

Distributed Quantum Computing

Erik Källman, RISE

QAS 2024



MONIQUE
CALISTI
MARTEL
INNOVATE

WHAT WE'LL DISCUSS:

CONVERGENCE
OF SEVERAL TECH
AREAS INTO THE
**COMPUTING
CONTINUUM**

↓

COMMON
RESEARCH +
INNOVATION
AGENDA



PEARSE
O'DONOHUE
EUROPEAN
COMMISSION

WE RELY ON
DIGITAL →

THE COVID
CRISIS
HIGHLIGHTED
THIS

WE NEED TO
GRASP THE
COMPUTING
CONTINUUM



MAX
LEMKE
EUROPEAN
COMMISSION

WHAT WE MEAN
WITH "CONTINUUM"

CLOUD/HPC

EDGE

IoT

ALL INTEGRATED

Continuum

"A coherent whole characterized as a collection, sequence, or progression of values or elements varying by minute degrees"

- Merriam webster

PRESS RELEASE | 5 December 2023 | Brussels | 8 min read

Commission approves up to €1.2 billion of State aid by seven Member States for an Important Project of Common European Interest in cloud and edge computing technologies

Challenges Compute Infrastructure

From a developer perspective ...



■ User experience

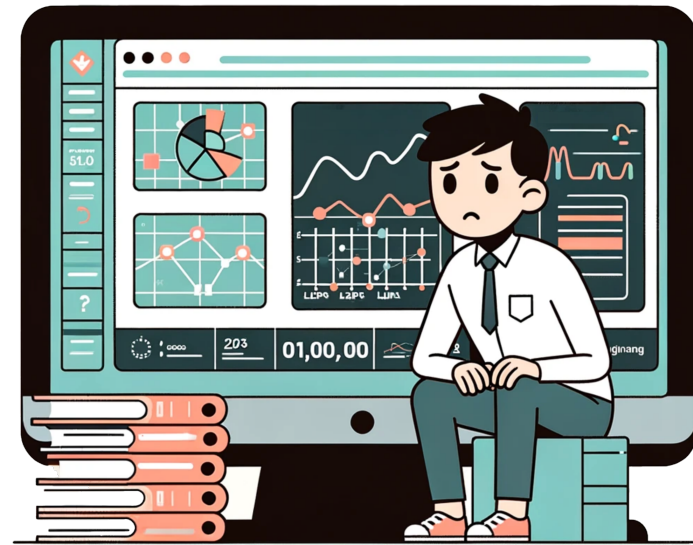
- Complex login process: SSH to a login node
- Setting up tunnels
- Mastering Slurm jobs
- When will my job run? Will someone kill my job?

■ Data management

- Determining data storage locations
- Manual data transfers can be time-consuming and error-prone

■ Integration issues

- Connecting HPC systems with cloud to streamline workflows?
- No APIs? Lack of automation tools (GitOps/CI/CD)
- Multi-factor authentication
- Sometimes no Internet access on compute nodes

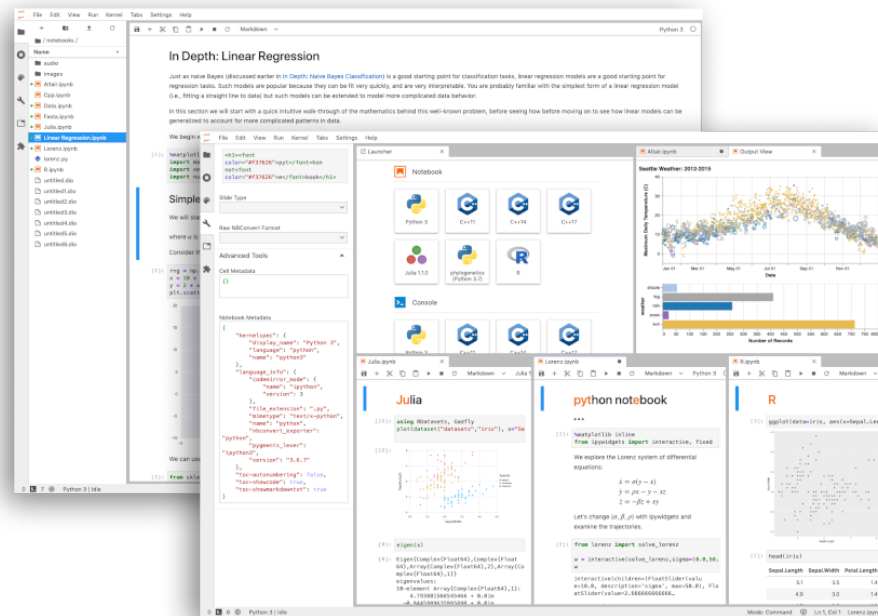


Generated by ChatGPT

High-Performance Computing

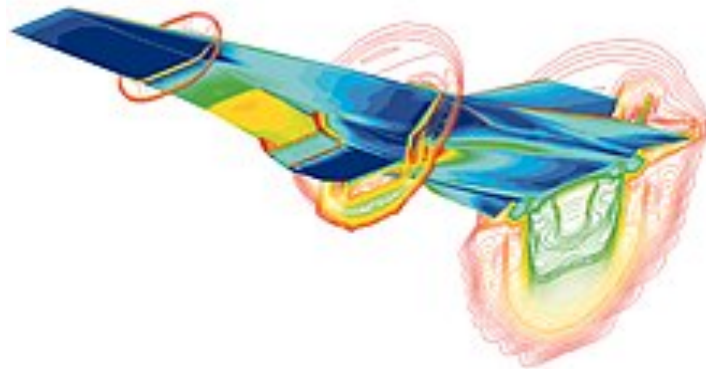
■ Scientific / Research workflows

- **Manual interaction:** Required to set up simulations or experiments
- **Valuable outcomes over efficiency:** Quickly obtain accurate and valuable research results
- **Exploratory:** Research workflows can be less predictable and require more hands-on adjustments
- **Batch processing:** Requiring manual scripting and queue management



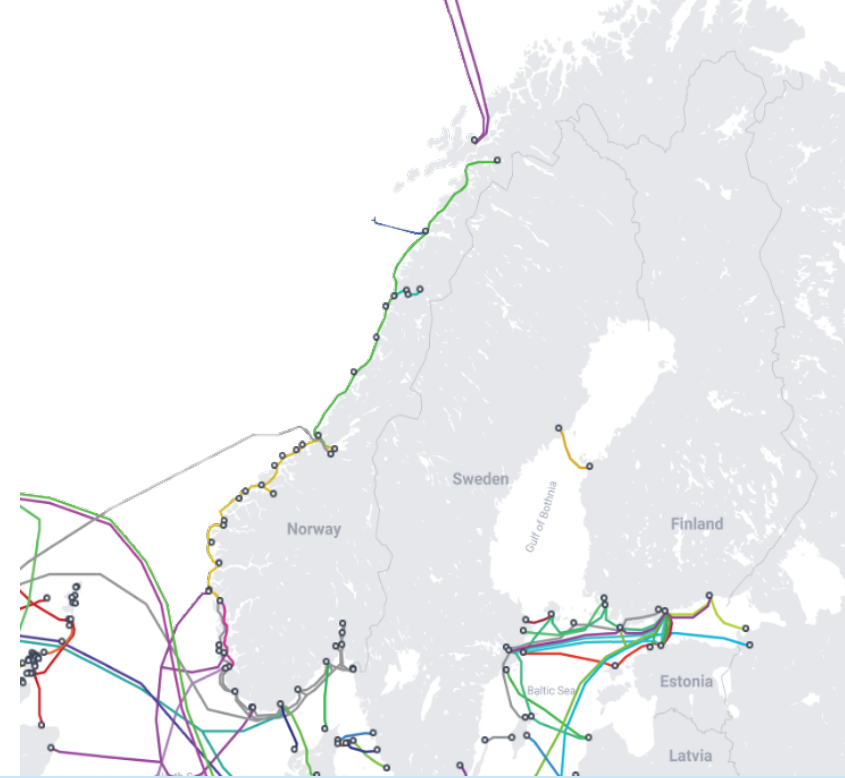
■ Why not use cloud platforms?

- Cloud platforms can be very complex and cumbersome to use for researchers
- Cloud platforms like Kubernetes are not designed for HPC workloads (*not optimized for performance*)



Problems with Cloud Computing

- Dependency on network access
- Vendor lock-in
- Compliance and Regulations
- Security and privacy concerns
- Digital sovereignty



Opinion **Russian politics**

[+ Add to myFT](#)

Putin knows that undersea cables are the west's Achilles heel

Moscow has invested in subsurface naval capabilities that hold the world's internet infrastructure at risk

EDWARD STRINGER

[+ Add to myFT](#)

Ideal for **scientific workflows, large-scale simulations, complex engineering computations**, and tasks requiring extensive computational power and high data throughput

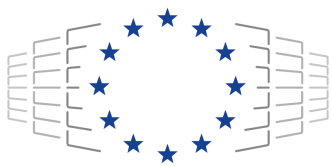
Ideal for **development, testing, and small-scale experimentation**. Suited for prototyping, debugging, and tasks that require immediate, hands-on access to computational resources

Ideal for **data storage, big data processing, machine learning, and production environments**. Optimized for scalable, distributed web services, and cost-effective resource management across global infrastructures

HPC

Local

Cloud



Local

- Link, share, and use local resources (laptops, gaming machines) into a *personal grid*

Compute Continuum

- Simplify cloud accessibility for HPC users
- Seamless migration to cloud after using EuroHPC

HPC

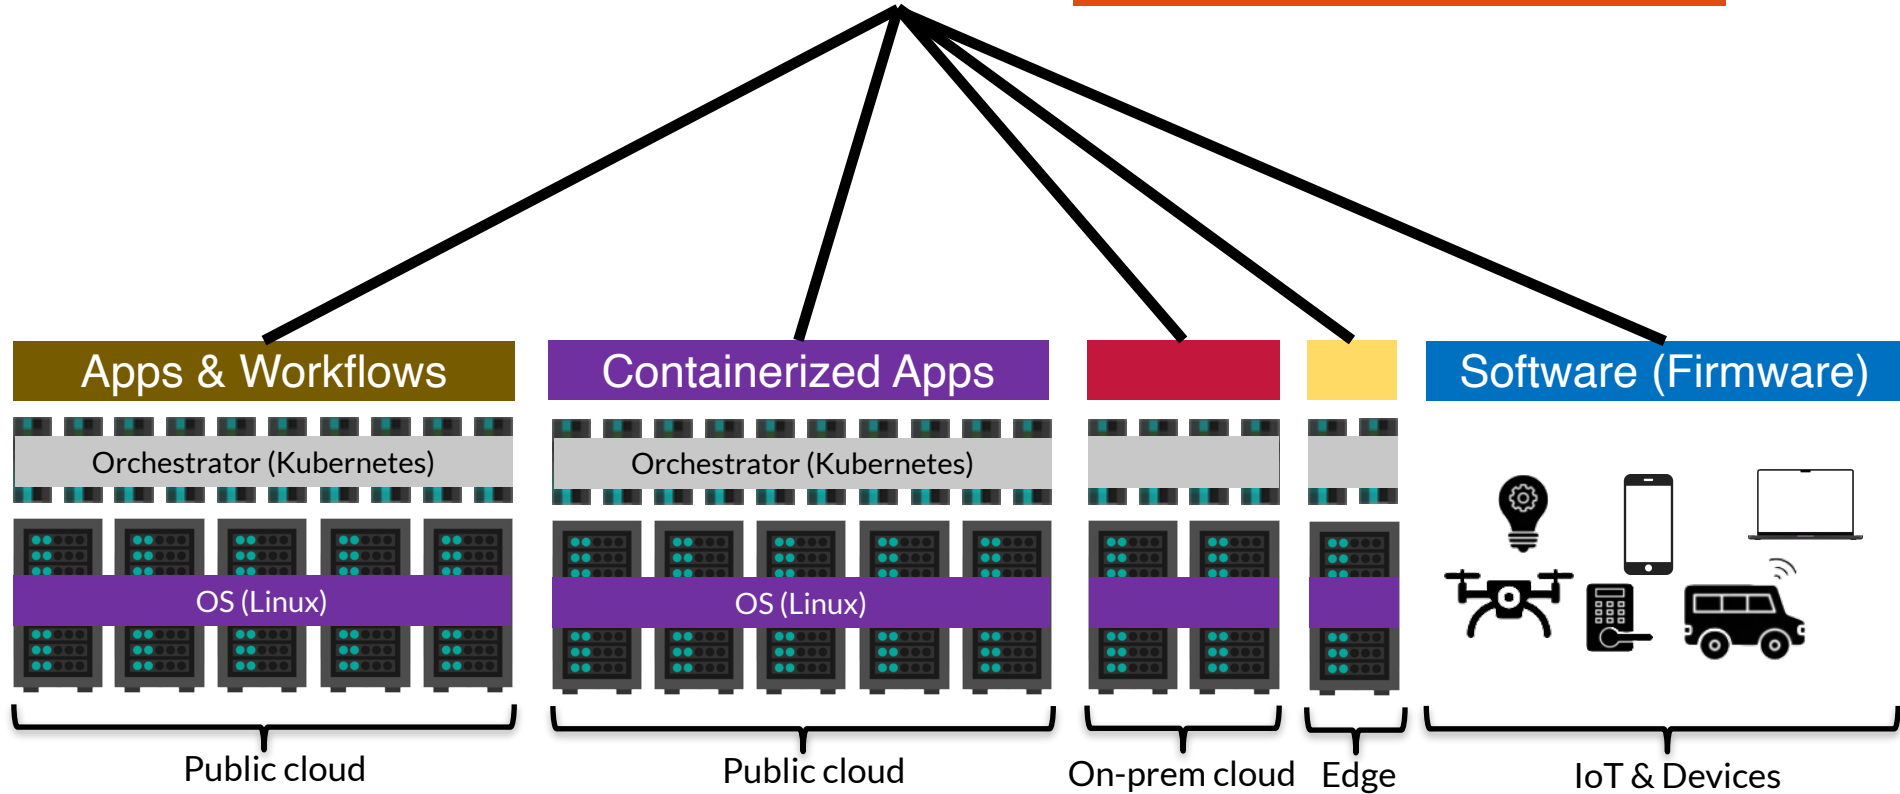


- Provide access to HPC with a modern API
- Access to “free” GPUs

Cloud



Complex Software





Less Complex Software

Unified API

Compute Continuum

Apps & Workflows

Orchestrator (Kubernetes)

OS (Linux)

Public cloud

Containerized Apps

Orchestrator (Kubernetes)

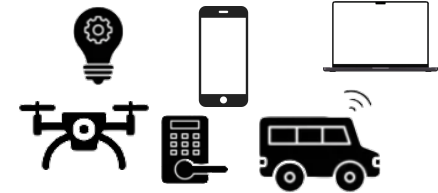
OS (Linux)

Public cloud

On-prem cloud

Edge

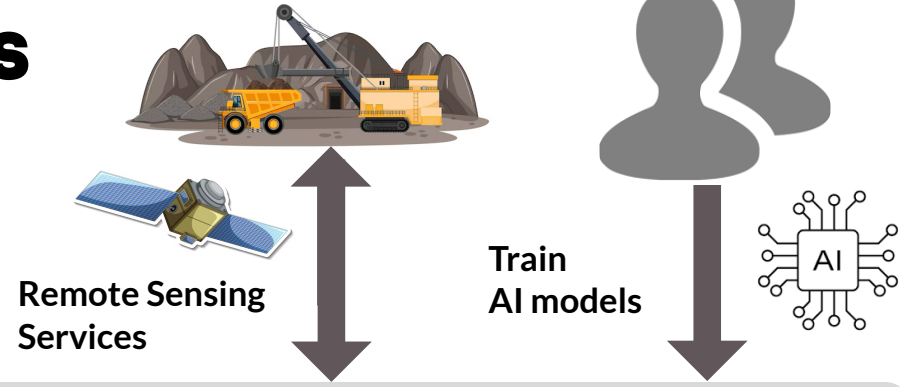
Software (Firmware)



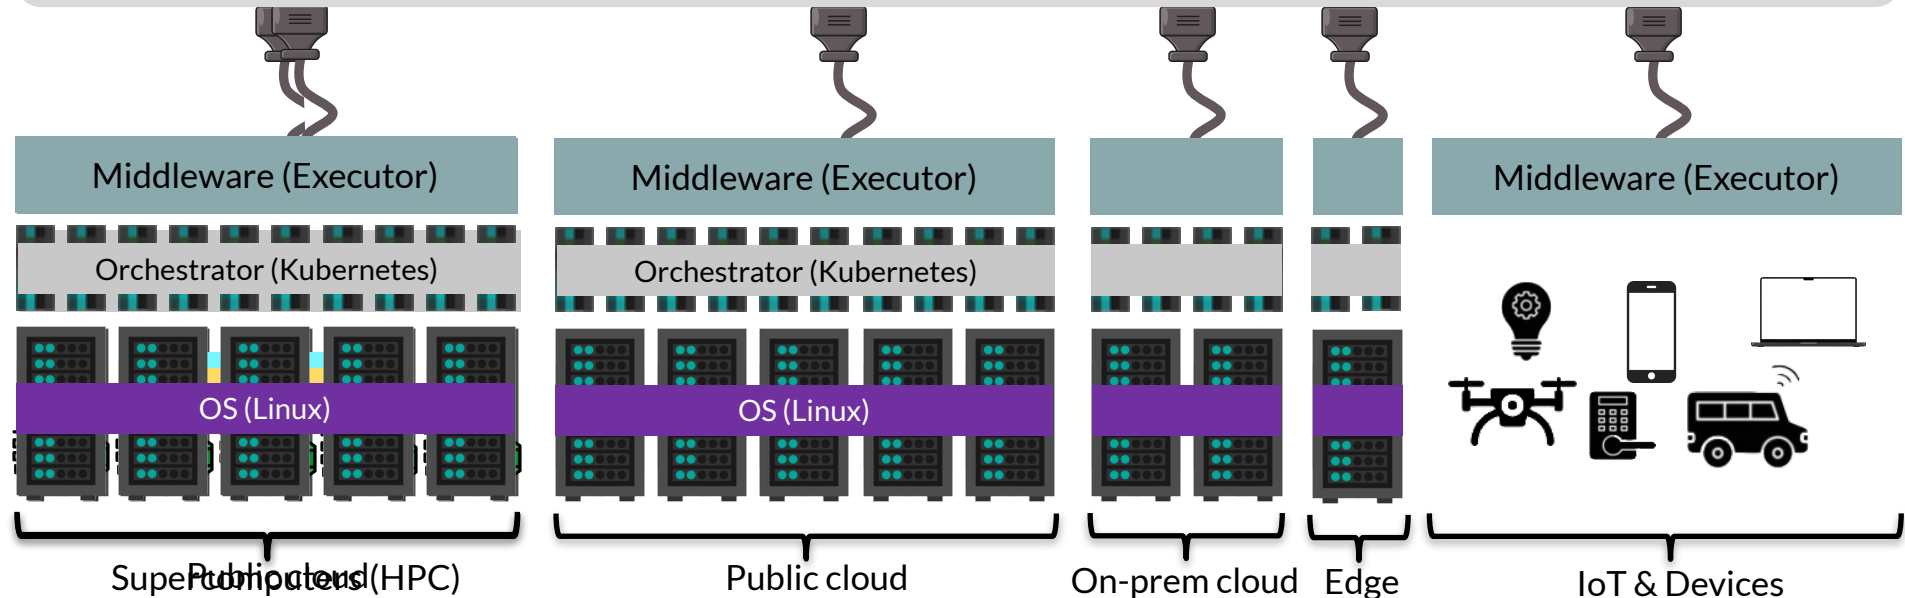
IoT & Devices

Meta-Operating Systems

A foundation for Compute Continuum



Meta-Operating System





ColonyOS

Unleashing Computational Power Everywhere!



What is ColonyOS?

In a rapidly digitalizing world, seamless interoperability and robust large-scale computing aren't just luxuries—they're *essential*. Yet, as we shift towards decentralized and diverse computing landscapes, developing cross-platform applications becomes a daunting task. Imagine a world where AI workloads can easily be developed and run seamlessly across any platform, including Cloud, Edge, and HPC.

Welcome to ColonyOS!

ColonyOS is an [open-source](#) research project developed by [RISE AB](#), and is used by [ENCCS](#) to foster greater High-Performance Computing (HPC) adoption. It is also used by [RockSigma AB](#) to implement a compute engine designed for seismic processing in underground mines. RockSigma AB has contributed to the development of ColonyOS.

[Read more](#)

[Getting started](#)

[Contact us](#)

Use Cases



Distributed Compute Engines

Implement distributed compute engines that optimize data processing across diverse platforms. Perform intensive computations on one platform and then effortlessly merge the



Streamlined HPC

ColonyOS offers modern APIs and cloud integration, expanding supercomputers' reach and accessibility. HPC Executors enables easy, platform-agnostic deployment of workloads. boostine global



Virtual Supercomputing

Harness and combine computational power of multiple disparate computing systems, whether HPC, cloud-based infrastructures, or other computing resources, to

<https://colonyos.io>



ColonyOS

3 followers Sweden

[Unfollow](#)

[Overview](#) [Repositories](#) 16 [Projects](#) [Packages](#) [Teams](#) [People](#) 3 [Settings](#)

Pinned

[Customize pins](#)

View as: **Public**

You are viewing the README and pinned repositories as a public user.

You can [create a README file](#) visible to anyone.

[Get started with tasks](#) that most successful organizations complete.

Discussions

Set up discussions to engage with your community!

[Turn on discussions](#)

People



[Invite someone](#)

Top languages

[Go](#) [JavaScript](#) [Python](#) [Julia](#) [TypeScript](#)

Most used topics

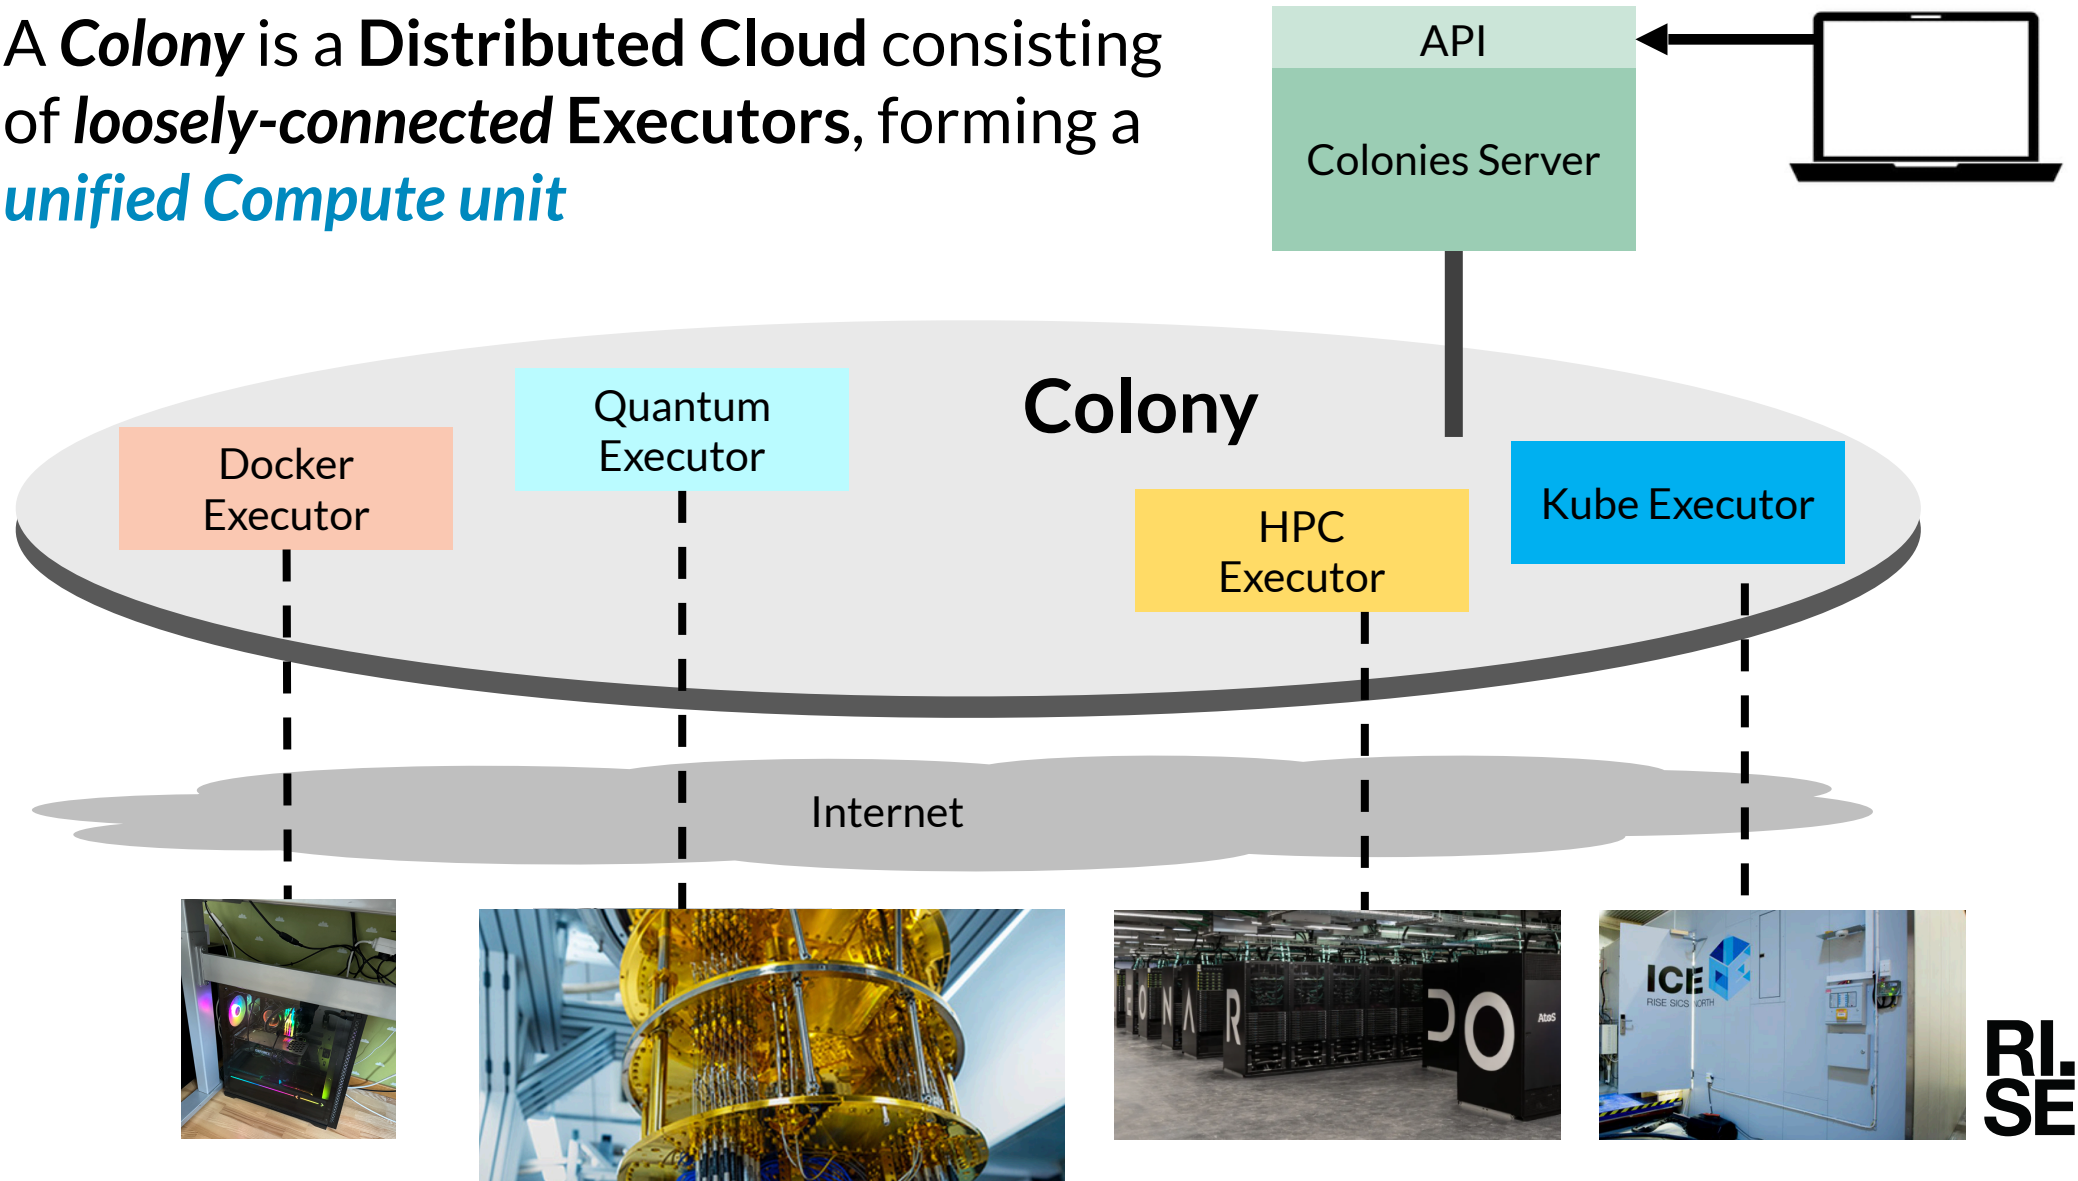
[Manage](#)

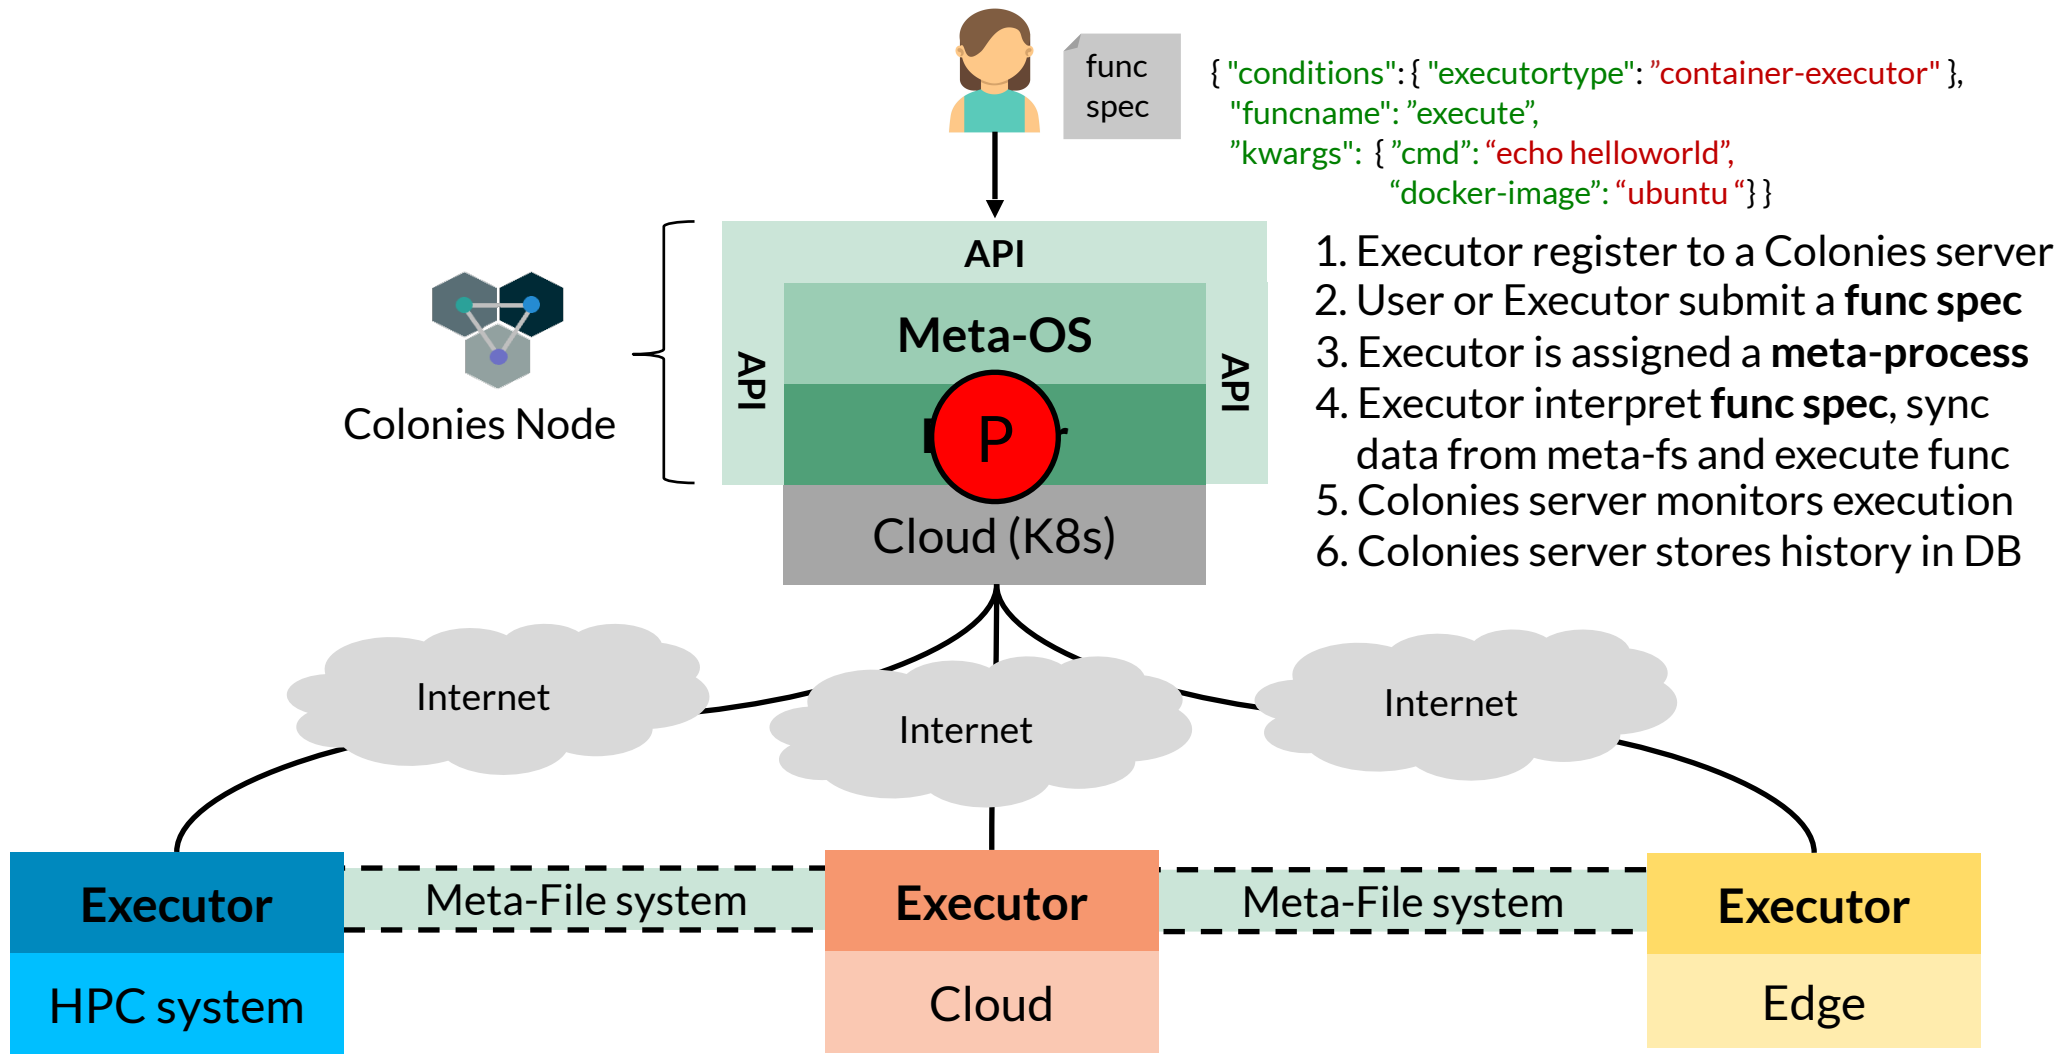
[distributed-systems](#) [edge-computing](#) [gridcomputing](#) [kubernetes](#) [config](#)



<https://github.com/colonyos>

A **Colony** is a **Distributed Cloud** consisting of *loosely-connected* **Executors**, forming a *unified Compute unit*

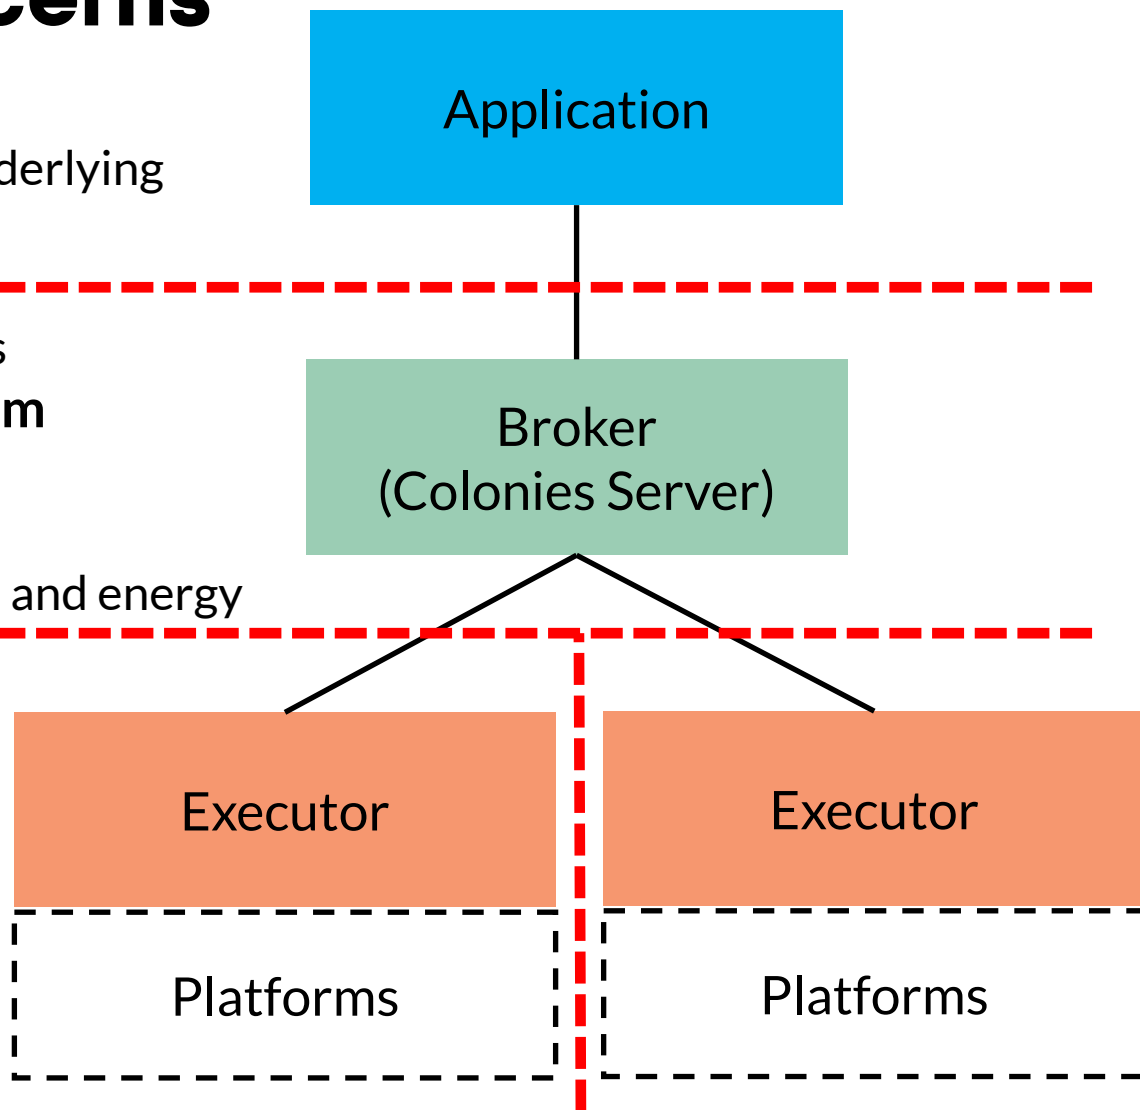




1. Executor register to a Colonies server
2. User or Executor submit a **func spec**
3. Executor is assigned a **meta-process**
4. Executor interpret **func spec**, sync data from meta-fs and execute func
5. Colonies server monitors execution
6. Colonies server stores history in DB

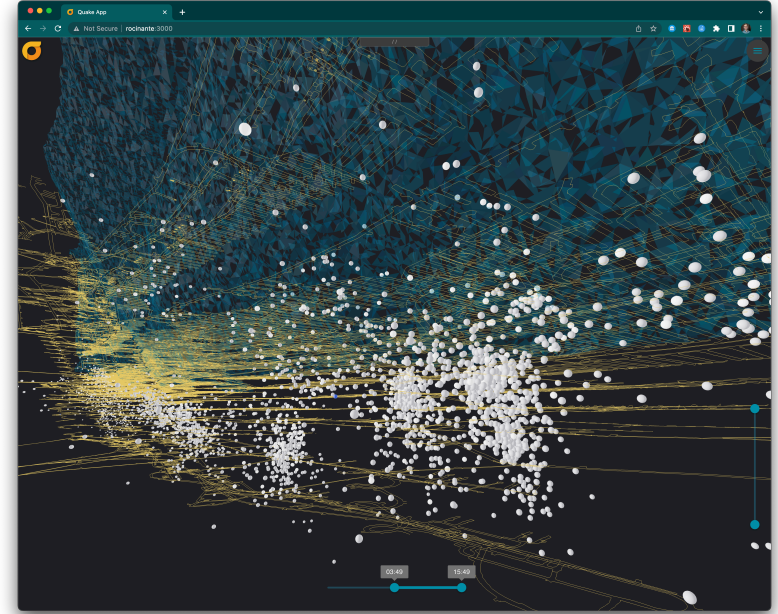
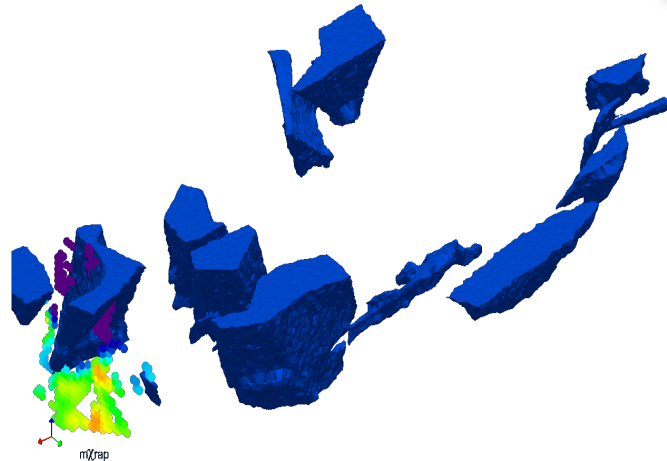
Separation of concerns

- Users describe meta function calls
 - Do not need to understand the underlying platforms
-
- Abstracts away complex platforms
 - Enables a **loosely coupled system**
 - Ledger
 - Dynamic allocation of resources
 - Optimize performance, scalability, and energy
-
- Executors are microservices designed to execute specific functions
 - Integrate with other platforms
 - System integrator
 - Reside anywhere on the Internet

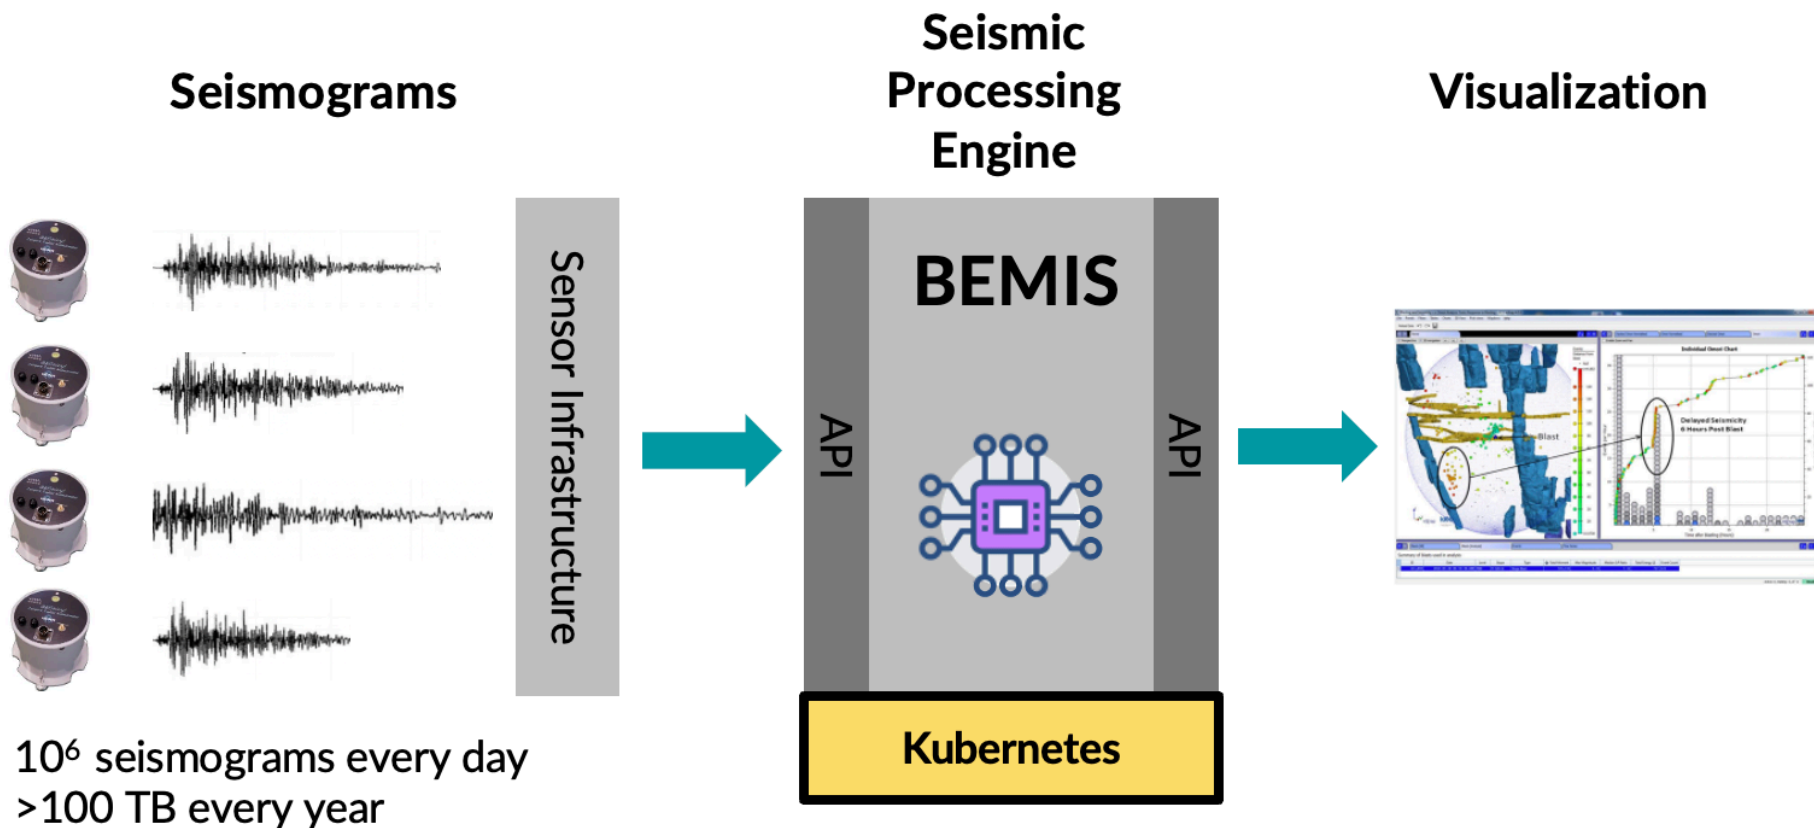


RockSigma AB

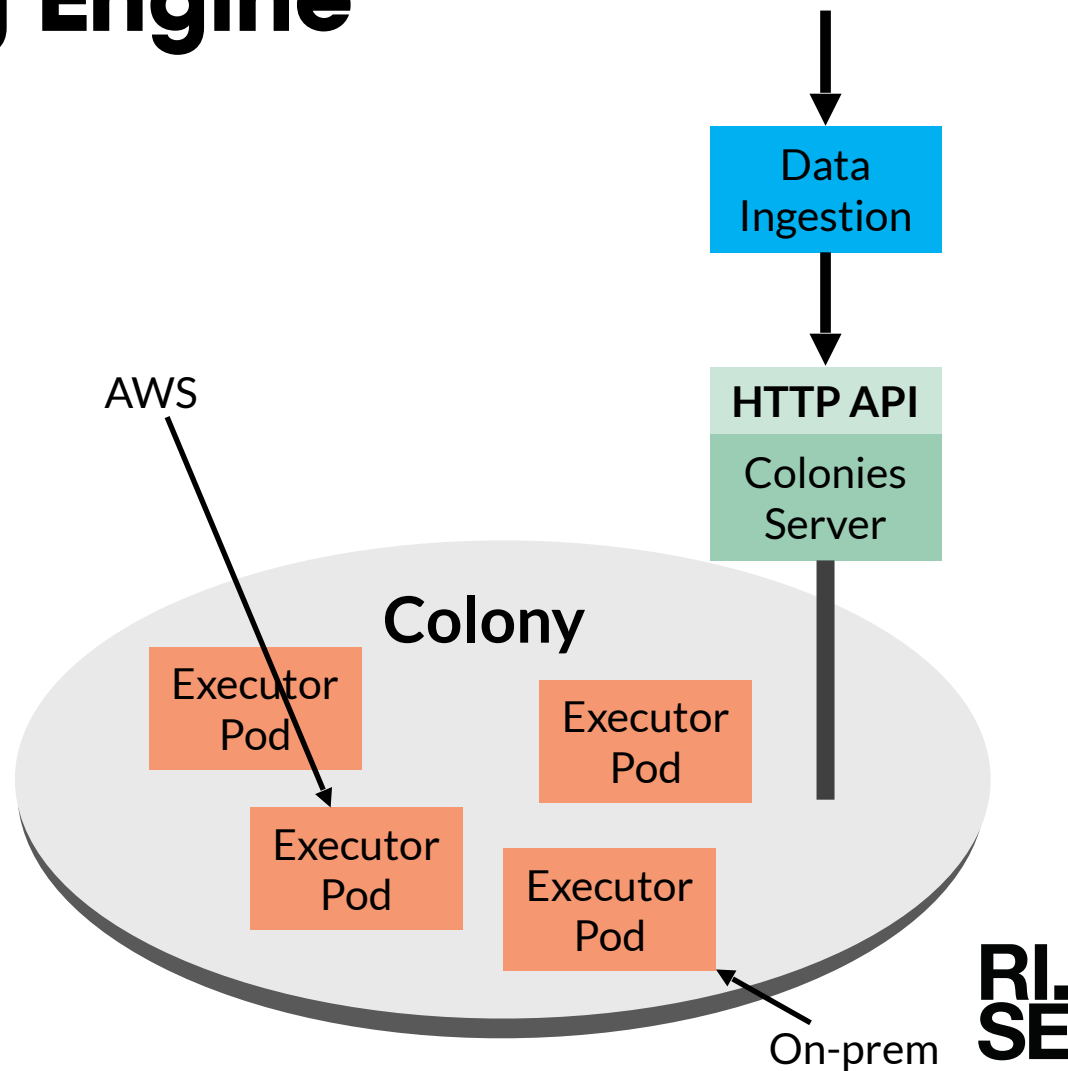
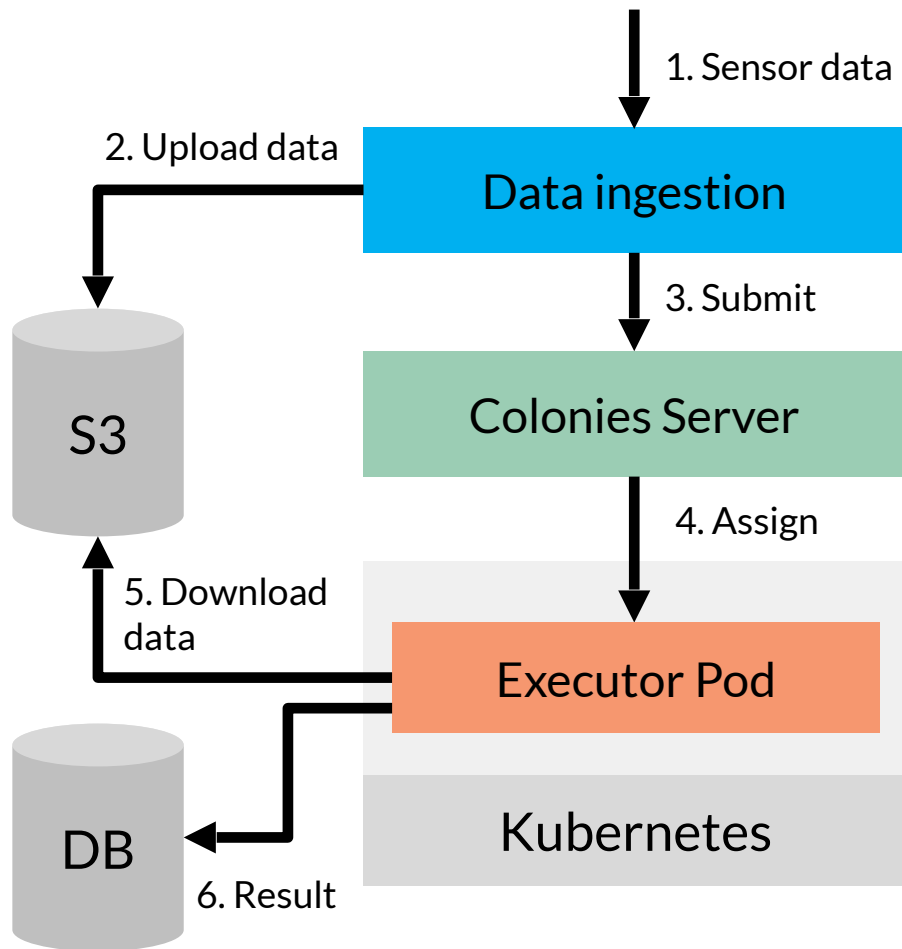
- Seismic processing underground mines
- Used by LKAB to analyze seismicity and process a massive amount of data from one the largest mines in the world (Kiruna/Malmberget)
- On-prem + cloud

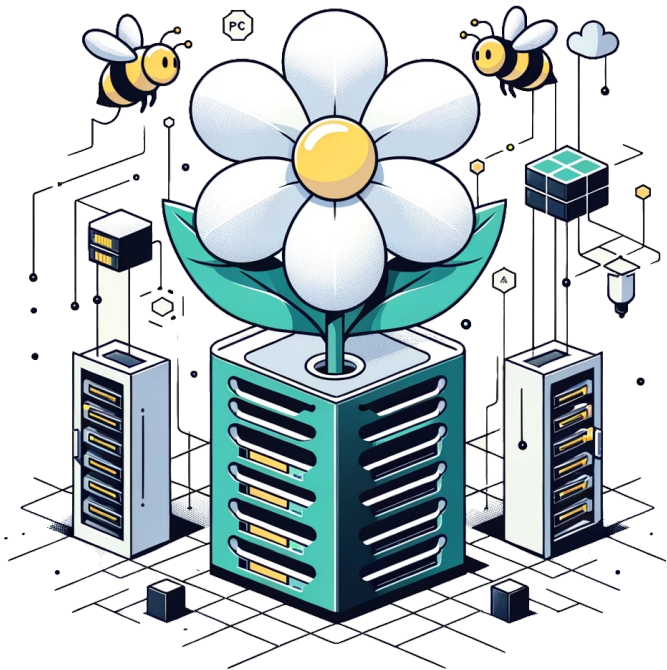


RockSigma AB



A Seismic Processing Engine





Pollinator

HTTP API

Colonies Server

Colony

LUMI HPC
Executor

Leonardo HPC
Executor

Local Executor

K8s Executor

Pollinator

Pollinator provides a **PaaS alike** user experience for ML development on HPC & K8s

Eliminates the need to learn Slurm, Kubernetes

Execution history (Ledger)

[illegible]

The image is a screenshot of the AWS SageMaker Studio console. On the left is a dark sidebar with navigation icons and labels: Dashboard, Users, Executors, Functions, Processes, Workflows, Filesystem, Cron, Generators, and Server. The main area on the right is titled 'Logs' and shows the output of a training process. The log text includes 'Process Id - 7', 'Pulling from job', 'Digest: sha256:', 'Status: Image id', '2024-09-24 10:12', 'To enable the f', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', '2024-09-24 10:12', 'XXXXXXXXXXXXXXXXXXXX', and 'Model: "model"'. Below this, a table-like structure shows the model architecture layers: 'Layer (type)' followed by 'input_1 (Input', 'conv2d (Conv2D', 'conv2d_1 (Conv2', 'max_pooling2d', 'dropout (Dropou', 'conv2d_2 (Conv2', 'conv2d_3 (Conv2', 'max_pooling2d_', 'dropout_1 (Drop', and 'conv2d_4 (f'.

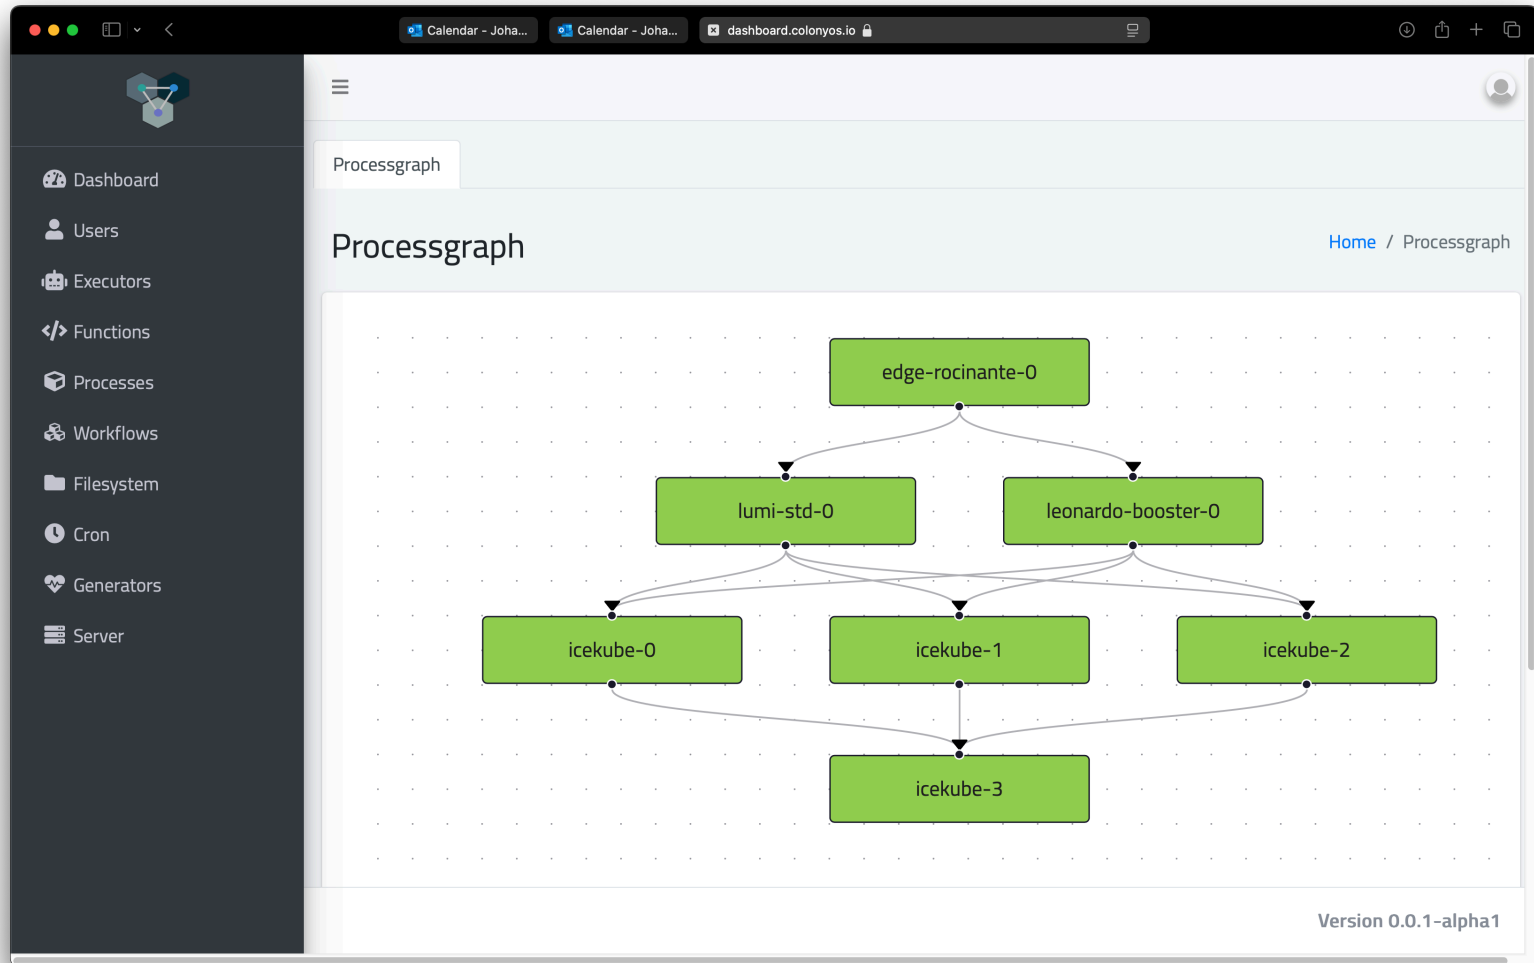
```
monad (~) >>> colonies process get -p 769e1d368f85a80843c8eb5e56665f1d8f3aa827b37cfe1cc02bf1ad9c2be098
```

Process	
Id	769e1d368f85a80843c8eb5e56665f1d8f3aa827b37cfe1cc02bf1ad9c2be098
IsAssigned	True
InitiatorID	bcaeac1a507036f7fed0be9d38c43ba973be7c0064d1b0b010ede2f088093b3f
Initiator	johan
AssignedExecutorID	7fecd3bcbe700bfc69d623bd75068f1e515a9f25102f30ff11e67caef41c287a
AssignedExecutorID	Successful
PriorityTime	1726740741634350173
SubmissionTime	2024-09-24 12:12:21
StartTime	2024-09-24 12:12:21
EndTime	2024-09-24 12:12:21
WaitDeadline	0001-01-01 00:53:28
ExecDeadline	2024-09-24 12:22:20
WaitingTime	29.346ms
ProcessingTime	17.591476s
Retries	0
Input	
Output	
Errors	

Function Specification	
Func	execute
Args	None
KwArgs	docker-image:johan/hackaton init-cmd: rebuild-imag...
MaxWaitTime	-1
MaxExecTime	599
MaxRetries	3
Label	

Conditions	
Colony	hpc
ExecutorNames	edge
ExecutorType	container-executor
Dependencies	
Nodes	1
CPU	10000m
Memory	15000Mi
Processes	0
ProcessesPerNode	0
Storage	0Mi
Walltime	600
GPUName	
GPUs	1

Cross-platform workflows



Files

main

Go to file

01-getting-started

README.md

cat.json

echo.json

gen_file.json

overview.png

overview.pptx

python.json

python_snapshot.json

02-colonyfs

03-python

04-faas

05-workflows

06-crons

07-pollinator

08-security

09-production

10-k8s-executor

11-hpc-executor

12-anomaly-detection

13-earth-observation

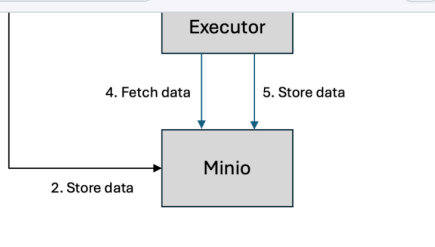
LICENSE

README.md

tutorials / 01-getting-started / README.md

PreviewCodeBlame330 lines (277 loc) · 23 KB

RawDownloadEdit



```
graph TD; Executor[Executor] -- "2. Store data" --> Minio[Minio]; Minio -- "4. Fetch data" --> Executor; Minio -- "5. Store data" --> Executor;
```

Setting up a development environment

The following commands will use Docker Compose to set up and configure a Colonies server, a TimescaleDB, a Minio server, and a Docker Executor. To set up a production environment, it is recommended to use Kubernetes.

Note! The `docker-compose.env` file contains credentials and configuration and must be sourced before using the Colonies CLI command.

On Mac or Linux type:

```
wget https://raw.githubusercontent.com/colonyos/colonies/main/docs/docker-compose.env;
source docker-compose.env;
wget https://raw.githubusercontent.com/colonyos/colonies/main/docs/docker-compose.up
docker-compose up
```

On Windows type:

```
wget https://raw.githubusercontent.com/colonyos/colonies/main/windowsenv.bat
wget https://raw.githubusercontent.com/colonyos/colonies/main/docs/docker-compose.up
docker-compose up
```

Note that all three commands must be types separately on Windows.

Press control-c to exit.

To remove all data, type:

<https://github.com/colonyos/tutorials>

