



# Monitoring an heterogenous Lustre environment

LUG 2015

[frederick.lefebvre@calculquebec.ca](mailto:frederick.lefebvre@calculquebec.ca)

# Plan

Site description

A bit of history

Real time utilization monitoring

Moving forward

# Who we are

Advanced computing center @Université  
Laval

Part of Calcul Québec & Compute Canada

Operate 2 Clusters & 4 Lustre fs

# Our infrastructure

1000+ lustre clients (2.5.3 & 2.1.6)

On 2 IB fabrics - 6 LNET routers

4 Lustre file systems (~2.5 PB)

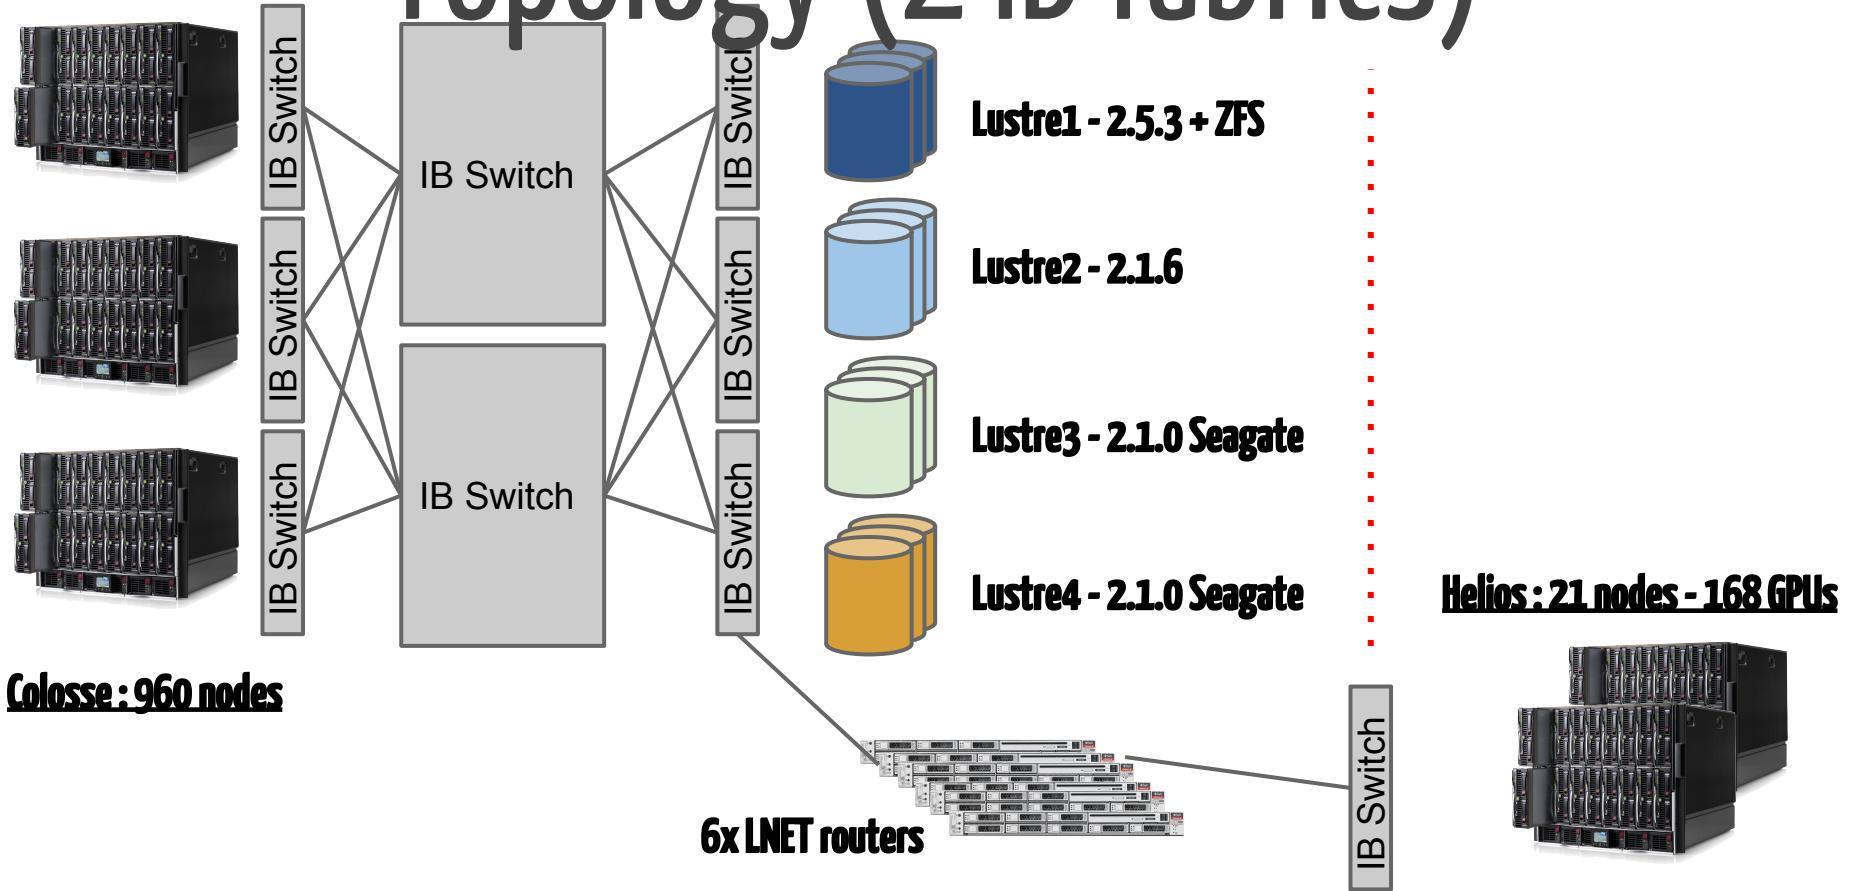
500 TB - Lustre 2.5.3 + ZFS - 16 OSSs

500 TB - Lustre 2.1.6 - 6 OSSs

1.4 PB - Lustre 2.1.0 (Seagate) - 12 OSS

120 TB - Lustre 2.1.0 (Seagate) - 2 OSS

# Topology (2 IB fabrics)



# An history of monitoring

We initially monitored because

We knew we should

We wanted to identify failures

Tried multiple options over the years

collectl

cerebro +lmt

# An history of monitoring (cont.)

Converged towards custom gmetrics scripts for Ganglia

Monitored IB  
Lustre IOPS and  
throughput



# An history of monitoring (cont.)

Converged towards custom gmetrics scripts for Ganglia

Monitored IB

Lustre IOPS and  
throughput

Worked reasonably well

But quality of our scripts was inconsistent

Some used a surprising amount of resources when running



# Issues

Users and staff keep asking questions

Why is it so long navigating my files ?

Why is my app so slow at writing results ?

Which filesystem should I use ?

We had most of the data “somewhere” but searching for  
it was always painful

# Issues (cont.)

## Data is everywhere

- Performance data from compute node is in mongoDB
- Lustre client and server data in Ganglia (RRD)
- Scheduler data in text log files

Easy to end up with a spiderweb of complex script to integrate all data

# New reqs

1. Give staff a near-realtime view of the filesystems utilisation and load
2. Provide end-users with data to identify their utilisation of the Lustre filesystems
  - a. Bonus point for cross-referencing with the scheduler's stats

# New reqs (cont.)

We want to keep “historical” data for future analysis and comparison

With the amount of storage now available, it didn’t make sense to keep losing granularity over time with RRD.

# Enter Splunk

We have been using Splunk to store and visualize system logs and energy efficiency data.

It's well suited to indexing a lot of unstructured data

# Plan

Collect data from /proc on all nodes and  
send it to the system's log with logger

The logs are indexed and visualized with  
Splunk

We use a separate index for performance data

# Data collection

Existing options to monitor Lustre tend to  
be tied to a specific backend (RRD or DB)

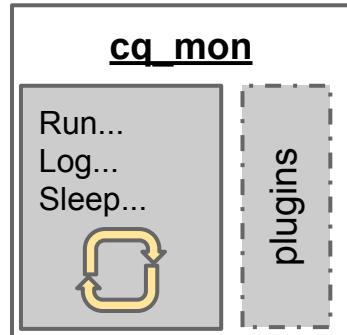
More flexibility is needed

We decided to write another one

We could have modified an existing application

# Overview

## On Lustre clients and servers



syslog

rsyslogd

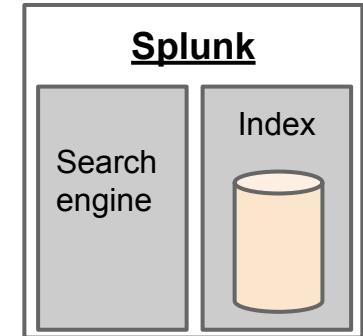
python-lustre  
python-psutil  
...

syslog

## On Splunk server



HTTPS



rsyslogd

write to file

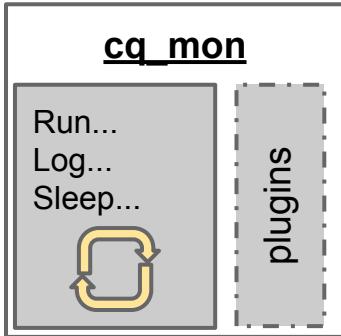
/var/log/...

# cq\_mon

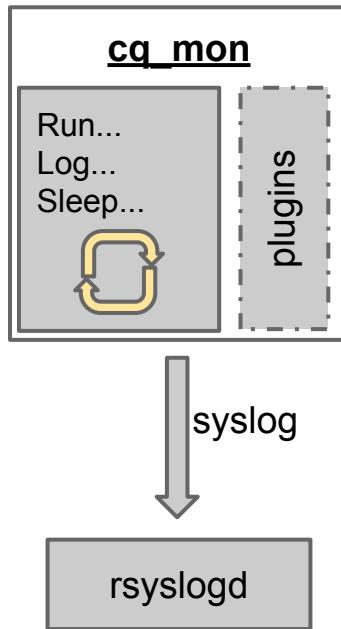
Written in Python

Configuration specifies which plugins to run

Runs each plugin at interval and log the results as ‘key=values’ to syslog



# cq\_mon



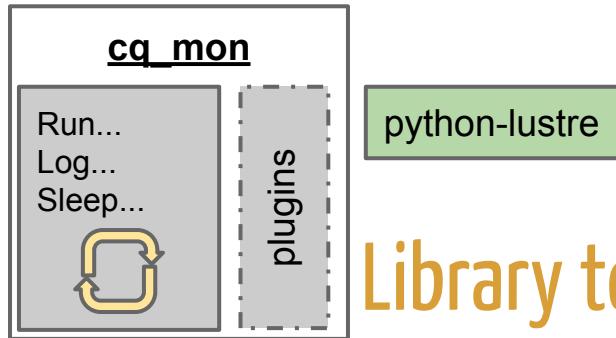
Written in Python

Configuration specifies which plugins to run

Runs each plugin at interval and log the results as 'key=values' to syslog

Add syslog filter so as not to 'pollute' system logs with performance data

# python-lustre\_util



Library to interact with lustre through /proc

Loads stats for different Lustre components make them available through an appropriate object (ie: lustre\_mdt, lustre\_ost, etc)

# python-lustre\_util (cont.)

## Work in progress

Tested with Lustre 2.1.6, 2.4.x, 2.5.x+ZFS and 2.1.0-xyratex

Available for download through Github: <https://github.com/calculquebec>

## Offer stats for:

MDS, OSS, MDT, OST

Client stats from the client and not the OSS/MDS

# Splunk dashboards

**Put together custom dashboards for our needs**

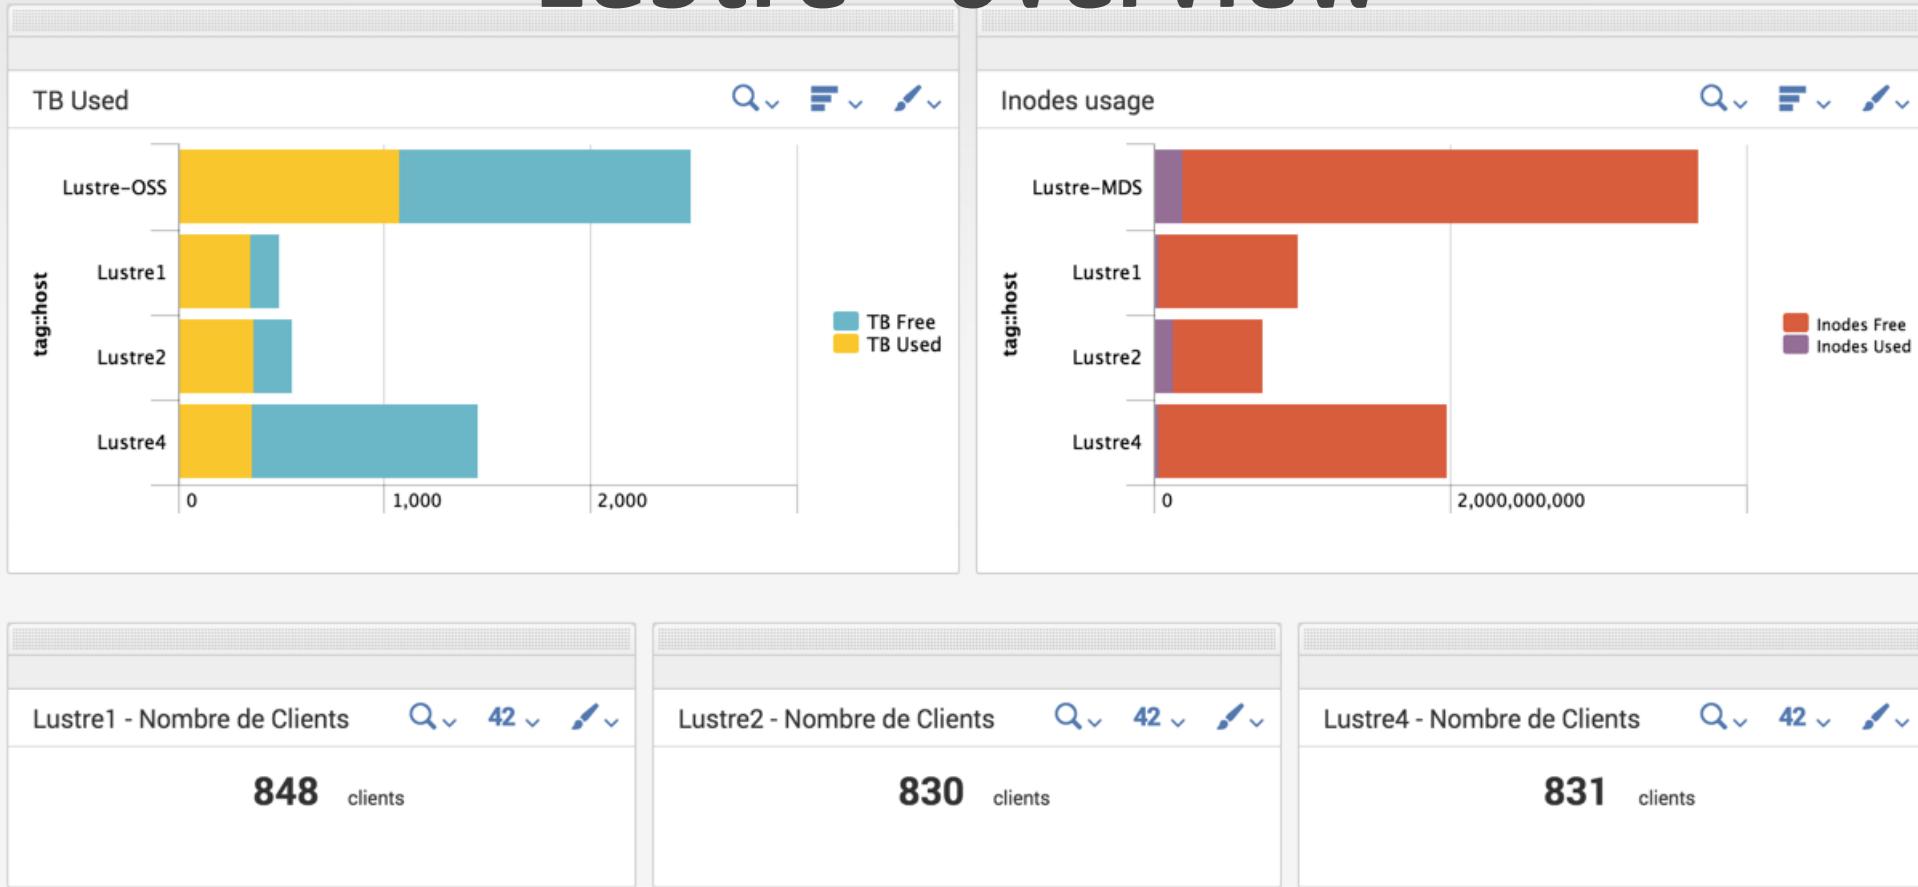
Lustre filesystems overview

1 detailed dashboard for each Lustre filesystem

Job report

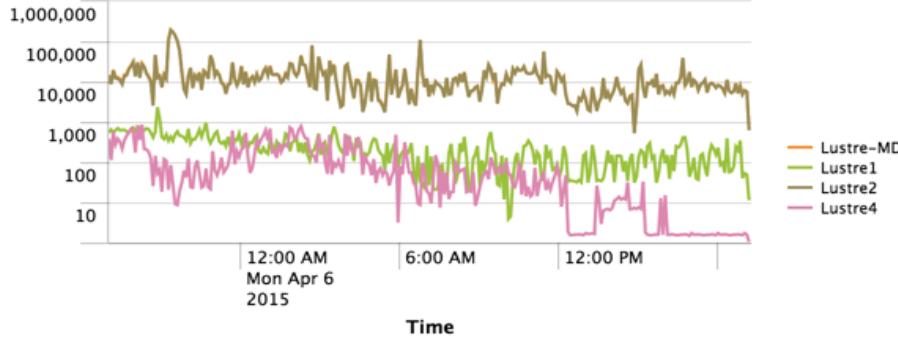
**Some reports are slower than we would like**

# Lustre - overview

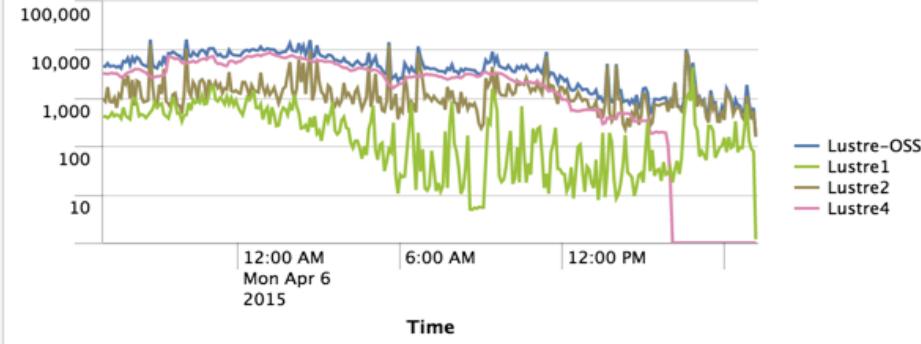


# Lustre overview (cont.)

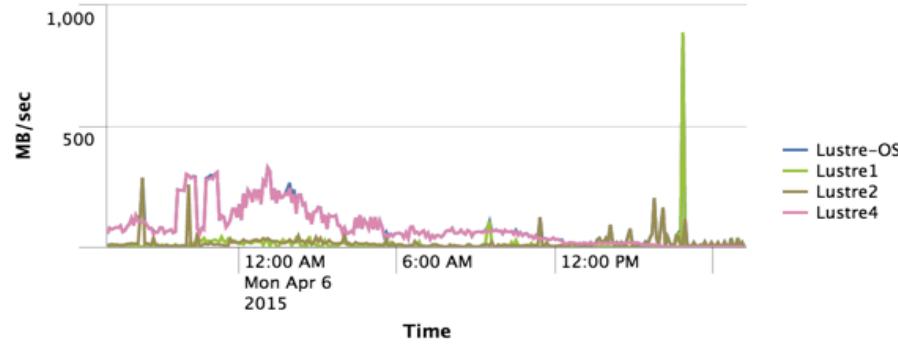
MDT IOPS



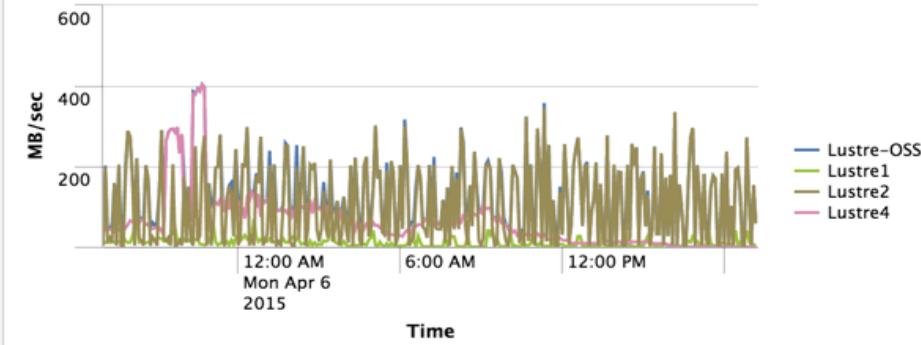
OSS IOPS



Throughput - Read (MB/sec)

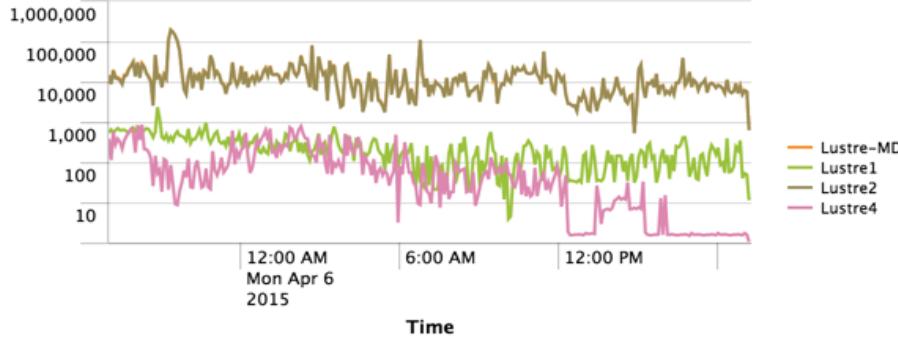


Throughput - Write (MB/sec)

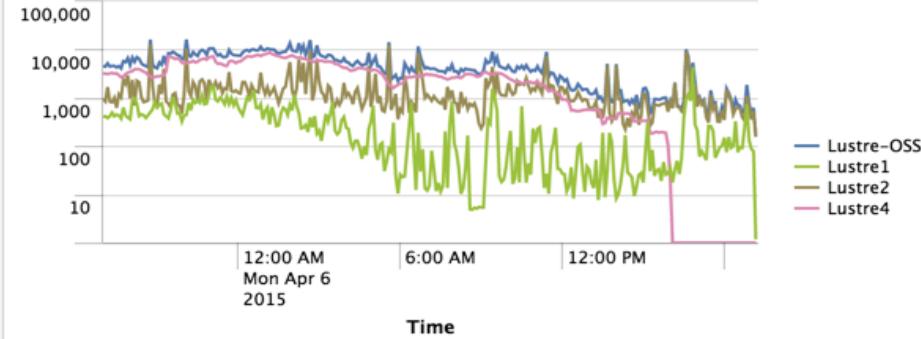


# Lustre overview (cont.)

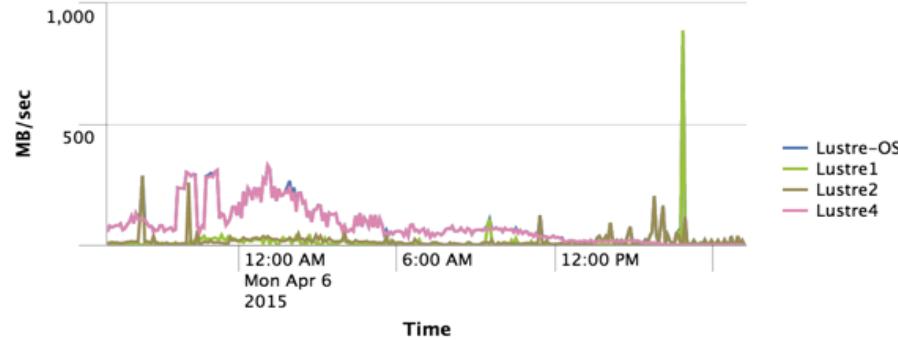
MDT IOPS



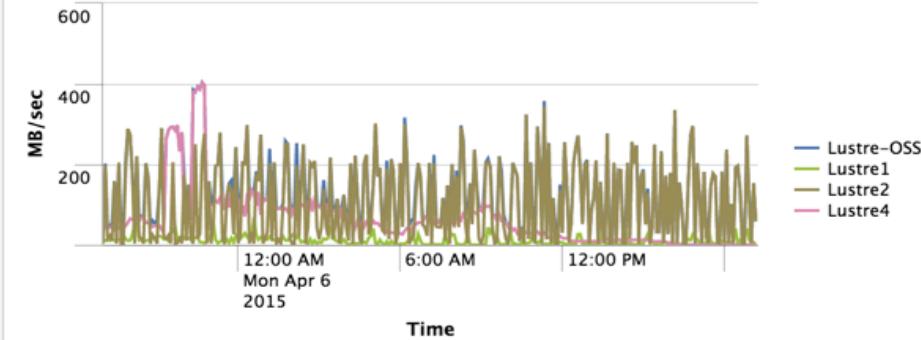
OSS IOPS



Throughput - Read (MB/sec)

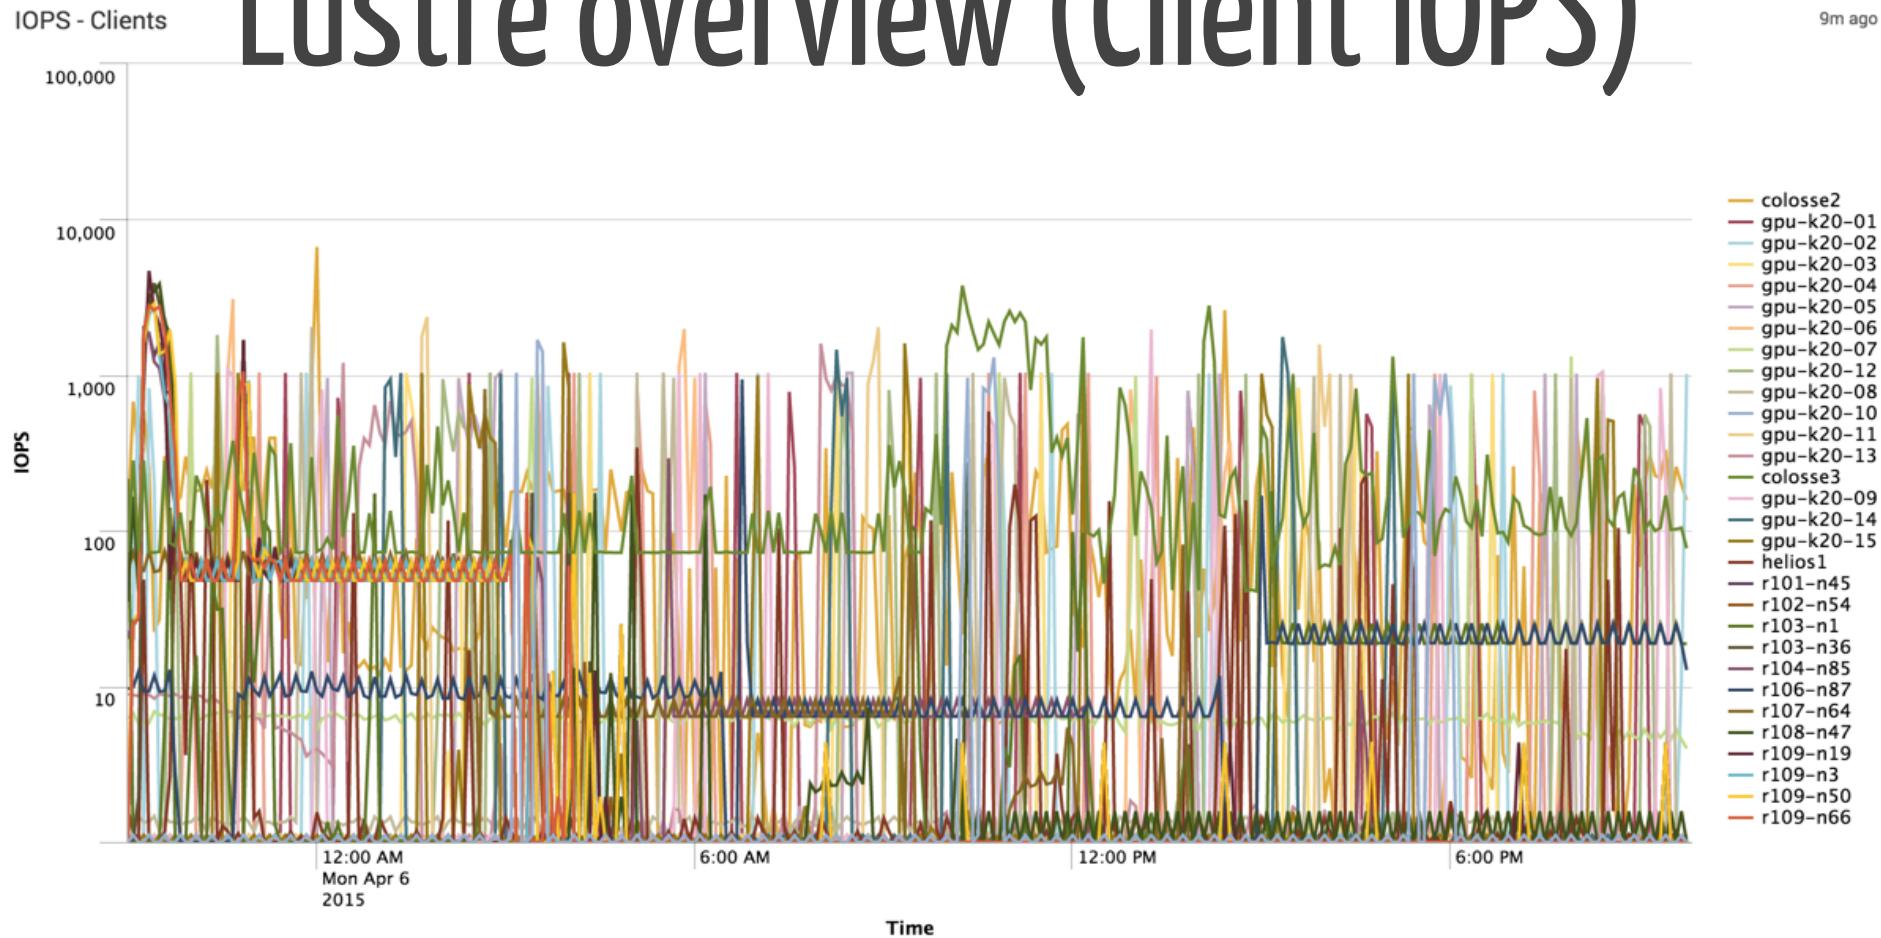


Throughput - Write (MB/sec)



# Lustre overview (Client IOPS)

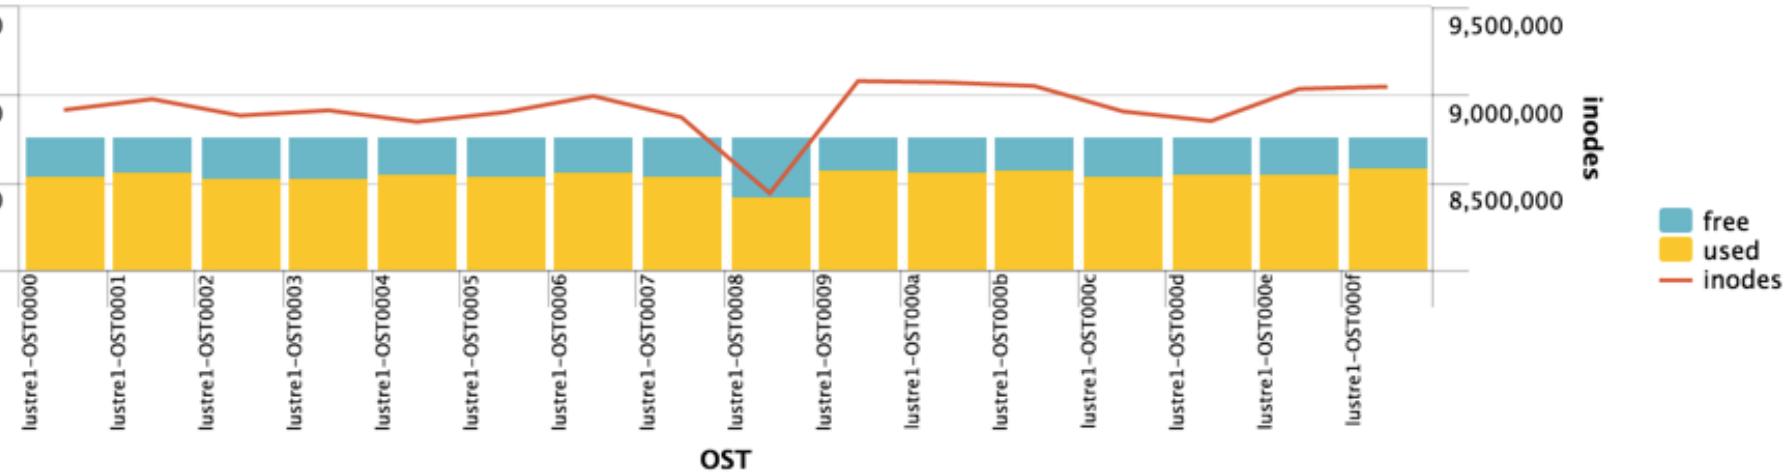
9m ago



# Per FS view (Utilization)

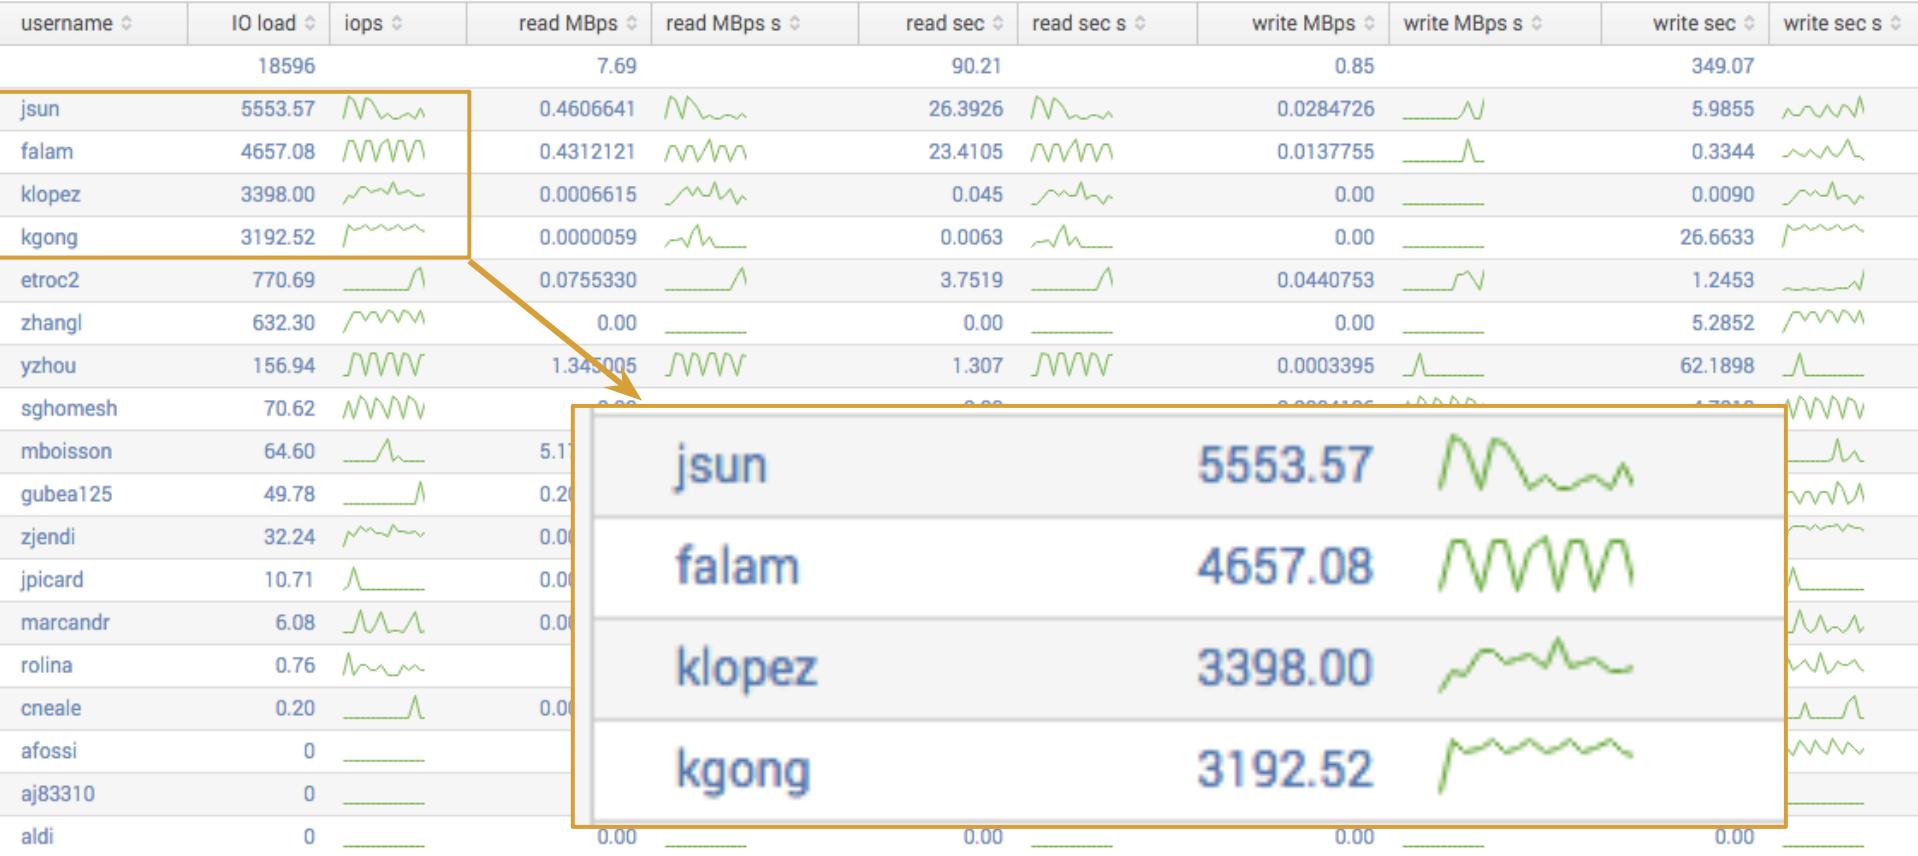
Disk usage by OST

15m ago



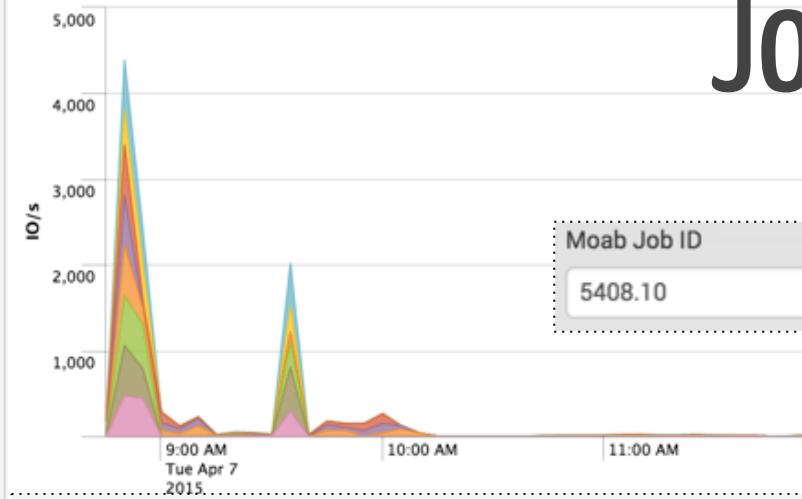
# Per FS view (top users)

Top users - last 15 min



# Job report

## IOPS Lustre2



Moab Job ID

5408.10

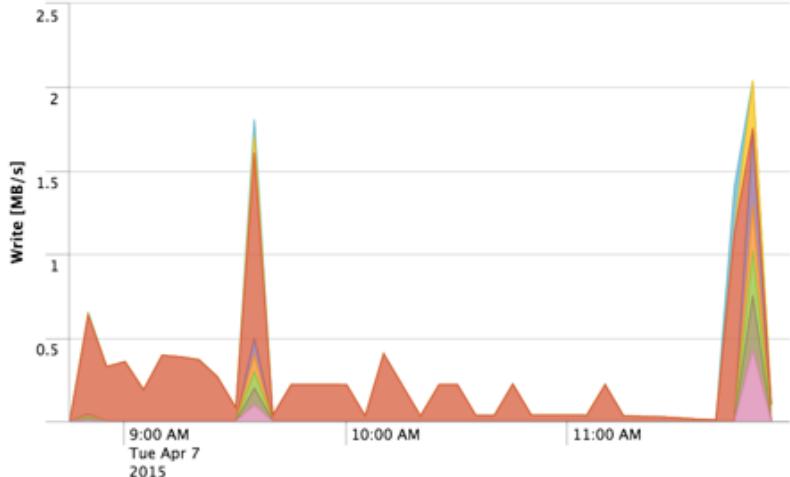
## Read MB/s Lustre2



Cluster

Helios

## Write MB/s Lustre2



## Modules used

module ⚙

StdEnv

apps/buildtools/20140527

apps/moab

apps/ray/2.3.0

compilers/intel/14.0

mpi/openmpi/1.6.5

# Limitations

The quantity of data collected is ‘scary’

Especially true for clients (~1MB/client per day)

Impact of running extra codes on compute nodes

Still uncertainty on how to monitor or visualize some components

ie: LNET routers, failovers, DNE setups, RPCs

# Looking ahead

Add support for multiple MDS/MDT

Collect client's data on the servers

Cross-reference statistics with data from  
RobinHood DB

Make the code more resistant to failovers

# Thank you

[frederick.lefebvre@calculquebec.ca](mailto:frederick.lefebvre@calculquebec.ca)



UNIVERSITÉ  
**Laval**

 Calcul Québec

The logo for Calcul Québec features a stylized infinity symbol composed of two interlocking curves, one dark blue and one light blue, with the word "Calcul" in a dark blue sans-serif font and "Québec" in a light blue sans-serif font below it.

compute \* calcul  
CANADA