4CeeD: Real-Time Operating Infrastructure for Capturing, Curating, Correlating, Coordinating and Distributing Materials-related Data

Klara Nahrstedt (klara@illinois.edu)
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

Background: Increasingly data-driven and interdisciplinary scientific research

- *Key enabling factor*: Network connected scientific instruments capable of real-time data capture

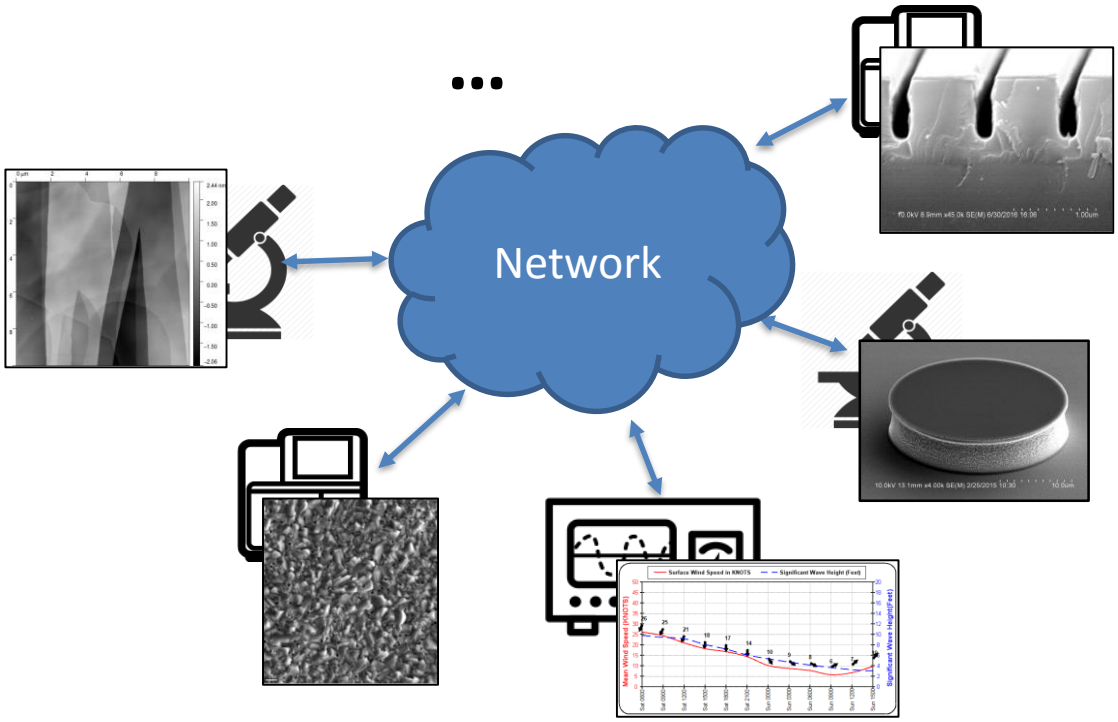

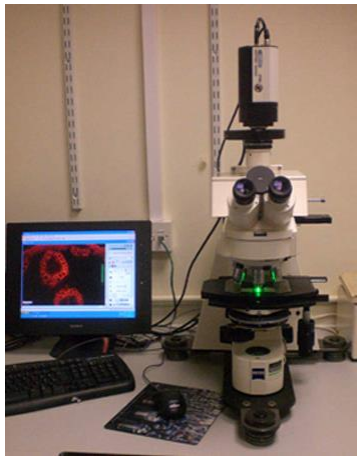

Digital microscope

Current situation in campus cyberinfrastructure

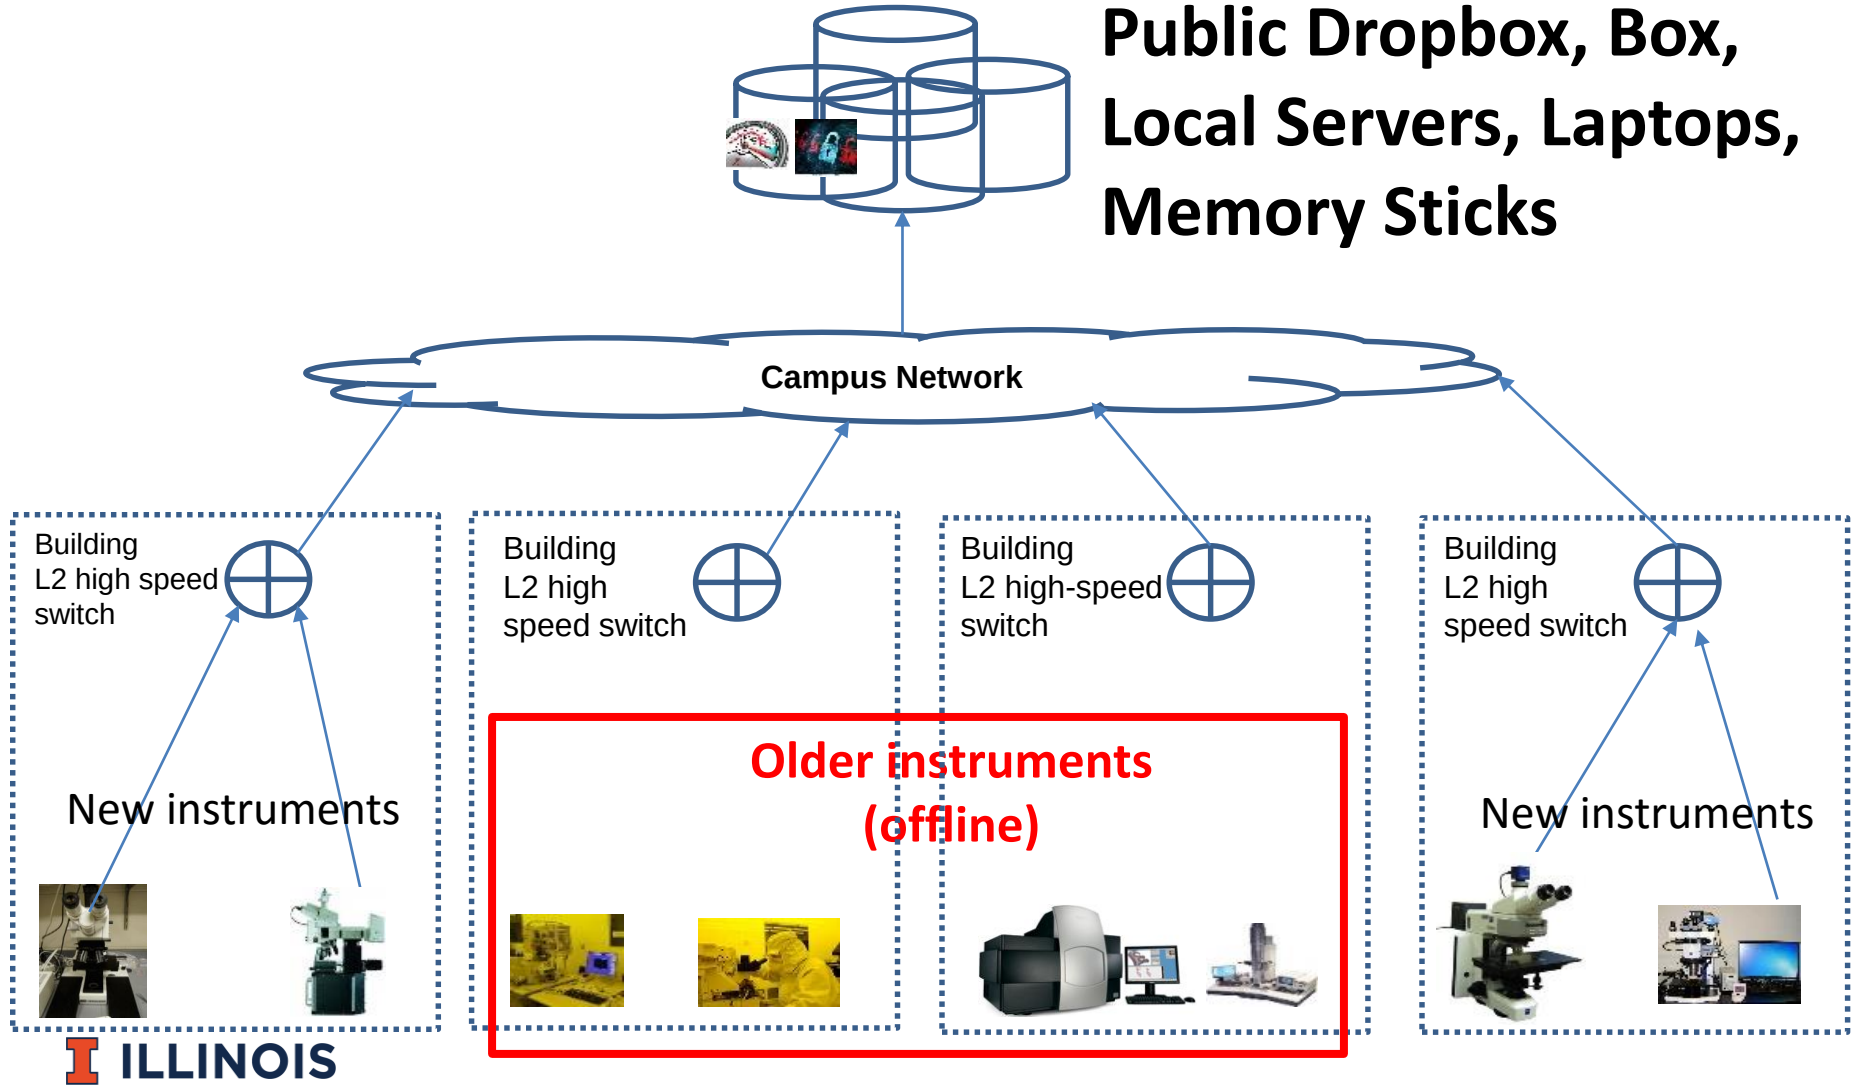

Example: Typical experimental process in material science research


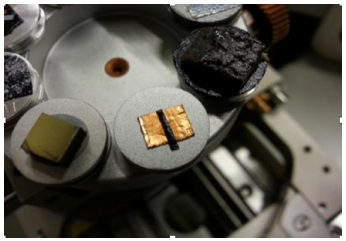

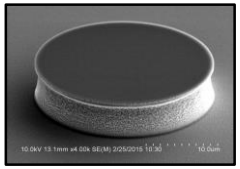

1. Create physical experimental sample

2. Prepare analytical sample from the physical experimental sample

3. Analyze the analytical sample

4. Convert digital analytical data

5.Transport the converted analytical data

6. Interpret the converted analytical data

7. Modify the procedure for creation of the physical experimental sample

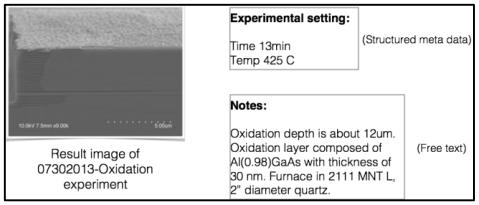


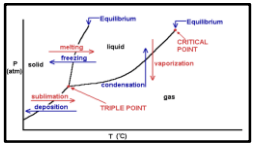


It typically takes 20 years to go from the discovery of new materials to fabrication of new and next-generation devices*

*Source: National Science and Technology Council's report, 2011

Motivation: Needs for advanced cyber-infrastructure for long-tail scientific data

- Related efforts mainly focus on ***homogenous, well-organized data*** in an offline or batch manner

- Much less effort has been on ***long-tail scientific data***:
  - Small/medium sized data sets collected during day-to-day research
  - "Dark data", e.g., unpublished data of failed experiments

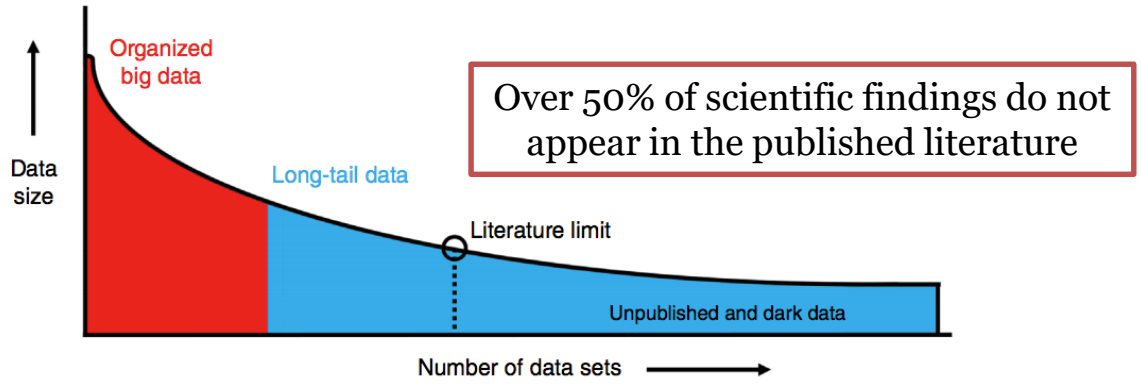

Long-tail scientific data

Source: Ferguson et al. *Nature neuroscience* 17.11 (2014)

Challenges in General

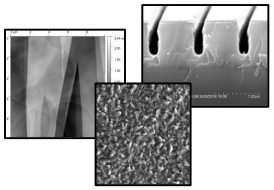

- Heterogeneous scientific data management and processing

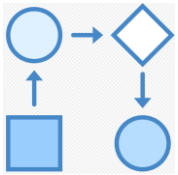

- Support ad hoc and complex data analysis workflows

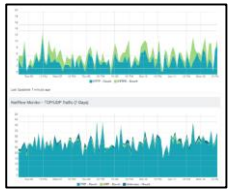

- Shorten time from digital capture to interpretation & insights

- Real-time data capture and acquisition

- Analytics support to gain insights from data

Challenges on Campuses

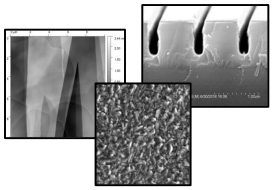

- Very diverse scientific instruments in Materials Research Lab (MRL), Micro-Nano-Technology Lab (MNTL), other labs

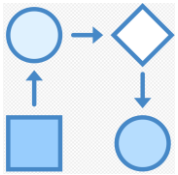

- Support very different user groups that collect and analyze data

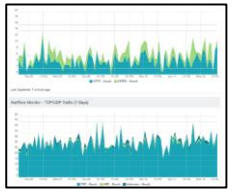

- Relations between students, faculty staff and academic cycles are different than in industry, impacting how insights are gained

- Rules on campuses regarding secure access to data and metadata in scientific labs vary

- Analytics tools support in scientific campus labs vary

Our approach

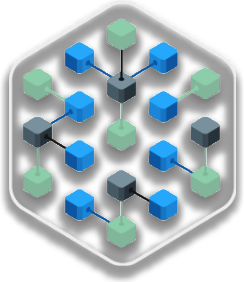

✓ Micro-service private cloud execution environment for instrument data curation and coordination (4CeeD)

✓ Data acquisition from aging instruments (BRACELET)

Long-tail scientific data processing challenges

- ***Challenges***: Support execution of heterogeneous types of data processing & analysis workflows

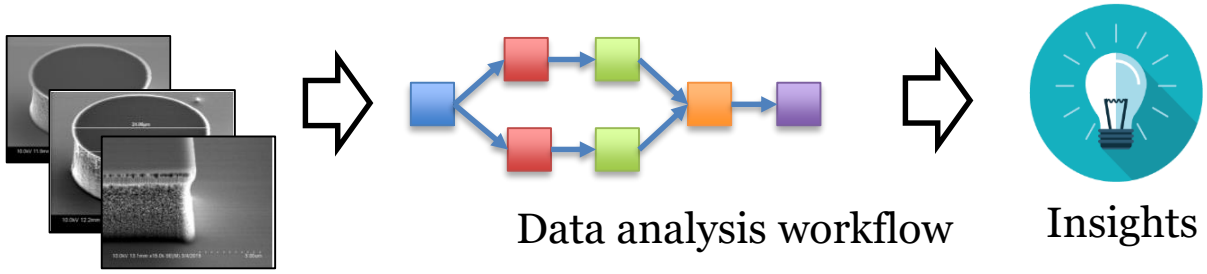

- Previous work often employs a monolithic approach in workflow implementation and execution
  - E.g.: Pegasus, Taverna, Kepler, etc.
  - Run on large-scale & homogeneous datasets

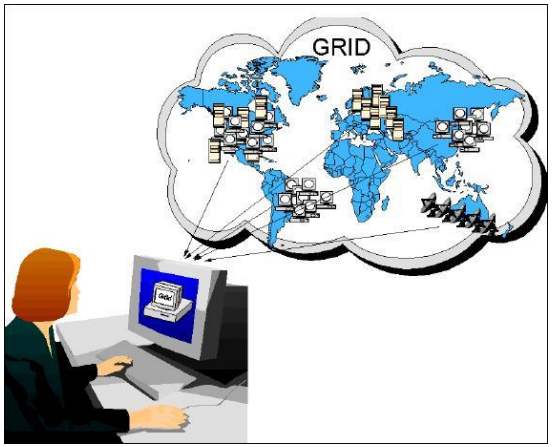

Executing workflows on grid infrastructure

Our Goal for Campus Cyberinfrastructure regarding New Scientific Instruments (1)

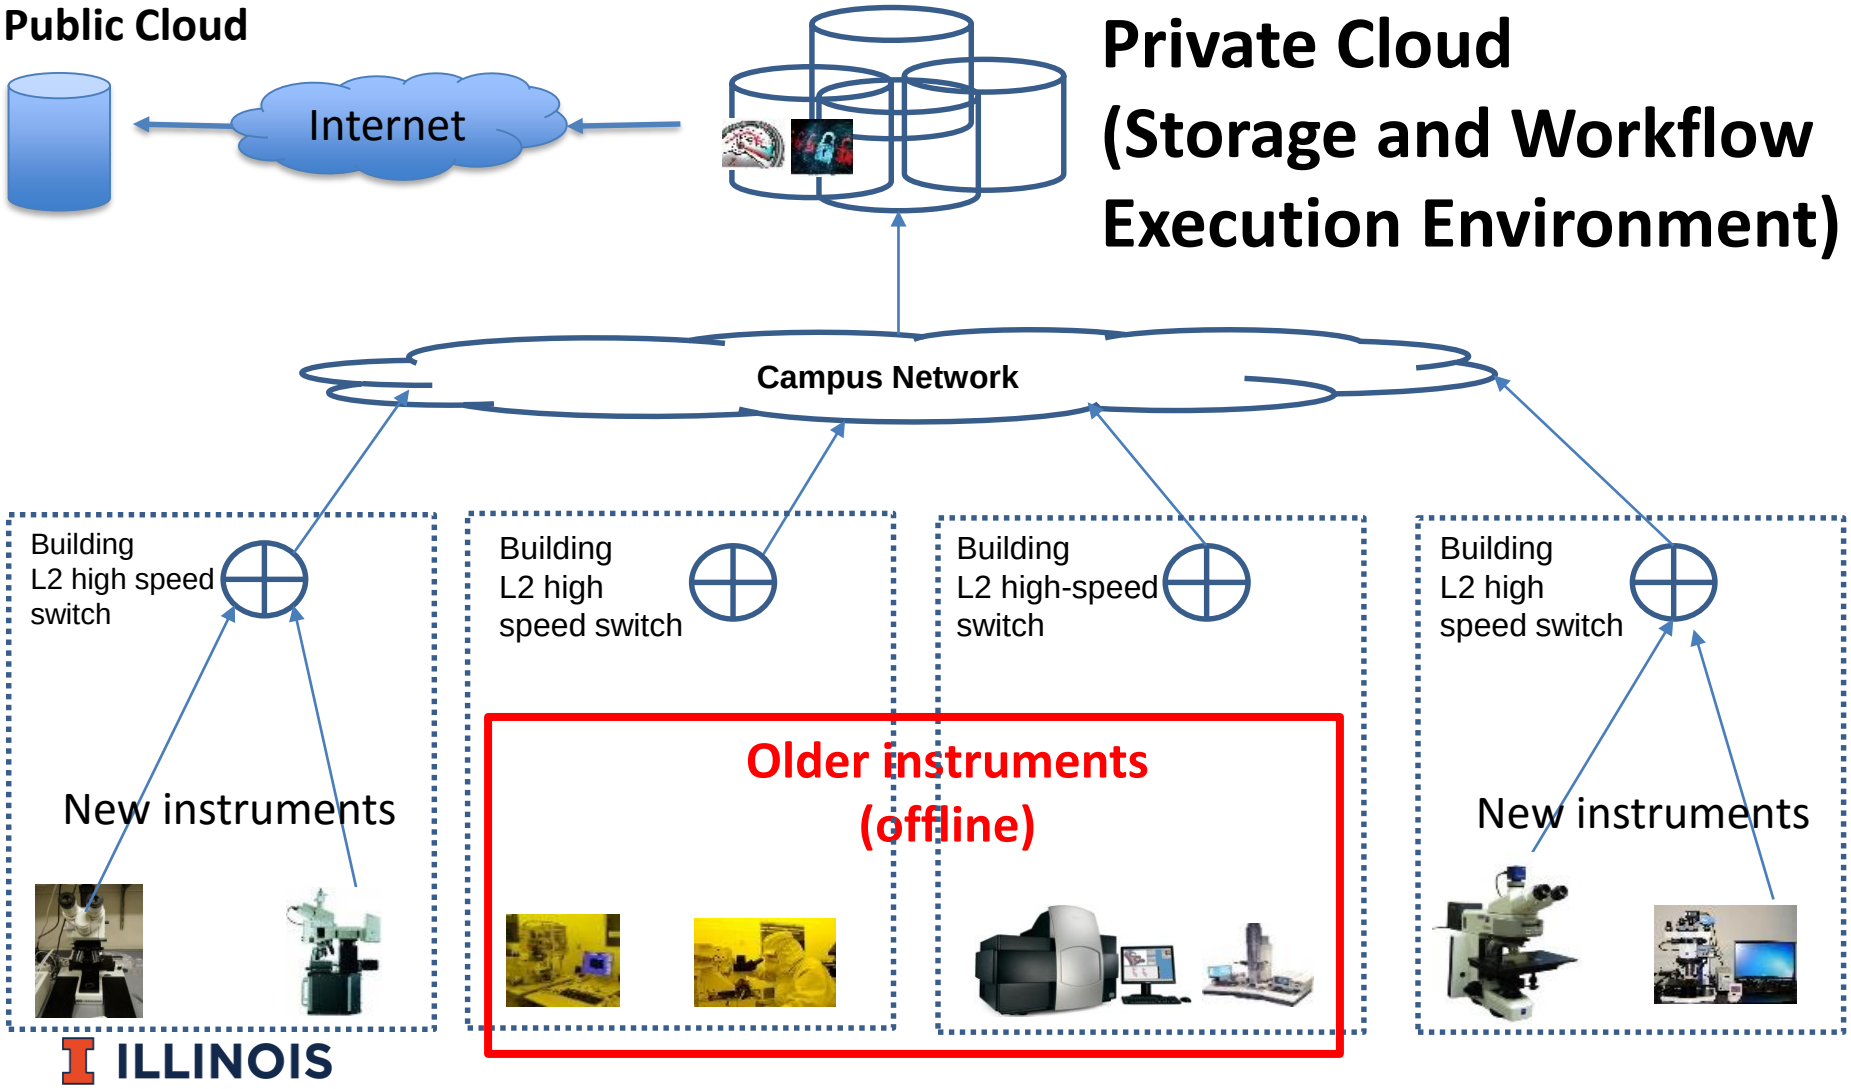

Our Approach: Micro-service execution environment in Private Cloud

- ***Micro-services over monoliths***: Each task is modeled as a micro-service
  - Use publish-subscribe middleware to connect between micro-services

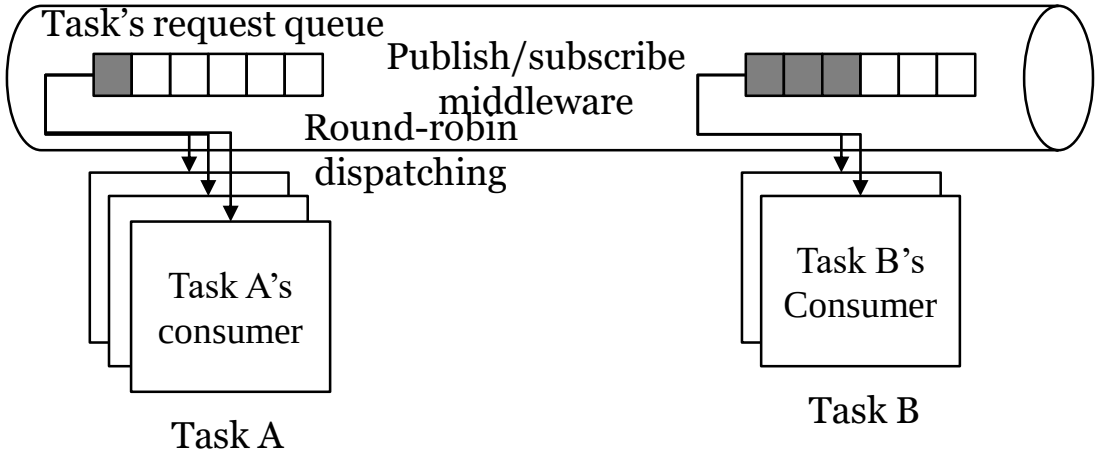

- Separate task dependencies from task implementation & deployment
  - Enable flexible workflow composition
  - Task-level resource provisioning

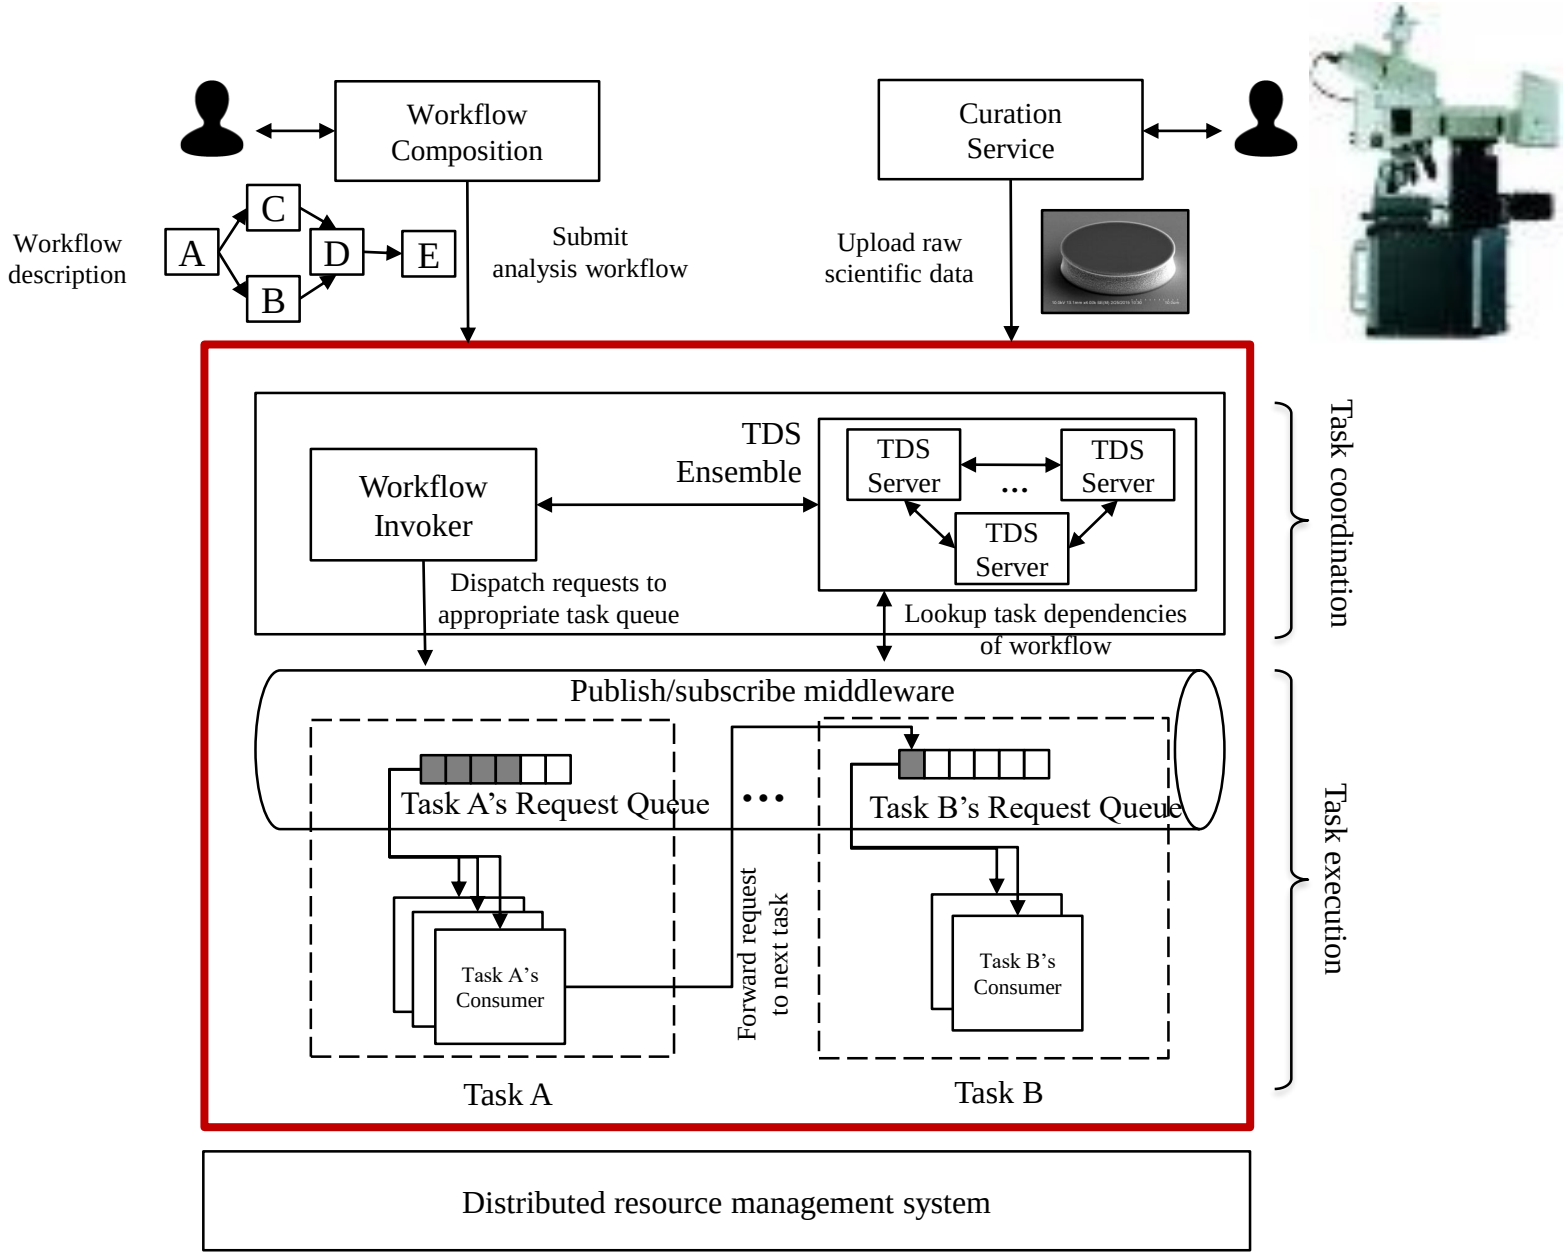

Phuong Nguyen *et al.*, *"Resource Management for Elastic Publish Subscribe Systems: A Performance Modeling-based Approach"* (IEEE CLOUD 2016)

Example: Executing scientific data processing workflow

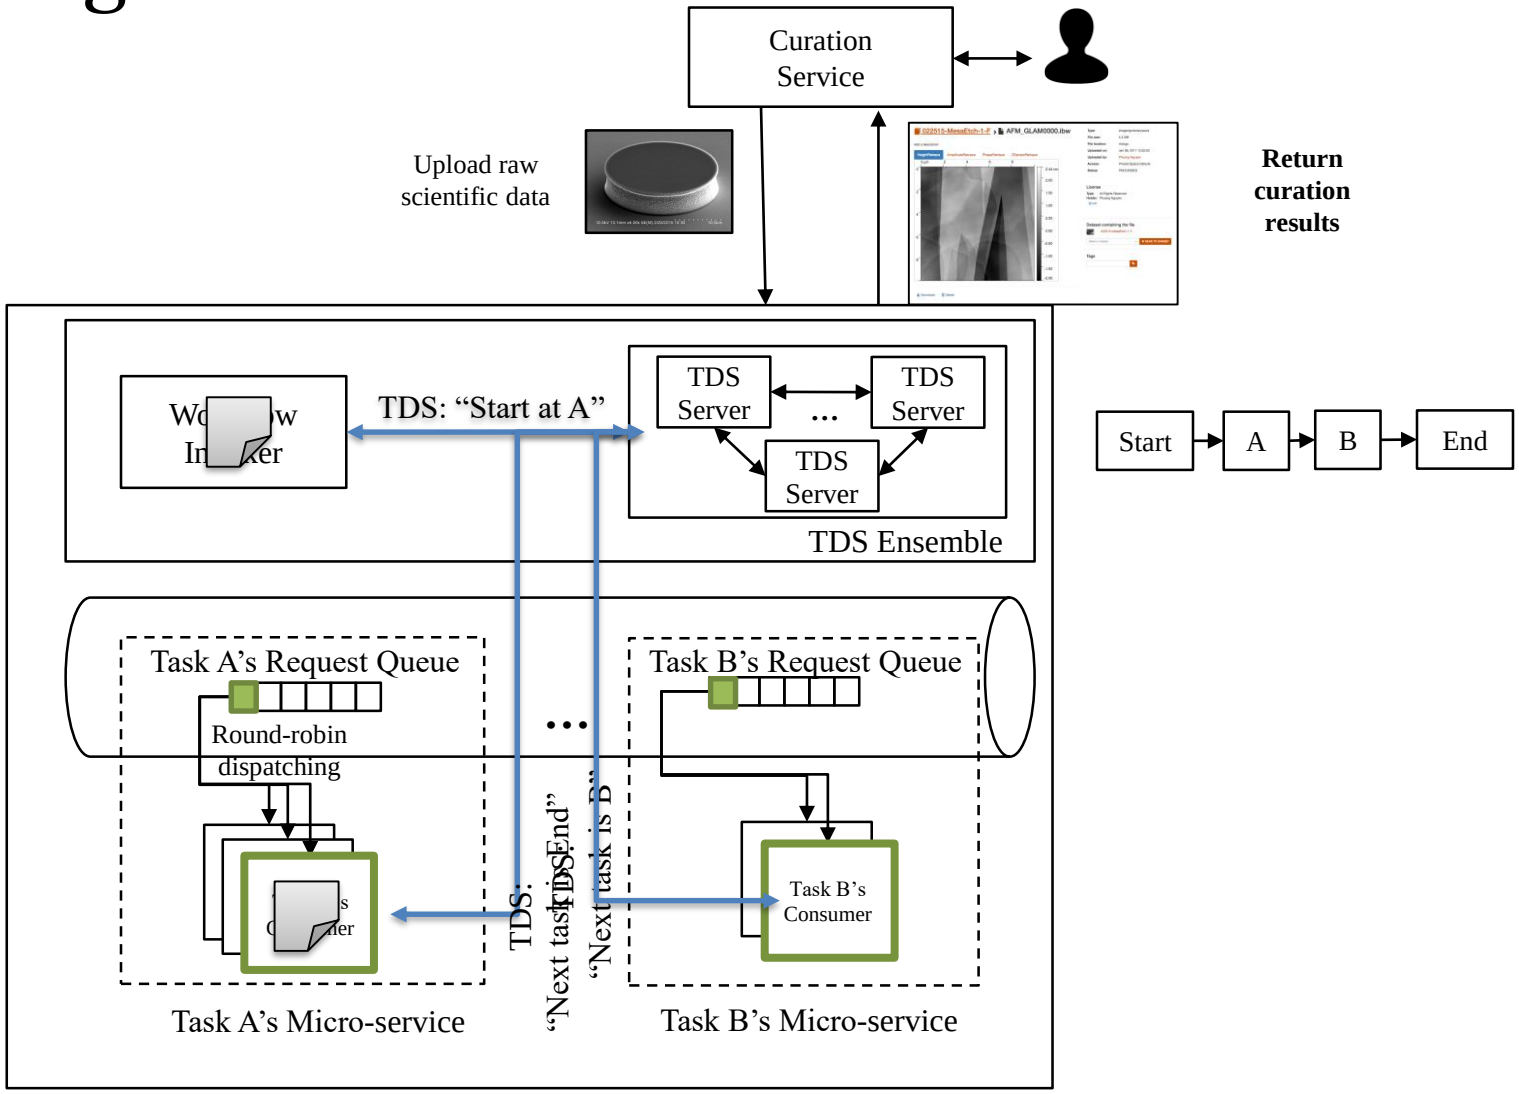

Adaptive micro-service system implementation

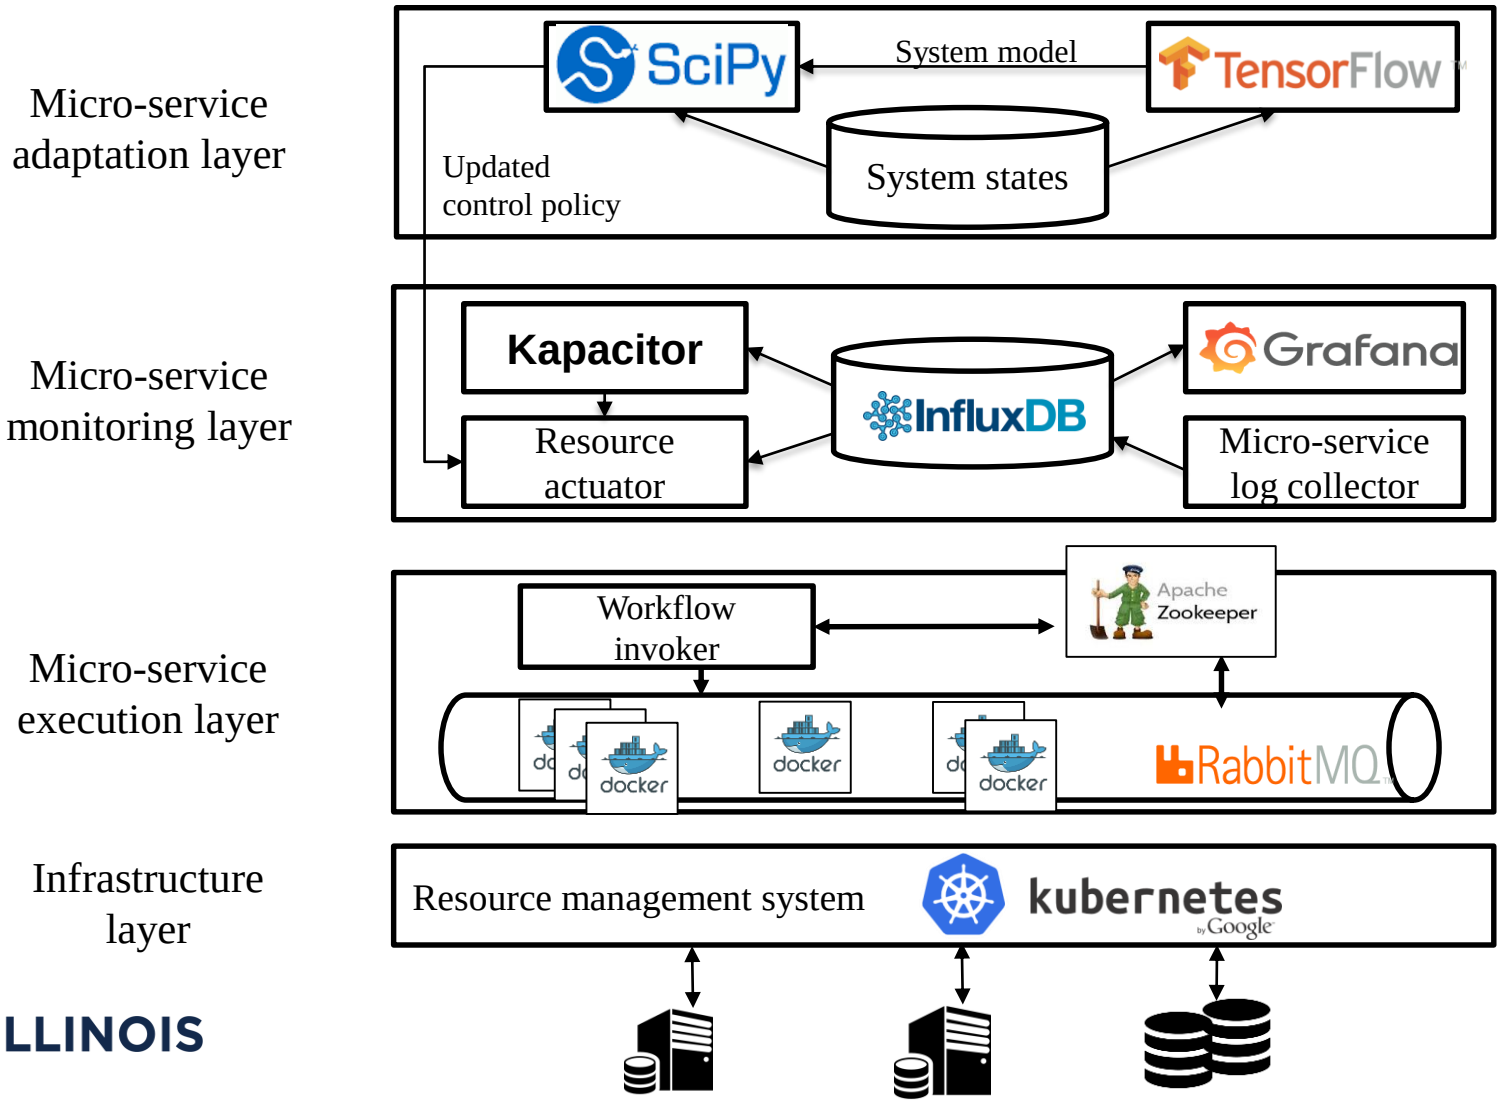

Evaluation #1: Micro-service resource adaptation

- Data processing workflows:
  - MDP: material data processing workflows (to process output of digital microscopy, such as DM3, AFM, etc.)
  - LIGO: analyze data to study stars and black holes

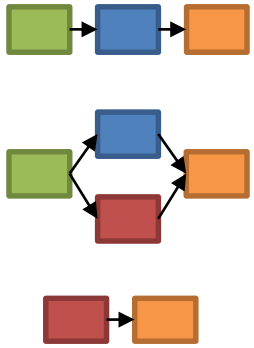

MDP workflows

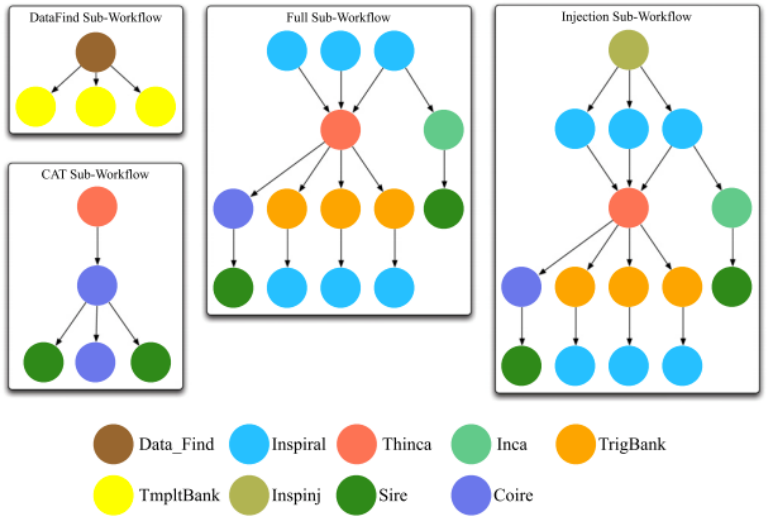

LIGO workflows

Effectiveness of neural network-based system identification

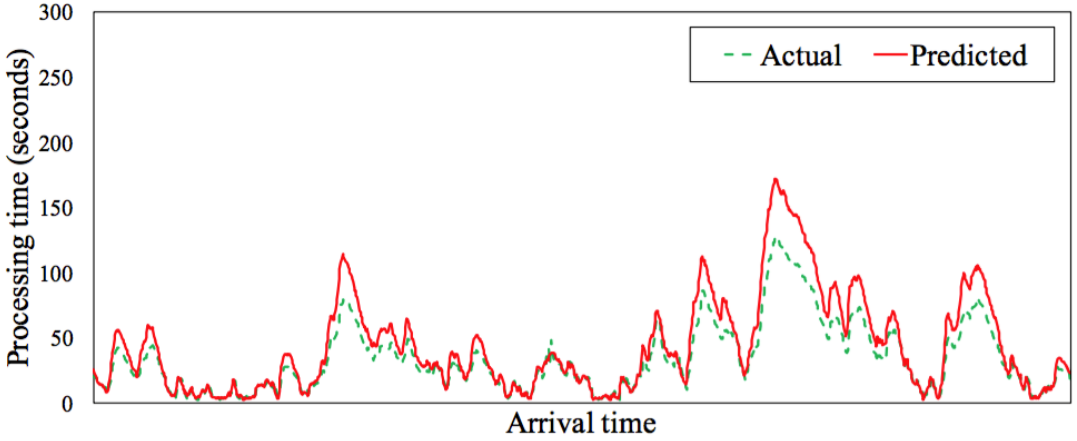

(a) MDP

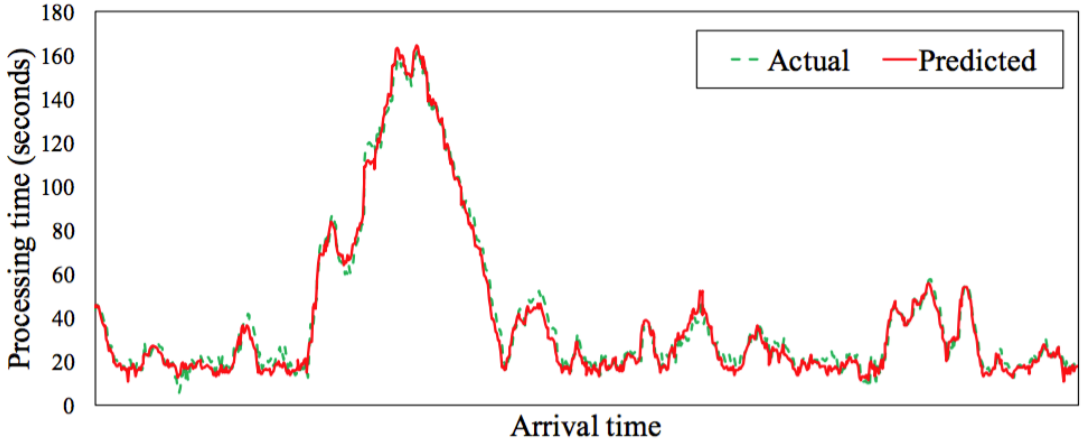

(b) LIGO

CASE STUDY: 4CeeD - Real-time Acquisition and Analysis Framework for Materials-related Cyber-Physical Environments

4CeeD: Extending micro-service infrastructure to material-related environment

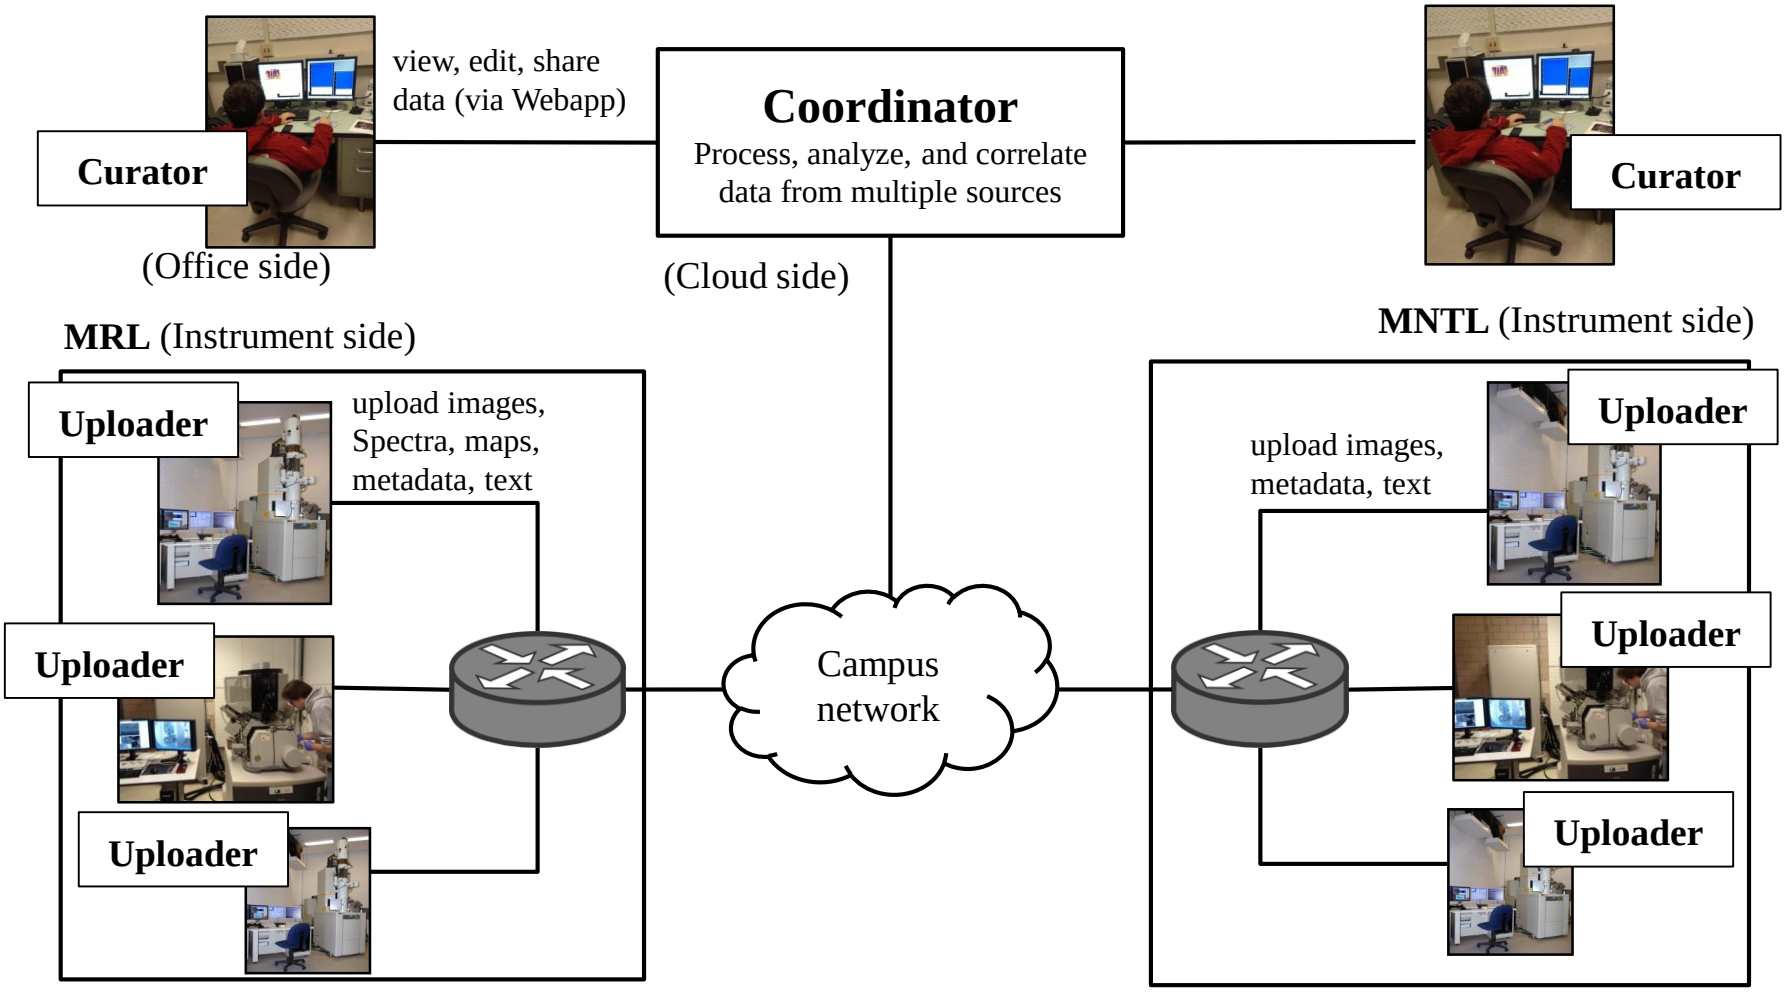

Phuong Nguyen *et al., "4CeeD: Real-time Acquisition and Analysis Framework for Materials-related Cyber-Physical Environments"*. In CCGrid 2017, **Best Paper Award**

4CeeD Uploader Service
(Simple and Speed-Up Usage at Microscopes)

Simple steps, with support for advanced usage

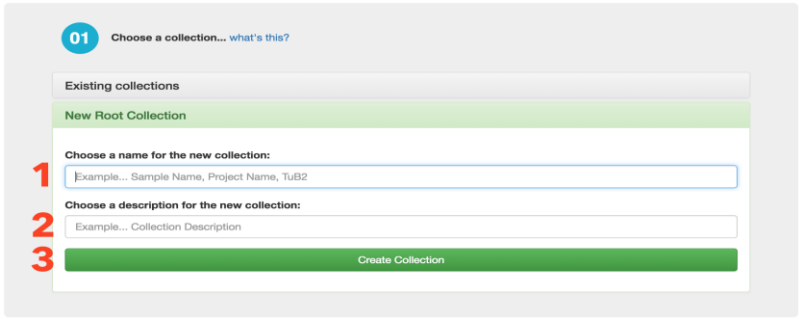

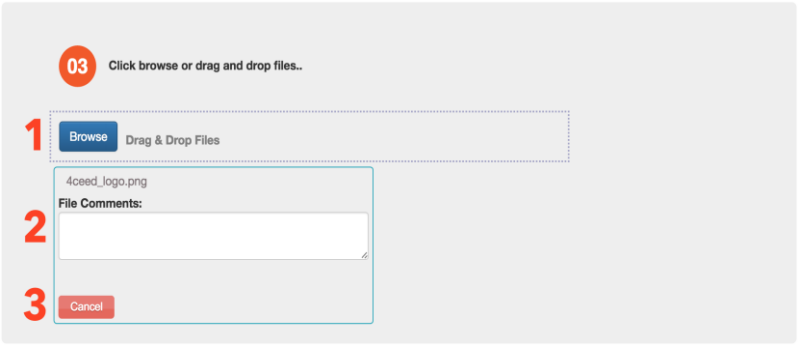

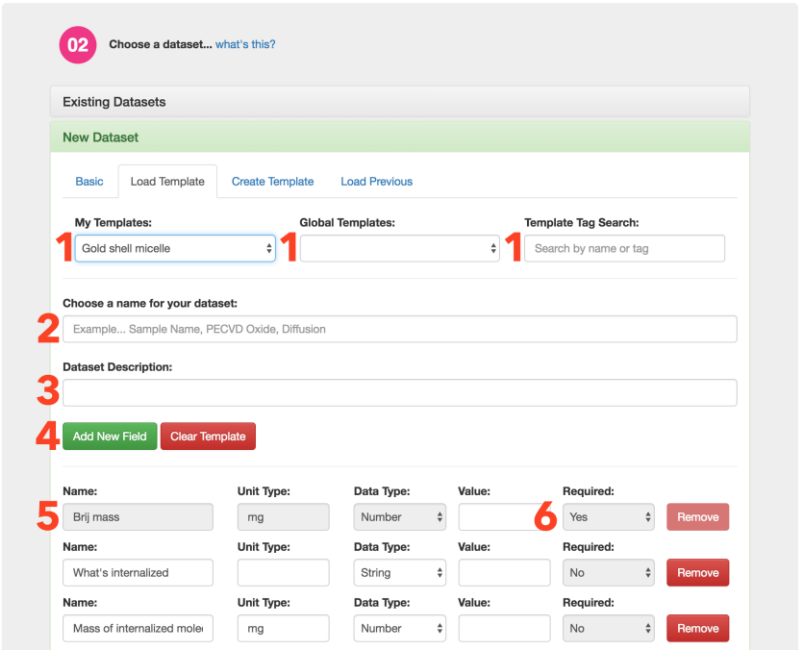

1. Choose or select a collection.

2. Load template and enter user defined metadata to create a dataset.

3. Upload files to cloud coordinator.

ILLINOIS

4CeeD: Two Types of Uploaders

- Smart Dropbox-like Uploaders
- Two types of Uploaders exist in 4CeeD to make it easier to import data.
  - **Standard Upload** allows for templates to be used
  - **Zip Upload** allows for amounts of data to be uploaded with file structure kept intact.
-

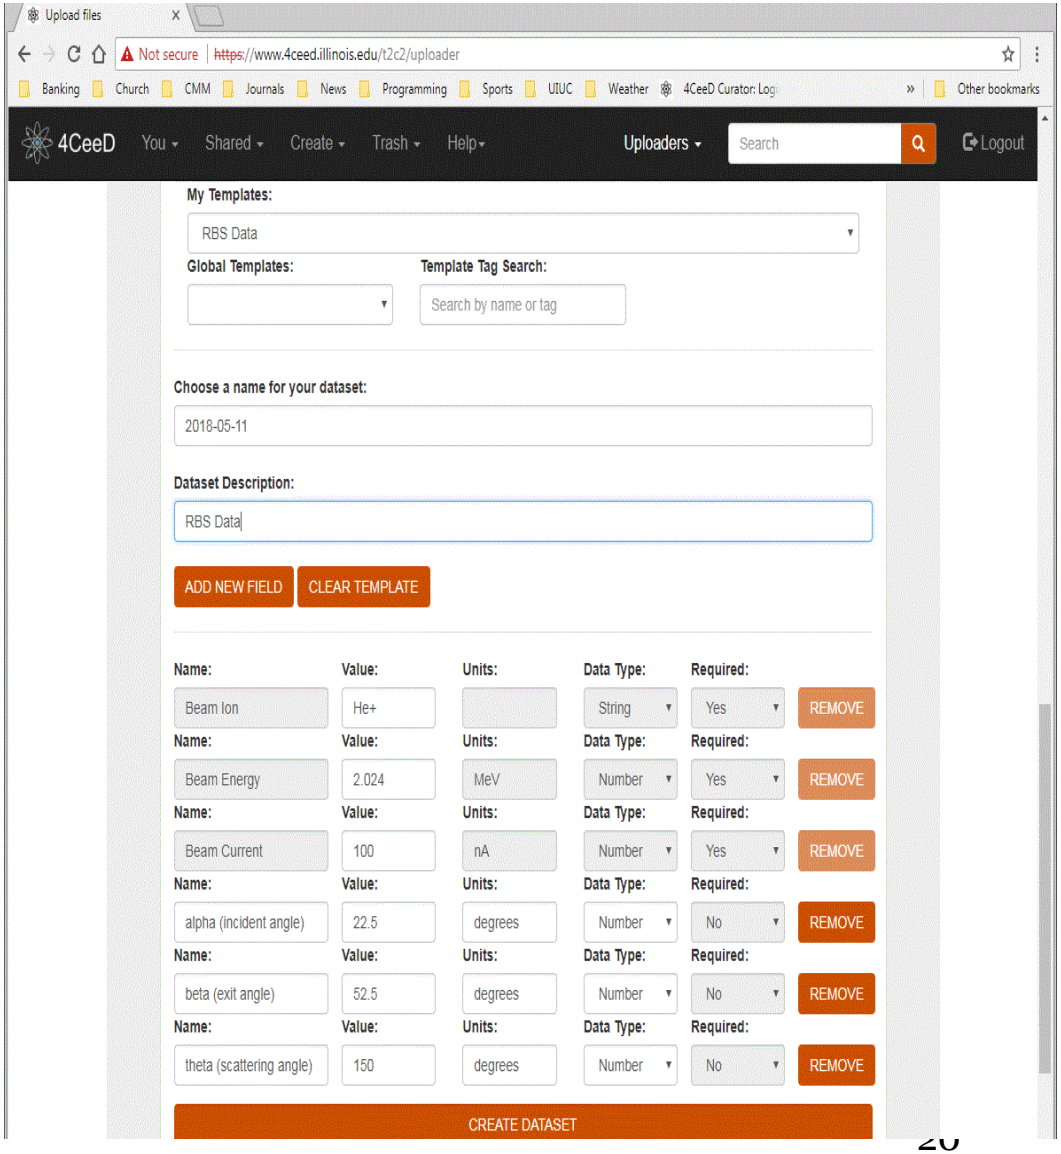

4CeeD Curator Service (Speed-Up Curation)

File View

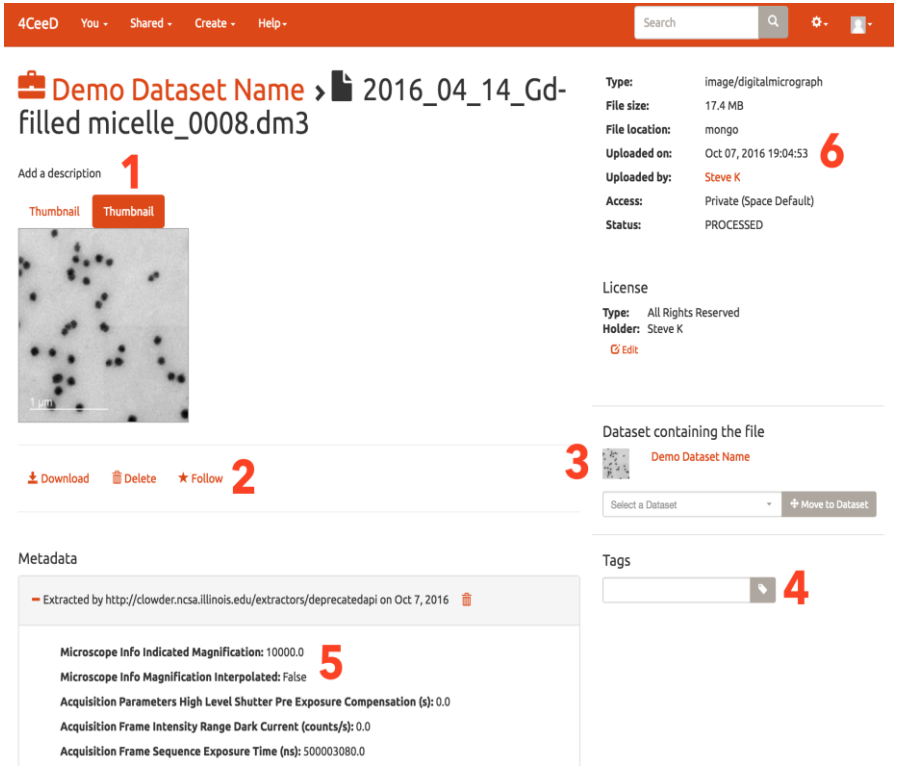

[Preview, annotate, download, extracted metadata]

Dashboard View

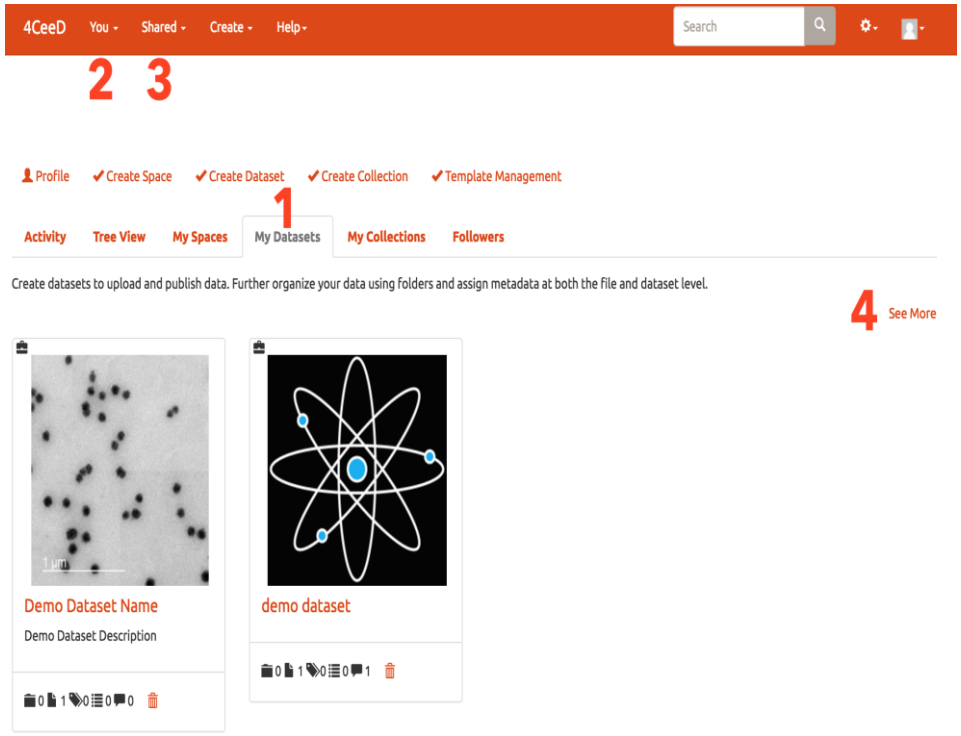

[Dashboard management]

ILLINOIS

4CeeD: Extractors as Micro-Services at Cloud Side

- TEM Extractor
  - Works with DM3 Files
- SEM Extractor
- AFM Extractor

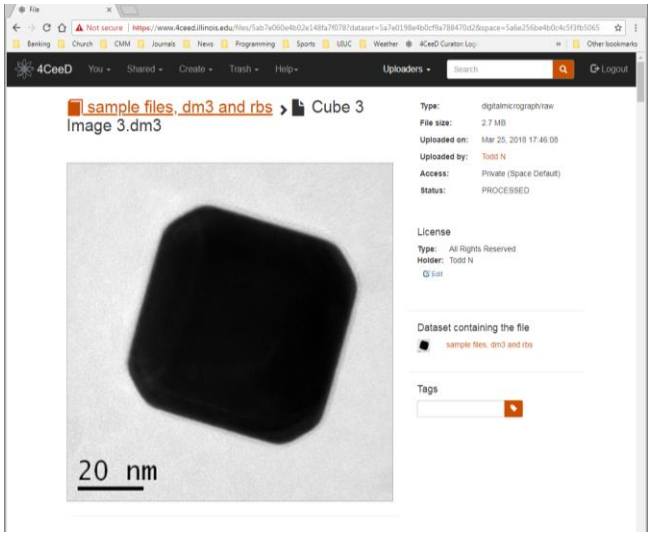

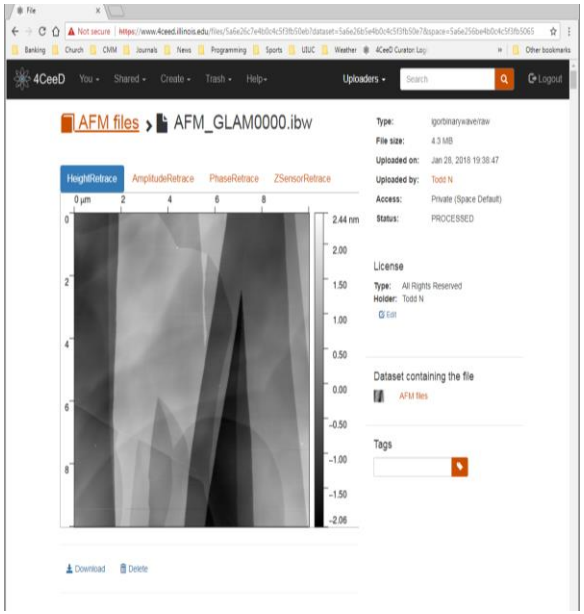

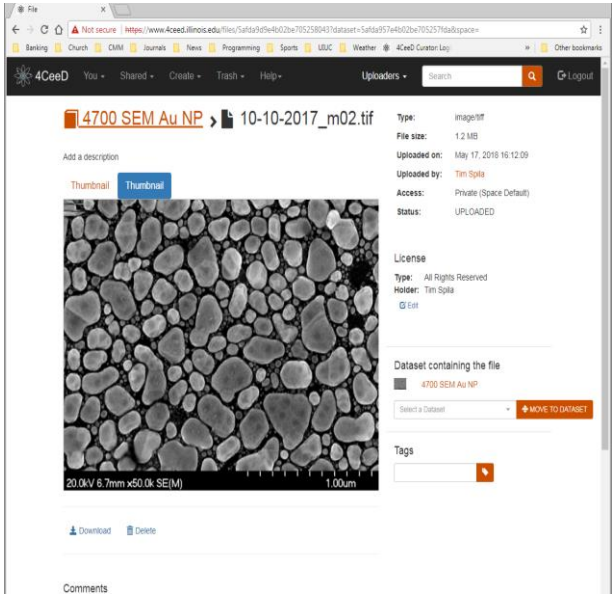

4CeeD Smart Data Management

4CeeD Data Model organizes projects into collections, datasets, and files.
These can then be shared in spaces. 4CeeD utilizes and modifies NCSA Clowder data management system.

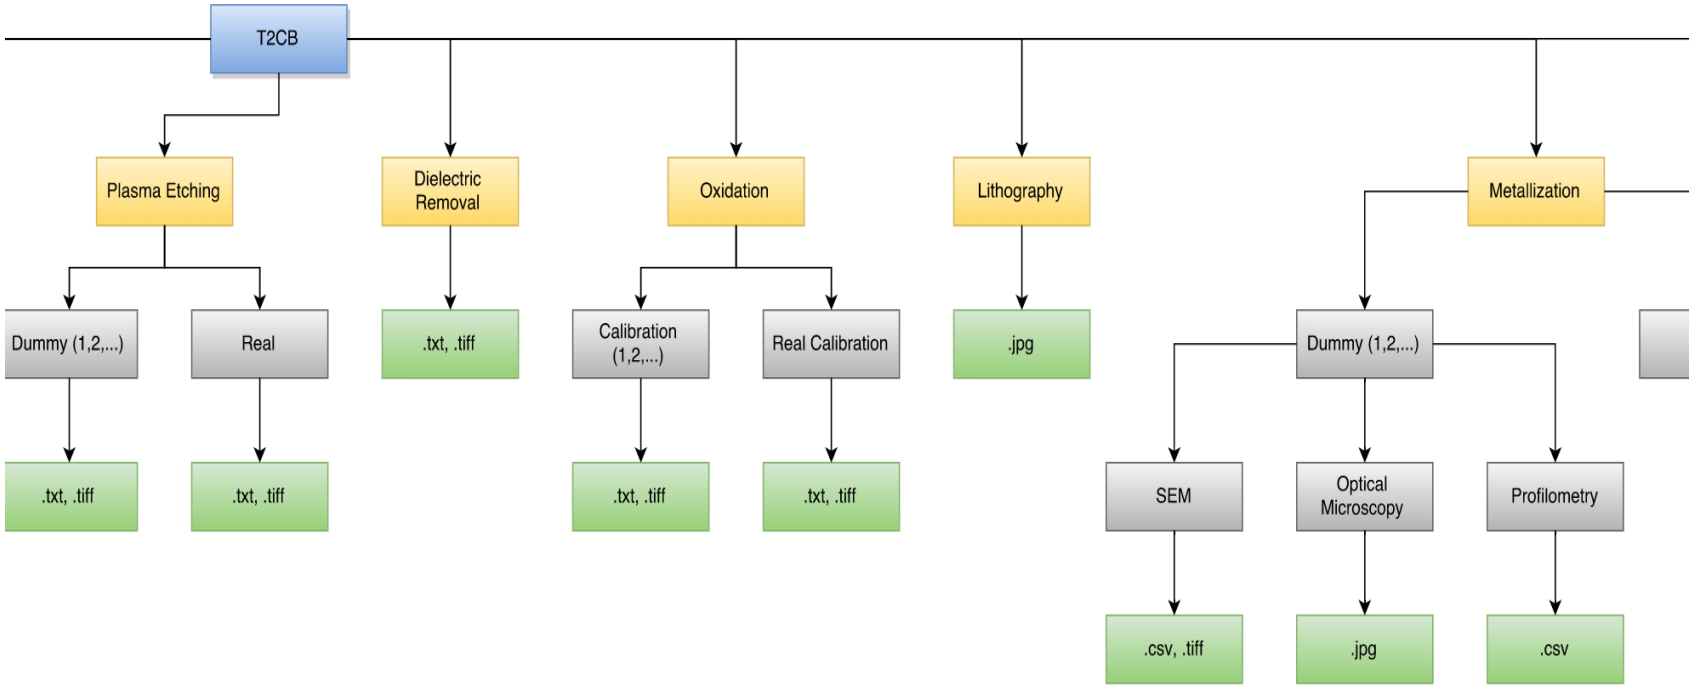

4Ceed++: Jupyter Notebook Integration

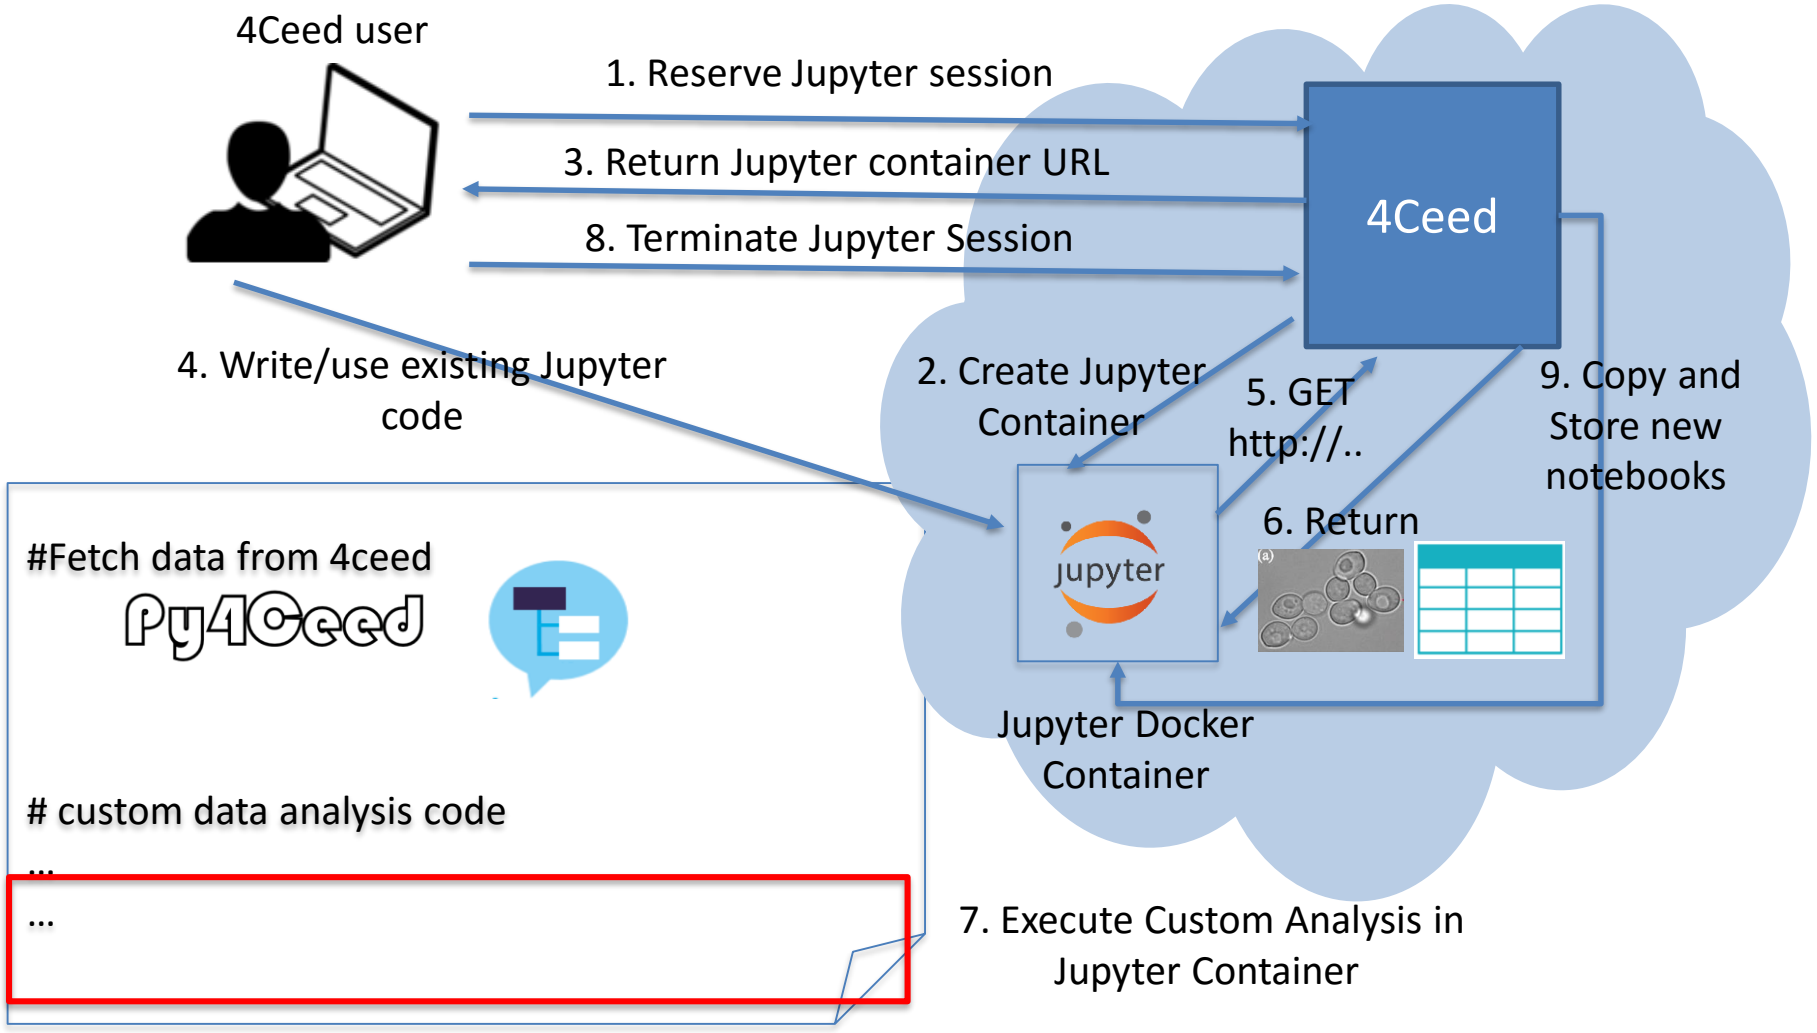

ILLINOIS

4CeeD Production System

Goals:

- Redundancy
- Availability
- Scalability

Storage Layer:

- **40 TB (20 TB per investor)**
- Replicated for redundancy

Compute Layer:

- Docker container

orchestration (Kubernetes)

- Single master

(High Available masters in future)

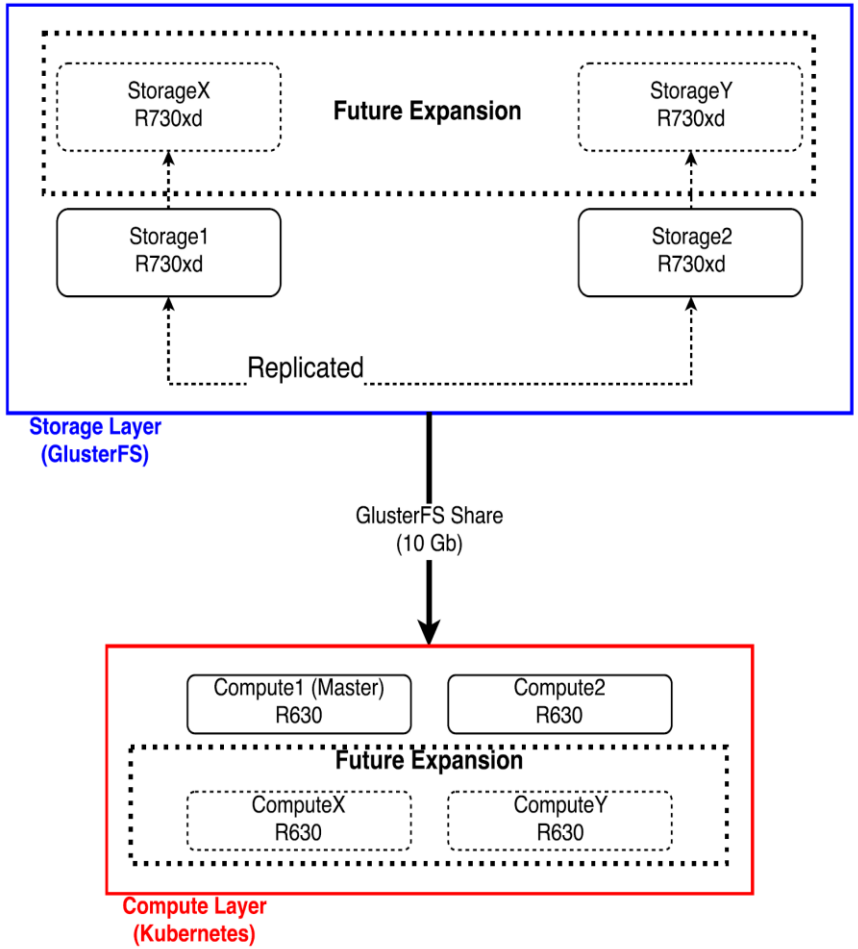

Our approach

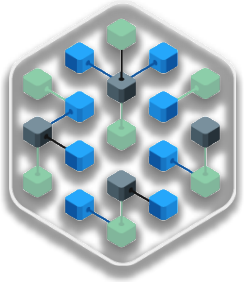

✓ Micro-service private cloud execution environment for instrument data curation and coordination (4CeeD)

✓ Data acquisition from aging instruments (BRACELET)

Current Status for Campus Cyberinfrastructure regarding Aging Scientific Instruments

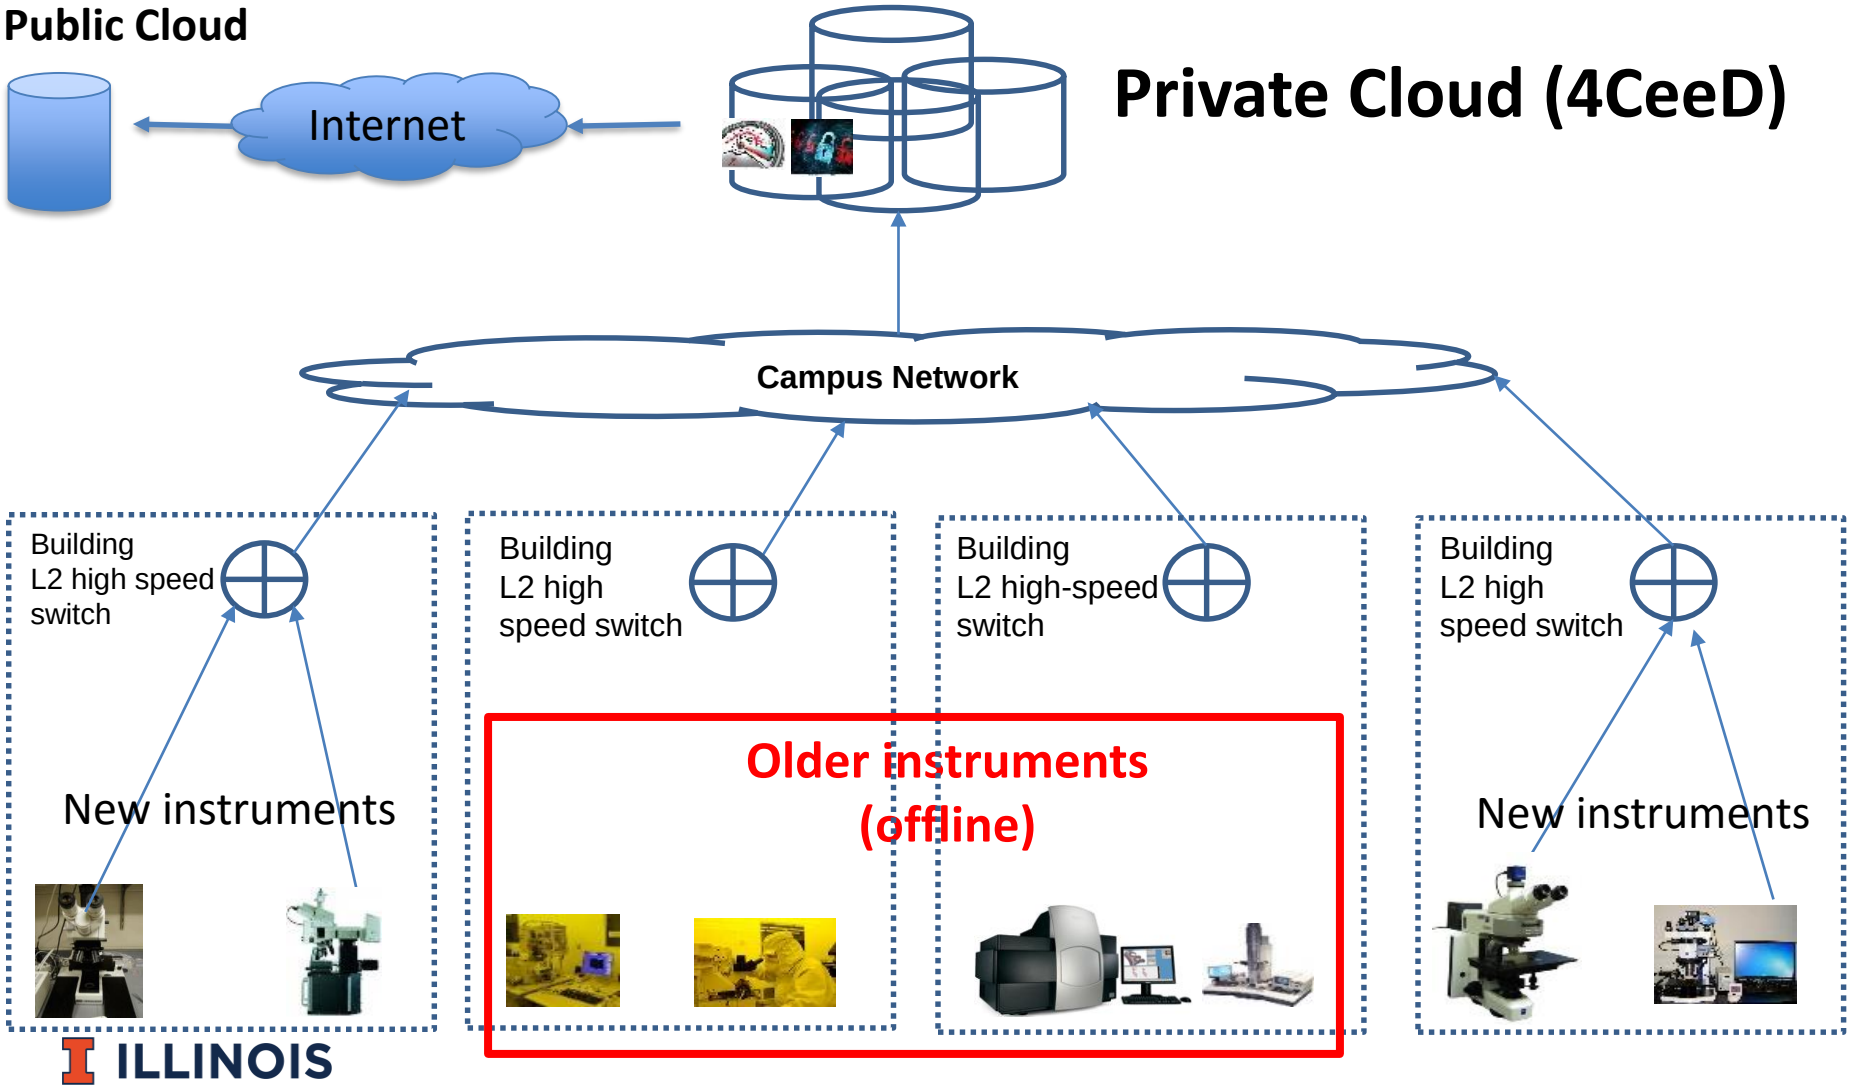

Our Goal for Campus Cyberinfrastructure regarding Aging Scientific Instruments (2)

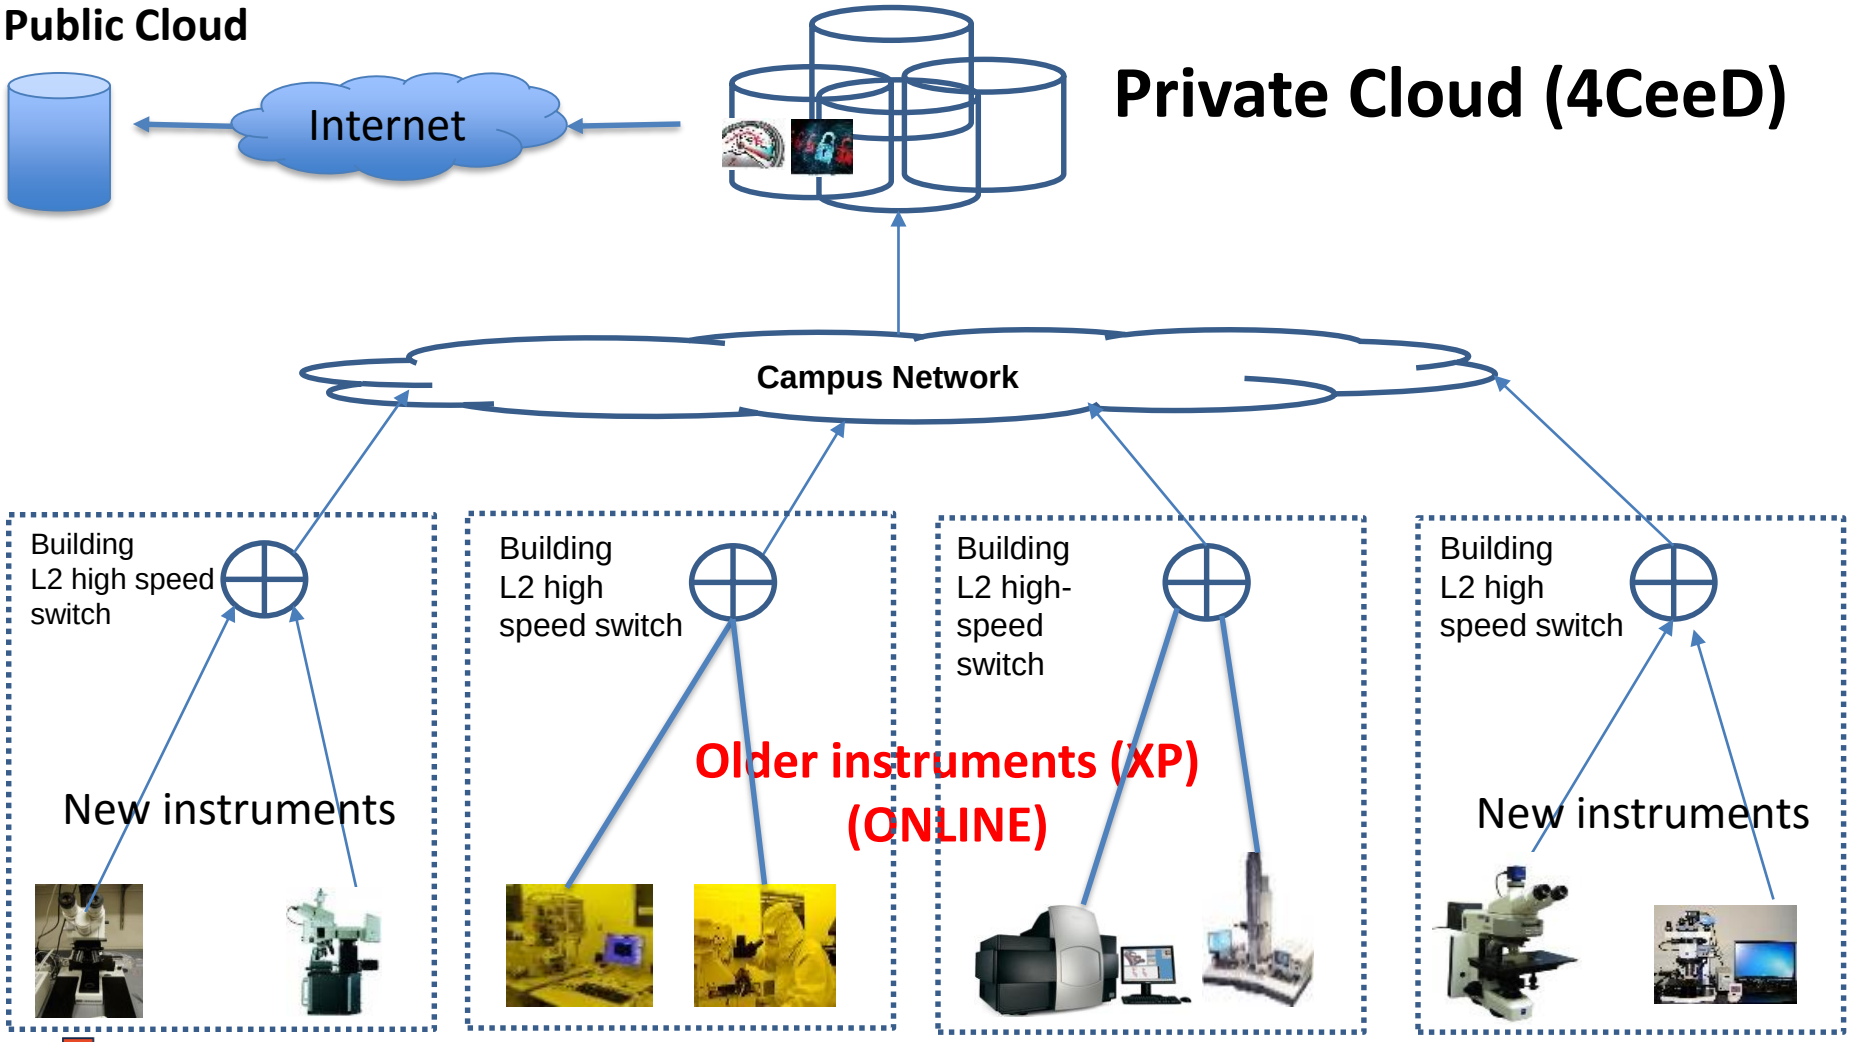

ILLINOIS

Challenges of connecting offline older instruments

- **Performance mismatch**: Older instruments' Windows XP runs network protocols at lower bandwidth speeds (10Mbps or 100Mbps)

- **Obsolete security**: Older devices and their OS systems cannot be patched, hence being vulnerable & taken offline

BRACELET: Putting edge device between older instruments and private cloud

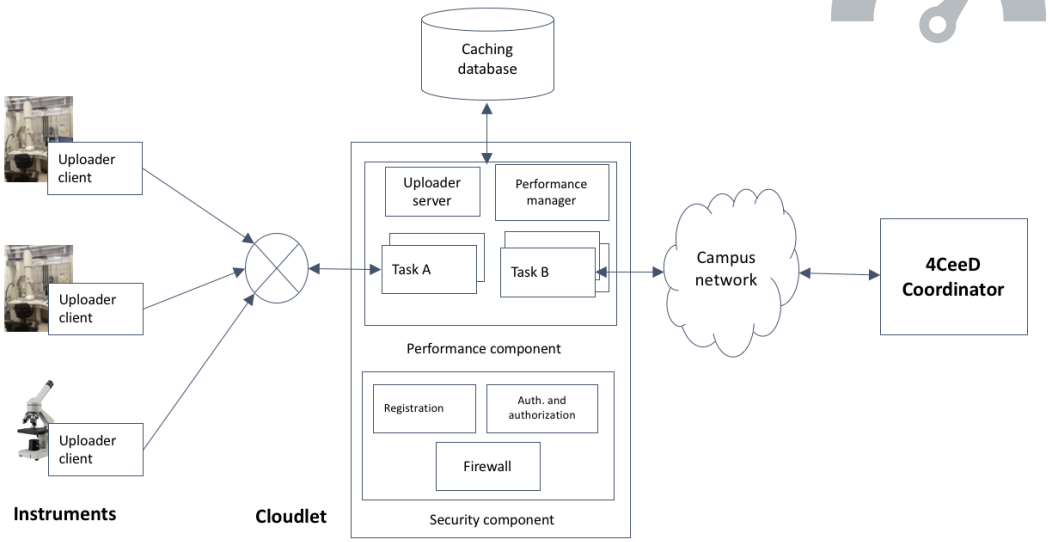

BRACELET in 3-tier architecture

Performance:

- Have two network interfaces configured at different speeds
- Traffic shaping & offloading between edges & cloud

Security:

- User & instrument registration
- Data encryption during upload
- Firewall to protect against external threats

BRACELET: Extending cloud-based architecture to the edges for seamless integration

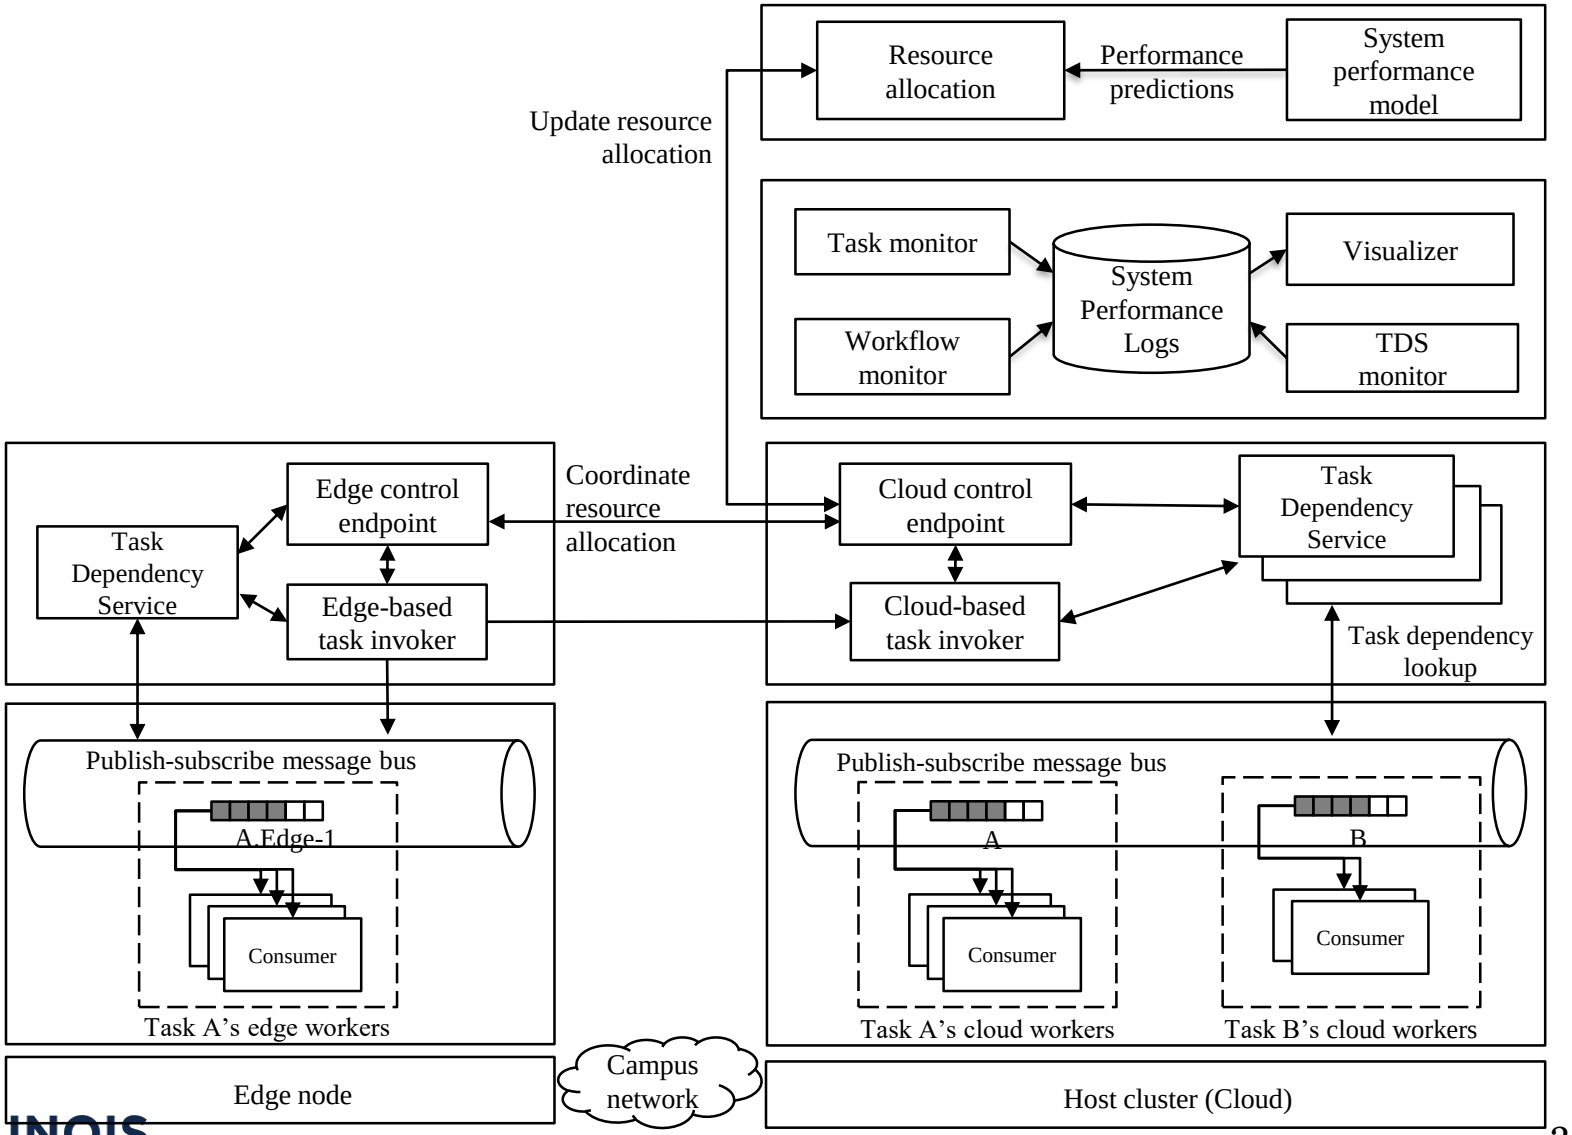

BRACELET's implementation & integration with 4CeeD

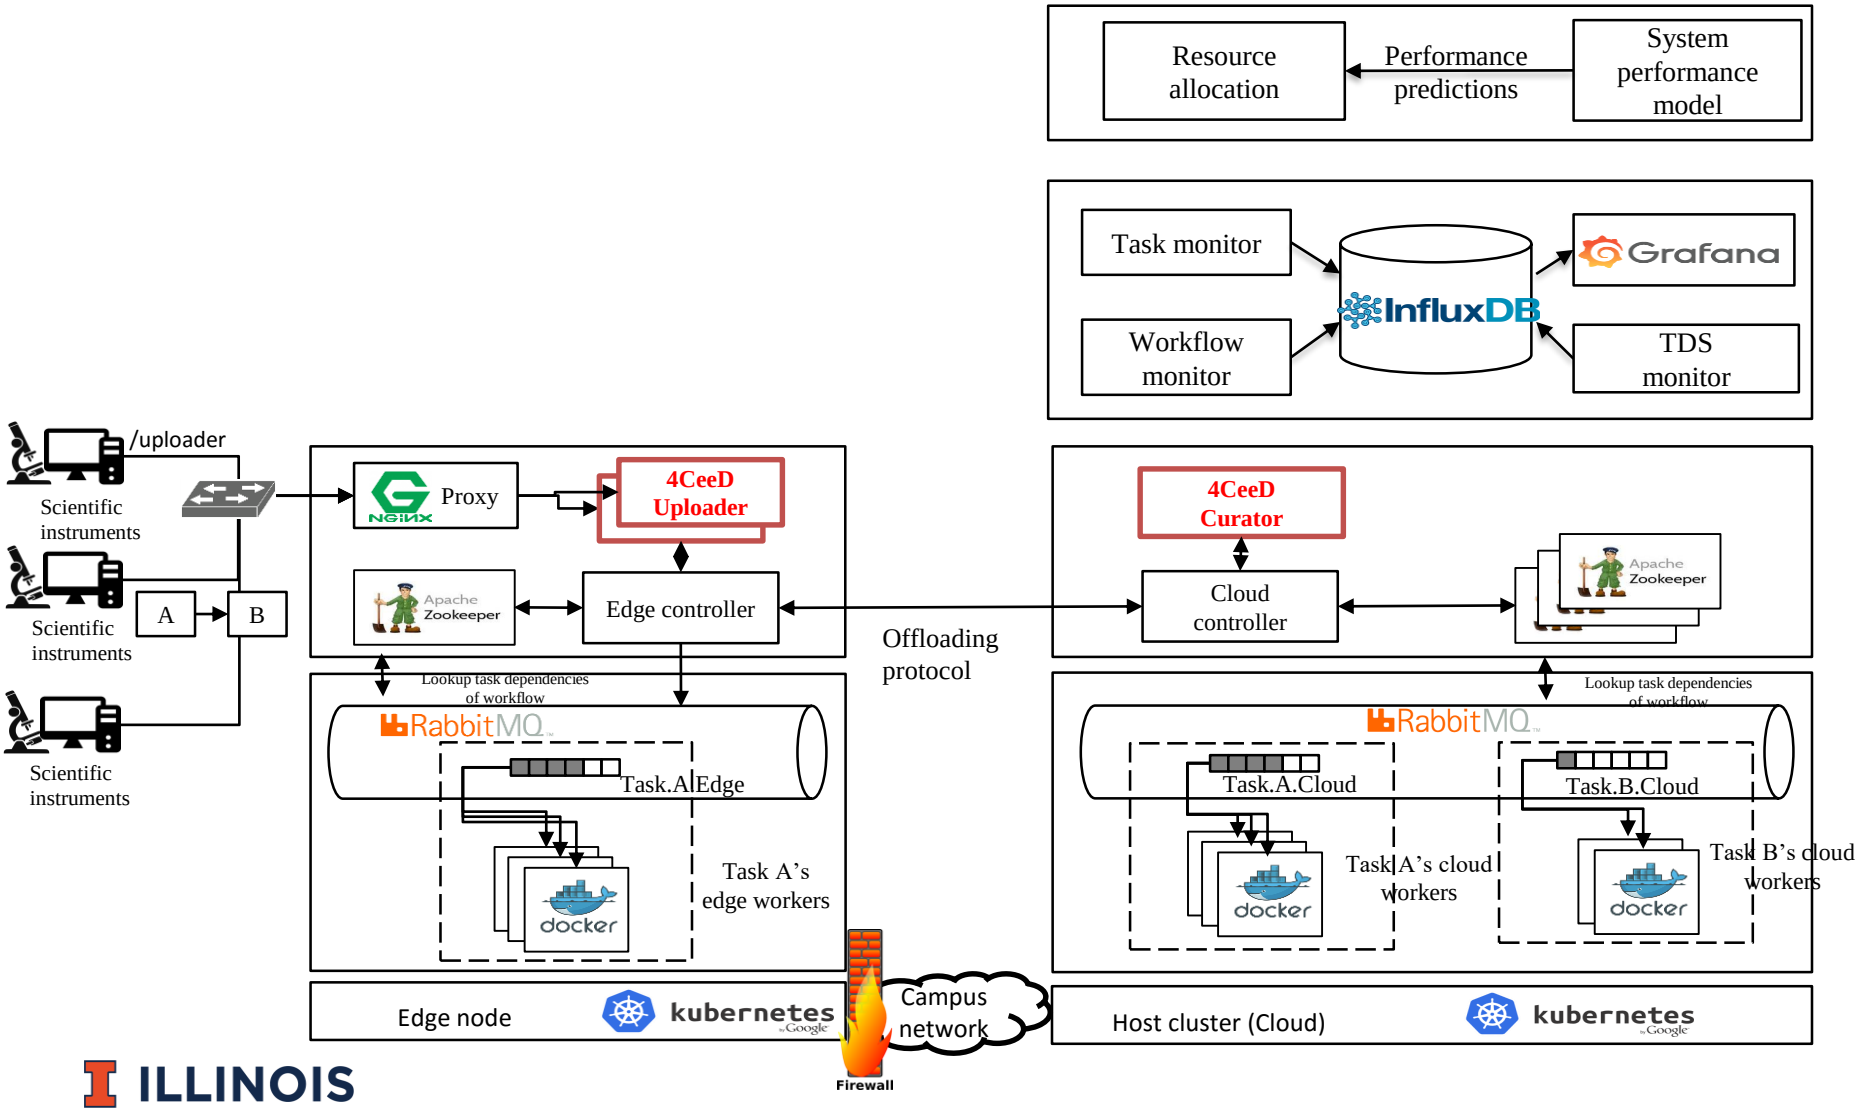

ILLINOIS

User authentication from instruments via BRACELET

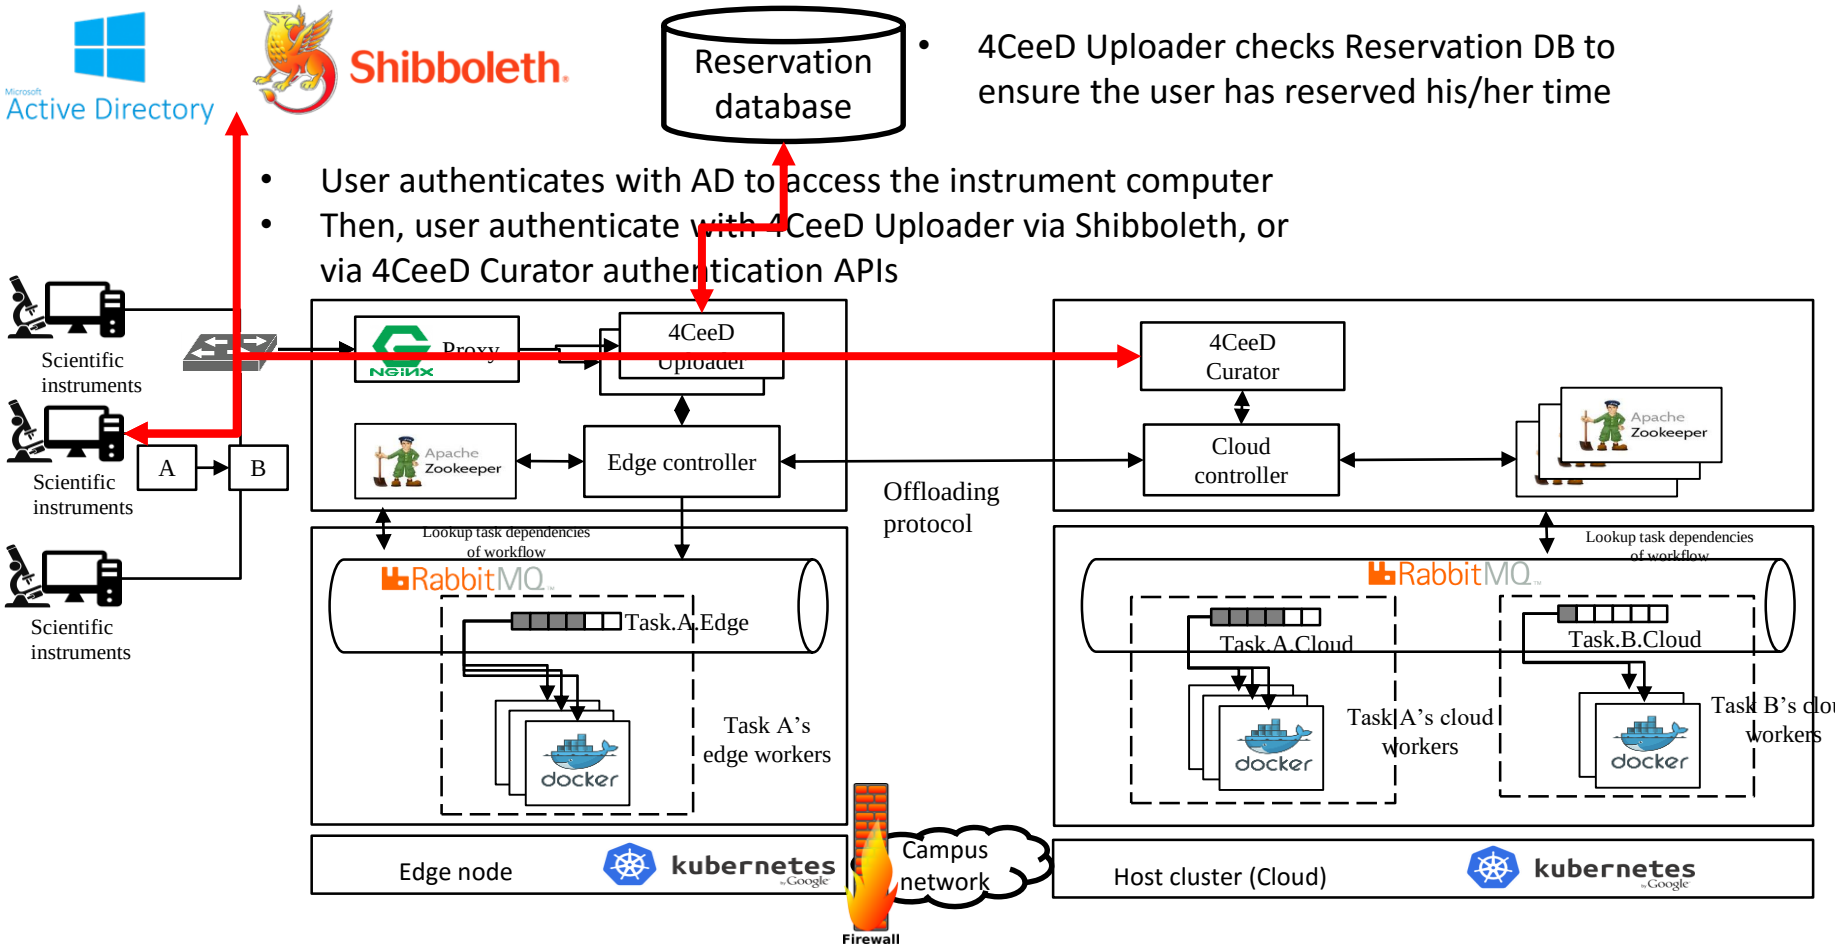

ILLINOIS

Computation offloading between edge & cloud

- After learning about the placement, the data consumer forwards processing requests to the next task (first task)
- After processing the request, the task forwards the request to the next task (if any)
- 4CeeD Uploader communicates with local Edge controller to learn about where to send request to
- Edge controller periodically communicates with cloud controller to update task placements

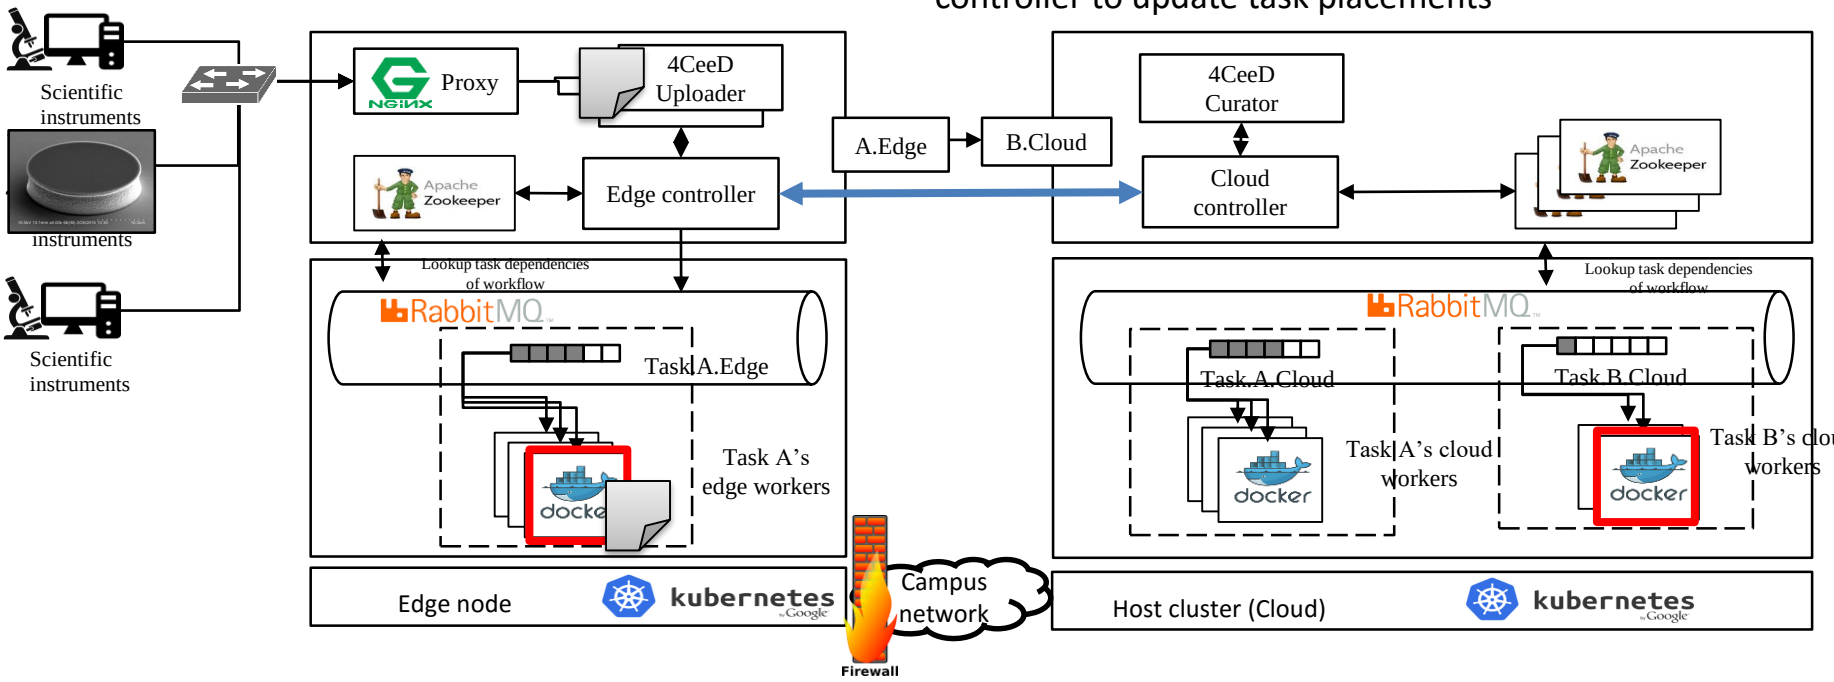

Evaluation of 4Ceed-BRACELET micro-service performance prediction

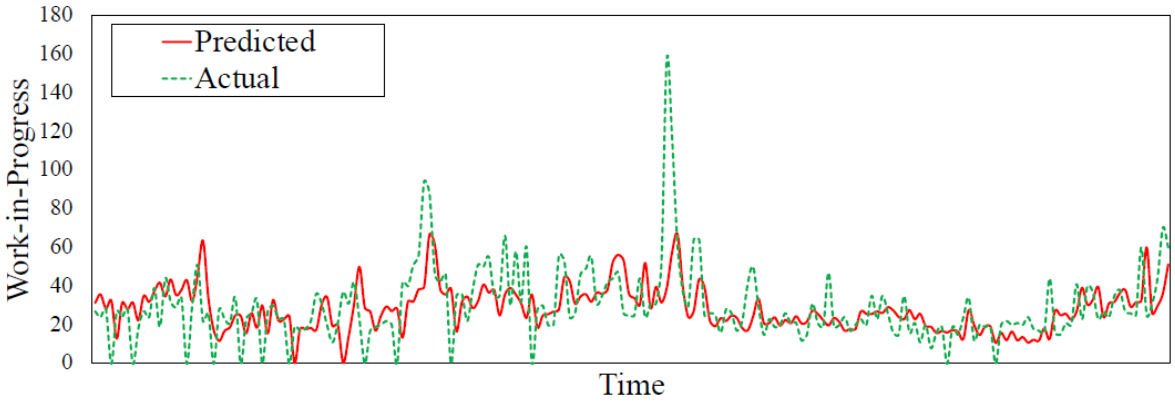

(a) Cloud-based micro-service D

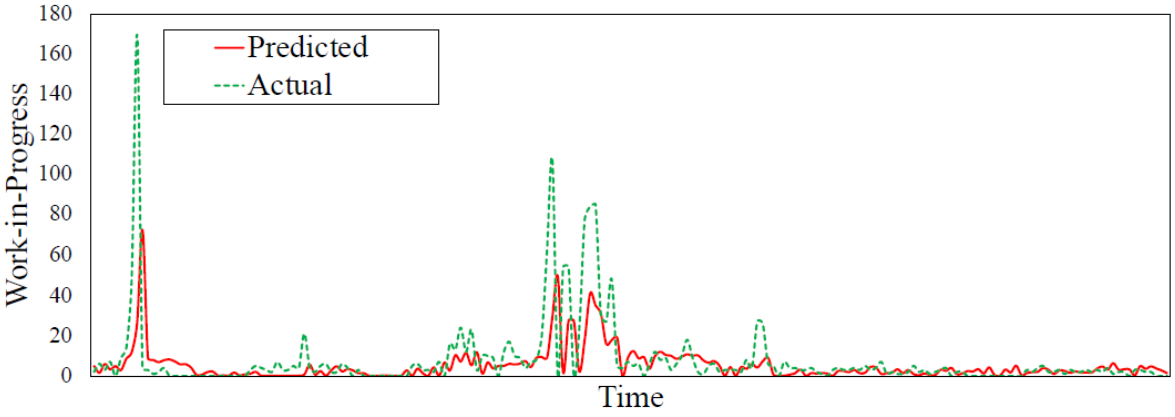

(b) Edge-based micro-service C on cloudlet E1

Software Availability

- All publication are available in IEEE Digital Libraries
- All software systems, 4CeeD, BRACELET are open source
- All Projects are described at https://t2c2.csl.illinois.edu/
- 4CeeD System is available for download https://github.com/4ceed
  - Contact: Steve Konstanty (stevek@illinois.edu) and Todd Nicholson (tcnichol@illinois.edu)
- Bracelet system is in testing phase. Software will be released by December 2019.
  - Contact: Steve Konstanty (stevek@Illinois.edu)

Lessons Learned

- We have explored novel cloud system approaches that lend themselves well for real-time and trustworthy materials-to-device data and metadata storage, management, and computing of workflows over these data.

- Lightweight micro-service cloud architecture for materials genomic challenge is the way to go, including the three tier approach (instrument - private cloud – public cloud).

- Bringing aging instruments online proved to be very challenging for very old OS systems

- Hardest Part is convincing experimentalists/scientists to use new data management techniques and new cyber-infrastructure
  - Continuous training and inclusion of modern data management systems into experimental instrumentation classes will be needed.

- To achieve sustainability, it is crucial to work with **college and campus IT** teams!

Acknowledgement

- Research and Development Team:
  - Phuong Nguyen (CS), Tarek Elgamal (CS), Zhe Yang (CS), Tuo Yu (CS), Xiaoyuan Wang (CS), Steve Konstanty (Senior Research Programmer), Todd Nicholson Research Programmer), Patrick Su (ECE), Robert Kaufman (ECE), Tommy O'Brien (ECE).

- NSF Funding: ACI DIBBS and CC Programs
  - *NSF ACI DIBBS Award 1443013 – 4CeeD (T2C2 Project)*
  - *NSF OAC CC Award 1659293 - BRACELET*

- Co-PIs and Collaborators:
  - Roy Campbell (CS/CSL/Engineering IT), Indranil Gupta (CS), Paul Braun (MRL), Brian Cunningham (MNTL/ECE), Greg Pluta (MNTL), Tim Spila (MRL), Michael Chan (Engineering IT), Tracy Smith (IT Tech Services), Kenton McHenry (NCSA), John Dallesasse (ECE/MNTL), Mark McCollum (MNTL), Gianni Pezzarossi (Engineering IT), Stuart Turner (Engineering IT), Laura Herriott (Engineering IT)

Publications

- Phuong Nguyen, Klara Nahrstedt, “Resource Management for Elastic Publish Subscribe Systems: A Performance Modeling-based Approach”, **IEEE International Conference on Cloud Computing (CLOUD 2016)**, San Francisco, CA, June 2016
- Phuong Nguyen, Steven Konstanty, Todd Nicholson, Thomas O'Brien, Aaron Schwartz-Duval, Timothy Spila, Klara Nahrstedt, Roy Campbell, Indranil Gupta, Michael Chan, Kenton McHenry and Normand Paquin, “4CeeD: Real-Time Data Acquisition and Analysis Framework for Material-related Cyber-Physical Environments”, IEEE/ACM 17th **IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing**. Madrid, Spain, May 14-17, 2017 (Best Paper Award)
- Phuong Nguyen, Klara Nahrstedt, “MONAD: Self-adaptive Micro-service Infrastructure for Heterogeneous Scientific Workflows”, 14th **IEEE International Conference on Autonomous Computing (ICAC 2017)**, July 17-21, 2017, Columbus, Ohio
- Phuong Nguyen, Tarek Elgamal, Steve Konstanty, Todd Nicholson, Stuart Turner, Patrick Su, Michael Chan, Klara Nahrstedt, Tim Spila, Kenton McHenry, John Dallesasse, Roy Campbell, “BRACELET: Edge-Cloud Micro-service Infrastructure for Aging Scientific Instruments”, **IEEE International Conference on Computing, Networking, and Communications (ICNC) 2019**, Hawaii, February 2019.
- Zhe Yang, Phuong Nguyen, Haiming Jin, Klara Nahrstedt, “MIRAS: Model-based Reinforcement Learning for Microservice Resource Allocation over Scientific Workflows”, **IEEE International Conference on Distributed Computing Systems (ICDCS 2019)**, July 2019, Dallas, TX.
- Zhe Yang, Patrick Su, Robert Kaufman, Steve Konstanty, John Dallesasse, Klara Nahrstedt, “SENSELET: Sensory Network Infrastructure for Scientific Lab Environments”, **ACM Practice and Experience in Advanced Research Computing Conference Series (PEARC 2019)**, Chicago, IL (Poster), August 2019