

HTCondor @ ScotGrid Glasgow Cluster Monitoring



Outline

- Glasgow Tier2 site description
- ARC-CE + HTCondor batch system
- Cluster Monitoring
- Outlook

ScotGrid Glasgow:

Gareth Roy (general admin),
Samuel Skipsey (storage),
Gordon Stewart (networking)
Emanuele Simili (computing)

Group leader:

Prof. David Britton



UKI-SCOTGRID-GLASGOW

- Part of the GridPP collaboration providing resources to the Worldwide LHC Compute Grid (WLCG).
- One of 19 institutions comprising 4 distributed Tier-2 sites (SCOTGRID, NORTHGRID, SOUTHGRID and LT2).
- Part of the SCOTGRID Distributed Tier-2 including Glasgow, Edinburgh and Durham Universities.



Current Capacity

At present (*), ScotGrid Glasgow consists of:

~ 6000 CPU cores

4.9 Pb storage (3.2 Pb CEPH + 1.7 Pb DPM)

160 Gb/s internal network bandwidth

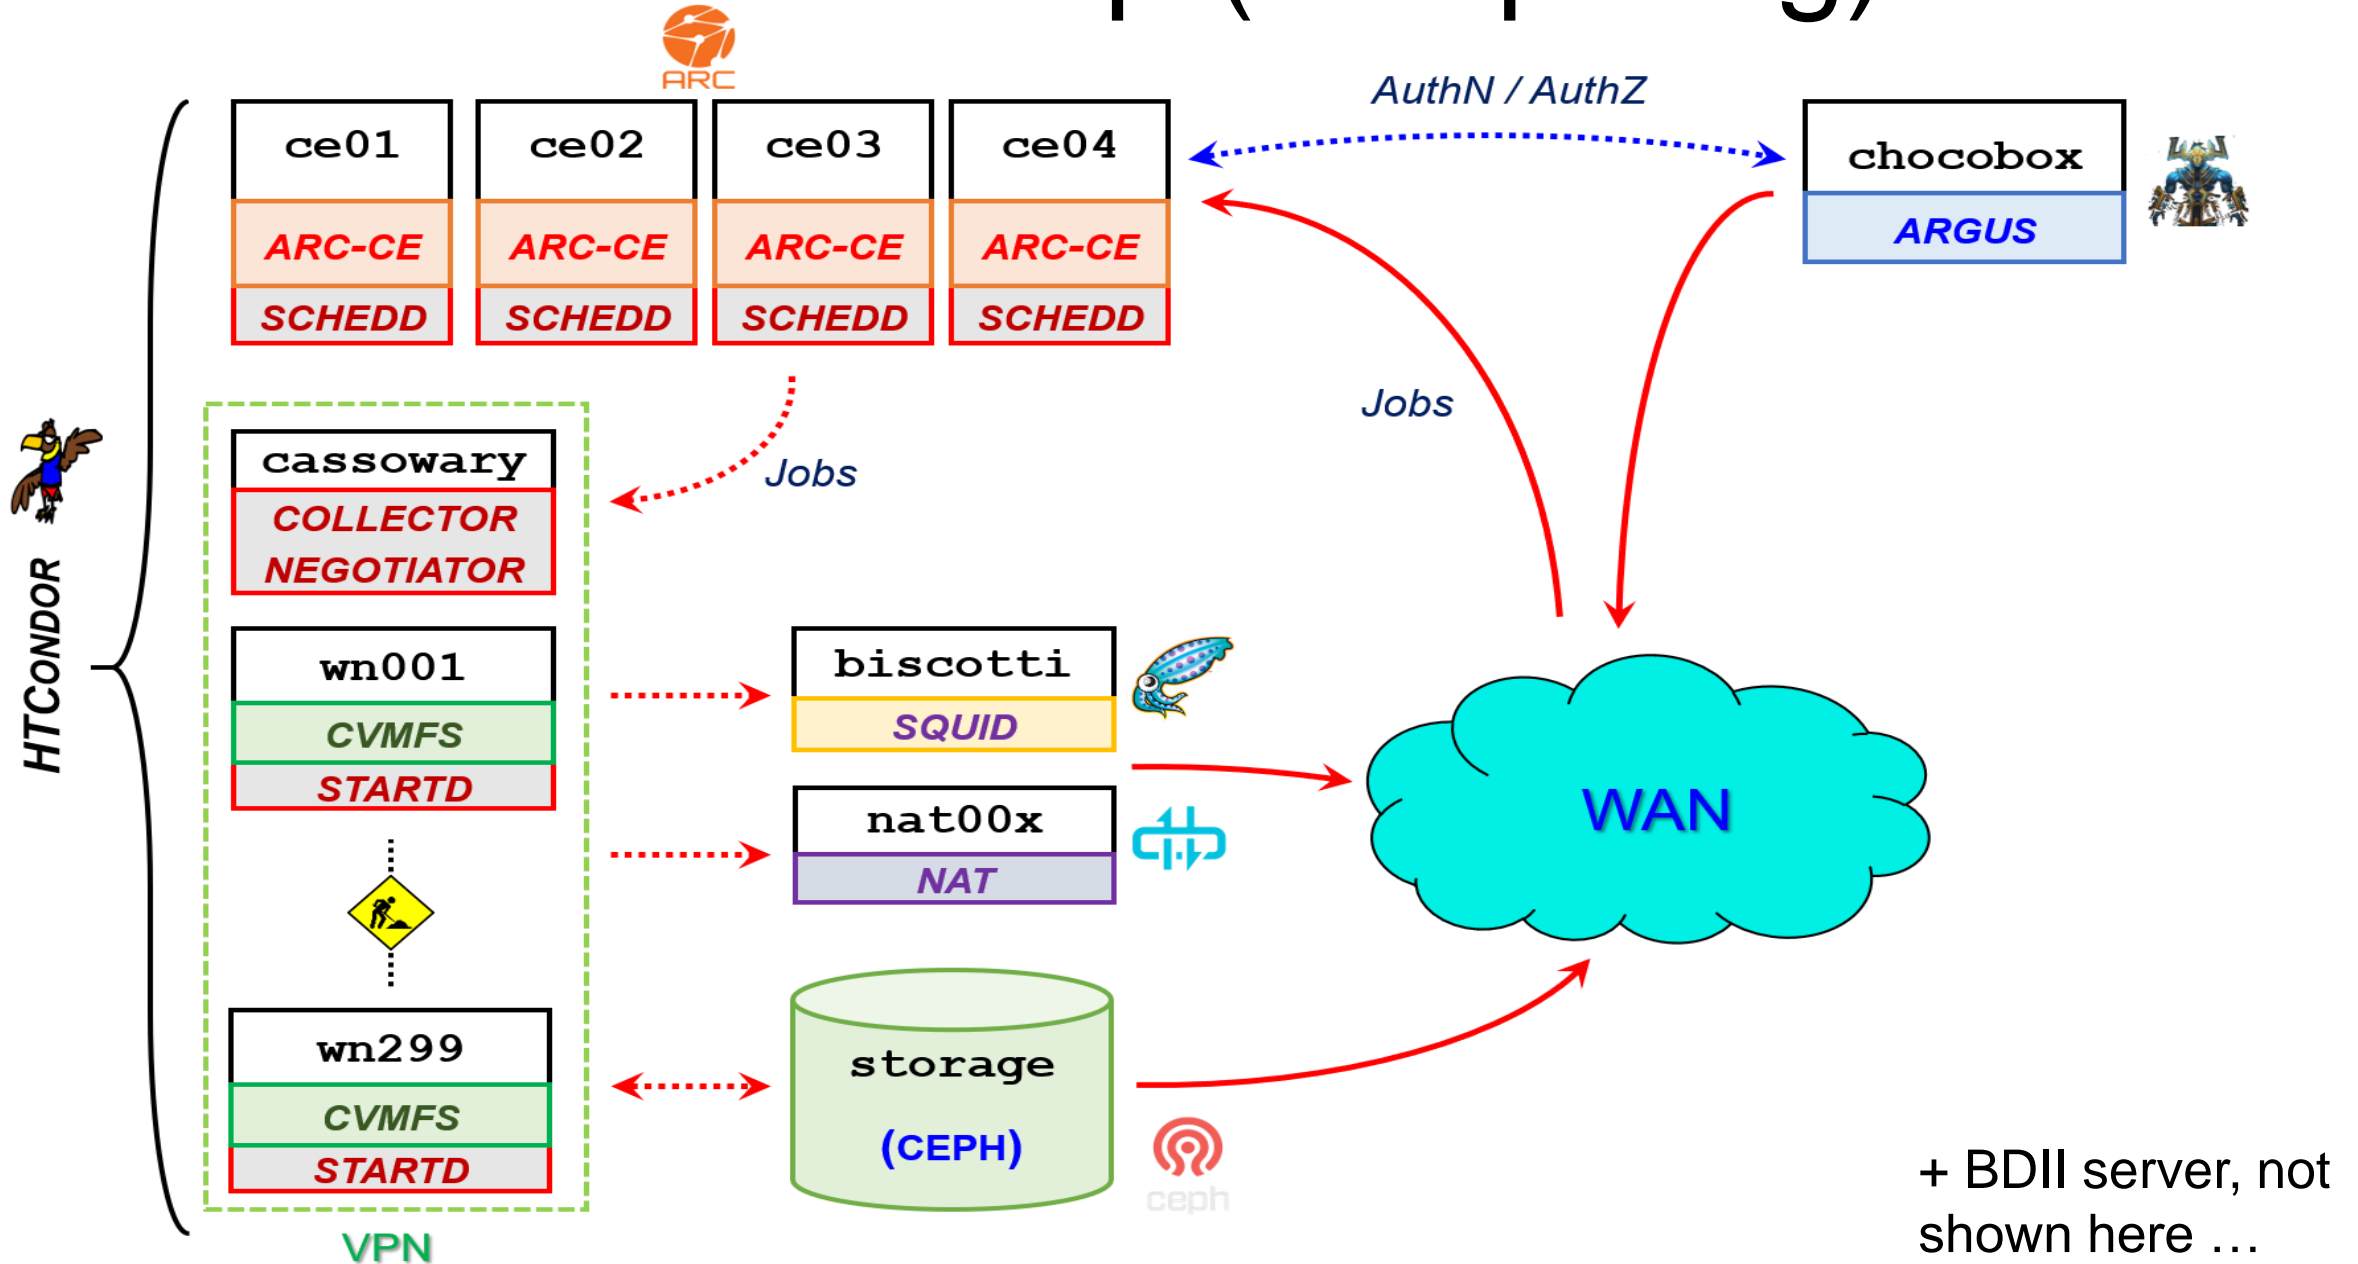
For a rough total HepSpec of about 61 KHS06

() we are undergoing a long sought relocation to our new fancy Data Center (Soughfield building) completed in late 2019.*

Relocation started slowly in early 2020, then the Covid-19 slowed down our efforts even more ...



Cluster Map (computing)



HTCondor Batch System

Our **HTCondor** batch system is built on **CentOS7**

- We use **PPE PiXiE** for provisioning (<https://github.com/danderson/netboot>)
- We use **Ansible** for all software install and configuration management

Name	Role	HTCondor Daemons	fqdn (int)	IPv4 (int)	fqdn (ext)	IPv4 (ext)
cassowary	Condor Manager	master , collector , negotiator , startd	cassowary.beowulf.cluster	10.x.x.x	no	no
ce01	ARC-CE	master , schedd	ce01.beowulf.cluster	10.x.x.1	ce01.gla.scotgrid.ac.uk	130.x.x.x
...						
ce04	ARC-CE	master , schedd	ce04.beowulf.cluster	10.x.x.4	ce04.gla.scotgrid.ac.uk	130.x.x.x
chocobox	ARGUS		svr029.beowulf.cluster	10.x.x.x	svr029.gla.scotgrid.ac.uk	130.x.x.x
wn001	WorkerNode	master , startd	wn001.beowulf.cluster	10.x.1.1	no	no
wn002	WorkerNode	master , startd	wn002.beowulf.cluster	10.x.1.2	no	no
...			...			
wn299	WorkerNode	master , startd	wn299.beowulf.cluster	10.x.2.44	no	no

HTCondor daemons installed ...

Node type	Daemons
Worker Node	MASTER, STARTD
Manager Node	COLLECTOR, MASTER, NEGOTIATOR, SCHEDD
CE Node	MASTER, SCHEDD

Cluster Monitoring

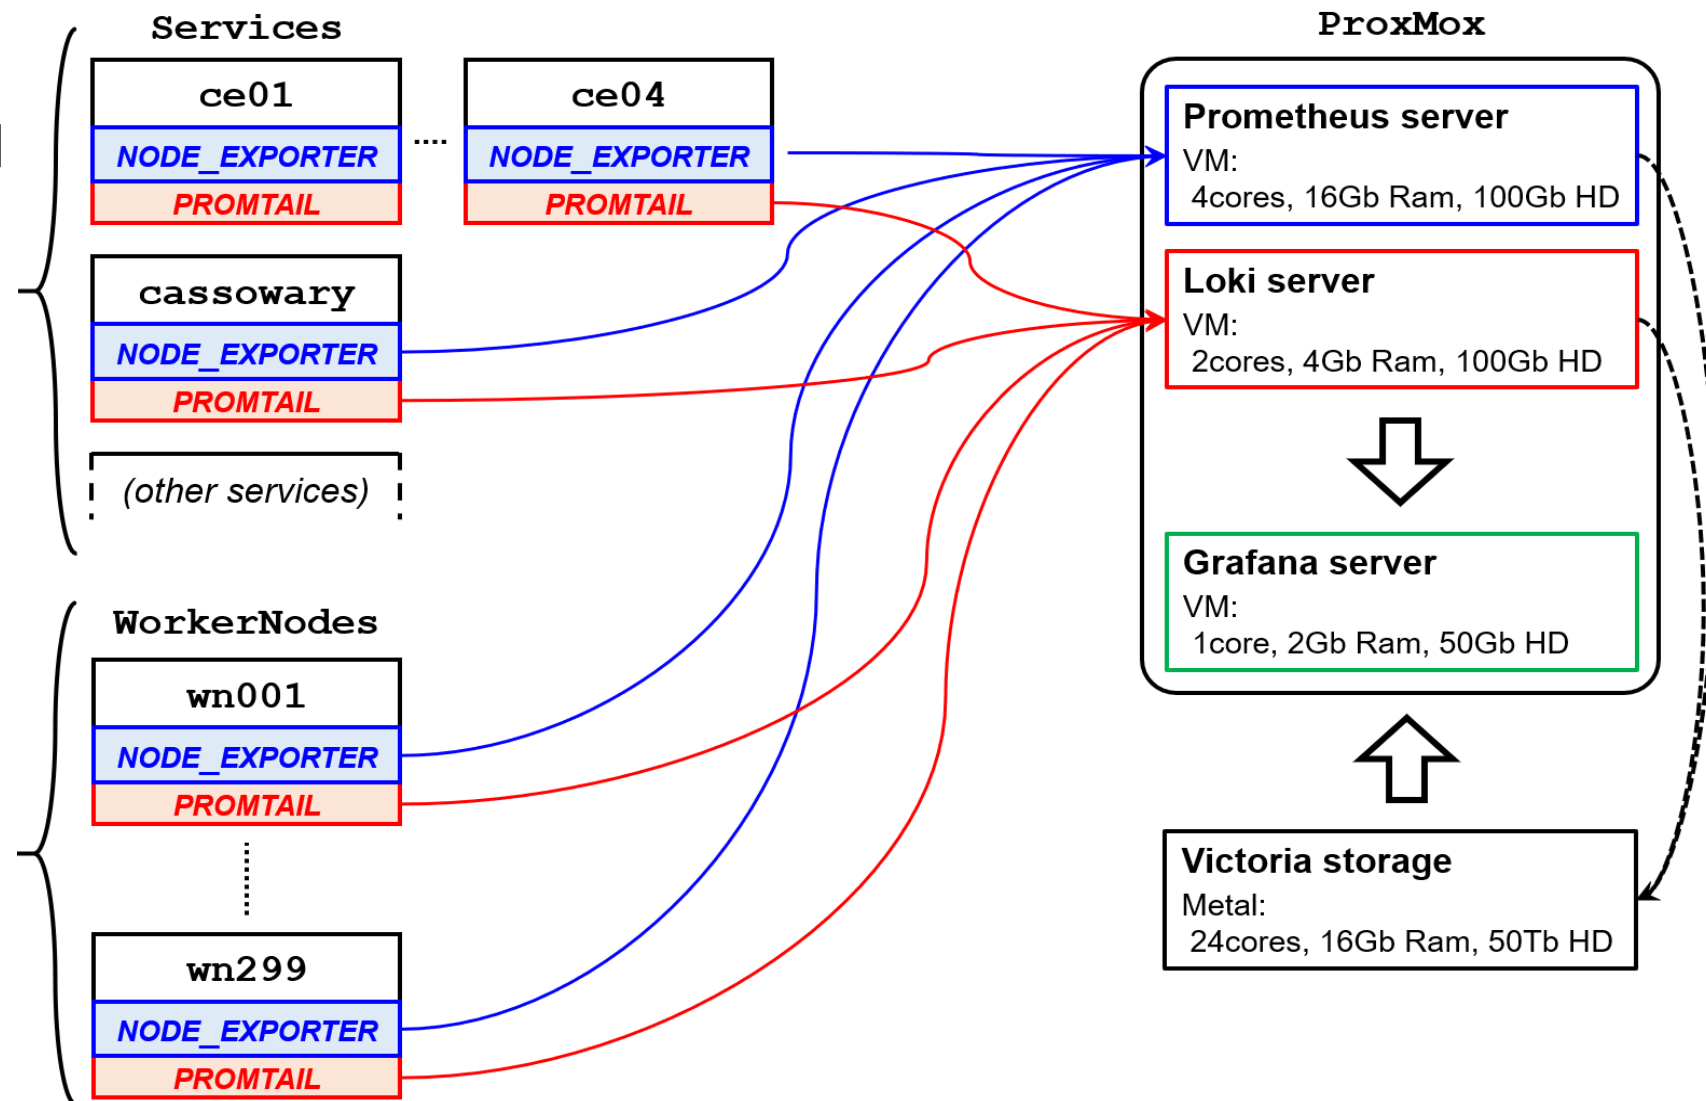
Our monitoring is built on **Prometheus** and **Grafana** ...

Metrics are exported by **node_exporter** and collected by **Prometheus** (on a dedicated VM).

Logs are exported by **PromTail** and collected by **Loki** (dedicated VM).

Grafana pulls data from both servers and provides the tools for querying and building colorful graphs and dashboards.

VictoriaMetrics archives the collected data into a large storage server. Archived data can be queried by Grafana.



HTCondor Monitoring

Nodes provide hardware metrics and job info (from HTCondor):

- **node_exporter** produces standard and custom metrics

syntax: `whatever_metric{slot=...,vo=...,...} value`

export folder: `/var/lib/node_exporter/textfile_collector/`

- few **custom scripts** periodically query HTCondor and format the output as above

`ce_get_info.sh`

→ runs as cron job on the ARC-Ces

`node_get_condorinfo.sh`

→ runs as cron job on every workernode

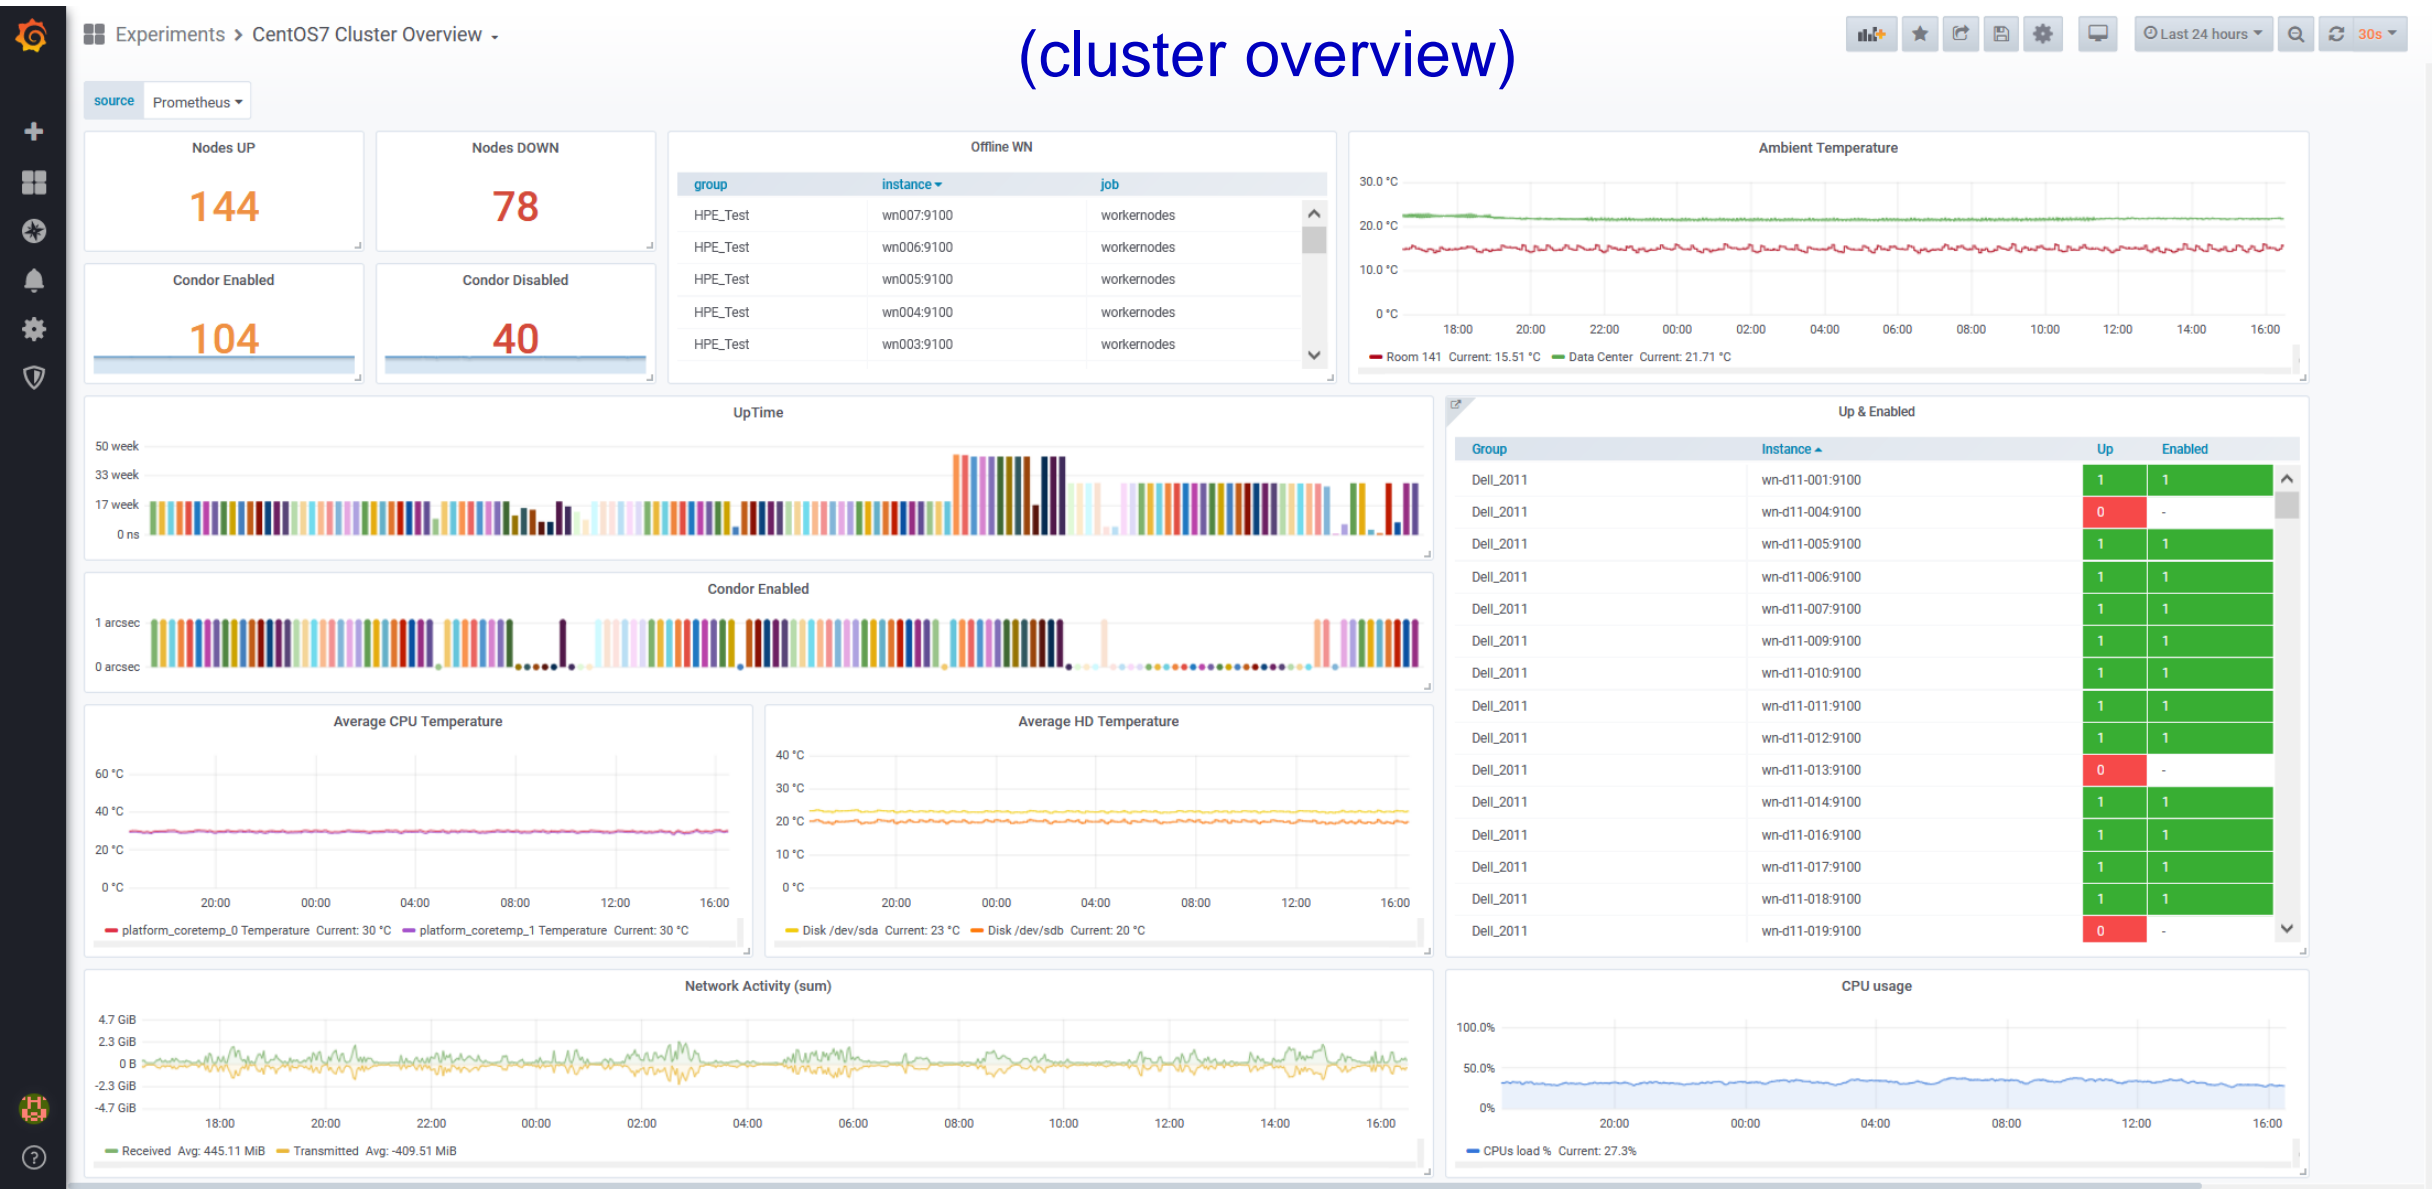


Basically, a call to **condor status -startd**
with **-autoformat:t** and parse the output ...

(Name CPUs Memory LoadAvg State Activity TotalJobRunTime RemoteUser ClientMachine TotalCPUs TotalMemory)

Grafana Dashboards

(cluster overview)



I keep this page constantly opened in my browser ...

Grafana Dashboards

Experiments > CentOS7 Nodes Monitor

(workernodes)

Last 24 hours

HTCondor

Condor Enabled



1 = Enabled, 0 = Disabled


Running Jobs

Time	VO	jobs
2020-09-19 13:15:59	ATLAS	40

Reboots


No data points

HTCondor Jobs/Slot by VO




Running jobs per slot and VO


Jobs run-time by VO




Load by VO & CE



CPU usage by VO

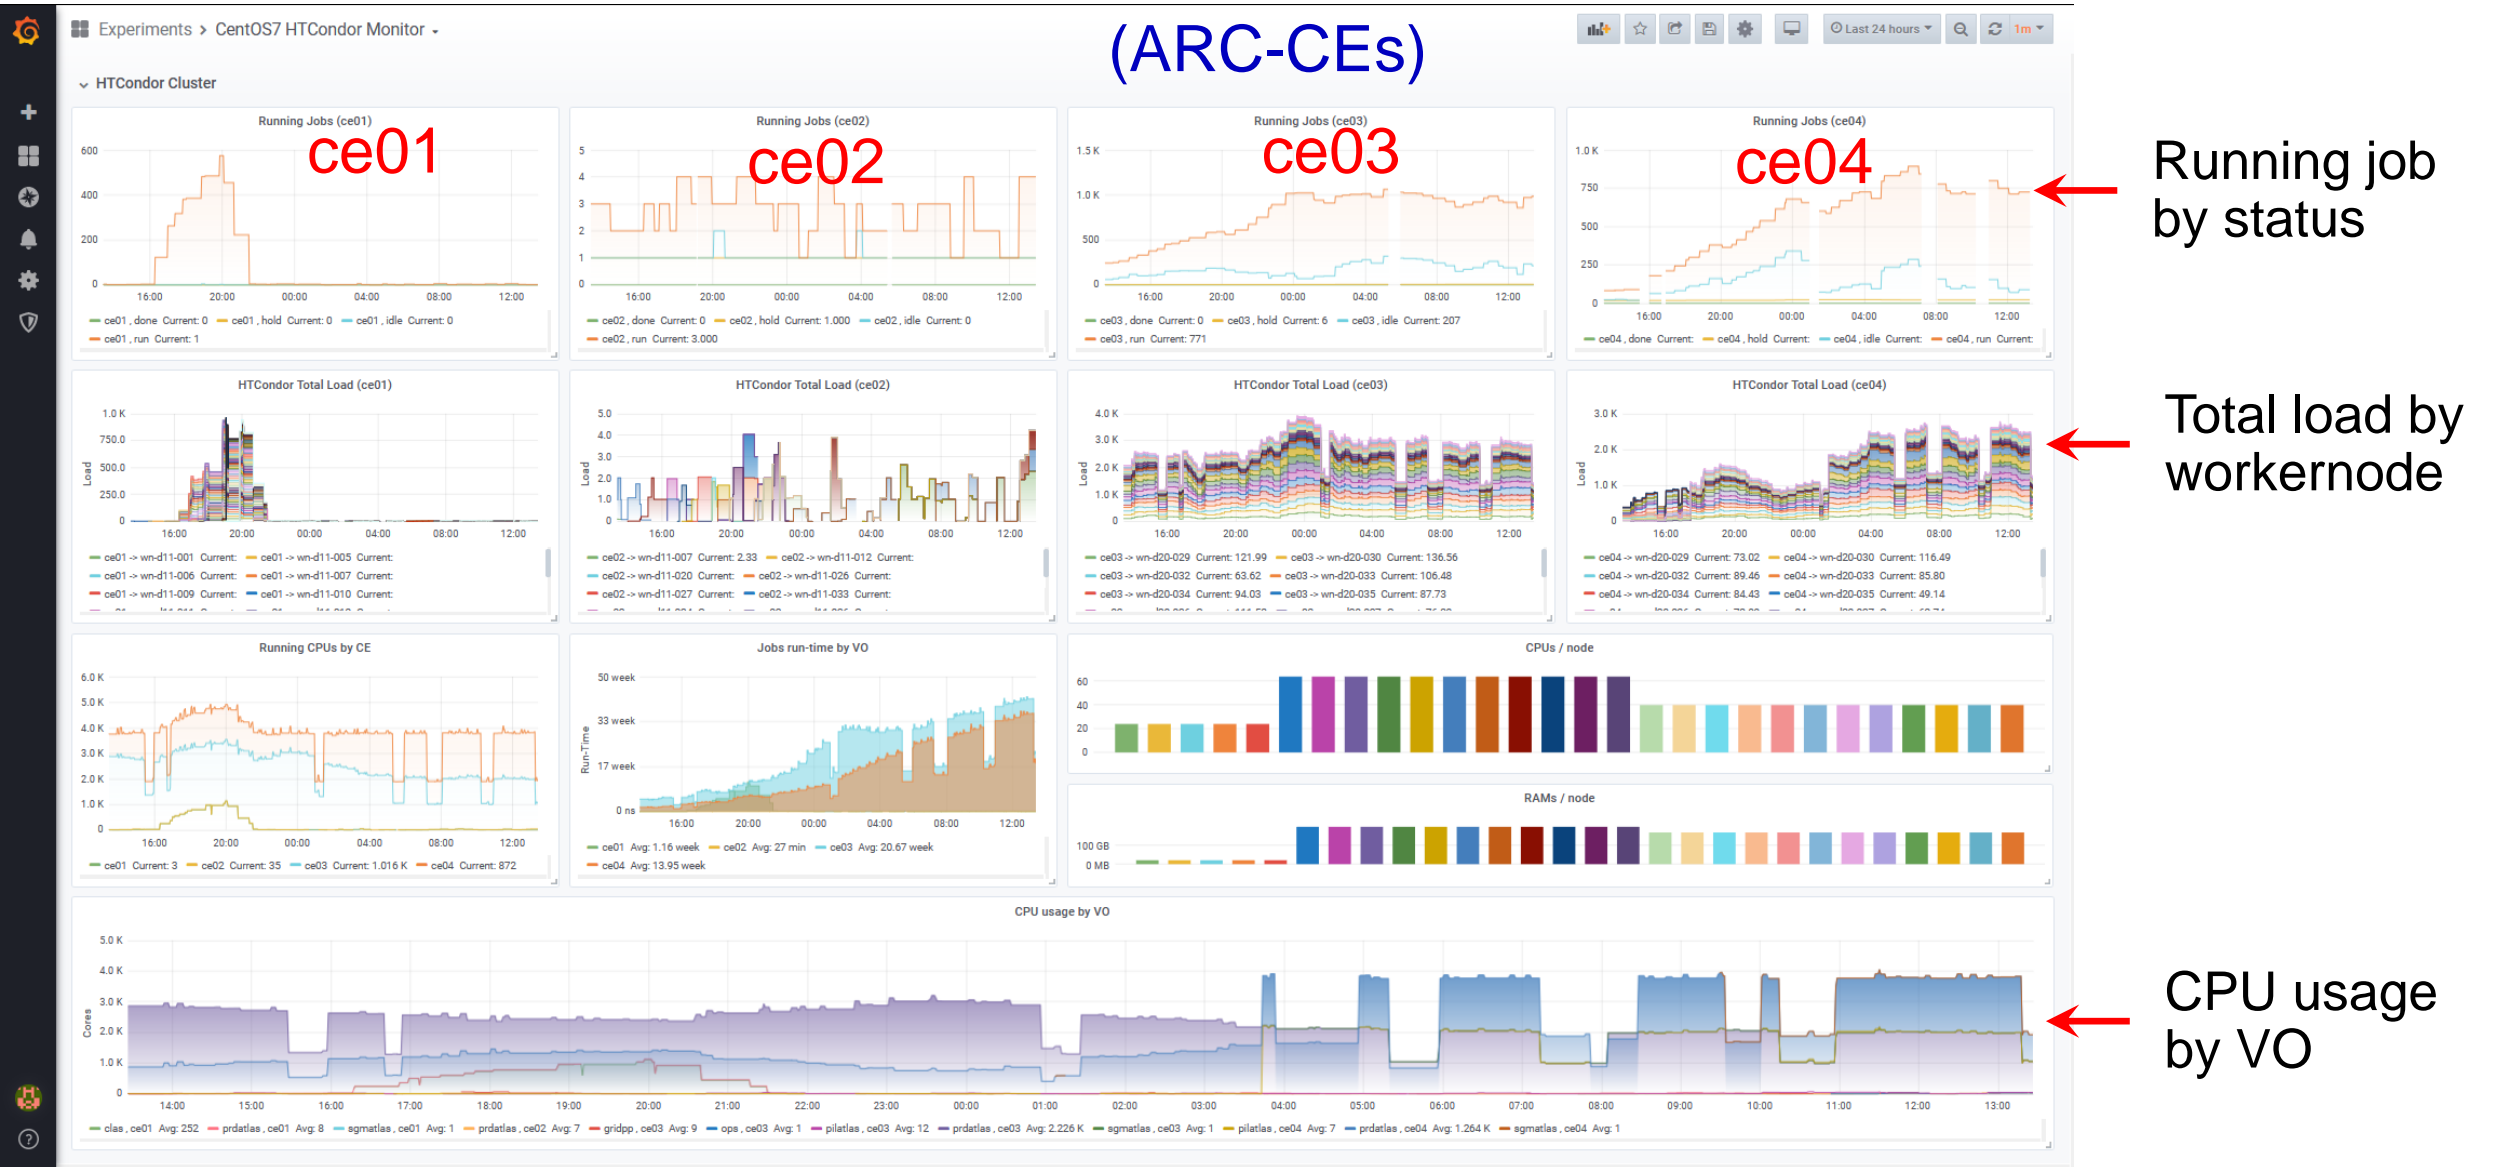


Memory Usage by VO



Over 50 graphs about CPU, Memory, Disks, Networking ...

Grafana Dashboards



Other custom dashboards cover specific services (ARC-CE, Squid, NAT) ...

HTCondor Logging

PromTail is installed on every machine and configured to specifically export service logs to a central server running **Loki** ...

```
1 # PromTail Configuration defines the rules for scraping local logs
2
3 server:
4   http_listen_port: 9080
5   grpc_listen_port: 0
6
7 # Positions
8 positions:
9   filename: /tmp/positions.yaml
10
11 # Loki Server URL
12 clients:
13   - url: http://{{ central_log_server }}:3100/loki/api/v1/push
14
15 scrape_configs:
16   ## Common Logs
17   - job_name: messages # log messages (/var/log/messages)
18     static_configs:
19       - targets:
20         - "{{ inventory_hostname }}:9080"
21         labels:
22           job: messages
23           host: "{{ inventory_hostname }}"
24           path: /var/log/messages
25   # ...
```



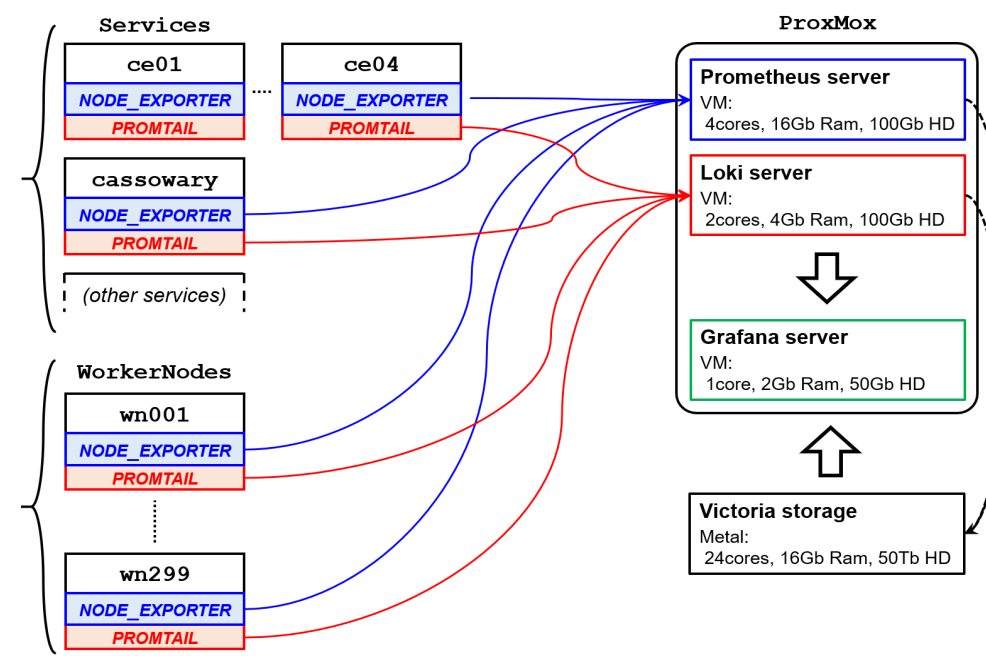
General Logs

```
27 ## Service Logs
28 - job_name: condor
29   static_configs:
30     - targets:
31       - "{{ inventory_hostname }}:9080"
32     labels:
33       job: condor
34       host: "{{ inventory_hostname }}"
35       path: /var/log/condor/*og
```



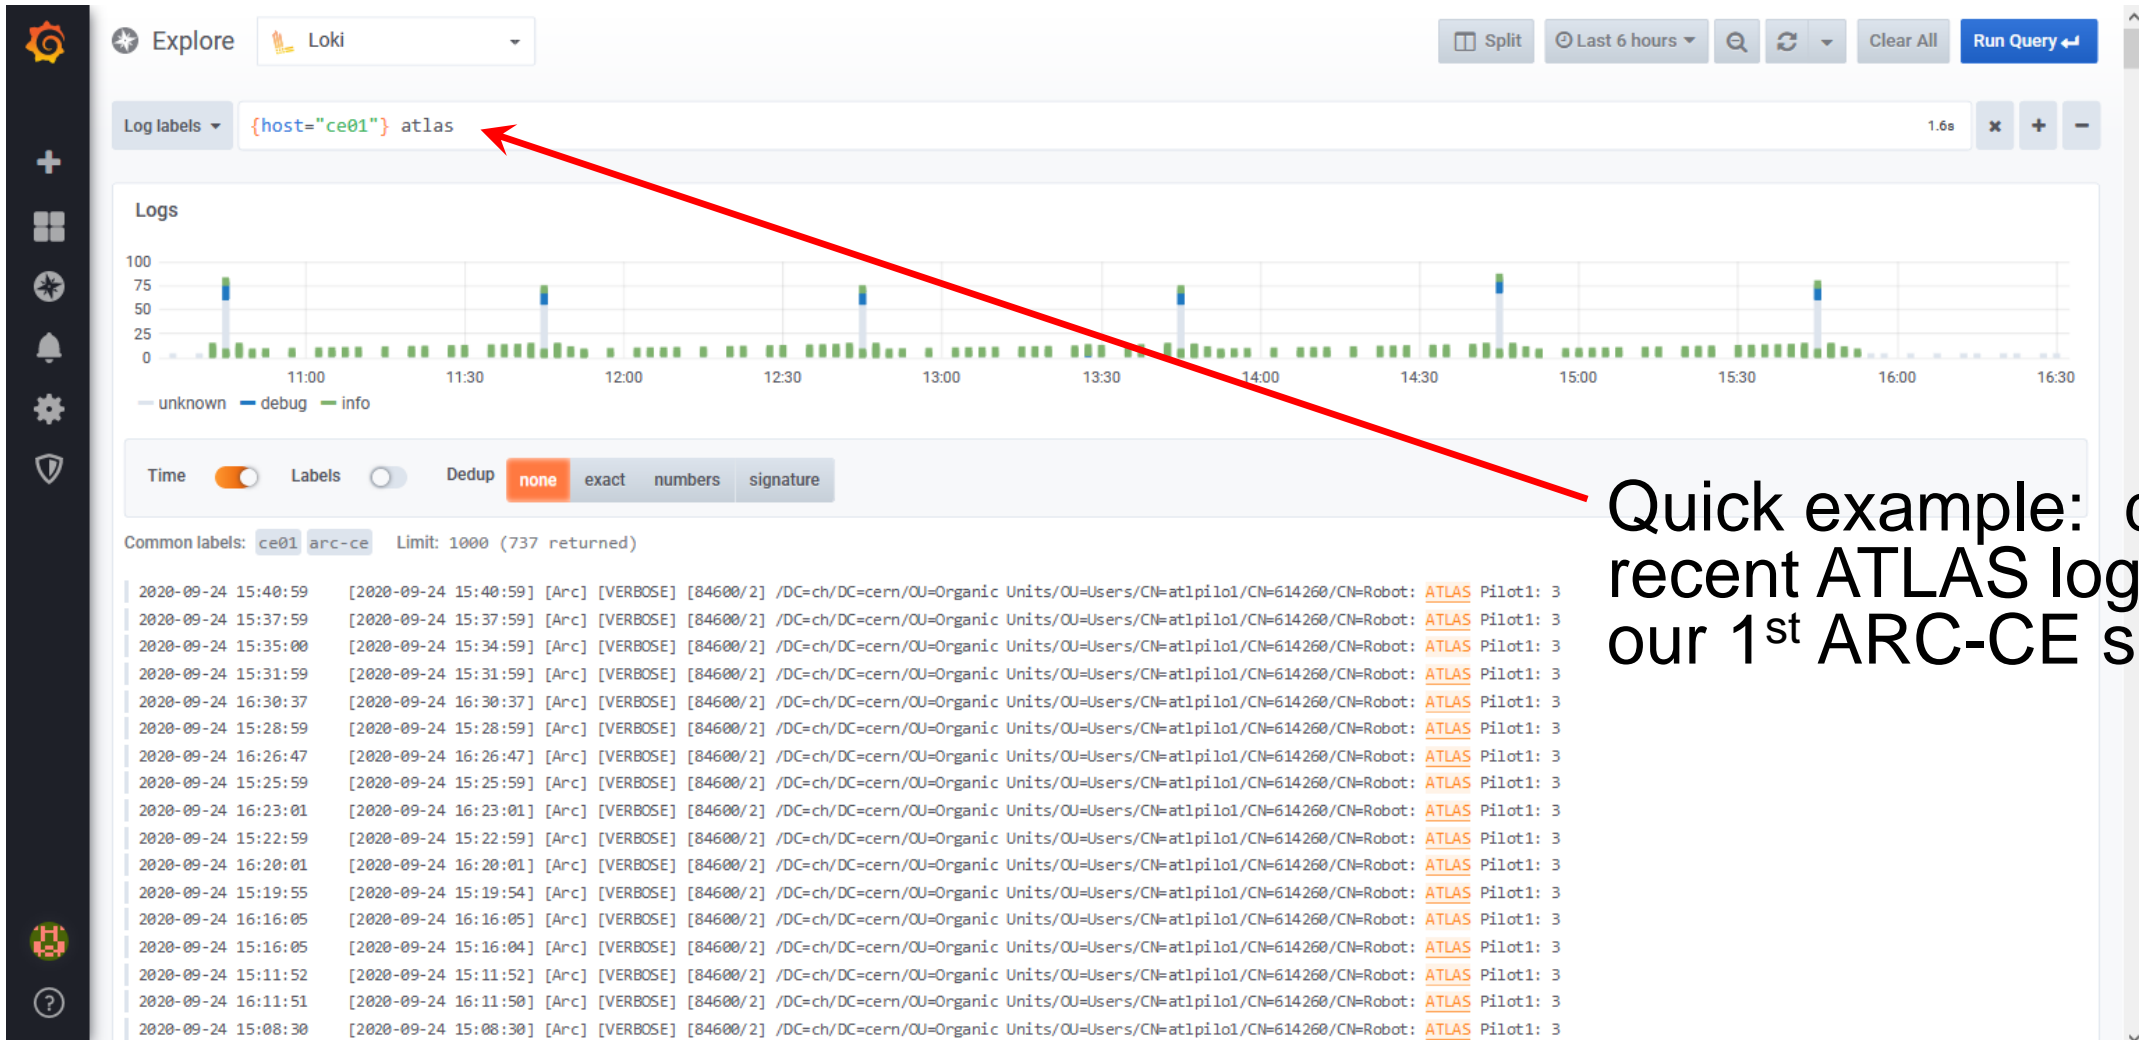
HTCondor logs:

/var/log/condor/



Logs Board


Logs can be queried also from the Grafana web interface (by host, service, content, ...)



Quick example: display recent ATLAS logs from our 1st ARC-CE server

ElasticFlow

We also have an **ElasticSearch** instance to analyze network flows from switches (*sFlow*)



... but this is another story.

Outlook

The monitoring system is relatively new (early 2020) and still under development:

- We already have an alerting system that can send emails/SMS in case of emergency (e.g., rising temperature).
- Next we are trying to identify more complex alerts based on combinations of metrics and define automated recovery actions (reboot / re-provisioning).
- Eventually, we will experiment with Machine Learning based Anomaly Detection and Recovery.
- We also started engaging other GridPP sites to come up with the best practice for monitoring and automated fault recovery.

And, something completely outside the scope of my talk:

- We have tried to install HTCondor-CE based on the Liverpool use-case, but it did not seem to work in our environment (SCHEDD and NEGOTIATOR located differently). Depending on peer pressure, we may give it another try in the future.

END