

SKACH Spring Meeting

Updates on Swiss SRCNet Infrastructure prototyping activities

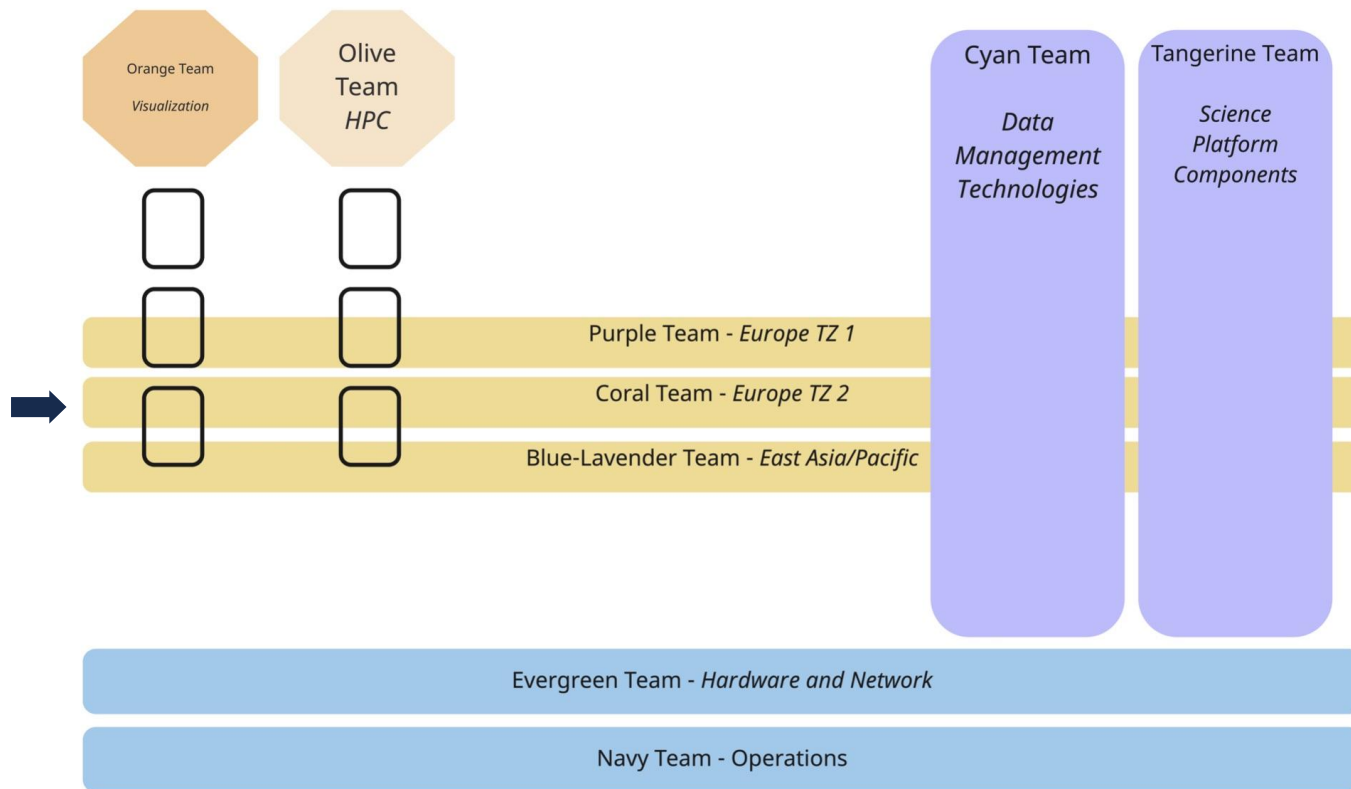
Pablo Llopis - EPFL / Coral team
1st June 2023




Agenda

- Introduction
- Infrastructure building blocks
- **Data management** solutions
 - HEP: **Rucio**
 - CADC: **Storage Inventory**
- Benchmarking

Introduction: Coral Team



Introduction: SRCNet Roadmap



Introduction: SRCNet Roadmap

Focus is on **storage**:
Mostly **data** and
metadata management
solutions.

4.1.2 SRCNet v0.1 Development

Version	SRCNet v0.1
Development Duration	March 2023-Nov 2024 = 20 months
Software Thread Versions	SRCNet Presentation Tier v0.1 Data Management v0.1 Metadata Management v0.1 AAI preliminary module
Description	<p>Preliminary version of the portal, enabling <u>discovery queries on the data</u>. That would also contain a basic local interactive analysis.</p> <p>A preliminary version of the <u>data management system</u>, to allow the creation of digital object replicas on the SRCNet node selected by a user.</p> <p>A preliminary version of the <u>metadata system to discover data</u> in the data lake</p> <p>Minimum viable product of the authentication/authorisation system, in connection with the <u>data management system</u> to control access policies</p>

Agenda

- Introduction
- **Infrastructure building blocks**
- Data management solutions
 - HEP: Rucio
 - CADC: Storage Inventory
- Benchmarking

Kubernetes + OpenStack as base platform



monitoring

Search



Workloads > Pods

Workloads N

- Cron Jobs
- Daemon Sets
- Deployments
- Jobs
- Pods**
- Replica Sets
- Replication Controllers
- Stateful Sets


Service

- Ingresses N
- Ingress Classes
- Services N


Config and Storage

- Config Maps N
- Persistent Volume Claims N
- Secrets N

CPU Usage



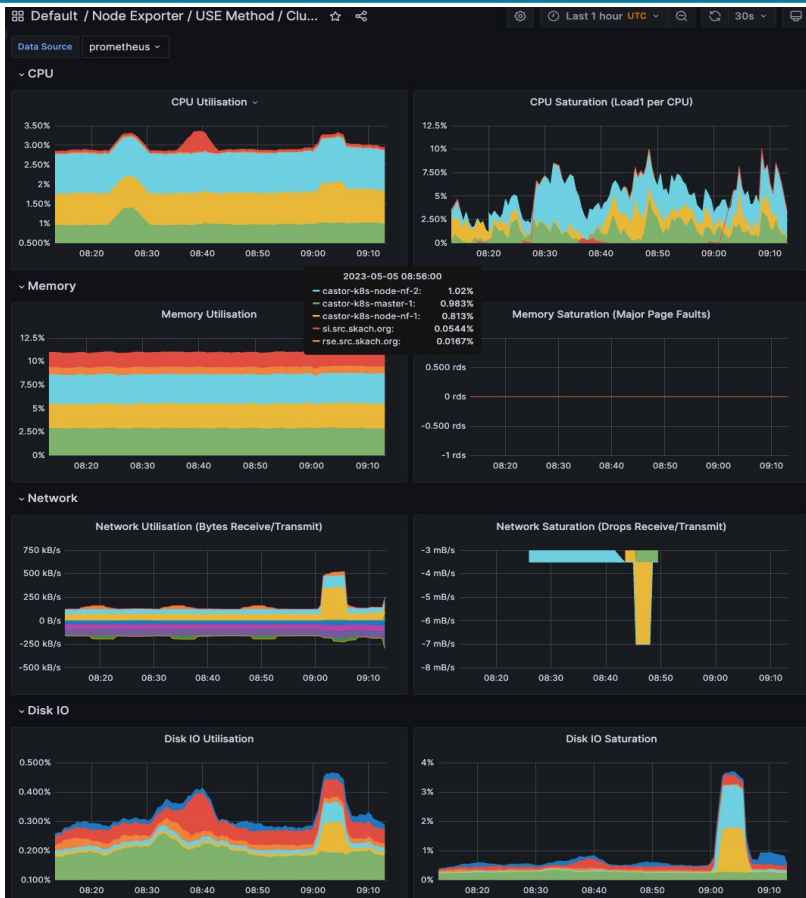
Memory Usage



Pods

Name	Images	Labels	Node	Status	Restarts	CPU Usage (cores)	Memory Usage (bytes)	Created ↑
● prometheus-k8s-1	quay.io/prometheus/prometheus:v2.43.0 quay.io/prometheus-operator/prometheus-config-reloader:v0.64.0	app.kubernetes.io/component: prometheus app.kubernetes.io/instance: k8s app.kubernetes.io/managed-by: prometheus-operator	castor-k8s-node-nf-2	Running	0	75.00m	467.06Mi	5.days.ago

Monitoring infra for prototype services



skach.org: Domain usage for infrastructure



Using scratch.org domain name for infrastructure.

TLS certificates thanks to Let's Encrypt.

skach.org: Domain usage for IVOA

Euro-VO Registry






MY RESOURCES

 Information

Results (3)



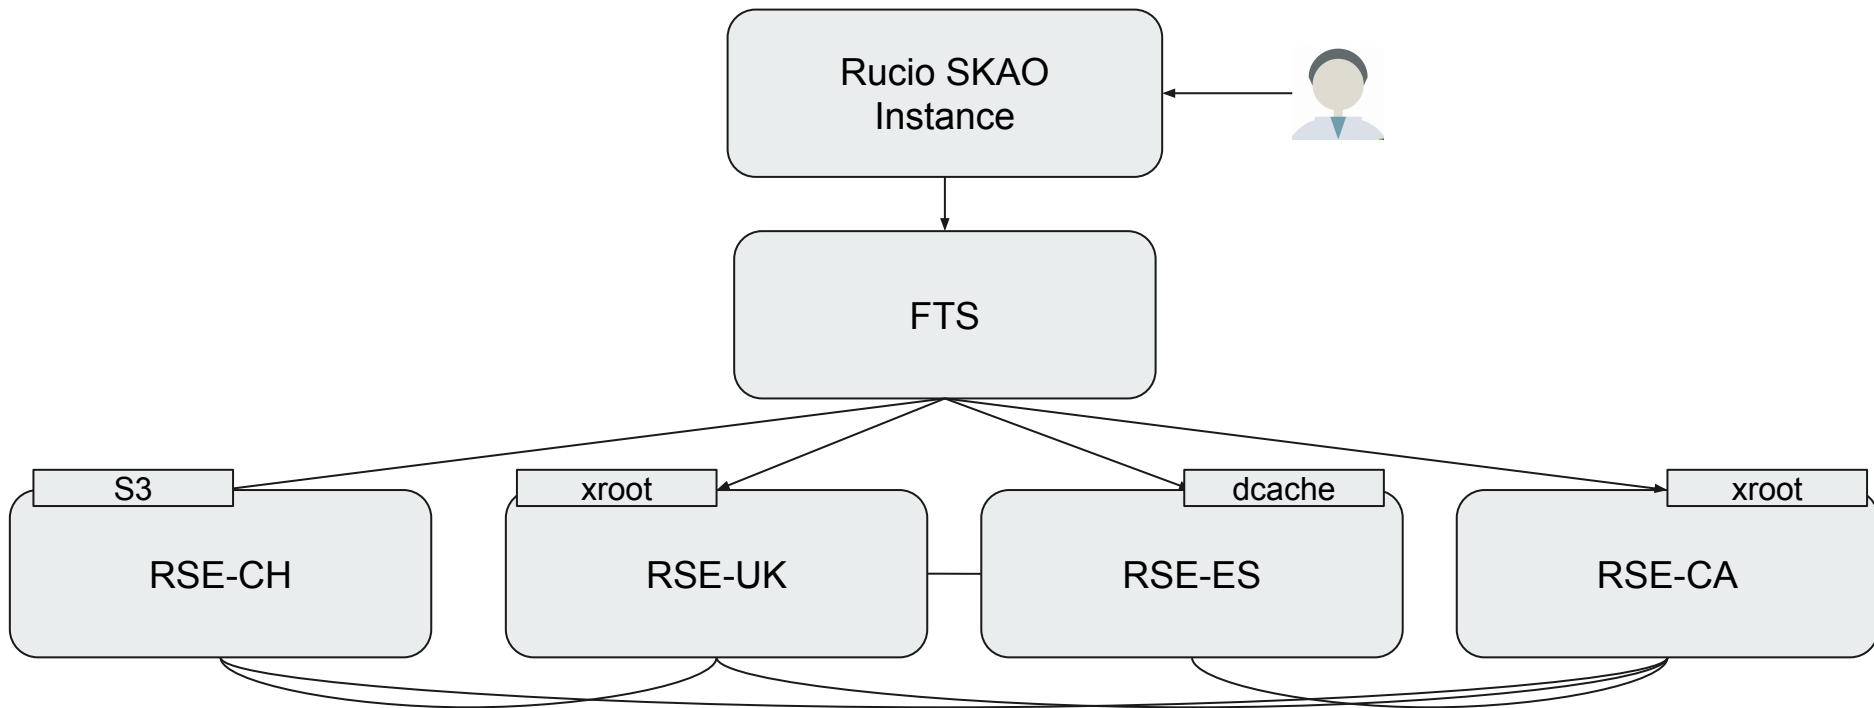
	Short Name	IVOID	Status	Title	Resource Type	Capabilities	Compliance	#total versions	Created
		ivo://src.skach.org	active	CH SRC IVOA Authority	vg:Authority			1	2023-04-19 09:58:24
		ivo://src.skach.org/luskan	active	Metadata service for SKA data lake by CADC; Implementation at CH SRC	vr:Service			1	2023-04-19 10:04:33
		ivo://src.skach.org/minoc	active	File service for SKA data lake by CADC; Implementation at CH SRC	vr:Service			1	2023-04-19 10:08:46

Needed for OpenCADC Storage Inventory deployment

Agenda

- Introduction
- Infrastructure building blocks
- **Data management solutions**
 - HEP: Rucio
 - CADC: Storage Inventory
- Benchmarking

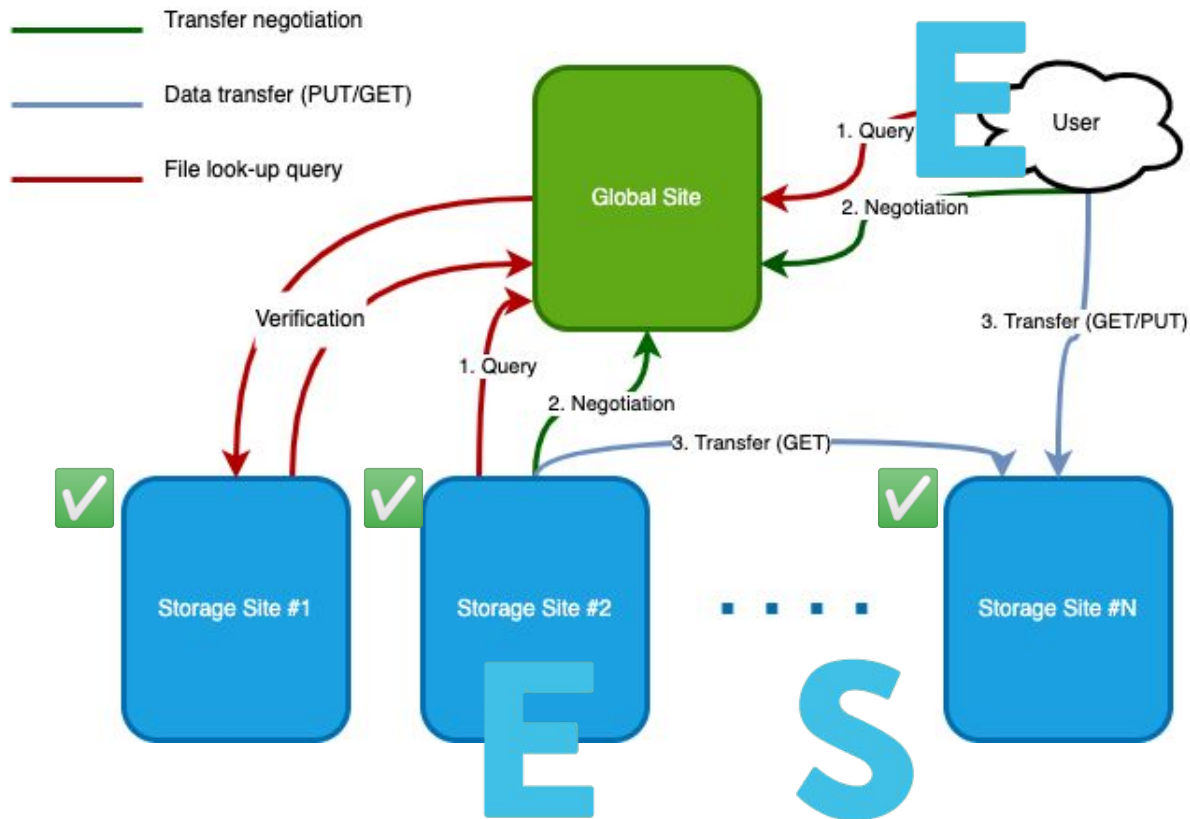
Data Lake prototype: Rucio



Data Lake prototype: Rucio

- Identified remaining challenges for Rucio
 - **Token expiration** for large data transfers or many parallel transfers
 - **Extra tooling** needed for SKA data products (e.g. SDC3)
 - Data transfer **performance issues** for heavily nested datasets - Rucio's DB becomes the bottleneck.
(credit: Cyan team)

OpenCADC Storage Inventory



Data Lake prototype: Storage Inventory


Data distribution plan for mini SRCNet

	SI Namespace	Dynamic	Count	TiB	SPSRC	SWSRC	SWISRC	???
1	cad:CGPS/		3338	0.061		✓	✓	
2	cad:VGPS/		118	0.027	✓		✓	
3	cad:VLASS/	✓	167592	0.106	✓	✓		
4	cad:WALLABY/		7733	0.002		✓	✓	
5	casda:RACS/		799	0.570	✓	✓	✓	
6	nrao:VLASS/	✓	162488	8.760	2/3	2/3	2/3	
7	ska:???							

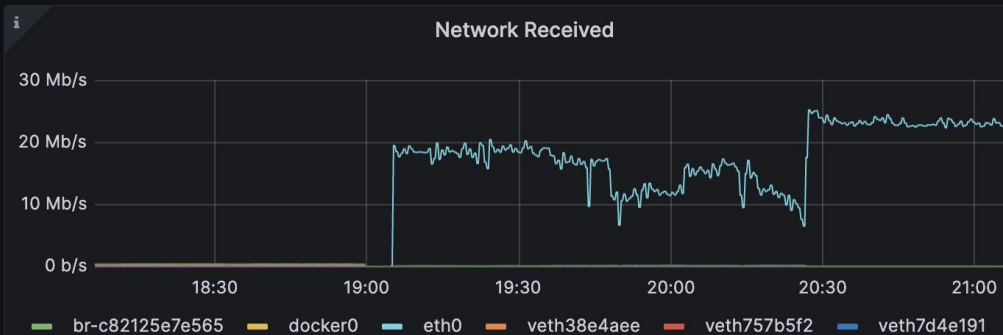
Data Lake prototype: Storage Inventory

Data movement and replication started just a few days ago!

~ Disk




~ Network



Data Lake prototype: Storage Inventory

Starting to **understand** how to **fine tune** this solution to scale **data transfer performance**.

Big thank you to the **CADC** folks!



Storage Inventory: tasks for near future

- ✓ ● Telescope data seeded and replicated in a mini SRCNet
- ✓ ● Anonymous data access
- 👷 ● Data transfer performance testing
- Authorised data access, group management system
- Namespace support
 - Data from related national activities such as from other telescopes could co-exist
- IVOA interfaces for data + metadata querying and discovery

Agenda

- Introduction
- Infrastructure building blocks
- Data management solutions
 - HEP: Rucio
 - CADC: Storage Inventory
- **Benchmarking**

Towards a global SRCNet perf score

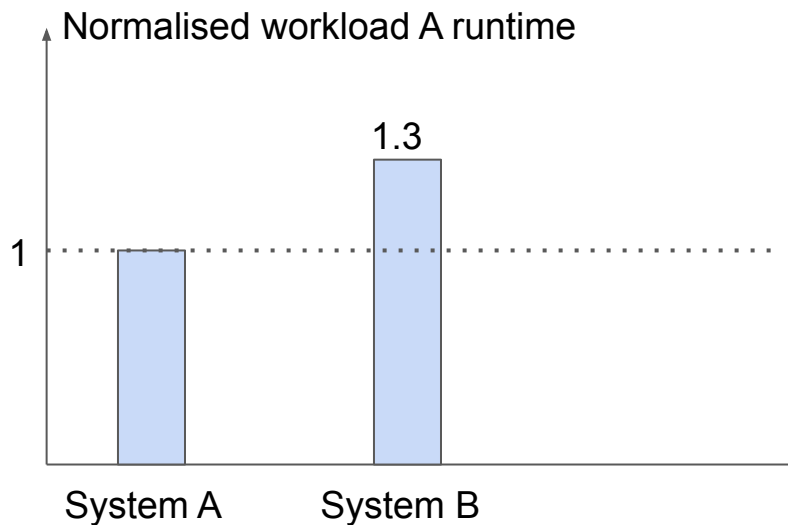
➔ *Can we define a **performance metric** that's **representative** of SRCNet workloads?*

If we had a **representative benchmark suite** of workloads that will be run on **SRCNet** nodes that produced a **performance score**, we could leverage it for:

- **Pledges**
- **Accounting**
- **Procurement** of new systems

Towards a global SRCNet perf score

Select a baseline system as a reference for wallclock runtime.



Reference:
System A



System B

Towards an SKA benchmark suite

➔ How to produce a single score out of many different benchmarks?

➔ How to identify workloads that are computationally similar?

Which SKA benchmarks are similar?

Are there any industry standard benchmarks that we could leverage?

SKACH Spring Meeting

Updates on Swiss SRCNet Infrastructure prototyping activities

Pablo Llopis - EPFL / Coral team
1st June 2023

