NEARDATA
Extreme Near-data Processing Platform

NEARDATA
EXTREME NEAR-DATA PROCESSING PLATFORM

List of participants

Participant no.	Participant organisation name	Country
1 (Coordinator)	Universitat Rovira i Virgili (URV)	Spain
2	Barcelona Supercomputing Center (BSC)	Spain
3	Technical University Dresden (TUD)	Germany
4	German Cancer Research Center (NCT)	Germany
5	European Molecular Biology Laboratory (EMBL)	Germany
6	Centre for Computational Medicine (SANO)	Poland
7	KIO networks (KIO)	Spain
8	SCONTAIN (SCO)	Germany
9	Dell (DELL)	Ireland
10	UK Health Security Agency (UKHS)	United Kingdom

NEARDATA

HORIZON-CL4-2022-DATA-01-05:

Extreme data mining, aggregation and analytics technologies and solutions

Provide better technologies, tools and solutions for **data mining** (searching and processing) of extreme data.

Extreme data is defined as data that exhibits one or more of the following characteristics, to an extent that makes current technologies fail: increasing **volume**, **speed**, **variety**; **complexity/diversity/multilinguality** of data; the **dispersed data sources**; **sparse/missing/insufficient** data/extreme variations in values).

The technologies and solutions are expected to discover and **distil meaningful, reliable and useful data** from heterogeneous and dispersed/scarce sources and deliver it to the requesting application/user **with minimal delay and in the appropriate format**.

NEARDATA

Extreme Near-Data Processing Platform

Why **Locality** ?

- Volume
- Privacy
- Low latency
- Hardware Acceleration

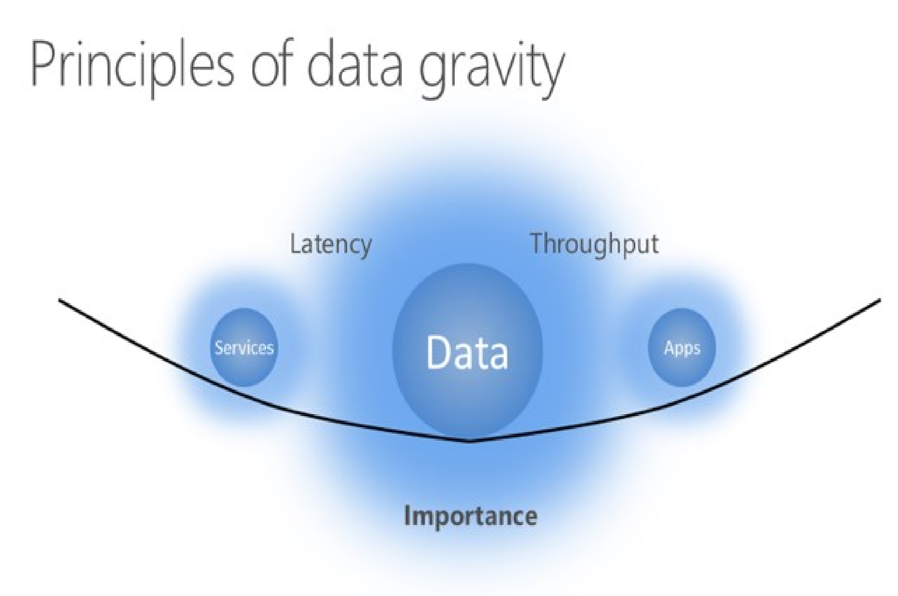

NEARDATA

Objectives

The main goal is to design an Extreme near-data processing platform to enable consumption, mining and processing of distributed and federated data without needing to master the logistics of data access across heterogeneous data locations and pools.

- **O-1 Provide high-performance near-data processing for Extreme Data Types:** The first objective is to create a novel intermediary data service (Data Plane / Xtreme DataHub) between Object Storage and Analytic platforms.
- **O-2 Support real-time video streams but also event streams that must be ingested and processed very fast to Object Storage:** The second objective is to seamlessly combine streaming and batch data processing for analytics.
- **O-3 Provide secure data orchestration, transfer, processing and access:** The third objective is to create a Data Broker service enabling trustworthy data sharing and confidential orchestration of data pipelines across the Compute Continuum.

NEARDATA

KPIs

- KPI-1 - Significant performance improvements (data throughput, data transfer reduction) in Extract-Transform-Load (ETL) phases validated with near-data connectors over extreme data volumes (genomics, metabolomics).
- KPI-2 - Significant data speed improvements (throughput, latency) in real-time video analytics validated using stream data connectors.
- KPI-3 - Demonstrated resource auto-scaling for batch and stream data processing validated thanks to data-driven orchestration of massive workflows.
- KPI-4 - High levels of data security and confidential computing validated using TEEs and federated learning in adversarial security experiments.
- KPI-5 - Demonstrated simplicity and productivity of the software platform validated with real user communities in International Health Data Spaces.

NEARDATA

Reference Architecture

INTERNATIONAL DATA SPACES ASSOCIATION

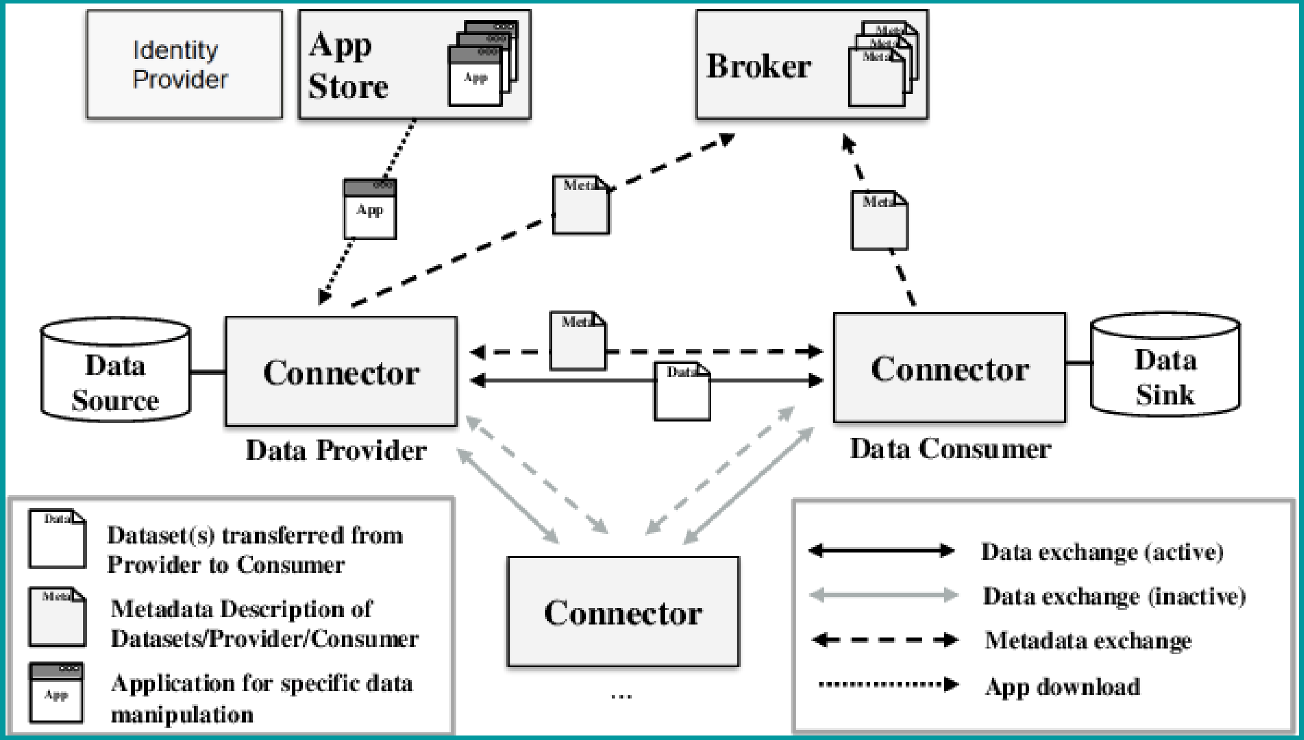

NEARDATA

NEARDATA Architecture

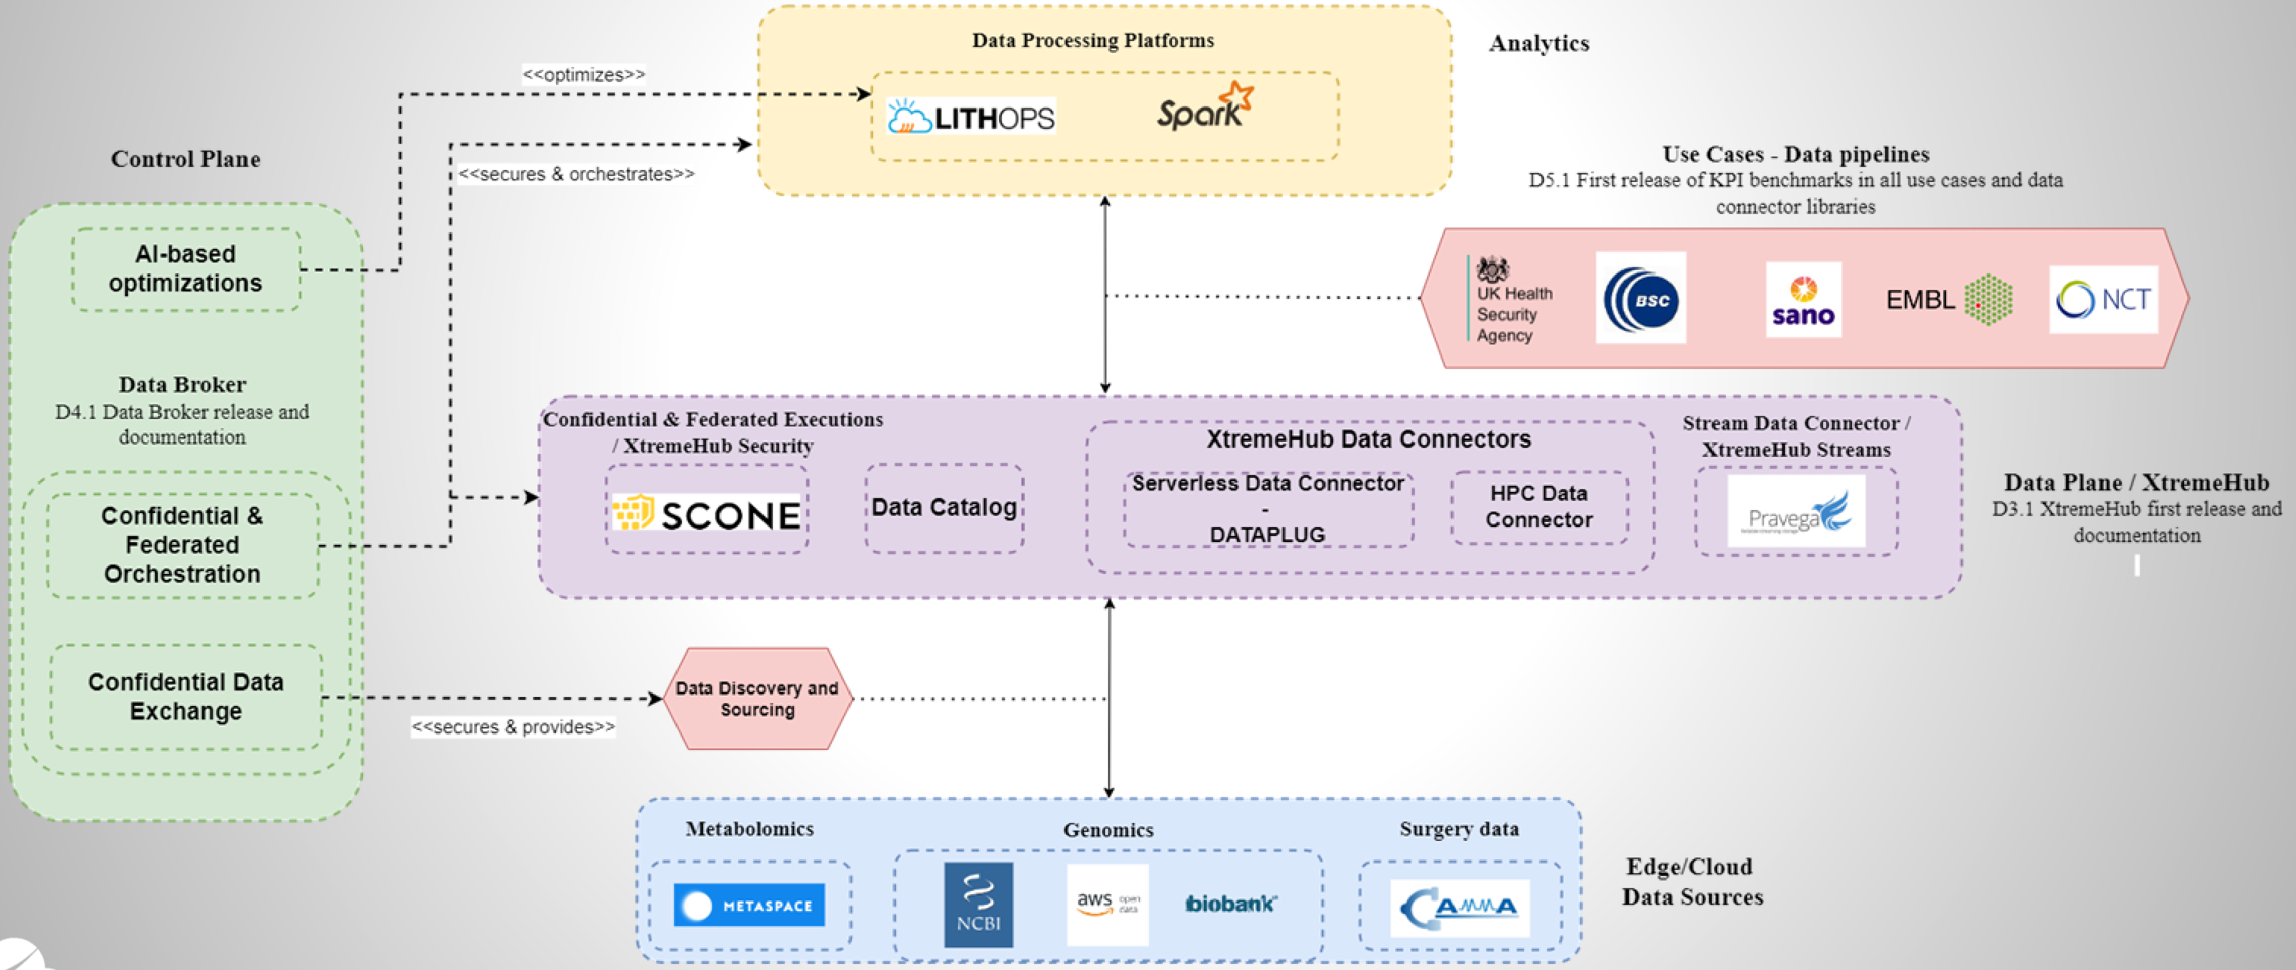

NEARDATA

NEARDATA Architecture

Analytics Module

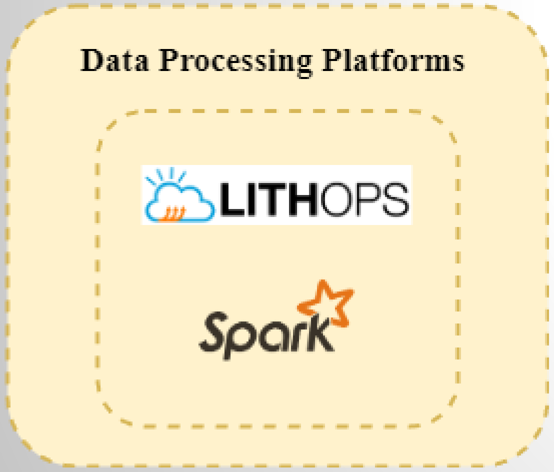

The **Analytics module** incorporates Data Analytics platforms capable of processing extreme data from mechanisms that ensure large-scale computation.

- **Lithops** is a Python multi-cloud distributed computing framework. It allows you to run unmodified local python code at massive scale in the main serverless computing platforms. Lithops delivers the user's code into the cloud without requiring knowledge of how it is deployed and run. Moreover, its multicloud-agnostic architecture ensures portability across cloud providers. [1, 2]

- **Apache Spark** is one of the most widely used data analytics platforms by data scientists capable of being deployed on multiple clusters adapting to the workload and scale of the data.

[1] G. Finol, G. París, P. García-López, and M. Sánchez-Artigas, "Exploiting inherent elasticity of serverless in algorithms with unbalanced and irregular workloads," Journal of Parallel and Distributed Computing, vol. 190, p. 104891, 2024.

[2] G. T. Eizaguirre and M. Sánchez-Artigas, "A seer knows best: Auto-tuned object storage shuffling for serverless analytics," Journal of Parallel and Distributed Computing, vol. 183, p. 104763,2024.

NEARDATA Architecture

Run Lithops Cloud

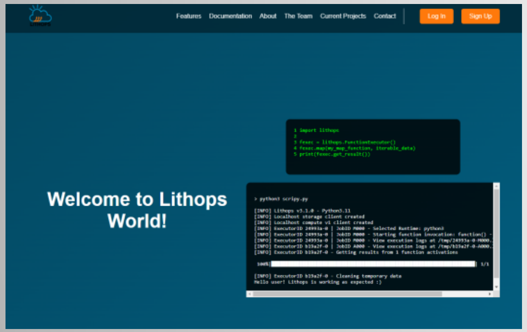

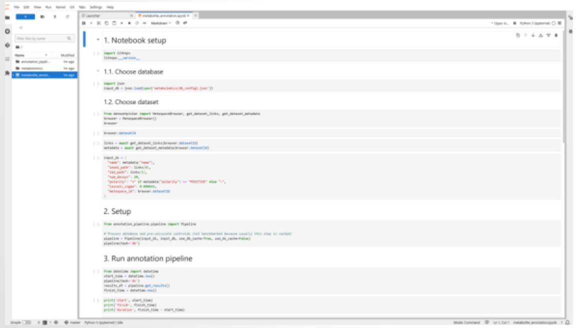

NEARDATA

- **Run.Lithops.Cloud** is a service that allows users to develop and execute code in the cloud. It is currently focused on simplifying the execution of **Lithops** pipelines.
- The platform allows the development and execution of pipelines, experiments, among others, through **Jupyter Notebooks** directly from a web browser.
- We plan to incorporate all project **pipelines** that take advantage of Lithops. Currently Metabolite annotation (metabolomics) and Serverless Variant Calling (genomics) are available.
- Integration with public data sources (METASPACE catalog).
- Real-time **Metrics**.

NEARDATA Architecture

Data Sources: Edge/Cloud Module

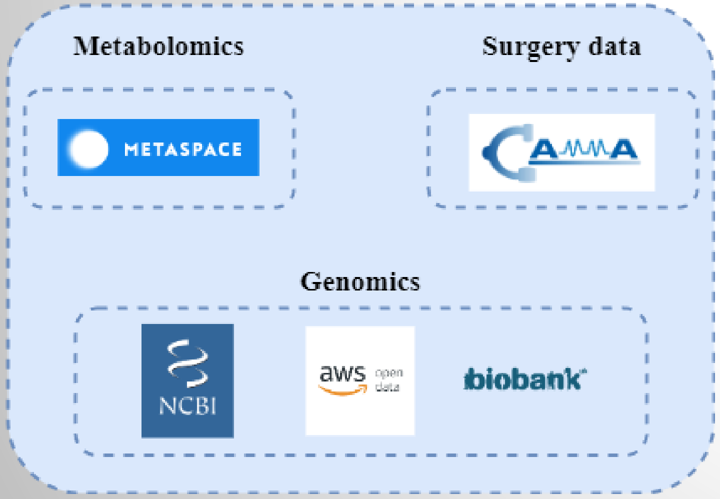

The **Data Sources module** incorporates different massive sources of health data specific to each format.

- **Metabolomics:**
  - The **METASPACE** platform hosts an engine for metabolite annotation of imaging mass spectrometry data as well as a spatial metabolite knowledgebase of the metabolites from thousands of public datasets provided by the community. METASPACE is implemented on top of **Lithops**.
  - **Metabolights** and **Metabolomics Workbench.**
- **Genomics:**
  - **National Center for Biotechnology Information** (NCBI), **UK Biobank** and **AWS Open Data**.
- **Surgery:**
  - **Camma** and **Endoscopic Vision Challenge**

NEARDATA Architecture

Data Plane: XtremeHub Module

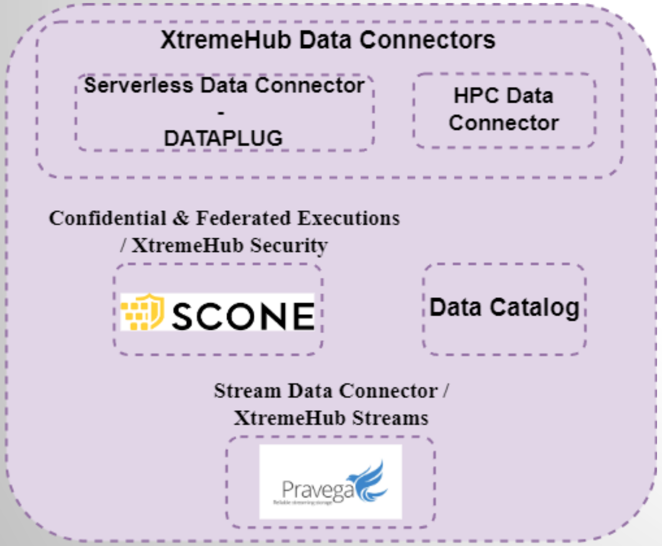

NEARDATA

The **Data Plane / XtremeHub platform module** presents different tools that facilitate the ingestion and management of massive unstructured data.

- **Confidential & Federated Execution (XtremeHub Security):**
  - **SCONE** is a platform to build and run secure applications with the help of Intel SGX (Software Guard eXtensions). In a nutshell, the SCONE objective is to run applications such that data is always encrypted, i.e., all data at rest, all data on the wire as well as all data in main memory is encrypted. [3, 4]
- **XtremeHub Data Connectors:**
  - Serverless Data Connector: Dataplug & HPC Data Connector
- **Stream Data Connector (XtremeHub Streams):**
  - **Pravega** is an open source distributed storage service implementing Streams. It offers Stream as the main primitive for the foundation of reliable storage systems: a high-performance, durable, elastic, and unlimited append-only byte stream with strict ordering and consistency. [5, 6]

[3] A. Galanou, K. Bindlish, L. Preibsch, Y.-A. Pignolet, C. Fetzer, and R. Kapitza, "Trustworthy confidential virtual machines for the masses," in Proceedings of the 24th International Middleware Conference, Middleware '23, (New York, NY, USA), p. 316–328, Association for Computing Machinery, 2023.

[4] G. P. Fernandez, A. Brito, A. P. P. Hartono, M. U. Sardar, and C. Fetzer, "Lld: A last-level defense for application integrity and confidentiality," in Proceedings of the IEEE/ACM 16th International Conference on Utility and Cloud Computing, UCC '23, (New York, NY, USA),Association for Computing Machinery, 2024.

[5] R. Gracia-Tinedo, F. Junqueira, T. Kaitchuck, and S. Joshi, "Pravega: A tiered storage system for data streams," in Proceedings of the 24th International Middleware Conference, Middleware '23,(New York, NY, USA), p. 165–177, Association for Computing Machinery, 2023.

[6] R. Gracia-Tinedo, F. Junqueira, B. Zhou, Y. Xiong, and L. Liu, "Practical storage-compute elasticity for stream data processing," in Proceedings of the 24th International Middleware Conference: Industrial Track, Middleware '23, (New York, NY, USA), p. 1–7, Association for Computing Machinery, 2023.

NEARDATA Architecture

Dataplug - Serverless Data Connector

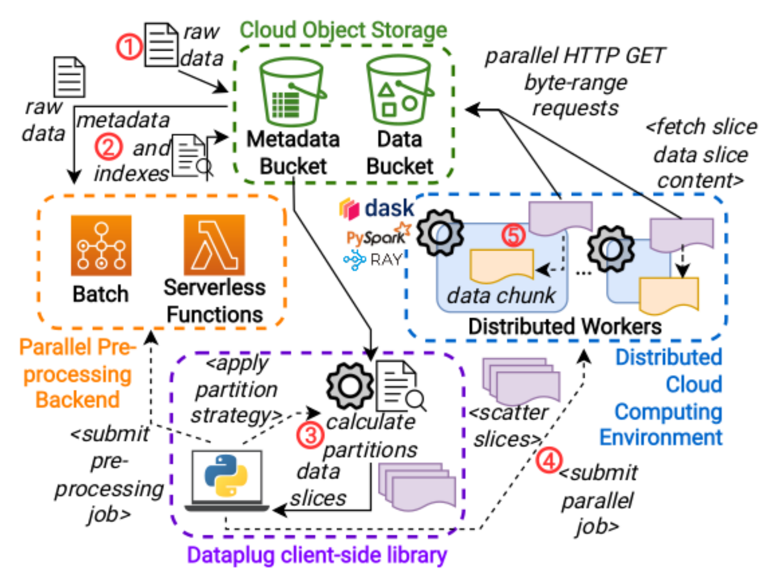

Dataplug is an extensible framework that implements the on-the-fly data partitioning model. Dataplug enables parallel access to unstructured scientific datasets efficiently. [7]

Dataplug reduces 200% data transfer and data duplication by avoiding the use of traditional static partitioning.

Supported scientific formats:

- Genomics: FASTA, FASTQGZIP, VCF
- Geospatial data: LIDAR
- Metabolomics: imzML

[7] A. Arjona, P. García-López and D. Barcelona-Pons, "Dataplug: Unlocking extreme data analytics with on-the-fly dynamic partitioning of unstructured data" in the 24th IEEE/ACM international Symposium on Cluster, Cloud and Internet Computing (CCGrid'24), May 6-9, 2024, Philadelphia, USA.

NEARDATA Architecture

HPC Data Connector

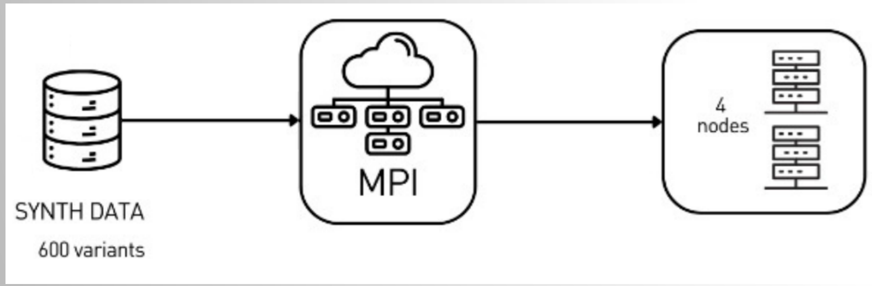

- New HPC Data Connector to leverage High-Performance Computing platforms using MPI.
- Allow use-cases to use supercomputing facilities, and most particularly MareNostrum supercomputer hosted by BSC-CNS.
- Developed for the genomic use-case MDR (BSC)

NEARDATA

NEARDATA Architecture

Data Plane: XtremeHub Module

Results I

- **Lithops:**
  - **KPI-3**: Incorporates an extensible storage and compute backend architecture that enables elastic and scalable cloud solutions to be designed according to the resources needed to run the workload.
- **Pravega:**
  - **KPI-1:** Pravega shows the highest throughput (350 MBps) compared to Kafka (330 MBps) and Pulsar (250 MBps) when using multiple segments.
  - **KPI-2:** The Pravega reader achieves both low end-to-end latency and high throughput compared to Kafka and Pulsar for the cases tested.
  - **KPI-3:** Pravega is the first streaming storage system that provides elastic streams: data streams that are automatically re-partitioned according to the ingestion load and the scaling policy.
  - **KPI-5:** Pravega can seamlessly handle streaming and batch serverless workloads with the same API, thus making it easier for developers to manage data in serverless functions compared to having to use a different storage system per workload

NEARDATA Architecture

Data Plane: XtremeHub Module

Results II

SCONE

- **SCONE:**
  - **KPI-4**: SCONE supports confidential execution aided by TEE (hardware enabler) of both Lithops and Flower ML, two systems developed in Python. Flower ML can run attested as a standalone application. Lithops attestation for its serverless mode is a work in progress.
- **Dataplug:**
  - **KPI-1:** Dataplug offers data transfer reduction (200%) and data throughput improvements in preprocessing tasks (x2,9 faster in FASTQGZip indexing and x3,7 faster in FASTQGZip fetching partitions).
- **HPC Data Connector:**
  - **KPI-1:** Initial experiments show that the MDR use-case implemented with MPI improves performance x5 times compared to the previous cloud-based solution with Apache Spark.
  - **KPI-3:** Future work - Integration with Lithops.

NEARDATA

NEARDATA Architecture

Control Plane: Data Broker Module

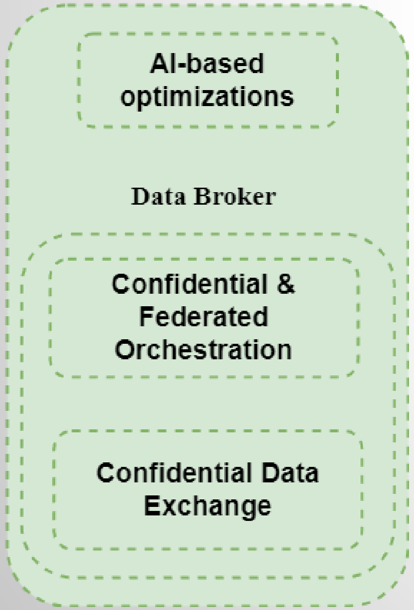

The **Control Plane** is the major front-end of the NEARDATA platform which includes both data discovery, governance and access but also optimized orchestration and declarative interconnection of heterogeneous data flows.

- The **Data Broker** is the cockpit of the NEARDATA platform, exposing and orchestrating all services in the Data and Control planes.
- The **AI-based Optimizer of Cloud/Edge Workflows** is a learning service that focuses on improving data-driven orchestration of workloads and pipelines defined in the Orchestration layer.

NEARDATA

NEARDATA Architecture

Control Plane: Data Broker Module

Results

- **SCONE – KPI 4:**
  - **Surgery Use-case:**
    - Examine how security such as confidentiality and integrity can be provided for applications such as federated learning by harnessing technologies such as Trusted Execution Environments (TEEs).
    - Assessing the performance degradation of the confidential federated learning application by analysing the impact of the network shield, configured through the Data Broker's CAS component, on execution time.
    - Assessing the impact on the initial start up and response time of functions in the federated learning application frameworks such as Flower.
  - **Metabolomics:**
    - Integration of the Confidential Compute Layer component of the Data Broker into Lithops and tested its early usage on an on-premise K8s cluster.

NEARDATA

NEARDATA Health Data Spaces

Metabolomics

EMBL

Genomics

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

UK Health Security Agency

sano

Surgomics

NCT

NEARDATA

Metabolomics Data Space

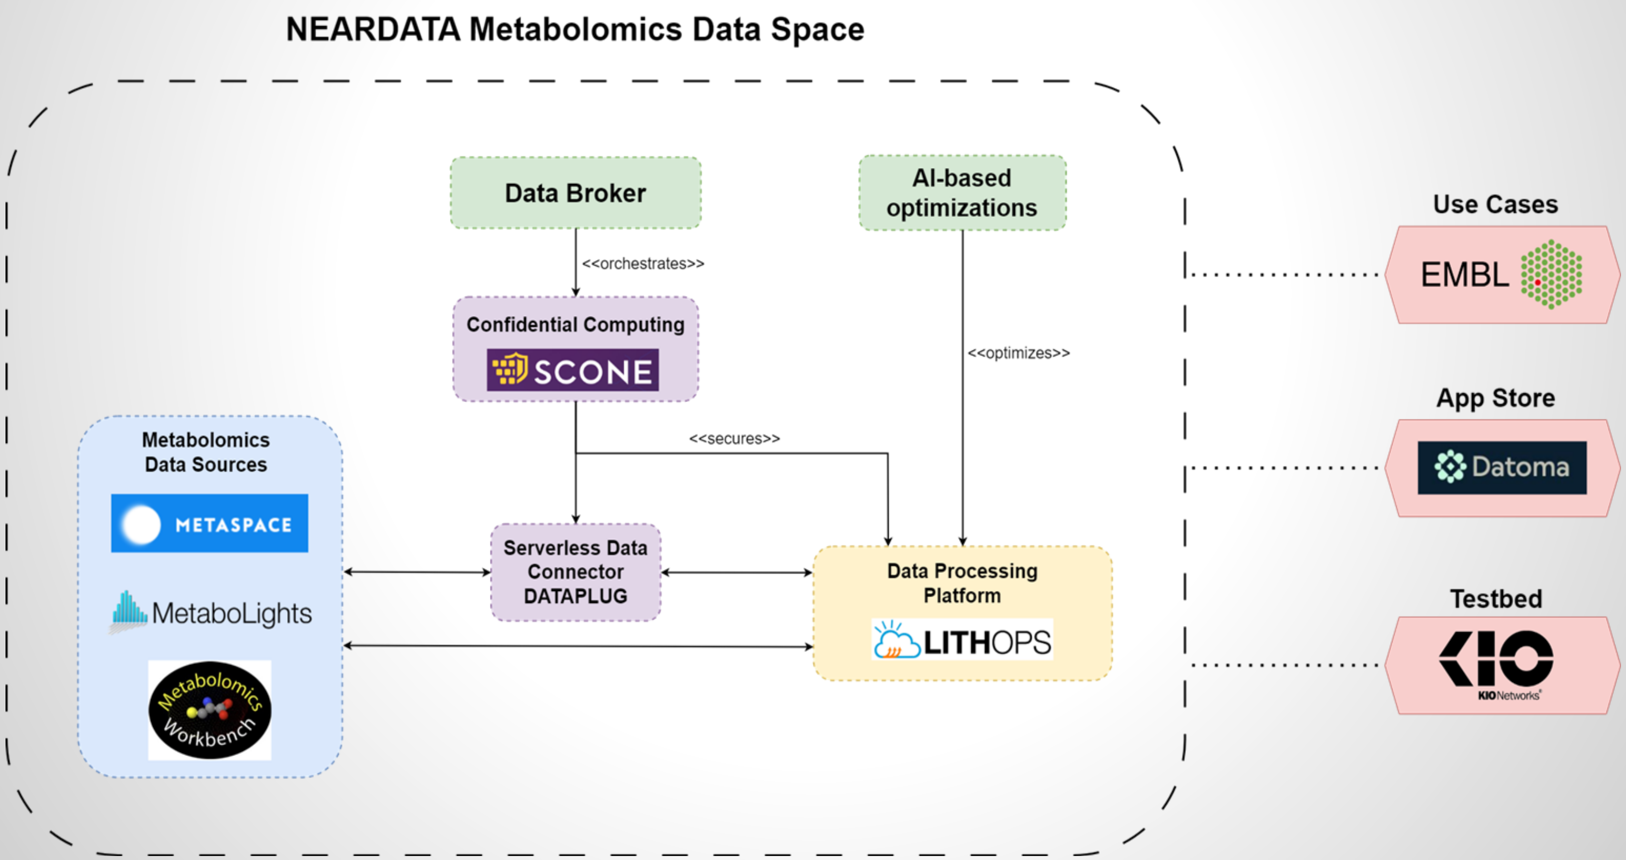

Metabolomics Data Space

Metabolomics Use-Case (EMBL)

- **Experiment 1:** Machine learning-based metabolite identification using Lithops and METASPACE.
- **Experiment 2:** Federated METASPACE

- **Experiment 1:**
  - **KPI-1:**
    - Dataplug offers partitioning strategies for metabolomics data formats (ImzML).
    - The ML-based version of the metabolite identification using Lithops eliminates the need for data transfer, as ML inference can be performed directly in METASPACE.
  - **KPI-3:** The implemented ML-version of metabolite identification allows for resource auto-scaling for datasets of the size ranging from under 1 GB to 20 GB.
  - **KPI-5:** The ML-based metabolite identification is already available to users on the production version of METASPACE and is already used by the METASPACE users.
- **Experiment 2:**
  - Integration of the Confidential Compute Layer component of the Data Broker into Lithops and tested its early usage on an on-premise K8s cluster.

B. Wadie, L. Stuart, C. M. Rath, B. Drotleff, S. Mamedov, and T. Alexandrov, “Metaspace-ml:Metabolite annotation for imaging mass spectrometry using machine learning,” bioRxiv, 2024.

G. T. Eizaguirre, D. Barcelona-Pons, A. Arjona, G. Vernik, P. García-López and T. Alexandrov, "Serverful Functions: Leveraging Servers in Complex Serverless Workflows (industry track)" in Middleware Industrial Track ’24, December 02–06, 2024, Hong Kong, China.

Genomics Data Space

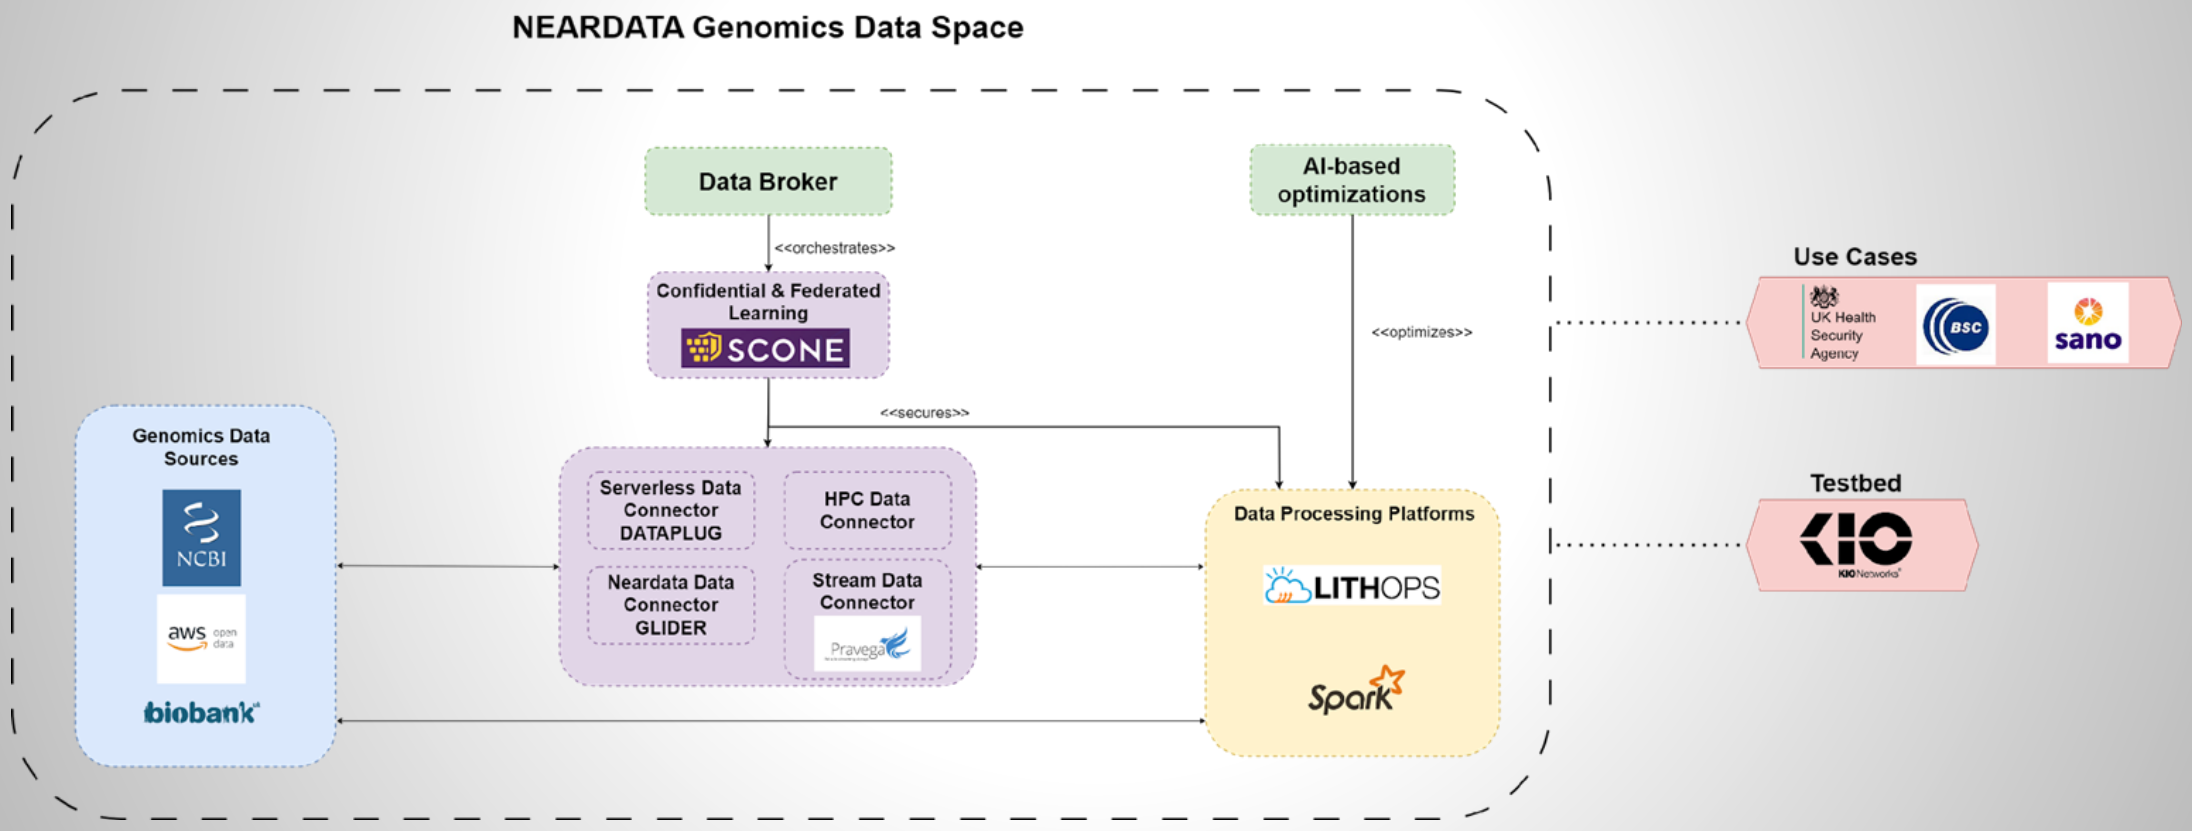

NEARDATA

Genomics Data Space New components

Genomics Use Case – UKHS
Glider: Neardata Data Connector

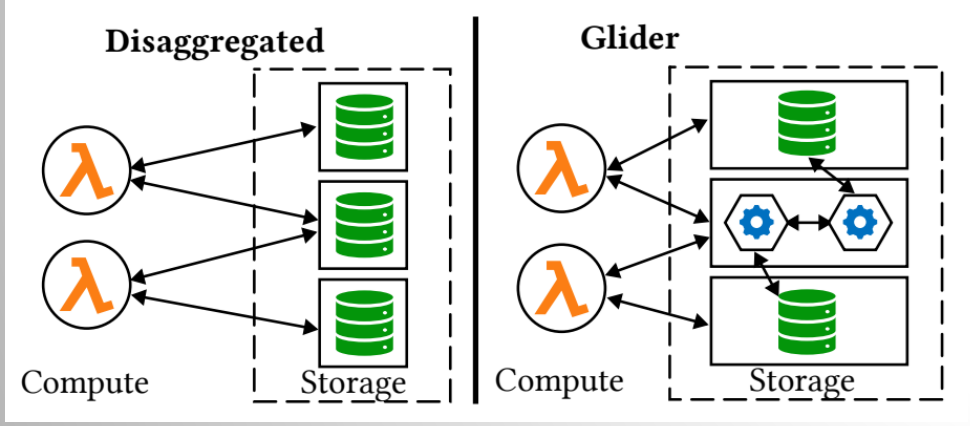

Glider is an ephemeral storage system incorporating near-data computation to reduce the volume of data transferred for our serverless architecture. [8]

Glider aims to improve communication between serverless computing stages, allowing data to "glide" smoothly through the processing pipeline rather than bouncing between different services.

[8] D. Barcelona-Pons, P. García-López, and B. Metzler, "Glider: Serverless ephemeral stateful near-data computation," in Proceedings of the 24th International Middleware Conference, Middle-ware '23, (New York, NY, USA), p. 247–260, Association for Computing Machinery, 2023.

NEARDATA

Genomics Data Space

Genomics Pathogens (Serverless Variant Calling) (UKHS)

- **Porting an HPC version of a genomics variant calling to a serverless architecture in the cloud.**

- **KPI-1:** Reduced data partitioning, data transfer (by 200%), and data duplication through the integration of the Dataplug connector, as we avoid re-uploading partitions to the object storage.

- **KPI-3:**
  - Moving the variant calling pipeline from HPC to serverless with Lithops demonstrates a further optimization of the resources needed for the massively parallelized pipeline thanks to the elasticity and flexibility of cloud services. Thanks to this, the serverless version is **x37.46** times faster than the HPC version.
  - Glider integration on the serverless version reduces execution time by **36%** with the full data. The Neardata data connector demonstrates a reduction of the data transfer ensuring the workflow orchestration from the distribution of the actions.

A. Arjona, A. Gabriel-Atienza, S. Lanuza-Orna, X. Roca-Canals, A. Bourramouss, T. K. Chafin,L. Marcello, P. Ribeca, and P. García-López, “Scaling a variant calling genomics pipeline with faas,” in Proceedings of the 9th International Workshop on Serverless Computing, WoSC ’23,(New York, NY, USA), p. 59–64, Association for Computing Machinery, 2023.

NEARDATA

Genomics Data Space New components

Genomics Use Case – BSC Lithops Singularity

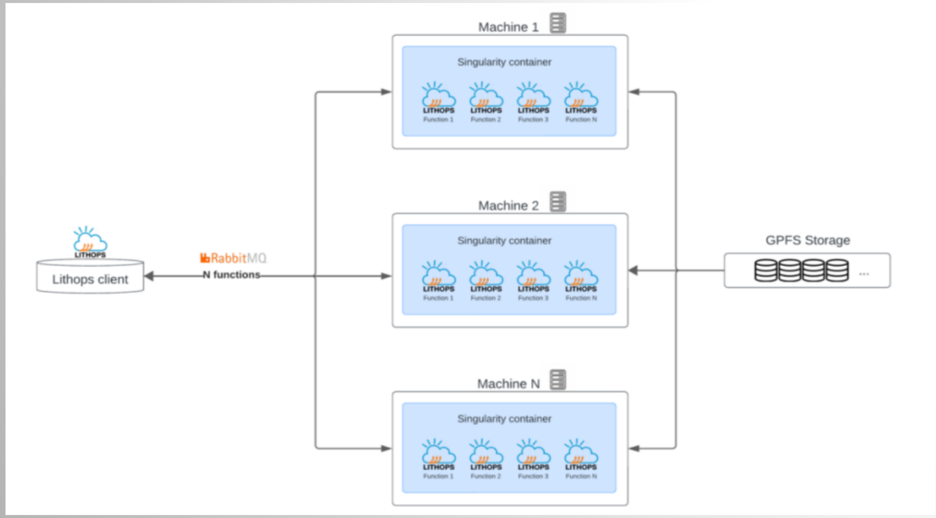

Cloud computing paradigms such as Serverless and Function-as-a-Service offering high parallelism capabilities, feasible deployment, and elasticity, are captivating to combine with the HPC power.

Previous attempts to bind HPC infrastructure and cloud elasticity have faced problems such as the need for user privileges or lack of full elasticity.

The new **singularity-based** architecture for **Lithops** overcomes those restrictions and allows its execution on HPC systems such as the MareNostrum (MN5) supercomputer. [9]

[9] A. Benavides-Arevalo, D. Coll-Tejeda, A. Call, P. García-López, R. Nou-Castell, "Enhancing HPC with Serverless Computing: Lithops on MareNostrum5" in Cloud-Edge Continuum (CEC) Workshop 2024, October 28th, 2024, Charleroi, Belgium.

NEARDATA

Genomics Data Space

Variant-Interactions (BSC)

- **Multi Dimensionality Reduction MDR:** use statistical methods to discover pairs of variants which, synergically, contribute to the development of T2D.

- **Genome-Wide discovery (GWD):** use machine learning methods to find groups of variants that, simultaneously, are associated with T2D.

NEARDATA

- **MDR:**
  - **KPI-1:** MPI version shows a speed-up of 5x compared to the Apache Spark version which translates into 5x more data ingestion capabilities. Experiments were performed on the MareNostrum supercomputer

- **GWD:**
  - **KPI-1**: We explored the adoption of a GPU component in our HPC data connector to further improve performance on heavy computational tasks. We have seen a potential speed-up of 2.1x times.
  - **KPI-3:** Leveraging Lithops Singularity can increase data ingestion capabilities, thus, improving data throughput. Initial experiments in local environments show a speed-up of 2.83x in performance.

G. Gómez-Sánchez, L. Alonso, M. Á. Pérez, I. Morán, D. Torrents, and J. L. Berral, "Exhaustivevariant interaction analysis using multifactor dimensionality reduction," Res. Sq., Oct. 2023.

G. Gómez-Sánchez, A. Call, X. Teruel, L. Alonso, I. Moran, M. Á. Pérez, D. Torrents, and J. L.Berral, "Challenges and opportunities for risc-v architectures towards genomics-based work-loads," in High Performance Computing (A. Bienz, M. Weiland, M. Baboulin, and C. Kruse,eds.), (Cham), pp. 458–471, Springer Nature Switzerland, 2023.

Genomics Data Space

Transcriptomics Use-Case (SANO)

- **Transcriptomics Atlas:** Porting an HPC version of the pipeline to the cloud.
- **Federated Learning for Human Genome Variation Analysis**.

- **KPI-1:**
  - Through the implementation of an early stopper for STAR aligner we can predict whether we should stop processing the alignment with just 10% of the total number of reads.
  - New generation of the STAR index on the newest genome release that resulted in 12-times faster pipeline.
  - The FL workflow yields a significant data transfer reduction and an improved data ingestion rates compared to centralized approaches due to inherent parallelism.
- **KPI-3:**
  - Index distribution solution has proved itself viable and scalable for the Transcriptomics Atlas architecture.
  - Using spot instances we are able to reduce compute costs by about 50%.
- **KPI-4:** Data transfer reduction due to FL nature guarantees a high level of privacy. Further integration with SCONE will only improve these guarantees.
- **KPI-5:** Increased regulatory compliance and data security due to FL nature increases the productivity of the platform.

P. Kica, S. Lichołai, M. Malawski, “Novel approaches toward scalable composable workflows in hyper-heterogeneous computing environments,” in Proceedings of the SC ’23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis, SC-W ’23, (New York, NY,USA), p. 2097–2108, Association for Computing Machinery, 2023.

NEARDATA

Surgomics Data Space

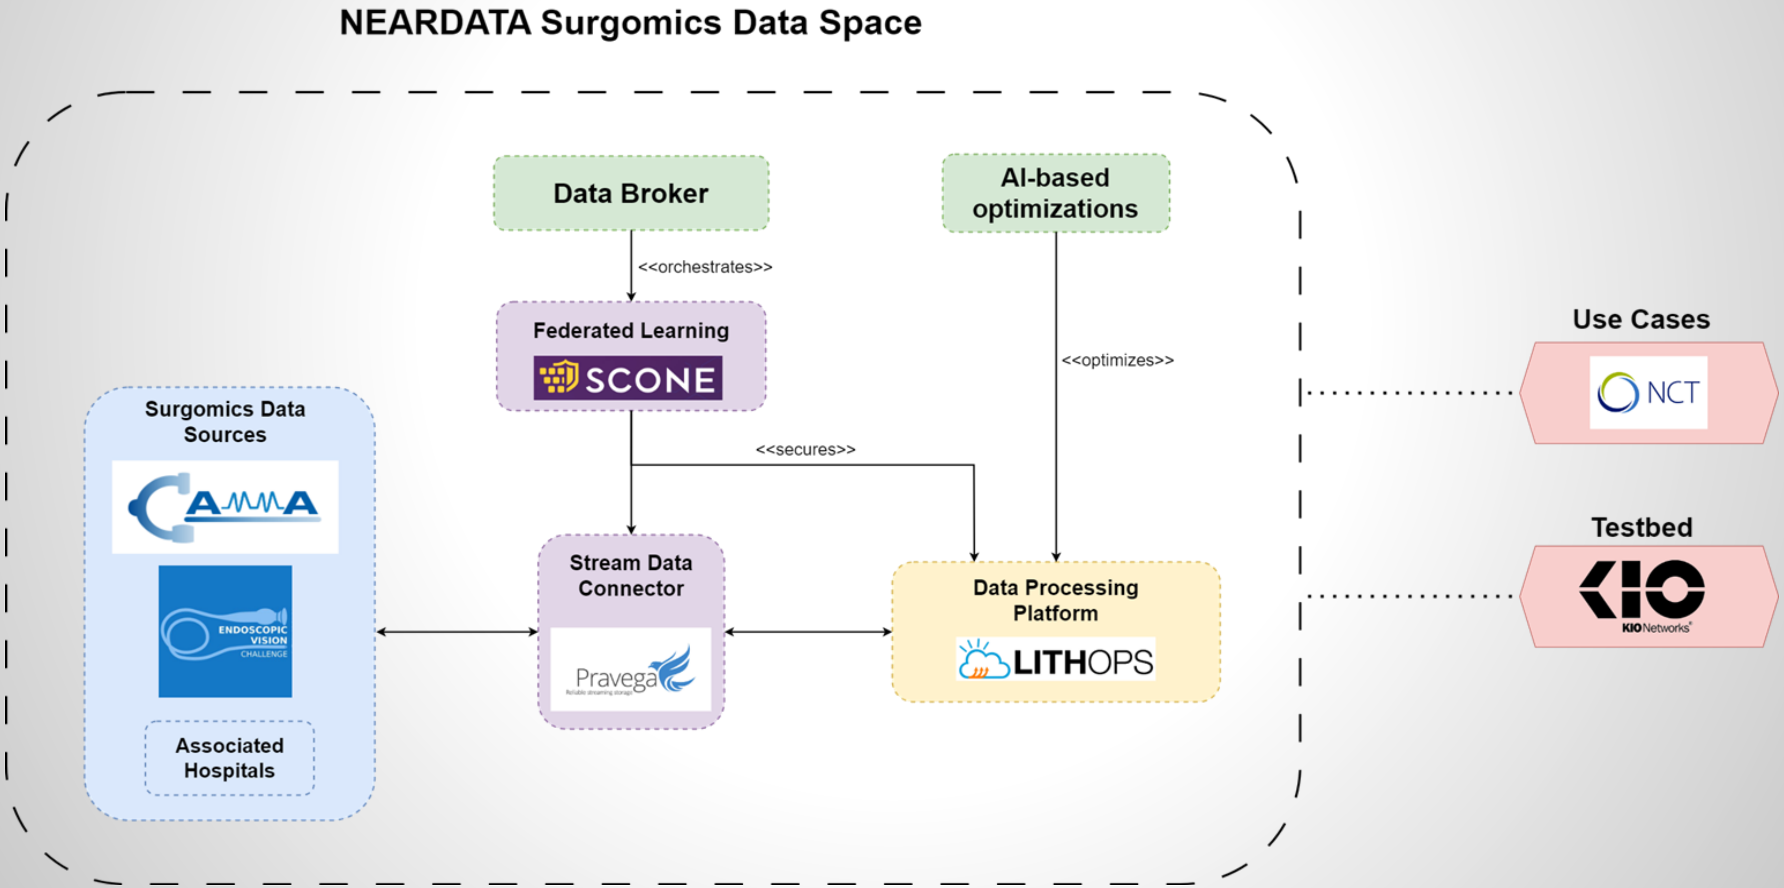

NEARDATA

Surgomics Data Space

Surgery Use-Case (NCT)

- **Federated Learning and development of a surgical video streaming application to handle multiple inference jobs.**

- **KPI-2:** With Pravega integration end-to-end IO latency represents only a fraction of the inference latency in the tested models (i.e., 45% lower).
- **KPI-3:** Combining federated with Scone will auto-scale the training of multiple clients with extreme-date achieving high performance. The orchestrator component auto-scales streaming analytics automatically leveraging Pravega's stream parallelism.
- **KPI-4:** Utilizing SCONE training data and code has been encrypted by making all training inside TEE enclaves. We guaranteed training computations ran correct code not modified by malicious clients. Preliminary results show such confidentiality is ensured.
- **KPI-5:** NCT estimates that using containerized AI inference applications jointly with Pravega/GStreamer may reduce the combined deployment and data management time for video analytics in 50%.

A. C. Jenke, S. Bodenstedt, F. R. Kolbinger, M. Distler, J. Weitz, and S. Speidel, “One model to use them all: training a segmentation model with complementary datasets,” International Journal of Computer Assisted Radiology and Surgery, vol. 19, pp. 1233–1241, June 2024.

G. Finol, A. Gabriel, P. García-López, R. Garcia-Tinedo, L.Liu, R. Docea, M. Kirchner and S. Bodenstedt, "StreamSense: Policy-driven Semantic Video Search in Streaming Systems" in Middleware Industrial Track ’24, December 02–06, 2024, Hong Kong, China.

NEARDATA

Software Outcomes

Component Name	Module	Type	Repository
Lithops	Analytics	Component	https://github.com/neardata-eu/lithops-hpc-singularity-and-k8s-backend
Dataplug	Data Plane	Component & Integration with Lithops (Native)	https://github.com/neardata-eu/dataplug
METASPACE	Data Sources & Data Plane	Component	https://github.com/metaspace2020/metaspace
METASPACE & Lithops	Data Sources, Data Plane & Analytics	Integration	https://github.com/metaspace2020/Lithops-METASPACE
Pravega	Data Plane	Component	https://github.com/pravega/pravega
SCONE	Data Plane	Component	https://github.com/neardata-eu/scone-artifacts
Glider	Data Plane	Specific NEARDATA Genomics Data Space Component	https://github.com/neardata-eu/glider-store

Components	Description	GitHub URL
Variant-Interactions use-case	MDR use-case source-code integrated with HPC Data Connector	https://github.com/neardata-eu/MPI-genomics-MDR
	MDR use-case with Apache Spark	https://github.com/neardata-eu/spark-mdr
Lithops backend	Lithops backend for singularity	https://github.com/neardata-eu/lithops-hpc-singularity-and-k8s-backend
Video-streaming benchmarks	Video-streaming benchmarks with Gstreamer and Pravega connectors	https://github.com/neardata-eu/video-streaming-benchmarks
Dataplug	Dataplug python framework	https://github.com/neardata-eu/dataplug
Genomics use-case	Variant calling source code	https://github.com/neardata-eu/serverless-genomics-variant-calling
	Glider	https://github.com/neardata-eu/glider-store
Surgery use-case	Federated Learning Source Code	https://github.com/neardata-eu/nct_tud_fl_demo
	Surgical Pravega GStreamer Demo	https://github.com/neardata-eu/nct_dell_demo_pravega
Transcriptomic Atlas use-case	Federated Learning for Human Genome Variation Analysis	https://github.com/SanoScience/neardata-fl-for-transciptomics
	Transcriptomic Atlas Pipeline	https://github.com/neardata-eu/transcriptomics-atlas-sano
Metabolomics use-case	METASPACE codebase including the implementation of ML-based metabolite identification (Experiment 1)	https://github.com/metaspace2020/metaspace
	ML models from Experiment 1	https://github.com/metaspace2020/metaspace/tree/master/metaspace/scoring-models

Communication and dissemination activities

Type of communication	Category of audience	Achievements
Scientific Publications	Scientific community	17 publications
Conferences and Workshops	Scientific community and Industry	44 events
Community Building	Scientific community, Industry	4 events
Meetings outside the consortium	Scientific community, Industry	23 meetings
Events for society	General Public	29 events
Press releases	General Public	6 publications

NEARDATA

Thank you

UNIVERSITAT ROVIRA i VIRGILI

BSC

TECHNISCHE UNIVERSITÄT DRESDEN

NCT

EMBL

DELL Technologies

KIO Networks

sano

SCONE

Department of Health

BDV BIG DATA VALUE ASSOCIATION

Funded by the European Union