




CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Computing Platforms and Infrastructure Program Update

Victor Holanda, CSCS

Darren Reed, UZH

January 22th, 2024



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Reporting time

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition
- Support on the Rucio+S3 prototyping with SKAO

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition
- Support on the Rucio+S3 prototyping with SKAO
- Testing the SKAO AAI prototype at CSCS

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition
- Support on the Rucio+S3 prototyping with SKAO
- Testing the SKAO AAI prototype at CSCS
- Participate as a facility partner in the SDC3 data challenge

What have we work on in 2023?

The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition
- Support on the Rucio+S3 prototyping with SKAO
- Testing the SKAO AAI prototype at CSCS
- Participate as a facility partner in the SDC3 data challenge
 - Congratulate the SKACH team - 10th Place
 - Allocated 15'000 node hours to them (94% utilization)

What have we work on in 2023?


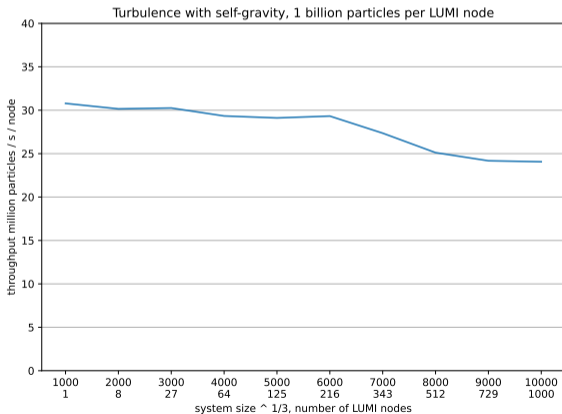
The Computing and Platforms contributions

- Deployed dedicated Kubernetes infrastructure to SKACH
 - Currently extending to include another development platform
- Following the SRCNet node requirements discussions
- Working on the Swiss SRC requirement definition
- Support on the Rucio+S3 prototyping with SKAO
- Testing the SKAO AAI prototype at CSCS
- Participate as a facility partner in the SDC3 data challenge
 - Congratulate the SKACH team - 10th Place
 - Allocated 15'000 node hours to them (94% utilization)
- Supporting the different synergies in collaboration with CTA

What have we work on?

The Computing and Platforms contributions

- Supported the submission of Lucio's Project to LUMI-G using SPH-EXA



What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What's new?

What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What's new?

- Work on getting vCluster running on Pawsey, Australia - MWA collaboration

What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What's new?

- Work on getting vCluster running on Pawsey, Australia - MWA collaboration
- Develop guidelines for Secure Software Development Life Cycle with SKAO

What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What's new?

- Work on getting vCluster running on Pawsey, Australia - MWA collaboration
- Develop guidelines for Secure Software Development Life Cycle with SKAO
- Work on the FirecREST integrated with JupyterHUB

What is currently planned to 2024?

An exciting future ahead of us

- We will continue working on the 2023 commitments

What's new?

- Work on getting vCluster running on Pawsey, Australia - MWA collaboration
- Develop guidelines for Secure Software Development Life Cycle with SKAO
- Work on the FirecREST integrated with JupyterHUB
- Extend the Kubernetes offering with a development cluster



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Moving to ALPS

The (current) plan

When and how are we moving?

- The migration should start in April



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move
 - New (non-GPU) projects are going to eiger
 - Move hybrid to User Lab



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move
 - New (non-GPU) projects are going to eiger
 - Move hybrid to User Lab
- Hardware will be different



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move
 - New (non-GPU) projects are going to eiger
 - Move hybrid to User Lab
- Hardware will be different
 - ARM based (NVIDIA Grace CPU)
 - New NVIDIA GPUs (NVIDIA Hopper GPU)



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move
 - New (non-GPU) projects are going to eiger
 - Move hybrid to User Lab
- Hardware will be different
 - ARM based (NVIDIA Grace CPU)
 - New NVIDIA GPUs (NVIDIA Hopper GPU)
- Software installation will have a different workflow



The (current) plan

When and how are we moving?

- The migration should start in April
 - Expect a busy machine - everyone will be migrating
 - Suggest you to write ReFrame tests for your applications and workflows
- Single stage move
 - New (non-GPU) projects are going to eiger
 - Move hybrid to User Lab
- Hardware will be different
 - ARM based (NVIDIA Grace CPU)
 - New NVIDIA GPUs (NVIDIA Hopper GPU)
- Software installation will have a different workflow
 - Based on Spack and Stackinator
 - Uses squashfs images





CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

SKACH Computing Platform

Gornergrat and beyond

- Same components as the User Lab at CSCS

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - We don't foresee any differences

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - We don't foresee any differences
- What does it change from the code developer point of view?

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - We don't foresee any differences
- What does it change from the code developer point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - There will be different architectures

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - We don't foresee any differences
- What does it change from the code developer point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - There will be different architectures
- What does it change from the system administrator point of view?

Gornergrat and beyond

- Same components as the User Lab at CSCS
 - Look and feel to be similar to the largest portion of our system
 - Failover opportunities
 - Overall project costs reductions
- We are currently identifying the gaps
- What does it change from the scientist point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - We don't foresee any differences
- What does it change from the code developer point of view?
 - It will depend heavily on the actual implementation of SRC nodes
 - There will be different architectures
- What does it change from the system administrator point of view?
 - Client services will have to be deployed using Nomad

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters
- We have identifying some gaps

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters
- We have identifying some gaps
- Aiming to make the tests portable - to be shared between SRCNet nodes

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters
- We have identifying some gaps
- Aiming to make the tests portable - to be shared between SRCNet nodes
 - Focus on system testing and on user environment
 - Application testing will be next

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters
- We have identifying some gaps
- Aiming to make the tests portable - to be shared between SRCNet nodes
 - Focus on system testing and on user environment
 - Application testing will be next
- The areas of the tests encompass:

Internal Hackathons

Gornergrat vCluster consistency - We are almost there

- Focus on writing tests to speed up the deployment of vClusters
- We have identifying some gaps
- Aiming to make the tests portable - to be shared between SRCNet nodes
 - Focus on system testing and on user environment
 - Application testing will be next
- The areas of the tests encompass:
 - OS checks
 - Individual components (e.g. SLURM, Filesystem, Sarus, FirecREST, etc)
 - Component integration
 - Workflow integration (e.g. K8s + vCluster)

What are the open questions for 2024?

Let's find the answers together

- MWA collaboration

What are the open questions for 2024?

Let's find the answers together

- MWA collaboration
 - What are the plans for integrating Swiss projects into MWA?

What are the open questions for 2024?

Let's find the answers together

- MWA collaboration
 - What are the plans for integrating Swiss projects into MWA?
 - How do manage to get Karaboo integrated and running with MWA data?

What are the open questions for 2024?

Let's find the answers together

- MWA collaboration
 - What are the plans for integrating Swiss projects into MWA?
 - How do manage to get Karaboo integrated and running with MWA data?
 - How do we integrate Bluebild into the pipeline?

What are the open questions for 2024?

Let's find the answers together

- MWA collaboration


- What are the plans for integrating Swiss projects into MWA?
- How do manage to get Karaboo integrated and running with MWA data?
- How do we integrate Bluebild into the pipeline?
- What are the plans for leveraging the MWA data for SPH-EXA simulations?



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you!
Questions?





CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Accounting

Resource Consumption per Project

2023 Q1

Project	Group Leader	Node type	Quota [nh]	Used [nh]	Used [%]
sk08	lmayer	HYBRID	100'000	100'099	100.1 %
sk04	mstutz	HYBRID	5'000	299	6.0 %
sk02	dkorber	HYBRID	5'000	262	5.2 %
sk09	jkneib	HYBRID	17'500	204	1.2 %
sk012	yrevaz	MULTICORE	5'000	151	3.0 %
sk10	etolley	HYBRID	5'000	89	1.8 %
sk05	lgehrig	HYBRID	5'000	45	0.9 %
sk014	mibianco	HYBRID	5'000	18	0.4 %
sk07	pllopiss	HYBRID	2'500	12	0.5 %
sk07	pllopiss	MULTICORE	2'500	11	0.4 %
sk015	phirling	HYBRID	5'000	2	0.0 %
sk012	yrevaz	HYBRID	5'000	0	0.0 %
sk01	hvictor	HYBRID	0	0	0.0 %
sk011	kshreyam	HYBRID	5'000	0	0.0 %
sk013	dpotter	MULTICORE	2'500	0	0.0 %
sk013	dpotter	HYBRID	2'500	0	0.0 %
sk016	fcabot	HYBRID	5'000	0	0.0 %
sk06	fschramk	MULTICORE	5'000	0	0.0 %
Total			182'500	101'192	55.4 %

Resource Consumption per Project

2023 Q2

Project	Group Leader	Node type	Quota [nh]	Used [nh]	Used [%]
sk08	lmayer	HYBRID	175'000	100'018	57.2 %
sk09	jpgkneib	HYBRID	8'750	3'346	38.2 %
sk015	phirling	HYBRID	5'000	529	10.6 %
sk05	lgehrig	HYBRID	10'000	216	2.2 %
sk10	etolley	HYBRID	5'000	139	2.8 %
sk012	yrevaz	MULTICORE	5'000	100	2.0 %
sk014	mibianco	HYBRID	5'000	73	1.5 %
sk04	mstutz	HYBRID	10'000	51	0.5 %
sk09	jpgkneib	MULTICORE	8'750	13	0.2 %
sk02	dkorber	HYBRID	5'000	1	0.0 %
sk016	fcabot	HYBRID	5'000	1	0.0 %
sk07	pllopiss	MULTICORE	2'500	0	0.0 %
sk01	hvictor	HYBRID	0	0	0.0 %
sk011	kshreyam	HYBRID	5'000	0	0.0 %
sk013	dpotter	MULTICORE	2'500	0	0.0 %
sk013	dpotter	HYBRID	2'500	0	0.0 %
sk06	fschramk	MULTICORE	5'000	0	0.0 %
sk07	pllopiss	HYBRID	2'500	0	0.0 %
Total			262'500	104'487	39.8 %

Resource Consumption per Project

2023 Q3

Project	Group Leader	Node type	Quota [nh]	Used [nh]	Used [%]
sk08	lmayer	HYBRID	175'000	176'195	100.7 %
sk014	mibianco	HYBRID	10'000	7'970	79.7 %
sk05	lgehrig	HYBRID	10'000	7'378	73.8 %
sk015	phirling	HYBRID	5'000	149	3.0 %
sk04	mstutz	HYBRID	10'000	100	1.0 %
sk017	framunno	HYBRID	5'000	0	0.0 %
sk01	hvictor	HYBRID	0	0	0.0 %
sk012	yrevaz	MULTICORE	5'000	0	0.0 %
sk013	dpotter	MULTICORE	2'500	0	0.0 %
sk013	dpotter	HYBRID	2'500	0	0.0 %
sk016	fcabot	HYBRID	5'000	0	0.0 %
sk018	pdenzel	HYBRID	5'000	0	0.0 %
sk07	pllopiss	MULTICORE	2'500	0	0.0 %
sk07	pllopiss	HYBRID	2'500	0	0.0 %
sk10	etolley	HYBRID	5'000	0	0.0 %
Total			245'000	191'791	78.3 %

Resource Consumption per Project

2023 Q4

Project	Group Leader	Node type	Quota [nh]	Used [nh]	Used [%]
sk08	lmayer	HYBRID	175'000	176'150	100.7 %
sk014	mibianco	HYBRID	15'000	5'993	40.0 %
sk19	lmachado	HYBRID	5'000	2'531	50.6 %
sk018	pdenzel	HYBRID	5'000	1'806	36.1 %
sk05	lgehrig	HYBRID	15'000	1'450	9.7 %
sk015	phirling	HYBRID	5'000	128	2.6 %
sk04	mstutz	HYBRID	10'000	11	0.1 %
sk016	fcabot	HYBRID	5'000	0	0.0 %
sk01	hvictor	HYBRID	0	0	0.0 %
sk012	yrevaz	MULTICORE	5'000	0	0.0 %
sk013	dpotter	MULTICORE	2'500	0	0.0 %
sk013	dpotter	HYBRID	2'500	0	0.0 %
sk017	framunno	HYBRID	5'000	0	0.0 %
sk020	mbredber	HYBRID	5'000	0	0.0 %
sk021	msargent	HYBRID	5'000	0	0.0 %
sk07	pllopiss	MULTICORE	2'500	0	0.0 %
sk07	pllopiss	HYBRID	2'500	0	0.0 %
Total			265'000	188'068	71.0 %