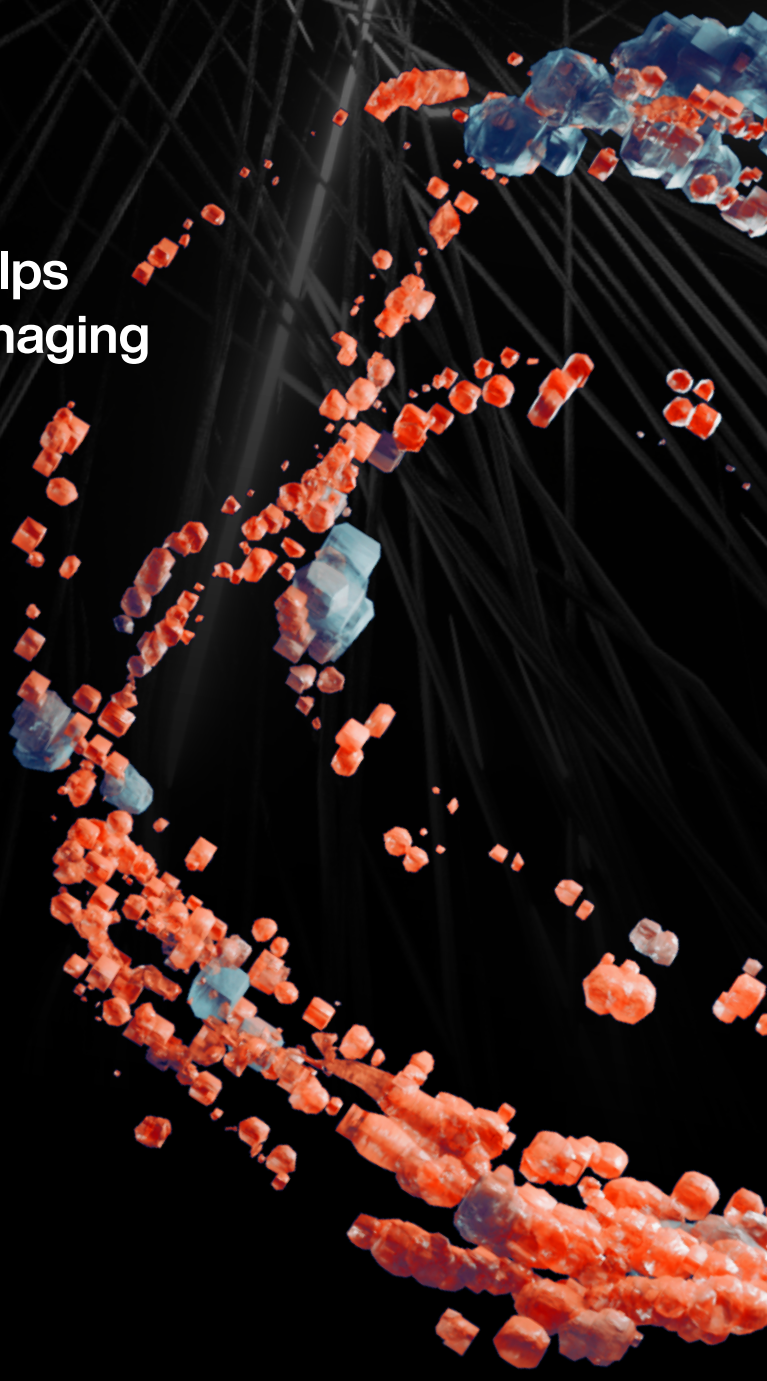
PRESENTATION NOTES
RMACC 2023: ARCITECTA × METADATA

Metadata

An Enabling Technology that Helps Relieve Some of the Burden Managing Data for HPC



Metadata is Awesome!

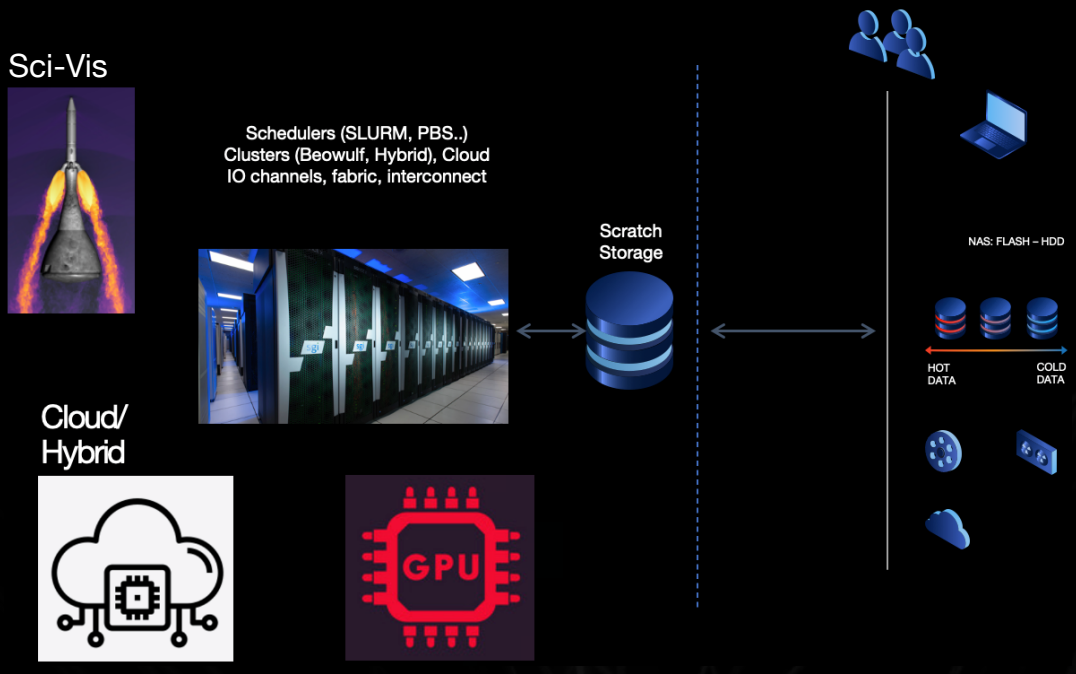
1. What is metadata, history, analog examples, why now?
2. Different types of metadata and how to efficiently manage it.
3. Metadata for compliancy and privacy protection
4. The “Golden Run”, why talk about metadata at an HPC conference? Hint: It is a key ingredient for efficient Data Lifecycle Management
5. Efficient HPC Scratch Storage use
6. Data Explosion and massive data sets mean more time spent on data wrangling, less money left over to run HPC jobs
7. Standards and Working Groups

ARCITECTA
OPERATING SYSTEMS FOR META+DATA

Arcitecta®, Mediaflux® and XODB® are registered trademarks of Arcitecta IP Pty. Ltd. in the USA and trademarks of Arcitecta IP Pty. Ltd. in Australia
© 2023 Arcitecta IP Pty. Ltd.
www.arcitecta.com
talk@arcitecta.acom

1. What is Metadata?

Simply: Data about Data



Images courtesy of NASA

Traditional View of HPC

Everyday Examples of Metadata

- Metadata has been in use for a couple of centuries, we are just moving to the digital version.
- Card catalog since 1862 at Harvard
- Great example of a “database”
- Catalogs by Titles, authors, subjects
- Search Keys
- iPod, FM song + artist, etc.

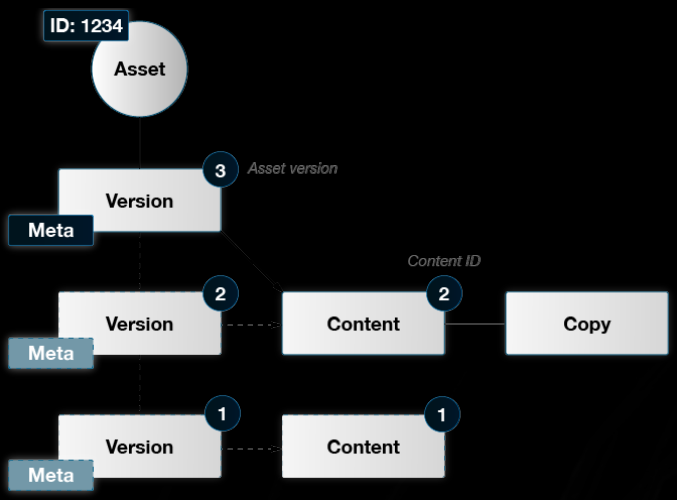


Image courtesy of New York Society Library

© Copyright 2023 Arcitecta Inc.
All rights reserved

Digital Object Model of an “Ideal Metadata Database”

- A binary optimized database for the metadata is managed independently of the content.
- Data is not “held hostage.”
- User defined Metadata and System Metadata.
- Let’s take a closer look...



2. Types of Metadata

“System Metadata”

- File name, size, create, access, modify time, ownership permissions, etc.
- In Unix / Linux world this information comes from inodes and is often used by Backup or HSM software.

Embedded File Metadata

- Typically parsed out via MIME type.

User-defined Metadata

- This enables data life cycle management, notes, accounting.

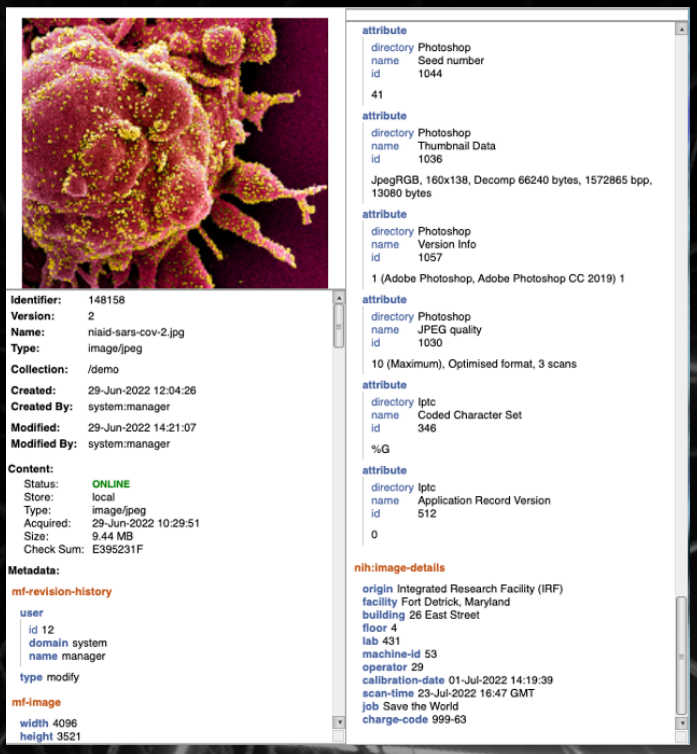
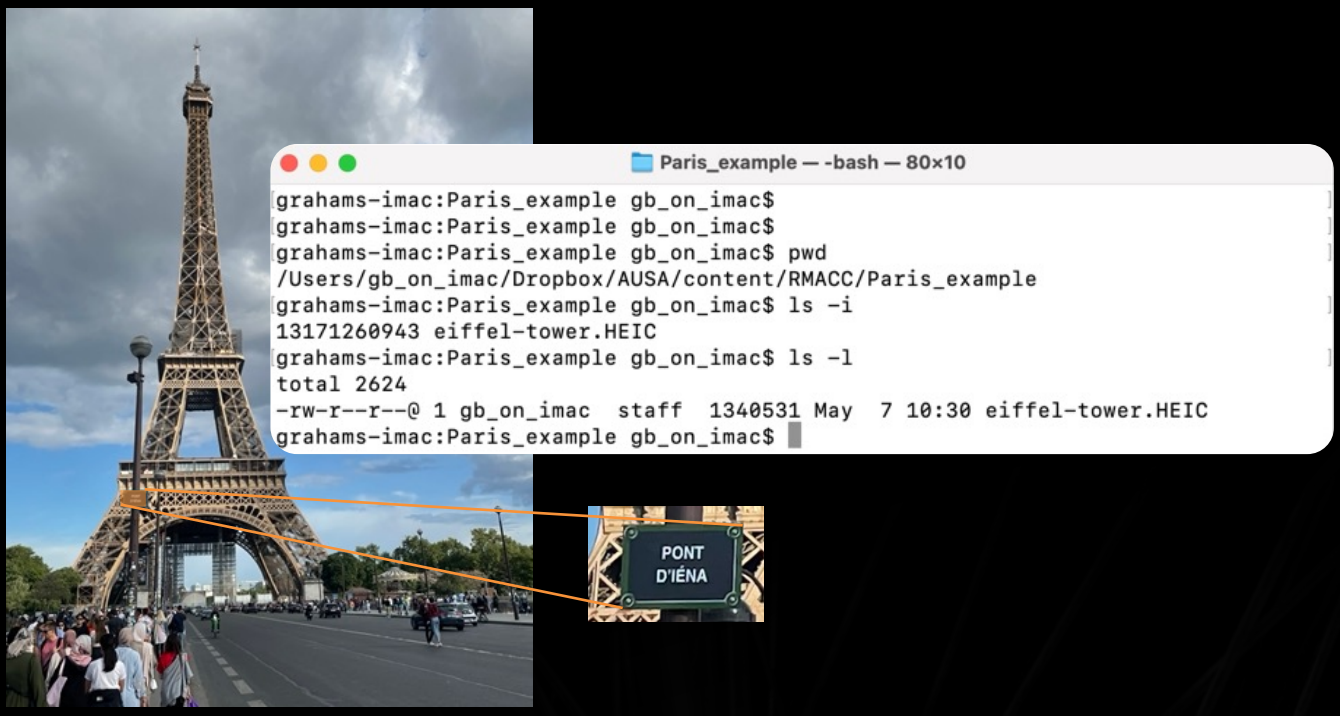
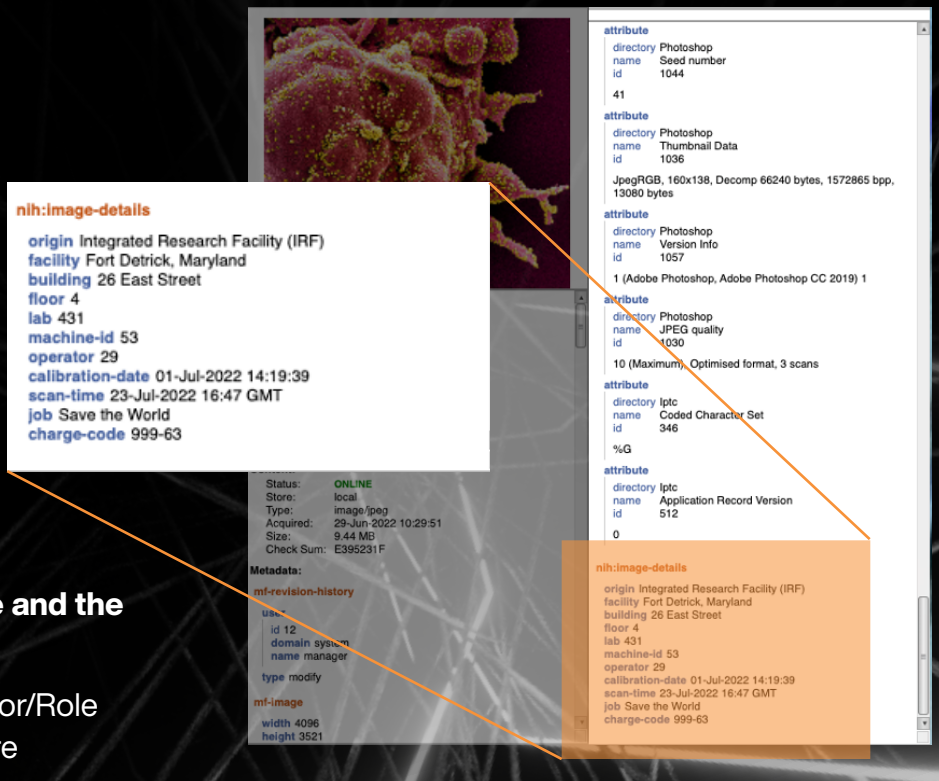


Image courtesy of NIH
National Institute of Allergy and Infectious Diseases (NIAD)

© Copyright 2023 Arcitecta Inc.
All rights reserved



Metadata Example



***Privacy Information goes here and the fields can Evolve**

- Information that influences Actor/Role access models can also go here

Image courtesy of NIH
National Institute of Allergy and Infectious Diseases (NIAD)

© Copyright 2023 Arcitecta Inc.
All rights reserved

3. Metadata for Compliancy and Privacy Protection

- Data Provenance
- Access Controls
- GDPR, HIPAA, CPRA
- What can you do? Build a system with Actor/Role model
- Update Metadata as laws/rules change
- Who needs this... Life Sciences, Defense, Financial... just about everyone!

Metadata Needs to Evolve.

Nobody knows what the future will bring — the perfect metadata now might not be perfect in 2 years!

Metadata is Awesome!

- Why? It helps you find and organize things!
- Optimally!
- When you need them, even automatically!
- Other names “tags”, “search keys”, fields

Feed the Beast

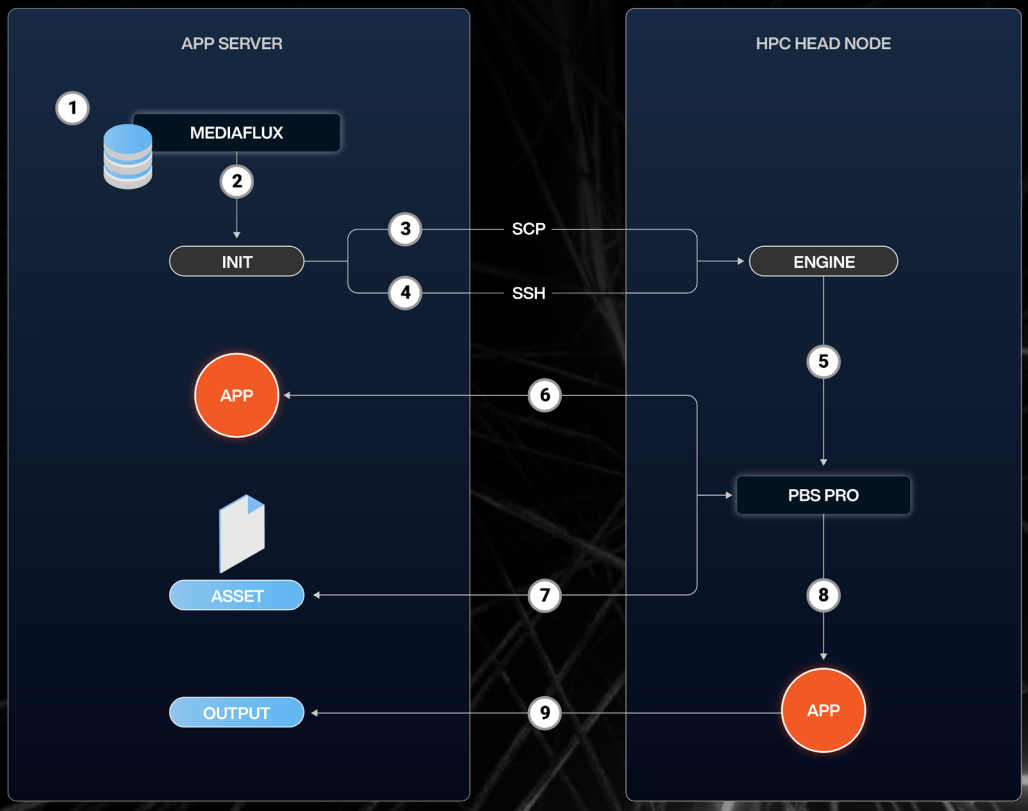
- You spent a fortune on an HPC Cluster
- Almost as much on Network infrastructure
- Tons on storage
- Why?
- To refine computing models to ever-increasing fidelity, and develop sophisticated new ones, data staging is often imperative.
- Non-trivial - art/science called “Data Wrangling”

© Copyright 2023 Arcitecta Inc.
All rights reserved

4. Putting it all Together: Data Lifecycle Management

Compute Services Framework

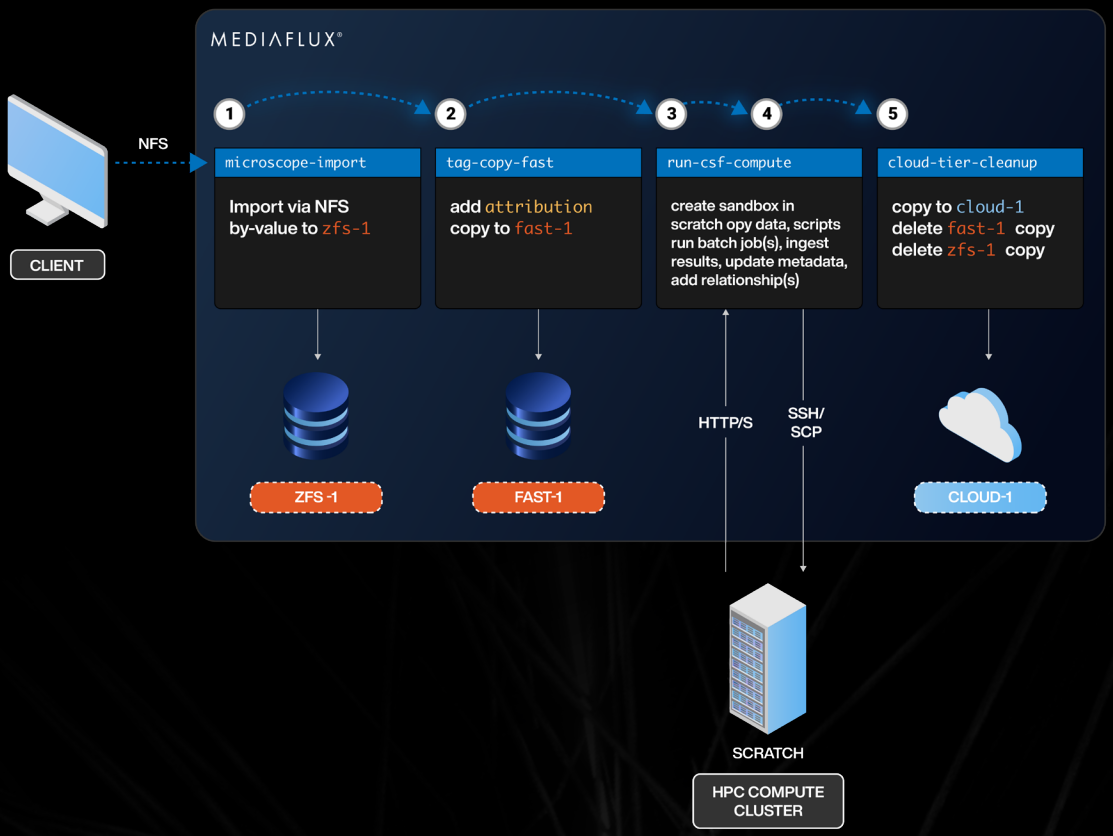
- Many use cases require external computation using applications and/or resources external to Mediaflux®.
- Examples include HPC applications, workflow and processing pipelines – PBS Pro, SLURM, ...
- Mediaflux® Compute Services provides a framework for configuring, deploying and executing compute applications
- Once configured, a compute service may be utilized by any authorized application and/or user interface by calling a single service within Mediaflux®.



Example Compute Services Workflow

© Copyright 2023 Arcitecta Inc.
All rights reserved

5. Efficient HPC Scratch Storage Use

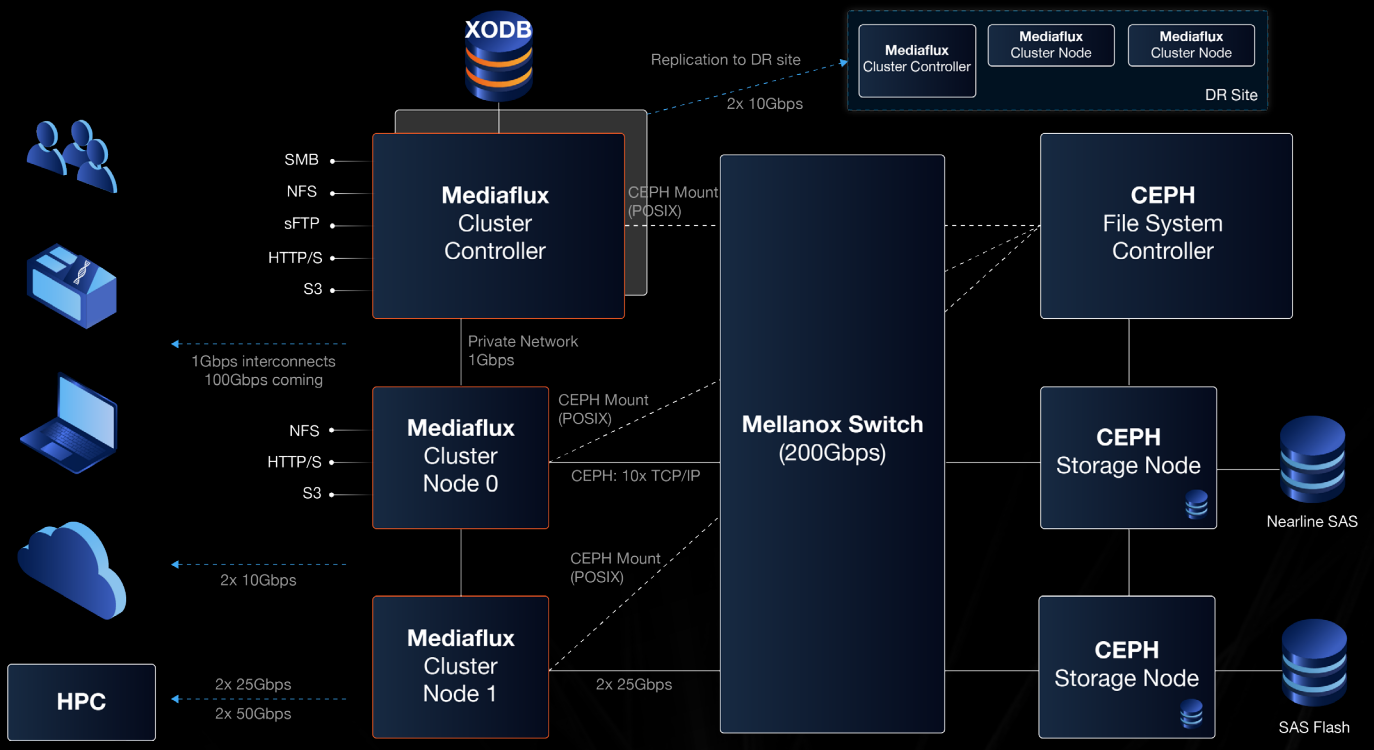


Data-Driven HPC Workflow

Data-Driven HPC Workflow

1. Data “lands” from instrument or other source in a directory in an NFS-mounted file system. Mediaflux can be configured to watch particular directories, and add any newly-appearing entries to recurring ingest.
2. Once ingested, incoming data is tagged automatically. In the above we add the tag “attribution”. We also create a content copy, again automatically, to protect the new data via redundancy.
3. A Mediaflux compute services workflow is utilized in the 4th step, to set up and run an HPC computation. First, all needed data is copied out to scratch. Then, a batch job is run by Mediaflux on the HPC compute cluster.
4. Once the HPC computation is complete Mediaflux ingests any/all results that need to be saved before the scratch area is cleared.
5. In the final step archival copies of data are made to cheaper deep storage, and copies on faster storage are removed to recover those resources for other purposes.

© Copyright 2023 Arcitecta Inc.
All rights reserved



University of Melbourne

University of Melbourne: Scale Up

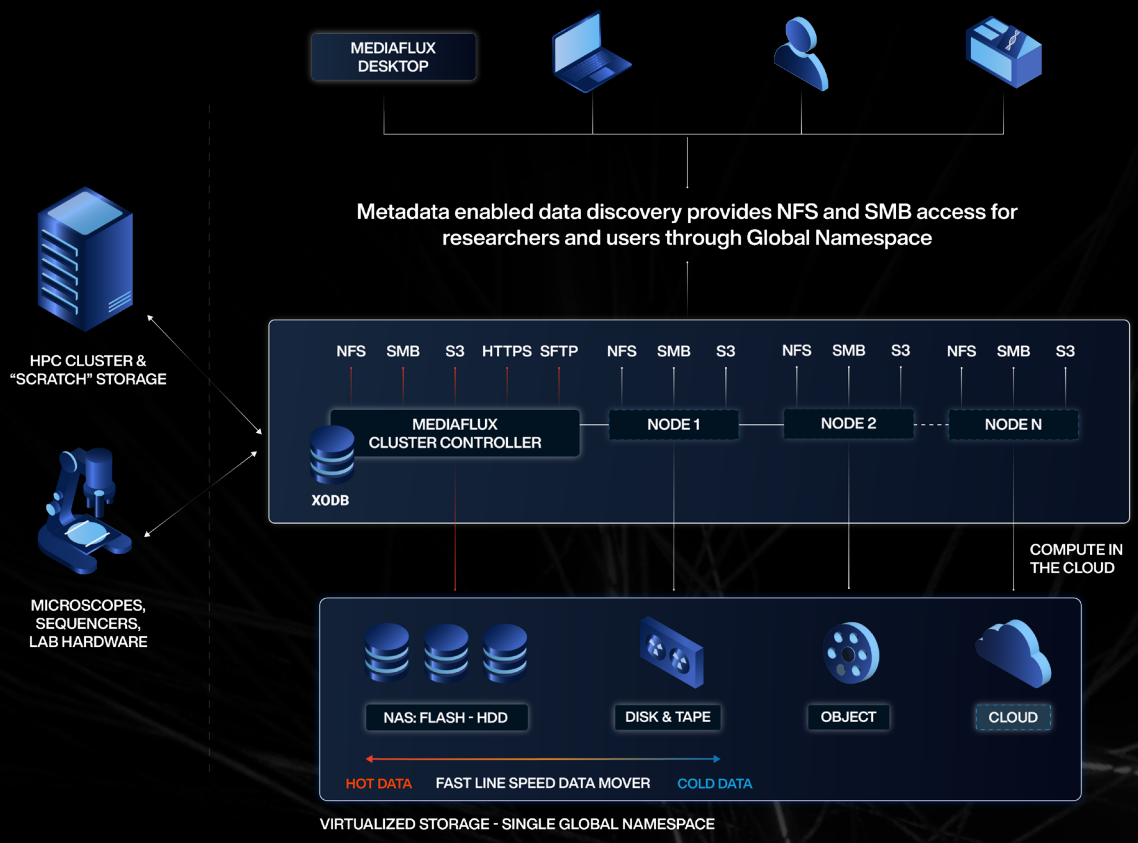
Purpose: To provide an enterprise-grade storage infrastructure that will scale out horizontally using Mediaflux and CEPH.

How it Works:

1. CEPH: Will present 2x mounts: a) Flash and b) Nearline SAS. Policy above will drive usage.
2. Larger files are better - consider gene sequence example - need to bundle (at client, or server-side). If at server side, then need server side “scratch” area - that can be used for staging data and then creating a larger bundle which is written to the CEPH cluster.
3. CEPH: Large file limit: 4 TB. Preferred max file size: 1 TB (for bundling we would want to bundle to a much smaller number - say max 1-10GB, say 2GB - easily fits in NVMe 10GB cache)
4. CEPH: Object size = 4MB, stripe size = 1MB, strip count = 4
5. I-nodes: We could cluster / group files behind the scene to reduce i-node counts on the backend file systems. CEPH has the concept of active i-nodes (files in use - current limit 60 million) and inactive i-nodes.

© Copyright 2023 Arcitecta Inc.
All rights reserved

6. Data Explosion and Massive Data Sets Mean More Time Spent on Data Wrangling, Less Money Left Over to Run HPC Jobs



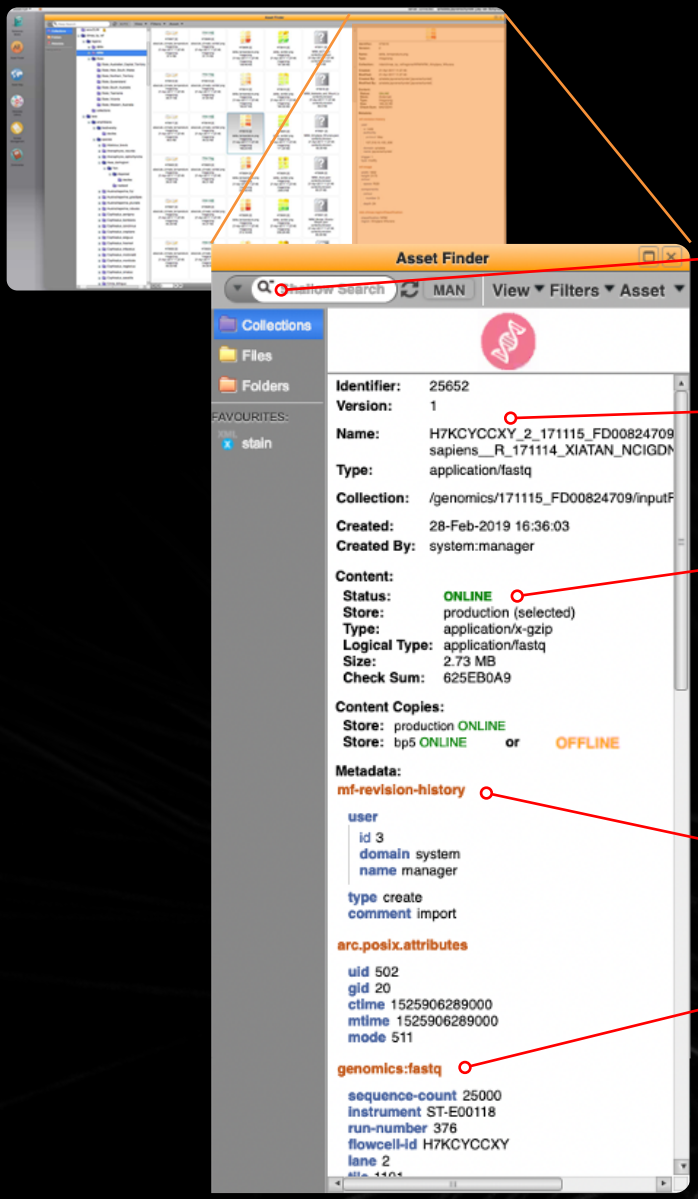
Hybrid Environment for HPC and Research Data Management

When technologies moved from analog to digital, data rates in virtually all industries exploded, for the most part exceeding Moore's law and for sure exceeding the rate at which storage prices go down.

© Copyright 2023 Arcitecta Inc.
All rights reserved

Access Via a Browser-based Interface—

Mediaflux® Desktop



File Browser GUI Asset Finder searches metadata fields of interest to filter and create compound queries

File system level information (metadata) such as: Name, Location, Creation Date, etc.

Content information, checksum, size and **data accessibility**
ONLINE – content is available on nearline storage
OFFLINE – content is on deep storage (on BlackPearl® Tape)
ONLINE+OFFLINE – content is available on both
MISSING – content is missing

Revision history, audit trails, provenance and **reproducibility**

Descriptive metadata extracted during the ingest process

Custom-built scraper that extracts specific metadata based on the data type. These plug-in content analyzers extract image type and resolution, scanner instrument information, project, study, population, PI, etc.

User generated metadata such as: annotations, labels, tags, comments, and workflow-specific actions

© Copyright 2023 Arcitecta Inc.
All rights reserved

Data is Everywhere

Our technology is at the core of some of the most demanding big data environments.

Life Sciences

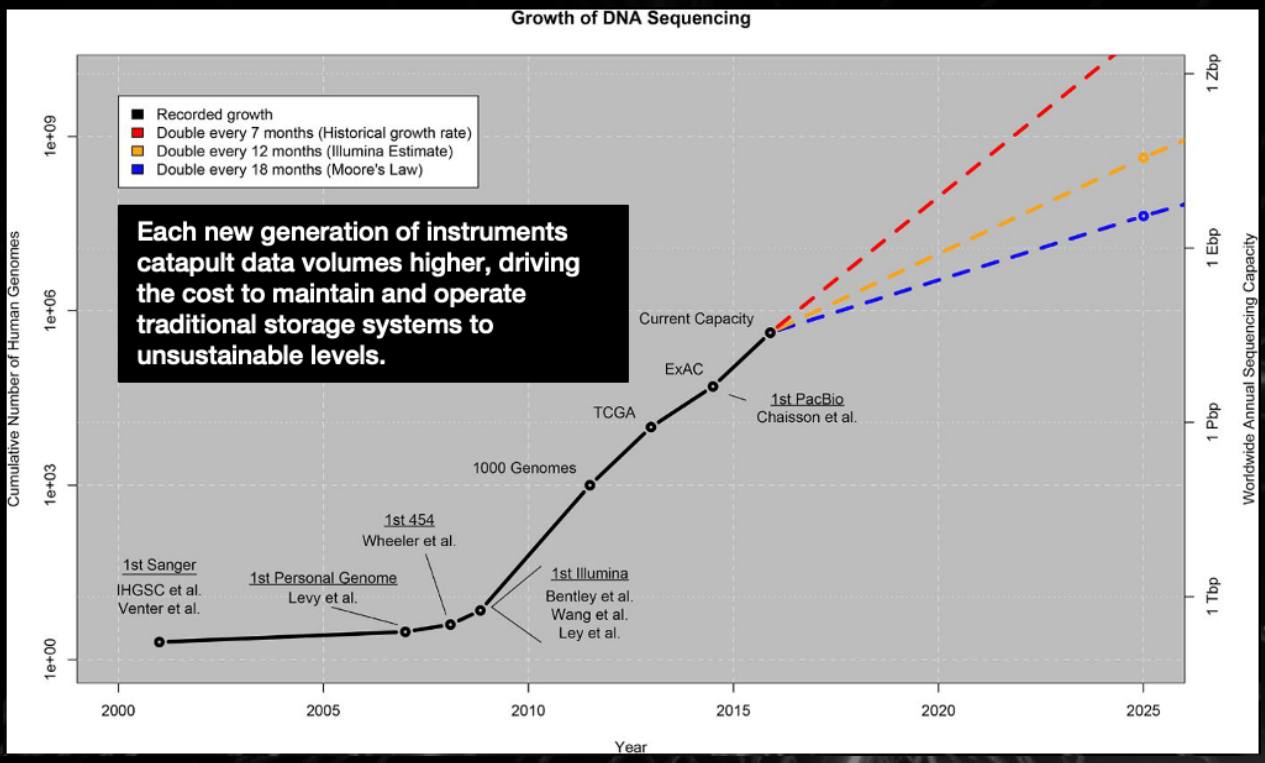
Research

Clinical

Geospatial

Media & Entertainment

Defence & Intelligence



A Known Problem: Exponential Data Growth in Genomics

Big Data: Astronomical or Genomical? - Scientific Figure on ResearchGate. Available from: https://www.researchgate.net/figure/Growth-of-DNA-sequencing-The-plot-shows-the-growth-of-DNA-sequencing-both-in-the-total_fig1_279863341

© Copyright 2023 Arcitecta Inc.
All rights reserved

7. Standards and Working Groups

“FAIR” Principles

From: https://www.go-fair.org/

Findable

“The first step in (re)using data is to find them. Metadata and data should be easy to find for both humans and computers. Machine-readable metadata are essential for automatic discovery of datasets and services, so this is an essential component of the FAIRification process.”

Accessible

“Once the user finds the required data, she/he needs to know how can they be accessed, possibly including authentication and authorisation.”

Interoperable

“The data usually need to be integrated with other data. In addition, the data need to interoperate with applications or workflows for analysis, storage, and processing.”

Reusable

“The ultimate goal of FAIR is to optimise the reuse of data. To achieve this, metadata and data should be well-described so that they can be replicated and/or combined in different settings.”

© Copyright 2023 Arcitecta Inc.
All rights reserved

The perfect metadata now might not be the perfect metadata in 2 years, 10 years or 20 years.

More Information

Data Management Tools, Whitepapers, Products, etc.
www.arcitecta.com

2022 Gartner report ”Strategic Roadmap for Storage”
Key take aways— Suggested action plans like: Appoint “data champions” to develop policy-aware data classification and retention strategies for business-critical applications and compliance and privacy issues.

2022 Whitepaper “Scientific Data” by Devan Ray Donaldson
Donaldson is PhD Professor at Luddy School of Informatics, Computing, and Engineering at Indiana University
https://www.nature.com/articles/s41597-022-01428-w

Citable ID’s etc.
https://libguides.library.nd.edu/research-data-services/make-citable

Data Management plans, etc.
https://researchdatamanagement.harvard.edu/data-management-plans
https://researchdata.princeton.edu/research-lifecycle-guide/data-management-plans
https://dataservices.research.unimelb.edu.au/services/40/

© Copyright 2023 Arcitecta Inc.
All rights reserved