generation

Transfer Throughput



vo

src_tier

src_country

src_federation

src_site

src_experiment_site

src_hostname

dst_tier

dst_country

dst_federation

dst_site

dst_experiment_site

dst_hostname

ipv6

Grafana plot 'disappears'? (Group by src_hostname)

Next Generation

Transfer Throughput

4 GB/s

3 GB/s

2 GB/s

1 GB/s

0 MB/s

Value

0.1.2.1.2.0.a.2.dynamic.cust.swisscom.net

10.0.150.25

10.1.11.107

10.1.11.108

10.1.11.112


10.1.11.115

10.1.11.116

Grafana plot 'disappears'? (Group by src_hostname)

Next Generation

Transfer Throughput



Next Generation


Transfer Throughput



Grafana monit grouped by src_site

Next Generation

Transfer Throughput



General issues with Grafana monitoring

- XRootD data no longer appearing in the WLCG monit since GLED retired
- Need to combine 'Shoveler' and 'dCache' XRootD monitoring
- Minimum binning is 1 hour - makes testing more difficult
- Data appearing in some dashboard plots and not others
- Cross-labelling between sites
- Too much 'UNKNOWN' in the plots

UNKNOWN info

#	data.count	1
t	data.dst_country	UNKNOWN
t	data.dst_country_code	UNKNOWN
t	data.dst_domain	UNKNOWN
t	data.dst_experiment_site	UNKNOWN
t	data.dst_federation	UNKNOWN
t	data.dst_hostname	UNKNOWN
t	data.dst_site	UNKNOWN

#	data.count	1
t	data.dst_country	Germany
t	data.dst_country_code	DE
t	data.dst_domain	gridka.de
t	data.dst_experiment_site	T1_DE_KIT
t	data.dst_federation	DE-KIT
t	data.dst_hostname	c02-119-115.gridka.de
t	data.dst_site	FZK-LCG2

- Sites not up to date in CRIC?
- How can it not know the destination hostname?

Testing: back-to-basics

Isolated file transfer tests


- Transferred various sized files 5 times each from RAL disk through a AAA proxy gateway and looked for the evidence in each of three monitoring systems:
 - RAL Vande (from host network); RAL kibana; Shoveler (see next slides)
- 2GB file:

```
[rrq44858@lcfgui06 ~]$ for ((i=0; i<5; i++)); do date; xrdcp -f root://ceph-svc20.gridpp.rl.ac.uk:1094//store/temp/ShovelerTest_2GB /dev/null; sleep 300; done
Mon 19 Aug 12:48:10 BST 2024
[1.863GB/1.863GB][100%][=====][70.64MB/s]
Mon 19 Aug 12:53:37 BST 2024
[1.863GB/1.863GB][100%][=====][79.47MB/s]
Mon 19 Aug 12:59:01 BST 2024
[1.863GB/1.863GB][100%][=====][79.47MB/s]
Mon 19 Aug 13:04:25 BST 2024
[1.863GB/1.863GB][100%][=====][79.47MB/s]
Mon 19 Aug 13:09:49 BST 2024
[1.863GB/1.863GB][100%][=====][82.93MB/s]
```

Shoveler monitoring (OpenSearch)

- Shoveler picks up the 5 transfers in OpenSearch
- Times match up with the transfer-end time

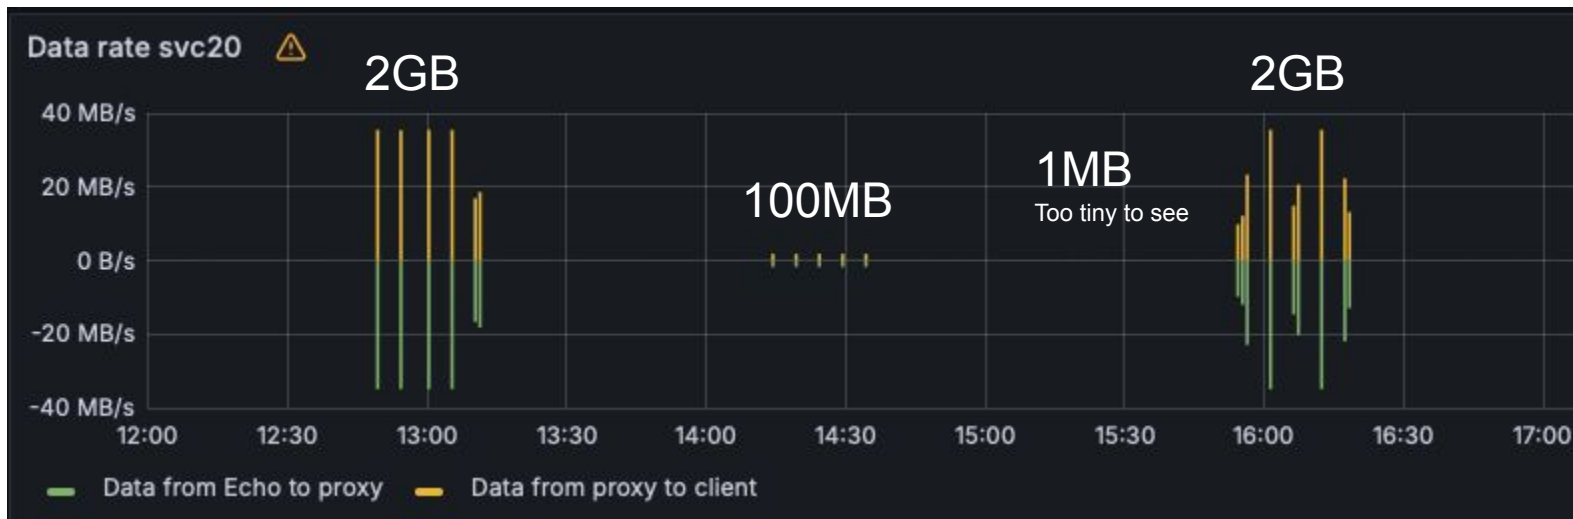
```
Mon 19 Aug 12:48:10 BST 2024  
[1.863GB/1.863GB][100%][====  
Mon 19 Aug 12:53:37 BST 2024  
[1.863GB/1.863GB][100%][====  
Mon 19 Aug 12:59:01 BST 2024  
[1.863GB/1.863GB][100%][====  
Mon 19 Aug 13:04:25 BST 2024  
[1.863GB/1.863GB][100%][====  
Mon 19 Aug 13:09:49 BST 2024  
[1.863GB/1.863GB][100%][====
```



Time (GMT)	data.dst_hostname	data.src_hostname	data.file_lfn	data.file_size	data.ipv6	data.is_transfer	data.operation_time	data.read_operations
Aug 19, 2024 @ 12:10:12.000	lgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/Shoveler Test_2GB	2,000,000,000	true	true	23	239
Aug 19, 2024 @ 12:04:49.000	lgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/Shoveler Test_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:59:25.000	lgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/Shoveler Test_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:54:01.000	lgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/Shoveler Test_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:48:37.000	lgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/Shoveler Test_2GB	2,000,000,000	true	true	26	239

RAL Vande transfers

- Internal monitoring from the RAL Vande (graph-generator) monitoring
- Data comes from the host network monitoring via telegraf
- Binning is 1 minute
- Rates are commensurate with transfer time (<30 seconds)



(6) transfers in RAL Kibana compared with Shoveler

- Kibana has same data stream as Shoveler (xrootd.monitor)
- Queries the 'fstream.close' message-type, as Shoveler uses

Time	xrd_XFR_read	xrd_derived_read_rate	xrd_derived_duration
August 19th 2024, 16:17:12.000	1.863GB	70.643MB	27.00
August 19th 2024, 16:11:44.000	1.863GB	68.565MB	27.82
August 19th 2024, 16:06:16.000	1.863GB	70.643MB	27.00
August 19th 2024, 16:00:49.000	1.863GB	61.527MB	31.00
August 19th 2024, 15:55:18.000	1.863GB	68.12MB	28.00
August 19th 2024, 15:53:15.000	536MB	67MB	8.00

Xrdcp output
27.4
28.3
27.4
30.4
28.1

RAL Kibana

Partial transfer

I killed this one on purpose

Time	data.ipv6	data.is_transfer	data.operation_time	data.read_operations	data.read_bytes_at_close
> Aug 19, 2024 @ 15:17:12.000	true	true	27	239	2,000,000,000
> Aug 19, 2024 @ 15:11:44.000	true	true	28	239	2,000,000,000
> Aug 19, 2024 @ 15:06:16.000	true	true	27	239	2,000,000,000
> Aug 19, 2024 @ 15:00:49.000	true	true	31	239	2,000,000,000
> Aug 19, 2024 @ 14:55:18.000	true	true	28	239	2,000,000,000
> Aug 19, 2024 @ 14:53:15.000	true	false	8	67	562,036,736

Shoveler
OpenSearch

Vector reads?

- CMS jobs typically do not download a whole file but stream the parts needed...often using “vector reads”
- I simulated this to see how Shoveler handles the different operation

```
chunks=[]
for i in range(1, 101):
    chunks.append((1000000*i, 10))
status, res = fd.vector_read(chunks)
```


data.file_lfn	data.is_transfer	data.file_size	data.operation_time	data.read_operations	data.read_vector_bytes	data.read_vector_count_max
/store/temp/ShovelerTest_2GB	false	2,000,000,000	0	1	1,000	100

↑
XRootD doesn't do sub-second operation timing


↑
Number of bytes read

↑
100 vector reads made

Shoveler rates compared with network rates



Transfer Throughput



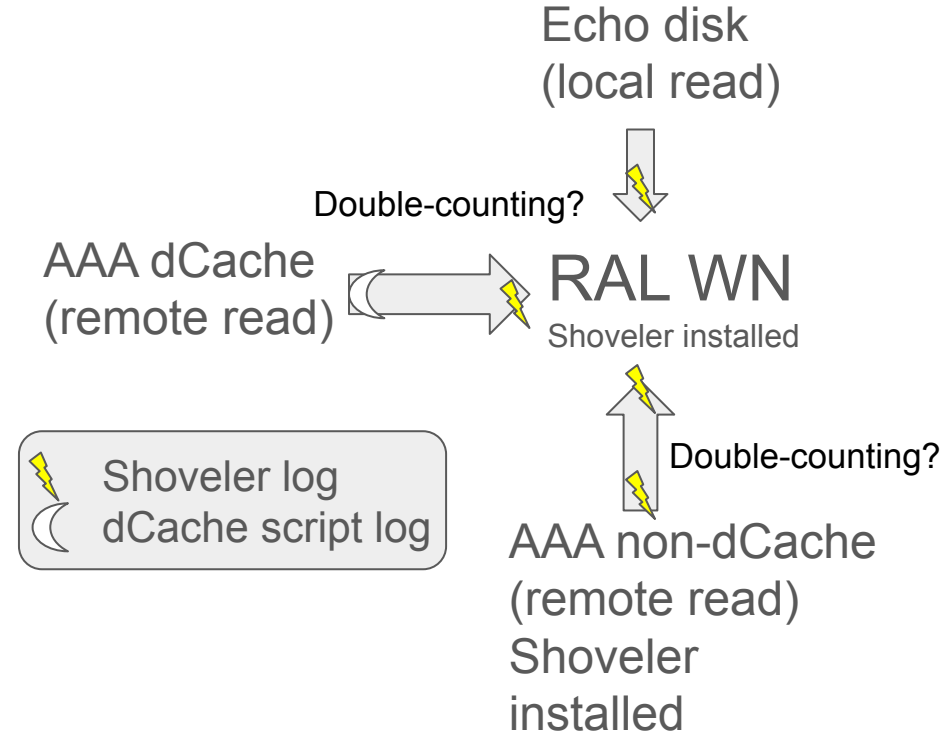
AAA access gateway proxy for remote sites reading from RAL disk

Additional work at RAL

- RAL is 'different' as usual :-)
- To make local reads from our batch farm to our disk storage, we use a lightweight 'gateway' on every WN
 - To monitor XRootD reads from the batch farm, every WN needs the Shoveler config
- Rolled out last week
- We were curious to see how Shoveler handles the number of machines
 - Although traffic should not be too heavy
 - Nevertheless the entire batch farm would be more connections than previously tested by Shoveler team

Additional work at RAL

- RAL is 'different' as usual :-)
- To make local reads from our batch farm to our disk storage, we use a lightweight 'gateway' on every WN
 - To monitor XRootD reads from the batch farm, every WN needs the Shoveler config
- Rolled out last week
- We were curious to see how Shoveler handles the number of machines
 - Although traffic should not be too heavy
 - Nevertheless the entire batch farm would be more connections that previous tested by Shoveler team




Monitoring RAL WNs (OpenSearch)

5,292 hits



Sep 11, 2024 @ 19:10:02.523 - Sep 12, 2024 @ 10:10:02.523 per Auto



Time	data.dst_hostname	data.src_hostname	data.dst_experiment_site	data.src_site
> Sep 12, 2024 @ 10:02:49.000	lcg2604.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:02:48.000	lcg2593.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:02:48.000	lcg2593.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:02:24.000	lcg2541.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:01:52.000	lcg2541.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:01:27.000	lcg2573.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:01:10.000	lcg2541.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL
> Sep 12, 2024 @ 10:00:47.000	lcg2604.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL

Monitoring RAL WNs (OpenSearch)

data.dst_hostname	data.src_hostname	data.dst_experiment_site	data.src_site	data.operation	data.read_bytes	data.read_vector_bytes	data.write_bytes	data.operation_time
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	157
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	140
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	120
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	40
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	34

Monitoring RAL WNs (OpenSearch)

data.dst_hostname	data.src_hostname	data.dst_experiment_site	data.src_site	data.operation	data.read_bytes	data.read_vector_bytes	data.write_bytes	data.operation_time
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	157
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	140
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	120
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	40
lcg2555.gridpp.rl.ac.uk	UNKNOWN	T1_UK_RAL	T1_UK_RAL	unknown	0	0	0	34

“XCache configuration required”

Future plans

- Consult with Borja on issues and anomalies found
- More plot comparisons!
- Higher stress-test
- Analyse additional sites that are involved in testing
 - Including the dCache sites
- Re-set up Shoveler on RAL WNs (XCache config) and other machines (external gateways to disk storage + stage out from RAL WN jobs) - can Shoveler handle the connections?

- A parallel talk and proceedings at CHEP24 on the validation
- Confident to use Shoveler for XRootD and AAA analysis by the end of 2024

Summary

- Testing is progressing
 - Some significant monitoring issues still to address
- XCache Shoveler needed for RAL WNs?
- Basic transfer tests look good - higher stress would be better
- More difficult to ensure every message is logged in a busy system
- I am not yet advocating a roll-out over many sites
 - But some sites have joined the testing - thanks!
- Hoping to make progress on the validation so we can roll-out at all CMS sites and actually start learning about our XRootD usage before the end of the year