Shoveler testing and validation for CMS

Katy Ellis, CHEP24, 22nd Oct 2024

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

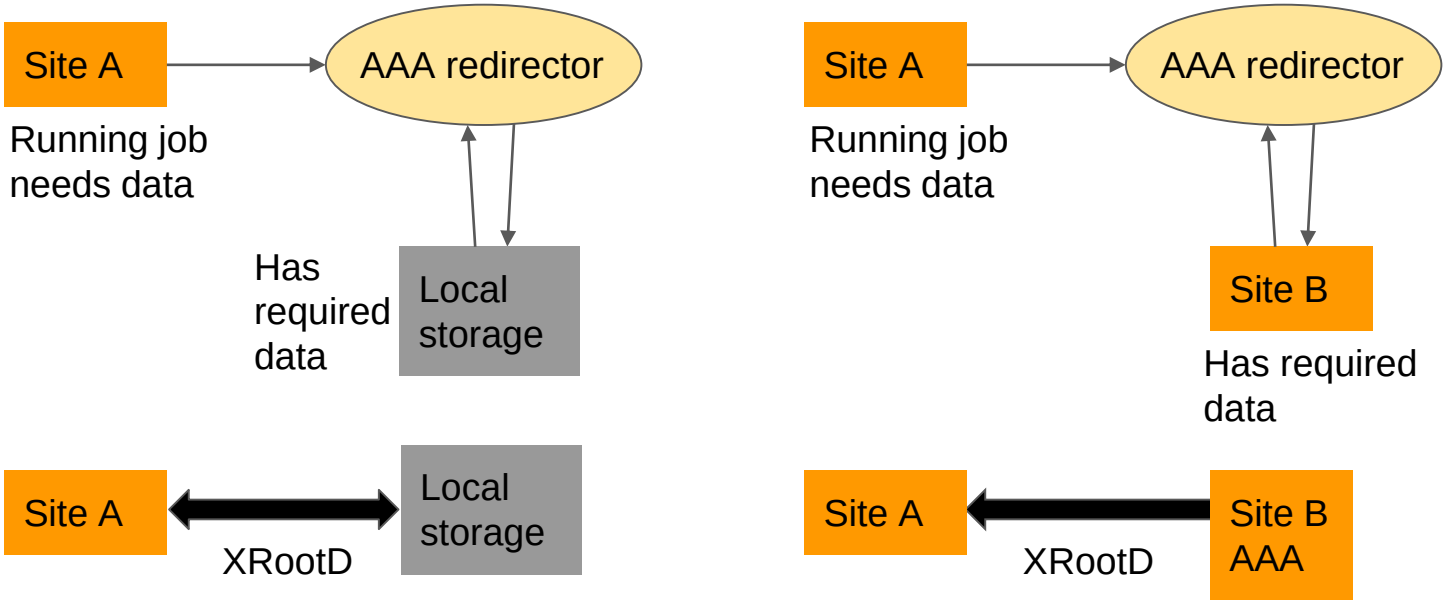

XRootD monitoring for WLCG

- 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)

Importance 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?
- To make improvements, we need accurate and informative monitoring
  - Hence the reason for this work
    - Accurate and complete monitoring
    - User-friendly experience
    - This talk gives current status of work in progress

Old XRootD monitoring - CMS

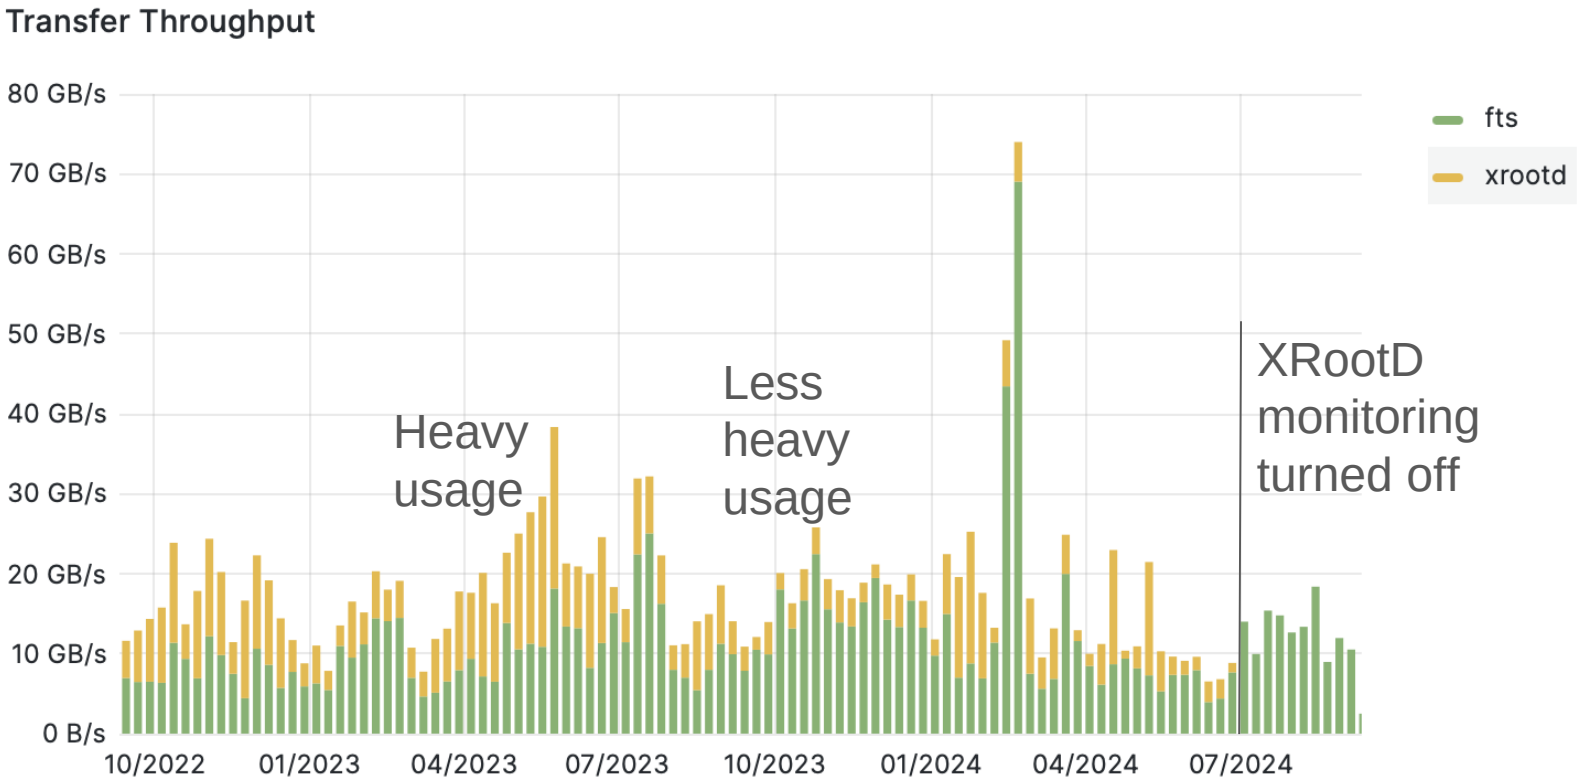

How Shoveler works

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

RAL example only - do not use!

```
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 shbveler-monit01.gridpp.rl.ac.uk:9993
```

New Shoveler machine at RAL (now using IP addr)

Existing RAL-based Kibana monitoring

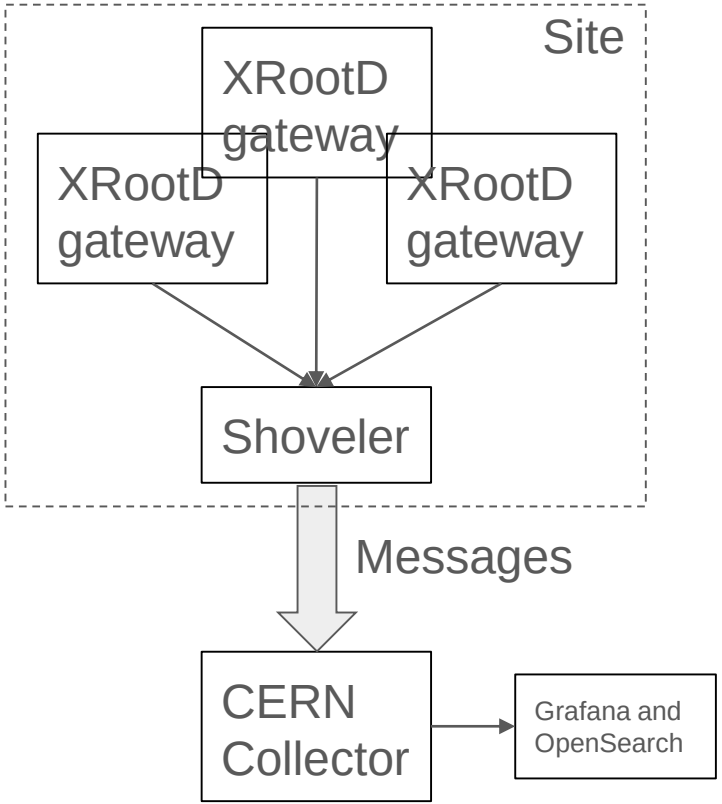

N.B. If the connection between Site and CERN is lost, data is stored on the Shoveler machine in /tmp/

The dCache approach

- dCache endpoints have a non-native XRootD implementation which not directly compatible with Shoveler
  - Some of CMS' largest sites run dCache, e.g. FNAL
  - A lot of traffic would be missed
- A script was written to match the Shoveler monitoring
  - Evidence that some sites are using this
- My current focus is on validating Shoveler, but clearly both are important for CMS

What is being monitored so far?

- CERN
- RAL and some other UK sites
  - >2 years testing RAL AAA gateways
  - Issues addressed:
    - Shoveler process dying silently (although appears still to be running)
    - Ability to identify transfer by VO
    - Lack of rate monitoring
- DESY and KIT (partial, using dCache script)
- Many US sites
- Not all sites - rollout campaign is on hold

How can I see the monitoring?

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

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

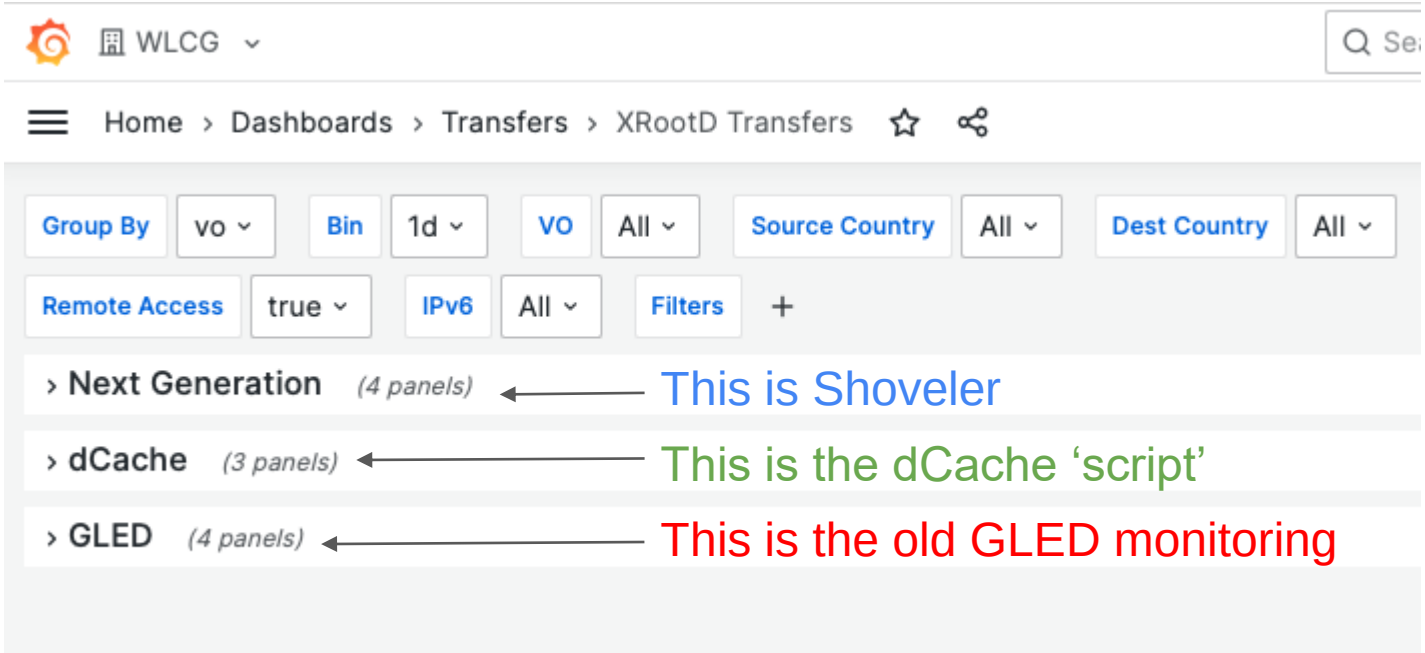

Observation: It would be nice to have these combined!

Shoveler monitoring in Grafana by VOs

A ticket to track Unknown VO is here

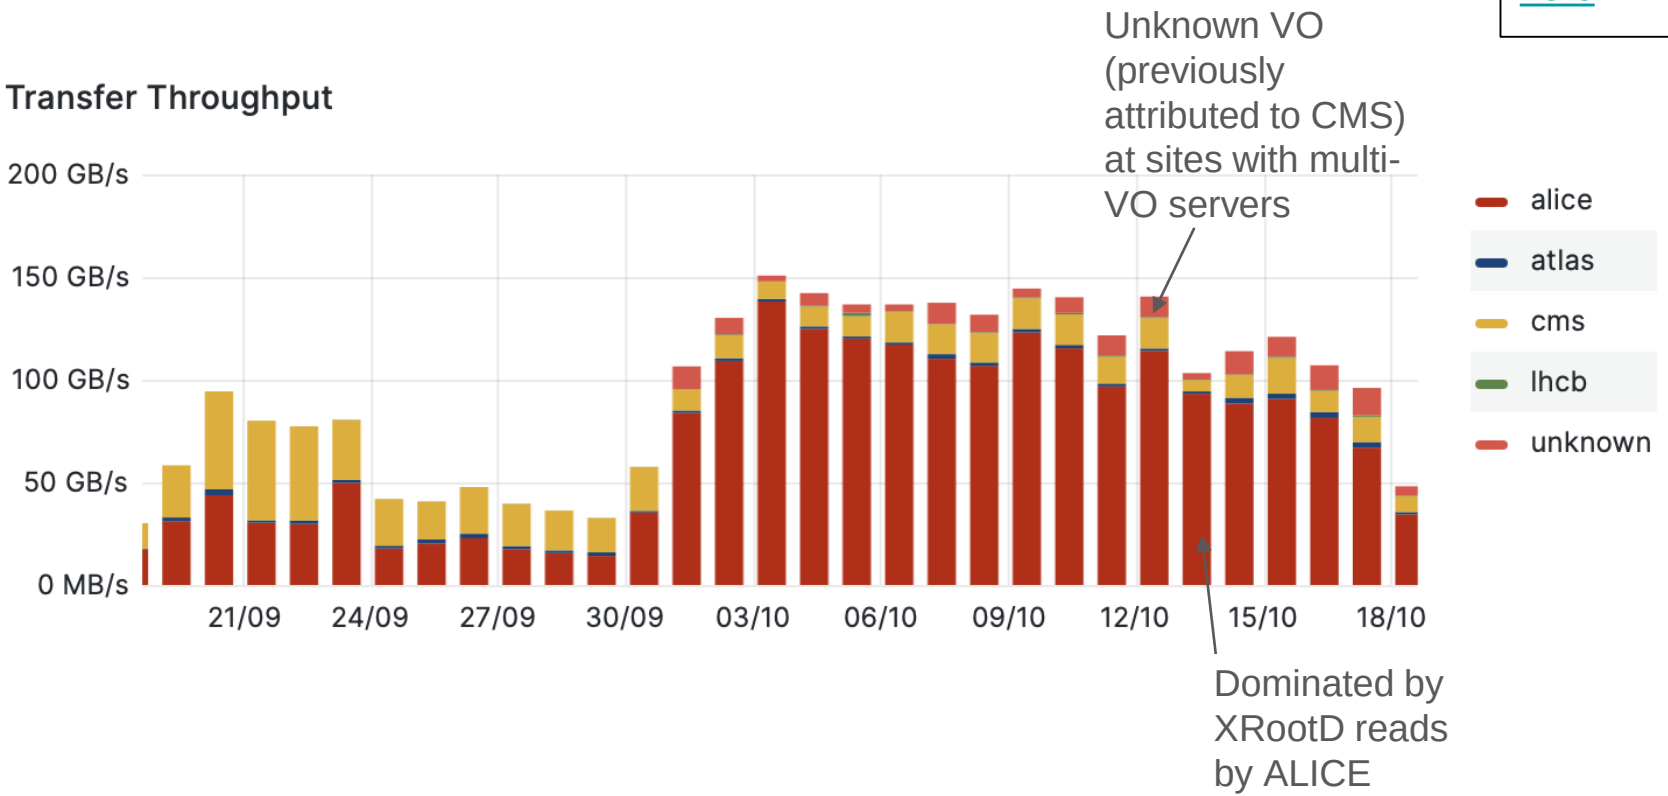

Inconsistent throughput/volume

Now fixed!

Supposing the throughput is true…and that we are using a 1-day bin-size…calculate the volume for the first bin

Volume = 140GB/s * 24 * 3600 = 12PB (not 40PB)

Grafana monit grouped by src_site

Next Generation

Transfer Throughput

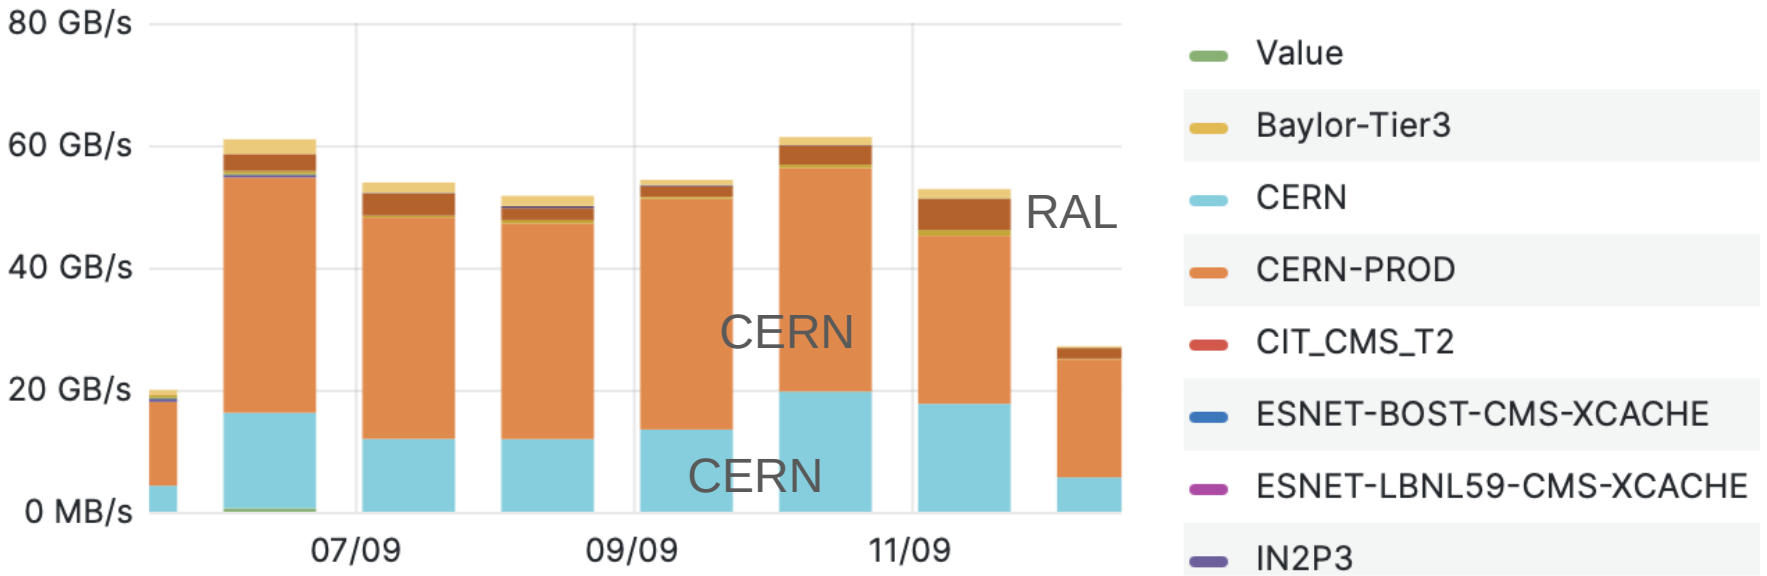

Grafana plot ‘disappears’? (Group by src_hostname)

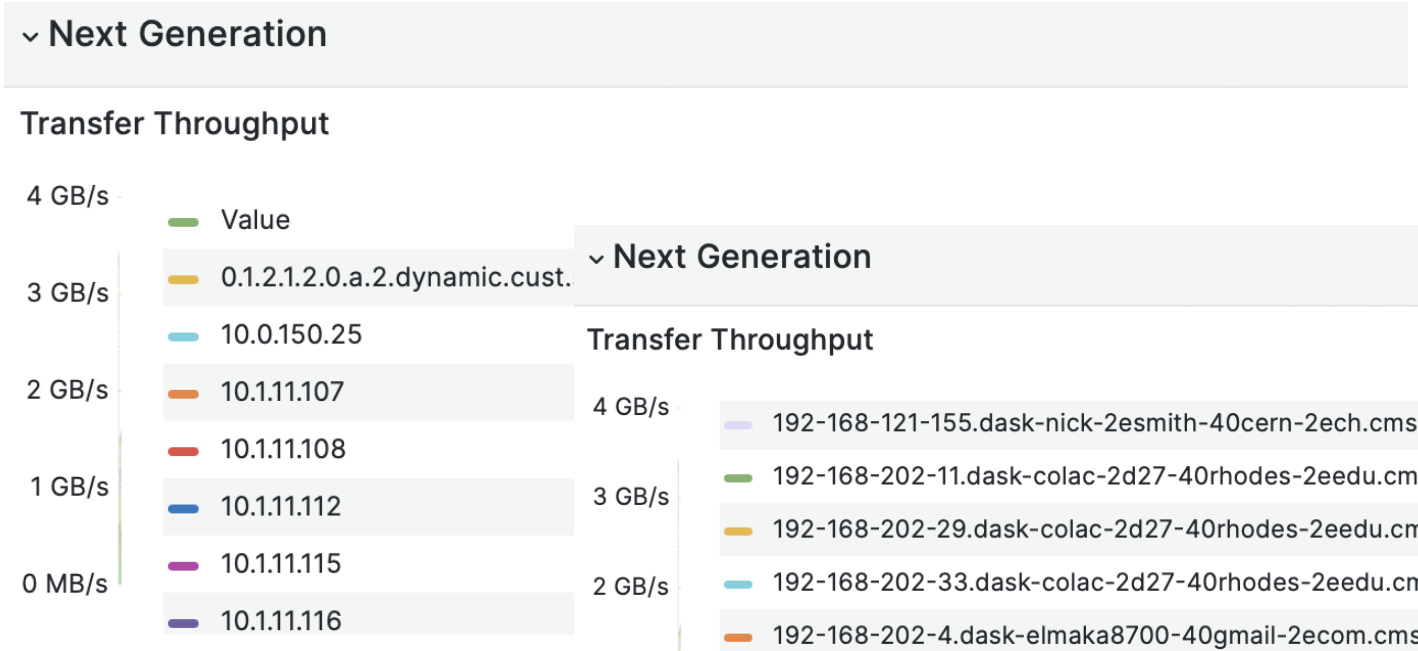

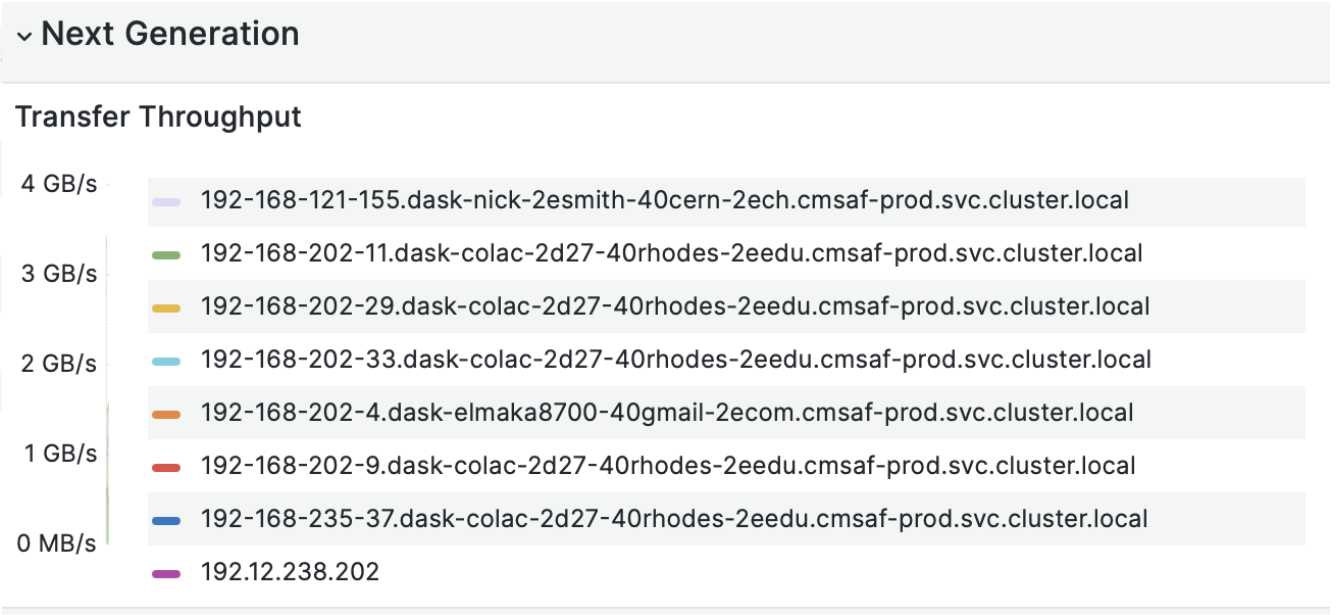

Shoveler monitoring in OpenSearch

Use index: monit_prod_xrootdng_enr*

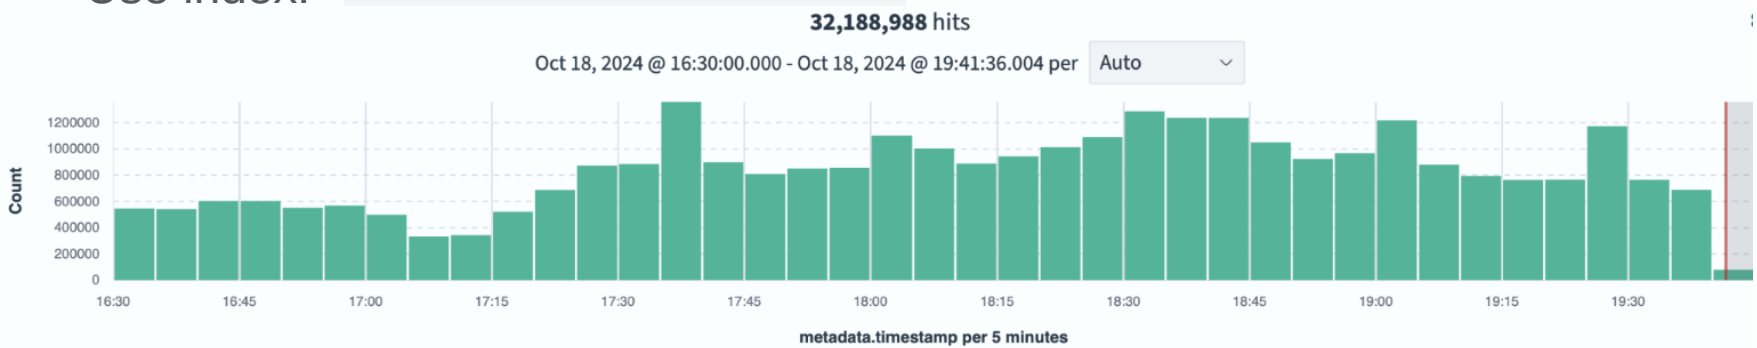

Time	data.src_hostname	data.server_site	data.dst_hostname	data.dst_site	data.remote_access
Oct 18, 2024 @ 19:41:03.000	g36n13.hep.wisc.edu	T2_US_Wisconsin	g32n16 ?	UNKNOWN	true
Oct 18, 2024 @ 19:40:52.000	p06636710b90882.cern.ch	CERN-PROD		UNKNOWN	true
Oct 18, 2024 @ 19:40:52.000	p06636710b90882.cern.ch	CERN-PROD	missing ?	UNKNOWN	true
Oct 18, 2024 @ 19:40:52.000	p06636710b90882.cern.ch	CERN-PROD		UNKNOWN	true
Oct 18, 2024 @ 19:40:52.000	p06636710y31446.cern.ch	CERN-PROD	p06636710u99343.cern.ch	CERN-PROD	false
Oct 18, 2024 @ 19:40:52.000	?	CERN-PROD	p06636710b90882.cern.ch	CERN-PROD	true

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

Shoveler monitoring in OpenSearch

```
/lhcb:prod/lhcb/LHCb/Collision17/BHADRON.MDST/00206514/0000/00206514_00001317_1.bhadron.mdst
missing directory
/lhcb:prod/lhcb/LHCb/Collision18/BHADRON.MDST/00210361/0002/00210361_00024568_1.bhadron.mdst
missing directory
/lhcb:buffer/lhcb/MC/2011/SIM/00239604/0000/00239604_00000848_1.sim
```

10.41.12.1

>	Sep 11, 2024 @ 10:49:51.000	/store/mc/RunIISummer20UL18NanoAODv9/ZZTo4L_TuneCP5_13TeV_powheg_pythia8/NANOAODSIM/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/NANOAODSIM/106X_mcRun2_asymptotic_preVFP_v11-v3/70000/6B22927A-6D17-9446-9A6A-E8002ACF96AC.root	DE-DESY-ATLAS-T2

Failure failure

- XRootD (and hence Shoveler) has no obvious way to monitor *streaming* 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
- CMS is *strongly* in favour of this
- Issue is now being considered here

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:

```
[rqq44858@lcgui06 ~]$ 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%][====
```

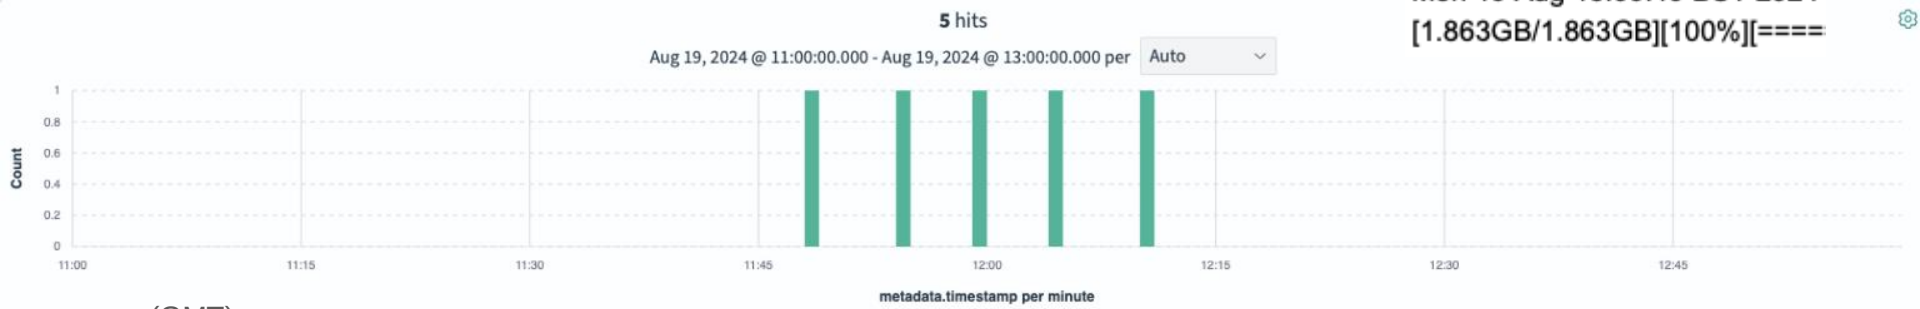

(GMT)

Time	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	lcgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/ShovelerTest_2GB	2,000,000,000	true	true	23	239
Aug 19, 2024 @ 12:04:49.000	lcgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/ShovelerTest_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:59:25.000	lcgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/ShovelerTest_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:54:01.000	lcgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/ShovelerTest_2GB	2,000,000,000	true	true	24	239
Aug 19, 2024 @ 11:48:37.000	lcgui06.gridpp.rl.ac.uk	ceph-svc20.gridpp.rl.ac.uk	/store/temp/ShovelerTest_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)

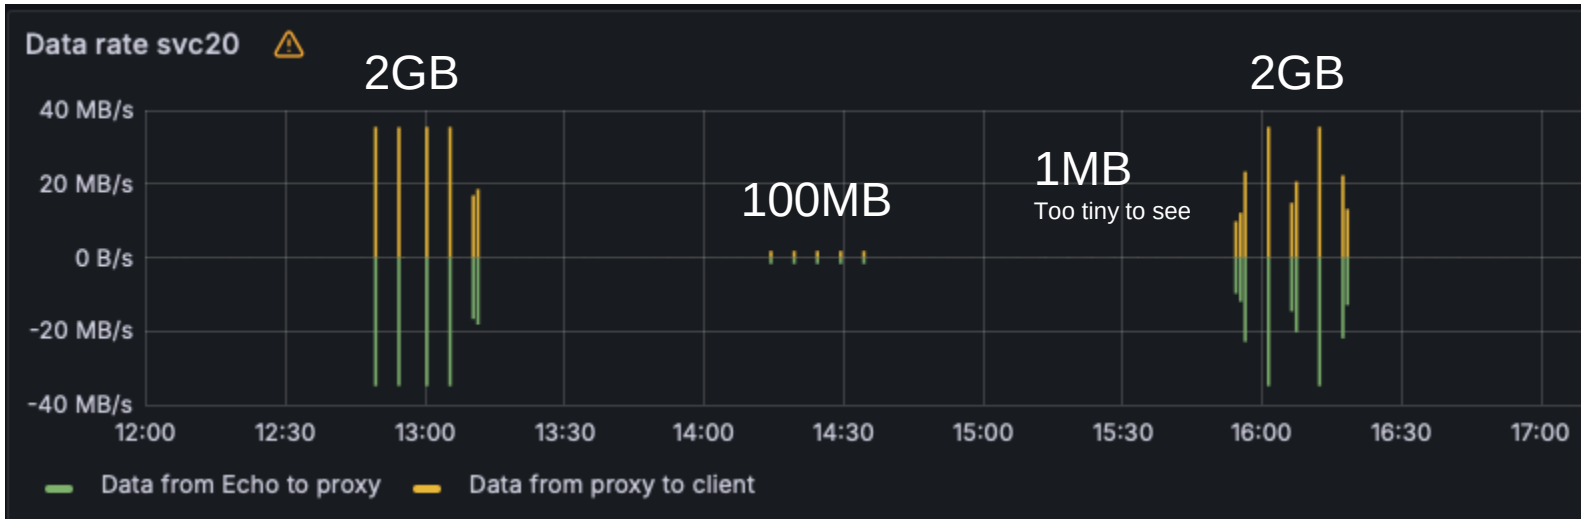

(6) transfers in RAL Kibana compared with Shoveler

- 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	Xrdcp output	xrd_derived_duration
August 19th 2024, 16:17:12.000	1.863GB	70.643MB	27.4	27.00
August 19th 2024, 16:11:44.000	1.863GB	68.565MB	28.3	27.82
August 19th 2024, 16:06:16.000	1.863GB	70.643MB	27.4	27.00
August 19th 2024, 16:00:49.000	1.863GB	61.527MB	30.4	31.00
August 19th 2024, 15:55:18.000	1.863GB	68.12MB	28.1	28.00
August 19th 2024, 15:53:15.000	536MB	67MB		8.00

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

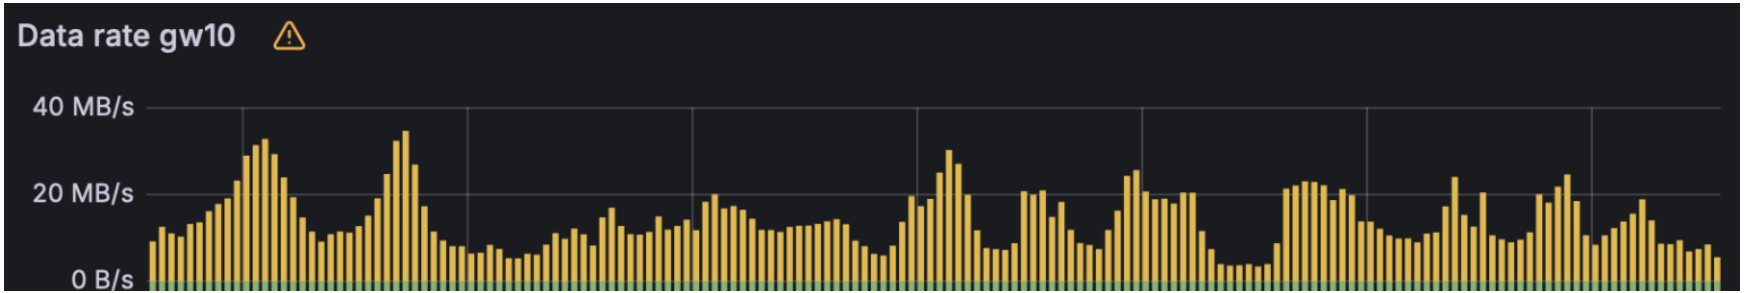

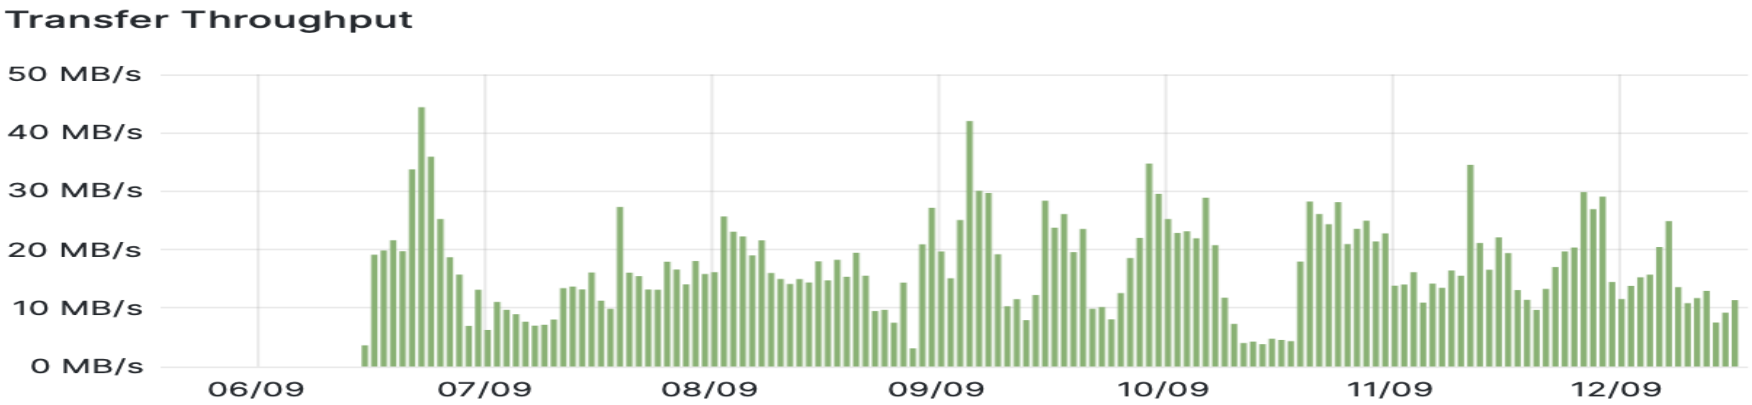

AAA access gateway proxy for remote sites reading from RAL disk

Summary

- Testing is progressing
  - Some significant ‘visualisation’ improvements still to address
  - More monitoring comparisons to be done
  - Including dCache sites
- Basic transfer tests look good - higher stress would be better
  - More difficult to ensure every message is logged in a busy system
- CMS strongly in favour of failure monitoring to be developed
- I am not yet advocating a roll-out over many sites
  - I don’t see the lack of failure monitoring as a blocker to this
  - But some sites have joined the testing - thanks!