



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

CRAY



Sonexion Monitoring/Metrics BoF
Matteo Chesi, Craig Flaskerud, Patti Langer

CUG 2017. CAFFEINATED COMPUTING

Redmond, Washington May 7-11, 2017

Overview

- **Purpose**

- Share and discuss storage monitoring in complex environments, what Caribou does today, what it can do for you tomorrow

- **Benefit/Value**

- Better, and improving, understanding of what's going on across your storage, fabric, and applications

Agenda

- **Motivation for project Caribou**
- **Use cases with demonstration**
- **Open discussion**

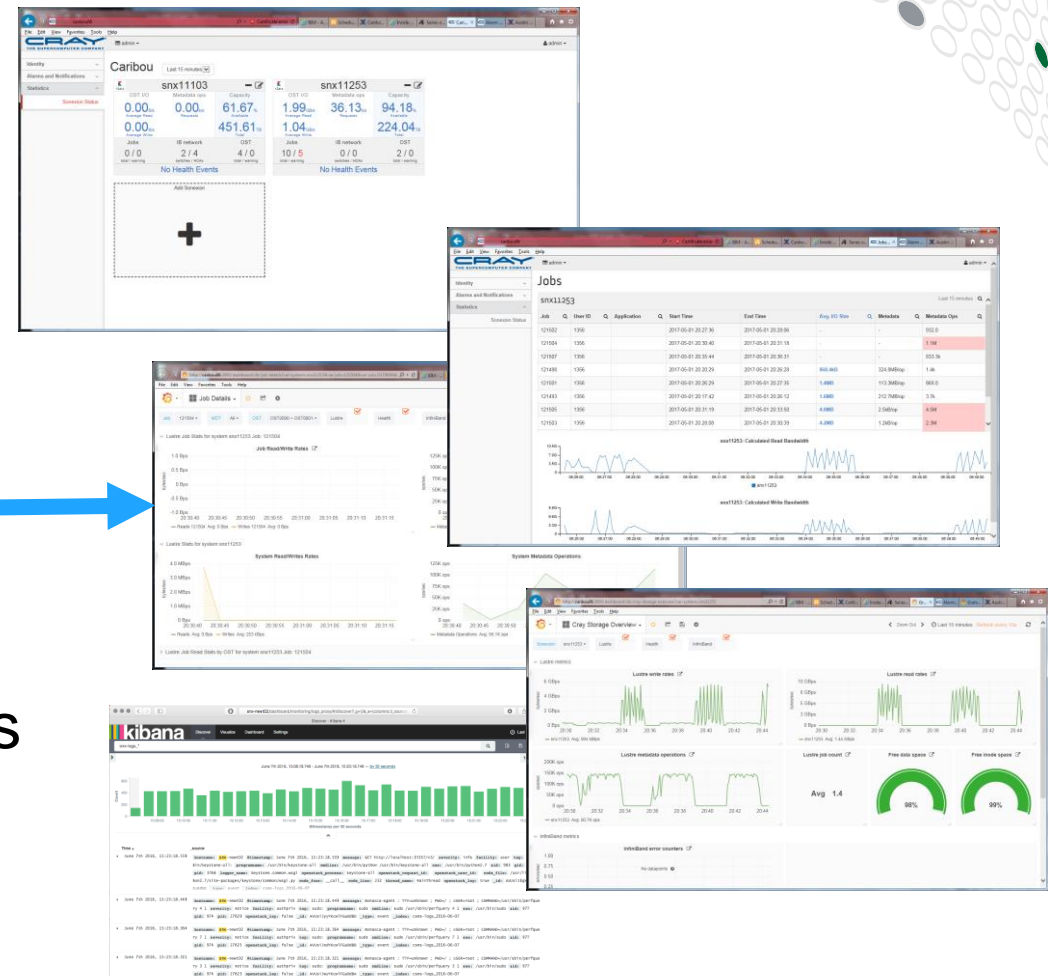
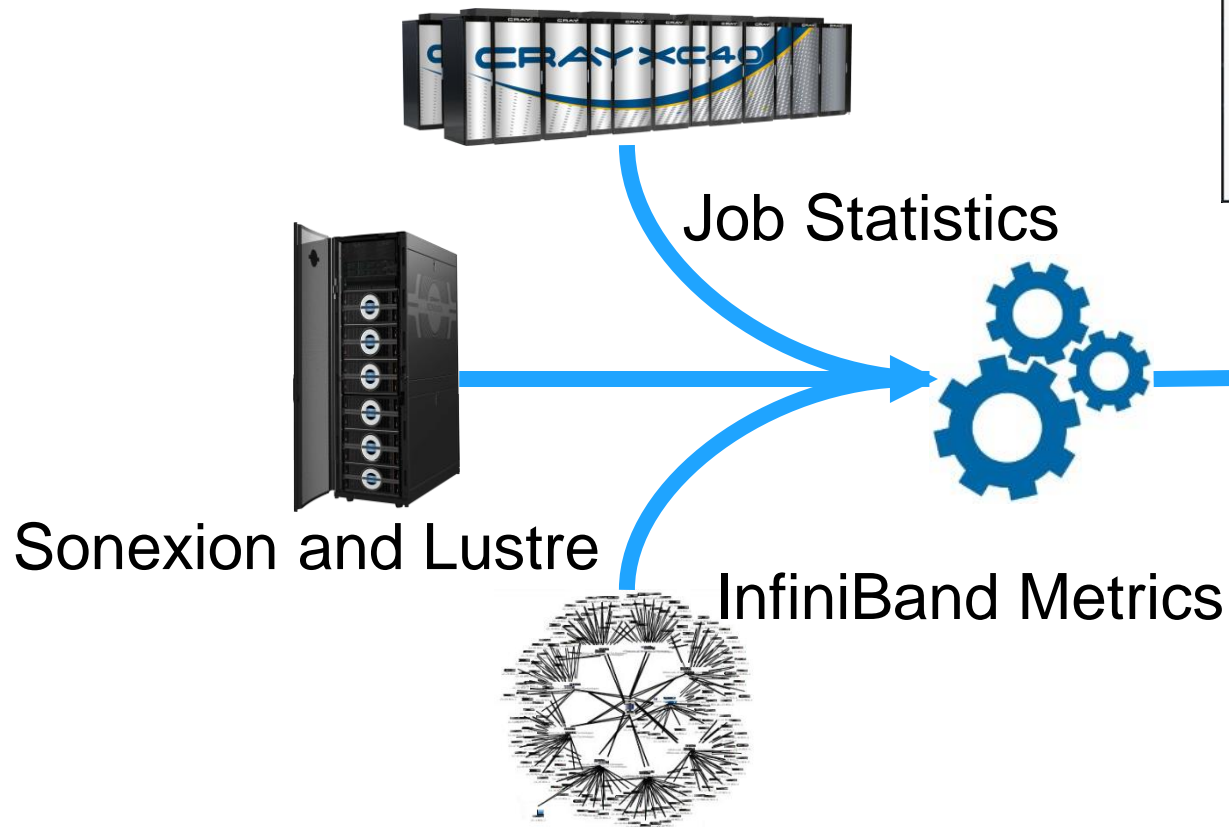
Motivation For Project Caribou

- **Provide visibility and insight**
 - Why did this application run the way it did (slow or fast)?
 - What metric information do I need and how do I get access to it?
 - How much storage do I really need for my users?
- **Enable understanding across a complex environment**
 - How do things impact each other across my environment?
 - How do I tie events together across the different entities?
 - What do I need to be better inform?

Project Caribou

- **Provide holistic approach to storage monitoring**
 - Job data...tied to storage and fabric...finally!
 - Collect, aggregate, and persist data from multiple data sources
 - Real-time and historical analysis
 - Visualization through graphs, and dashboards
 - Workflows for troubleshooting, administration, planning
- **Make your life easier**
 - Enable informed decisions from measured data
 - Reduce time to problem resolution
 - Data driven input for resource utilization and capacity planning

Project Caribou Components



COMPUTE

STORE


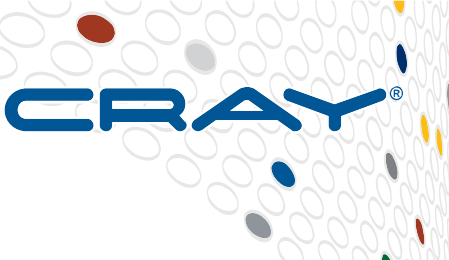
ANALYZE

CRAY INC

Exercise:

What do you do today and how can we make it easier?

Demonstration : Poor Performing Job




COMPUTE

STORE

ANALYZE

Demonstration: Sonexion Performance Issue



COMPUTE

STORE

ANALYZE

CRAY INC

Demonstration : Events, Notifications, Logs



The image shows a Grafana dashboard for 'Cray Storage Overview' and a Kibana log viewer. The Grafana dashboard displays metrics for two storage nodes: snx11103 and snx11253. It includes line graphs for Lustrite write and read rates, Lustrite metadata operations, and InfiniBand error counters. It also features gauges for Lustrite job count, free data space (49%), and free inode space (88%).

The Kibana log viewer shows the following log entries:

```

    April 29th 2017, 16:05:25.074 - May 2nd 2017, 16:05:25.074 -- by hour
    Count
    4
    3
    2
    1
    0
    2017-04-30 01:00    2017-04-30 13:00    2017-05-01 01:00    2017-05-01 13:00    2017-05-02 01:00    2017-05-02 13:00
    @timestamp per hour

    Time
    _source
    May 2nd 2017, 16:04:06.973  msg: IB topology on CA: m1x4_0, Port: 2 has diverged from the baseline. @timestamp: May 2nd 2017, 16:04:06.973 cmd: /usr/sbin/ibnetdiscover -C m1x4_0 -P 2 --diff /etc/monasca/agent/ibnetdiscover_cache.dat_m1x4_0_2 details: < vendorid=0x2c9 < devid=0xcb20 < sysimguid=0x7cfe900300be6350 < switchguid=0x7cfe900300be6350(7cfe900300be6350) < Switch 36 "S-7cfe900300be6350" # "SwitchIB Mellanox Technologies" base port 0 lid 21 lmc 0 < [2] "H-7cfe900300ca2c46"[1](7cfe900300ca2c46) # "snx11221n004 m1x5_0" lid 28 4xED
    May 1st 2017, 13:34:33.199  msg: OK. IB Topology transitioned to baseline @timestamp: May 1st 2017, 13:34:33.199 cmd: /usr/sbin/ibnetdiscover -C m1x4_0 -P 2 --diff /etc/monasca/agent/ibnetdiscover_cache.dat_m1x4_0_2 details: ret urn_code: 0 _id: AVVFSutyeA0v6P25zMu _type: IB_topology_error _index: event-ibtopology_2017-05-01
    May 1st 2017, 13:31:32.943  msg: IB topology on CA: m1x4_0, Port: 2 has diverged from the baseline. @timestamp: May 1st 2017, 13:31:32.943 cmd: /usr/sbin/ibnetdiscover -C m1x4_0 -P 2 --diff /etc/monasca/agent/ibnetdiscover_cache.dat_m1x4_0_2 details: < vendorid=0x2c9 < devid=0xcb20 < sysimguid=0x7cfe900300be6350 < switchguid=0x7cfe900300be6350(7cfe900300be6350) < Switch 36 "S-7cfe900300be6350" # "SwitchIB Mellanox Technologies" base port 0 lid 21 lmc 0 < [2] "H-7cfe900300ca2c46"[1](7cfe900300ca2c46) # "snx11221n004 m1x5_0" lid 28 4xED
  
```

Open Discussion

- **What tools do you use today?**
- **Does Project Caribou look interesting?**
- **What additional workflows are critical to you?**
- **What are the typical causes of performance issues?**

Open Discussion

- **Who are the consumers of performance data?**
- **What is the preferred method for problem alerts?**
- **How big is your site?**
- **User experience feedback**

Summary

- **It's all about enabling understanding of application storage performance**
- **With Project Caribou we hope to make your life easier**
- **For more information, attend Caribou talk**
 - Project Caribou : Monitoring and Metrics for Sonexion
 - Thursday, May 11 at 10:30am
 - Technical Session 25B
 - Salon 2

Legal Disclaimer

Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publically announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, and URIKA. The following are trademarks of Cray Inc.: APPRENTICE2, CHAPEL, CLUSTER CONNECT, CRAYPAT, CRAYPORT, ECOPHLEX, LIBSCI, NODEKARE, REVEAL, THREADSTORM. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT, and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used in this document are the property of their respective owners.



Q&A

CUG.2017.CAFFEINATED COMPUTING

Redmond, Washington May 7-11, 2017

- Identity ▾
- Alarms and Notifications ▾
- Statistics ▲

Sonexion Status

Caribou

Last 15 minutes ▾

kibana

snx11031

-

OST I/O	Metadata ops	Capacity
364.68 MB/s <small>Average Read</small>	0.00 k/s <small>Requests</small>	98.60 % <small>Available</small>
444.91 MB/s <small>Average Write</small>		169.07 TB <small>Total</small>
Jobs	IB network	OST
9 / 1 <small>total / warning</small>	11 / 16 / 1 <small>switches / HCAs / warning</small>	2 / 0 <small>total / warning</small>
No Health Events		

kibana

snx11104

-

OST I/O	Metadata ops	Capacity
0.00 B/s <small>Average Read</small>	0.00 k/s <small>Requests</small>	96.09 % <small>Available</small>
0.00 B/s <small>Average Write</small>		225.80 TB <small>Total</small>
Jobs	IB network	OST
0 / 0 <small>total / warning</small>	6 / 8 / 1 <small>switches / HCAs / warning</small>	2 / 0 <small>total / warning</small>
No Health Events		

Add Sonexion

Identity

Alarms and Notifications

Statistics

Sonexion Status


Jobs

snx11031


Last 30 minutes

Job	User ID	Application	Start Time	End Time	Avg. I/O Size	Metadata Ratio	Metadata Ops
6056184792581554675					-	-	13.0
6056184792581554677					-	-	13.0
6056184792581554678					-	-	13.0
657915	22459		2017-05-05 14:06:29		1.0kB	0.1B	17.0M
657914	22459		2017-05-05 14:02:46	2017-05-05 14:09:08	1.0MB	4.9GB	7.0
657912	22459		2017-05-05 13:58:46	2017-05-05 14:06:26	1.0MB	5.1GB	7.0
657907	22459		2017-05-05 13:52:59	2017-05-05 13:59:09	1.0MB	604.3MB	7.0
657909	22459		2017-05-05 13:54:47	2017-05-05 14:01:59	1.0MB	1.0GB	7.0

snx11031: Calculated Read Bandwidth



snx11031: Calculated Write Bandwidth



- Identity
- Alarms and Notifications
- Statistics
- Sonexion Status


Jobs

snx11031 Last 15 minutes


Job	User ID	Application	Start Time	End Time	Avg. I/O Size	Metadata Ratio	Metadata Ops
657915	22459		2017-05-05 14:06:29		1.0kB	0.1B	17.0M
657917	22459		2017-05-05 14:09:10	2017-05-05 14:14:57	1.0MB	2.6GB	7.0
657913	22459		2017-05-05 14:00:46	2017-05-05 14:07:47	1.0MB	5.3GB	12.0
657914	22459		2017-05-05 14:02:46	2017-05-05 14:09:08	1.0MB	4.9GB	7.0
657912	22459		2017-05-05 13:58:46	2017-05-05 14:06:26	1.0MB	8.3GB	12.0
657918	22459		2017-05-05 14:13:57		1.0MB	5.8GB	7.0
657916	22459		2017-05-05 14:07:50	2017-05-05 14:13:55	1.0MB	960.0MB	7.0
657919	22459		2017-05-05 14:14:59		1.0MB	3.2GB	7.0

Job 657915

Avg. I/O Size



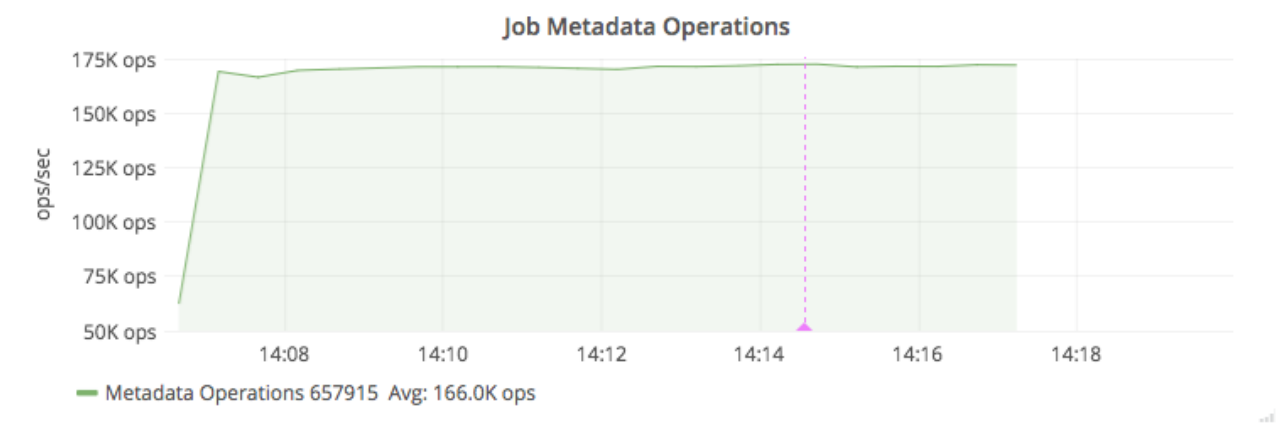
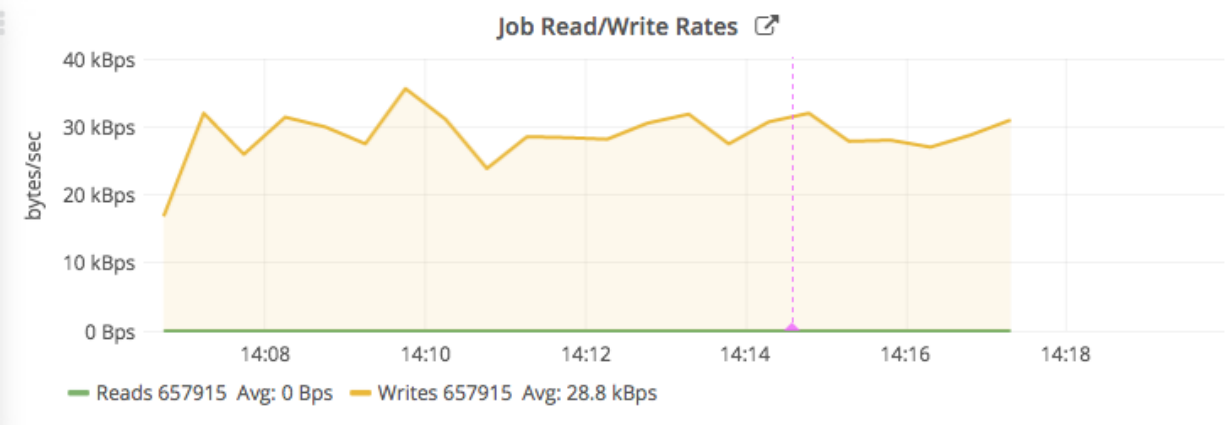
Metadata Ratio





Job 657915
MDT All
OST OST0001
Lustre
Health
InfiniBand

▼ Lustre Job Stats for system snx11031 Job: 657915



> Lustre Stats for system snx11031


> Lustre Job Read Stats by OST for system snx11031 Job: 657915

> Lustre Job Write Stats by OST for system snx11031 Job: 657915


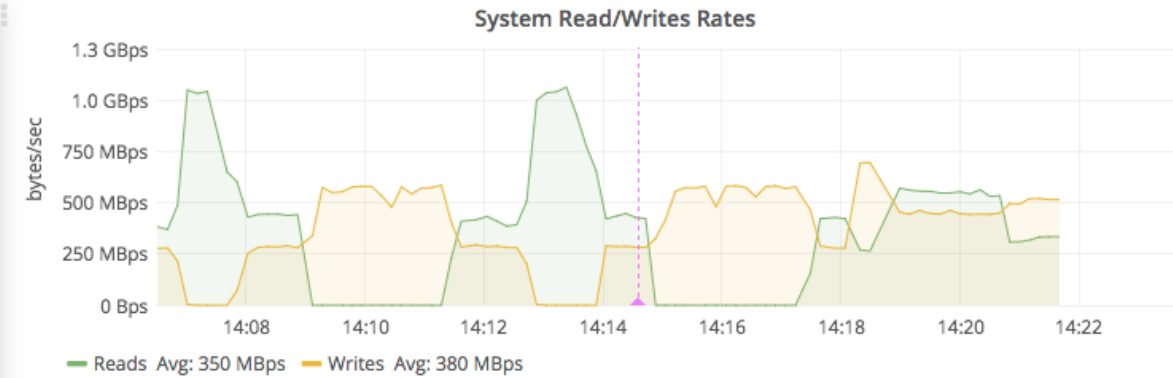
> Breakdown Metadata Operations for system snx11031 Job: 657915



Lustre Job Stats for system snx11031 Job: 657915



Lustre Stats for system snx11031



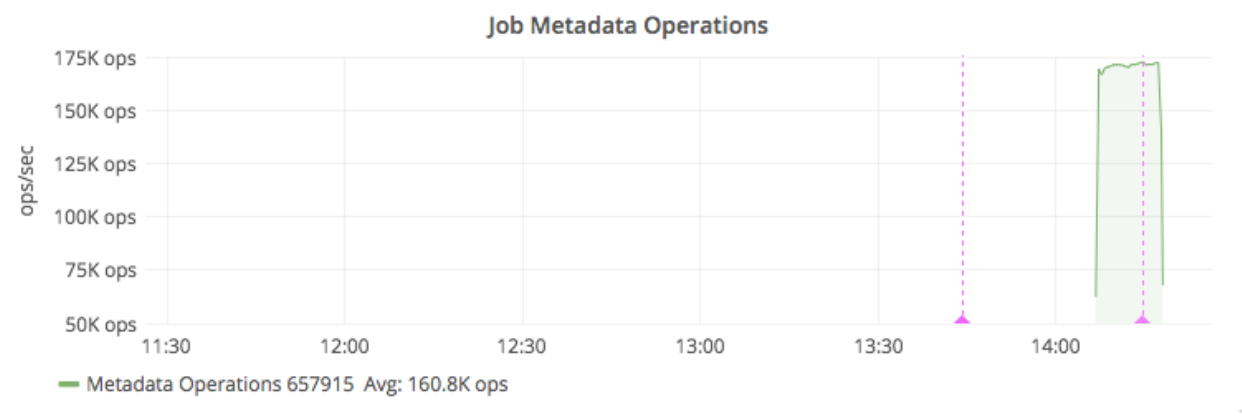
> Lustre Job Read Stats by OST for system snx11031 Job: 657915

> Lustre Job Write Stats by OST for system snx11031 Job: 657915

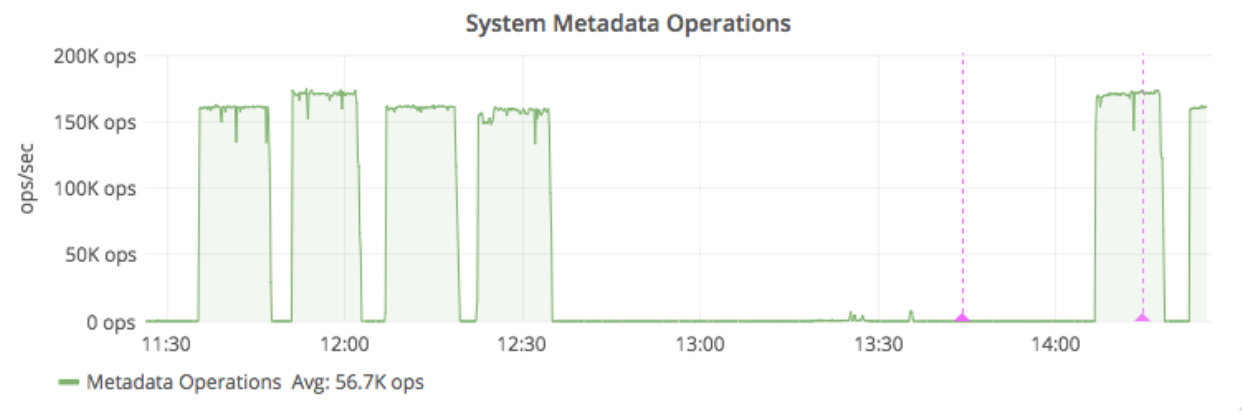
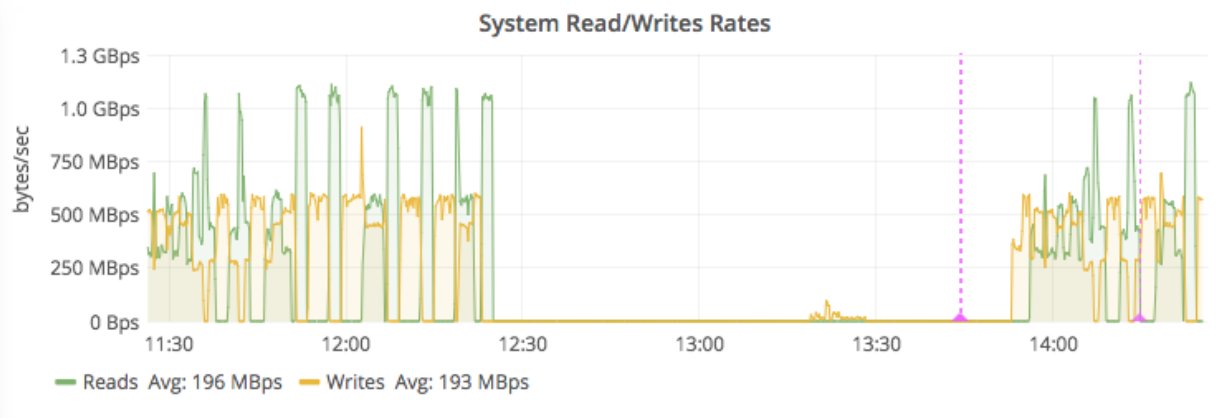
> Breakdown Metadata Operations for system snx11031 Job: 657915



▼ Lustre Job Stats for system snx11031 Job: 657915



▼ Lustre Stats for system snx11031



> Lustre Job Read Stats by OST for system snx11031 Job: 657915

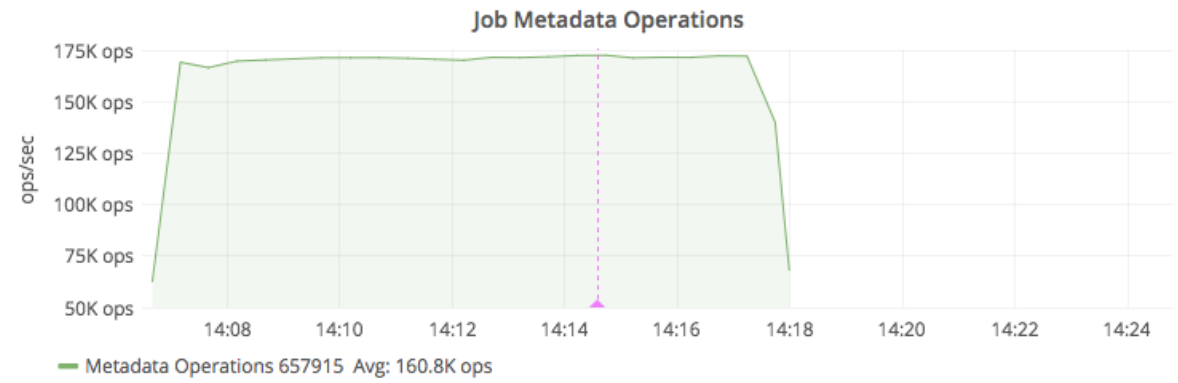
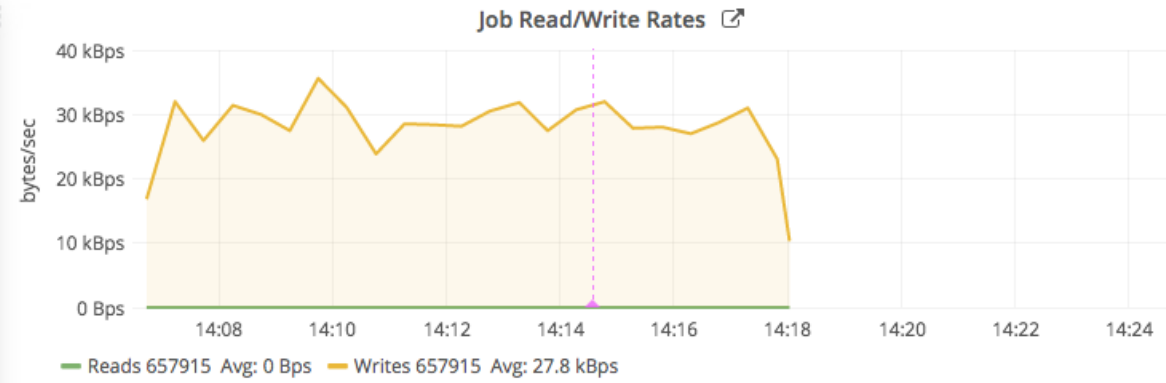
> Lustre Job Write Stats by OST for system snx11031 Job: 657915

> Breakdown Metadata Operations for system snx11031 Job: 657915



Job 657915 MDT All OST OST0001 Lustre Health InfiniBand

Lustre Job Stats for system snx11031 Job: 657915



Lustre Stats for system snx11031

Lustre Job Read Stats by OST for system snx11031 Job: 657915

Lustre Job Write Stats by OST for system snx11031 Job: 657915




Breakdown Metadata Operations for system snx11031 Job: 657915

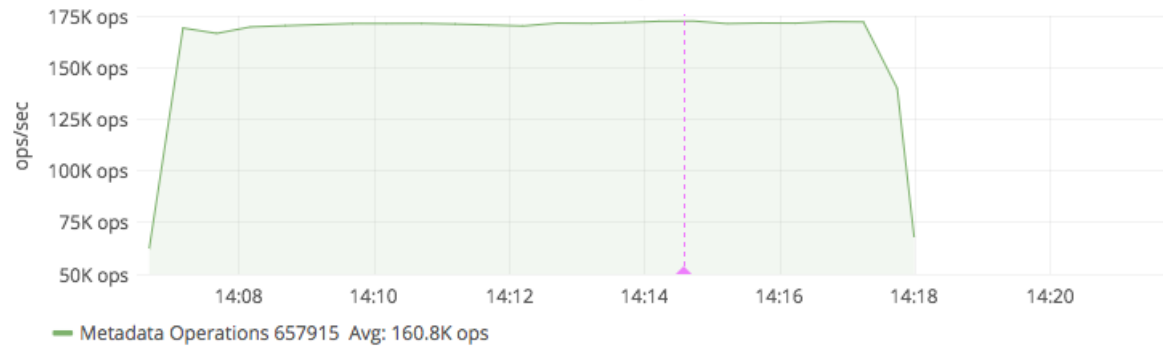


▼ Lustre Job Stats for system snx11031 Job: 657915

Job Read/Write Rates



Job Metadata Operations




> Lustre Stats for system snx11031

> Lustre Job Read Stats by OST for system snx11031 Job: 657915

> Lustre Job Write Stats by OST for system snx11031 Job: 657915

▼ Breakdown Metadata Operations for system snx11031 Job: 657915



- Identity
- Alarms and Notifications
- Statistics

Sonexion Status

Caribou

Last 15 minutes

kibana

snx11031

[-] [edit]

OST I/O	Metadata ops	Capacity
295.33 MB/s Average Read	120.16 k/s Requests	98.65 % Available
413.34 MB/s Average Write		169.07 TB Total

Jobs	IB network	OST
8 / 2 total / warning	11 / 16 / 0 switches / HCAs / warning	2 / 0 total / warning

No Health Events

kibana

snx11104

[-] [edit]



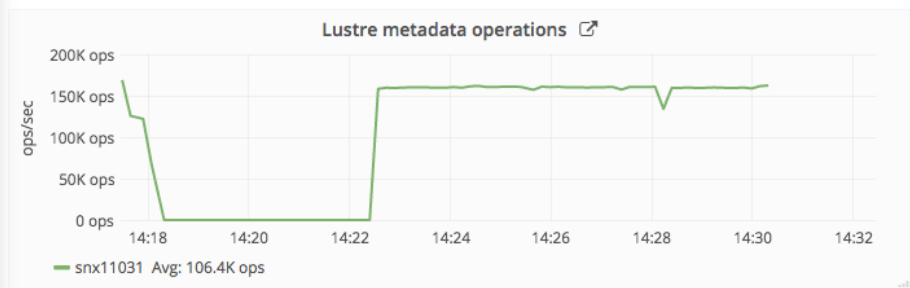
OST I/O	Metadata ops	Capacity
0.00 B/s Average Read	0.00 k/s Requests	96.09 % Available
0.00 B/s Average Write		225.80 TB Total

Jobs	IB network	OST
0 / 0 total / warning	6 / 8 / 0 switches / HCAs / warning	2 / 0 total / warning

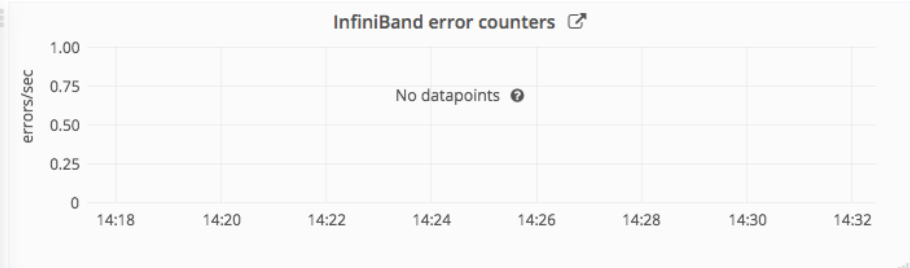
No Health Events

Add Sonexion

Lustre metrics



InfiniBand metrics





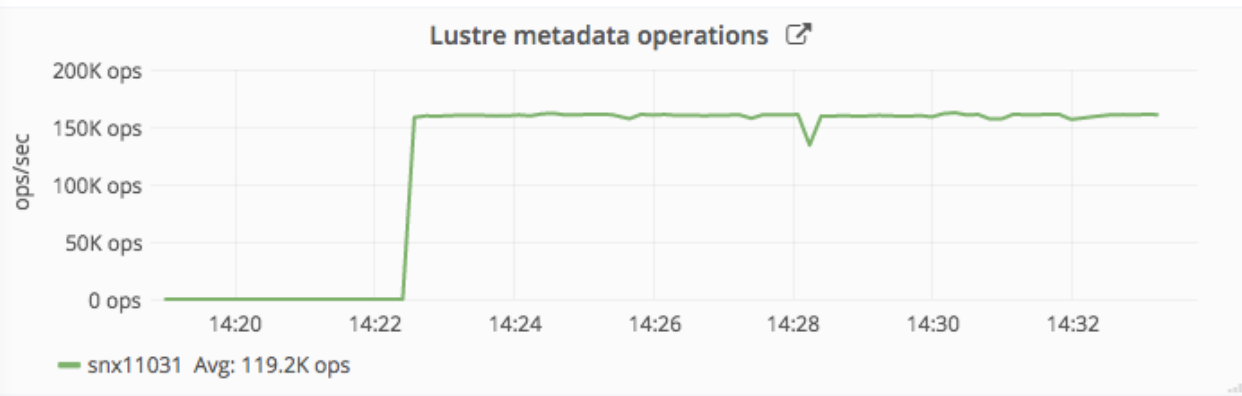
- Lustre metrics**
- 1 vs X ☆
 - JG ☆
 - Lustre Disk Space/Inode Usage ☆
 - Lustre Metadata Metrics ☆
 - Lustre Write-Read Metrics ☆

- InfiniBand metrics**
- InfiniBand Data Metrics ☆
 - InfiniBand Error Counters ☆


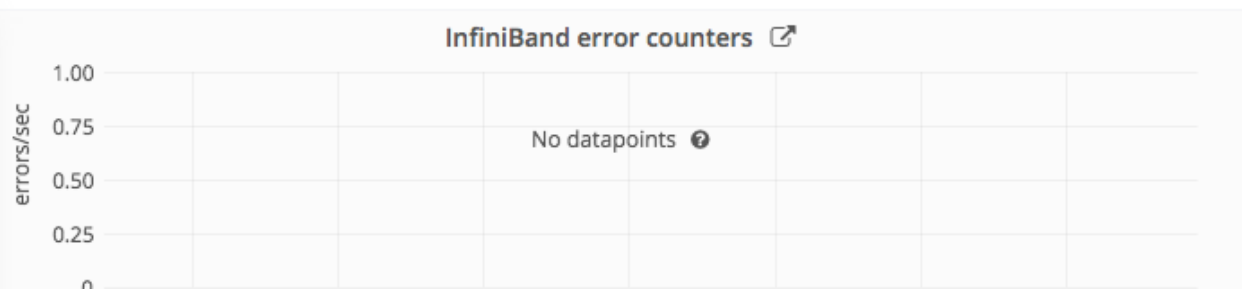
- Sonexion health**
- Sonexion Server Metrics ☆



▼ Lustr metrics



▼ InfiniBand metrics



- Identity
- Alarms and Notifications
- Statistics

Sonexion Status

Caribou

Last 15 minutes

kibana

snx11031

[-] [edit]

OST I/O	Metadata ops	Capacity
295.33 MB/s Average Read	120.16 k/s Requests	98.65 % Available
413.34 MB/s Average Write		169.07 TB Total
Jobs	IB network	OST
8 / 2 total / warning	11 / 16 / 0 switches / HCAs / warning	2 / 0 total / warning
No Health Events		

kibana

snx11104

[-] [edit]

OST I/O	Metadata ops	Capacity
0.00 B/s Average Read	0.00 k/s Requests	96.09 % Available
0.00 B/s Average Write		225.80 TB Total
Jobs	IB network	OST
0 / 0 total / warning	6 / 8 / 0 switches / HCAs / warning	2 / 0 total / warning
No Health Events		

Add Sonexion



Sonexion

snx11031

MDT

All

Lustre




Health



InfiniBand



Lustre metadata snx11031



	min	max	avg	current
All MDTs	0 ops	162.4K ops	131.9K ops	160.6K ops
getattr MDT0000	0 ops	40.4K ops	33.0K ops	39.5K ops
open MDT0000	0 ops	40.4K ops	33.0K ops	39.5K ops
close MDT0000	0 ops	40.4K ops	33.0K ops	39.5K ops
setattr MDT0000	0 ops	40.4K ops	33.0K ops	39.5K ops
mkdir MDT0000	0 ops	2 ops	0 ops	0 ops
rename MDT0000	0 ops	0 ops	0 ops	0 ops
quotactl MDT0000	0 ops	0 ops	0 ops	0 ops
getxattr MDT0000	0 ops	0 ops	0 ops	0 ops
rmdir MDT0000	0 ops	0 ops	0 ops	0 ops
unlink MDT0000	0 ops	0 ops	0 ops	0 ops

+ ADD ROW

Identity

Alarms and Notifications

Statistics

Sonexion Status

Jobs


snx11031

Last 15 minutes


Job	User ID	Application	Start Time	End Time	Avg. I/O Size	Metadata Ratio	Metadata Ops
657922	22459		2017-05-05 14:22:20	2017-05-05 14:34:34	1.0kB	0.2B	13.7M
657924	22459		2017-05-05 14:24:30	2017-05-05 14:30:08	1.0MB	1.8GB	8.0
657921	22459		2017-05-05 14:18:06	2017-05-05 14:24:25	1.0MB	6.4GB	7.0
657919	22459		2017-05-05 14:14:59	2017-05-05 14:22:17	1.0MB	15.2GB	7.0
657920	22459		2017-05-05 14:18:06	2017-05-05 14:24:26	1.0MB	5.6GB	8.0
657923	22459		2017-05-05 14:24:30	2017-05-05 14:30:09	1.0MB	2.1GB	7.0

Job 657922

Avg. I/O Size





Metadata Ratio





▼ Lustre Job Stats for system snx11031 Job: 657922




> Lustre Stats for system snx11031

> Lustre Job Read Stats by OST for system snx11031 Job: 657922

> Lustre Job Write Stats by OST for system snx11031 Job: 657922

▼ Breakdown Metadata Operations for system snx11031 Job: 657922



Identity

Alarms and Notifications

Alarm Definitions

Alarms

Notifications

Statistics

Alarm Definitions

Filter

+ Create Alarm Definition

× Delete Alarm Definitions

<input type="checkbox"/>	Name	Description	Notifications Enabled	Actions
<input type="checkbox"/>	Caribou_Sonexion_Metric_Health	Check for Incoming metrics	True	Edit Alarm Definition
<input type="checkbox"/>	Caribou_OST_Free_Space	The free disk space of a specific OST is under 5%	True	Edit Alarm Definition
<input type="checkbox"/>	Caribou_Sonexion_snmpwalk_timeout	There is a timeout using SNMPwalk to a Sonexion	True	Edit Alarm Definition
<input type="checkbox"/>	Root_Partition_inode_Usage	The use of the configured inodes of the root partition is over 95%	True	Edit Alarm Definition
<input type="checkbox"/>	Var_Filesystem_Disk_Space_Usage	The disk space of the /var filesystem partition is over 95%	True	Edit Alarm Definition
<input type="checkbox"/>	Root_Partition_Disk_Space_Usage	The disk space of the root partition is over 95%	True	Edit Alarm Definition
<input type="checkbox"/>	Var_Filesystem_Disk_Partition_inode_Usage	The use of the configured inodes of the /var filesystem partition is over 95%	True	Edit Alarm Definition
<input type="checkbox"/>	Caribou_IB_Topology	IB Topology Change Detection	True	Edit Alarm Definition
<input type="checkbox"/>	Caribou_OST_Free_Files	The free inode space of a specific OST/MDT is under 5%	True	Edit Alarm Definition
<input type="checkbox"/>	Keystone_Services_are_down	All the Keystone services are not operating properly	True	Edit Alarm Definition

Displaying 10 items

Next Page

Identity

Alarms and Notifications

Alarm Definitions

Alarms

Notifications

Statistics

Alarm D

Edit Alarm Definition

Description:

The Name field is used to identify the notification method.

The Expression field which if true, triggers a notification to be sent. See [Alarm Expressions](#) for how to write an expression.

The Alarm Actions field contains the list of Notification that should be sent when transitioning to an ALARM state.

Name *

Caribou_IB_Topology

Expression *

(max(cray_ib.topology, 60) > 0 times 1)

Apply function to metrics

- individually
- as a group

Description

IB Topology Change Detection

Severity

Critical

Notifications Enabled

Notifications

Name	Type	Address	
DEFAULT_EMAIL	EMAIL	storage-alert@cscs.ch	X

Cancel

Save

+ Create Alarm Definition

x Delete Alarm Definitions

Notifications Enabled

Actions

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

True

Edit Alarm Definition

Displaying 10 items

Next Page

- Identity
- Alarms and Notifications
- Statistics

Sonexion Status

Caribou

Last 15 minutes

kibana

snx11031

-

OST I/O	Metadata ops	Capacity
295.33 ^{MB/s} <small>Average Read</small>	120.16 ^{k/s} <small>Requests</small>	98.65% <small>Available</small>
413.34 ^{MB/s} <small>Average Write</small>		169.07 ^{TB} <small>Total</small>
Jobs	IB network	OST
8 / 2 <small>total / warning</small>	11 / 16 / 0 <small>switches / HCAs / warning</small>	2 / 0 <small>total / warning</small>
No Health Events		

kibana

snx11104

-

OST I/O	Metadata ops	Capacity
0.00 ^{B/s} <small>Average Read</small>	0.00 ^{k/s} <small>Requests</small>	96.09% <small>Available</small>
0.00 ^{B/s} <small>Average Write</small>		225.80 ^{TB} <small>Total</small>
Jobs	IB network	OST
0 / 0 <small>total / warning</small>	6 / 8 / 0 <small>switches / HCAs / warning</small>	2 / 0 <small>total / warning</small>
No Health Events		

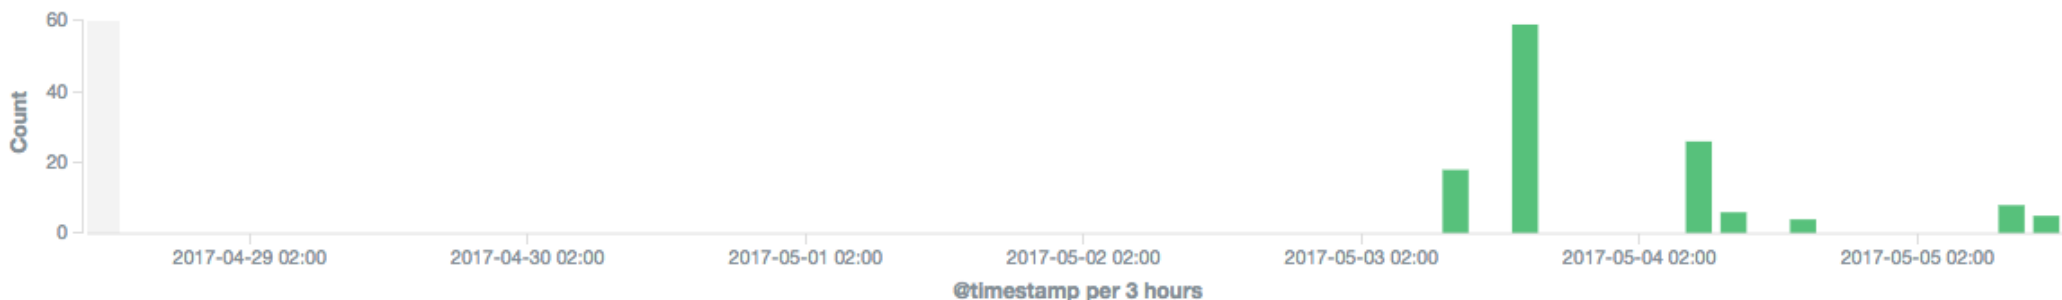
Add Sonexion

(Lustre OR LustreError) AND programname: kernel



Lustre messages 126 hits

April 28th 2017, 14:47:04.904 - May 5th 2017, 14:47:04.904 — [by 3 hours](#)



Time	hostname	message
▶ May 5th 2017, 13:30:46.000	snx11031n002	Lustre : MGS: haven't heard from client e22c841c-2bec-504c-52c2-b9965d39d905 (at 148.187.4.123@o2ib6002) in 682 seconds. I think it's dead, and I am evicting it. exp ffff8807da01b000, cur 1493983846 expire 1493983396 last 1493983164
▶ May 5th 2017, 13:30:45.000	snx11031n004	Lustre : snx11031-OST0000: haven't heard from client 5c7228ff-5db5-167c-feb6-082924c5d296 (at 148.187.4.123@o2ib6002) in 681 seconds. I think it's dead, and I am evicting it. exp ffff8804ba1d6c00, cur 1493983845 expire 1493983395 last 1493983164
▶ May 5th 2017, 13:30:45.000	snx11031n004	Lustre : Skipped 1 previous similar message
▶ May 5th 2017, 13:30:45.000	snx11031n003	Lustre : snx11031-MDT0000: haven't heard from client 5c7228ff-5db5-167c-feb6-082924c5d296 (at 148.187.4.109@o2ib6001) in 681 seconds. I think it's dead, and I am evicting it. exp ffff88102f1fec00, cur 1493983845 expire 1493983395 last 1493983164
▶ May 5th 2017, 13:30:42.000	snx11031n005	Lustre : snx11031-OST0001: haven't heard from client 5c7228ff-5db5-167c-feb6-082924c5d296 (at 148.187.4.109@o2ib6001) in 678 seconds. I think it's dead, and I am evicting it. exp ffff8807e6647c00, cur 1493983842 expire 1493983392 last 1493983164
▶ May 5th 2017, 09:05:00.000	snx11031n002	Lustre : MGS: Connection restored to 84d68edb-d49a-7036-67e1-bd73e72f45ea (at 148.187.6.103@o2ib2)
▶ May 5th 2017, 09:05:00.000	snx11104n004	Lustre : snx11104-OST0000: Connection restored to 16a1cf0a-030c-41eb-b47b-4b29c2e8d82c (at 148.187.6.103@o2ib2)

snx-logs_*

Selected Fields

- hostname
- message

Available Fields

Popular

- programname
- @timestamp
- _id
- _index
- _type
- cmdline
- exe
- facility
- gid
- openstack_log
- pid
- remote_log
- severity
- tag
- uid