mdx: A platform for the data-driven future

Toyotaro Suzumura
Information Technology Center
The University of Tokyo

Motivation

mdx

- **Accelerating inter-disciplinary research and collaborations**
  - Data science has become a key driver to advance science and technology, but more collaborations among various domain experts in both academia and private sectors should be formulated to solve complex problems in real-world society (e.g. material discovery, carbon-neutral society).
- **How can we accomplish this ?**
  - Agility : Need a national-wide cloud platform (IaaS and PaaS/SaaS) that enables researchers to promptly set up their own data analytics environments by allowing them to install/configure their required analytics softwares, and publish data repository services as needed
  - Connecting with Edge Devices (e.g. Electron Microscopes, IoT sensors) via High-Speed Network:
  - Coupling with HPC environments :
    - Some applications need big computing powers for such applications as large-scale simulations and large-scale AI/ML training thus the platform should be coupled with supercomputers when needed.
- **Why can't we rely sololy on supercompters or public clouds?**
  - These requirements can not be fulfilled by current types of supercomputers mainly targeting big science, and also public clouds from private sectors

Status in the U.S.

- Ian Foster, Daniel Lopresti, Bill Gropp, Mark D. Hill, Katie Schuman. "***A National Discovery Cloud: Preparing the US for Global Competitiveness in the New Era of 21st Century Digital Transformation***"
  - 2021 Apr https://arxiv.org/abs/2104.06953
  - The need for national discovery cloud:
    - *DOE and NSF supercomputers provide access to powerful simulation capabilities, but with access limited to small communities.*
    - *With a few notable exceptions, AI-ready datasets for research use are lacking. Commercial clouds are accessible to anyone with a credit card, but there is little of the coordination needed to create nationally useful discovery cloud services.*

mdx: A Platform for the Data-Driven Future

- Target is to leverage data utilization at all over Japan making full use of high performance R&E network "SINET" an R&E network of Japan operated by NII (National Institute of Informatics)
- Project supported by the Japanese government
- Currently jointly being operated by:
  - 9 National Universities (Tokyo, Hokkaido, Tohoku, Tsukuba, Tokyo Tech, Nagoya, Kyoto, Osaka, Kyushu)
  - NII (National Institute of Informatics)
  - AIST (National Institute of Advanced Industrial Science and Technology)
- Invite universities and public research institutes of all over Japan to use the platform for industry-academia and local government-academia collaboration activities.
- Production-level operation has been started since March 2021

mdx as On-demand Collaborative Platform

Conduct data analytics from various locations via SINET or Mobile SINET

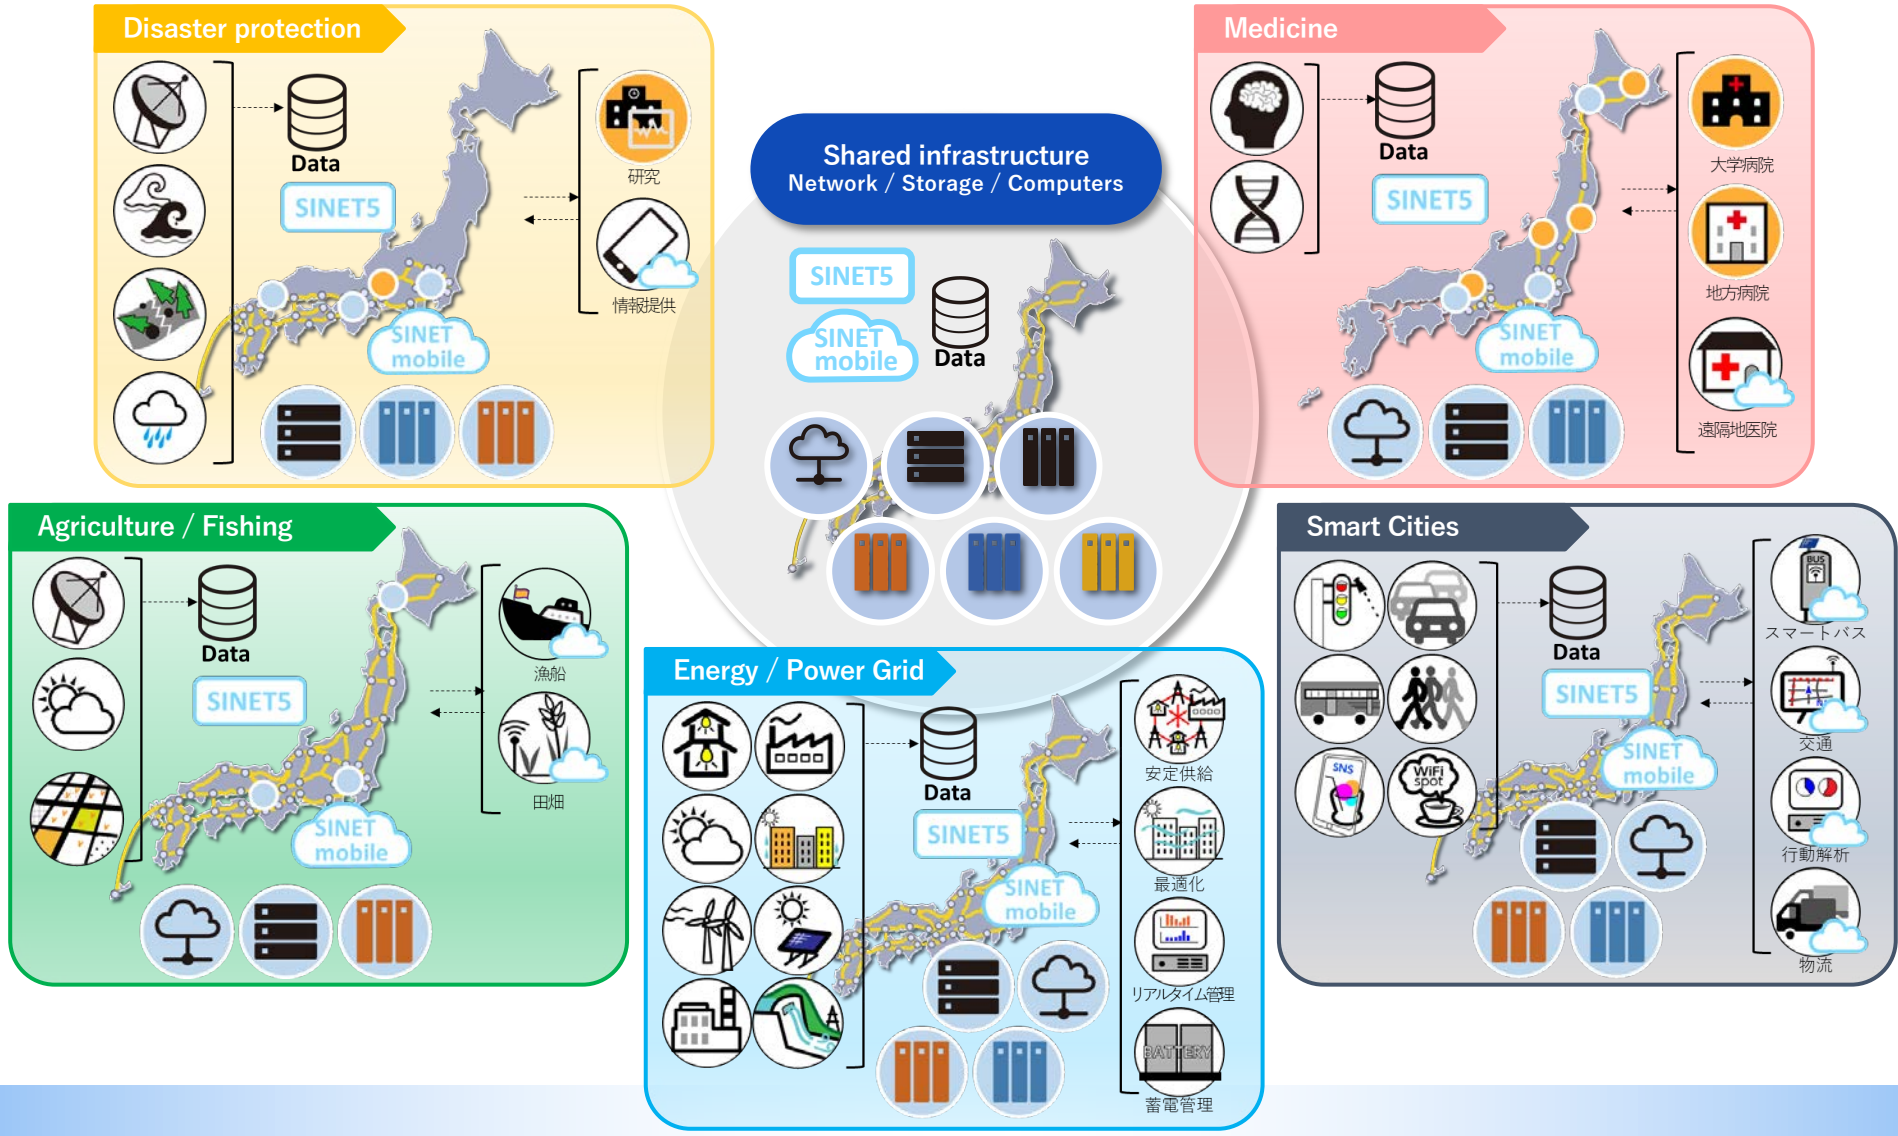

mdx : Providing Secure and On-demand Virtual Platform
(Optionally) with access to Supercomputers and Edge Devices

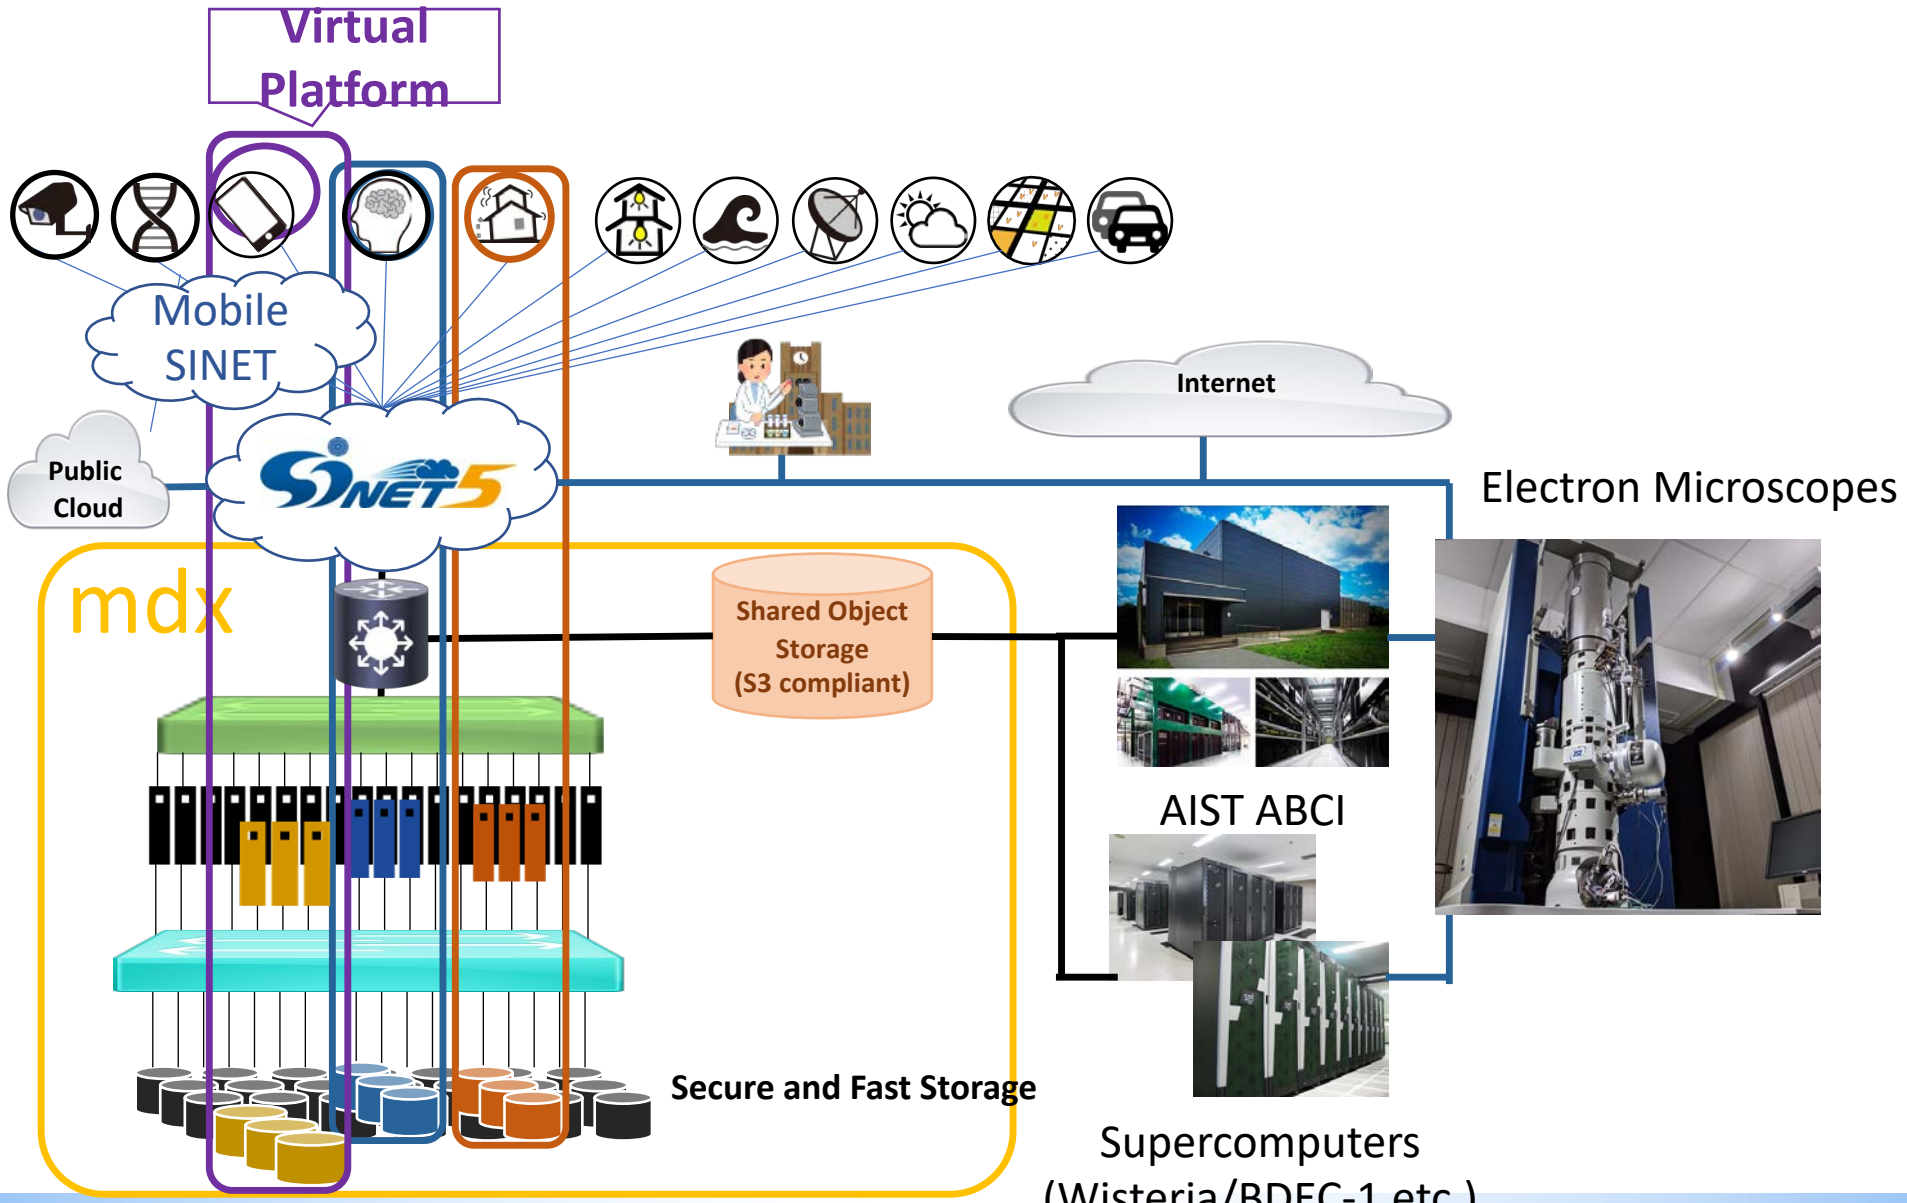

mdx: A Platform for the Data-Driven Future

- Agility : Provide **a rapid PoC environment** to accelerate R&D data utilization as well as industry-academia collaboration projects.
  - Sharing: Shared platform for various data utilization activities
  - Network: Combine a high-performance wide-area academic network called “SINET” with high performance computing and storage infrastructure
- Seamless Integration with Edge Devices:
  - Users can use wide bandwidth low latency “slices”
  - Wide-area virtual infrastructure isolated from “the internet”
  - Connect edge devices with high performance computing and storage infrastructure and supports real-time data processing
- Matching: Will provide matching function of:
  - Data providers: various data and their owners, and
  - Data scientists and researchers who have skills/tools to analyze data

Hardware Overview

mdx

- **Facility**
  - < **2.0 MW** including Cooling, <170 m^2
  - Same location with Wisteria/BDEC-01, same campus with AIST's ABCI supercomputer
- **Compute nodes (CPU)**
  - 368 nodes : Each node : Intel Xeon (IceLake-SP, 38 cores) x 2 CPU sockets/node
  - 2.1 Peta flops (double precision)
  - Total memory bandwidth : 150 TB/sec
- **Compute nodes (GPU)**
  - 40 nodes, Intel Xeon (IceLake-SP) x2 socket
  - NVIDIA A100 x 8 GPUs/node
  - 6.4 Peta Flops (FP64), 6.7 PF (FP32), 100 PF (FP16),
  - Total memory bandwidth: 496 TB/sec
- **Storage**
  - Fast Storage with NVMe SSD: 1.0 PB, 250 GB/sec
  - Large Storage with HDD: 16.3 PB. 157 GB/sec
  - Shared Object Storage (S3): 10 PB, 63 GB/sec
- **File System** : Lustre

- **Network**
  - Frontend (Juniper) : 25 Gbps Ethernet
    - 100G to SINET
      - 400G to Wisteria/BDEC-01
  - Storage, RDMA (Mellanox/NVIDIA) : 100G Ethernet with RoCEv2
  - Overlay with EVPN-VXLAN
- **Software, etc.**
  - VM & Container (VMware vSphere)
  - IaaS like management
  - High security, high availability

New building in Kashiwa II Campus

Resource Management and Scheduling

- Each institution is responsible for approving project applications and their required hardware resources

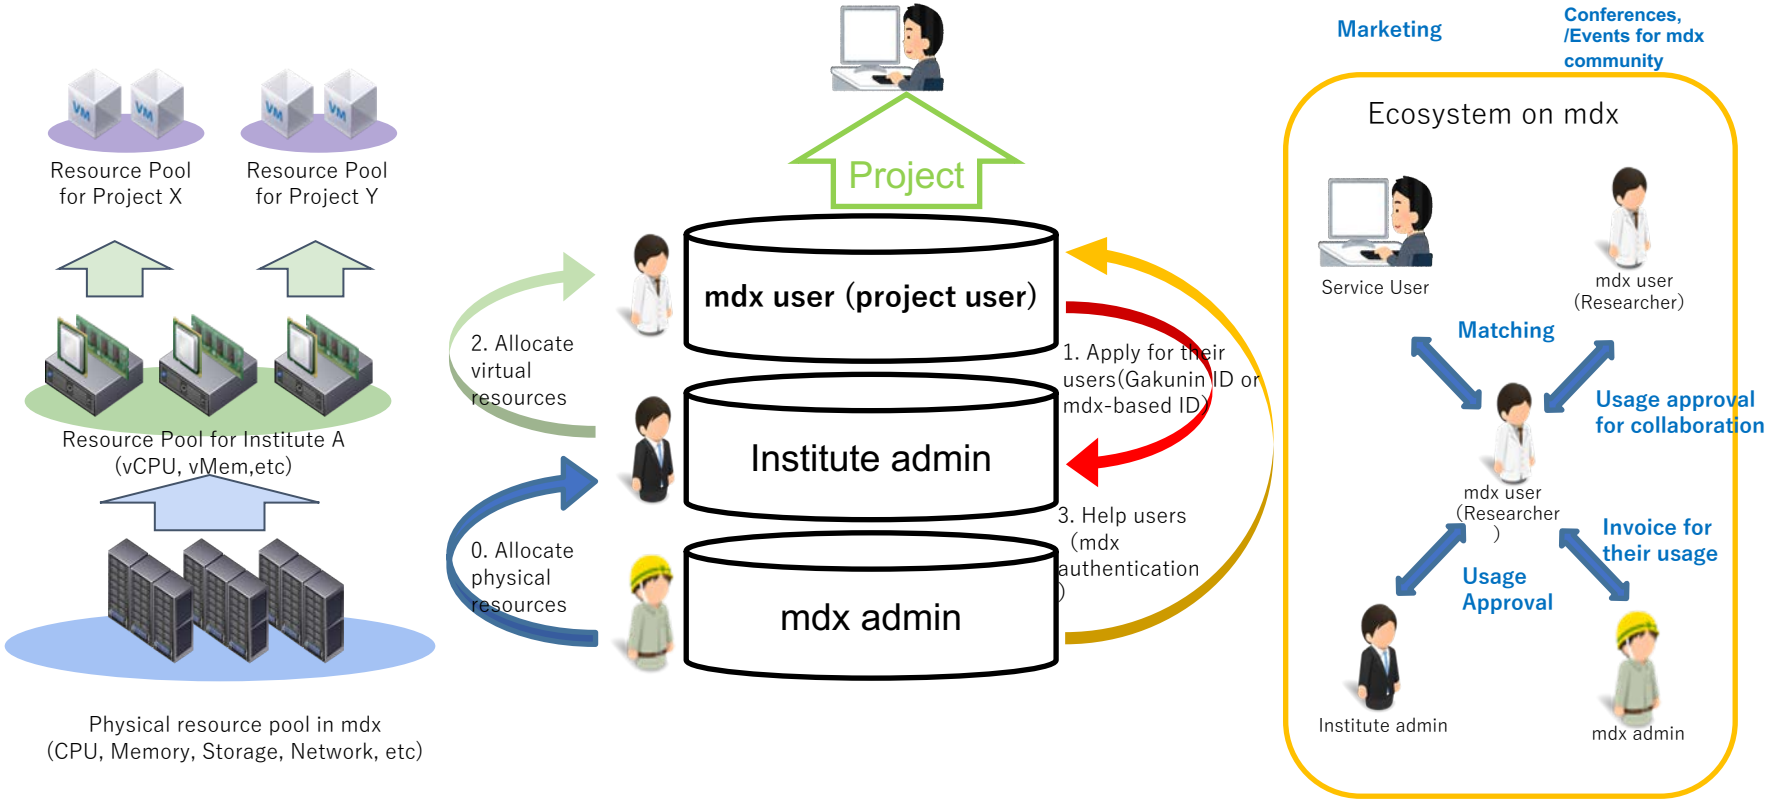

Data Community: Matching Data Providers and Data Scientists

- **Data Community on mdx**
  - As of Jan 2022, we are currently in the midst of designing a PaaS level platform that establishes a community between data providers and data scientists
  - The alpha version: January 2023 (Plan)
- **Requirements**
  - **Data providers** can easily upload their data to ”mdx” by specifying the spec of data and data usage conditions (e.g. only for research purpose)
  - **Data scientists** can easily find data based on their interests, and launch Jupyterlab

Data Providers

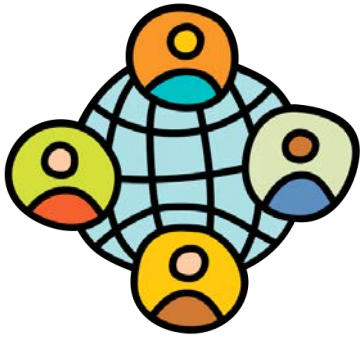

Data Scientists

Example) Kaggle-like community

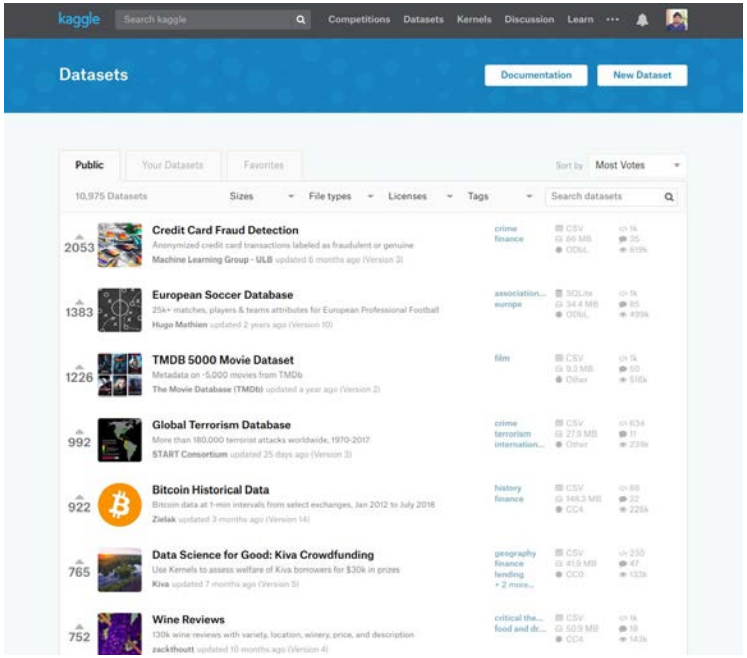

Usecase: Material Science and Engineering

- Key Challenges
  1. **Big data** from laboratory instruments
     - Image data from electron microscope
     - up to 1TB / experiment
  2. **High computational-power requirement**
     - Physics simulations (e.g., First-principle calculation)
     - Machine learning / data analysis
  3. **Flexible & secure data management**
     - Share confidential data among co-researchers
     - Publish open data on archives (e.g., ARIM Japan https://www.nanonet.go.jp/pages/arim/index.html)

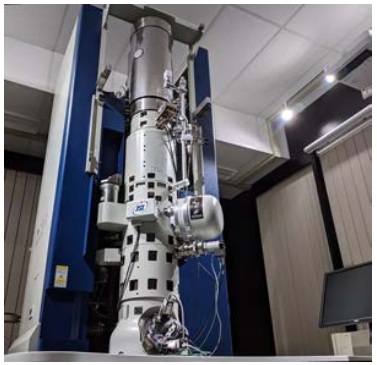

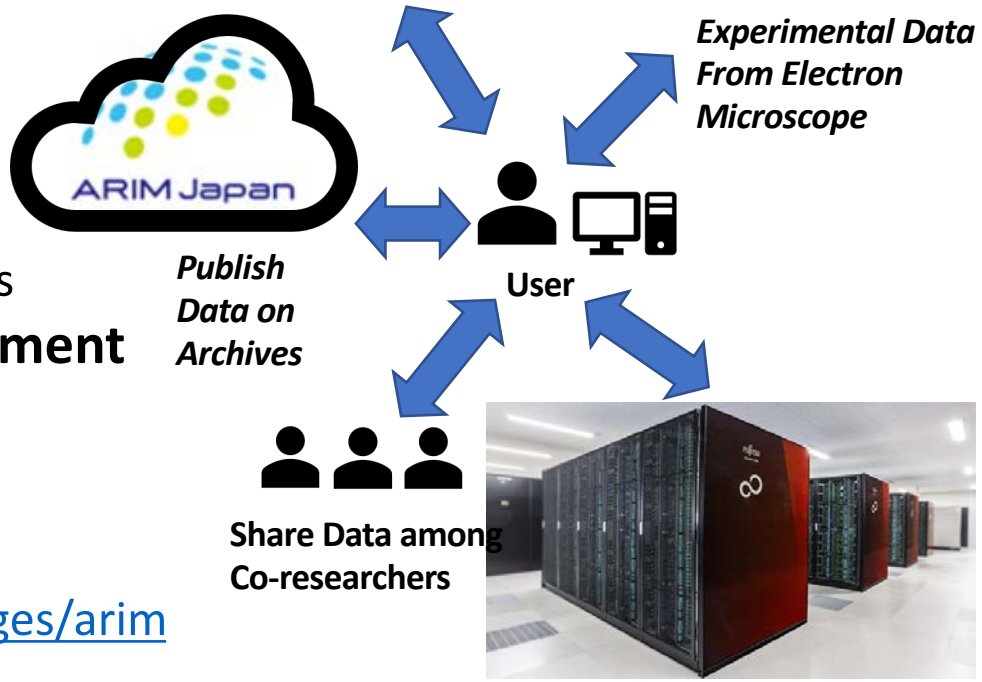

Usecase: Material Science and Engineering

Traditional ways

- Machine learning / data analysis on local servers
- Get data from laboratory instruments via portable storage
- Access supercomputers for First-principle calculation
- Send data to co-researchers
- etc.

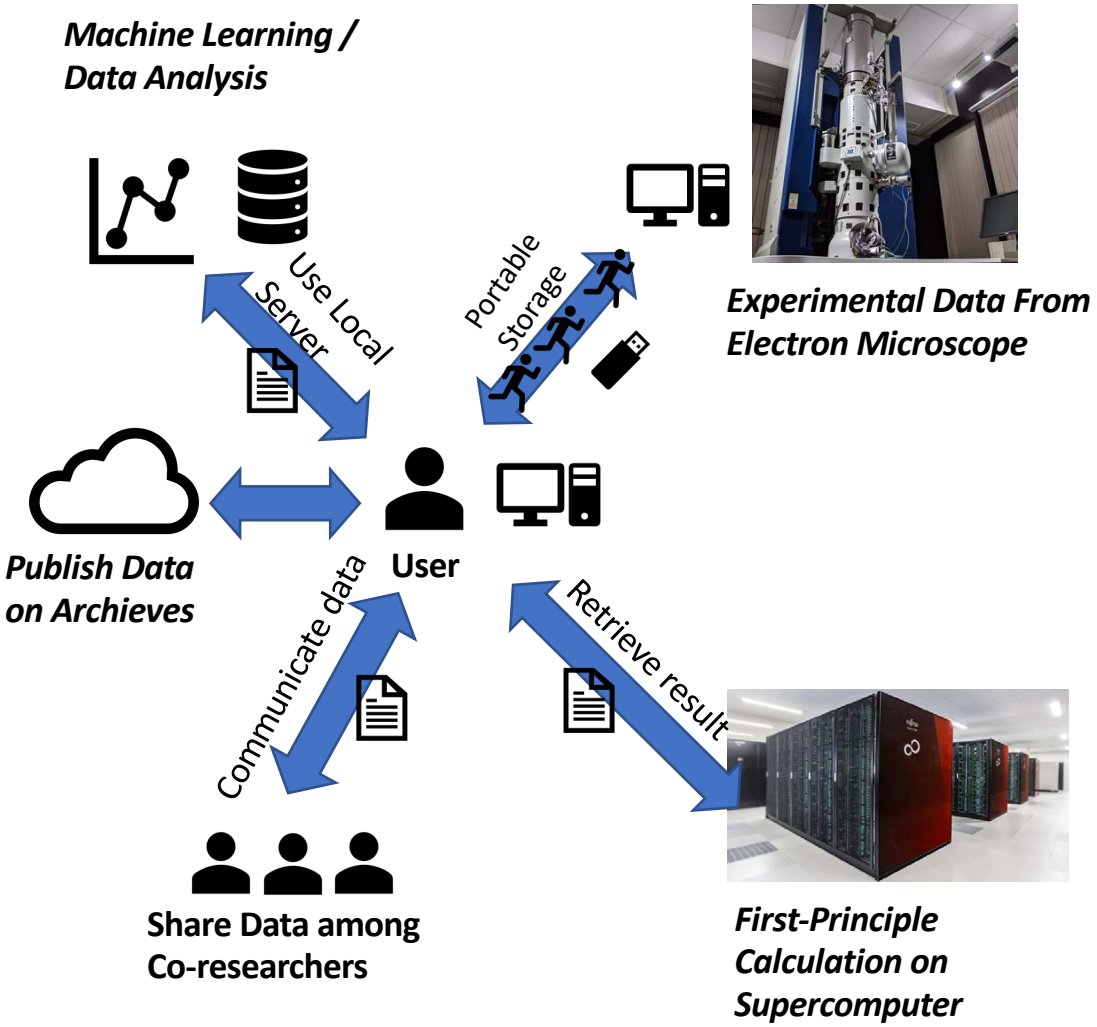

Usecase: Material Science and Engineering

System Integration with mdx:

- Install **mdx** as the extension of **"local sever"** with high performance and large storage
  - Customizable VM environment (GPU-available)
  - Store all data in the large storage on mdx (Lustre)
- **Secure & High-performance Inter-University Network** (SINET)
  - Data from laboratory instrument are via SINET (or intra-university network)
  - Seamless workload extension to academic supercomputers
    - Machine learning (GPUs on mdx) ⇔ First-principle calculation (Supercomputer)
- **Publish data via Amazon S3 protocol**

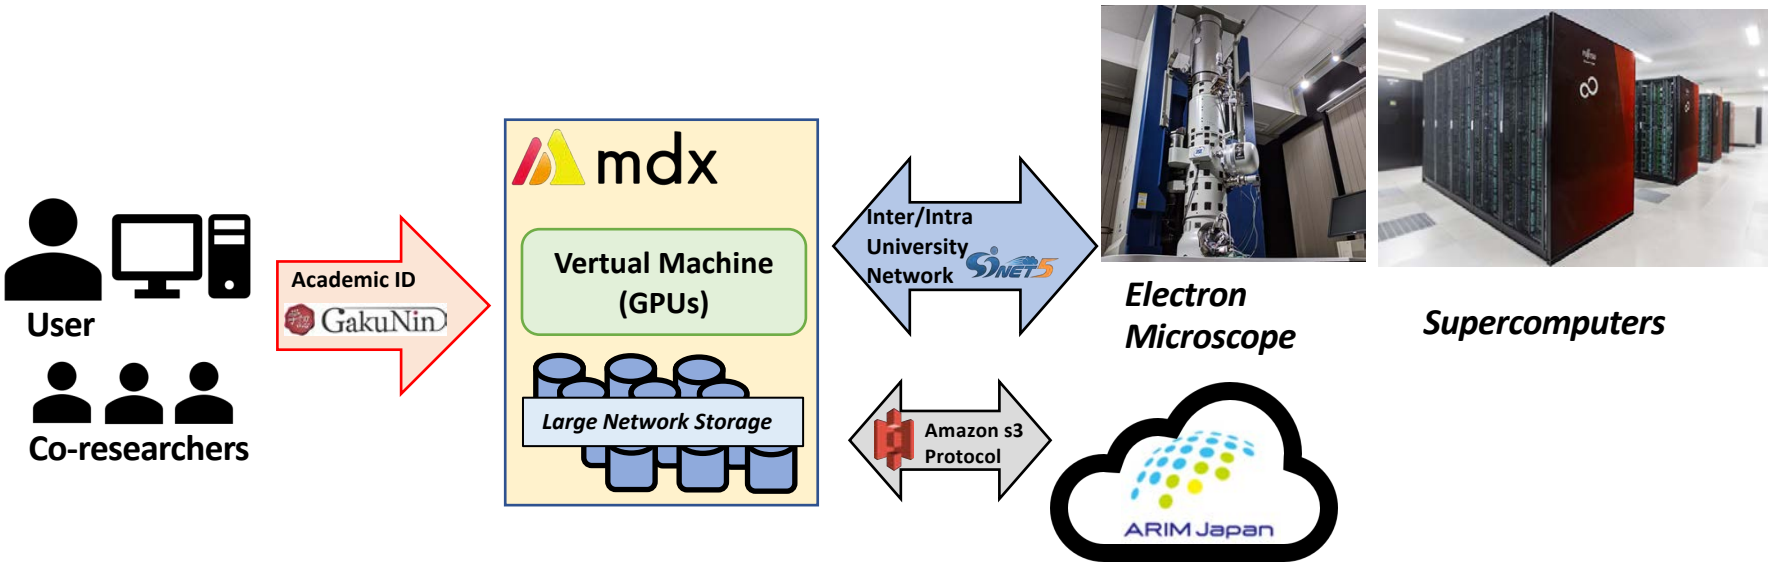

Comparison to Other Services

	OS-level Management	Customizable Resources	High Computational Power	Large-Scale Accessible Storage	Secure & High-End Academic Network	Open Data Publication
mdx	✓	✓	✓	✓	✓	✓
Local Lab PC/Server	✓				✓	
Supercomputer			✓		✓	
Enterprise Cloud	✓	✓		✓		✓

mdxは、高性能な計算機と大容量のストレージを備え、国立情報学研究所が運用する学術情報ネットワークSINET5（2022年度から次期システムに更新予定）と連携することで、広域からのデータ収集機能と、データ集積・処理機能を、企業や自治体との共同研究も含めた全国の大学・公的研究機関が関与する様々なデータ活用の取組に提供し、さらにはデータ活用のコミュニティーを形成して分野・セクタを横断した連携を触媒するハブとなることを目指します。

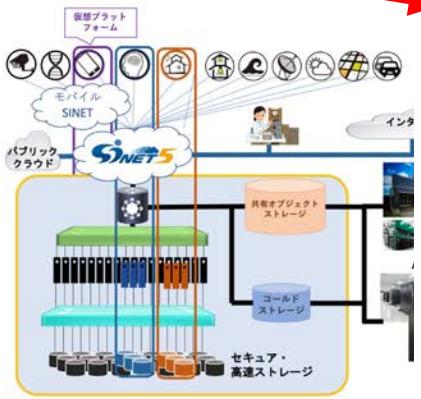

mdxについて

2021年に稼働を予定している、データ利活用、データ科学、に重点を置いた計算基盤を紹介します。

運営組織

mdxは、データ活用社会創成プラットフォーム共同研究基盤（共同研究基盤）の構成機関で運用されています。

問い合わ

お問合せは、郵便または電子メールにてお願いいたします。

https://mdx.jp/

of Launched Projects on mdx since October 2021

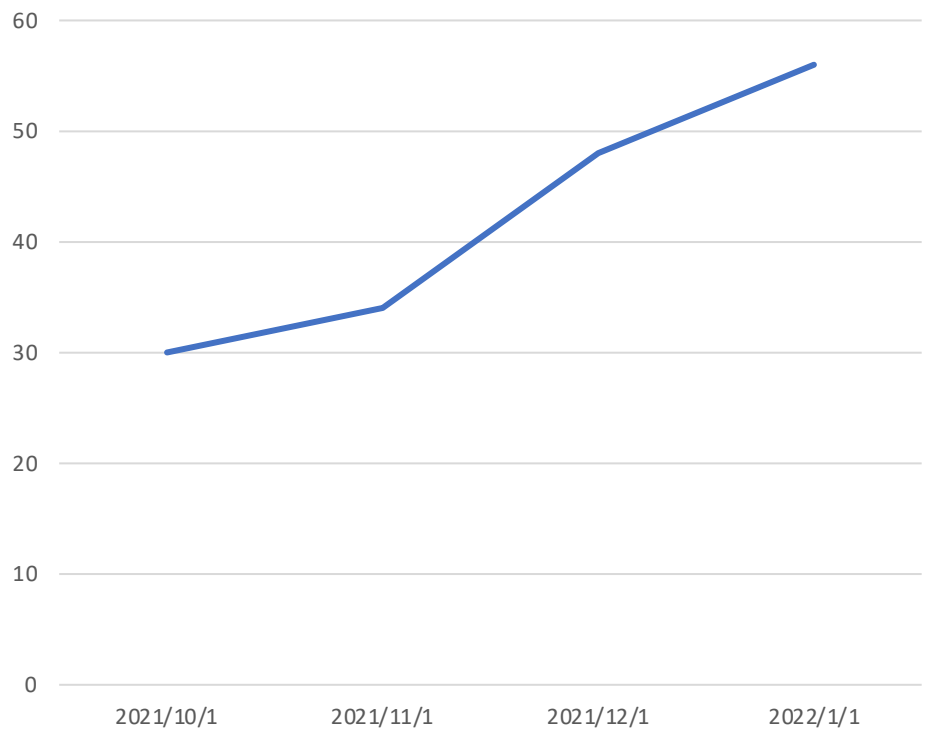

Word Cloud from Project Names

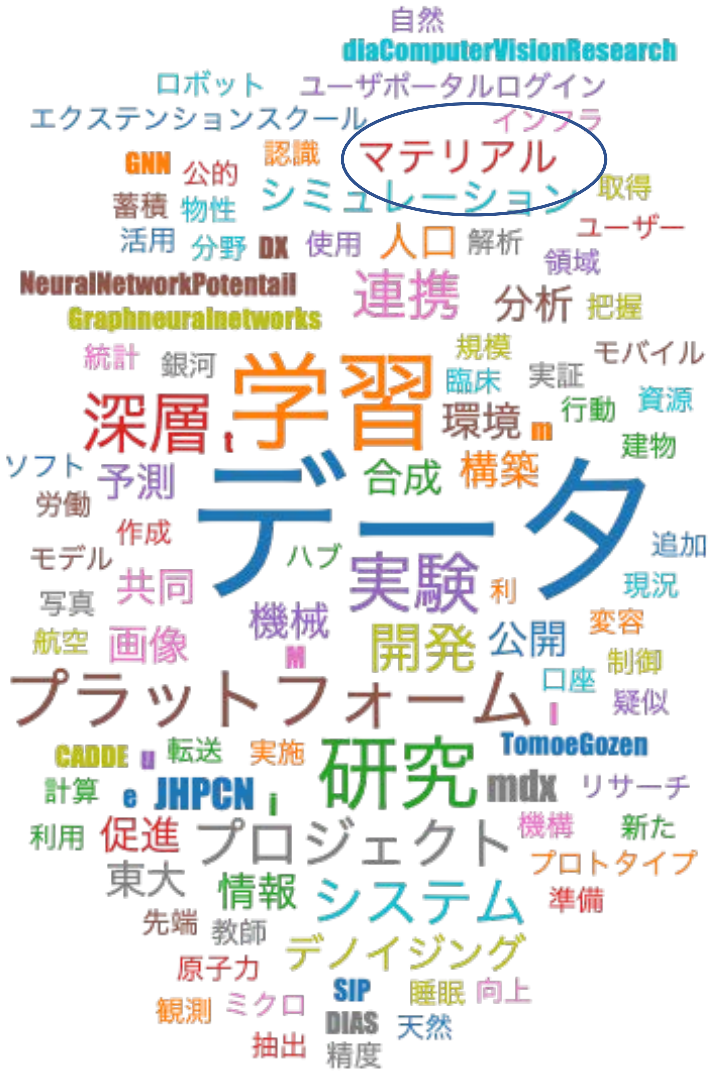

JHPCN

Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures

学際大規模情報基盤共同利用・共同研究拠点

Search English (en) Go

About JHPCN | 8 Centers | Themes | Pick up | User's Page | Committees

The Joint Usage/Research Center for Interdisciplinary Large-sc
Infrastructures is made up of eight centers equipped with supe
centers are the Information Initiative Center of Hokkaido Univer
Center of Tohoku University, the Information Technology Center
Tokyo, the Global Scientific Information and Computing Center
Technology, the Information Technology Center of Nagoya Univ
Center for Computing and Media Studies of Kyoto University, th
Osaka University, and the Research Institute for Information Tec
University. This is a network-type joint usage and collaborative r
core institution is the Information Technology Center of the Uni
Center began as a program of the Japanese Ministry of Educat
Science & Technology (MEXT) in April 2010.

The formal title of this program is "Joint Usage/Research Cente
Large-scale Information Infrastructures," but it is also known as
Performance Computing and Networking plus Large-scale Data
Information Systems (JHPCN)". This alternative name indicates
carried out in four areas:
very large-scale numerical computation, very large-scale data pr
capacity network technology, and very large-scale information s

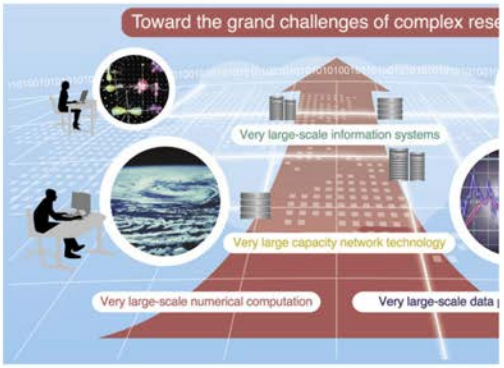

The objective of the program is to promote the continuous pro
research and to lay the groundwork for interdisciplinary collabo

Pamphlet

共同研究申し込み JOINT RESEARCH

シンポジウム SYMPOSIUM

お問い合わせ Contact Us

2022Project Dates

2021/11/26(Fri)	公募案内開始
2021/12/9(Thu)	課題応募受付開始
2022/1/6(Thu)	課題応募受付締切
Mid 2022/3	採否結果通知
2022/4/1(Fri)	共同研究開始
2022年7月	JHPCN第14回シンポジウム
2023/3/31(Fri)	共同研究期間終了
2023年7月	JHPCN第15回シンポジウム

2021Project Dates

2020/11/16(Mon)	公募案内開始
2020/12/10(Thu)	課題応募受付開始
2021/1/6(Wed)	課題応募受付締切【厳守】
2021/1/15(Fri)	紙媒体の課題申込書提出期日
Mid March	採否結果通知
2021/4/1(Thu)	共同研究開始
2021年7月8日・9日	JHPCN第13回シンポジウム 共同研究内容紹介
2021/12/9(Thu)	課題応募受付開始
2022/3/31(Thu)	共同研究期間終了
Mid 2022/7	JHPCN第14回シンポジウム

https://mdx.jp/