GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

Universität Göttingen, GWDG

HPS

Julian Kunkel

Smarter Data Processing in the DECICE and KISSKI projects

Outline

1 About Us

2 KISSKI Project

3 DECICE Project

4 Conclusion

Prof. Dr. Julian Kunkel

- Deputy Head of GWDG - HPC
- Group Leader Computing Working Group
- Professor at the University of Göttingen
  - Institute for Computer Science
  - Research Group **High-Performance Storage**
- Research goals
  - Data-driven workflows
  - Storage & Parallel file systems
  - Performance analysis of parallel apps and I/O
  - Performance portability
  - Application of machine learning methods
  - Data reduction techniques
  - Management of cluster systems
  - Software engineering of scientific software

Research Group: High-Performance Storage

Research goals

- Data-centric Input/Output architectures
- Efficient execution of data-driven workflows
- Autonomous storage systems making intelligent decisions

Noteworthy involvement

- The Virtual Institute for I/O vi4io
- The Journal of High-Performance Storage J∞HPS https://jhps.vi4io.org/
- Development of IOR/MDTest/MDWorkbench
- The IO500 Benchmark IO500 https://io500.org

HPS https://hps.vi4io.org

The GWDG: service organization and data and IT service center

- works in conjunction with Universität Göttingen and Max Planck Society
- carries out independent research in the field of computer science
- provides support in preparing future professionals for a career in IT
- employs about 200 experts in 8 working groups

GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

https://www.gwdg.de

The Computing Working Group (AG-C)

Mission

AG-C provides scalable solutions for resource-intensive applications

Areas of responsibility

- Planning and operation of HPC systems
- Advice on use
- Hosting and Housing of HPC-Systemen
- Research on a wide range of HPC-related topics
- Supervision of project work and theses
- Training courses on the use of HPC systems

AG Computing

Prof. Dr. Julian Kunkel

Group Head, GF HPC

Dr. Christian Boehme

Deputy Group Leader

Dr. Stefanie Mühlhausen

Research, Teaching

Dr. Bernhard Bandow

CIDBN

Dr. Tim Ehlers

Administration, Security

Dr. Vanessa End

NHR, Doku, Web

Alexander Goldmann

Community Manager, PR

Biermann, Johannes

NHR, Digital Humanities

AG Computing

Hauke Gronenberg

ML/DL, ForestCare

Christoph Hottenroth

DLR

Hendrik Nolte

Data Lakes, Sec. Workflow

Dr. Nils Kanning

DLR, Administration

Ruben Kellner

Training

Azat Khuziyakhmetov

DLR, Administration

Dr. Christian Köhler

NHR, Physics

Sebastian Krey

NHR, Administration

AG Computing

Michael Langfermann

DLR, Administration

Tino Meisel

NHR, Data Science

Marcus Merz

NHR, Monitoring

Rosemarie Meuer

DLR

Dr. Martin Leandro Paleico

NHR, Bioinformatics

Pavan Kumar Siligam

NHR, Earth Sys. Science

Dorothea Sommer

ML/AI, ForestCare

Timon Vogt

NHR, Admin., Monitoring

Outline

1 About Us

2 KISSKI Project

3 DECICE Project

4 Conclusion

Project Motivation

Critical Infrastructures

- Critical and sensible infrastructures are (among others)
  - Energy
  - Medicine
- Requires special care due to
  - Privacy regulations
  - Security concerns
  - Real-time response requirements
- Data management is a key challenge

AI Service Center

KISSKI: AI Service Centre for Sensitive and Critical Infrastructures

- Goal: research how to establish a AI service center
  - Establish the service center after the funding period
- Funded with 17M EUR for 3 years (5 partners)
- Critical infrastructure requirements: security, privacy, reliability

Project partners

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN
IN PUBLICA COMMODA SEIT 1737

GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

aQua
ZUKUNFT DURCH QUALITÄT

CIDAS
Campus-Institut Data Science

Fraunhofer IEE

UNIVERSITÄTSMEDIZIN GÖTTINGEN : UMG

Leibniz Universität Hannover

Overview of the Structure

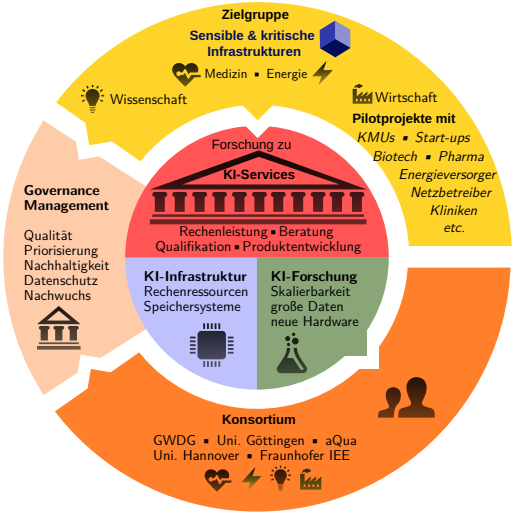

Project Motivation

The Long Way to AI

Getting into AI is **hard**

One needs

- **Data** to train
- **Hardware** to run on
- **Expertise** to optimize
- **Infrastructure** to serve

KISSKI provides it **all!**

Project Motivation

What We Provide Related to I/O

- **Data**
  - Trainings data
  - Pre-trained models
- **Expertise**
  - Training
  - Consulting
  - Support
- **Hardware**
  - High-Performance systems
  - GPUs, FPUs, Graphcore, ...
- **Infrastructure**
  - VAST storage
- We have a comprehensive service list (primarily: medicine, energy)

Project Motivation

What We Provide Related to I/O

- **Ensure Data Sovereignty**
  - Trainings and inference data is protected by GDPR
  - This generally prohibits the use of a shared infrastructure
    - which is a hard requirement due to cost concerns
  - Solution: Encrypt sensitive data - preventing admin access to data
- **Transparent Encryption**
  - The encryption has to be
    - transparent for the user and the application
    - highly scalable across multiple nodes
    - offer high performance in streaming, small file, and random IO
    - compatible with any (or most available) storage systems to prevent a lock-in

Project Motivation

What We Provide Related to I/O

- **Flat Namespaces**
  - Flat namespaces are a common problem for metadata servers
    - This workload is very common in ML projects
  - Too many targets per data item breaks a hierachical structure (like POSIX)
  - Required: Efficient way to select data based on semantic metadata
- **Provenance**
  - Derived artifacts might be indistinguishable by domain specific metadata
  - Data lineage is key for comprehensibility
  - In a hierachical organization this might be encoded in the path
  - Required: Find a novel way to find/present data based on provenance

Outline

1 About Us

2 KISSKI Project

3 DECICE Project

4 Conclusion

DECICE Project

Device-**E**dge-**C**loud **I**ntelligent **C**ollaboration fram**E**work

- EU Horizon project (ID 101092582) Start Nov. 2022 (6 M € financing)
- Topic ID: HORIZON-CL4-2022-DATA-01
- Development of a **Cognitive Cloud**
- AI-enabled computing continuum from Cloud to Edge
- Cooperation between
  - Universities
  - Data centers
  - Research institutes
  - Industry companies

DECICE Partners

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN
IN PUBLICA COMMODA SEIT 1737

GWDG
Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen

E4 COMPUTER ENGINEERING

KTH VETENSKAP OCH KONST

University of Stuttgart Germany

HUAWEI

FORSCHUNG Burgenland
RESEARCH & INNOVATION

SYNYO

MARMARA ÜNİVERSİTESİ

top-ix
WE DO CONNECTIONS

ALMA MATER STUDIORUM A.D. 1088

nag®

TRI

Motivation

- Cloud computing has become major factor
  - Improved productivity
  - Scalability via public clouds
- Similarly Edge computing is gaining relevance
  - Location awareness
  - Ultra-low latency
- Produce data that requires HPC for processing
  - On-premise
  - Specialized hardware

$\Rightarrow$ System and hardware landscapes are complex

Our Proposition

AI-Enabled Cloud to Edge Computing Continuum

- Handle complex landscapes in one system
  - APIs for users and admins
- Utilize AI automatically adapt to changes
  - Optimize load balancing for performance and energy efficiency for storage and compute
- AI learns from monitoring and simulations
  - Verify that AI performs better than naive solution

Applications are Complex

- Many applications consist of various components considered "independent"
  - Utilizes cloud, edge, and HPC compute and storage resources
- Governance: Involve multiple administrative domains

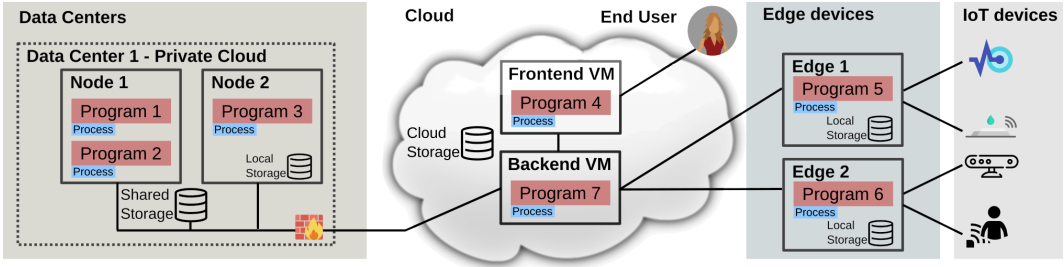

Goals

1. Development of a cloud management framework
   - Based on Open source
   - Allow application deployment across HPC, Cloud and Edge
2. Provide systems for administration/DevOps
   - Manage services, networks, compute resources, infrastructures
3. Maintain a digital twin of the compute continuum
   - Used for AI-forecasting and simulation
4. Automate decision making for application placement and optimization
   - Improve application performance, reliability, energy efficiency and cost

Activities Related to I/O

- Develop an I/O scheduler for efficient data management
  - for data allocation, data movement
  - considering storage characteristics and faults
  - utilizing AI technology
- Modelling and monitoring of system performance
- Integration of compute and data schedulers

Activities Related to I/O

Unify Usage of Cloud Storage

- HPC applications still often use POSIX filesystem
  - Lustre, BeeGFS, GPFS, etc.
- Cloud environments often use object storages
  - S3, Swift
- Exploring methods to unify both worlds:
  - Built up dedicated S3 storage layer for HPC applications
    - Explore opportunities particularly in small file io (SeaWeedFS)
  - Explore usability of joint storage systems (S3 and POSIX)
    - like Vast, or WekaFS
  - Explore options to mimic POSIX filesystem (e.g. s3fs)
    - Which consistency is really required?
    - Eventual consistency vs. strong consistency
    - How is the scalability across multiple nodes

Outline

1 About Us

2 KISSKI Project

3 DECICE Project

4 Conclusion

Conclusion

- Opportunities can be exploited by utilizing machine learning
- KISSKI researches an AI service center for critical infrastructures
- DECICE creates an I/O scheduler for the compute continuum
- We are happy to collaborate in any of the topics

The Decice project received funding from the European Union's Horizon 2022 research and innovation programme under grant agreement No **101092582**.

Disclaimer: This material reflects only the author's view and the EU-Commission is not responsible for any use that may be made of the information it contains.