Data Management Systems for Monitoring Workloads

Data Over Distance Symposium

Jessie Gaylord

July 19, 2018

LLNL-PRES-754789
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Outline

1. Introduction
2. Data Management goals
3. Current projects
4. Conclusion

Geophysical Monitoring Program

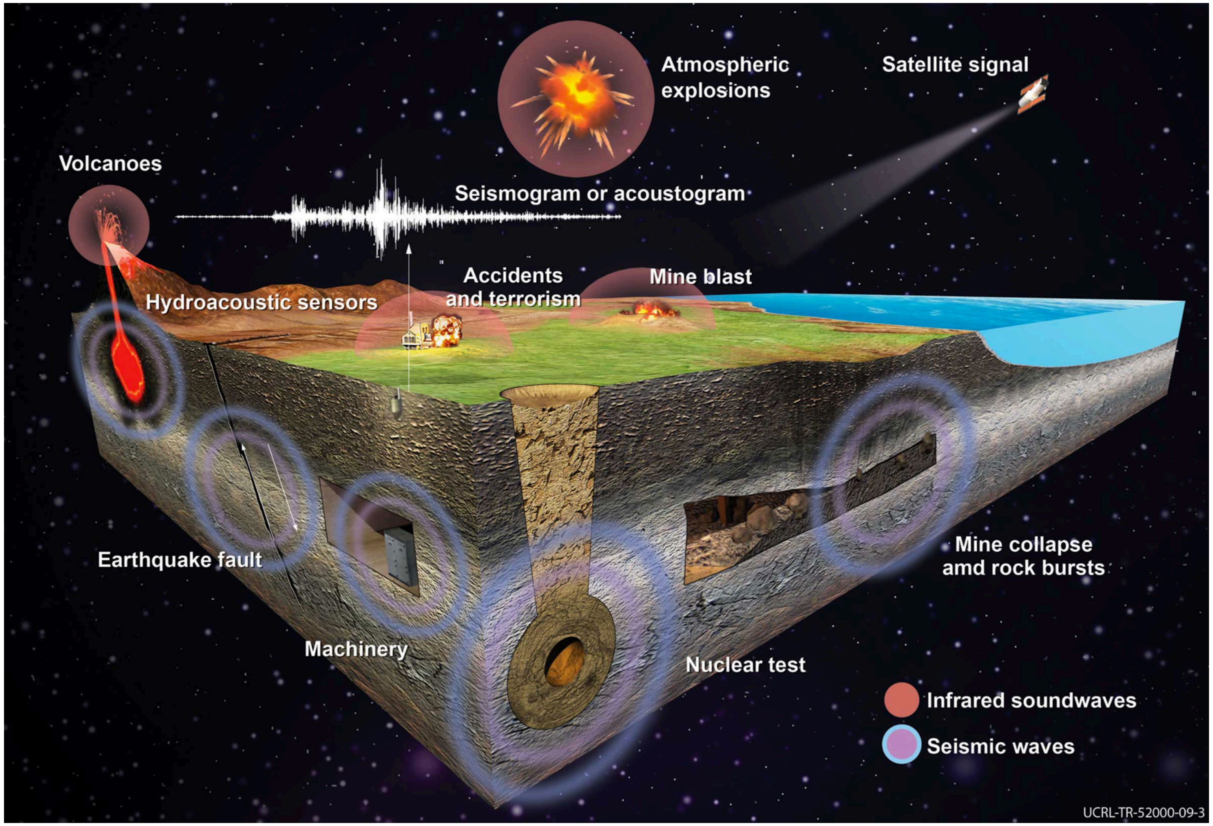

Justin Barno

Doug Dodge

Jessie Gaylord

Doug Knapp

Kathleen Li

Steven Magana-Zook

Stan Ruppert

Danielle Sullivan

Judy Thomas

- Supporting basic science, non-proliferation, nuclear monitoring, hazard analysis, and international outreach

Current Projects

- NA22 DNN R&D Ventures

MINOS

LYNM

LANL-LLNL-PNNL-SNL-MSTS-NNSA

- Internally funded containerized infrastructure to support LLNL research in Data Science

Why have Data Management?

- Data is a valuable asset
  - Quality
  - Curation and archiving
- Get the most out of research investments
  - Efficiency
  - Reuse and interoperability
- Enable reproducible research
  - Provenance
  - Traceability
- Support future efforts
  - Fully characterized data

Effective Data Management is critical to scientific success.

More Incentives

1. It supports critical science
2. Publishers require it
3. It is U.S. policy
4. Our sponsors require it
5. People will thank you

Science

SPRINGER NATURE

AGU

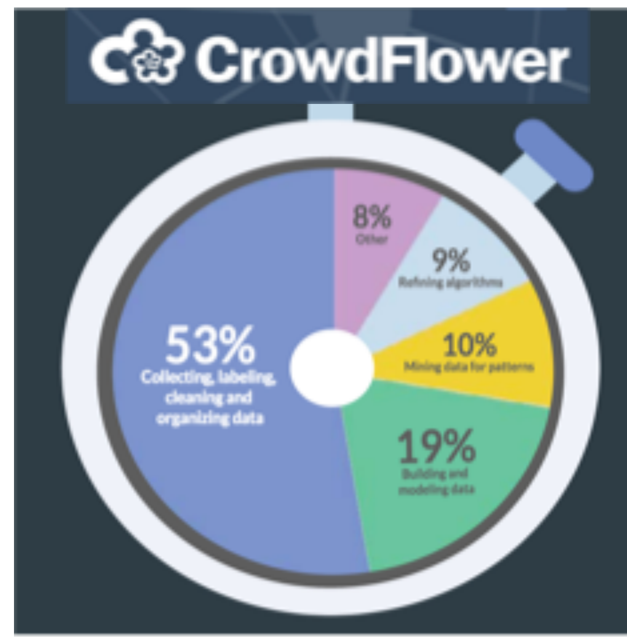

Monitoring Data Management Challenges

- Diverse datasets:
  - Parametric and measurement data
  - Wide variety of domains, formats, sizes, sampling frequencies, etc.
  - Batch and streaming data
  - Temporary field deployments and manual collections
  - Legacy data
  - Sensitive data
- Diverse users:
  - Multi-lab collaborations
  - Independent research groups
  - Multiple domain experts
- Diverse use cases:
  - Data discovery
  - Single mode analytics and data fusion
  - Novel research approaches

Data Management System Goals

- Careful metadata management
- Quality assessments for completeness and accuracy
- Tracking provenance and traceability
- Planning storage, organization, access, and security
- Short and long term data curation
- Support a Data Management Plan (DMP)

"Effective data management has the potential to increase the pace of scientific discovery and promote more efficient and effective use of government funding and resources."

- DOE Policy for Digital Research Data Management

Data Management Plan

- Determine the scope
  - Data sources
  - Data products
  - Potential for re-use
  - Data lifecycle
- Document the data
  - Description of data
  - How data were produced
  - File formats
  - Size
  - Source and POC
  - Stability
- Establish metadata standards
  - Persistent identifiers
  - Field glossary
  - Time and location synchronization
  - Provenance and traceability
  - Quality assurance and validations
- Make a file organization plan
  - Directory organization
  - File naming conventions
  - Encryption and compression
- Set policies for data access and sharing
  - Sharing platform
  - Access controls
  - Distribution
  - Sensitive data
  - Classification
- Plan archiving and preservation
  - Retention policy
  - Long term curation
  - Backup plan
  - Security
  - Storage platform
  - Support and maintenance
- Estimate costs and resources

Context

Conventional Monitoring Data Flow

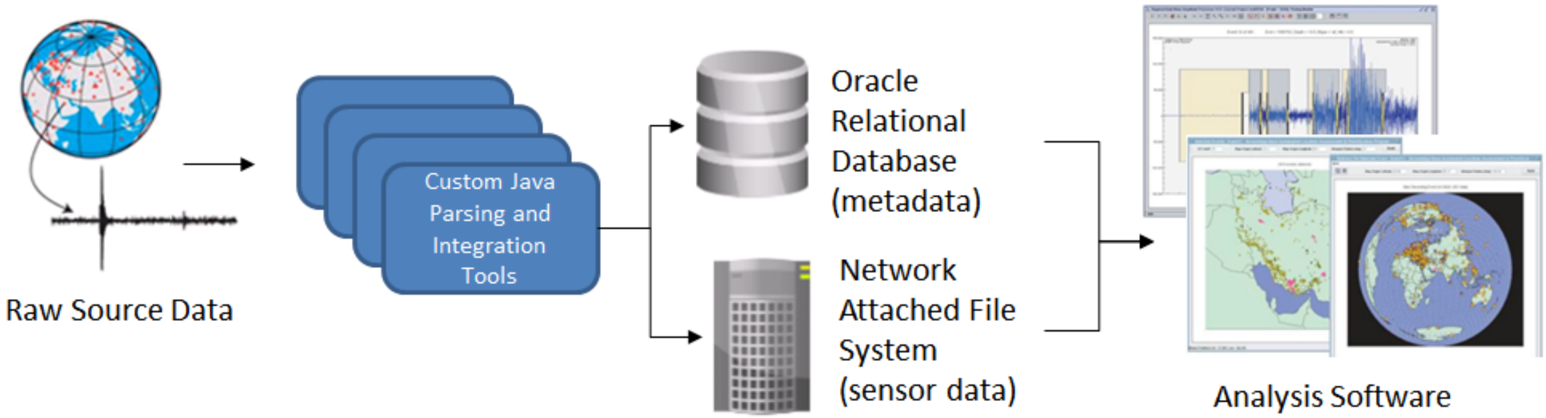

- Raw input files are parsed into:
  - RDMS for parametric and metadata
  - File system for sensor payloads
- RDMS records have pointers to sensor data files
- Analysis software queries the database to get the pointers and retrieves the sensor data for processing on client machine(s)

Ingestion rates are driven by:

- Increased data availability
  - Cheaper hardware
  - Increased coverage
  - Dynamic field deployments

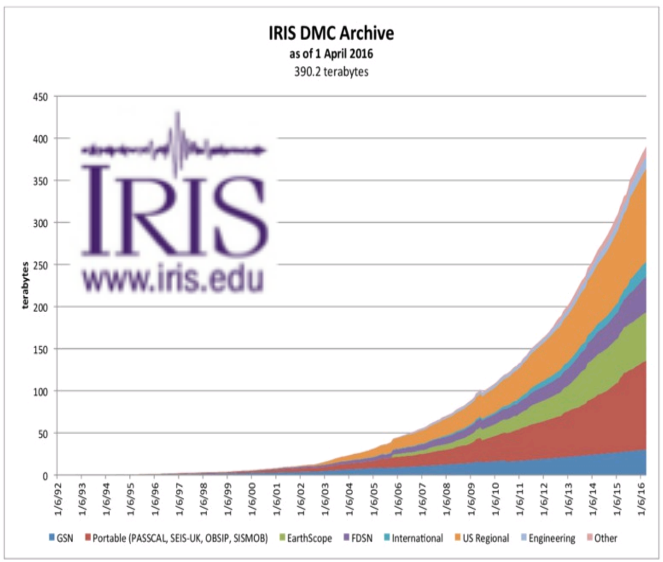

- Scientific need
  - Lower monitoring thresholds
  - Whole catalog analytics
  - Multi-physics data fusion
  - Automated data classification
  - Near real-time processing streams
  - Previously impossible research
  - Novel sensor platforms

Ubiquitous data enables new research and discovery.

“Big Data” Frameworks

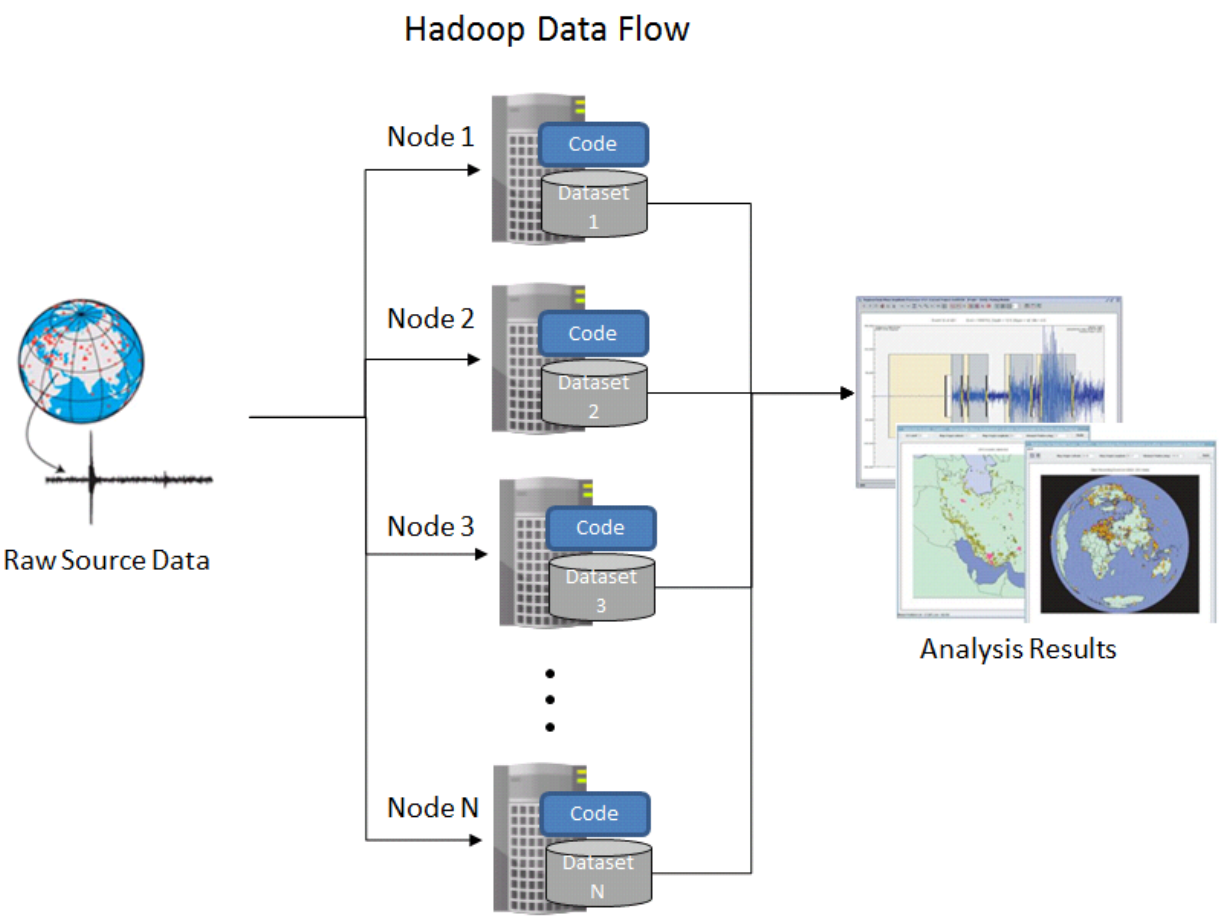

- HDFS file system distributes data over a cluster of commodity hardware
- Horizontally scalable
- Optimized for:
  - Massive datasets
  - Semi-structured text data
  - Streaming data
  - Data analysis workloads
- Data is processed locally on each node
- Reduces I/O
- MapReduce or Spark plus APIs for Java, Python, Scala, etc.

Prior Research

- 19x speedup on correlations of 300M waveforms

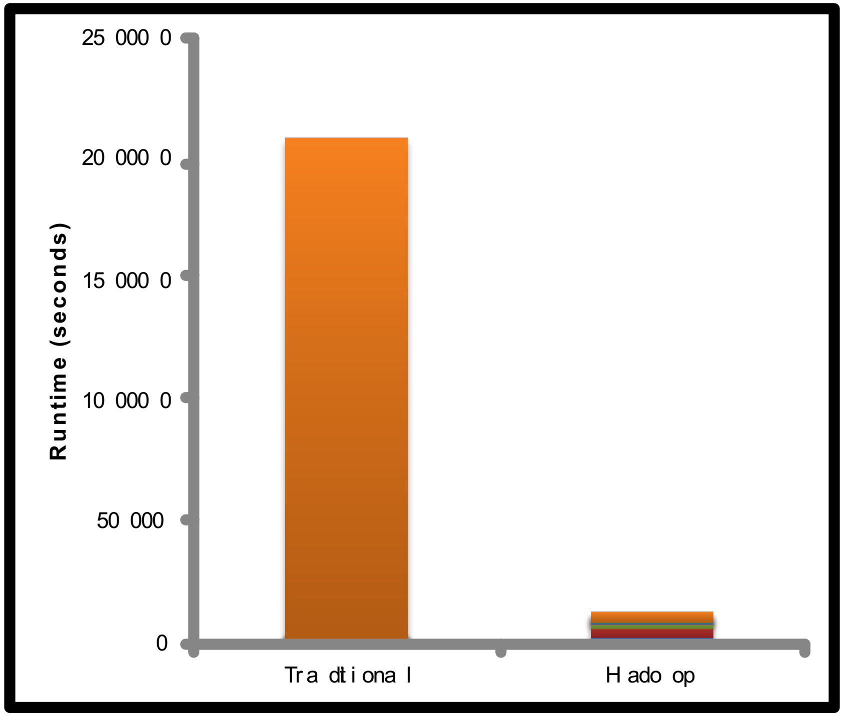

Addair, T. G., D. A. Dodge, W. R. Walter and S. D. Ruppert, Large-Scale Seismic Signal Analysis with Hadoop, *Computers and Geosciences*, May 2014

- Verified horizontal scaling on QC of 43TB of continuous data

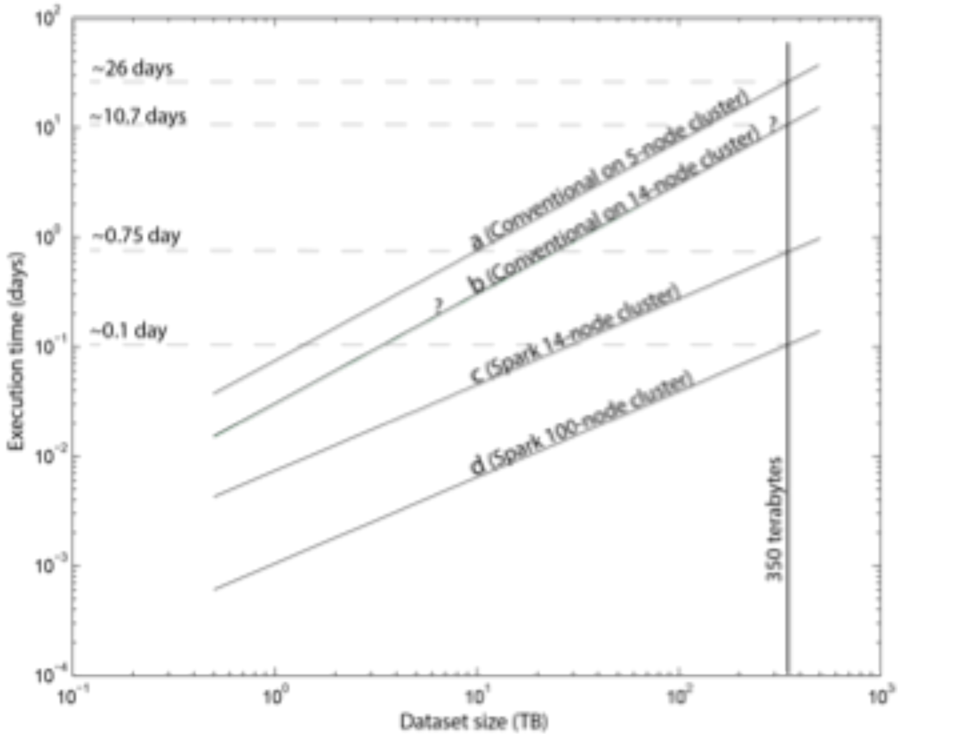

Magana-Zook, S., Gaylord, J., Knapp, D., Dodge, D., Ruppert, S., Large scale seismic waveform quality metric calculation using Hadoop, Computers & Geosciences (Volume 94, September 2016)

Layered Architecture

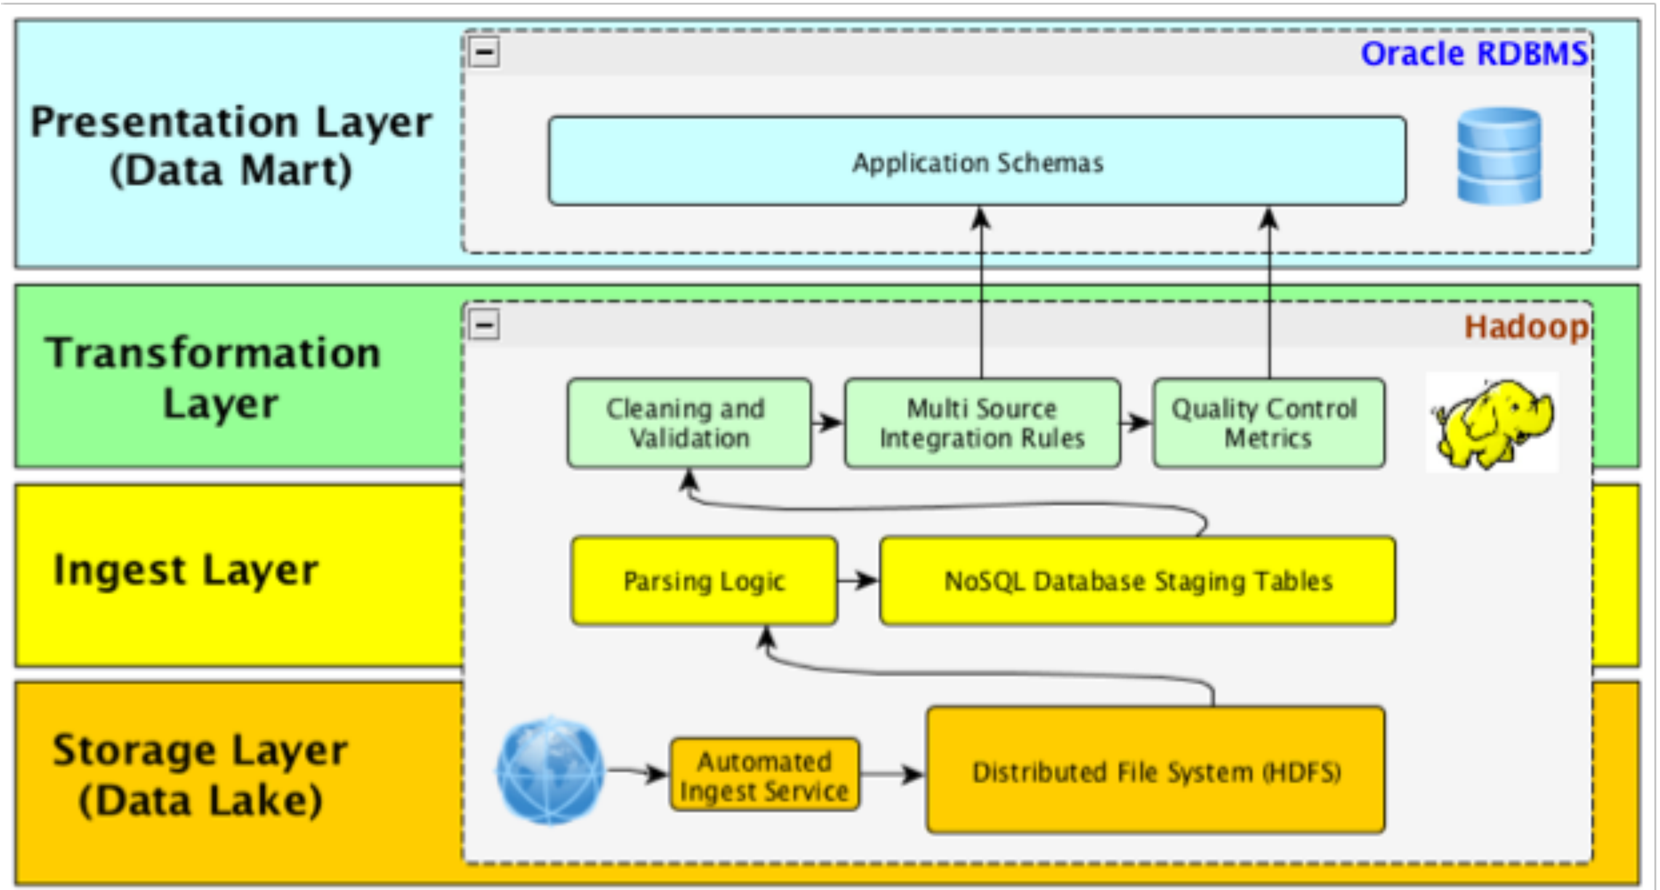

MINOS
Multi-Informatics for Nuclear Operations Scenarios

Signals of interest ➔ Data Fusion ➔ Generalization

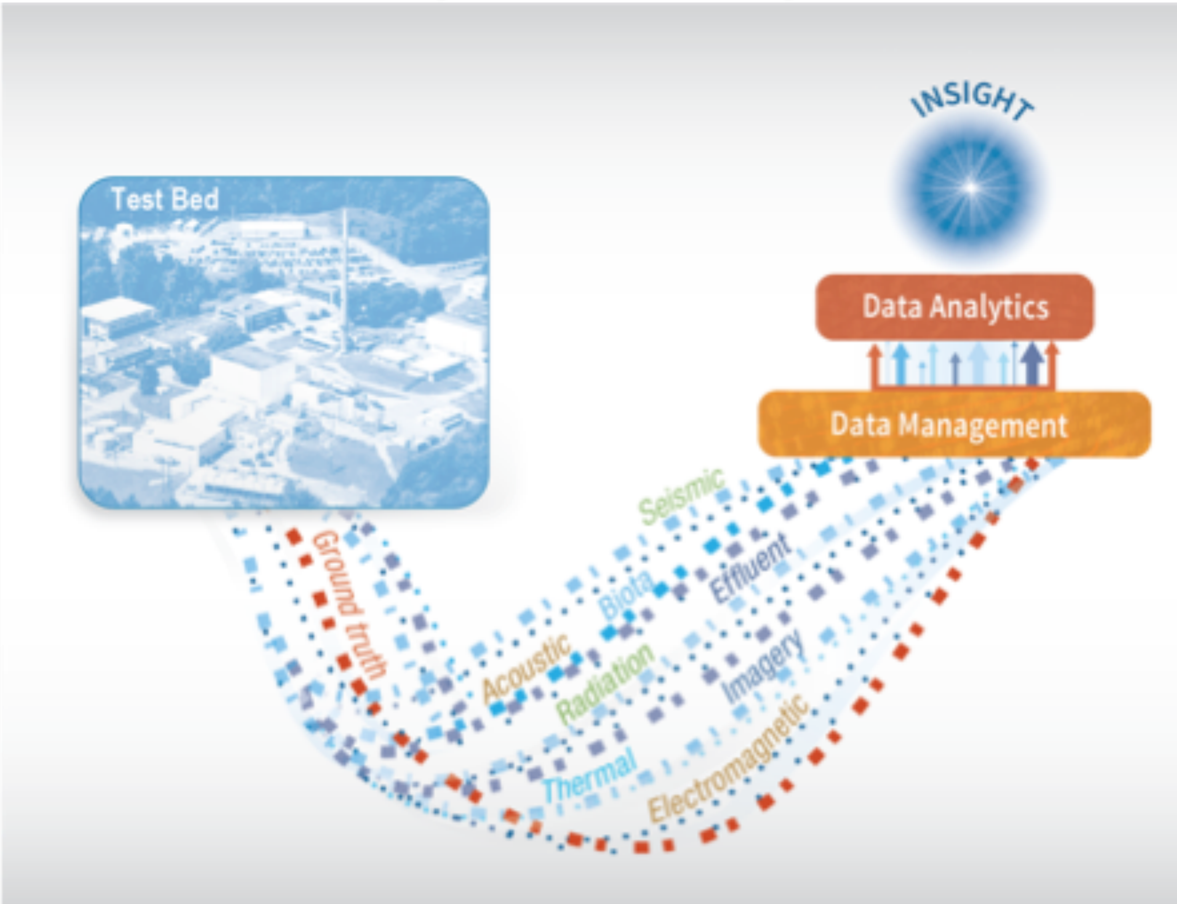

Test Bed
Operations of interest and ground truth

Data Collection
Measurements with different types of sensors

Data Management
Support of data storage, curation, access, evaluations

Data Analytics
Tools to leverage measurements for nonproliferation

MINOS Data Management Strategy

Key scientific goals:

1. Persistent field ingestion
2. Archiving, sharing, and curation
3. Analytics enabling platforms for monitoring data
   - Automation, Flexibility, Scalability

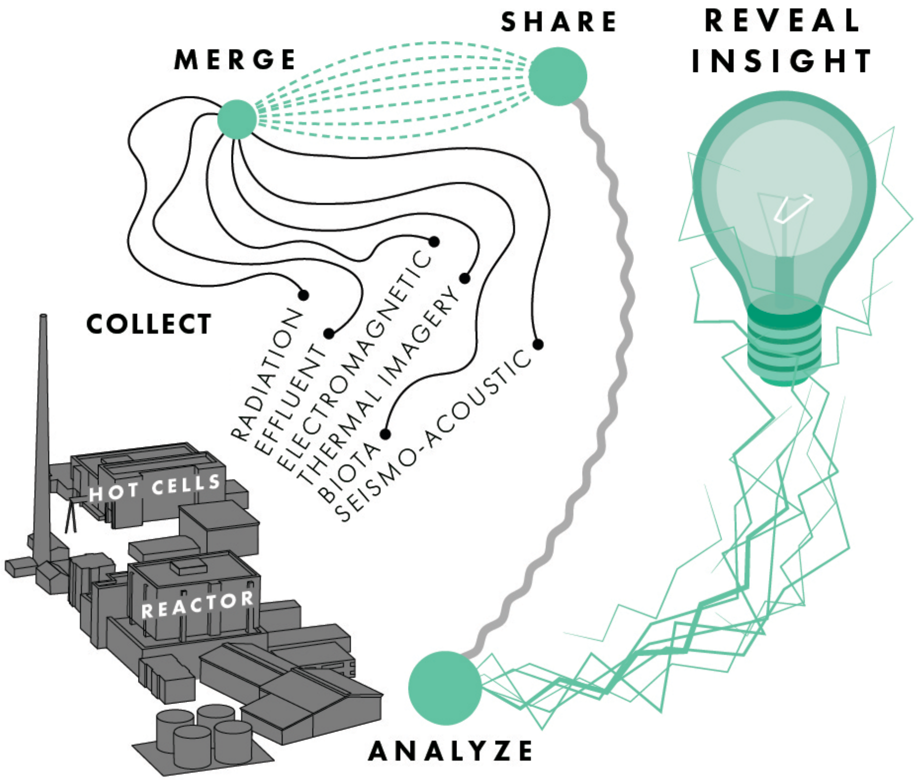

Task 3.1 Data Management Research

Task 3.2 Interim Storage Network

Task 3.3 Long Term Storage and Portal

Task 3.4 Summer Scholar Project

MINOS Use Case

- Venture – 10 labs, 25 Projects
  - Data system spans two labs
  - Disparate research groups for
    - Collections
    - Ground Truth
    - Analytics
  - Multi-modal streaming and batch data
  - OUO and export controlled data
  - Usage policy and hold-backs

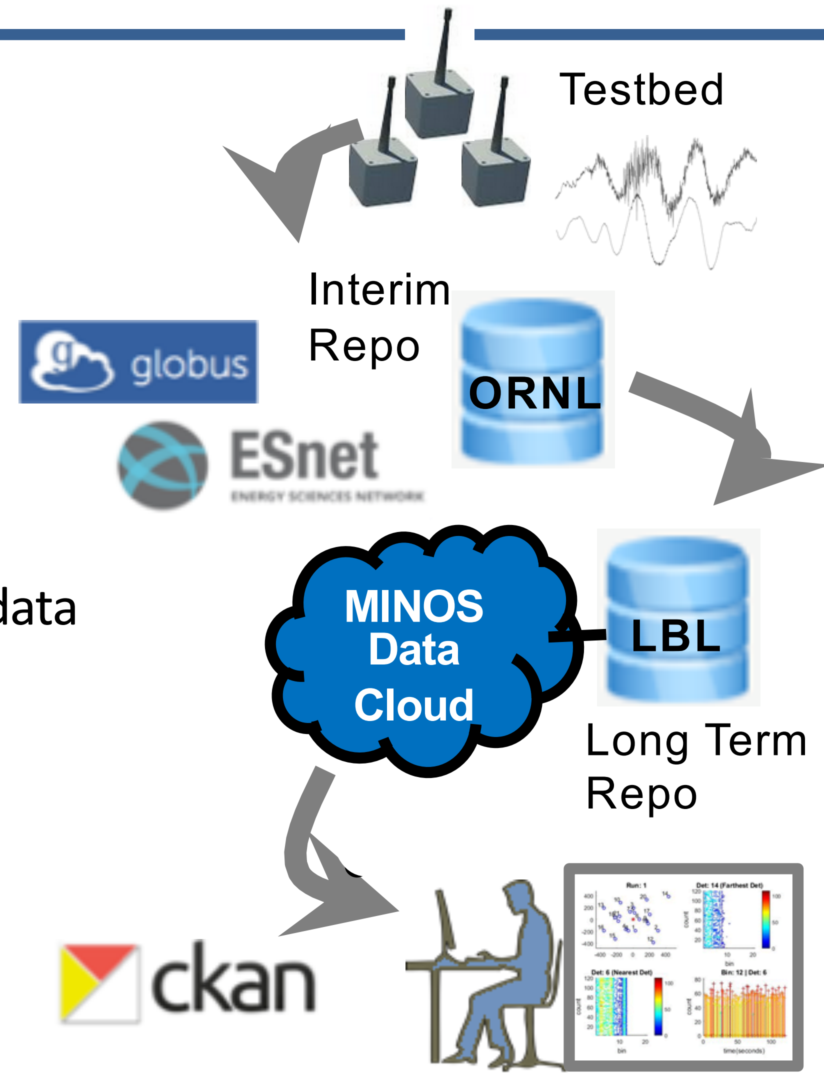

syncplicity SharePoint ckan

Monitoring with Big Data: A Scalable Data Ingestion Framework

- Fully automated
  - Streaming and batch data ingestion
  - Graceful handling of dynamic and unplanned data formats
  - Quality metrics and confidence measures
  - Spatial and temporal data integration
- Domain agnostic user interface
  - Context and device specific metadata
  - Integrated discovery over diverse datasets
- Up to multi-petabyte data holdings
- Horizontal scalability and performance
- Limited overhead for researchers and developers

Scalability and flexibility are key objectives

Other Project Goals

- Make information accessible about ingestion and data elements
- Make data transformations transparent, preserve provenance, and ensure data reproducibility
- Provide dynamic introspection into data as it is ingested for quality control and troubleshooting
- Plan for data drift: units of measure, precision, text length, date formats, source formats, format versions, etc.
- Isolate ingest from usage - remove forward coupling that implicitly ties applications to data sources
- Leverage existing tools to limit custom code, reduce complexity, and efficiently manage pipeline infrastructure

Dataflow Managers

- Control processes within the system, possibly including:
  - Collection
  - Storage
  - Parsing
  - Quality Control
  - Reformatting
  - Updating the data catalog
- May buffer intermediate data streams
- Several frameworks and utilities exist
  - Low level messaging queues (e.g. Kafka)
  - Dataflow and workflow utilities (e.g. Flume, Sqoop, Oozie)
  - Dataflow frameworks (e.g. NiFi, StreamSets)
- Or a custom solution (e.g. Java)

NoSQL Databases

- Horizontally scalable data stores
- Schema-on-read flexibility
- Improves performance with "eventual" consistency
- Not necessarily SQL, but usually SQL-like
- Many different types
  - Document (MongoDB)
  - Key-Value (Accumulo, Amazon S3 Dynamo)
  - Columnar (Cassandra, Hbase)
  - Graph (Neo4J)

Extensibility

- Plugin run-time container architecture
- Promotes modularity and availability
- For parser, QC, and analytics libraries

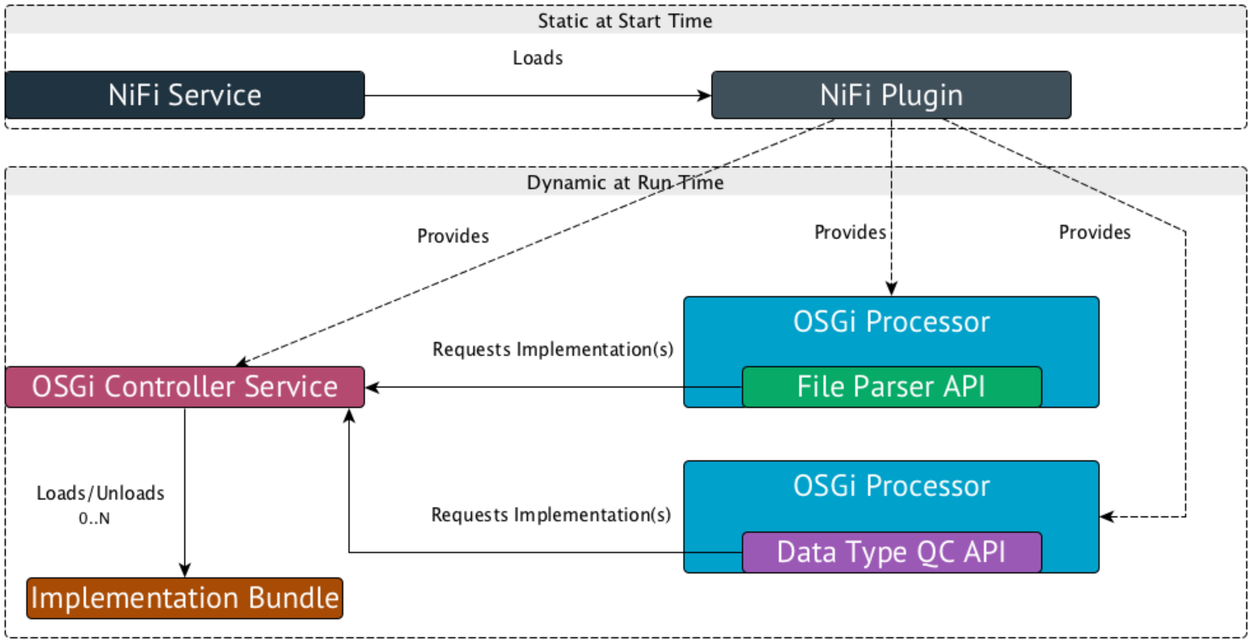

Containers

- Designed for code portability:
  - Contain everything code needs to run and nothing else
  - Host only provides container environment and OS
  - Package Once Deploy Anywhere! (as long as it's on Linux)
- Enables sharing entire technology stacks:
  - Includes all run-time dependencies and libraries
  - Reduces coupling between host and applications
- Supports flexibility in scale and infrastructure

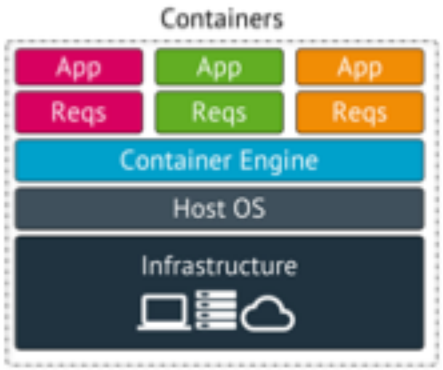

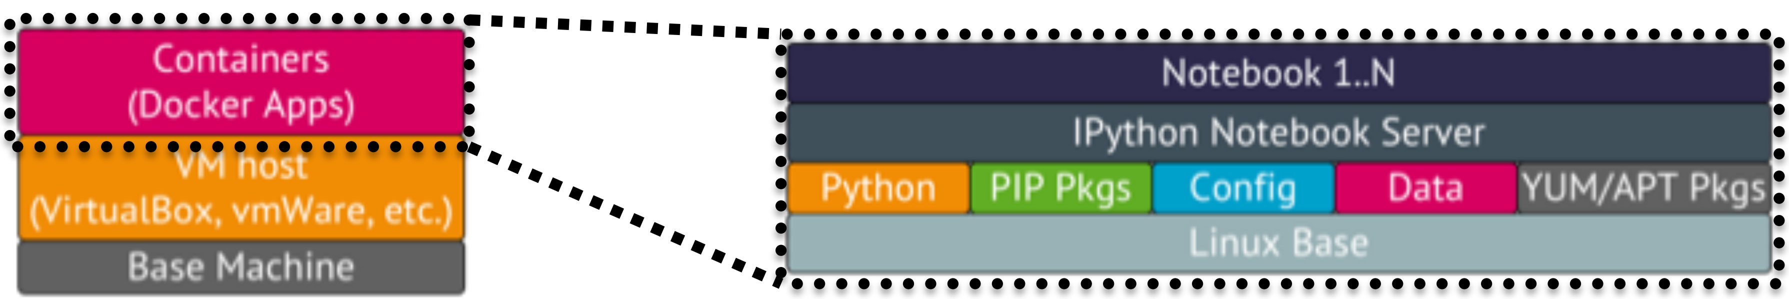

Challenges

- Technology stack is evolving rapidly
- Open source code can be poorly documented and/or buggy
- Systems usually include several technologies and can be complex
- Technologies are often optimized for commercial analytics workloads
  - Text datasets
  - Large files
- New schema and storage strategies are required
  - Multiple storage formats and poly-stores may be necessary
  - Joins can be extremely costly
  - Choose keys very carefully to avoid skew
  - Schema evolution, query types, and compression must be considered
- Data profiles, access patterns, and anticipated workloads can force difficult architecture tradeoffs

Scalable Ingestion Architecture

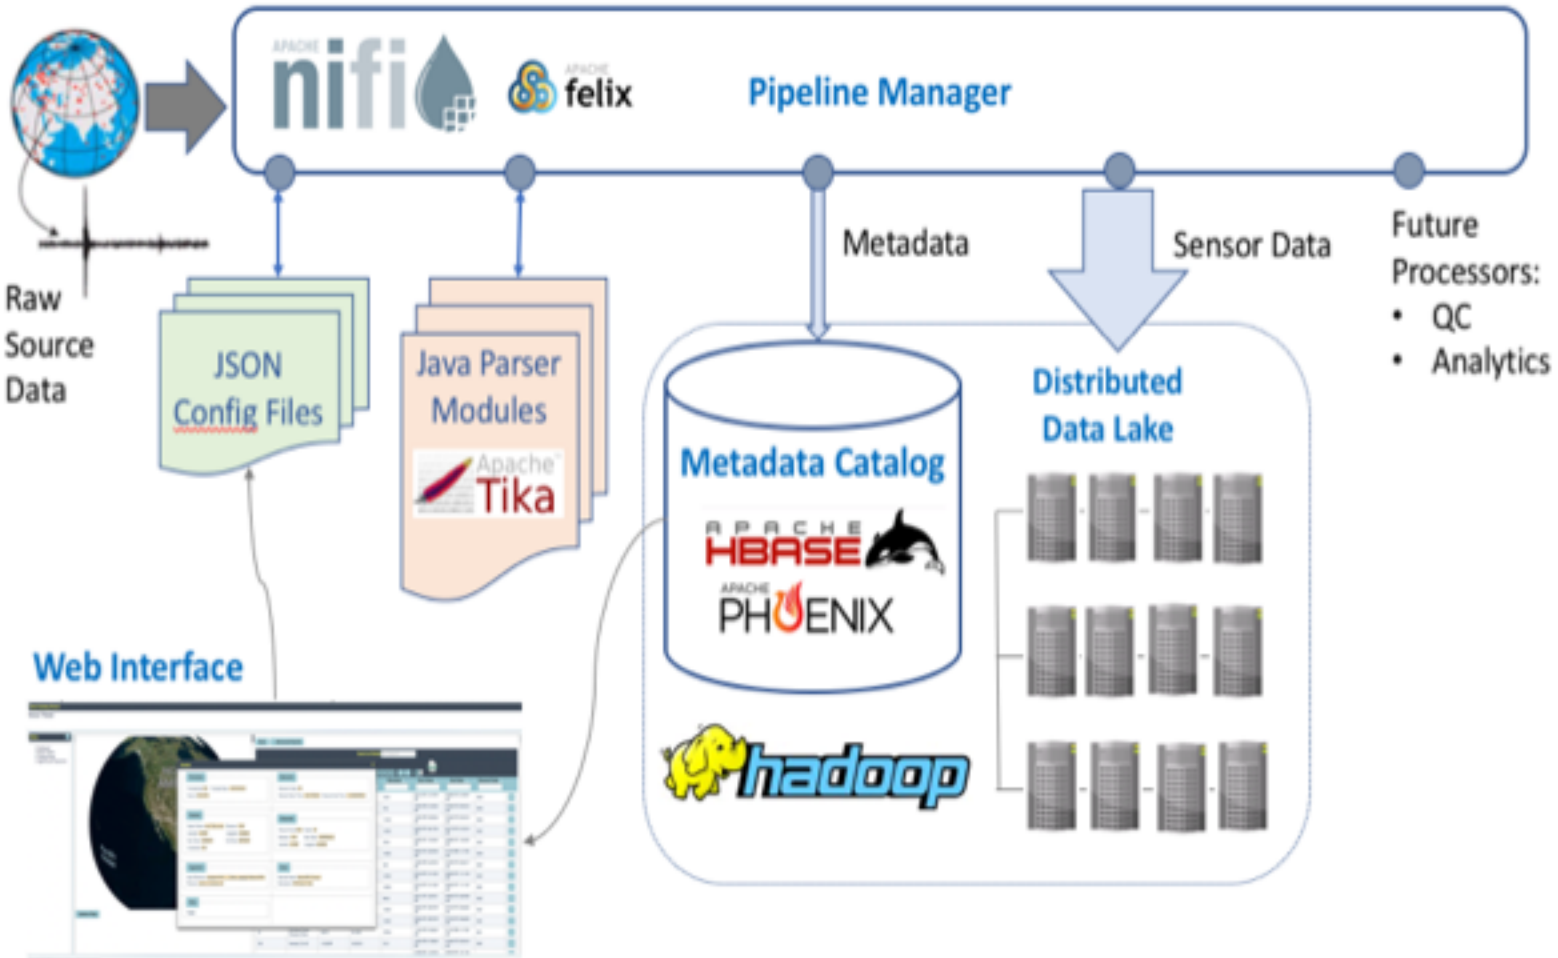

LYNM
Low Yield Nuclear Monitoring

How can we drive down monitoring thresholds?

- A venture of ventures
- 16 work groups
- Cross-cutting “Sprints”
  - Legacy Data
  - Simulation coupling
  - Integrated field experiments

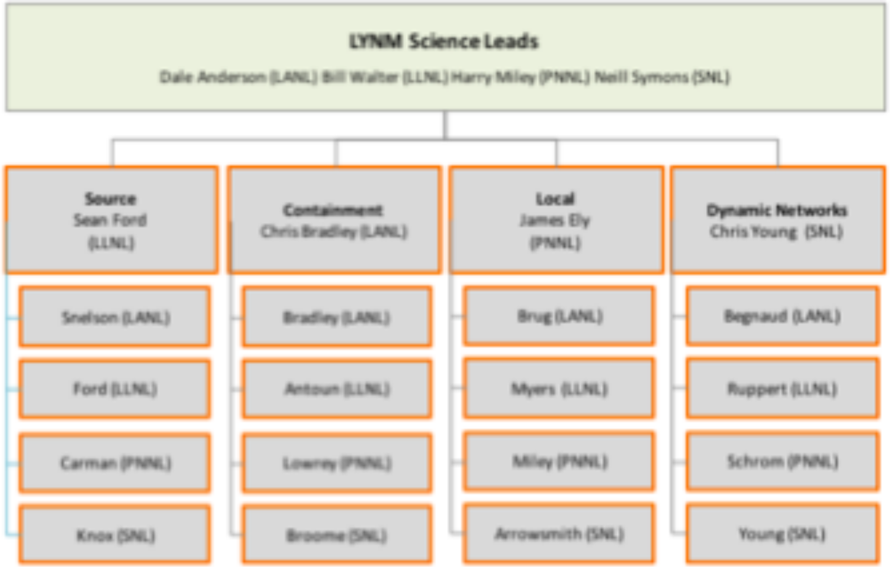

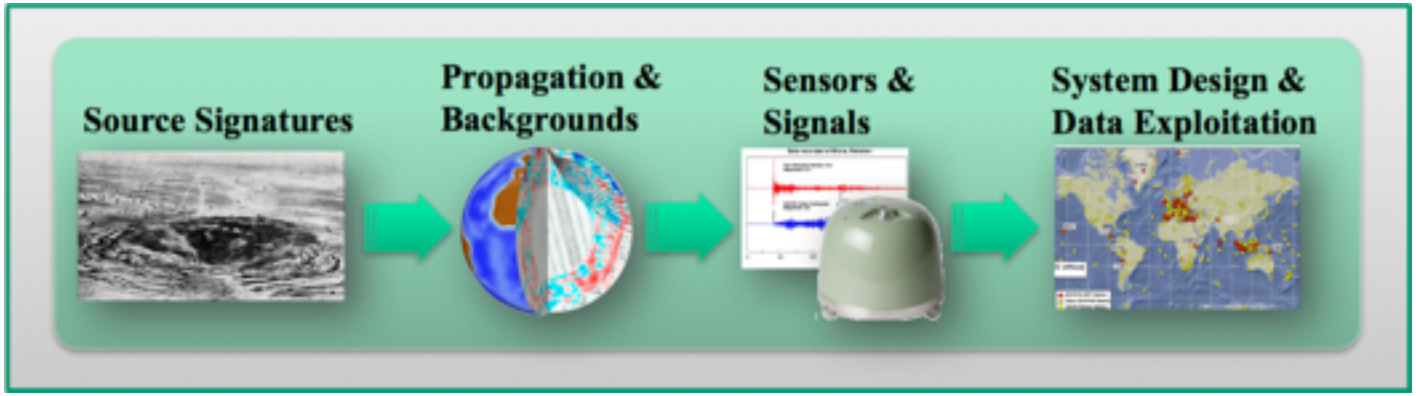

Potential DM Strategies

Less set-up
(harder to use)

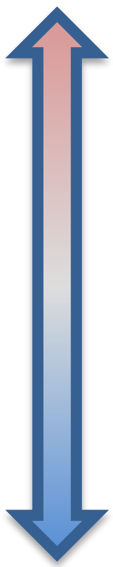

More set-up
(easier to use)

- Conventional
  - Each lab manages their own
  - Discovery and sharing is limited
  - No added overhead
- Federated
  - Catalogs are shared
  - Local administration and control
  - Better supports sharing
- Unified
  - Catalogs are combined
  - Supports sharing and data fusion
  - Requires a centrally managed repository

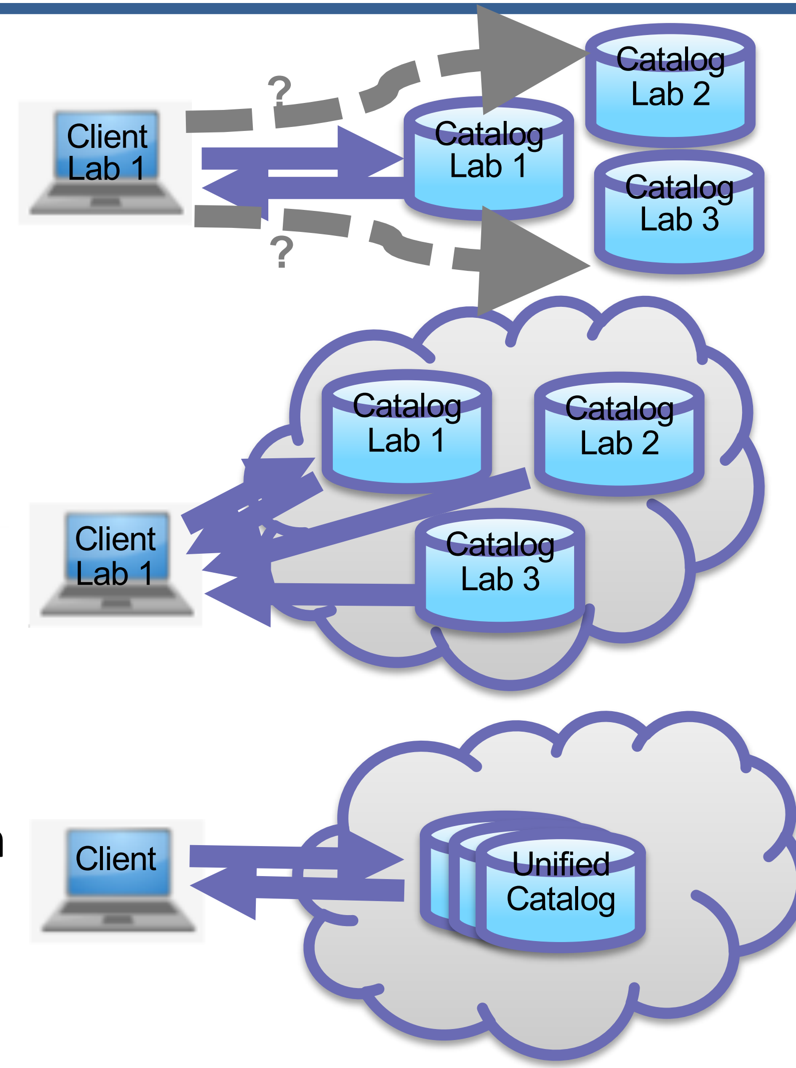

ESGF
Earth System Grid Federation

- Developed at LLNL
- 40 petabytes of climate data, over 700,000 datasets
- Federated architecture with internationally distributed peer-to-peer nodes
- 2017 R&D 100 award
- Working on extensions for bio and energy data

esgf.llnl.gov

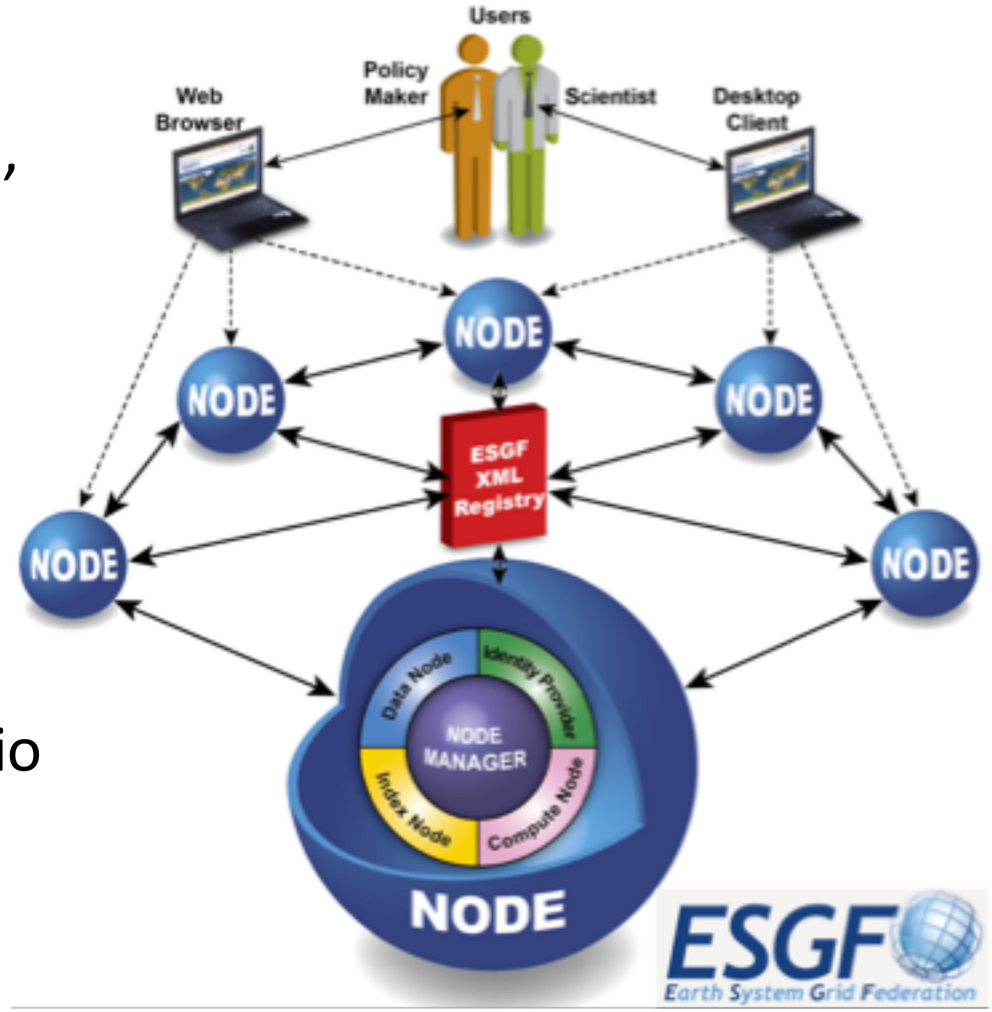

Internal Scientific Capability Project

- Many Data Science compute needs do not fit the HPC model
  - Workloads are IO and memory bound vs. compute bound
  - Persistent storage is necessary
  - Bi-directional, multi-network access and web servers are often required
  - A significant diversity of software frameworks/services are used
  - Flexibility to modify the stack easily and rapidly is often critical
  - Computing happens at all classification levels
- System administration is a significant overhead for programs, and security compliance requirements are only increasing
- Programmatic compute resources are often either scaled for peak/future loads or insufficient for new research opportunities
- Several programs rely on containers and orchestration to manage deployments to heterogeneous environments and complex dependencies

Vision

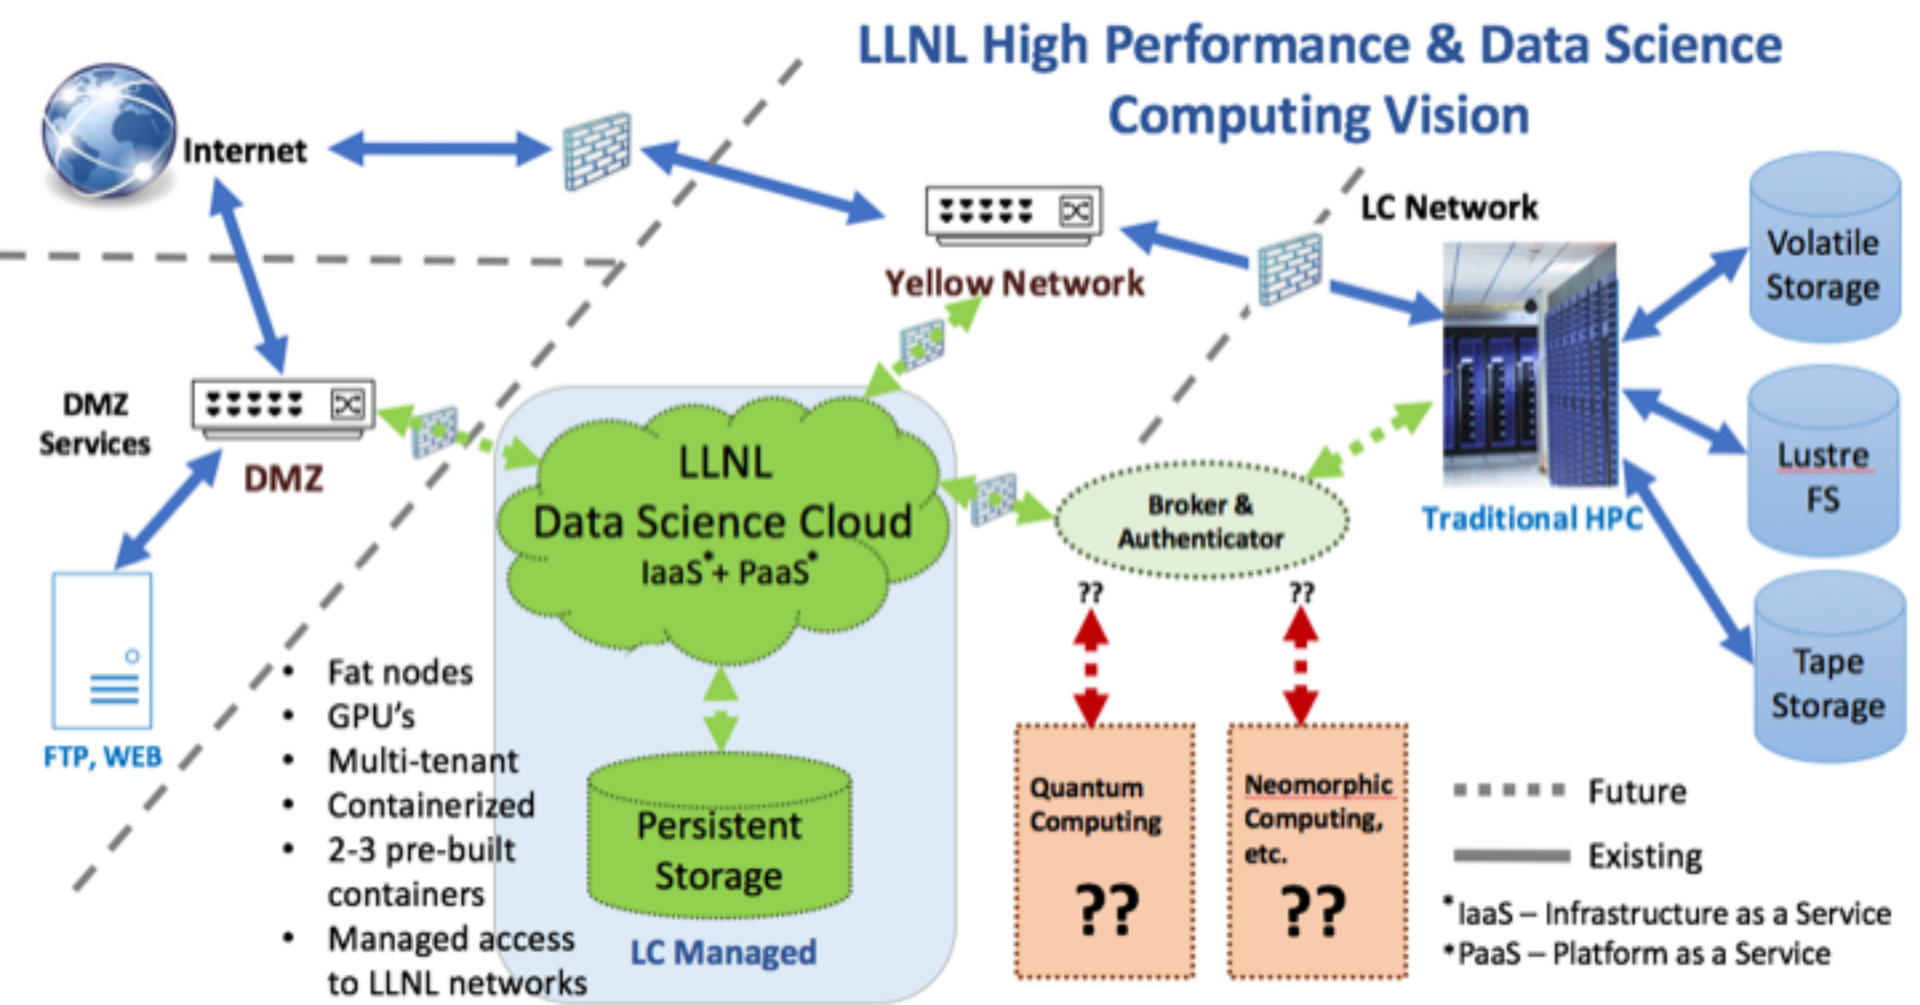
LLNL High Performance & Data Science Computing Vision
Internet
LC Network
Yellow Network
Volatile Storage
DMZ Services
DMZ
LLNL Data Science Cloud
IaaS* + PaaS*
Broker & Authenticator
Traditional HPC
Lustre FS
FTP, WEB
??
??
Tape Storage
Fat nodes
GPU's
Multi-tenant
Containerized
2-3 pre-built containers
Managed access to LLNL networks
Persistent Storage
LC Managed
Quantum Computing
??
Neomorphic Computing, etc.
??
Future
Existing
*IaaS – Infrastructure as a Service
*PaaS – Platform as a Service

Remaining Work

- Consolidate use cases
- Develop in-house container experts
- Multi-tenancy and container security
- Attribution and costing
- Partner with business services
- System ownership

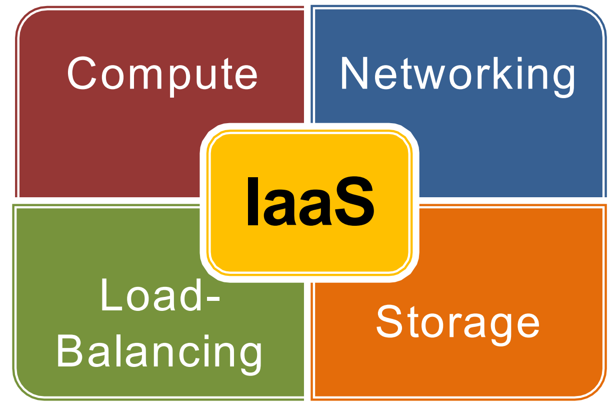

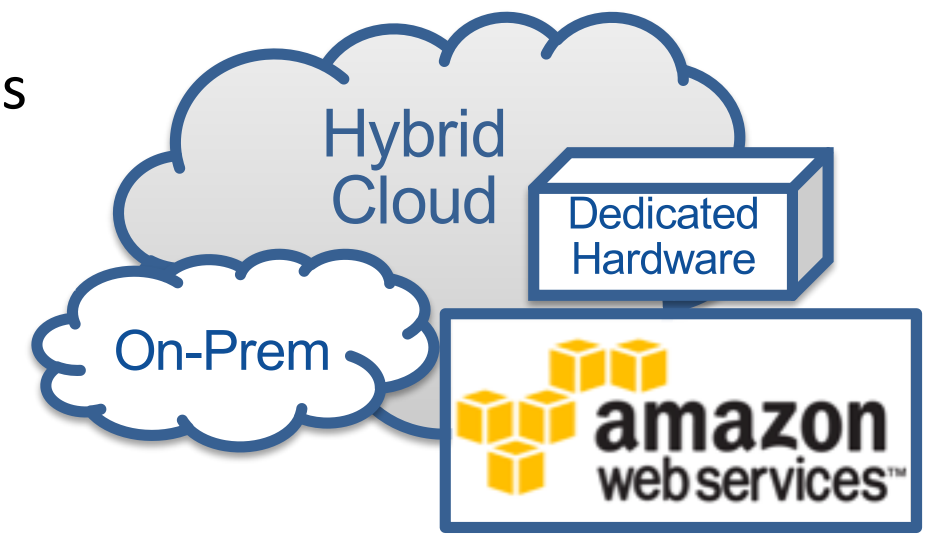

Conclusion

- Data Management is a key enabling capability for data science
- Monitoring data can be streaming or batched, legacy or new, loosely structured or managed, KBs to TBs, multi-modal, and unpredictable
- Systems and users may be distributed across clusters and sites
- Architectures and use cases vary

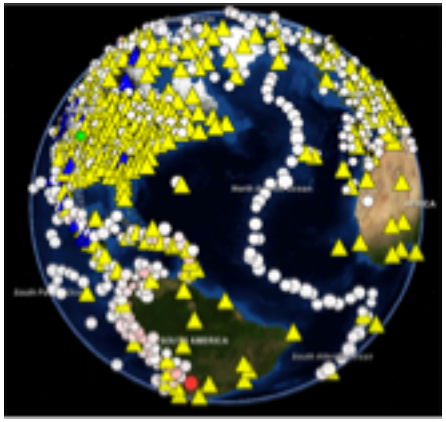

- What does this data mean?
- Where did this data come from?
- How good is this data?
- What am I permitted to do with this data?
- How can I get a copy?

Data is a critical asset that supports all of our work.

Thank you

Lawrence Livermore
National Laboratory