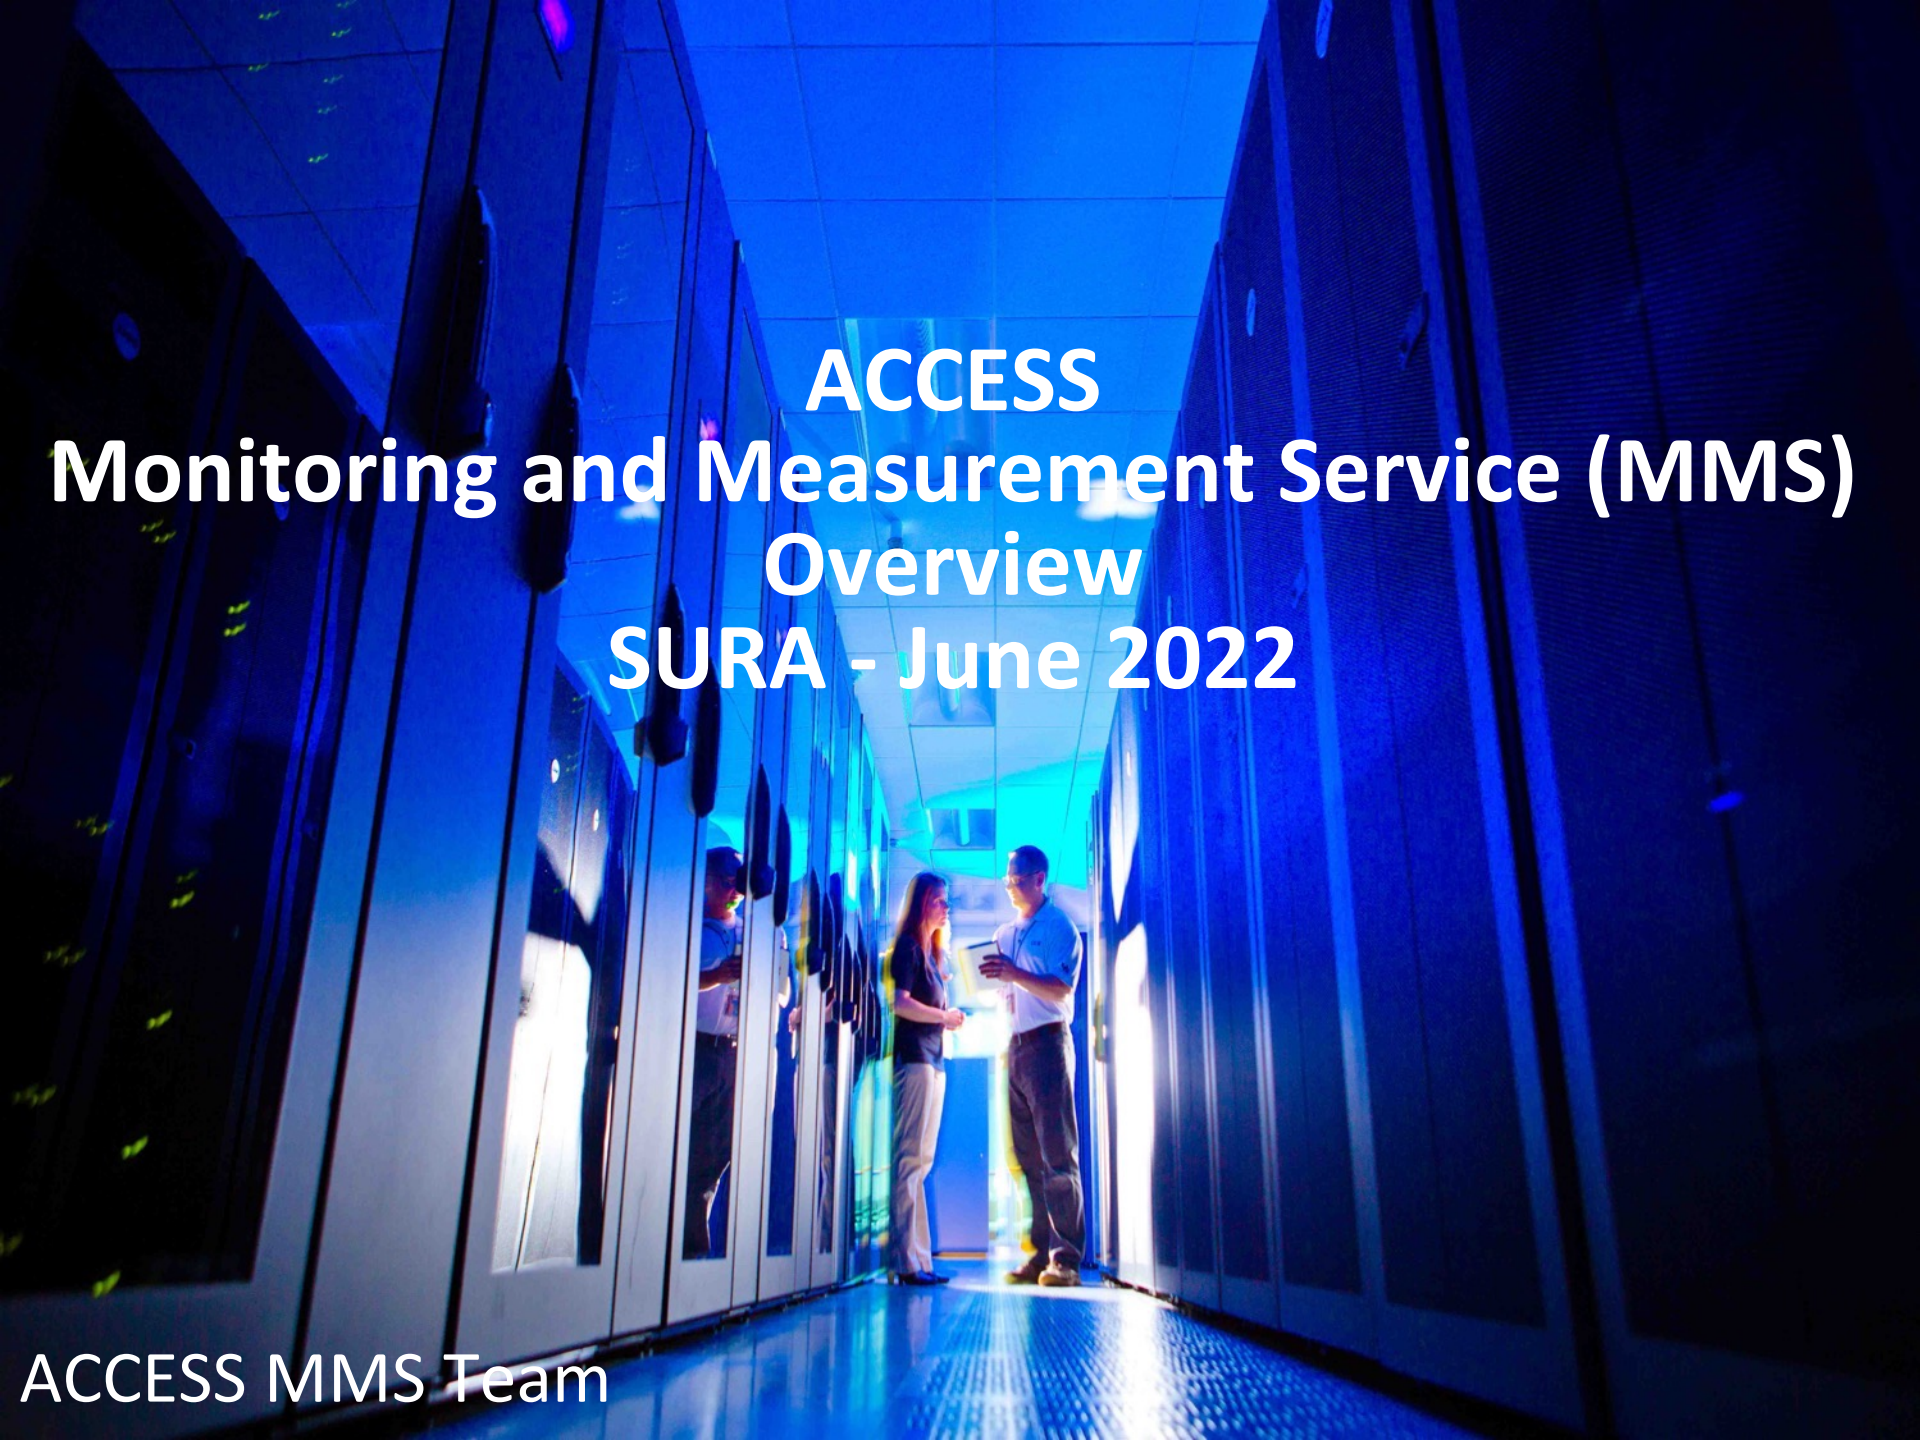

ACCESS
Monitoring and Measurement Service (MMS) Overview
SURA - June 2022

ACCESS MMS Team

Outline

- **ACCESS Monitoring and Measurement Service (MMS)**
  - ACCESS MMS Team
- **Why a CI Monitoring and Measurement Service?**
  - Potential for High Impact
  - Stakeholder Benefits
  - Data/Services Available to ACCESS Tracks
- **ACCESS MMS – What's New**
  - Data Analytics Framework
  - CI Simulator
  - Monitoring CI Ecosystem
  - Application Power Monitoring
  - Network Integration
  - Value Analytics Realm
  - Open XDMoD Development
  - ACCESS Auditing
  - Service Model
- **Discussion**

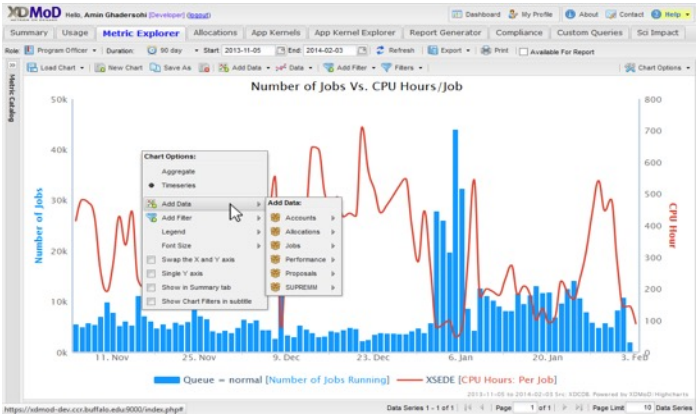

XDMoD
METRICS ON DEMAND

ACCESS Monitoring and Measurement Service

- **Comprehensive framework for CI system management**
- **Understand and optimize resource utilization and performance**
  - Provide instantaneous and historical information on utilization
  - Measure Quality of Service of CI systems and applications
  - Measure and improve job and system level performance
  - Inform computing system upgrades and procurements
- **XDMoD (XD Metrics on Demand) tool**
  - Analytics Framework for XSEDE/ACCESS
- **Open XDMoD*: Open Source version for CI centers**
  - Used to measure and optimize performance of HPC centers
  - 250+ academic, governmental, & commercial installations worldwide
  - https://open.xdmod.org/

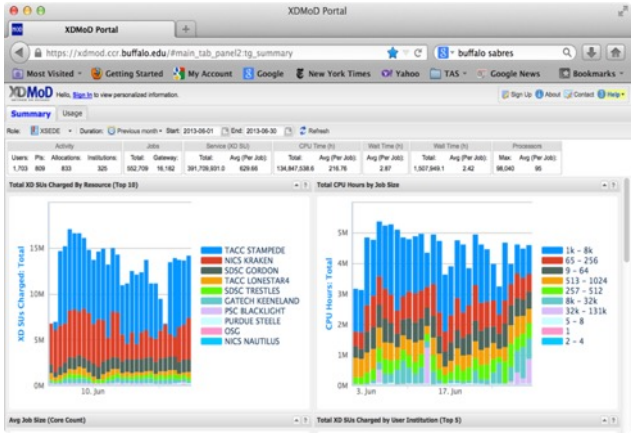

XDMoD
METRICS ON DEMAND

ACCESS MMS Team

- **University at Buffalo – Center for Computational Research**
  - Matt Jones, Bob DeLeon, Joe White, Jeff Palmer, Nikolay Simakov, Ryan Rathsam, Gregary Dean, Hannah Taylor, Conner Saeli
- **Roswell Park**
  - Tom Furlani (PI)
- **TACC**
  - Bill Barth, Stephen Harrell, Matt Cawood (performance monitoring)
- **Tufts**
  - Abani Patra (performance monitoring)
- **Indiana**
  - Jennifer Schopf (Netsage)
- **Case Western**
  - Vipin Chaudhary (application anomaly detection)
- **SDSC**
  - Shava Smallen (CloudBank)

XDMoD
METRICS ON DEMAND

Why a CI Monitoring and Measurement Service?

- **Monitoring and Audit of NSF funded CI**
  - Technology Audit Service (TAS): 2010 – 2015
  - XD Metrics Service (XMS): 2015 – 2022
- **The landscape before TAS/XMS**
  - Accounting data incomplete and only available quarterly
  - Job level performance data not available
  - No external measure of Quality of Service (QoS)
  - RP reporting to the central database was inconsistent
  - RP quarterly and annual reports done manually
  - CI planning and analysis difficult due to lack of accessibility of historical data (utilization, allocation, etc)

NPACI

TeraGrid™

XSEDE
Extreme Science and Engineering Discovery Environment

NSF

XDMoD
METRICS ON DEMAND

Overall CI Performance Matters

- **CI systems are typically oversubscribed**
  - Improving application or system performance will improve overall job throughput
  - Free's up otherwise wasted CPU cycles for useful work
- **Small improvements in system performance can have high impact**
  - Every 1% increase in system performance on the resources provided through XSEDE translates into the ability to allocate an additional 101 M CPU hours annually
  - Corresponds to a savings of $5M*

* Assuming a rate of $0.05 per CPU hour

XDMoD
METRICS ON DEMAND

Benefits for Stakeholders

- **PI and End User**
  - Account management, resource selection, application tuning, improved throughput
- **Systems Administrator**
  - System diagnostic and performance optimization, application tuning
- **Computational Scientist and Support Specialist**
  - Tool to facilitate work with end-users to improve job performance and throughput
- **CI Center Director**
  - Comprehensive resource management and planning tool
  - Return on Investment Metrics
- **External Reviewers**
  - Tool for data driven review for verification of best practices and project goals
- **NSF Senior Leadership**
  - Measure the effectiveness of supported programs
  - Inform deployment of future systems to fulfill unmet need

XDMoD
METRICS ON DEMAND

Data/Services Available to ACCESS Tracks

- Job accounting data
- Application performance data
- Allocations data
  - Users, Resources
- Gateways usage data
- Quality of Service data
  - Resource providers
  - Applications
- Job efficiency data
- User data
  - Job efficiency, usage, application performance, ……
- Resource workload analysis
- Data analytics framework
- Networking data
- Workflows

XDMoD
METRICS ON DEMAND

So What's New?

- **ACCESS MMS Tasks**
  - Data Analytics Framework
  - CI Simulator
  - Monitoring CI Ecosystem
  - Application Power Monitoring
  - Network Integration
  - Value Analytics Realm
  - Open XDMoD Development
  - ACCESS Auditing
  - Service Model

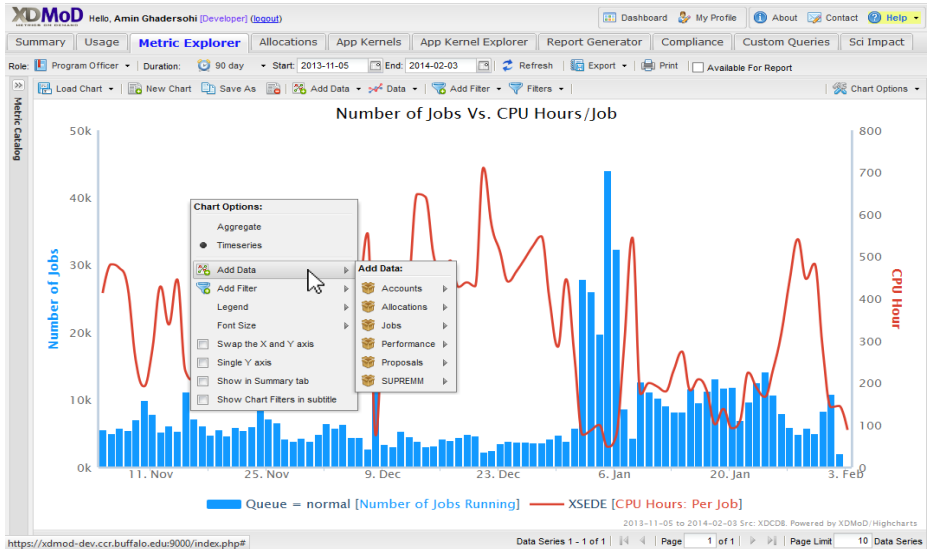

XDMoD METRICS ON DEMAND

Data Analytics Framework

- Provide analytic framework for direct access to rich depository of performance and utilization data in XDMoD data warehouse
- No reliance on XDMoD web portal
- GUI Built on Jupyter notebooks
- Python and R interfaces for automated workflows and expert users

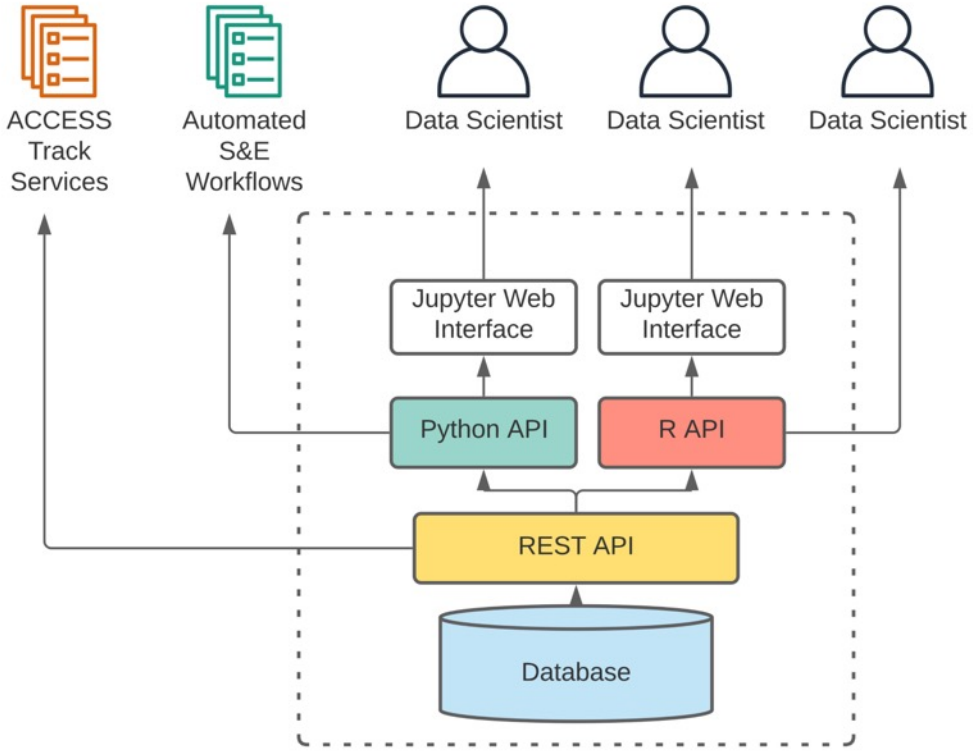

Data Analytics Framework Prototype

My Interactive Sessions - CCR | XDMoD Analytics - Jupyter No

ondemand.ccr.buffalo.edu/node/cpn-k08-34-02.cbls.ccr.buffalo.edu/39813/notebooks/XDMoD%20An...

jupyter **XDMoD Analytics** Last Checkpoint: 2 minutes ago (autosaved) Logout

File Edit View Insert Cell Kernel Widgets Help Trusted Python 3

Run Code

In [15]:
```
# Use the python API to get timeseries usage data from the OnDemand data realm in XDMoD
with xdw.DataWareHouse('https://metrics-dev.ccr.buffalo.edu:9002') as x:
    ondemandusage = x.timeseries(start = '2021-01-17', end = '2021-02-17',
                                 realm='OnDemand', metric='user_count', dimension='none')
```

In [33]:
```
# Add a Day name column
ondemandusage['Weekday Name'] = ondemandusage.index.strftime("%w\n%a")
```

In [45]:
```
ax = ondemandusage.boxplot(by="Weekday Name")
ax.set_ylabel('Number of Active Users')
plt.suptitle('')
```

Out[45]: Text(0.5, 0.98, '')

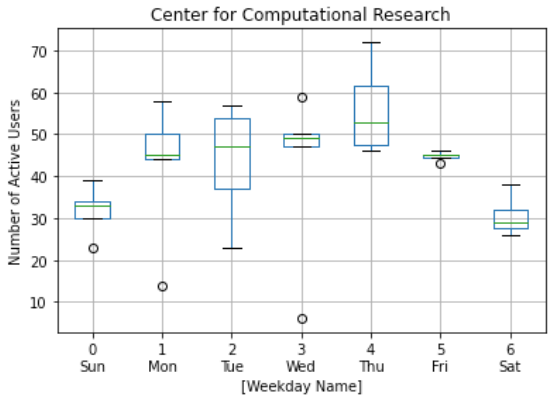

NSF

IoD DEMAND

CI Simulator

- **Develop a CI Simulator to model the NSF CI ecosystem**
- **Can be used to predict the response CI ecosystem to:**
  - New CI systems
  - Changes in the operational parameters of existing systems

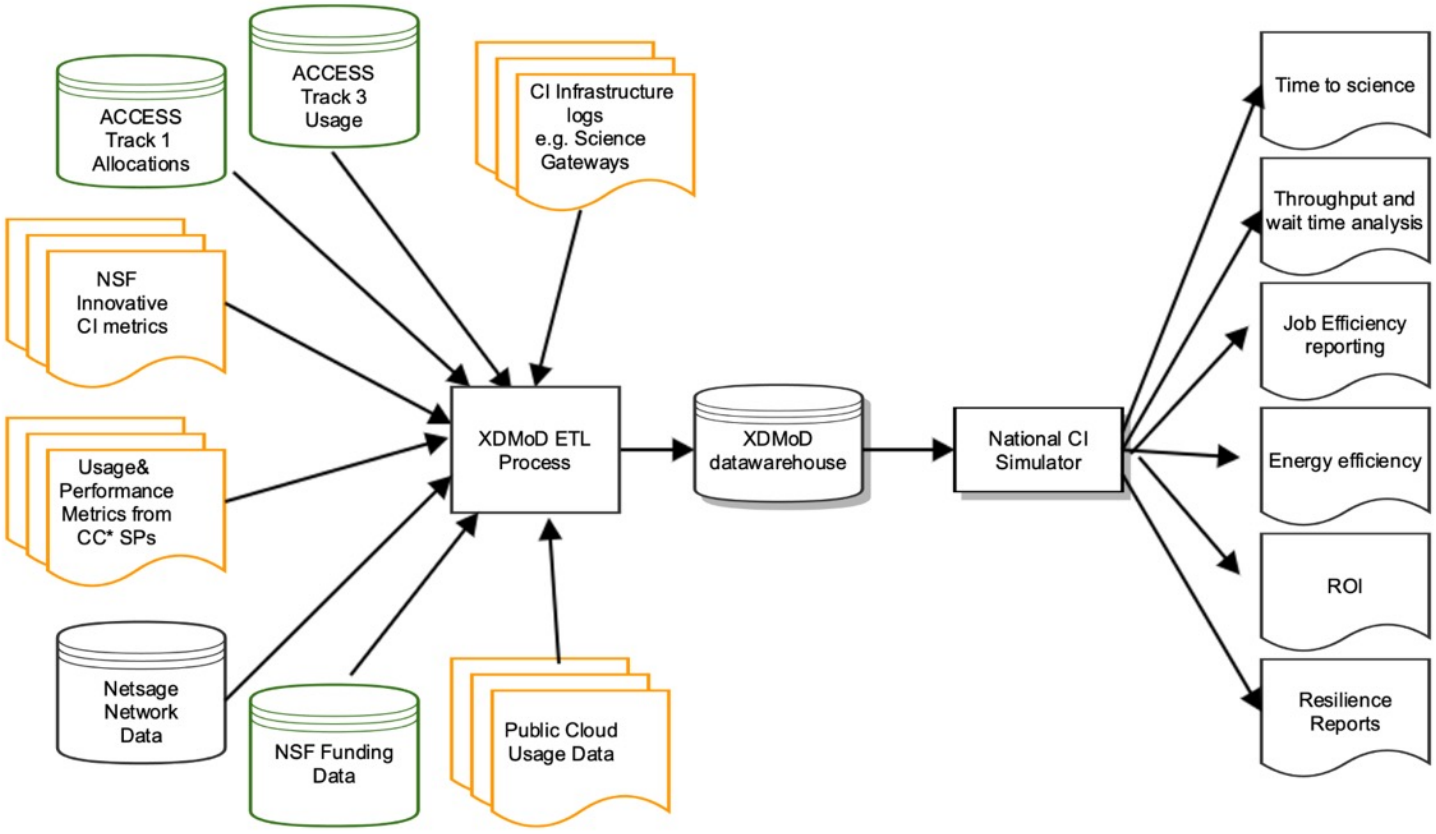

CI Simulator

- **Specific questions can be posed and the CI simulator can provide a quantitative estimate of response of CI ecosystem**
- **Example use cases**
  - Which of the following options would have the greatest impact on reducing the time to science:
    - one large resource,
    - two smaller resources,
    - or additional public cloud resources?
  - Simulating the relative impact of increased gateway resources versus MRI or CC* resources, which can then be used to guide funding decisions.

XDMoD
METRICS ON DEMAND

Monitoring the CI Ecosystem

- **Current XDMoD monitoring framework employed to better understand national capacity and capability class CI systems**
- **However, this represents a continually shrinking fraction of the national computational CI ecosystem which includes:**
  - campus-based systems, high throughput computing, science gateways, and public clouds

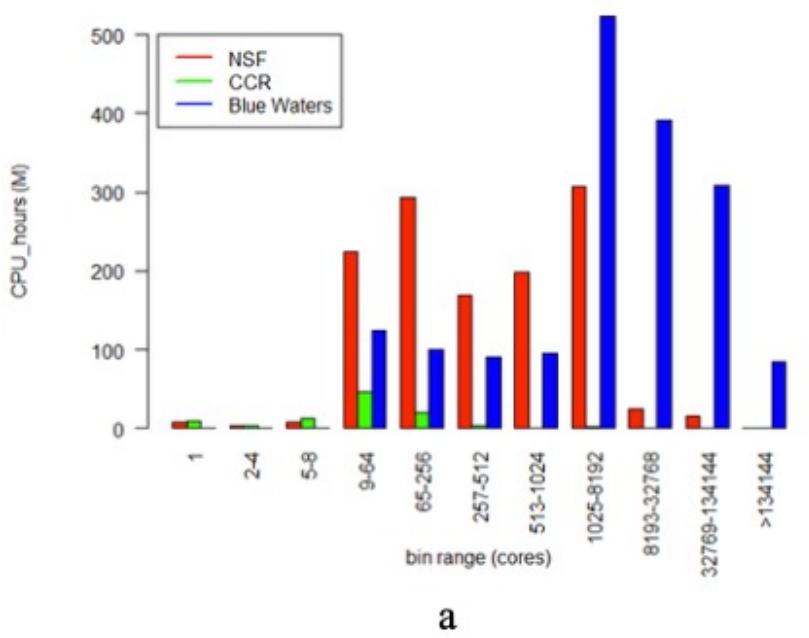

a

b

Monitoring CI Ecosystem

- **Development of a monitoring system for CI computational ecosystem**
  - national-level CI
  - NSF-funded public cloud usage
  - campus-level CI
  - network utilization
  - workload-type systems (OSG and Gateways)

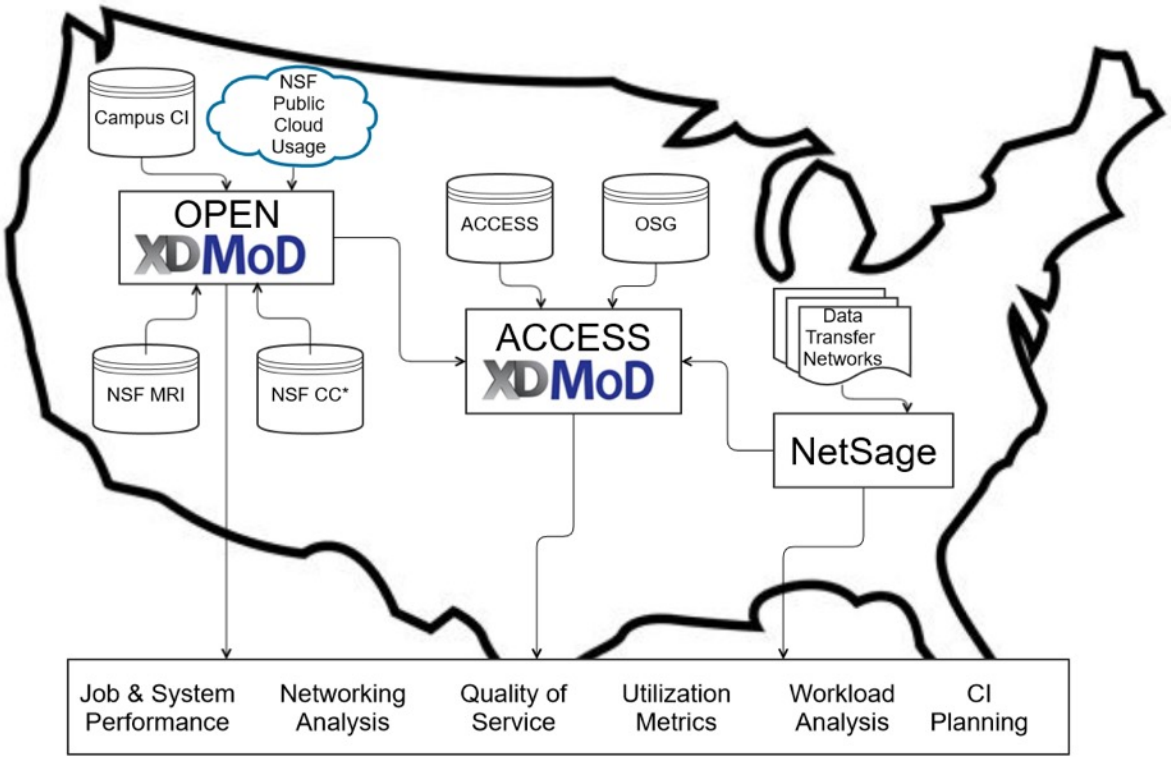

NSF National Computational CI

XDMoD METRICS ON DEMAND

Application-Level Power Monitoring

- **Provide a detailed workload analysis of time to solution versus power consumed for applications**
- **Useful for new and emerging technologies**
  - Will the energy savings realized by a more energy efficient architecture be offset by longer running times for the job mix?
- **Application developers can tune their applications to achieve better energy efficiency on different architectures**

XDMoD
METRICS ON DEMAND

Application-Level Power Monitoring

- **Example: Energy efficiency comparison for NAMD and GROMACS**
  - CPU - Intel Gold-6230,
  - GPU-NVIDIA V100
  - Raspberry Pi4 (ARM-based)

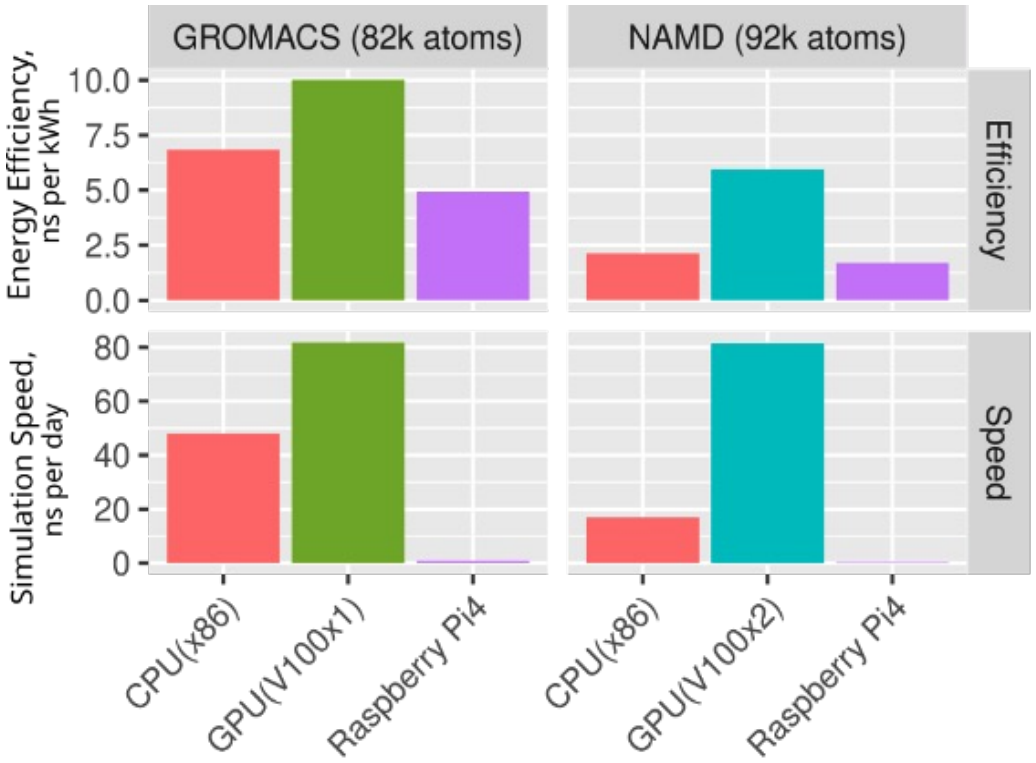

GPU versions are faster and more energy efficient

XDMoD METRICS ON DEMAND

Application Performance Diagnostics (Case Western Reserve University)

- **Effective application monitoring requires full stack monitoring from the front-end, user experience, to the back-end CI to provide complete visibility into every aspect of application performance**
  - Various layers such as infrastructure, network, logs, containers, databases, and clouds.
- **Develop machine learning based tools to monitor application performance metrics**
  - Develop framework with lightweight agents for real-time
  - Train models using time-series data from XDMoD
  - Target two widely used applications

XDMoD
METRICS ON DEMAND

Network Data Integration

- **Indiana University – Dr. Jennifer Schopf**
- **Motivation: Data movement increasingly important for computation, simulation and large-scale experimental facilities**
- **Use NetSage to track data movement**
  - Collect SNMP, flow data, and perfSONAR active testing data
- **Use XDMoD to correlate data movement with computational workload**

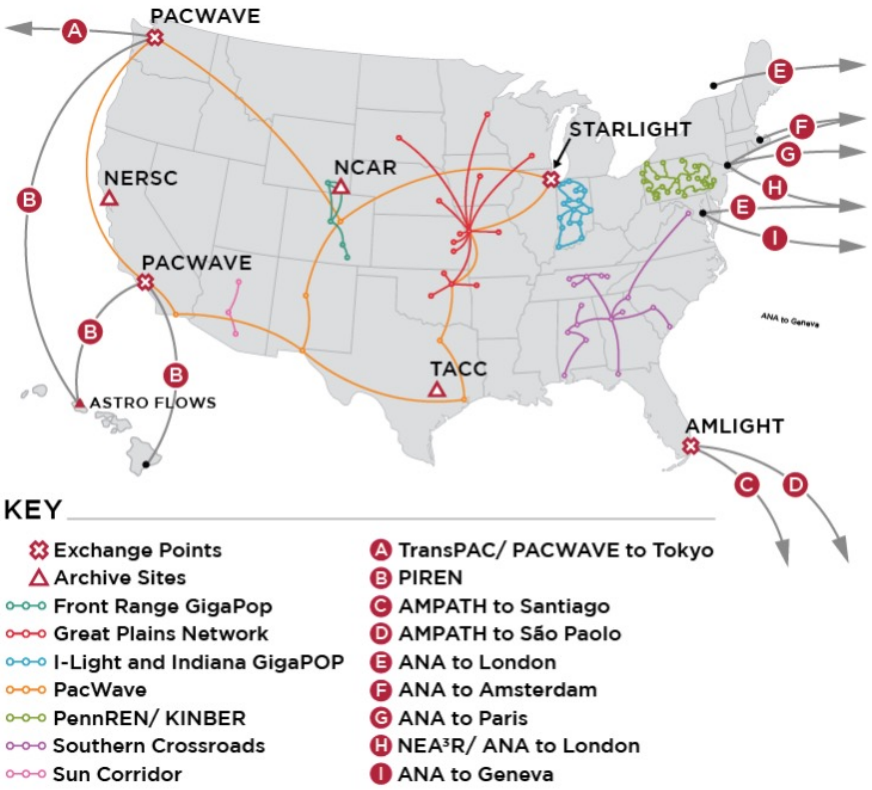

XDMoD
METRICS ON DEMAND

Value Analytics Realm in XDMoD

- **Help provide ROI metrics to support campus investment in CI**
- **XDMoD Value Analytics Realm will correlate system usage with PI funding**
- **Proposed sources of campus funding data**
  - university sponsored programs office
  - direct feeds from NIH, NSF, etc

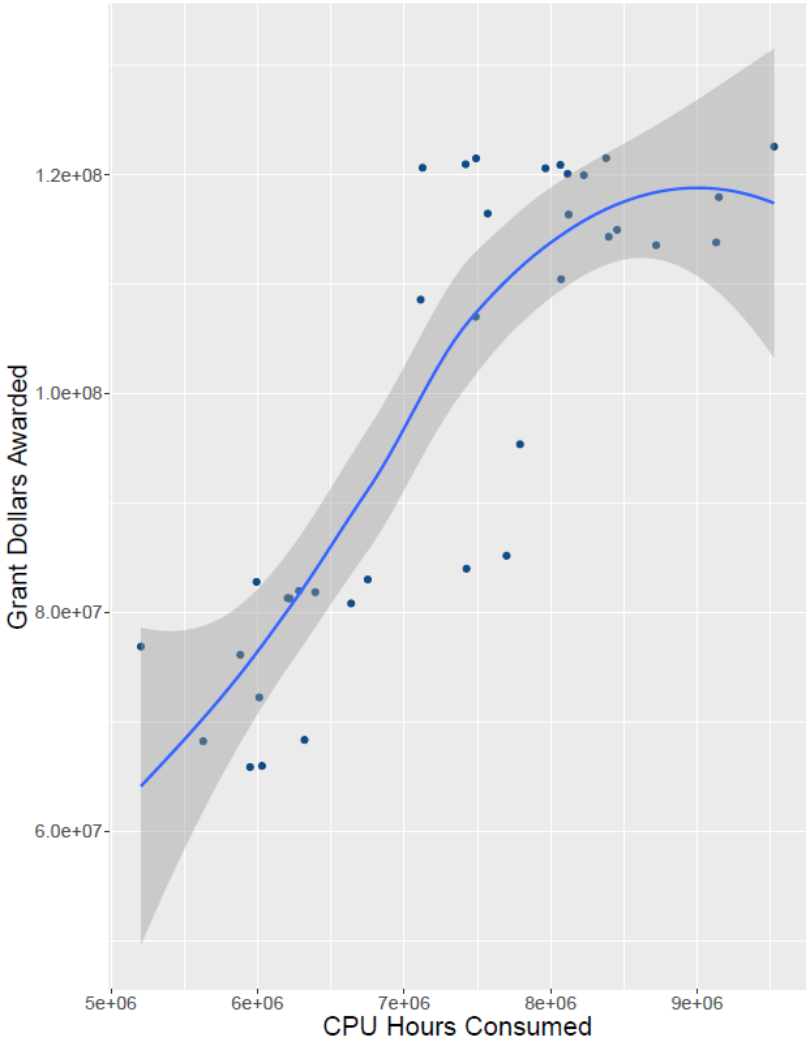

XDMoD
METRICS ON DEMAND

XDMoD & Open XDMoD Development

- **TACC_Stats development**
  - keep current, add new technologies and new features
- **Additional features including:**
  - Added GPU support
  - Expanded storage metrics
  - On Demand integration
  - ColdFront integration
- **Application Kernel Development**
  - Support for cloud environments
  - Support for workflows
  - Develop Network usage Application Kernel

XDMoD
METRICS ON DEMAND

ACCESS Auditing

- **ACCESS MMS will continue carrying out workload analyses on capability and capacity class systems as directed by NSF**
- **Similar to analyses carried out by XMS on Blue Waters and NSF Innovative Resources**
  - Blue Waters Final Report: https://arxiv.org/abs/1703.00924
  - NSF Innovative Resources Final Report: http://arxiv.org/abs/1801.04306

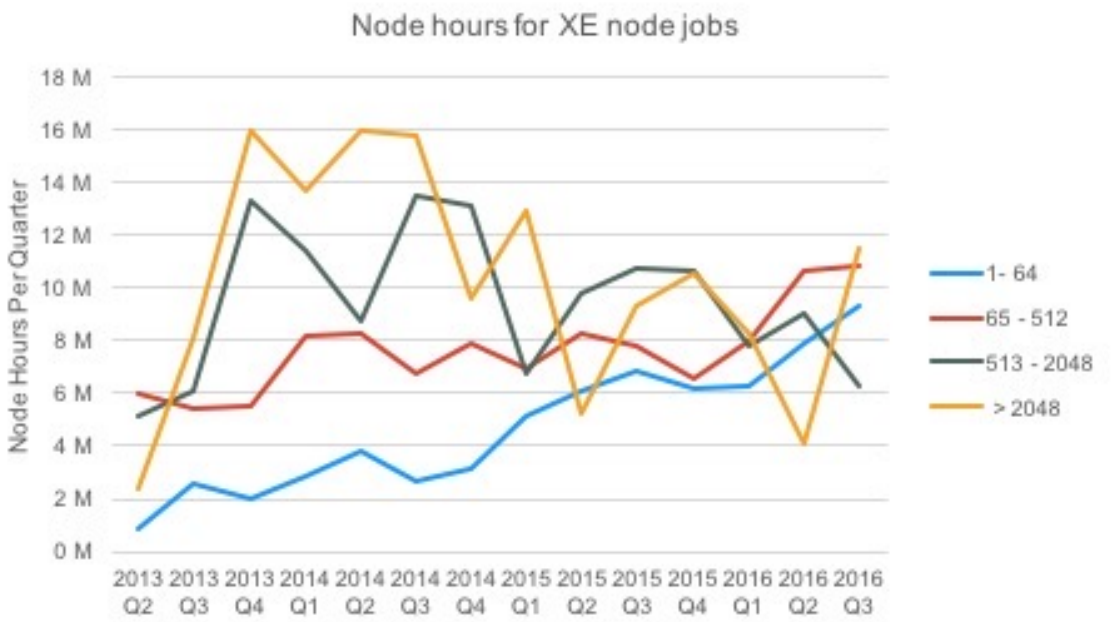

Service Model – Data Service

- **Embrace the FAIR principles for scientific data**
  - **Findable**
    - The ACCESS XDMoD data warehouse will have a searchable metadata index that is both human and machine readable, publishing data and service assets
  - **Accessible**
    - Data warehouse assets will be available through standard REST calls. Metadata will be publicly available (Findable)
  - **Interoperable**
    - Data warehouse exports will made available in commonly used machine readable formats including JSON and CSV
  - **Reusable**
    - MMS encourage reuse of XDMoD's data assets throughout the research community and will continue to do so

XDMoD
METRICS ON DEMAND

Service Model - ACCESS Tracks

- **Track 1 Allocations Service**
  - Historical CI usage for individual users/projects and overall CI resource usage
- **Track 2 End User Support Service**
  - Provide detailed information about end users historical jobs, workflow and network usage to help support and respond to customers' queries.
- **Track 3 Operations Service**
  - Provide information on ongoing ACCESS operations including QoS metrics
- **Track 5 Technology Translation**
  - Work with Track 5 and ACCESS ACO to determine and implement the proper reporting metrics
- **ACCESS ACO**
  - Provide information to support its management, oversight and coordination of the ACCESS awardees

XDMoD
METRICS ON DEMAND

Federated XDMoD

- **Federated XDMoD allows integration of multiple Open XDMoD instances into a single monitoring and reporting entity**
- **Can monitor each instance separately and/or look at the federation totals**

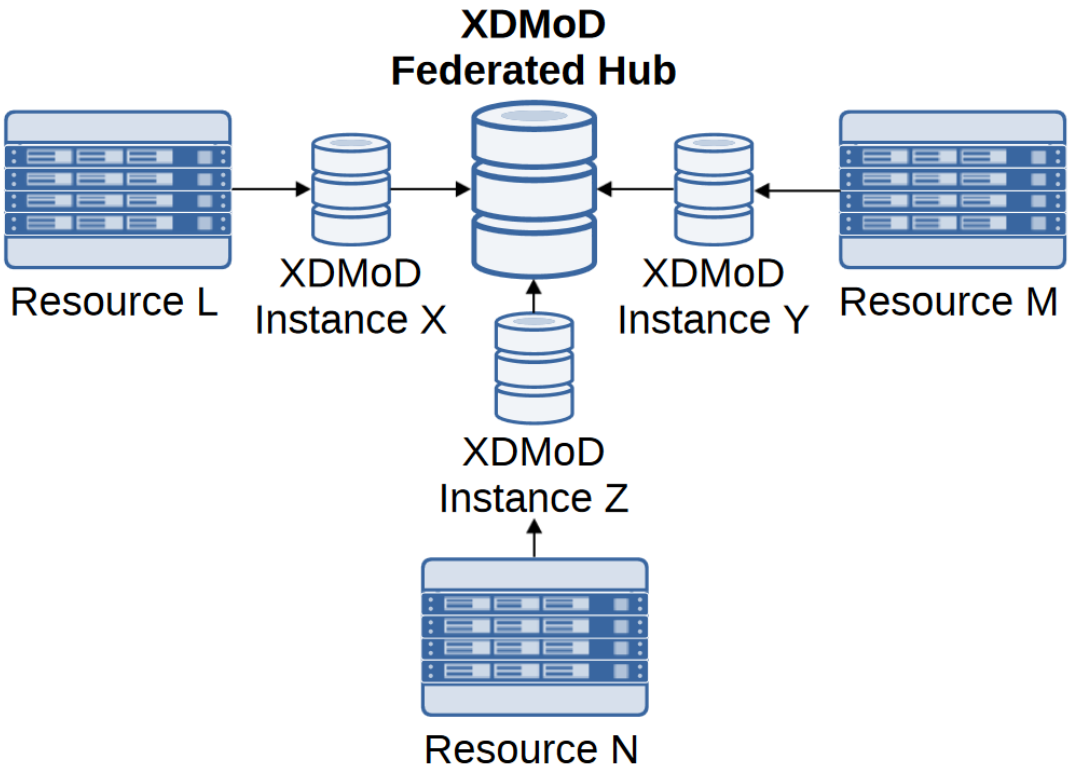

ACCESS MMS Service Model Schematic

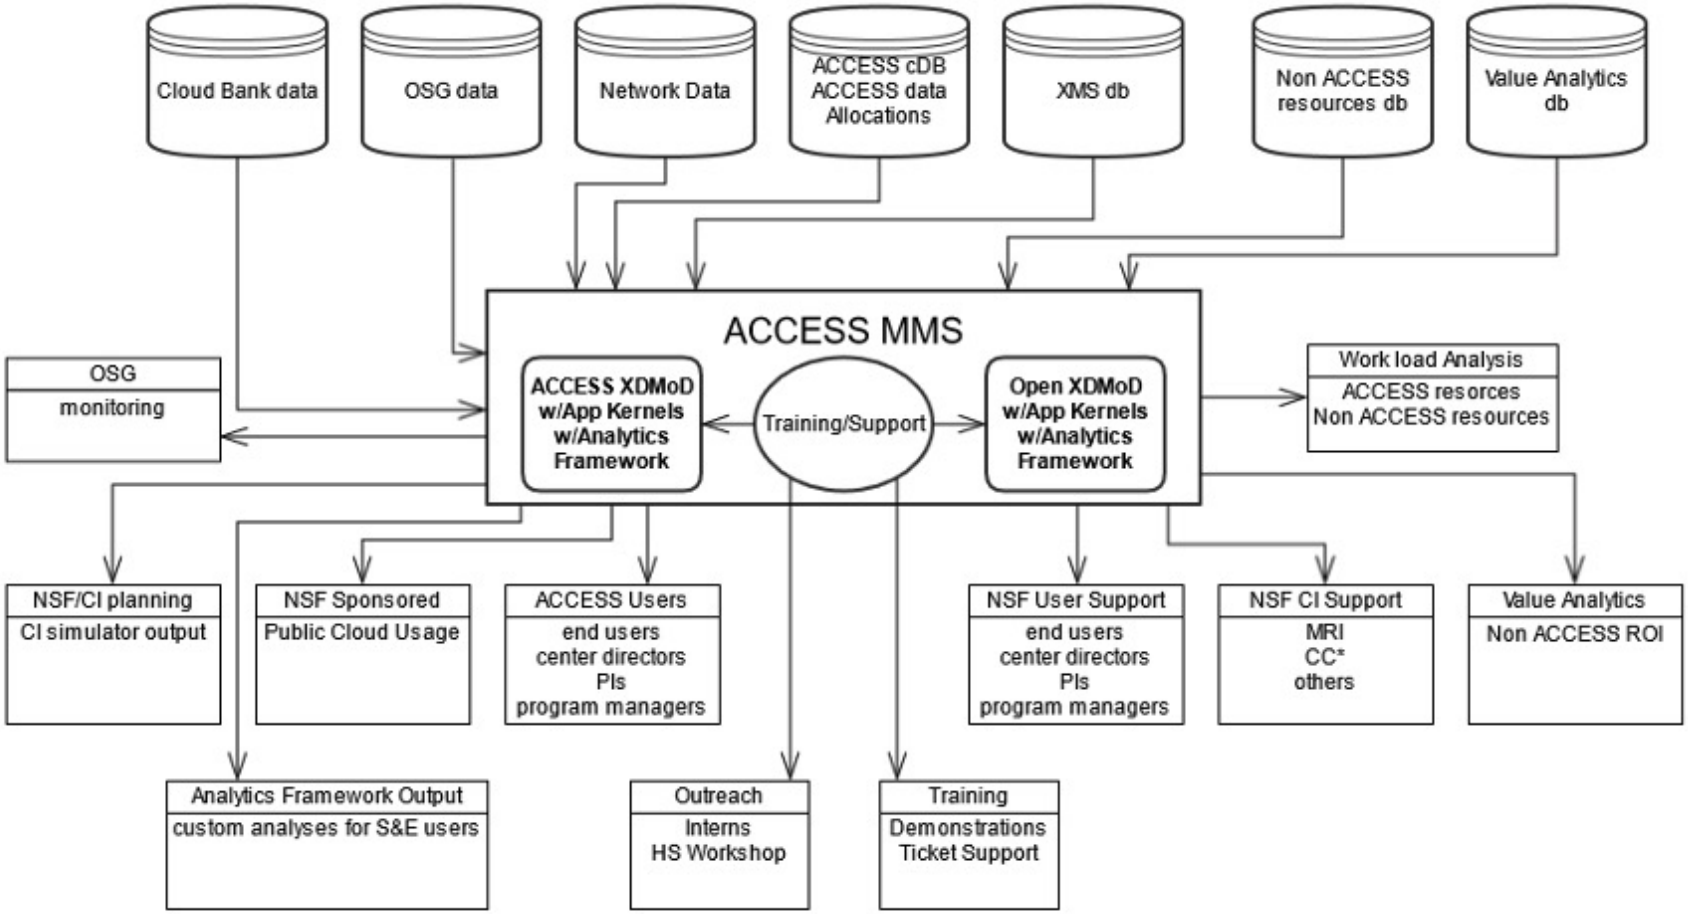

XDMoD
METRICS ON DEMAND

Discussion

XDMoD
METRICS ON DEMAND