Challenges for MODA in a leadership public data repository

Patrick Widener, Ph.D
Alex May
Tatiyanna Singleton
Olga Kuchar, Ph.D

Data Lifecycle Technologies Group
National Center for Computational Sciences
Oak Ridge National Laboratory

Presented by Christian Engelmann, ORNL

ORNL is managed by UT-Battelle LLC for the US Department of Energy

This work used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725

Introduction

- **Constellation** is a large-scale data repository hosted at the Oak Ridge National Laboratory's Leadership Computing Facility (OLCF)
- Data repositories are critical components of data management
  - Computational science at OLCF produces and consumes immense volumes of data
  - Sharing data is only practically achievable for DOE user facilities through repositories like Constellation

Why this is a MODA story

- Constellation is hosted by OLCF
  - Shared storage and network resources
  - OLCFs mission is oriented to computation
- We have to justify our use of shared resources
  - To date we have flown under radar
  - Changes in available resources and increased demand
  - This means collecting metrics
- We also must understand the data we steward
  - Avoid legal and operational problems
  - We view *curation* as a data analytics challenge, primarily due to scale

Why do we need Constellation?

- Dataset sizes/structures are in a separate equivalence class
  - Constellation is biased toward datasets with thousands of individual files and/or very large individual files
  - Total holdings > 7PB; largest dataset 3 PB; largest single file 17TB
- OLCF has large storage, but…
  - Mostly dedicated to scratch space for mod/sim codes
  - Heavy demand, limited metadata tooling, no external exposure
  - Storage allocated to projects with finite lifetimes

Constellation system architecture

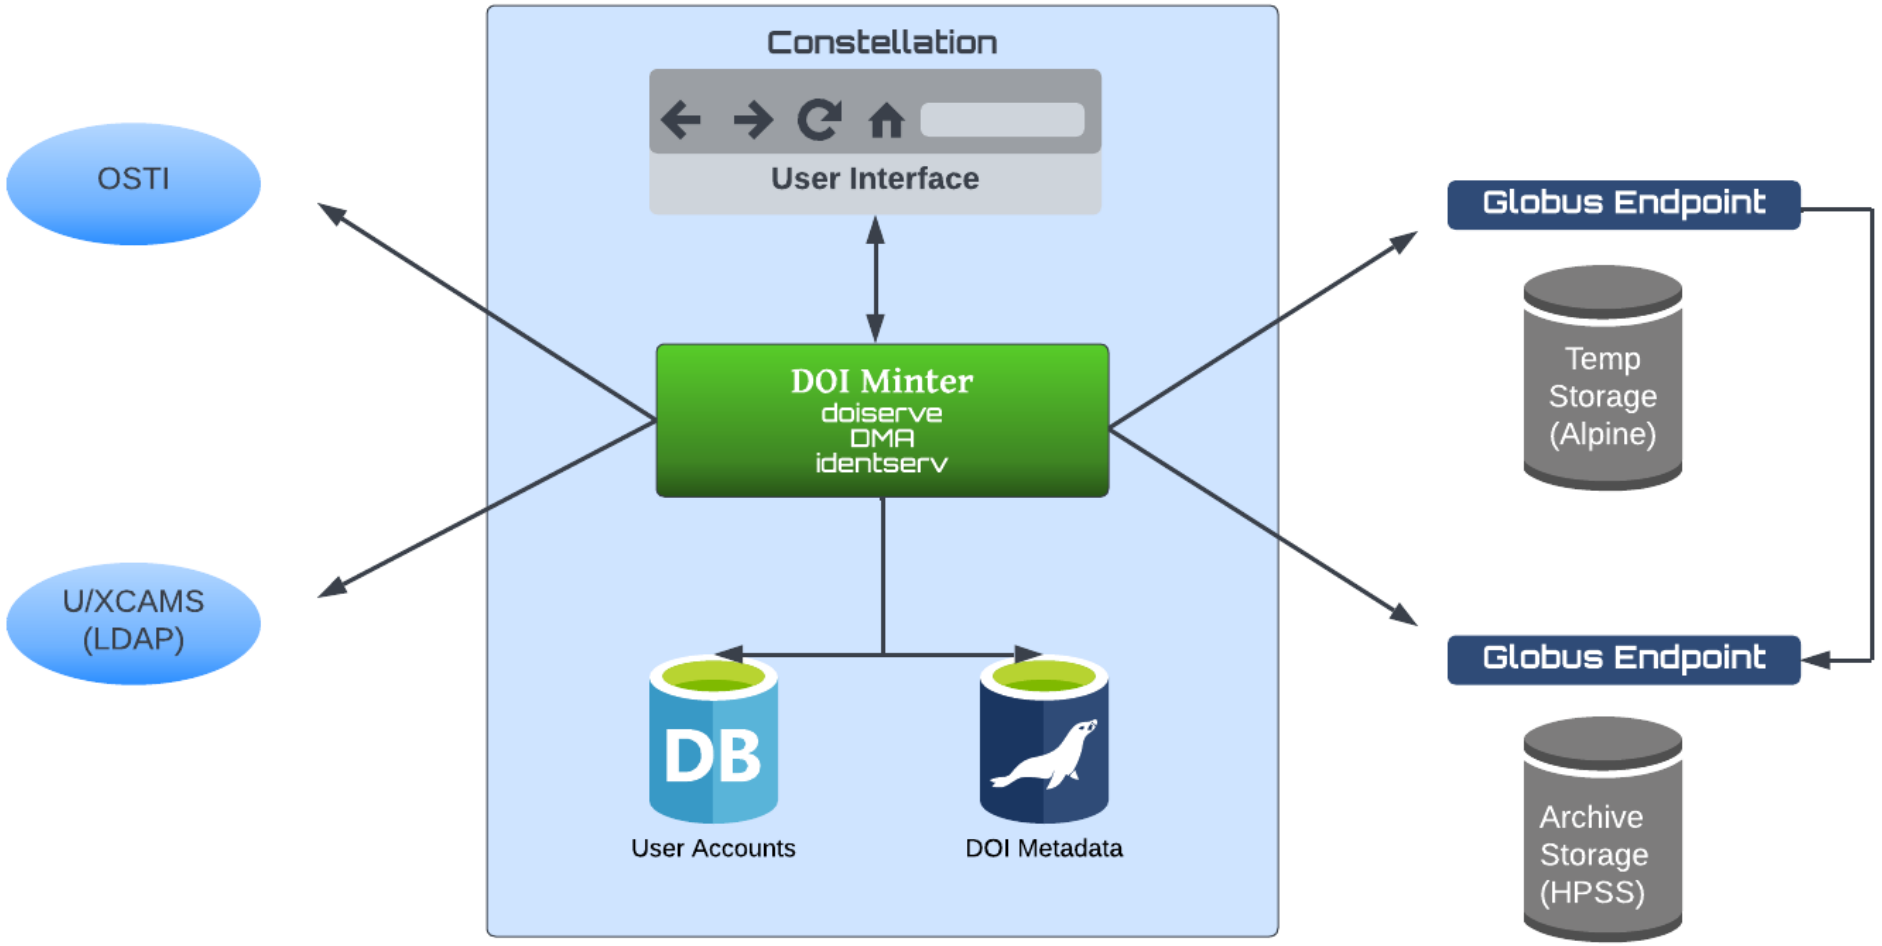

Constellation – User Perspective

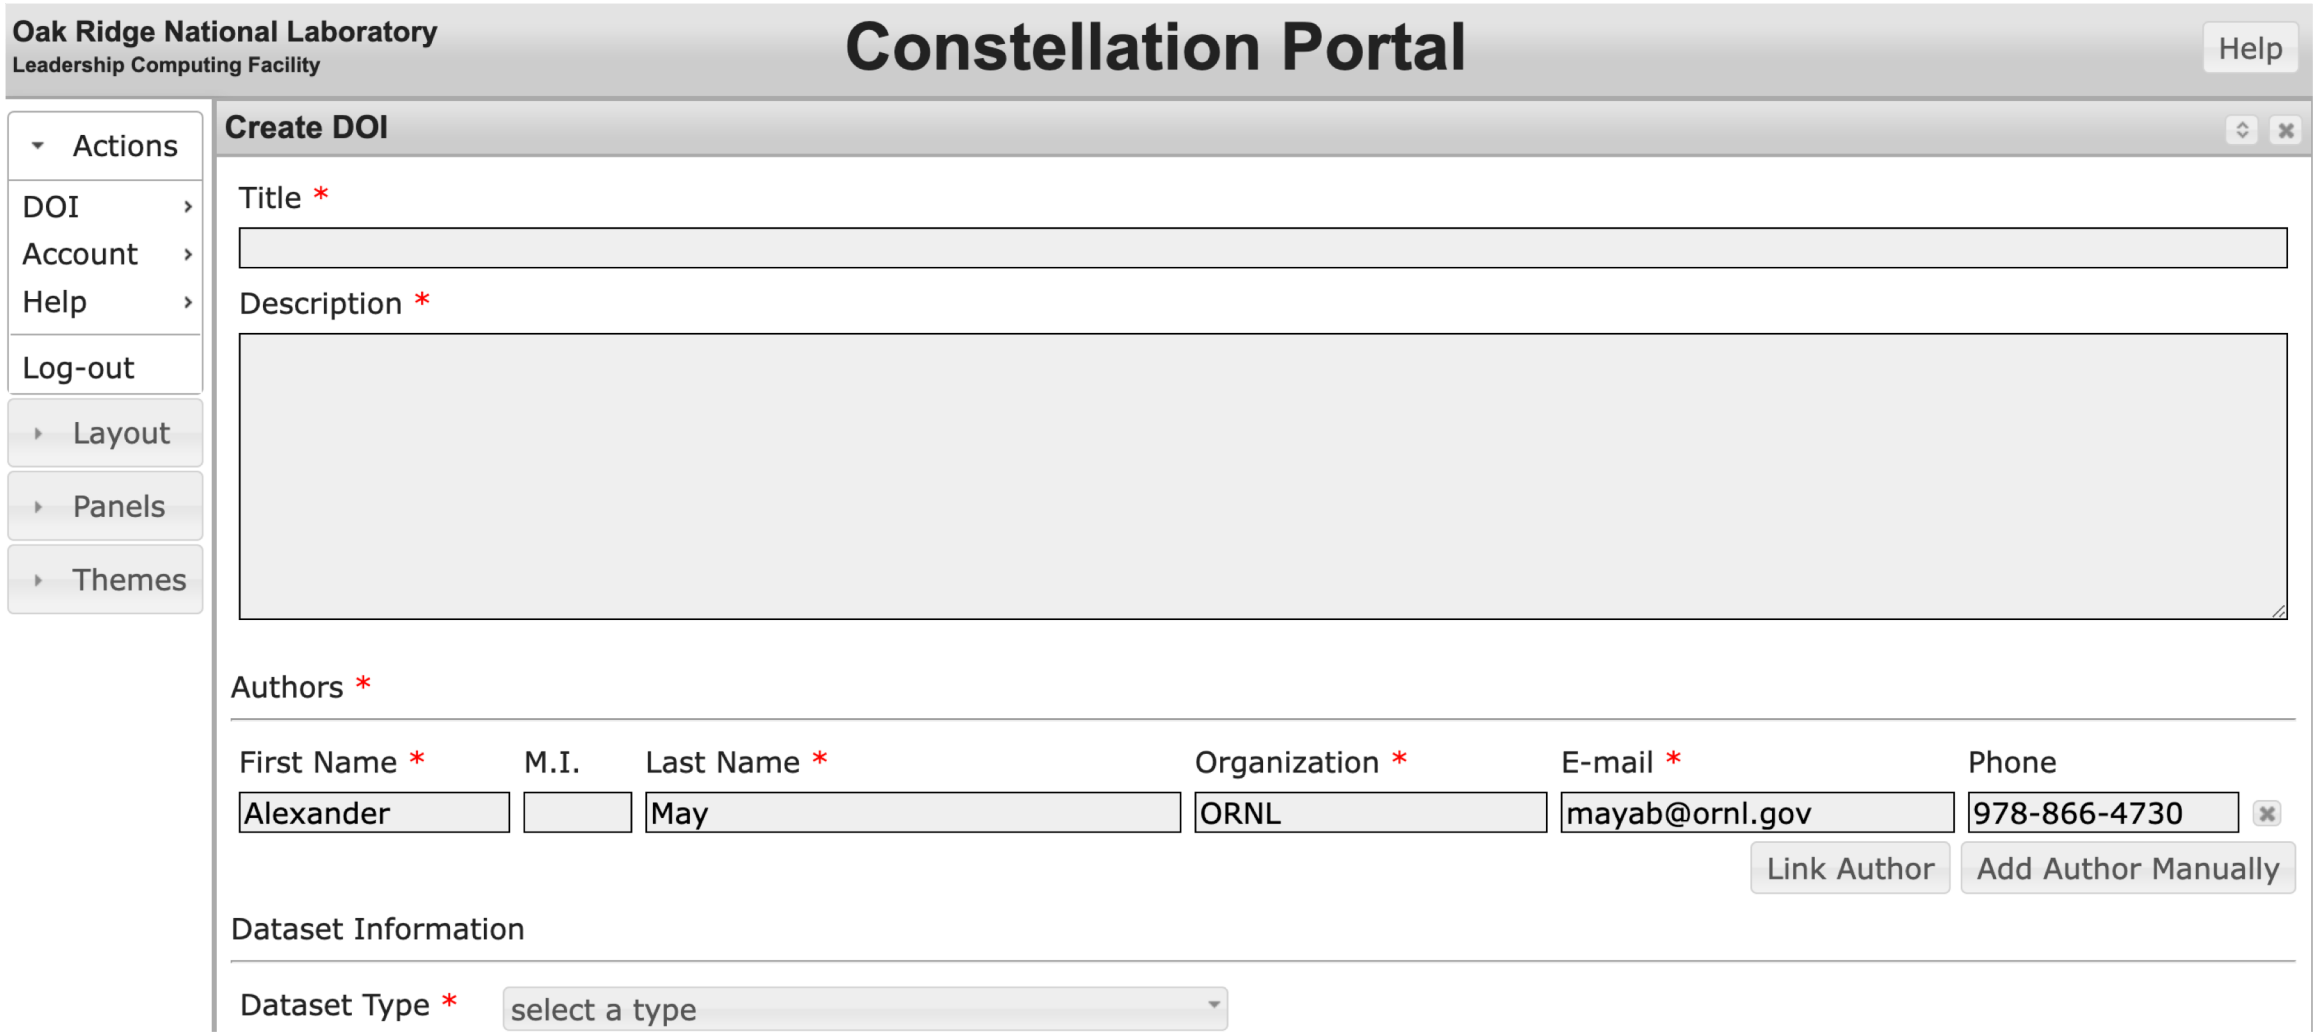

Constellation – User Perspective

General Information

Number:	10.13139/ORNLNCCS/1872748
Title:	Neutron computed tomography and high speed imaging of single hole gasoline direct injector
Description:	Data sets from neutron computed tomography and high speed neutron imaging of a single hole gasoline direct injector performed at HFIR CG-1D cold neutron imaging instrument.
Created:	6/17/2022, 1:07:22 PM
Published:	6/21/2022, 8:48:43 AM

Authors

Wissink, Martin wissinkml@ornl.gov

Dataset Details

Dataset Type:	ND Numeric Data
Subjects:	32 ENERGY CONSERVATION, CONSUMPTION, AND UTILIZATION;42 ENGINEERING
Keywords:	neutron imaging, tomography
Product Nos.:	
Software Needed:	

Dataset Files

- High speed imaging
- README.txt

Constellation *in situ* at OLCF

- Constellation has no dedicated resources
  - Not even storage
- Application containers run on a shared OpenStack cluster
- Temporary storage for intake is a Lustre PFS
  - Shared; heavily used; provides scratch space for Frontier jobs
- Long-term storage on tape
  - Currently HPSS – shared across OLCF
  - Migrating to IBM Spectrum Archive – also shared across OLCF
- These are production resources tightly controlled by operations groups

Monitoring challenges

- Monitoring OpenStack is straightforward
  - But Constellation is not a compute-bound service
- Instrumenting production storage resources is not straightforward
  - Basic statistics are available: disk space used, effective bandwidth
  - Globus giveth and taketh away
    - Abstracts away administrative, network, reliability issues
    - Obscures performance information, conflates it with other contributing factors
  - Advanced monitoring (e.g. energy cost of storage) will be difficult
    - Hope for HW/SW vendor support
    - As storage appliances proliferate this problem gets worse

What metrics can we collect?

- Primary "scientific impact" metric is download requests
  - Tracked in SQL database
  - Microsoft BI for reporting
- We report out common and not-so-common items
  - Dataset popularity
  - Correlation of downloads with project timelines and publication activity
  - Denial-of-service and intrusion attempts
    - ORNL is a popular target and Constellation is part of the attack surface
- This is work-in-progress

Downloads per year

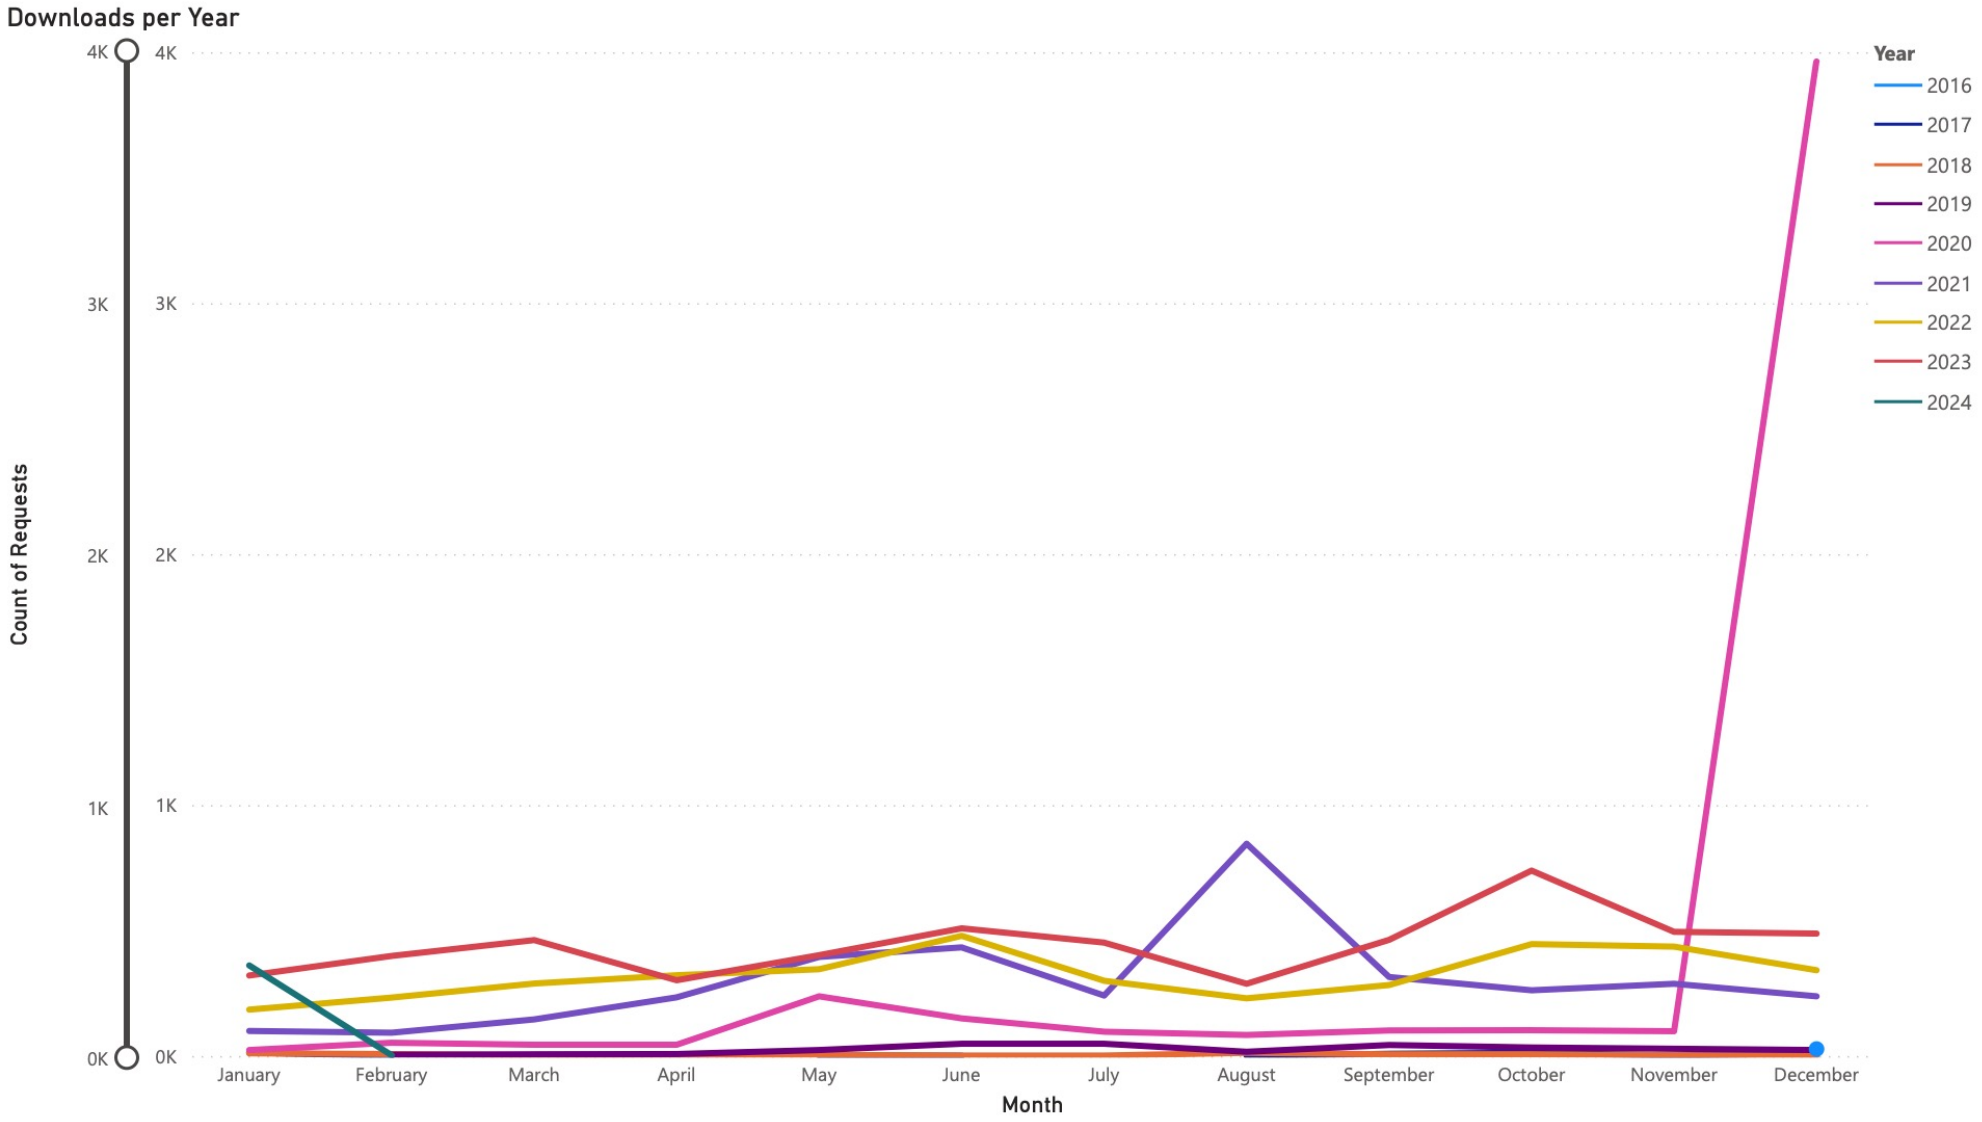

OAK RIDGE National Laboratory | 80

Open slide master to edit

Curation: Understanding the data

- Typical curation: inspect, attach metadata and descriptions
  - Protects against content-based issues (PII, illegal material)
  - Best-effort by humans, random sampling of large datasets
  - Not just best-practice; required by DOE and US Federal regulation
- Constellation data has 3 of the 4 Vs of data
  - Volume, Velocity, Variety
  - Sources range from ORNL scientific instruments to OLCF academic users
  - Data types range from numeric data to images to text, often combined
- *Infeasible for any reasonable number of curators to manually inspect datasets at our scale*
- Automation and AI support will be necessary

What kinds of support will be necessary?

- In general, size prevents curators from "beholding" an entire dataset at once
  - Summarization and visualization tools to assist in understanding
  - Quickly alerting curators to issues before datasets are formally accepted
  - Sampling of datasets to avoid transferring data which will need to be modified/deleted
    - Transfers of multi-PB datasets can take weeks
- Automated metadata extraction/creation
  - Required for repository standards, FAIR data
- These are open research areas for us and others in the data curation community

Curation Toolkit

- We are researching a collection of interoperating components
  - AI-enabled sensitivity platform to highlight problematic content
  - Dashboard which presents file formats and issues needing human review
  - Automated metadata extraction
    - Alignment of extracted metadata to existing controlled vocabularies
    - Necessary for federated discovery according to DOE standards
- We'll combine these tools with analysis of metrics
  - Automated decisions about which of the available storage resources to use
  - Resource planning as OLCF evolves

Conclusion

- Constellation's challenges will sharpen
  - Scientific data management becomes a more widespread issue
  - Dataset sizes continue to increase
    - Already have 6 PB additional data incoming this year
  - Interoperability with other OLCF, DOE initiatives
    - INTERSECT, IRI – management of data from instrument to storage
    - DOE High Performance Data Facility – announced but yet to take shape
  - Continued need to operate in OLCF without specialized resources
- We are working on MODA to
  - Help justify allocations of shared resources
  - Assist with intractable curation requirements
  - Plan for the future: the shape of OLCF will likely change