A solution for BigData integrate with Ceph

inwinSTACK / Chief Architect
Thor Chin

Agenda

- Why we need Ceph?
- Why we need to integrate Ceph with Big Data?
- How to integrate Ceph with Big Data?
- What's Alluxio?
- Alluxio User List
- Architecture for Ceph integrate with Big Data
- Example

inwinSTACK inc. All rights reserved

Why we need Ceph?

- Distributed storage system
  - Fault tolerant , no SPoF
- X86 Commodity hardware
  - Saving you costs, giving you flexibility
- Large scale – Incremental expansion
  - 10s to 1000s of nodes
- Unified storage platform
  - Scalable object , Block , File system.
- Open source – No vender lock-in
- Automatically balance the file system
- Data security
  - with 2 or more copies in different physical store media.

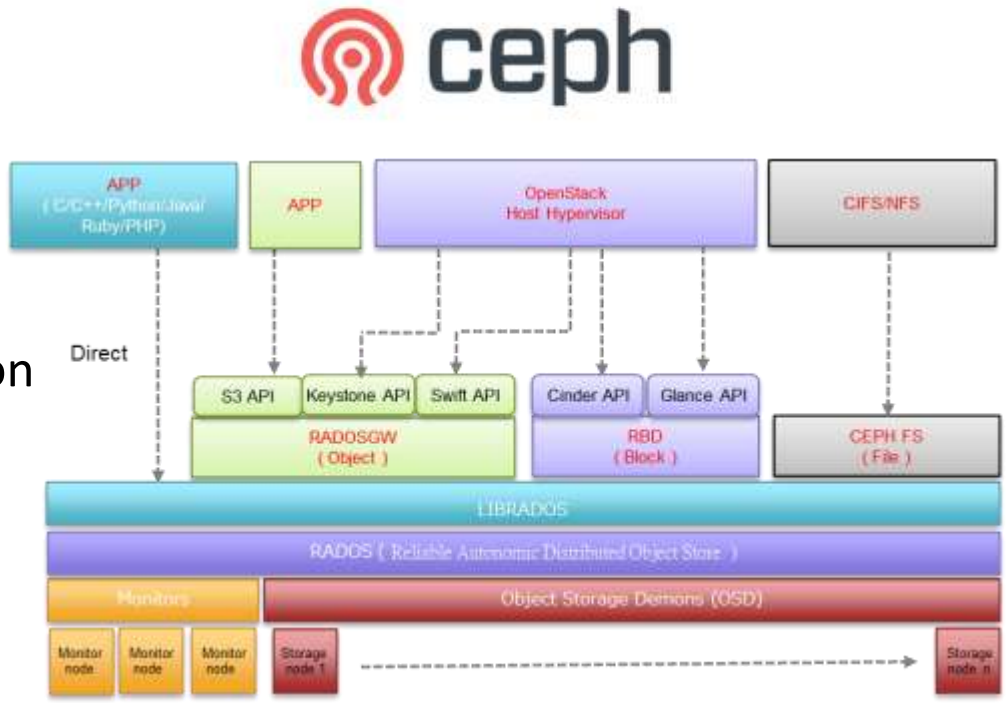

inwinSTACK inc. All rights reserved

Why we need to integrate Ceph with Big Data?

- Separate Application and Data
- Better Architecture to solve Hadoop Name Node HA Issue
- Scalable storage solution
- High Performance

inwinSTACK inc. All rights reserved

How to integrate Ceph with Big Data?

Through an Interface

Alluxio: A Virtual Distributed Storage System

inwinSTACK inc. All rights reserved

What's Alluxio

Unified Access	Thriving Community	High Performance
Alluxio is the unified access point for your data. Applications only need to connect with Alluxio to access data in all your under storage systems.	Start from UC Berkeley Open Source Project – Tachyon since 2013 Apr. released 16 versions and rename to Alluxio release 1.0 version in 2016 Feb.	Alluxio's memory-centric architecture enables data access orders of magnitudes faster than existing solutions.
	Fastest growing open source project in big data history with more than 200 contributors from over 50 organizations.	

inwinSTACK inc. All rights reserved

Alluxio User List

- Alibaba
- Alluxio
- Arimo
- Atigeo
- Baidu
- Barclays
- Huawei
- IBM Research
- Intel
- Neusoft
- Qunar
- Radicalbit
- UC Berkeley AMPLab
- Ultra Tendency

Ref : http://www.alluxio.org/documentation/en/Powered-By-Alluxio.html

inwinSTACK inc. All rights reserved

Architecture for Ceph integrate with Big Data

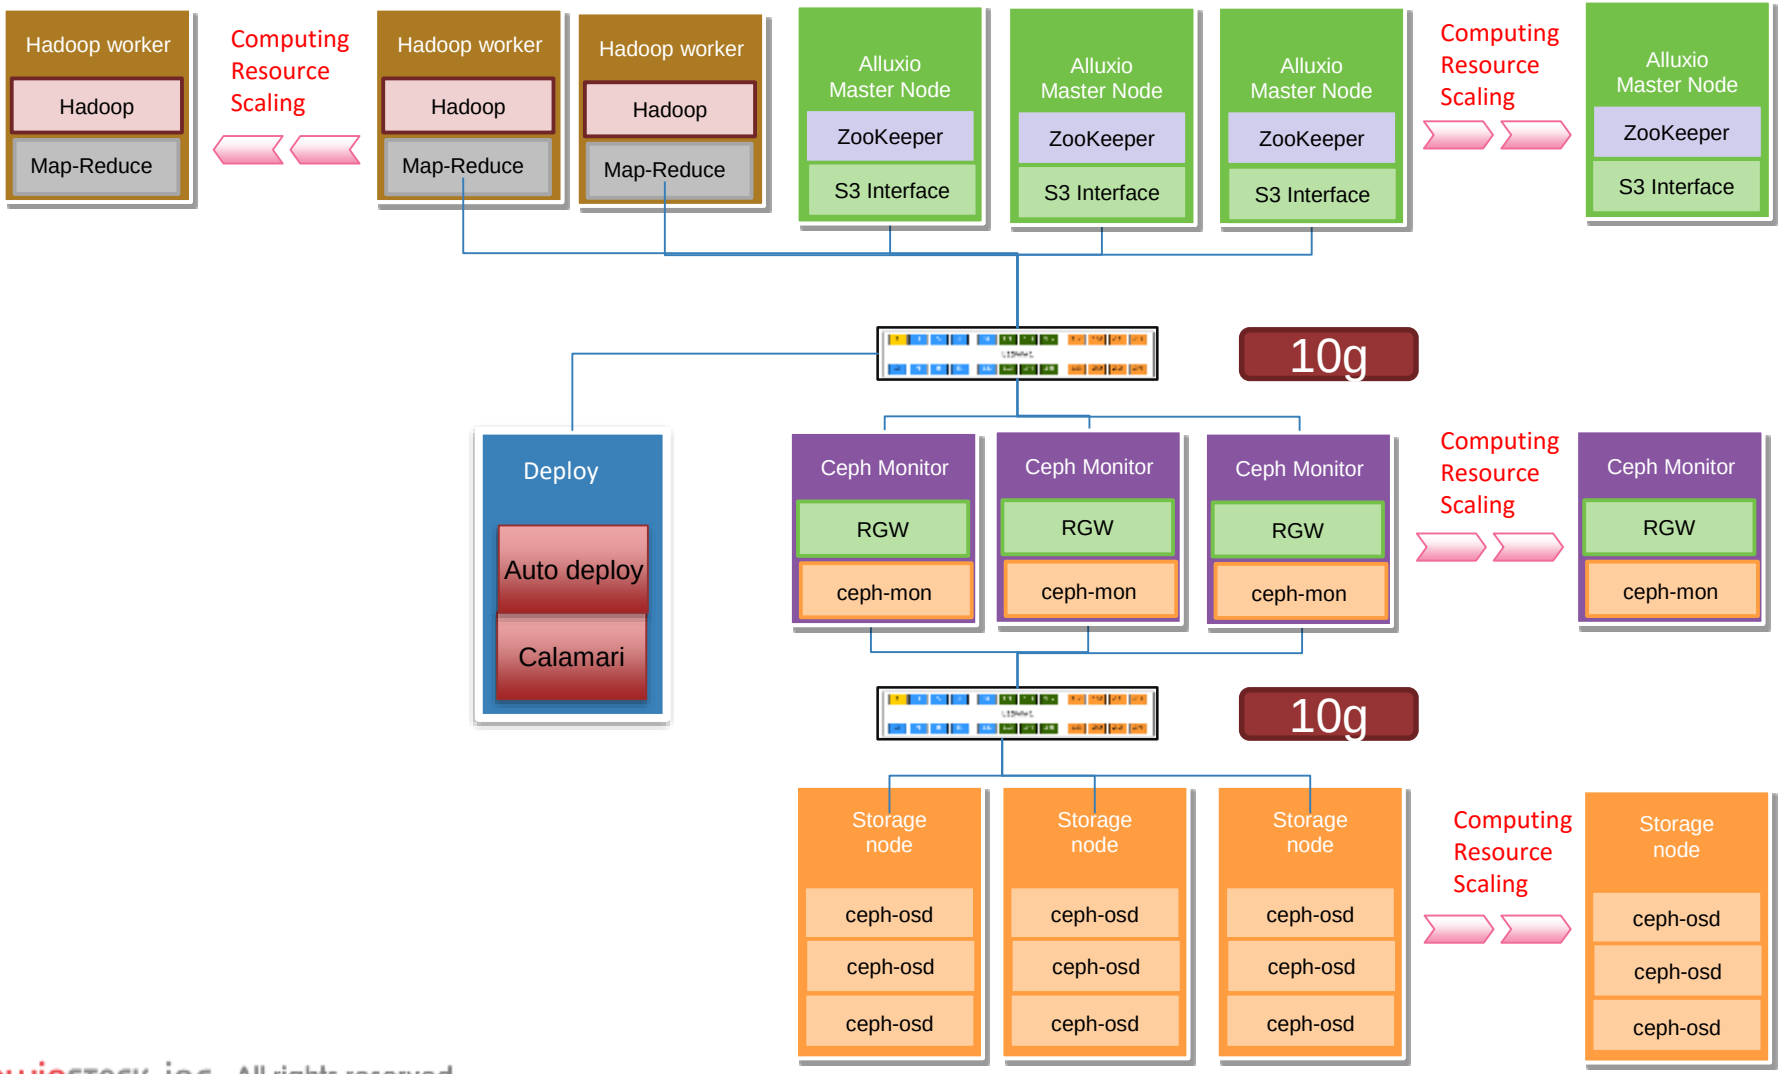

inwinSTACK inc. All rights reserved

Why we choose RGW rather than CephFS?

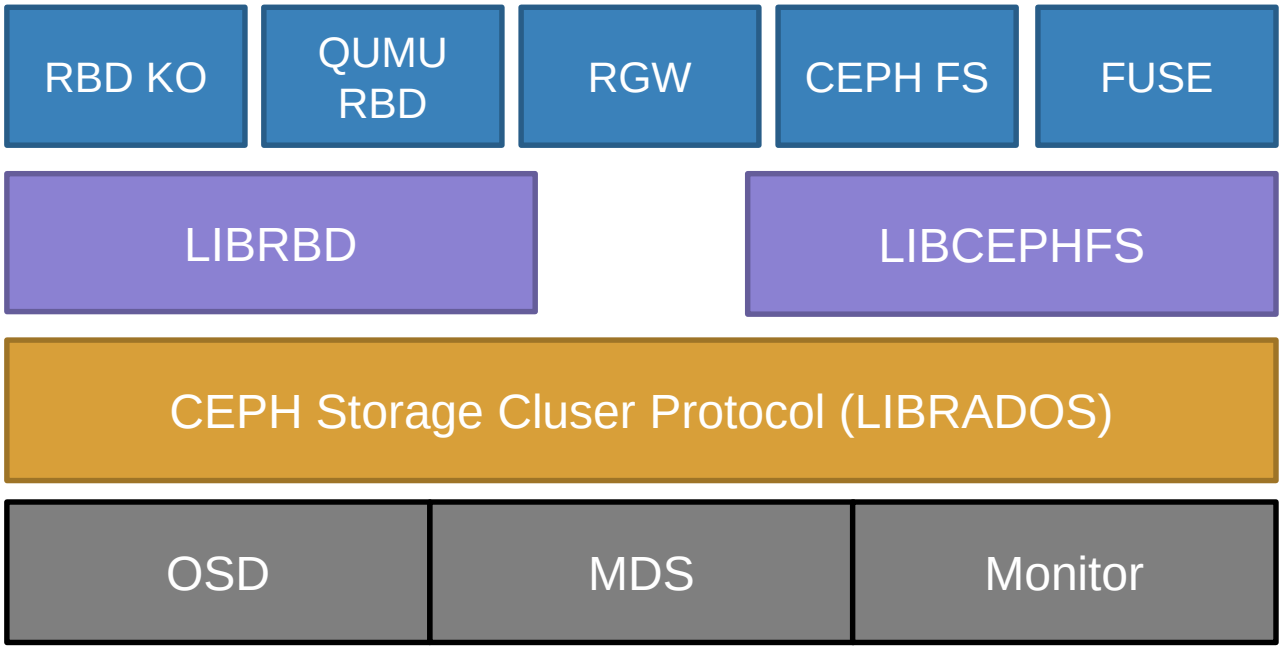

inwinSTACK inc. All rights reserved

Example

inwinSTACK inc. All rights reserved

Alluxio Overview

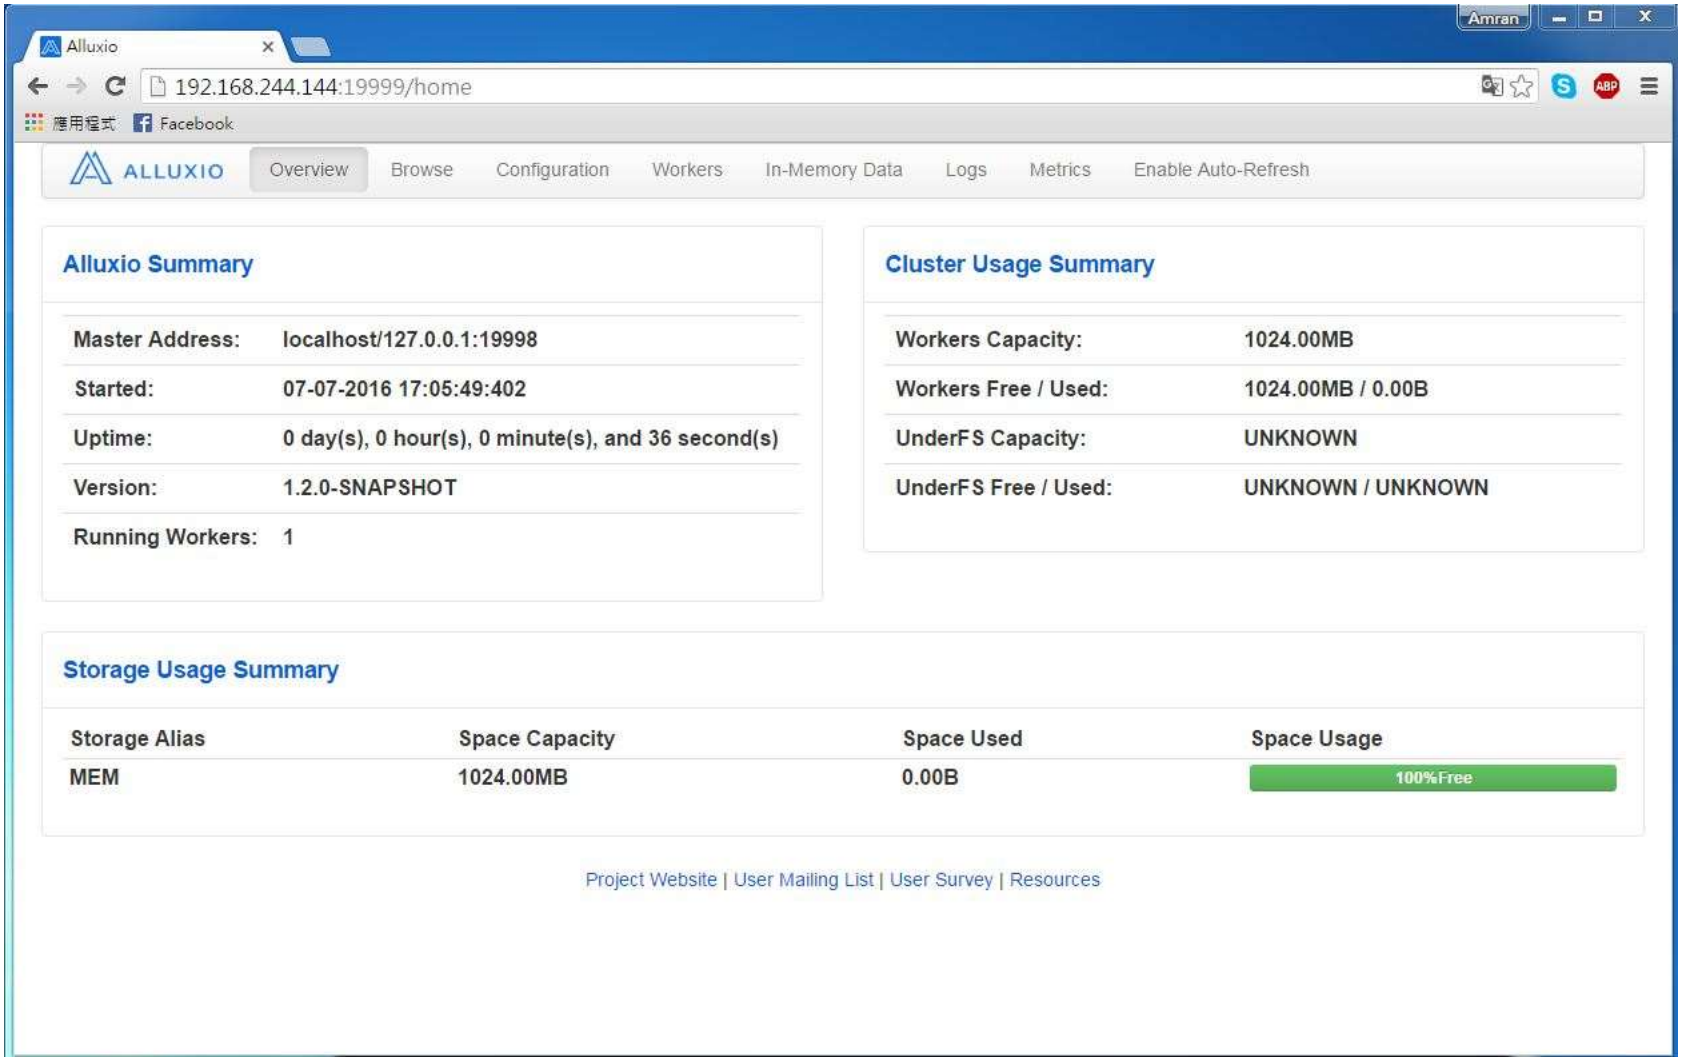

inwinSTACK inc. All rights reserved

Enable Auto Refresh

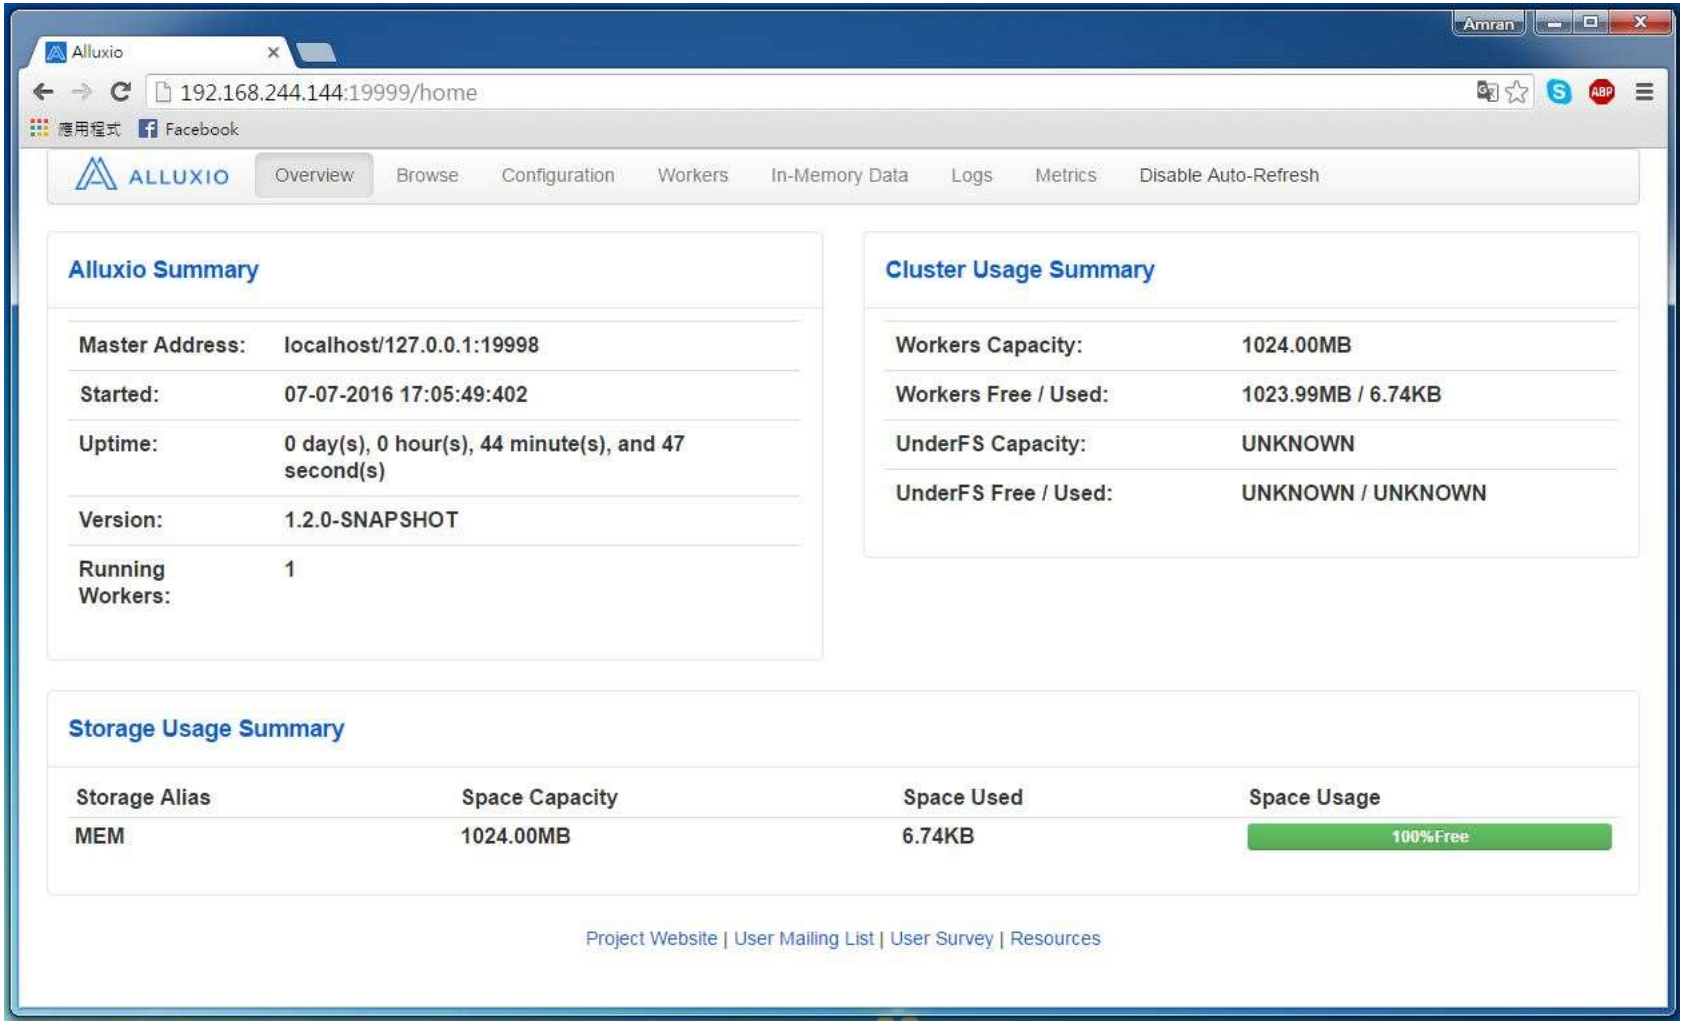

inwinSTACK inc. All rights reserved

System Configuration

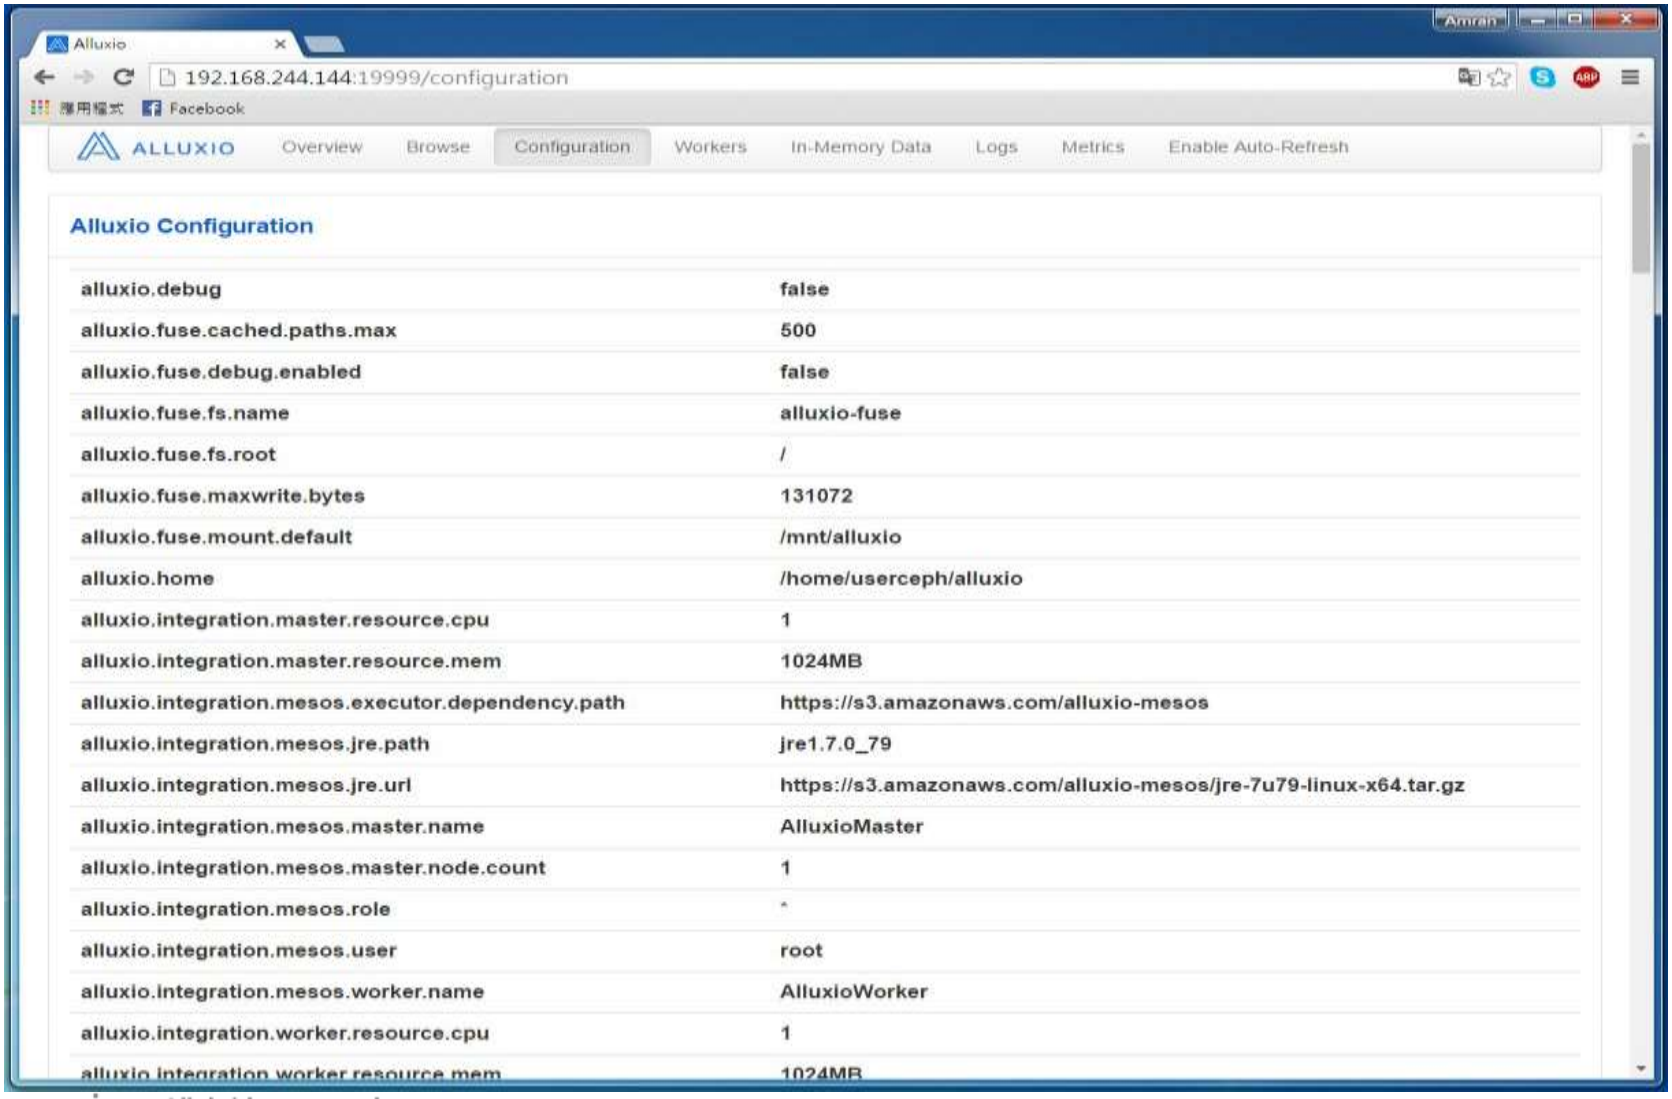

inwinSTACK inc. All rights reserved

Workers

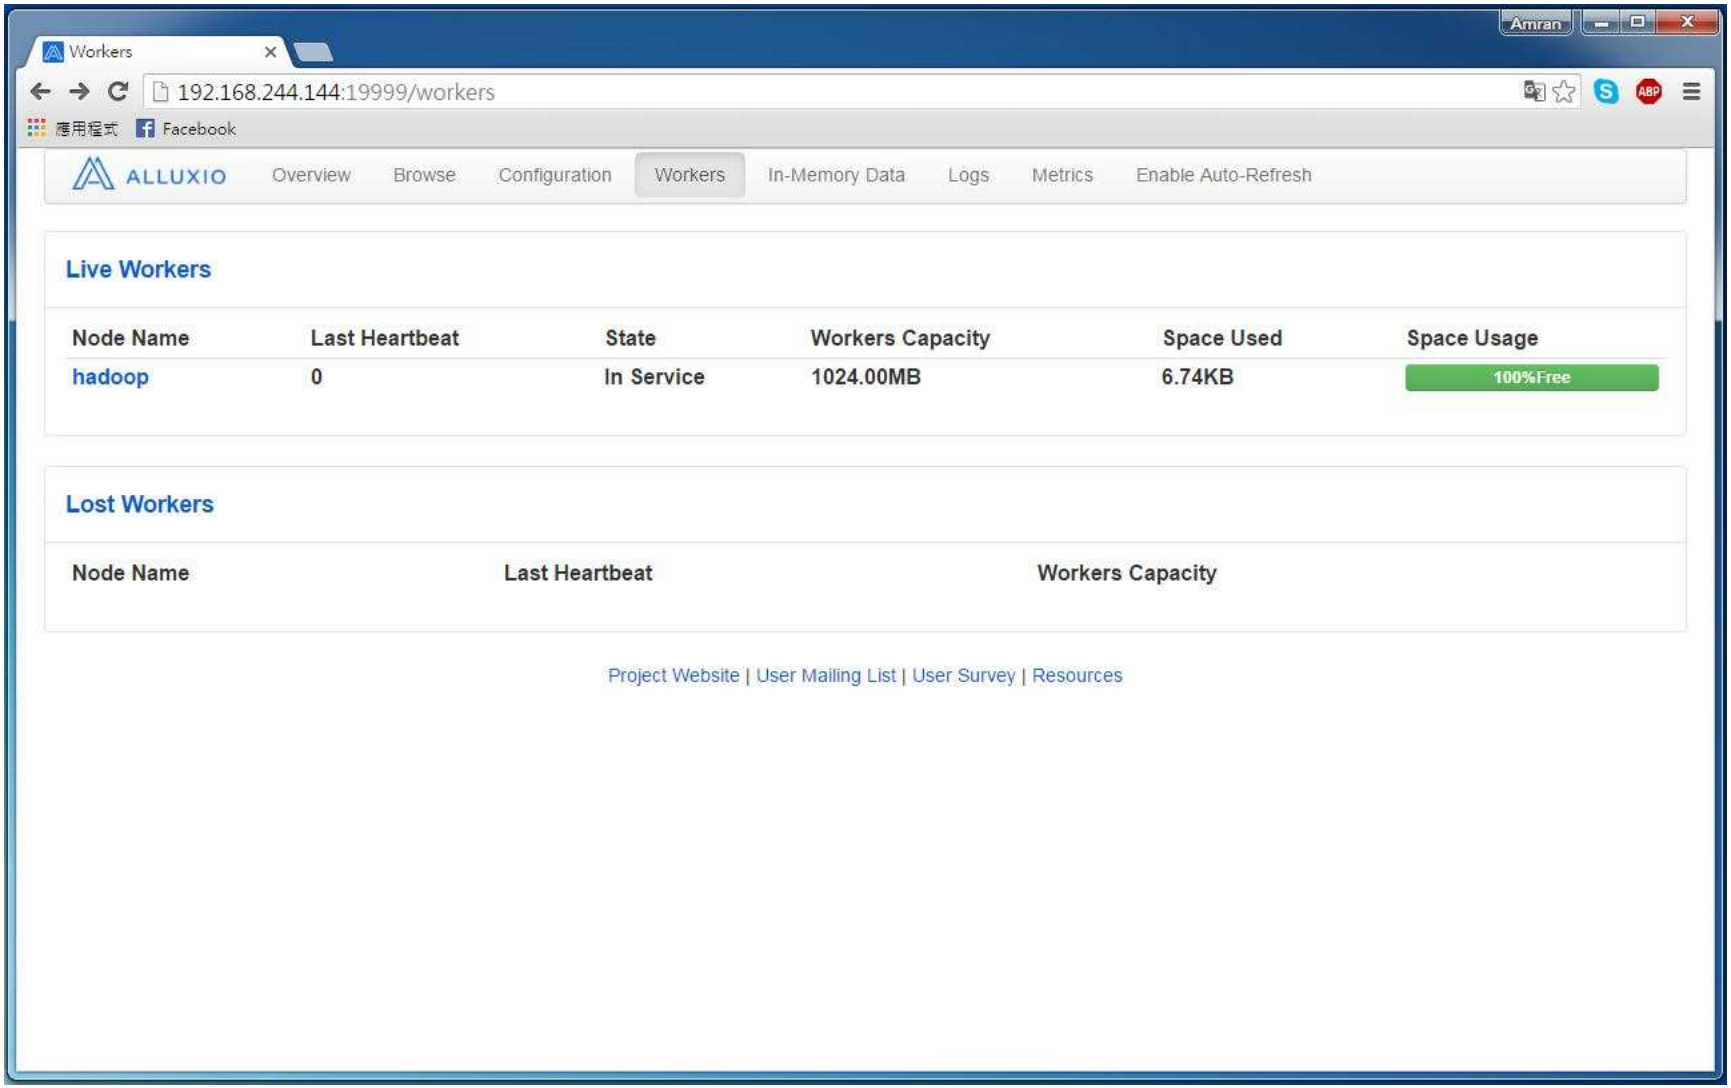

inwinSTACK inc. All rights reserved

Browse File System

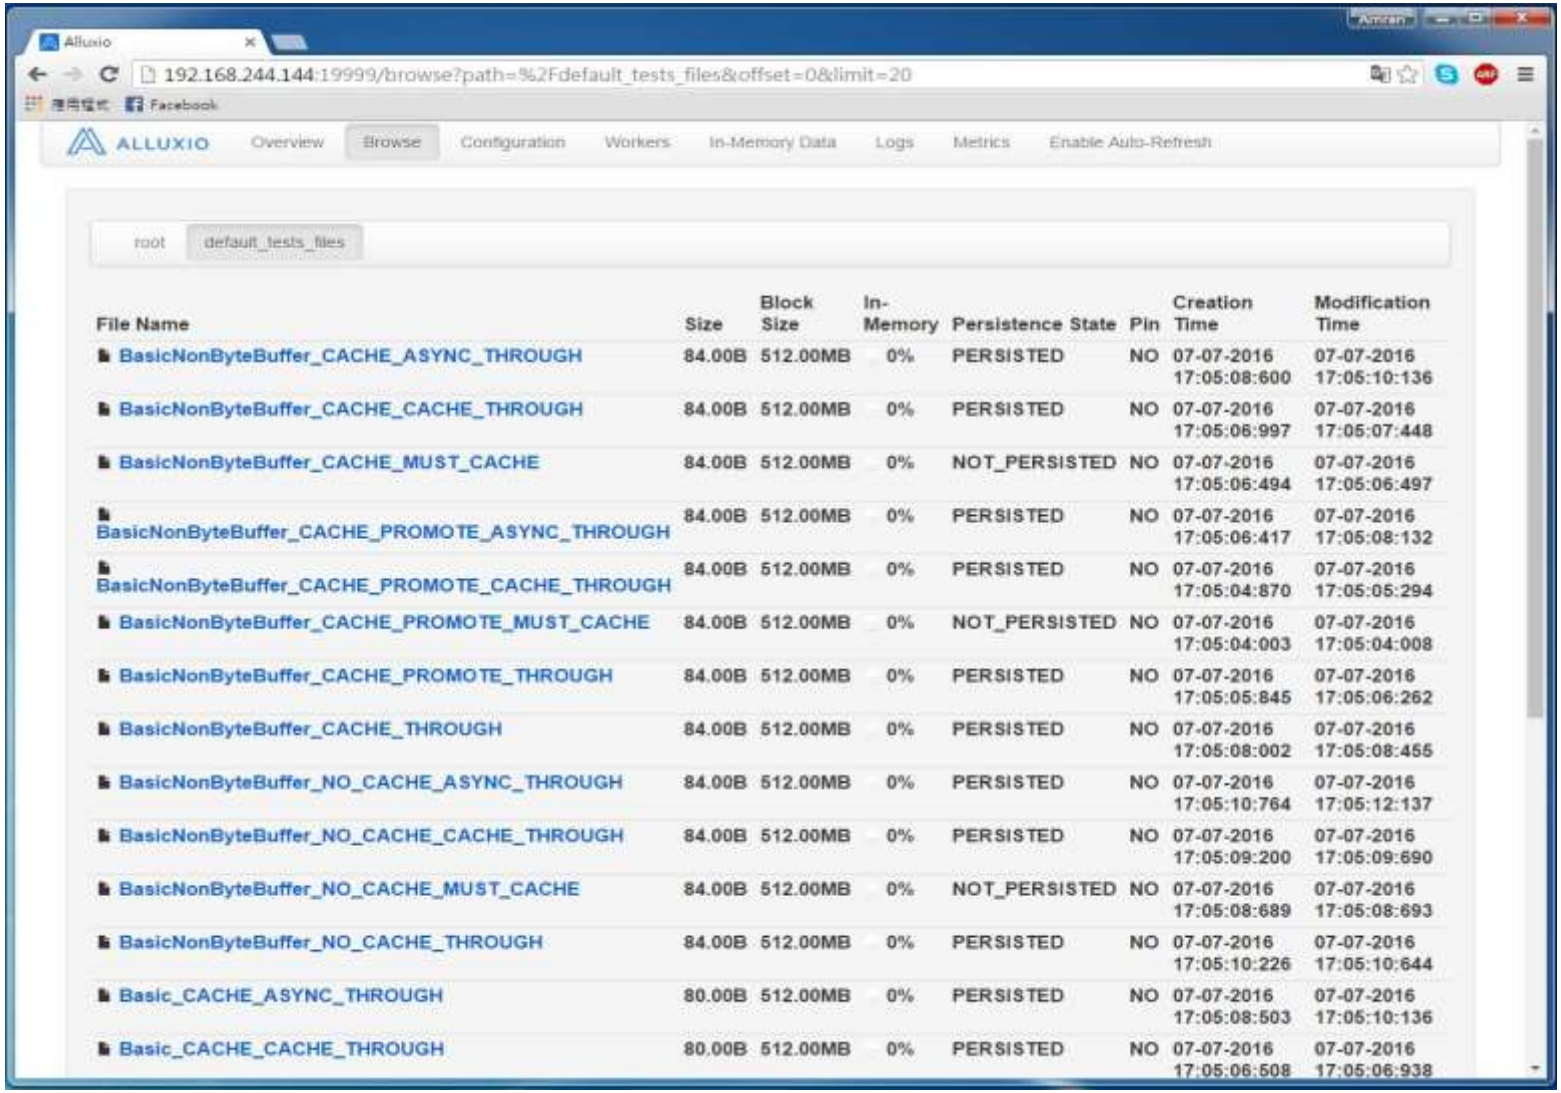

File Name	Size	Block Size	In-Memory	Persistence State	Pin	Creation Time	Modification Time
BasicNonByteBuffer_CACHE_ASYNC_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:08:600	07-07-2016 17:05:10:136
BasicNonByteBuffer_CACHE_CACHE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:06:997	07-07-2016 17:05:07:448
BasicNonByteBuffer_CACHE_MUST_CACHE	84.00B	512.00MB	0%	NOT_PERSISTED	NO	07-07-2016 17:05:06:494	07-07-2016 17:05:06:497
BasicNonByteBuffer_CACHE_PROMOTE_ASYNC_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:06:417	07-07-2016 17:05:08:132
BasicNonByteBuffer_CACHE_PROMOTE_CACHE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:04:870	07-07-2016 17:05:05:294
BasicNonByteBuffer_CACHE_PROMOTE_MUST_CACHE	84.00B	512.00MB	0%	NOT_PERSISTED	NO	07-07-2016 17:05:04:003	07-07-2016 17:05:04:008
BasicNonByteBuffer_CACHE_PROMOTE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:05:845	07-07-2016 17:05:06:262
BasicNonByteBuffer_CACHE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:08:002	07-07-2016 17:05:08:455
BasicNonByteBuffer_NO_CACHE_ASYNC_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:10:764	07-07-2016 17:05:12:137
BasicNonByteBuffer_NO_CACHE_CACHE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:09:200	07-07-2016 17:05:09:690
BasicNonByteBuffer_NO_CACHE_MUST_CACHE	84.00B	512.00MB	0%	NOT_PERSISTED	NO	07-07-2016 17:05:08:689	07-07-2016 17:05:08:693
BasicNonByteBuffer_NO_CACHE_THROUGH	84.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:10:226	07-07-2016 17:05:10:644
Basic_CACHE_ASYNC_THROUGH	80.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:08:503	07-07-2016 17:05:10:136
Basic_CACHE_CACHE_THROUGH	80.00B	512.00MB	0%	PERSISTED	NO	07-07-2016 17:05:06:508	07-07-2016 17:05:06:938

inwinSTACK inc. All rights reserved

Add Wordcount for Testing

$ ./bin/alluxio fs copyFromLocal /home/hduser/myfile /wordcount/myfile

```
hduser@alhadoop:~/alluxio-1.0.1$ ./bin/alluxio fs copyFromLocal /home/hduser/myf
ile /wordcount/myfile
Copied /home/hduser/myfile to /wordcount/myfile
hduser@alhadoop:~/alluxio-1.0.1$
```

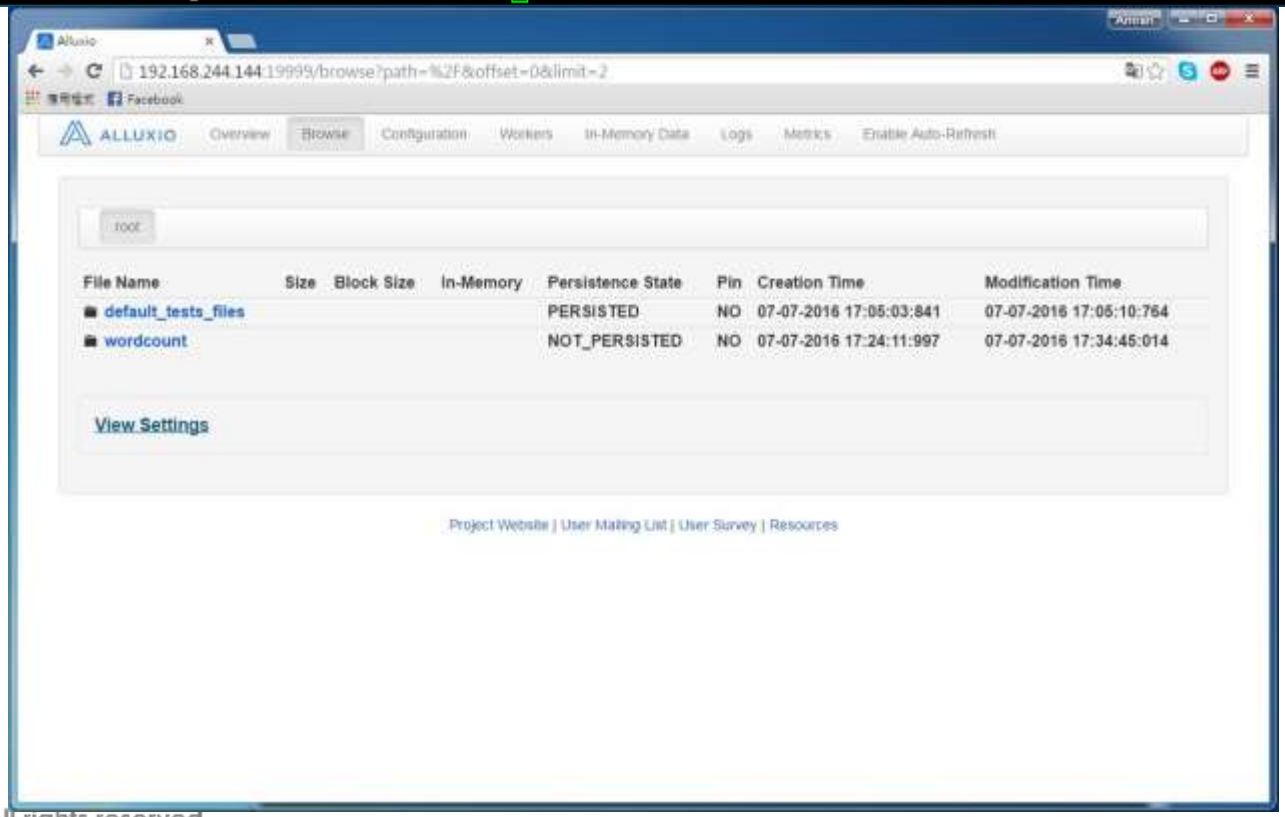

inwinSTACK inc. All rights reserved

Alluxio Wordcount output

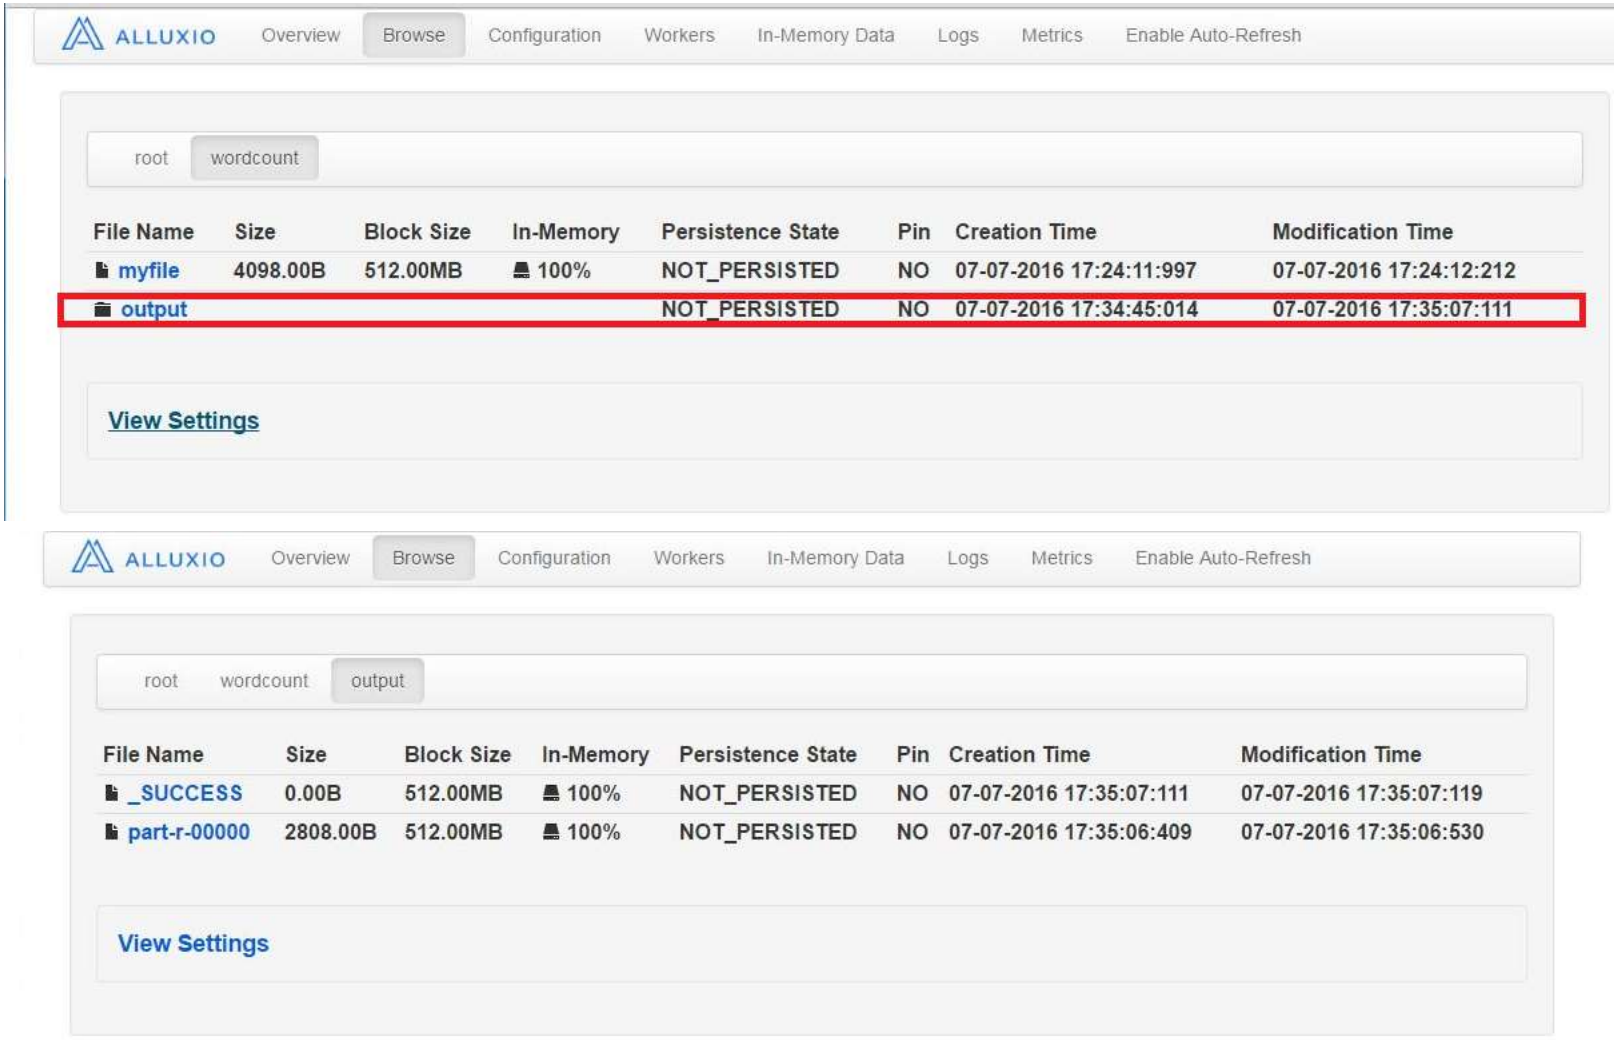

inwinSTACK Inc. All rights reserved

Alluxio Wordcount output (Cont.)

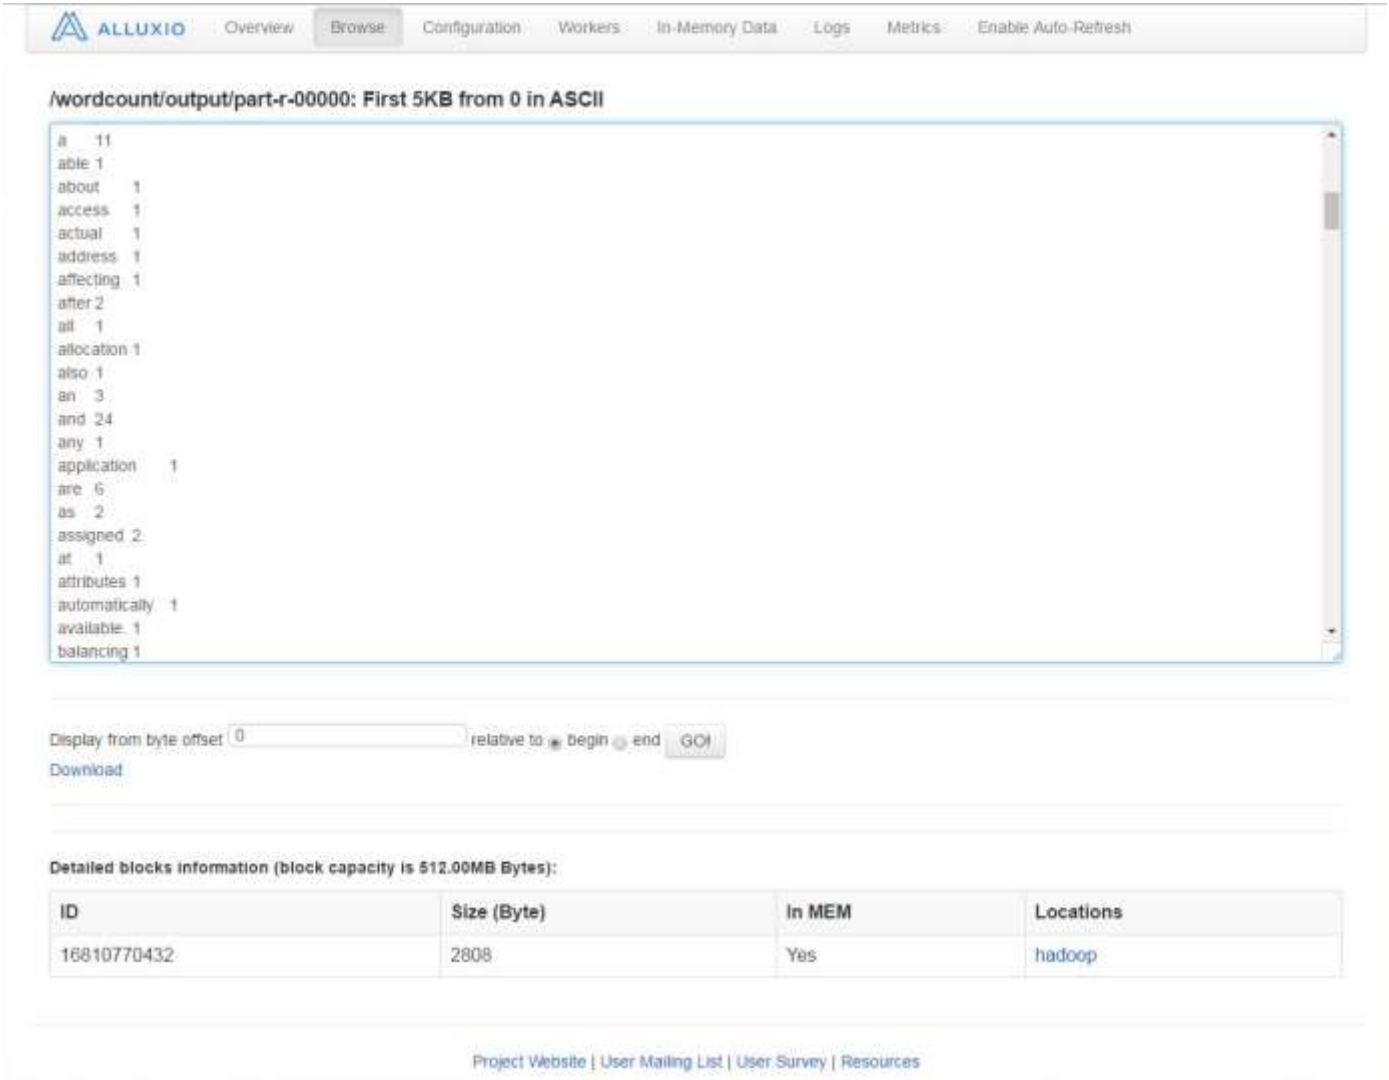

inwinSTACK inc. All rights reserved

In-memory Data

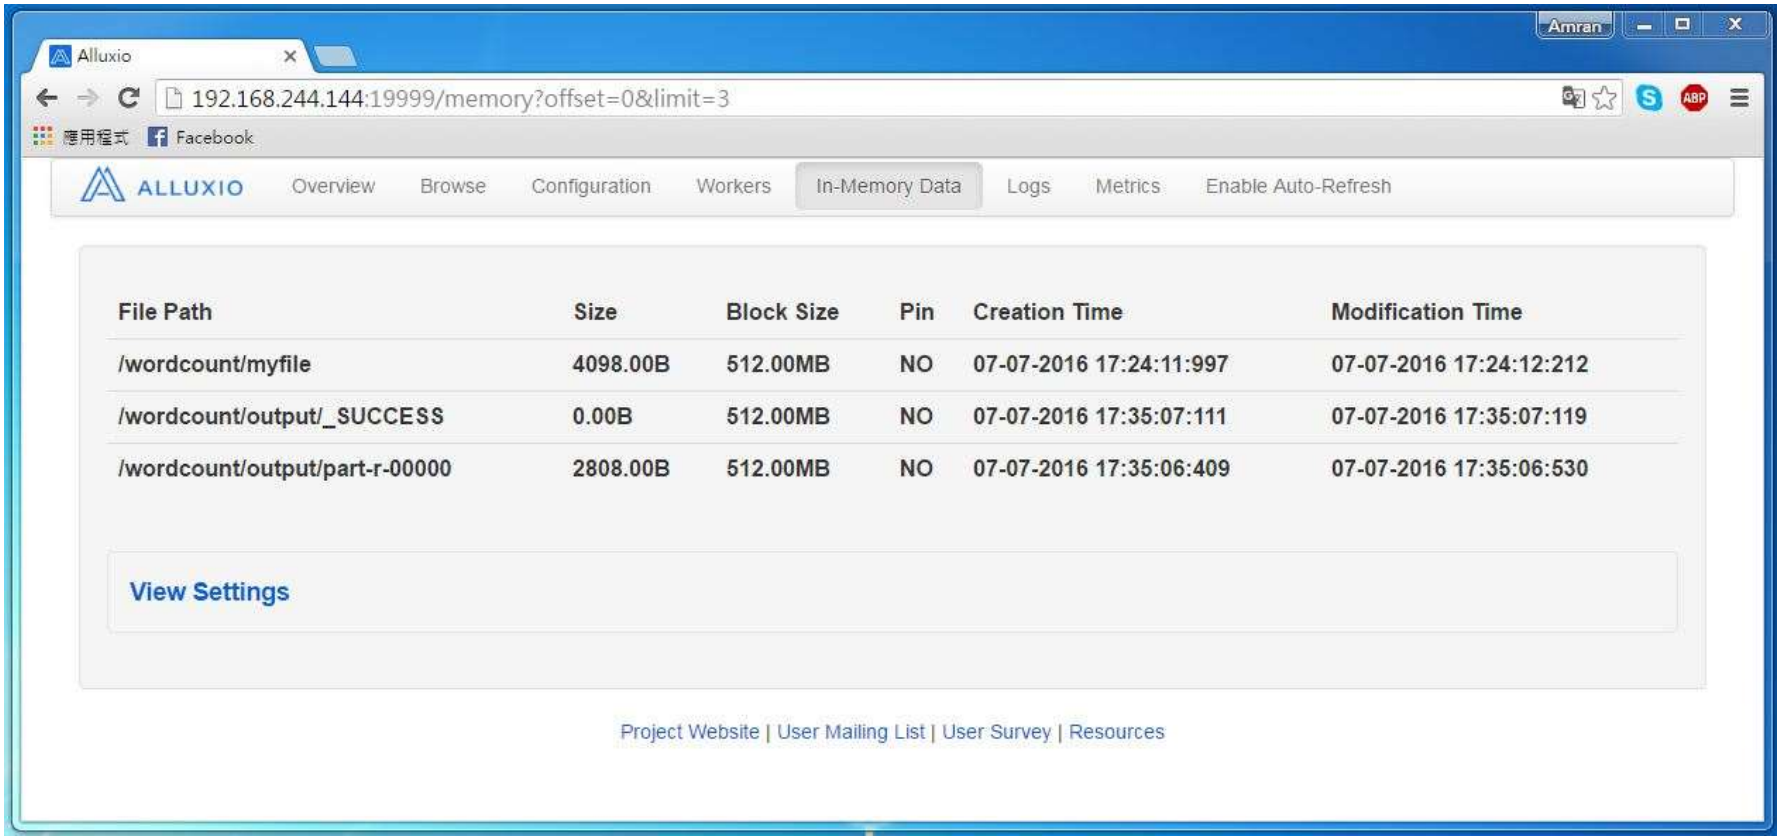

inwinSTACK inc. All rights reserved

Logs

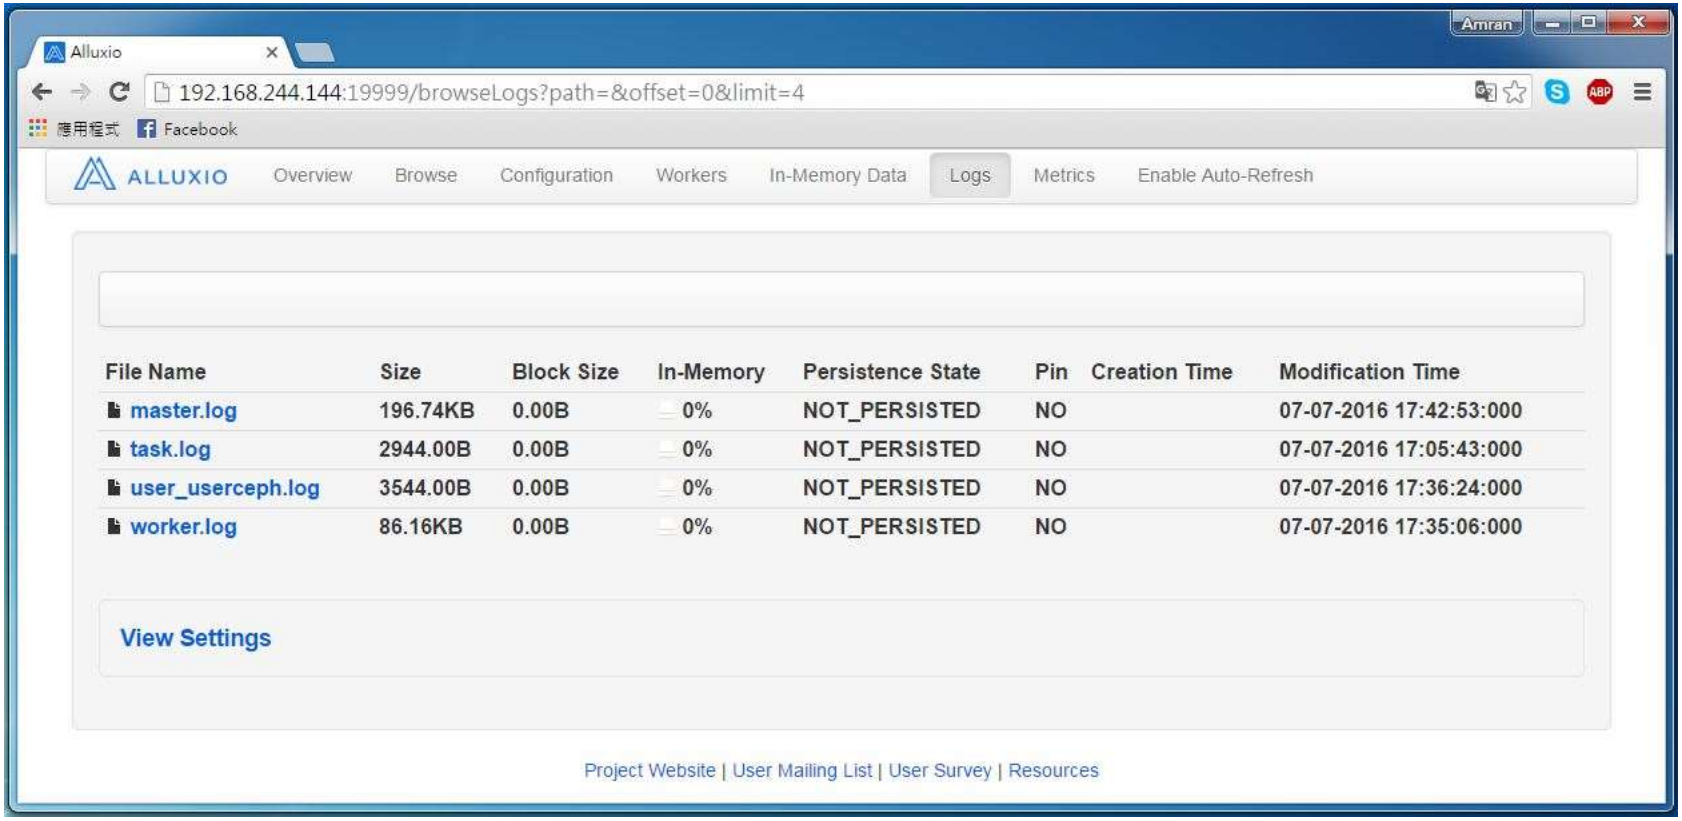

inwinSTACK inc. All rights reserved

Metrics

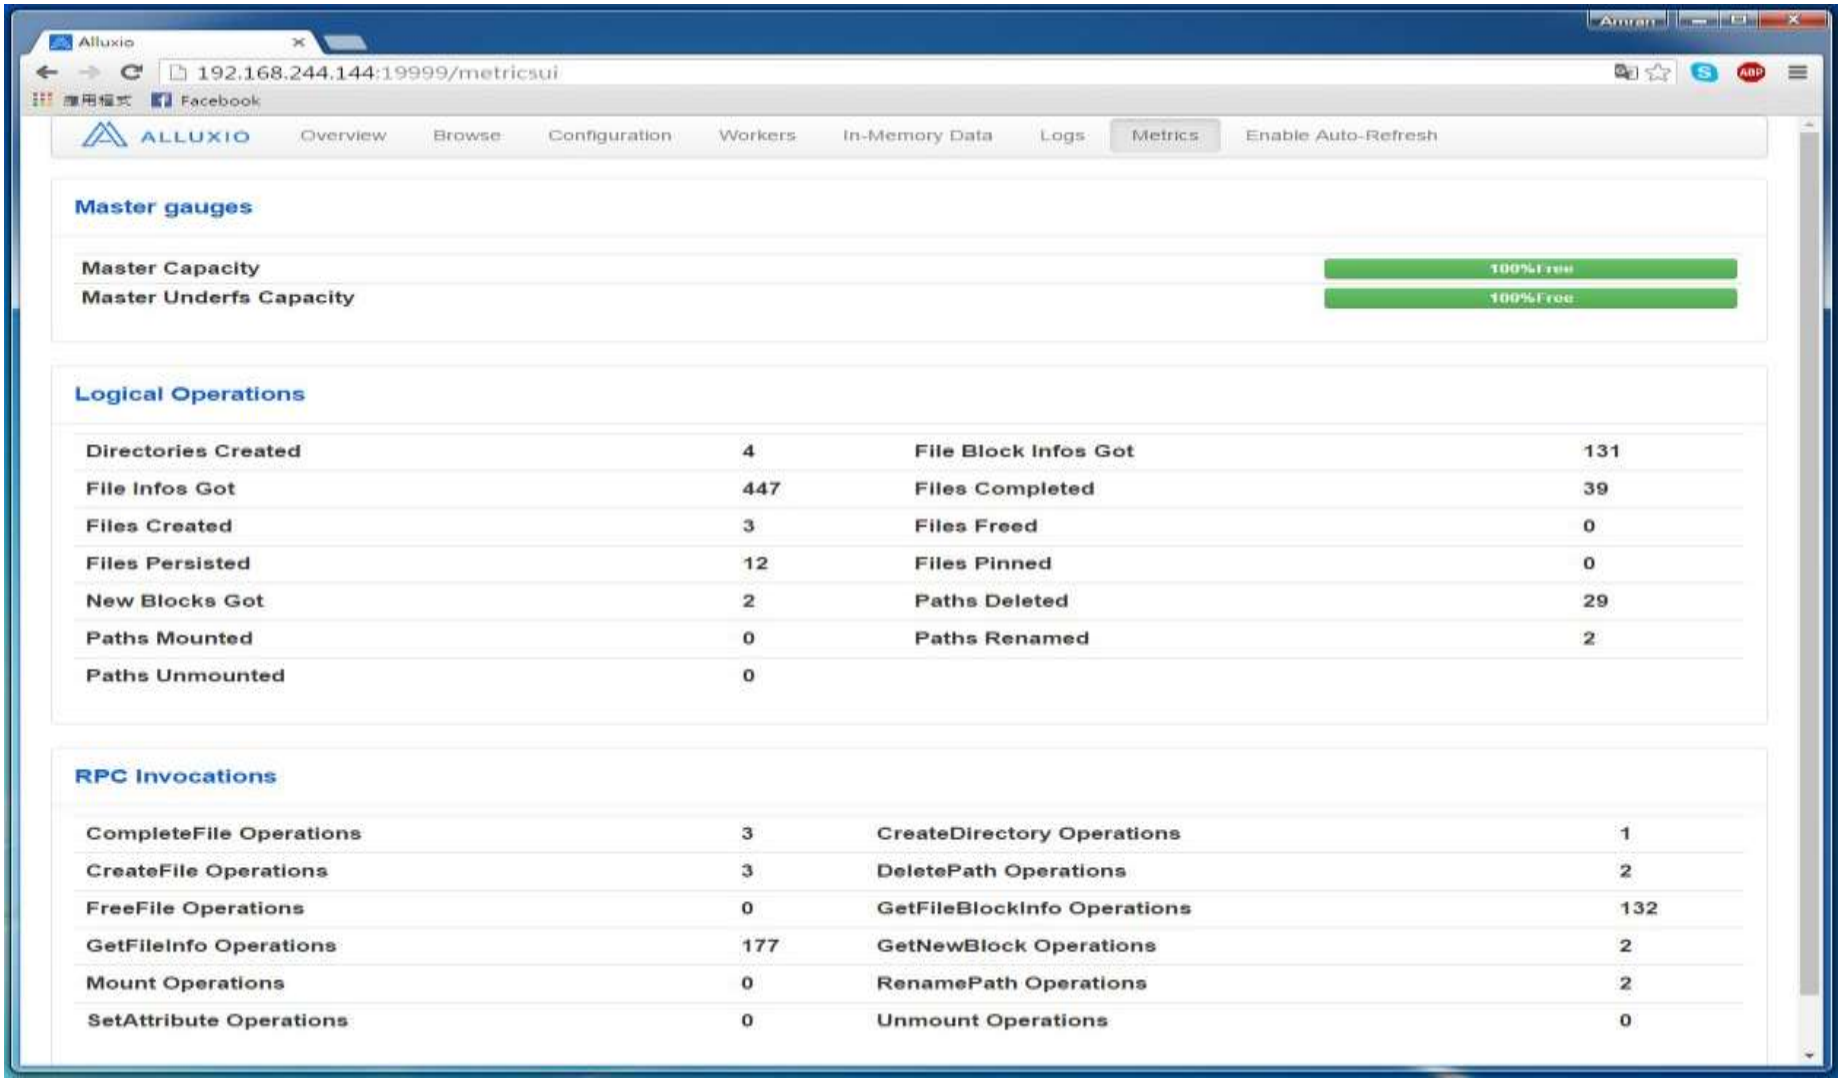

inwinSTACK inc. All rights reserved

Ceph APAC Road show 2016

8/24 集思台大會議中心

ceph days

TRANSFORM STORAGE IN PERSON

Join the expert Ceph team, Ceph's customers and partners, and the Ceph community as we discuss how Ceph, the massively scalable, open source, software-defined storage system, can radically improve the economics and management of data storage for your enterprise.

inwinSTACK inc. All rights reserved

迎棧科技股份有限公司

inwinSTACK
G.O.D. Great OpenStack Delivery

www.inwinstack.com