

First Graphs, Results and Questions from EOS metrics

Philipp Zigann

CERN IT-DSS-DT

19th Feb. 2013



- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)



- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)



Event/Entry

An EOS log entry that is written at the point of closing a file and describes what happened since opening it.

User

Defined by the user name within the EOS log file field td [e.g.: atlas001]
(ruid is not unique)

Example EOS Log Entry

```
log=7677503c-adc7-11e1-9083-003048f0e00c&path=/eos/atlas/atl...Ele.root&ruid=38112&rgid=1307&td=username.12459:127@lxplus309&host=lxfsrg15a07.cern.ch&lid=6291730&fid=45557244&fsid=2246&ots=1338760799&otms=547&cts=1338760890&ctms=654&rb=615562&wb=0&srb=368145245672&swb=0&nrc=70&nwc=0&rt=28.48&wt=0.00&osize=5671631075&csize=5671631075
```

Parameters

- File information
- User identification
- Number of seeked, written, read bytes (and used calls)
- Open and close time
- Waiting time for I/O

No information about a single read/write call!



- 1.-14. January 2013 (2 weeks)

Param	ATLAS	CMS
Entries	12,871,560	9,641,880
Reads	9,395,903	7,821,420
Writes	3,322,552	1,763,927
Users	476	678

- 1 Introduction
- 2 File Updates (ATLAS/CMS)**
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)



File Updates

- An update happens if:
 - File size (open) larger than 0 bytes
 - Number of written bytes larger than 0
 - Write ratio smaller than 1

Parameter	ATLAS Jan	ATLAS Nov	CMS
All write entries	3.3 Million	2.9 Million	1.8 Million
Updating entries	2081	13424	0
-Without reopens	86	65	0
-Of which VertexMon	1987	10390	0
-Of which log files	0	3032	0
-Of which other	94	2	0

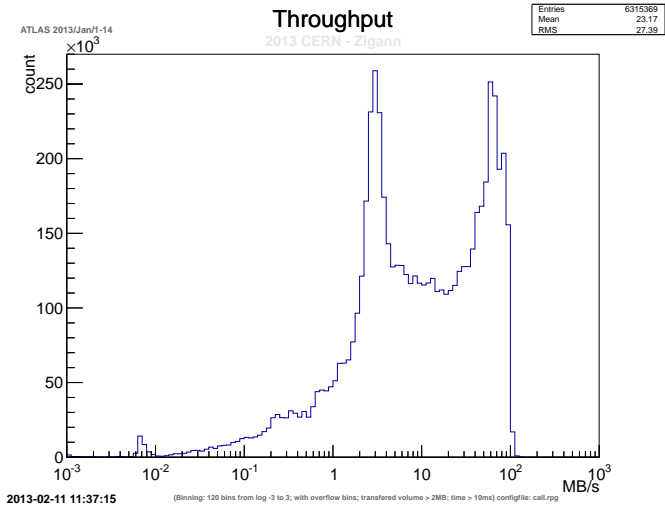
Do we really need the possibility of updates?



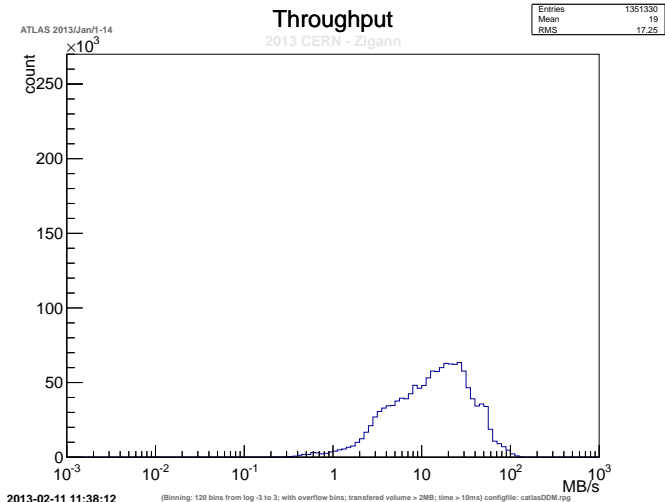
- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)**
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)



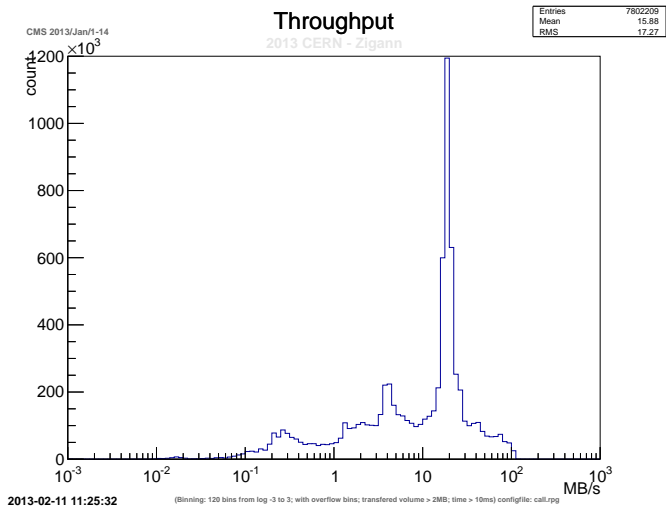
Throughput - ATLAS all



Throughput - ATLAS File Transfer Accounts (DDM)



Throughput - CMS all



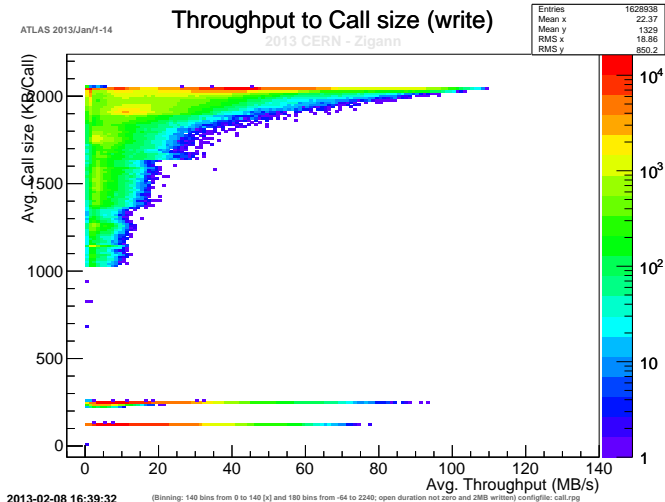
- Where do the peaks come from? Software related?
- Why did the 80MB/s peak in Atlas grow compared to previous measurements?
- Is there potential to increase throughput?



- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)**
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)

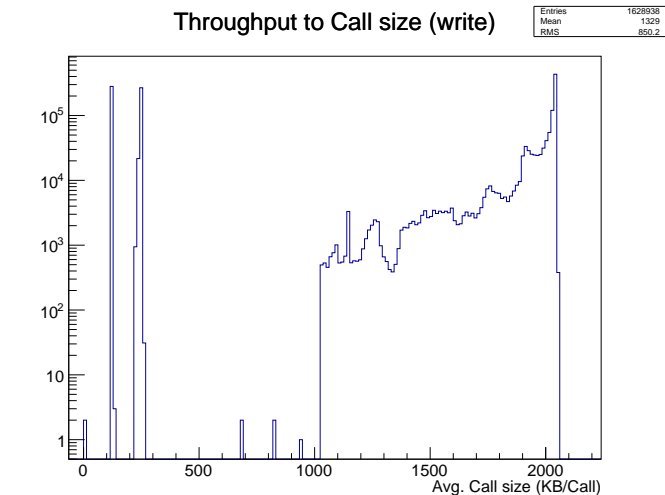


Throughput vs Call Size - Writing

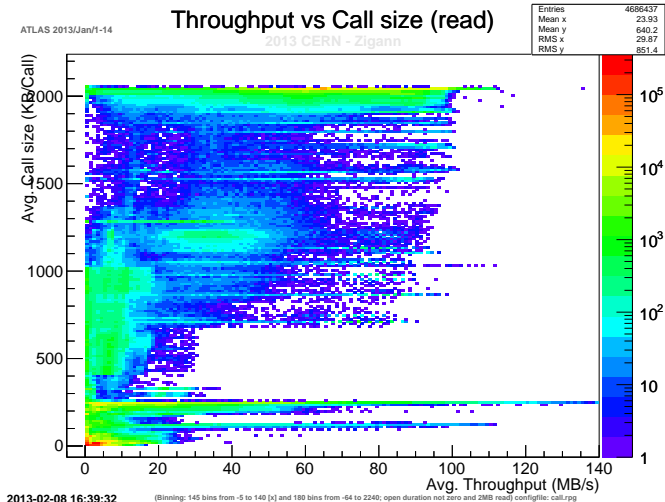


Call Size (Projection)

Throughput to Call size (write)

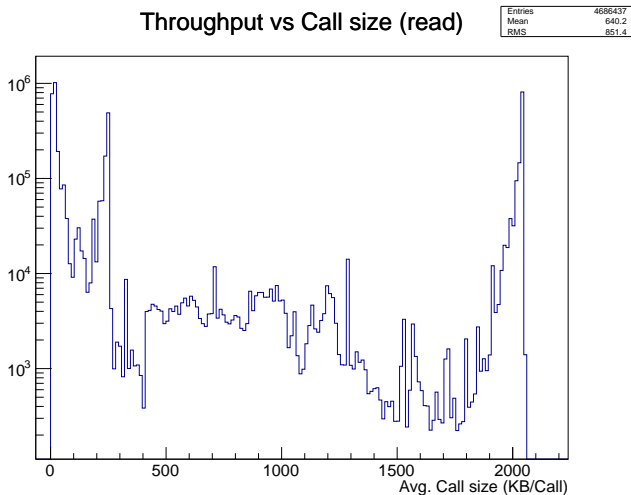


Throughput vs Call Size - Reading



Call Size (Projection)

Throughput vs Call size (read)



2013-02-08 16:39:32



Throughput vs Call Size - Questions

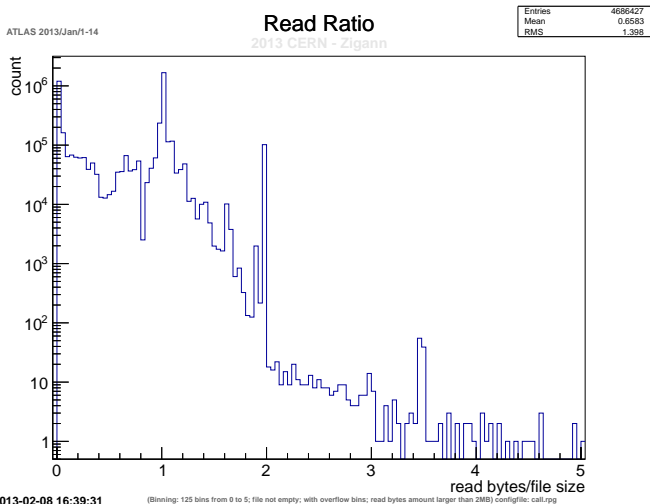
- 2 throughput peaks in 2MB/Call of ATLAS \Rightarrow use case? just CPU bound vs I/O bound?
- Can the average throughput be increased with more sync write access or larger buffer size?



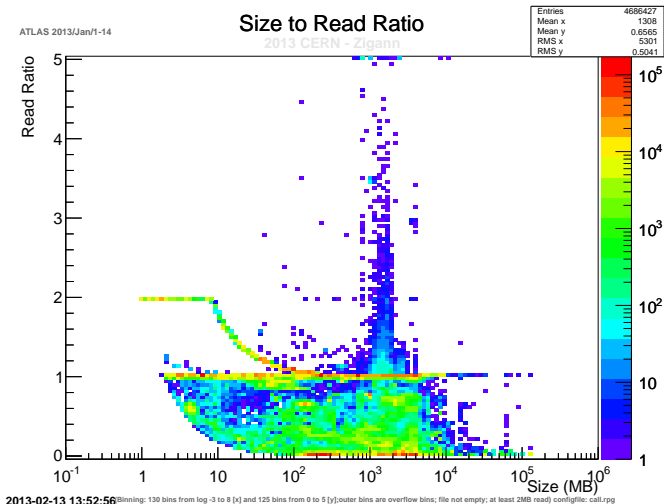
- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)**
- 6 Read Ratio (CMS)



Read Ratio



Read Ratio Vs Size



- Useful use case for 2 GB large files that are read more than 3 times?



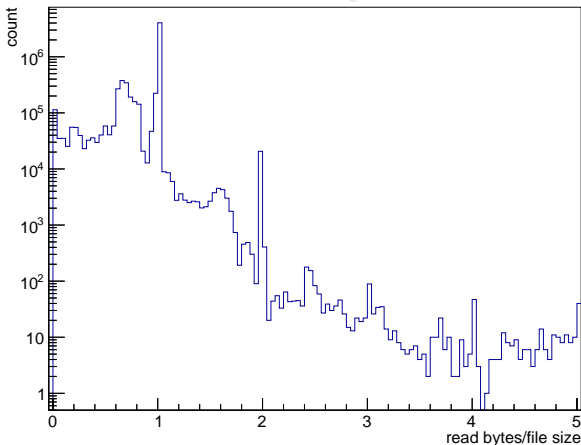
- 1 Introduction
- 2 File Updates (ATLAS/CMS)
- 3 Throughput (ATLAS/CMS)
- 4 Throughput vs Call Size (ATLAS)
- 5 Read Ratio (ATLAS)
- 6 Read Ratio (CMS)**



CMS 2013/Jan/1-14

Read Ratio 2013 CERN - Zigann

Entries	6592900
Mean	0.8661
RMS	1.629

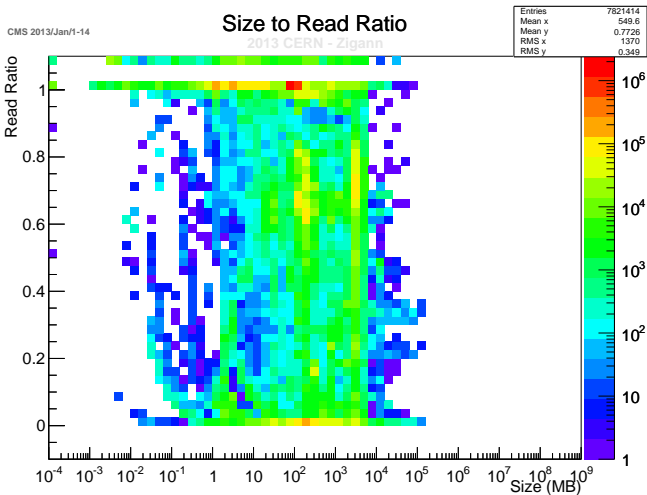


2013-02-08 16:56:05

(Binning: 125 bins from 0 to 5; file not empty; with overflow bins; read bytes amount larger than 2MB) configfile: call.rpg



Read Ratio Vs Size



- Are there known use cases for the peak around a read ratio of 0.7?



Thanks!

E-Mail

philipp.zigann@cern.ch

TWiki

<https://twiki.cern.ch/twiki/bin/view/IT/ZigannMeetingIOPattern>

