



Huawei Cloud Storage - Initial benchmark results

Maitane Zotes Resines, CERN IT

Openlab Minor Review Meeting

27. March 2012

CERN, Geneva

- Huawei Setup
- Introduction to S3
- Benchmark
- Upload Tests
- Download Tests
- Summary
- Future plans



- 3 Racks
- OBS and SOD boxes
- Total size of 768TB
- 3 copies redundancy
- 2 x 10GB network connection
- Total file size configurable



DSS

Huawei setup




CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it



- Simple Storage Service
 - Simple web services interface
 - Store and retrieve data
 - At anytime and anywhere
- Accounts and buckets
- Multiple datacenter support
- Aims
 - Scalability
 - High availability
 - Commodity costs
 - High aggregate performance



- Python Script
- Amazon AWS Python Library (awspylib)
- Operations
 - Upload
 - Download
 - Download range
- Multithread
- Multibucket



```
downloadFile=file10M.txt-
uploadFile=/tmp/file10M.txt
operation=downloadfile
container=cern.download.10mb
containerSize=50
url=olhw-s3.cern.ch
user=
passwd=
outputFile=download-1bucket-10M-Thread20
threads=20
fileSize=4000
dataSize=2000
```




UPLOADS



Threads	1	4	8	12	16	20
Files/Sec	7.3 ± 0.5	25.1 ± 1.2	33.9 ± 4.4	36.4 ± 8.1	36.7 ± 8.9	42.9 ± 8.7
MB/Sec	0,03 ± 0.002	0,09 ± 0.005	0,13 ± 0.02	0,14 ± 0.03	0,14 ± 0.035	0,17 ± 0.03
Threads	24	40	70	100	150	200
Files/Sec	43 ± 7.1	47.3 ± 6.7	49.6 ± 3.5	50.9 ± 4.8	52.3 ± 6.3	51.2 ± 3.4
MB/Sec	0,17 ± 0.03	0,18 ± 0.02	0,19 ± 0.01	0,19 ± 0.02	0,2 ± 0.02	0,19 ± 0.01

Threads	1	4	8	12	16	20
Files/Sec	7.3 ± 0.5	25.1 ± 1.2	33.9 ± 4.4	36.4 ± 8.1	36.7 ± 8.9	42.9 ± 8.7
MB/Sec	0,03 ± 0.002	0,09 ± 0.005	0,13 ± 0.02	0,14 ± 0.03	0,14 ± 0.035	0,17 ± 0.03
Threads	24	40	70	100	150	200
Files/Sec	43 ± 7.1	47.3 ± 6.7	49.6 ± 3.5	50.9 ± 4.8	52.3 ± 6.3	51.2 ± 3.4
MB/Sec	0,17 ± 0.03	0,18 ± 0.02	0,19 ± 0.01	0,19 ± 0.02	0,2 ± 0.02	0,19 ± 0.01

Threads	1	4	8	12	16	20
Files/Sec	7.3 ± 0.5	25.1 ± 1.2	33.9 ± 4.4	36.4		2.9 ± 8.7
MB/Sec	0,03 ± 0.002	0,09 ± 0.005	0,13 ± 0.02	0,14 ±		± 0.03
Threads	24	40	70	100	150	200
Files/Sec	43 ± 7.1	47.3 ± 6.7	49.6 ± 3.5	50.9 ± 4.8	52.3 ± 6.3	51.2 ± 3.4
MB/Sec	0,17 ± 0.03	0,18 ± 0.02	0,19 ± 0.01	0,19 ± 0.02	0,2 ± 0.02	0,19 ± 0.01





Threads	1	4	8	12	16	20
Files/Sec	1.7 ± 0.04	6.1 ± 0.1	9.1 ± 0.3	10.3 ± 0.1	10.7 ± 0.09	10.8 ± 0.3
MB/Sec	17 ± 0.4	61 ± 1	91 ± 3	103 ± 1	107 ± 0.9	108 ± 3
Threads	24	40	70	100	150	200
Files/Sec	10.9 ± 0.3	11 ± 0.07	11.1 ± 0.07	10.9 ± 0.07		
MB/Sec	108 ± 3	110 ± 0.7	111 ± 0.7	109 ± 0.7		

Threads	1	4	8	12	16	20
Files/Sec	1.7 ± 0.04	6.1 ± 0.1	9.1 ± 0.3	10.3 ± 0.1	10.7 ± 0.09	10.8 ± 0.3
MB/Sec	17 ± 0.4	61 ± 1	91 ± 3	103 ± 1	107 ± 0.9	108 ± 3
Threads	24	40	70	100	150	200
Files/Sec	10.9 ± 0.3	11 ± 0.07	11.1 ± 0.07	10.9 ± 0.07		
MB/Sec	108 ± 3	110 ± 0.7	111 ± 0.7	109 ± 0.7		

**CLIENT
MEMORY
CONSUMPTION**


Threads	1	4	8	12	16	20
Files/Sec	1.7 ± 0.04	6.1 ± 0.1	9.1 ± 0.3	10.3 ± 0.1	10.7 ± 0.09	10.8 ± 0.3
MB/Sec	17 ± 0.4	61 ± 1	91 ± 3	103 ± 1	107 ± 0.9	108 ± 3
Threads	24	40	70	100	150	200
Files/Sec	10.9 ± 0.3	11 ± 0.07	11.1 ± 0.07	10.9 ± 0.07	10.9 ± 0.07	10.9 ± 0.07
MB/Sec	108 ± 3	110 ± 0.7	111 ± 0.7	109 ± 0.7	109 ± 0.7	109 ± 0.7




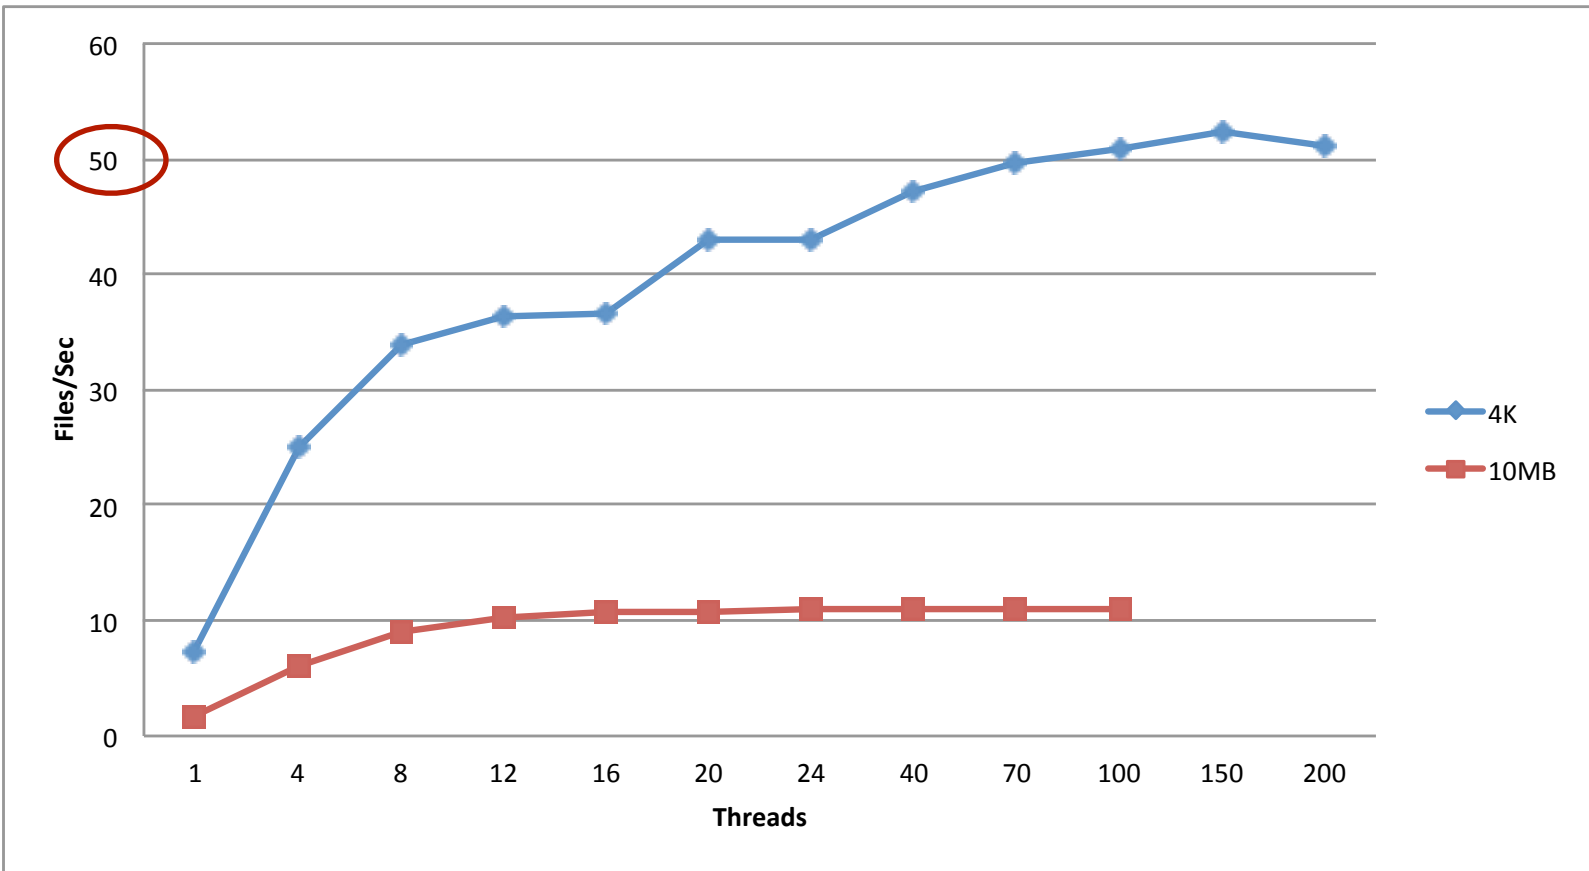
Threads	1	4	8	12	16	20
Files/Sec	1.7 ± 0.04	6.1 ± 0.1	9.1 ± 0.3	10.3 ± 0.1	10.7 ± 0.09	10.8 ± 0.3
MB/Sec	17 ± 0.4	61 ± 1	91 ± 3	103 ± 1	107 ± 0.9	108 ± 3
Th		40	70	100	150	
Files		± 0.07	11.1 ± 0.07	10.9 ± 0.07		
MB/Sec	108 ± 3	110 ± 0.7	111 ± 0.7	109 ± 0.7		



CLIENT BANDWIDTH



CLIENT MEMORY CONSUMPTION


























CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it







DOWNLOADS



Threads	1	4	8	12	16	20
Files/Sec	51.7 ± 7.2	186.8 ± 9.7	358.2 ± 14.7	423.5 ± 6.6		
MB/Sec	0,2 ± 0.02	0,72 ± 0.03	1,39 ± 0.05	1,66 ± 0.006		
Threads	24	40	70	100	150	200
Files/Sec						
MB/Sec						

Threads	1	4	8	12	16	20
Files/Sec	51.7 ± 7.2	186.8 ± 9.7	358.2 ± 14.7	423.5 ± 6.6		
MB/Sec	0,2 ± 0.02	0,72 ± 0.03	1,39 ± 0.05	1,66 ± 0.006		
Threads	24	40	70	100	150	200
Files/Sec						
MB/Sec						

**CLIENT
SOCKET
CONSUMPTION**

Threads	1	4	8	12	16	20
Files/Sec	51.7 ± 7.2	186.8 ± 9.7	358.2 ± 14.7	423.5 ± 6.6		
MB/Sec	0,2 ± 0.02	0,72 ± 0.03	1,39 ± 0.05	1,66 ± 0.006		
Threads	24	40	70	100	150	200
Files/Sec						
MB/Sec						



**CLIENT
SOCKET
CONSUMPTION**

Threads	1	4	8	12	16	20
Files/Sec	51.7 ± 7.2	186.8 ± 9.7	358.2 ± 14.7	423.5 ± 6.6		
MB/Sec	0,2 ± 0.02	0,72 ± 0.03	1,39 ± 0.05	1,66 ± 0.006		
Threads	24	40	70	100	150	200
Files/Sec						
MB/Sec						

620 FILES/SEC

**CLIENT
SOCKET
CONSUMPTION**


Threads	1	4	8	12	16	20
Files/Sec	2,2 ± 0.16	7,7 ± 0.18	8,7 ± 0.47	9,5 ± 0.03	9,4 ± 0.06	9,3 ± 0.06
MB/Sec	22 ± 1.6	77 ± 1.8	87 ± 4.7	95 ± 0.3	94 ± 0.6	93 ± 0.6
Threads	24	40	70	100	150	200
Files/Sec	9.28 ± 0.07	9,8 ± 0.06	9,7 ± 0.07	9,7 ± 0.03		
MB/Sec	96 ± 0.7	98 ± 0.6	97 ± 0.7	97 ± 0.3		

Threads	1	4	8	12	16	20
Files/Sec	2,2 ± 0.16	7,7 ± 0.18	8,7 ± 0.47	9,5 ± 0.03	9,4 ± 0.06	9,3 ± 0.06
MB/Sec	22 ± 1.6	77 ± 1.8	87 ± 4.7	95 ± 0.3	94 ± 0.6	93 ± 0.6
Threads	24	40	70	100	150	200
Files/Sec	9.28 ± 0.07	9,8 ± 0.06	9,7 ± 0.07	9,7 ± 0.03		
MB/Sec	96 ± 0.7	98 ± 0.6	97 ± 0.7	97 ± 0.3		



Threads	1	4	8	12	16	20
Files/Sec	2,2 ± 0.16	7,7 ± 0.18	8,7 ± 0.47	9,5 ± 0.03	9,4 ± 0.06	9,3 ± 0.06
MB/Sec	22 ± 1.6	77 ± 1.8	87 ± 4.7	95 ± 0.3	94 ± 0.6	93 ± 0.6
Threads	24	40	70	100	150	200
Files/Sec	9.28 ± 0.07	9,8 ± 0.06	9,7 ± 0.07	9,7 ± 0.03		
MB/Sec	96 ± 0.7	98 ± 0.6	97 ± 0.7	97 ± 0.3		


**CLIENT
MEMORY
CONSUMPTION**





Threads	1	4	8	12	16	20
Files/Sec	2,2 ± 0.16	7,7 ± 0.18	8,7 ± 0.47	9,5 ± 0.03	9,4 ± 0.06	9,3 ± 0.06
MB/Sec	22 ± 1.6	77 ± 1.8	87 ± 4.7	95 ± 0.3	94 ± 0.6	93 ± 0.6
Threads	40	70	100	150	200	
Files/Sec			9,7 ± 0.07	9,7 ± 0.03		
MB/Sec	96 ± 0.7	98 ± 0.6	97 ± 0.7	97 ± 0.3		



CLIENT BANDWIDTH



CLIENT MEMORY CONSUMPTION




























CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

20



- Upload limit of 50 Hz (per bucket)
 - 118Hz (5 buckets)
- Client limits
 - Download limit of 425 Hz
 - Not useful to raise number of buckets
- Issues encountered
 - Huawei issues
 - Bucket performance
 - Benchmark issues
 - Delayed closing of client sockets produces socket shortage
 - Client bandwidth limit
 - Client memory limit

- Client with better resources
- Multibox
- Multibyte range reads
- ROOT
- Multiclient
- Bigger files





Huawei Cloud Storage - Initial benchmark results

Maitane Zotes Resines, CERN IT

Openlab Minor Review Meeting

27. March 2012

CERN, Geneva