

# Grid Site Monitoring and Log Processing using ELK

Alexandr Mikula

Institute of Physics  
of the  
Czech Academy of Sciences

Coauthors: D. Adamová, M. Adam, J. Chudoba, V. Říkal, J. Švec



Akademie věd  
České republiky



# Computational centre of IoP CAS

- Projects

- Tier 2 centre for ATLAS & ALICE
- Other projects NOVA (OSG site), Auger, CTA,...
- Czech national grid (3 clusters for computing)

- Software

- OS Scientific Linux 6 + some SL5, Debian 7 & 8
- Grid services as DPM, XrootD, BDII,...
- Configuration by Puppet + some with CFEngine
- Mostly latest stable packages from distribution and EPEL



# Computational centre of IoP CAS

- Installation by PXE and updates managed with Spacewalk
- Hardware (370+ devices)
  - WN's - mixed variety of HW, total of 227 nodes, 5000 cores
  - Storage - 12 dpmpools (3PB), 5 xrootds (1.25PB), few local nfs export servers (100's TB)
  - Another **4 xrootds offsite (0.728 ms/10 km away) at INP Řež (~340TB)**
  - Other necessary site services as DNS, DHCP, ...

# Monitoring tools

- Monitoring

- Nagios with NRPE, Munin, Ganglia, Netflow, Observium, Spacewalk, Puppet dashboard, custom scripts, and grid tools as Panda,...
- It is mix of virtual and real machines of varying age and performance

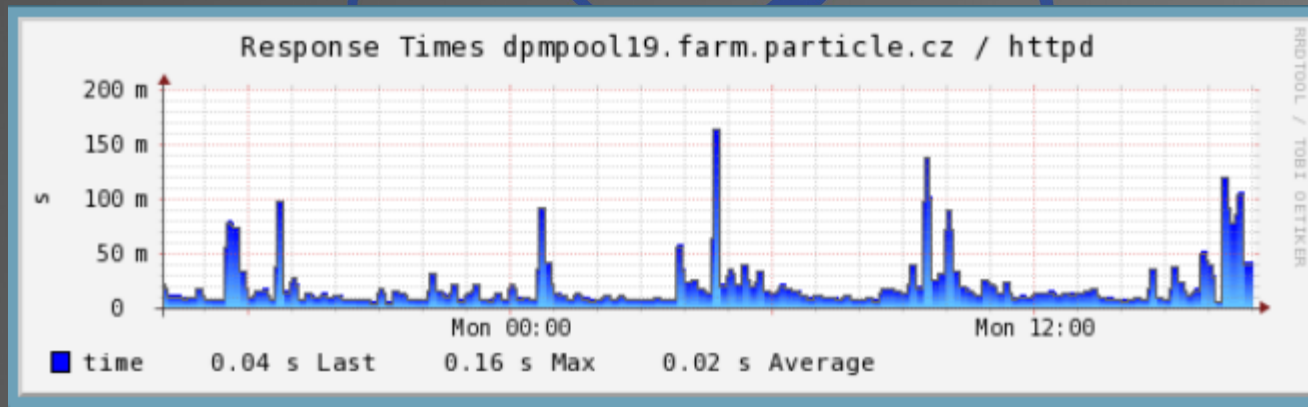




# Nagios

# Nagios

- Our monitoring workhorse
- Used mainly for service availability checking
- Extended with custom and added scripts
































# Nagios

# Nagios

- Used mainly through CheckMK interface
- Custom report sending script

CRIT					
Host	Service	Status detail	Age	Checked	Icons
elastic02.farm.particle.cz	Check_Elastic_Search	(Service Check Timed Out)	5 min	5 min	  
rubus10.farm.particle.cz	PAKITI_NSCA	(Service Check Timed Out)	27 hrs	27 hrs	    
ibis13.farm.particle.cz	DISK_ROOT	DISK CRITICAL - free space: / 6 GB (47% inode=85%):	2016-06-18 21:58:44	4 min	  
UNKN					
Host	Service	Status detail	Age	Checked	Icons
rubuk02.farm.particle.cz	MEGARAID	UNKNOWN: error: /opt/MegaRAID/MegaCli /MegaCli64 does not exist	2016-06-18 17:53:45	4 min	  
merab.fzu.cz	CCISSRAID	RAID UNKNOWN - /usr/sbin /hpacucli did not execute properly : /usr/local /sbin/check_cciss_raid.sh: line 203: /usr/bin/sudo: Permission denied	2016-06-09 03:32:56	4 min	  
merab.fzu.cz	CONFIGURATOR	NRPE: Unable to read output	2016-06-09 03:23:36	4 min	    
saul.fzu.cz	CONFIGURATOR	NRPE: Unable to read output	2016-04-27 19:46:46	4 min	    
		RAID UNKNOWN - /usr/sbin /hpacucli did not execute			




# Munin



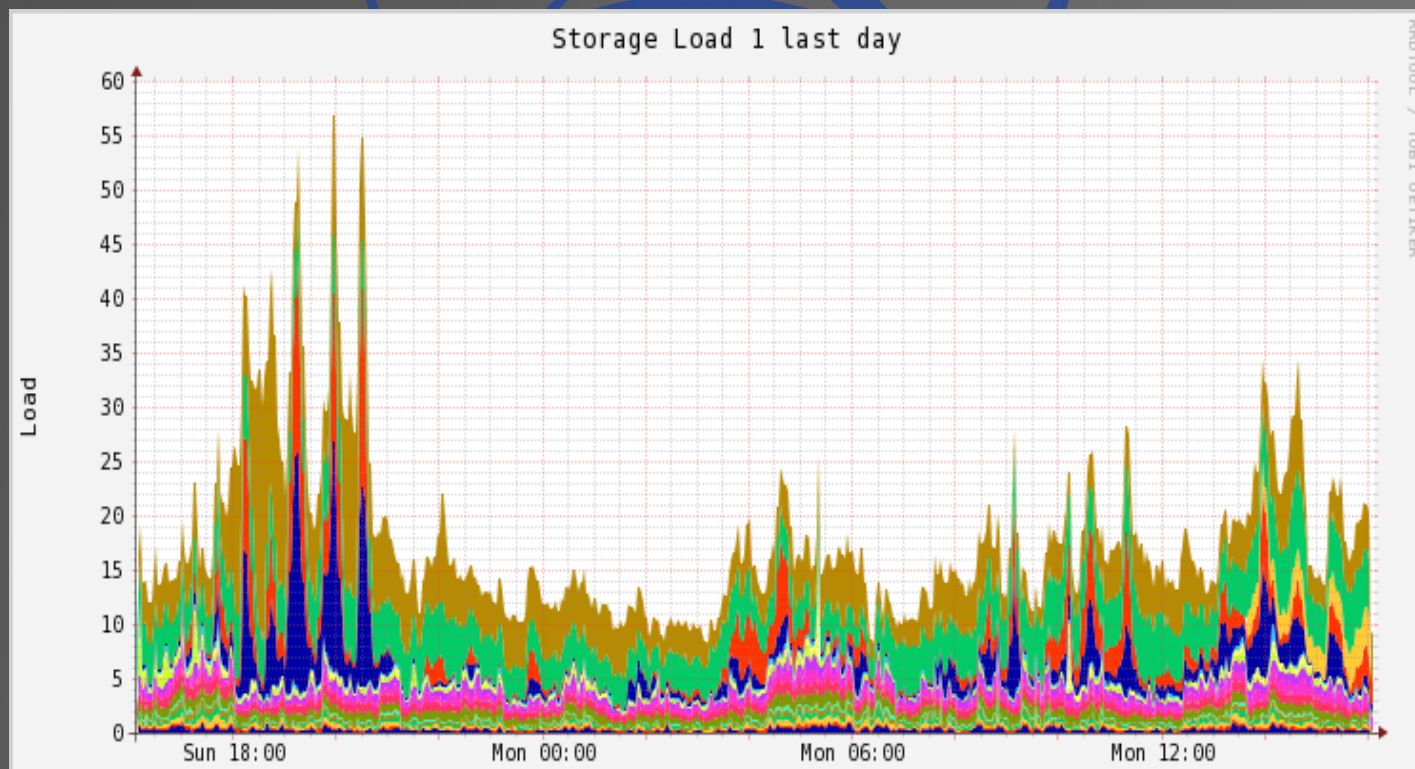
- Main plotting tool of our farm



# Ganglia



- Great alternative to Munin
- Easy to create aggregate graphs






# Observium




- Observium

- After initial set-up it finds all network appliances itself
- Used for monitoring router and switches
- Very handy on debugging network problems


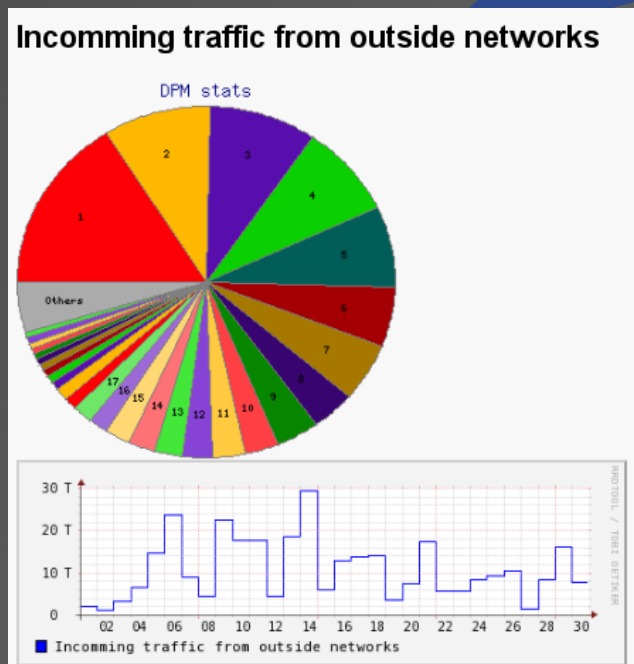


# Netflow



- Netflow

- Used only on chosen nodes, since it needs kernel module which needs to be compiled
- Visualisation of from/to where traffic flows



# Elk





# ELK stack



## Highlights:

- Highly scalable solution for text processing
- Best way for 21<sup>st</sup> century log analysis
- Secure (paid version)
- Java based (platform independent)
- REST api
- Built in redundancy

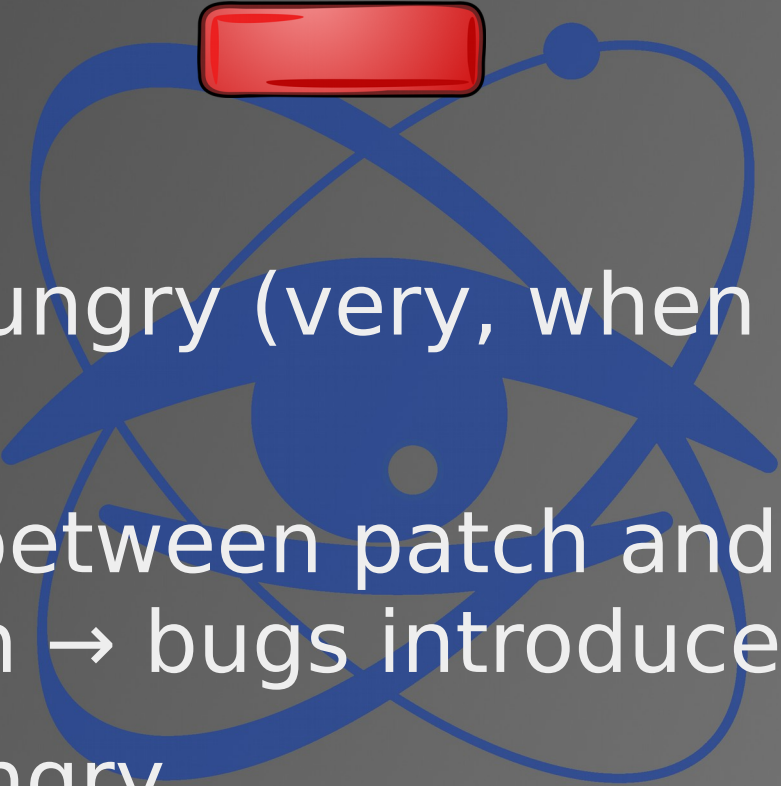







# ELK stack



## Negatives:

- Java
  - Resource hungry (very, when compared to rsyslog)
  - Small gap between patch and its stabilisation → bugs introduced very fast
  - Storage hungry
- 
- 
- 
- 

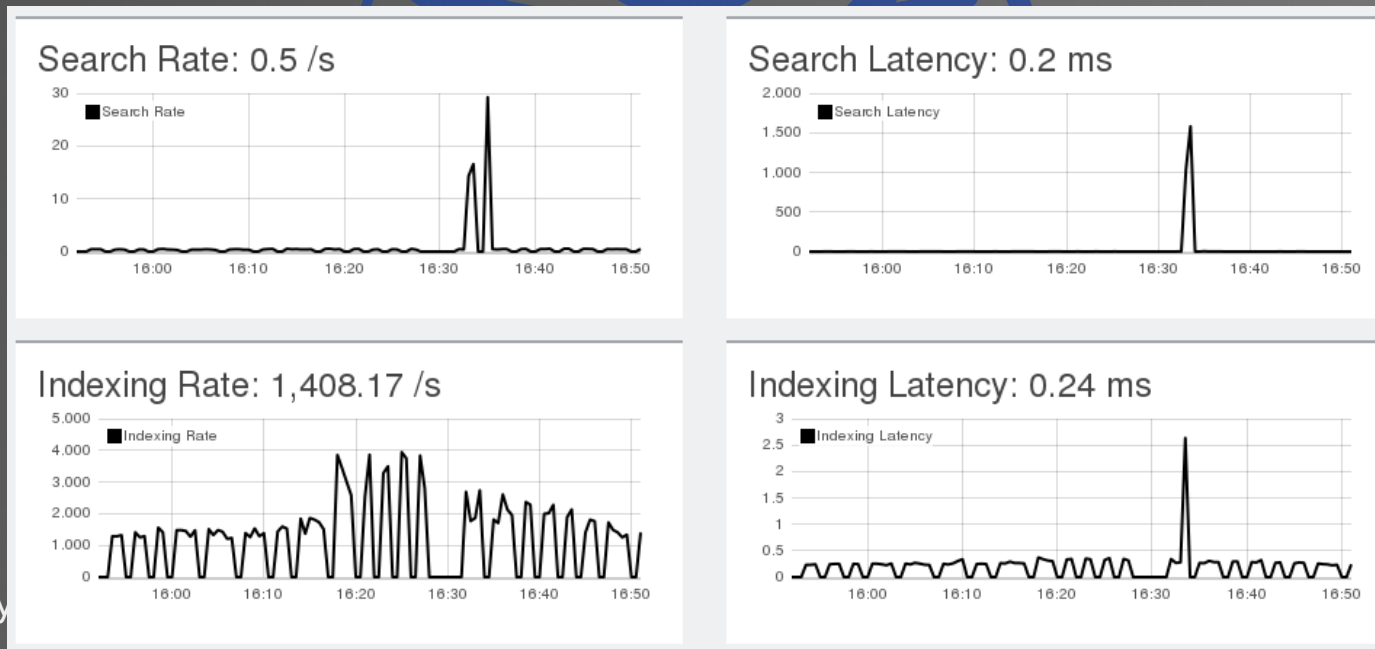


# ELK stack at our site

- 7 Logstash/ElasticSearch nodes
- One Kibana/ES master node (virtual)
- ~18TB of space
- 144GB ram for ES HEAP
- 68 cores
- Build on old HW, from 8 to 5 years old

# ELK statistics

- Processing messages from 361 devices
- ~40.000.000 entries per day
- ~30GB of space consumed daily
- Usually less than 50% HEAP used



# Data sources

- Rsyslog
  - All non puppetized machines
  - Network devices
- Filebeat
  - Tails log files
  - Ships messages to Logstash (using SSL)



# Data processing

- Logstash daemon
  - Many input codecs (ES, rsyslog, beats, ...)
  - Powerful REGEX matching
  - Scalability
  - Indexes data into ES



# Data storage

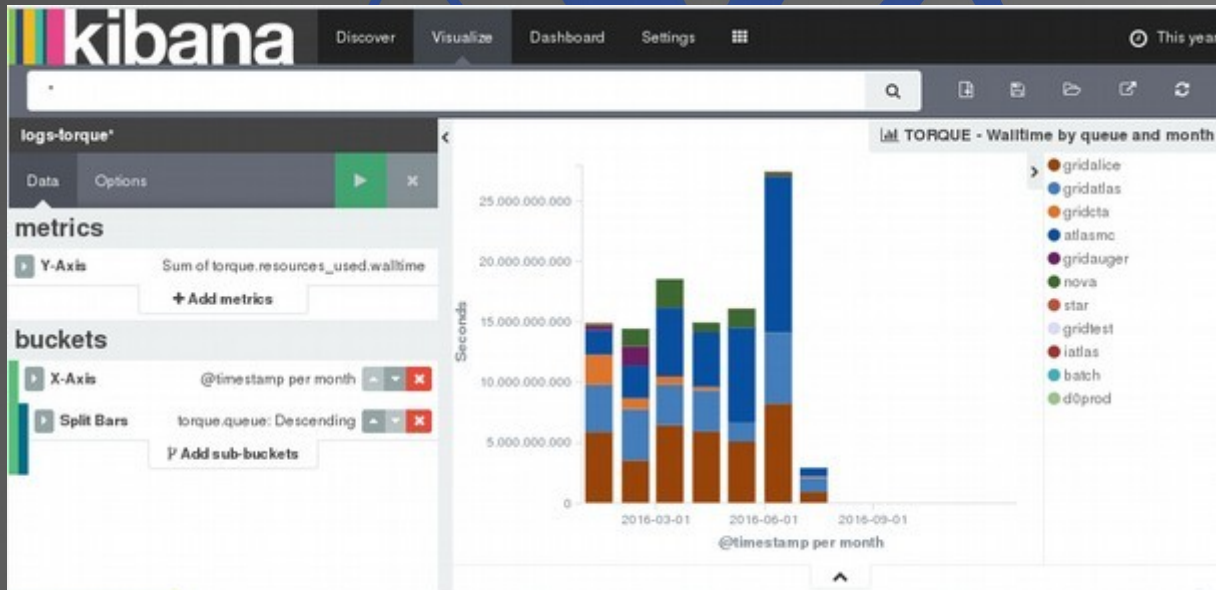
- Elasticsearch
  - Accepts data structured as JSON
  - Runs full text analysis on (default) all fields
  - Stores it compressed on drive
  - Backbone of ELK






# Data querying

- Directly through REST api
- Kibana frontend
  - Full text searching
  - Creating graph templates for data queries

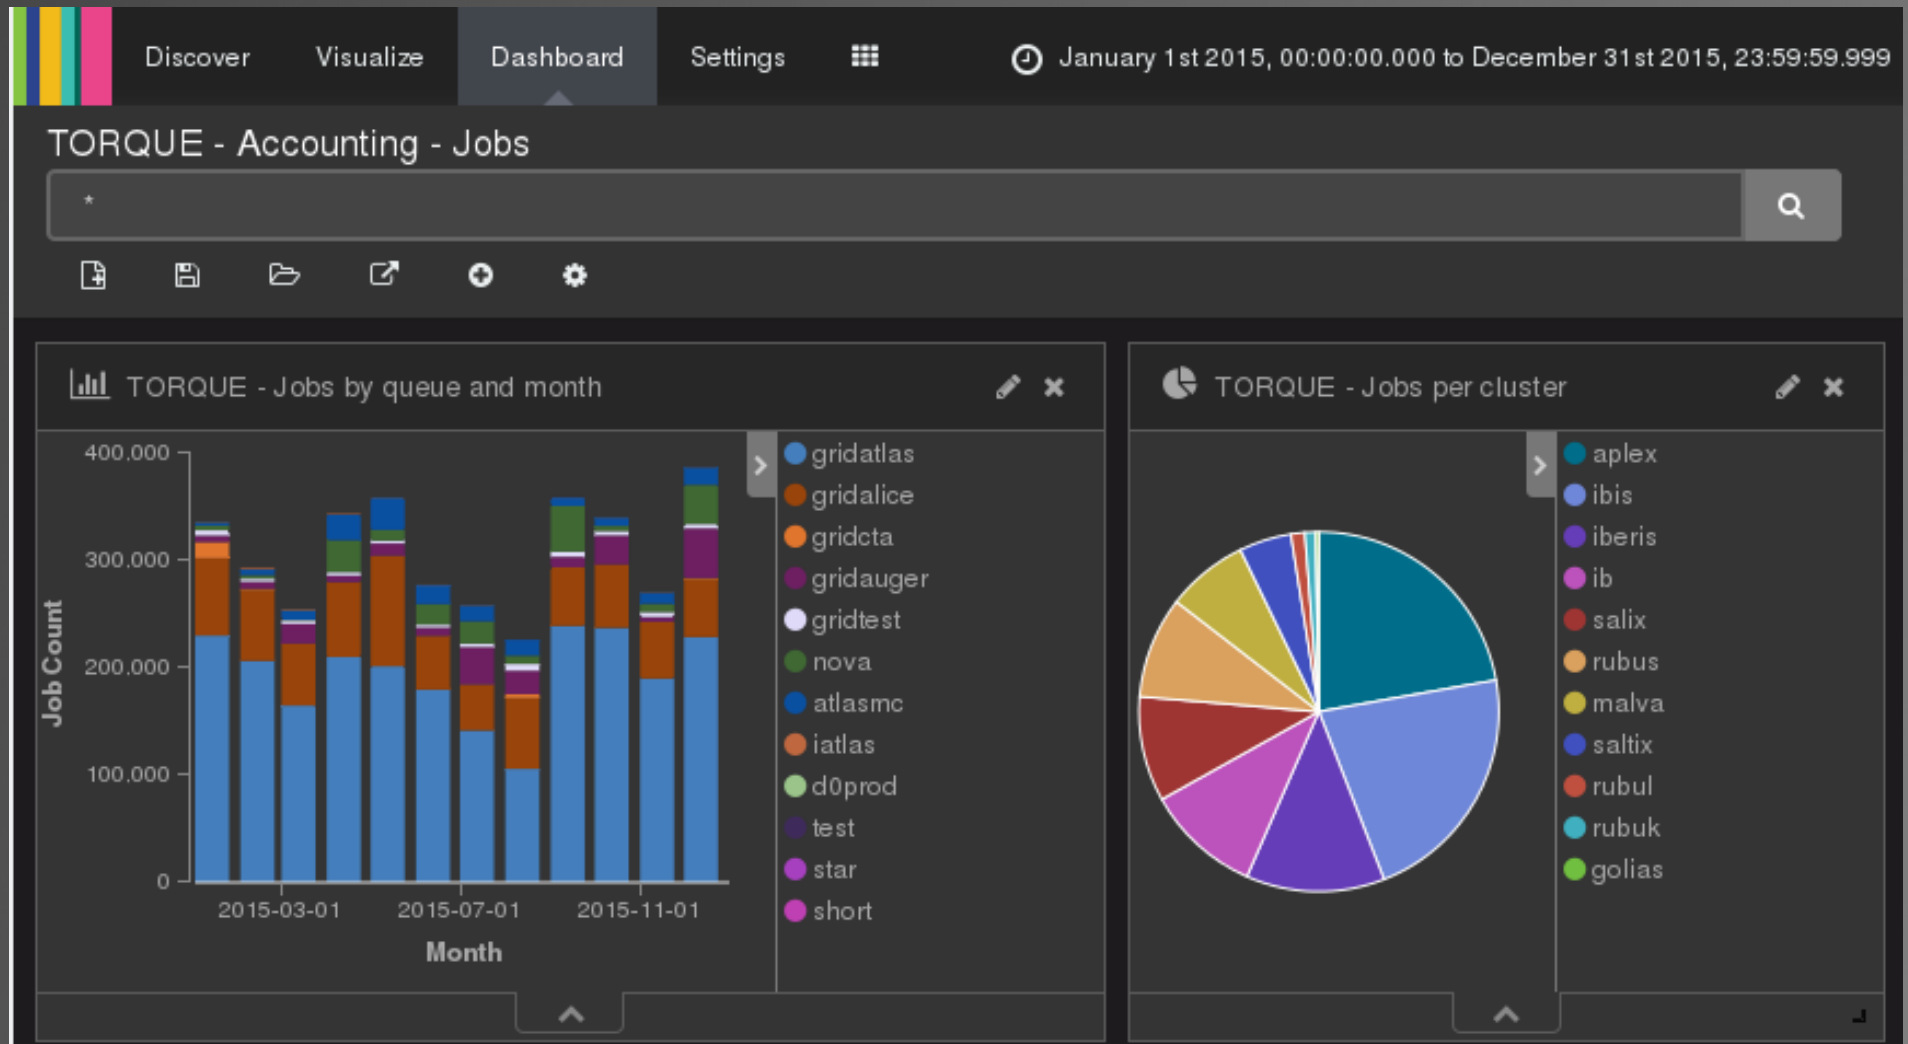


# Gotchas

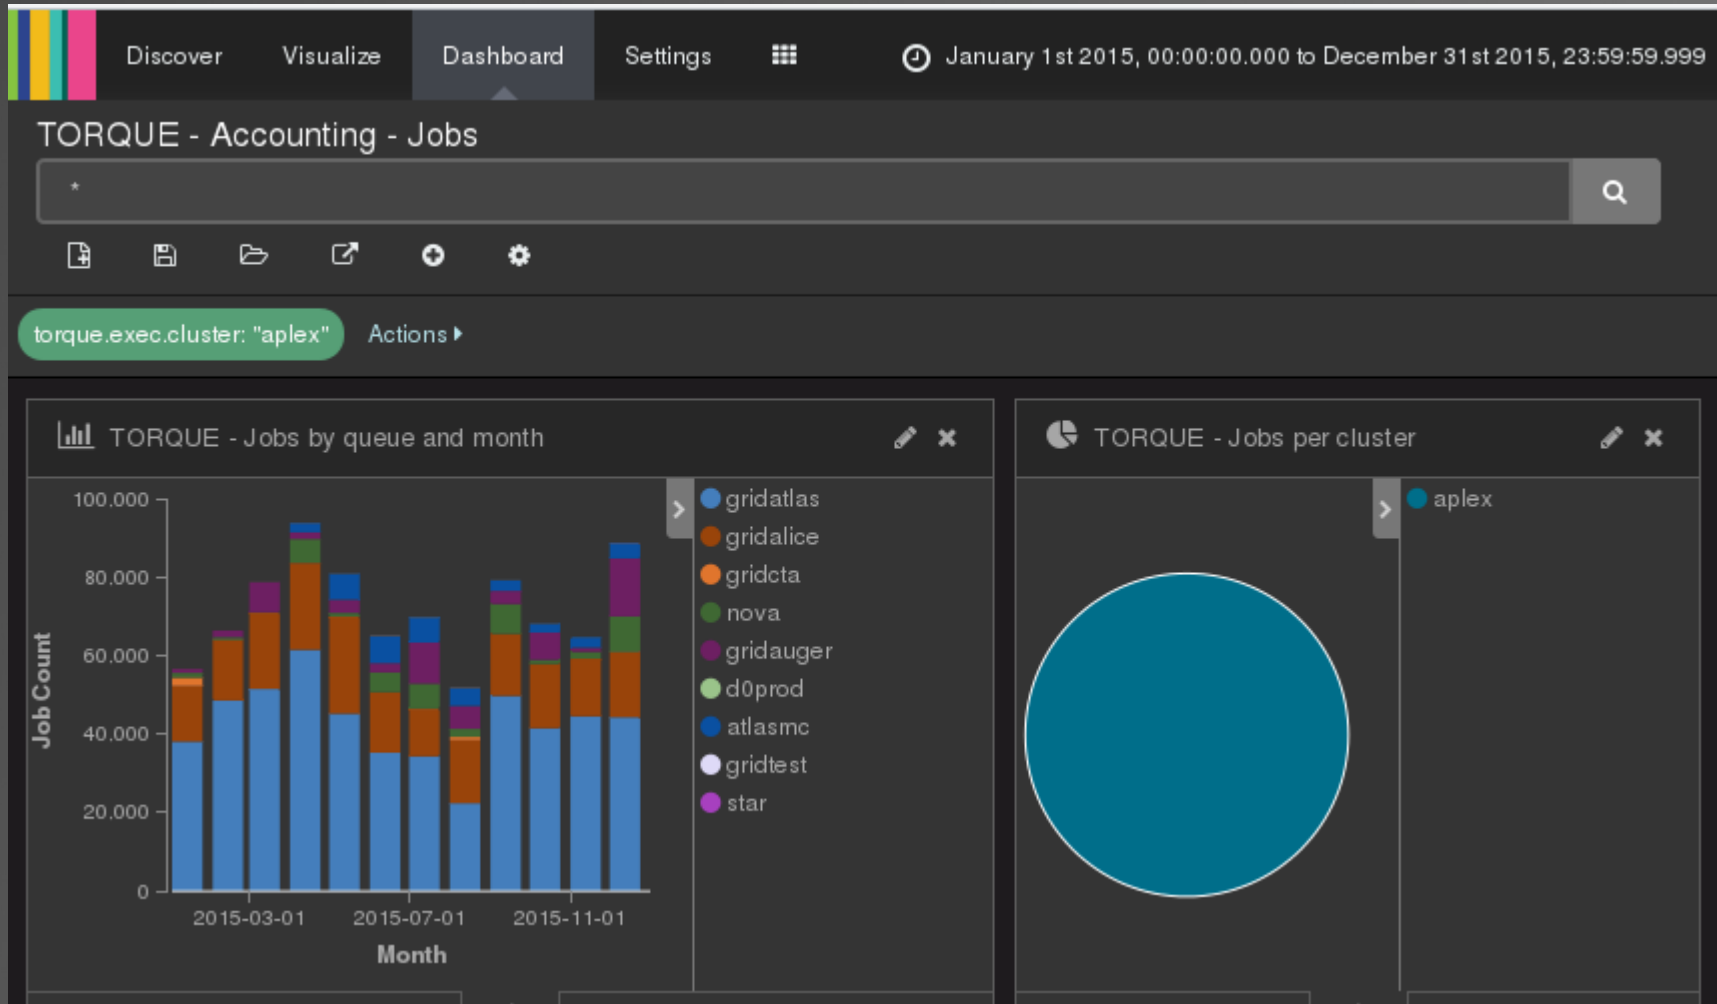
- Definition of variable types
- Lot of performance tuning options
- Kibana time-outs for resource intensive queries




# One click data filtering



# One click data filtering



# Few interesting outputs from ELK





# Few interesting outputs from ELK

