April 8th - 10th, 2014

DataDirect™ NETWORKS
INFORMATION IN MOTION™

LUG14

Lustre Log Analyzer

Kalpak Shah

DataDirect Networks

Lustre Log Analysis Requirements

DataDirect NETWORKS
INFORMATION IN MOTION

- Need scripts to parse Lustre debug logs
  - Only way to effectively use the logs for debugging
- No structured approach to analyzing and understanding Lustre logs
- Event correlation between multiple clients and Lustre servers can be very difficult
- MDS / OSS RPC Tracing capabilities
- Event notifications based on configurable triggers
- High level of expertise and Lustre knowledge needed to use logs effectively
- Lustre logs can grow to millions of lines across servers – sometimes in GBs

Who needs this?

DataDirect NETWORKS
INFORMATION IN MOTION

- Developers
  - A handy tool to understand and analyze Lustre logs would be most beneficial for new and experienced developers
- Support
  - Once logs are received from customer, support would like to identify suspicious issues quickly
  - Need ability to search and index large logs quickly and effectively
- Administrators
  - While debugging specific issues onsite, admins would like to monitor the logs (easily) and get alerts

Some options today

- **llanalyze**
  - Perl script to indent and color code logs
  - Extract specific subsystem log lines
  - Show RPC patterns
  - Limited use – and not maintained or improved anymore
- **Simple Event Correlator**
  - Useful for event log monitoring
  - Does not index the logs so searching is slow and in-effective
- **Splunk**
  - Good GUI for log management and analysis
  - Indexes log data so can be searched
  - Commercial tool – unlikely to be used by developers and support

Some specific log analysis needs

- Toggle subsystems on / off
  - mdc, mds, osc, ost, obdclass, obdfilter, llite, ptlrpc, portals, lnd, ldlm, lov
- Analyze logs across multiple servers in a single view
- RPC debugging
- Search for specific strings / errors / warnings
- View logs across nodes filtered between certain timestamps
- Ability to add custom logic / searches
- Email alerts based on certain conditions

Logstash

- logstash is a tool for managing events and logs. You can use it to collect logs, parse them, and store them for later use
- It helps if you take logs and other event data from your systems and move it into a central place
- Using Elasticsearch as a backend datastore - Logstash acts as the workhorse, creating a powerful pipeline for storing, querying and analyzing your logs.
- It has built-in inputs, filters, codecs and outputs, you can harness some powerful functionality with a small amount of effort.

Logstash Components – Log Collection Pipeline

DataDirect NETWORKS
INFORMATION IN MOTION

► Input

- Inputs are the mechanism for passing log data to Logstash.
- **file**: reads from a file on the filesystem
- **syslog**: listens on the well-known port 514 for syslog messages and parses it

► Filters

- Filters are used as intermediary processing devices in the Logstash chain. They are often combined with conditionals in order to perform a certain action on an event, if it matches particular criteria
- **drop**: drop an event completely, for example, *debug* events
- **clone**: make a copy of an event, possibly adding or removing fields.

► Output

- Outputs are the final phase of the Logstash pipeline
- **elasticsearch:** save your data in an efficient, convenient and easily queryable format
- **file:** writes event data to a file on disk.

Logstash Configuration for Lustre logs

```
input {
        file { path => "/tmp/lctl/debug_kernel" }
      }
 filter {
      lustre_debug { }
    }
output {
        elasticsearch { host => elasticsearch-
server }
      }
```

Custom parsing logic to parse lctl debug_kernel logs. It is a Ruby file

Elasticsearch

- Elasticsearch is a flexible and powerful open source, distributed, real-time search and analytics engine.
- Elasticsearch gives you the ability to move easily beyond simple full-text search
- Provides scalability to our log indexing and storage capabilities

Kibana

- We now have the data in elasticsearch, but looking at JSON documents would be worse than Lustre logs
- Kibana is an open source (Apache Licensed), browser based analytics and search dashboard for ElasticSearch
- It has a very nice interface to build graphs, charts and much more based on data stored in an elasticsearch index.
- Let's see how we can use it for our Lustre log analysis requirements

Email alerts

- Email alerts can be configured in Logstash using their conditionals configuration
- Use regular expression searches which will be used by Logstash while parsing logs to trigger alerts

Searching for logs of a particular subsystem across nodes

DataDirect NETWORKS
INFORMATION IN MOTION

In Logstash filtering, we can replace these subsystem codes with names

Search Any Text

Change Time Frame & auto-update window

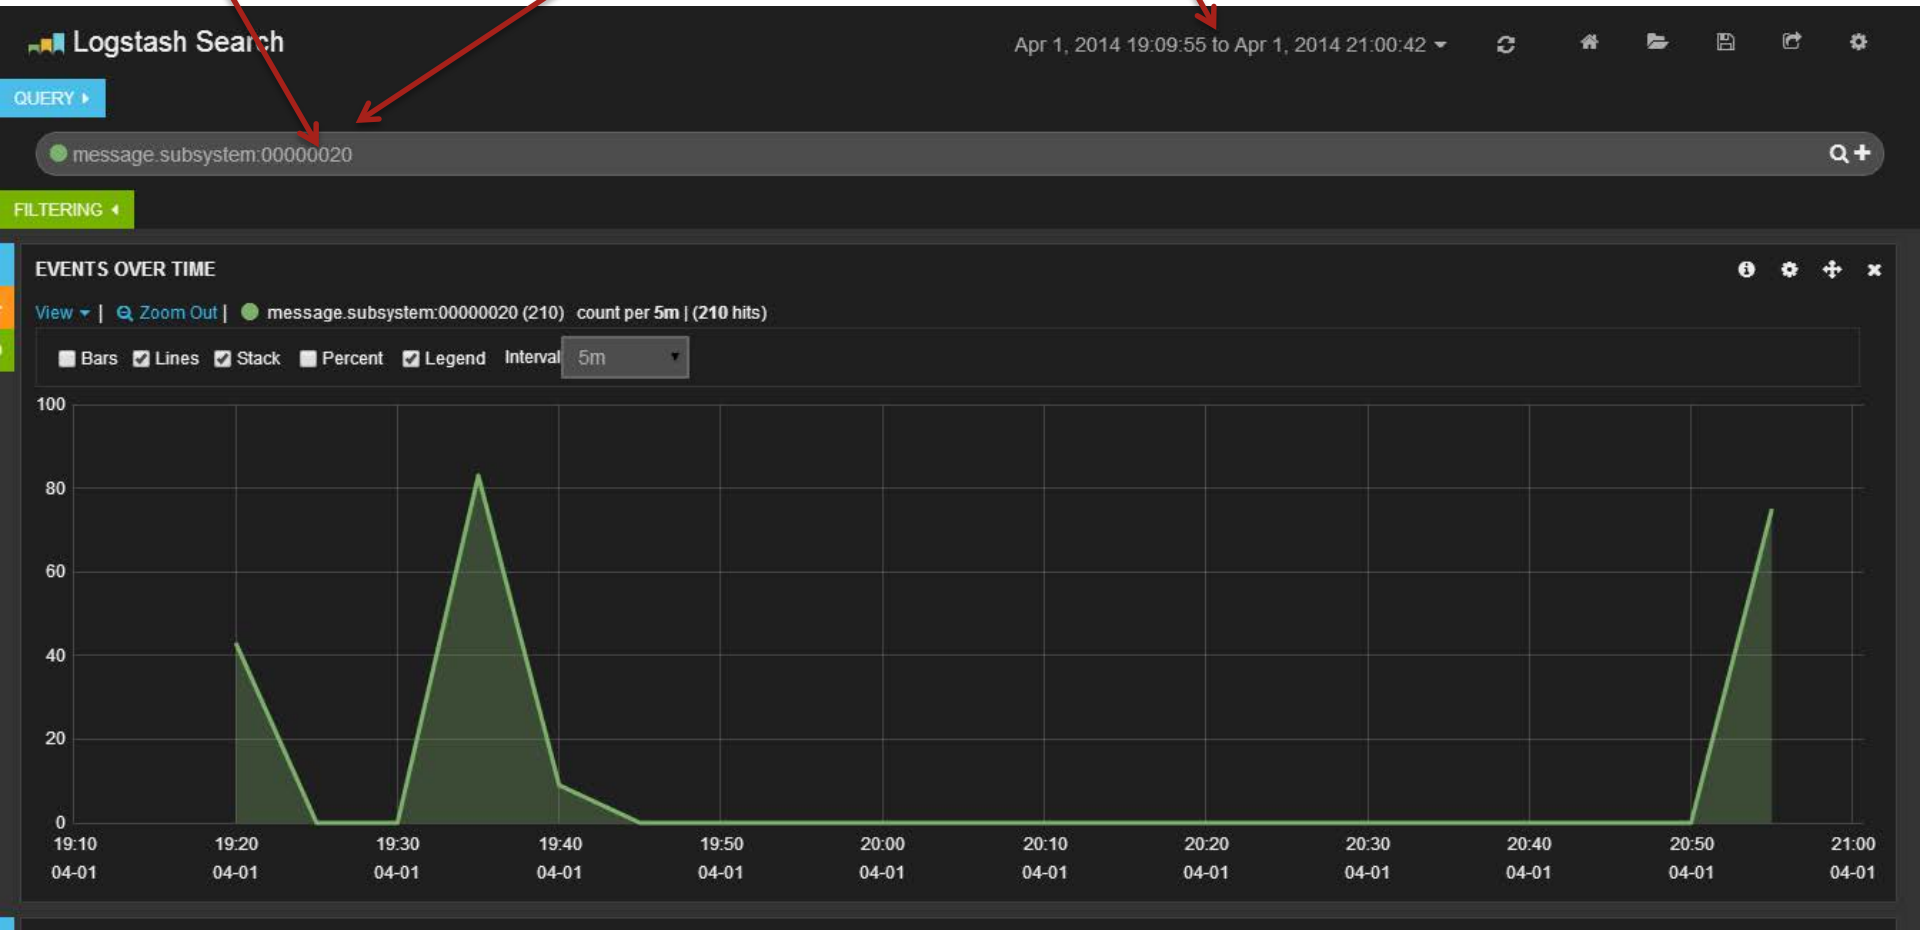

Graphical view for your text searches

DataDirect NETWORKS
INFORMATION IN MOTION

- Can be useful to determine log spikes across timeline

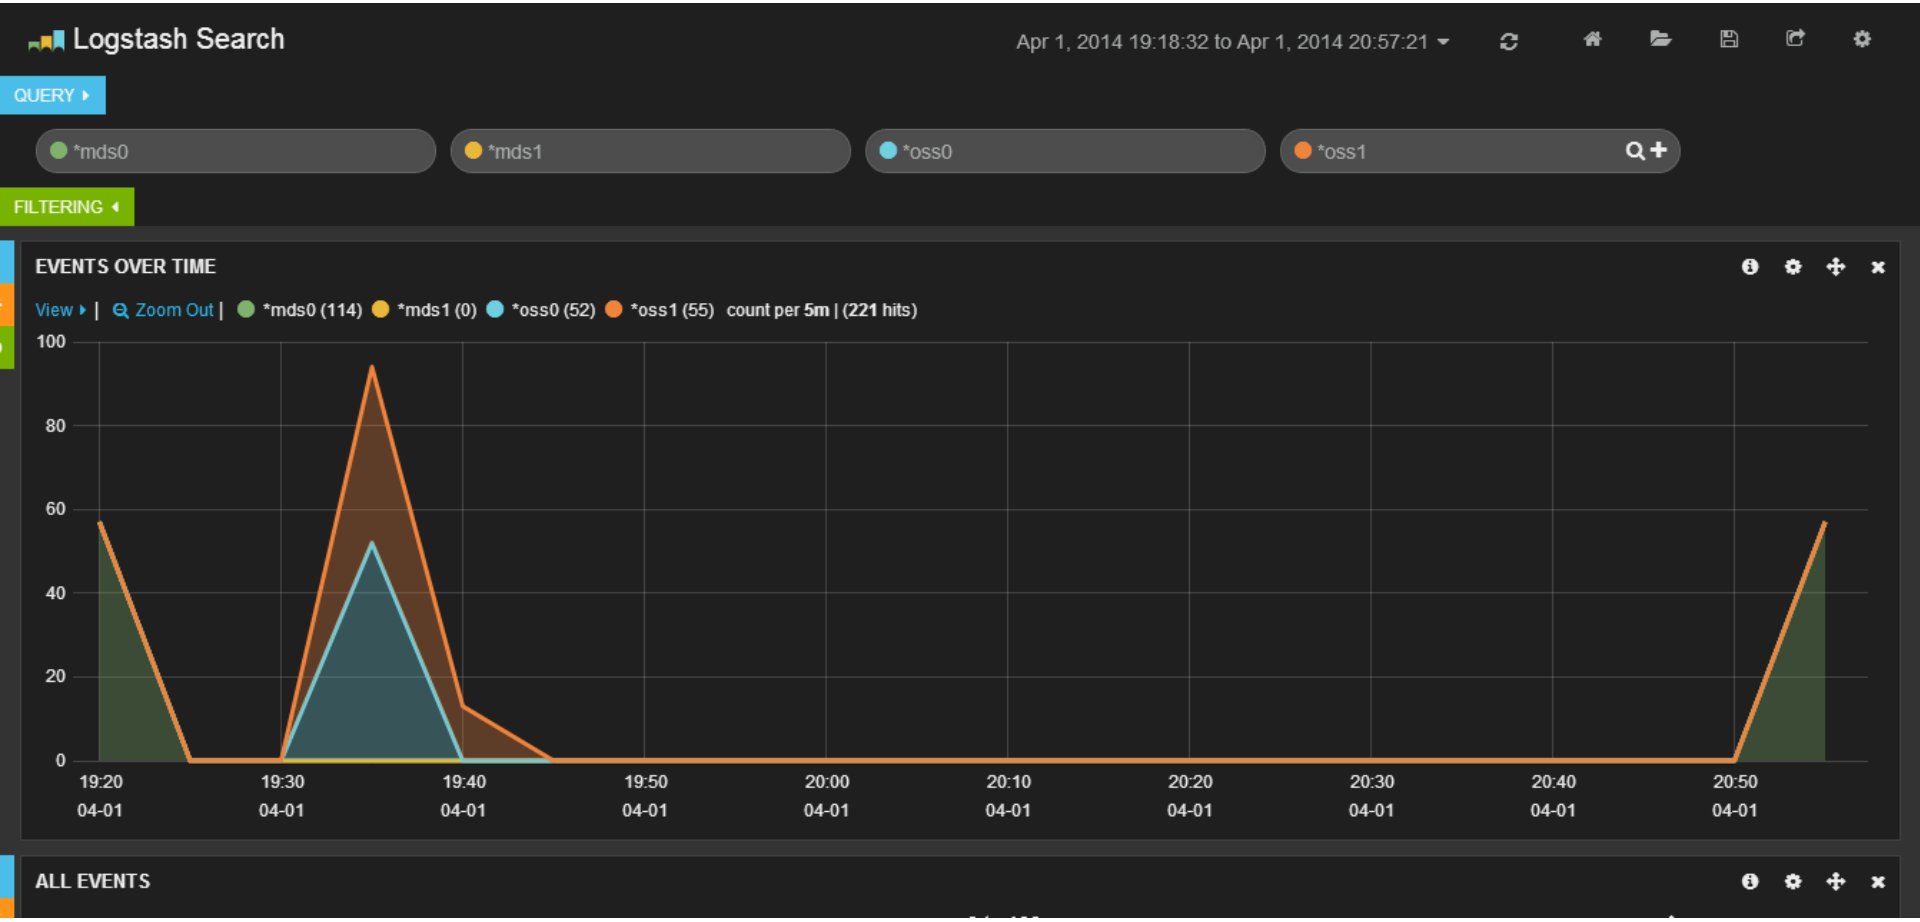

View actual logs & select columns

DataDirect NETWORKS
INFORMATION IN MOTION

- Viewing actual logs from multiple text search & select columns to view

Logstash Search

Apr 1, 2014 19:18:32 to Apr 1, 2014 20:57:21

QUERY

*mds0 *mds1 *oss0 *oss1

FILTERING

GRAPH

ALL EVENTS

Select columns to view

0 to 100 of 221 available for paging

Fields

All (44) / Current (13)

Type to filter...

@timestamp
@version
_id
_index
_type
host
message.debug_mask
message.file
message.function
message.host_pid
message.input_message
message.line
message.message1

host	message.message1	message.file	message.function
mds0	lcfg (null	mgs_llog.c	record_base()
mds0	lcfg v-MDT0000 0xcf001 mdt v-MDT0000_UUID (null	mgs_llog.c	record_base()
mds0	lcfg v-MDT0000-mdtlov 0xcf001 lov v-MDT0000-mdtlov_UUID (null	mgs_llog.c	record_base()
mds0	not reading header from 0-byte log	llog_osd.c	llog_osd_read_header()
mds0	Writing lov(v-MDT0000-mdtlov	mgs_llog.c	mgs_write_log_lov()
mds0	writing new mdt v-MDT0000	mgs_llog.c	mgs_write_log_mdt()
mds0	Set index for v-MDT0000 to 0	mgs_llog.c	mgs_set_index()
mds0	updating v-MDT0000, index=0	mgs_handler.c	mgs_target_reg()
mds0	not reading header from 0-byte log	llog_osd.c	llog_osd_read_header()
mds0	IR: current version is 2	mgs_nids.c	mgs_nidtbl_init_fs()
mds0	log v to resid 0x76/0x2 (v)	mgc_request.c	mgc_name2resid()
mds0	Creating new db	mgs_llog.c	mgs_find_or_make_fsdb()

Multiple search terms

- Logs across various oss nodes with multiple search

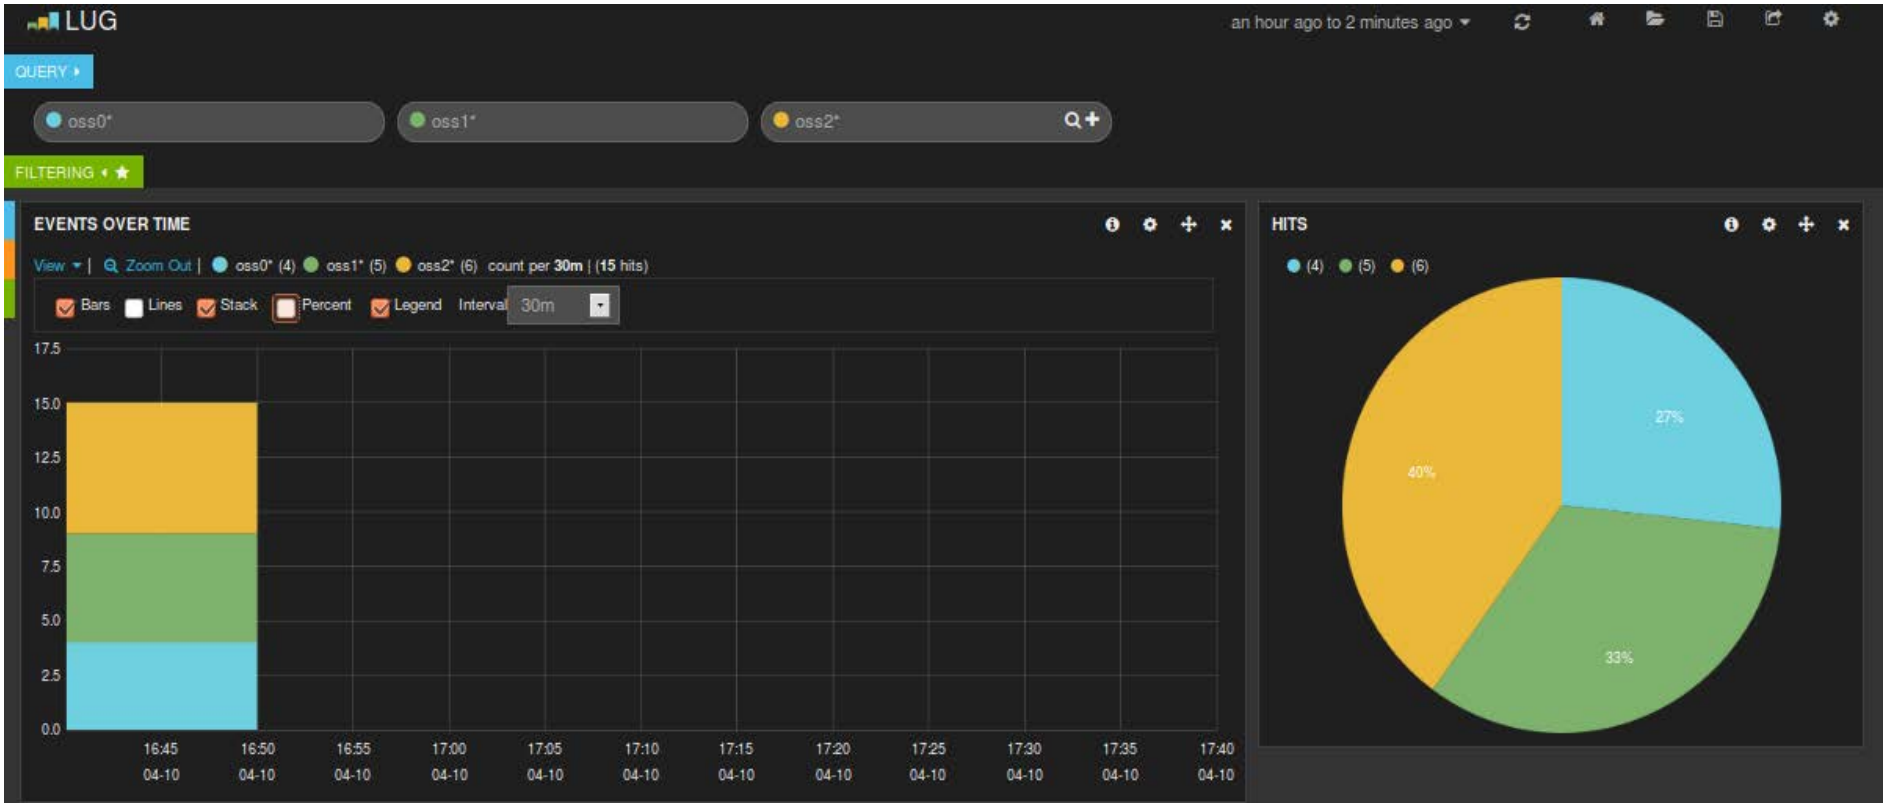

Search by timestamp & text filter

DataDirect NETWORKS
INFORMATION IN MOTION

Ex: search for all ldlm logs with LDLM_ENQUEUE event in last 3 mins on 2 specific nodes

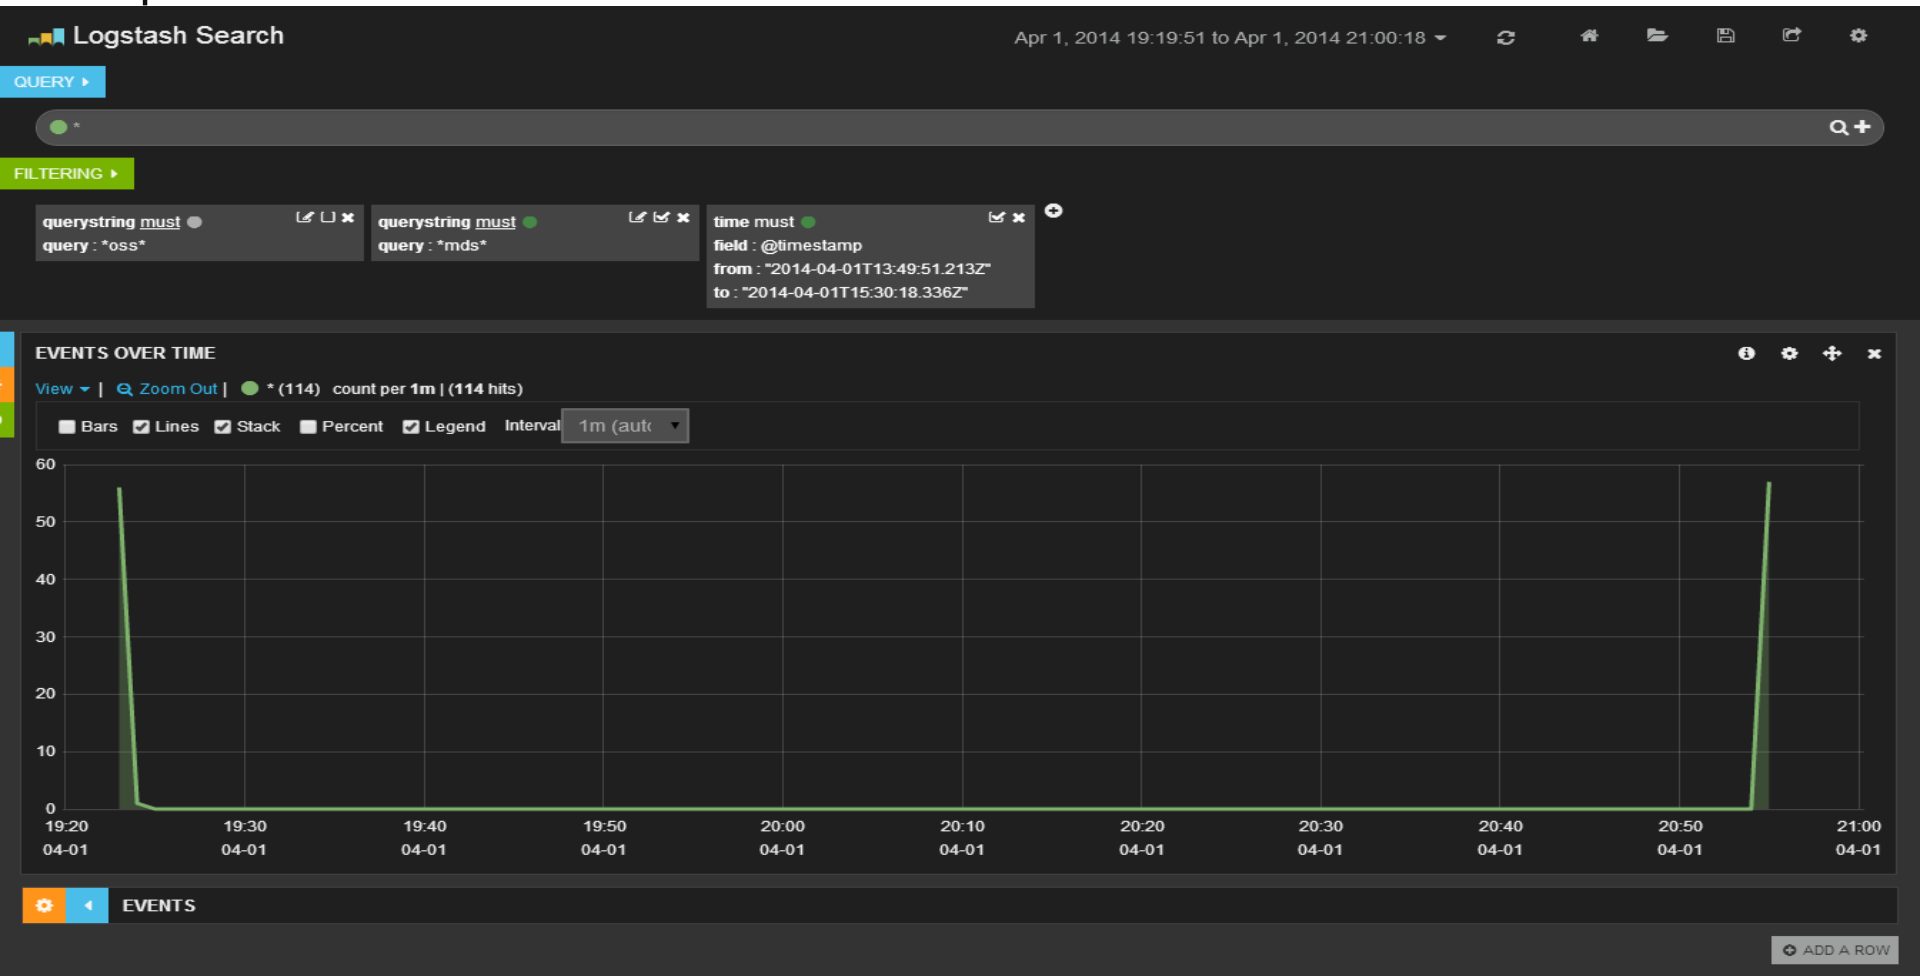

Search for specific transaction ID across logs of multiple nodes

DataDirect NETWORKS
INFORMATION IN MOTION

Lustre Log Search

Apr 1, 2014 19:19:51 to Apr 1, 2014 21:00:18

QUERY

*x1464182906028048

FILTERING

GRAPH

ALL EVENTS

0 to 1 of 1 available for paging

message.XID	message.message1	message.line	message.function	message.file
x1464182906028048	req from PID 5322 waiting for recovery: (FULL != DISCONN)	1431	ptlrpc_send_new_req()	client.c

0 to 1 of 1 available for paging

Fields

All (44) / Current (24)

Type to filter...

- @timestamp
- @version
- _id
- _index
- _type
- host
- message.debug_mask
- message.file
- message.filesystem
- message.function
- message.host_pid
- message.input_message
- message.line
- message.message1
- message.pid
- message.request_address
- message.request_mode
- message.request_type
- message.sec.used
- message.smp_processor_id
- message.stack_size
- message.subsystem
- message.transaction_id
- message.XID

ADD A ROW

Custom graphs

- Custom graphs can be added and configured to suit requirements
- For example below custom graph shows variance in number of logs across every 5 mins
- If variance is high, it indicates that number of log lines has jumped by 1366% across two 5 minute intervals - which could indicate a problem

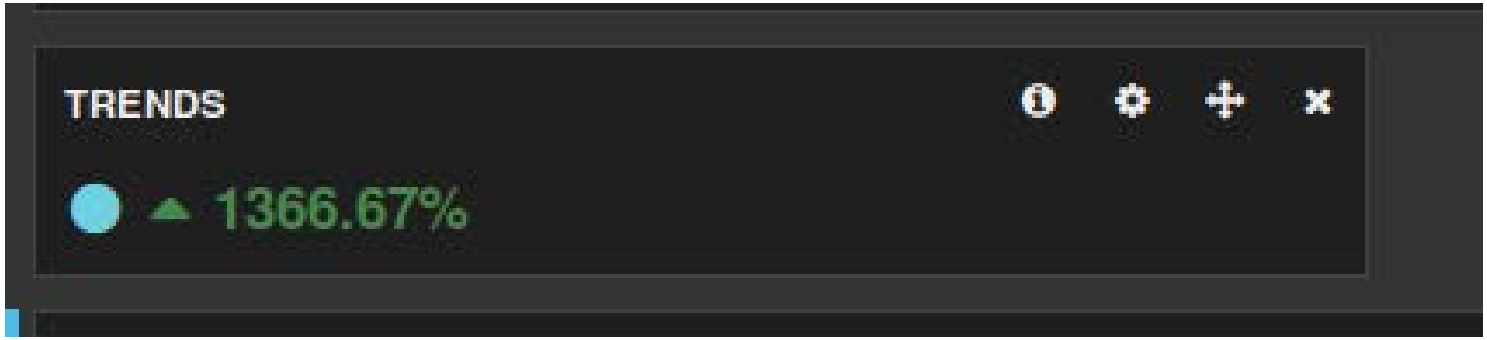

Multiple searches with column filtering

DataDirect NETWORKS
INFORMATION IN MOTION

- For example show me logs of debug_mark 0x80 which contains logs for function class_process*

Lustre Log Search

Apr 1, 2014 19:19:51 to Apr 1, 2014 21:00:18

QUERY

message.debug_mask:00000080

message.function:"class_process

FILTERING

GRAPH

ALL EVENTS

0 to 98 of 98 available for paging

Fields

All (44) / Current (18)

Type to filter...

Note These fields have been extracted from your mapping. Not all fields may be available in your source document.

- @timestamp
- @version
- geoip
- geoip.location
- host
- host.raw
- message.debug_mask
- message.debug_mask.raw
- message.file

message.message1	message.line	message.function	message.file	message.debug_mask
processing cmd: cf001	1107	class_process_config()	obd_config.c	00000080
attach type osd-ldiskfs name: MGS-osd uuid: MGS-os...	366	class_attach()	obd_config.c	00000080
finished setup of obd MGS-osd (uuid MGS-osd_UUID)	550	class_setup()	obd_config.c	00000080
Adding new device MGS-osd (ffff880307940038)	357	class_newdev()	genops.c	00000080
processing cmd: cf003	1107	class_process_config()	obd_config.c	00000080
OBD: dev 0 attached type osd-ldiskfs with refcount...	442	class_attach()	obd_config.c	00000080
connect: client MGS-osd_UUID, cookie 0xff90d6f228a...	1141	class_connect()	genops.c	00000080
processing cmd: cf001	1107	class_process_config()	obd_config.c	00000080
attach type mgs name: MGS uuid: MGS	366	class_attach()	obd_config.c	00000080
OBD: dev 1 attached type mgs with refcount 1	442	class_attach()	obd_config.c	00000080
Adding new device MGS (ffff880314796078)	357	class_newdev()	genops.c	00000080
processing cmd: cf003	1107	class_process_config()	obd_config.c	00000080

DataDirect™ NETWORKS
INFORMATION IN MOTION™

Questions?

LUG14

Installing Logstash

DataDirect NETWORKS
INFORMATION IN MOTION

- Prerequisites
  - Java
- Installing Logstash
  - $ curl -O https://download.elasticsearch.org/logstash/logstash/logstash-1.4.0.tar.gz
  - $ tar zxvf logstash-1.4.0.tar.gz
  - $ cd logstash-1.4.0
- Lets try out
  - $ bin/logstash -e 'input { stdin { } } output { stdout {} }'
    - hello world
    - 2013-11-21T01:22:14.405+0000 0.0.0.0 hello world

Installing Elasticsearch

- $ curl -O https://download.elasticsearch.org/elasticsearch/elasticsearch/elasticsearch-1.0.1.tar.gz
- $ tar zxvf elasticsearch-1.0.1.tar.gz
- $ cd elasticsearch-1.0.1/
- .$ /bin/elasticsearch