

SRCNet - CHSRC deployment and demo

Lukas Gehrig (FHNW), Pascal Herzog (FHNW), Carolina Lindqvist (EPFL), Pablo Llopis Sanmillan (EPFL)



Agenda

- What happened since the last SKACH meeting
- Highlights from new features
 - Migrating the monitoring stack from Thanos to Mimir
 - Developing and running STARS on Kubernetes
 - PanDA progress
 - Updates on global services maintained by CHSRC
 - Other improvements for SRCNet and CHSRC
- Future plans



PI planning news



- PI29 planning in Perth, AUSSRC
- Workshop week - very useful but unlikely to appear again
- PI33 planning will be in Switzerland (details mentioned in Rohinis presentation)
- Minor team changes, trying a new way of working with features/JIRA
- PI goals
 - Preparing for v.0.2
 - Integrating DevOps principles into SRCNet (CI/CD, integration tests...)
 - Documentation declutter



Image credits: Pablo Llopis 📷

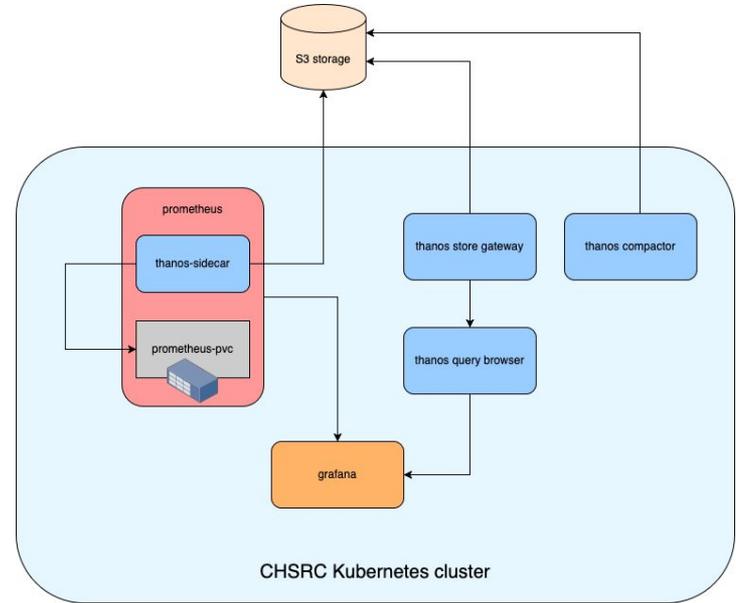


SRC|Net

SKAO Regional Centre Network

Migrating the monitoring stack - the old stack

- Old stack based on Thanos
- Storing data in S3 but not being able to visualise it in Grafana
- Debugging not successful



Migrating the monitoring stack - evaluation

- Choosing Grafana Mimir after evaluation
 - Fits the existing Prometheus-Grafana stack
 - Widely used tool
- Coincided with a object storage (S3) migration at CSCS
- Collaboration with INSRC for the setup



Migrating the monitoring stack - Mimir

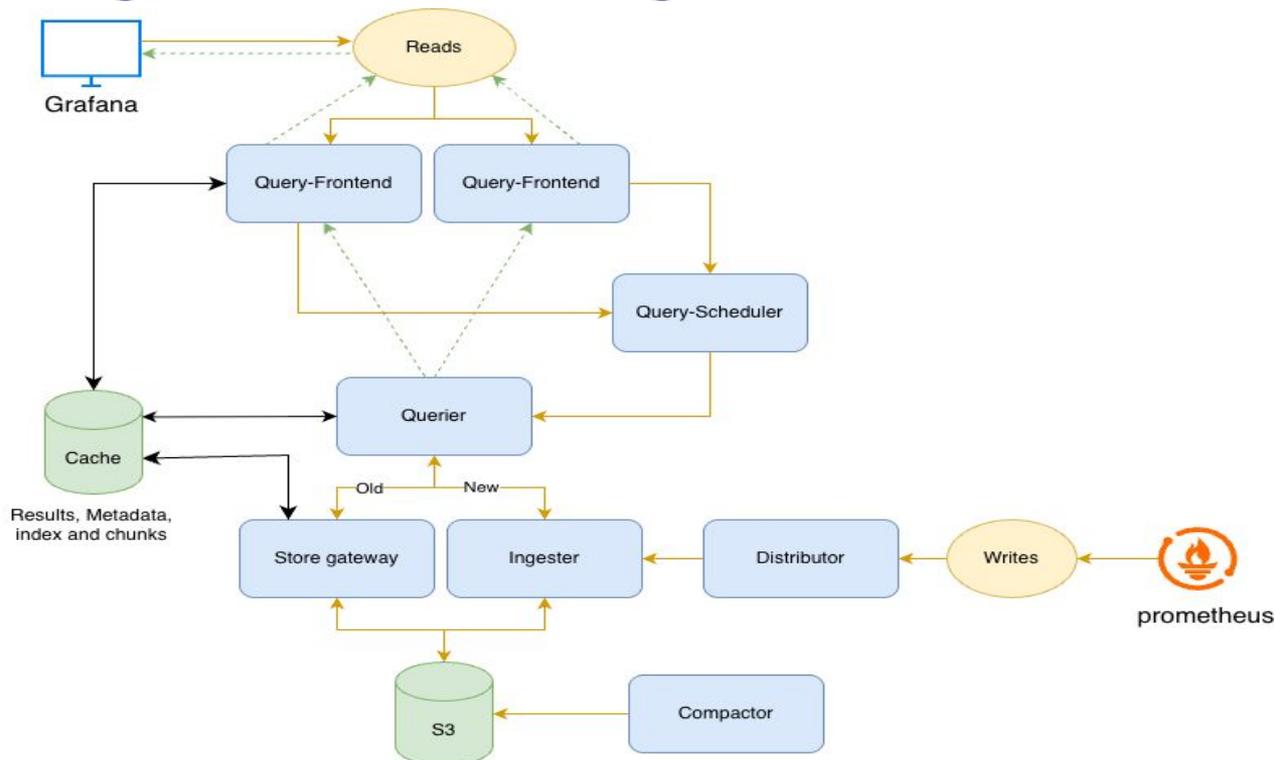


Diagram credit: Hrushikesh Kothem, NEON team

Migrating the monitoring stack - conclusions

- Mimir has a similar architecture as Thanos (based on the same underlying technology)
- There is only a high availability option that can be easily deployed
- This leads to higher resource consumption in the CHSRC clusters
- Overall Mimir integrates well with Grafana



STARS as a representative compute score

We're developing a benchmark suite comprised of workloads **representative of SRCNet.**

Outputs a **STARScore.**

Can be run: locally via Docker, Singularity, or on Kubernetes.

SRCNet compute is highly heterogeneous → STARS enables comparison of compute capacity.

Useful for pledging, accounting, and system testing.



PanDA progress

PanDA is the “umbrella” batch system across SRCNet.
Enables single entry point for large-scale, **data-aware execution** of jobs across different sites, no matter which batch system they’re using.

Initial prototype contains the following sites:



More sites to follow in the following months.



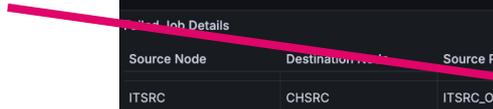
Global Services - System testing

Documentation



Node	Status	Value
ESPSRC	passed	4
	failed	3

Debug info



Source Node	Destination Node	Source RSE	Destination RSE	Response Body	Status	Step
ITSRC	CHSRC	ITSRC_OACT_XRD	CHSRC_XRD_PROD	Data not ready at requested site	200	Check replication job
CHSRC	CNSRC	CHSRC_XRD_PROD	CNSRC-SHAO-T1	unexpected HTML content	200	Perform Soda cutout
ITSRC	SKAOSRC	ITSRC_OACT_XRD	STFC-STO-M	Data not ready at requested site	200	Check replication job
ITSRC	CANSRC	ITSRC_OACT_XRD	CASRC_XRD	unknown	200	Check replication job
ESPSRC	CNSRC	ESPSRC-STO-M	CNSRC-SHAO-T1	Data not ready at requested site	200	Check replication job

New bugs



- NEON-549 System Testing: Different canfar versions may cause the S... PI29 1
- NEON-541 System testing: System tests always selects 0th image at S... PI29 0

... and more
SRC|Net



Global Services - other updates

- Grafana query time optimization/debugging
 - Allowing users to query for metrics over a long time frame (months/years)
- New local cluster monitoring dashboard (WIP)
 - Getting an overview of the cluster status without having to log in
 - Sharing cluster info with members who cannot access the cluster



Other improvements in CHSRC

- Thanks to CSCS for unifying the dev/prod environment VMs
 - This highlighted a need for changing a hacky shared volume setup that we maintained for a while.
- Moving from sharing a volume across namespaces to using static Ceph volumes
- Version upgrades across services
 - Release management process work ongoing in SOG



Future plans

- Panda adoption and prototyping across SRCNet
- Regularly running STARS and recording metrics
- Running DSCs at CHSRC
- Improving the stability of the cluster further
- Nice-to-have:
 - Partially automated version upgrades
 - Admins reviewing an automated MR instead of puzzling together helm chart changes



Thanks & Questions?

