iRODS

CONSORTIUM

iRODS

Executive Overview
August 12, 2014

Agenda

- What is iRODS?
- Who Uses iRODS?
- What Can iRODS Do?
- The Future of iRODS

WHAT IS iRODS?

What is iRODS?

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

iRODS is the underlying technology for the world's preeminent genomic research institutes. iRODS is an infinitely configurable data janitor. iRODS is the kind of technology you need to host everyone's unstructured data. iRODS is a powerful data migration tool. iRODS is the technology that underpins the iPlant Data Store. iRODS is a data preservation technology. iRODS is a fundamental technology for CineGrid. iRODS is a tool for providing fine-grained privacy and security controls. iRODS is extensible: iRODS has command-line clients, APIs for numerous programming languages, and web clients. iRODS supports new plug-ins for storage resources, authentication mechanisms, microservices, and network prot

iRODS is Open Source

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

The iRODS Consortium exists to ensure the sustainability of iRODS by:

- Ensuring that iRODS source code remains freely available for use and modification.
- Promoting the adaptation of iRODS to a variety of hardware and software platforms.
- Supporting continued development of core iRODS features.
- Facilitating interaction among members of the iRODS developer community.
- Providing a forum for key stakeholders to guide ongoing development of iRODS.

iRODS is Middleware

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

mid•dle•ware *`midl,we(ə)r*

noun software that acts as a bridge between an operating system or database and applications, especially on a network

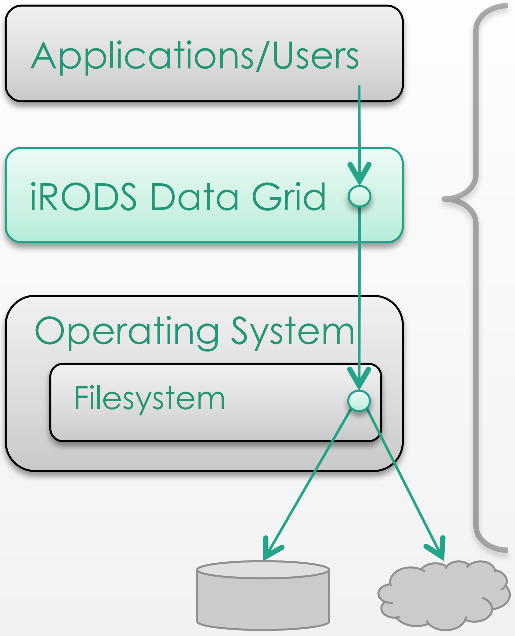

(Heterogeneous) Storage Systems

iRODS Middleware Layer

- abstracts out the low-level I/O
- provides a uniform interface to heterogeneous storage systems (POSIX **and non-POSIX**)

iRODS lets system administrators roll out an extensible data grid **without** changing their infrastructure.

Data is accessed using familiar APIs.

Data Discovery

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

Every iRODS data grid also has a metadata catalog, called the iCAT. The iCAT is used by iRODS to locate data, manage provenance, and to enable automation and access control.

The iCAT also permits user-defined metadata. Altogether, this metadata supports:

- Data discovery based on parameters such as user-defined tags, modification date, outcomes of automation activity.
- Capturing workflows as raw data is processed and used.
- Automation and access control policies.

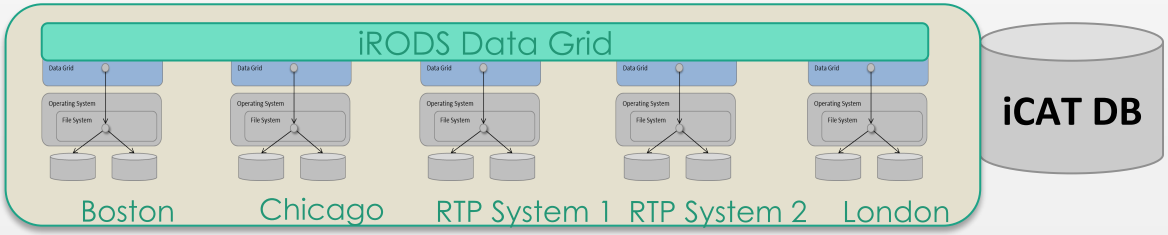

Example Metadata:

Logical Name (iRODS path):
/RDDept/LabX/Flow/Study1

Physical Name (Unix path):
/London/var1/proj/labx/stuff

Lab PI: Jane Doe
Date: 12/1/2010
Time: 01:45:12

Title:
Proliferation optimization studies

Data Source: Flow Cytometer
Assay Conditions: Data captured
…

Workflow Automation

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

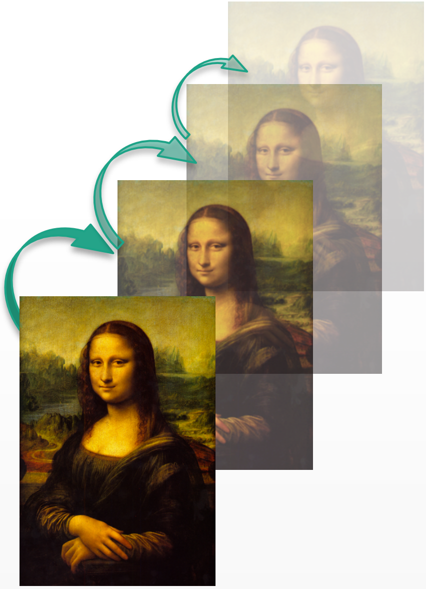

With iRODS, **any agent** can initiate **any action** upon **any trigger.**

This powerful capability allows administrators to automate policies such as:

- Validating checksums every time a new file is placed in a folder.
- Backing up a set of files every second Thursday.
- Archiving data that hasn't been accessed in over 1 month.
- Logging each time a file is replicated or destroyed.
- Permitting a file to be accessed by multiple independently defined user groups.

These operations can be **distributed** to the storage resource or client.

Secure Collaboration

iRODS is open source data grid middleware for…

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

iRODS presents centralizes distributed storage systems under a unified namespace.

Administrators can control how the grid is presented to users and implement replication, load-distribution, and archiving policies that are completely transparent to the user.

Independent grids can be federated with one another to allow controlled access to remote grids or grids operated by separate workgroups.

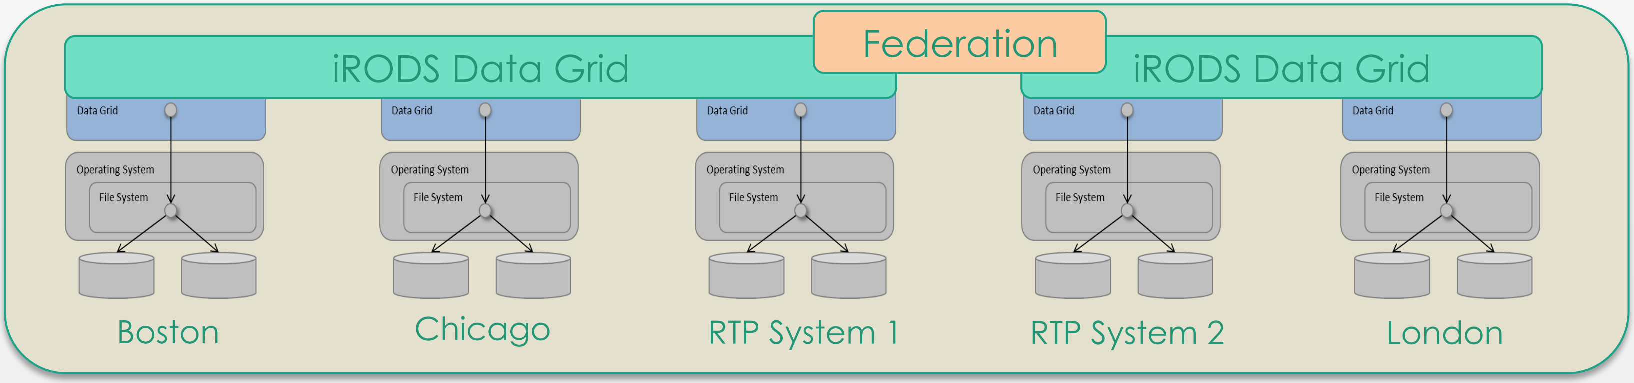

Data Virtualization

iRODS is open source data grid middleware for...

- Data Discovery
- Workflow Automation
- Secure Collaboration
- Data Virtualization

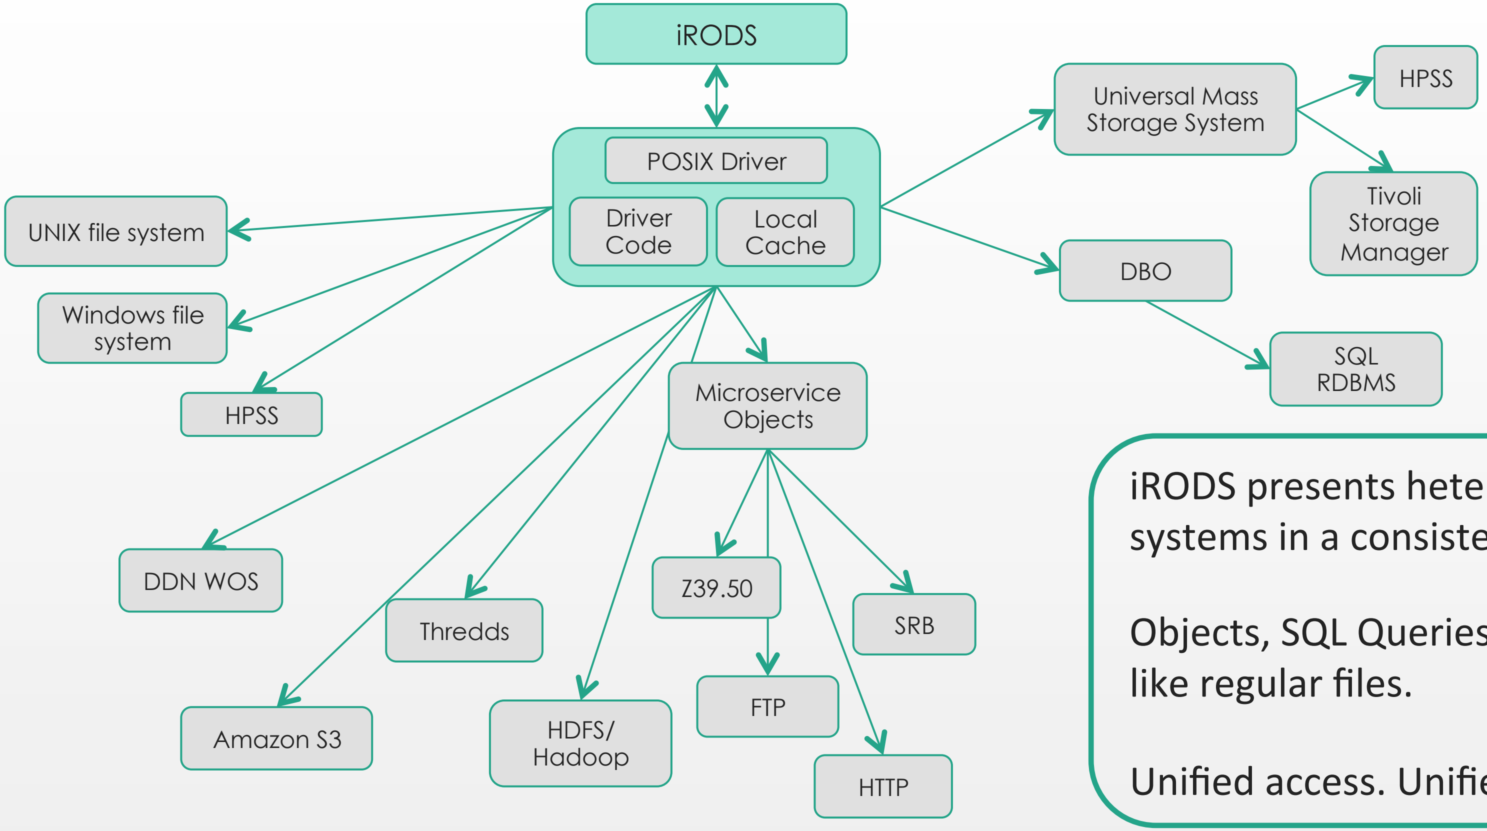

iRODS presents heterogeneous storage systems in a consistent, familiar format.

Objects, SQL Queries, URLs all accessed like regular files.

Unified access. Unified control.

iRODS is Extensible

iRODS has a pluggable architecture.

Existing plug-ins support a variety of hardware, communication technologies, database technologies, and storage topologies. Templates are available for new, custom plug-ins.

Command line, web clients, and numerous other clients are available for iRODS. Generic APIs allow developers to build efficient access to iRODS in to their software.

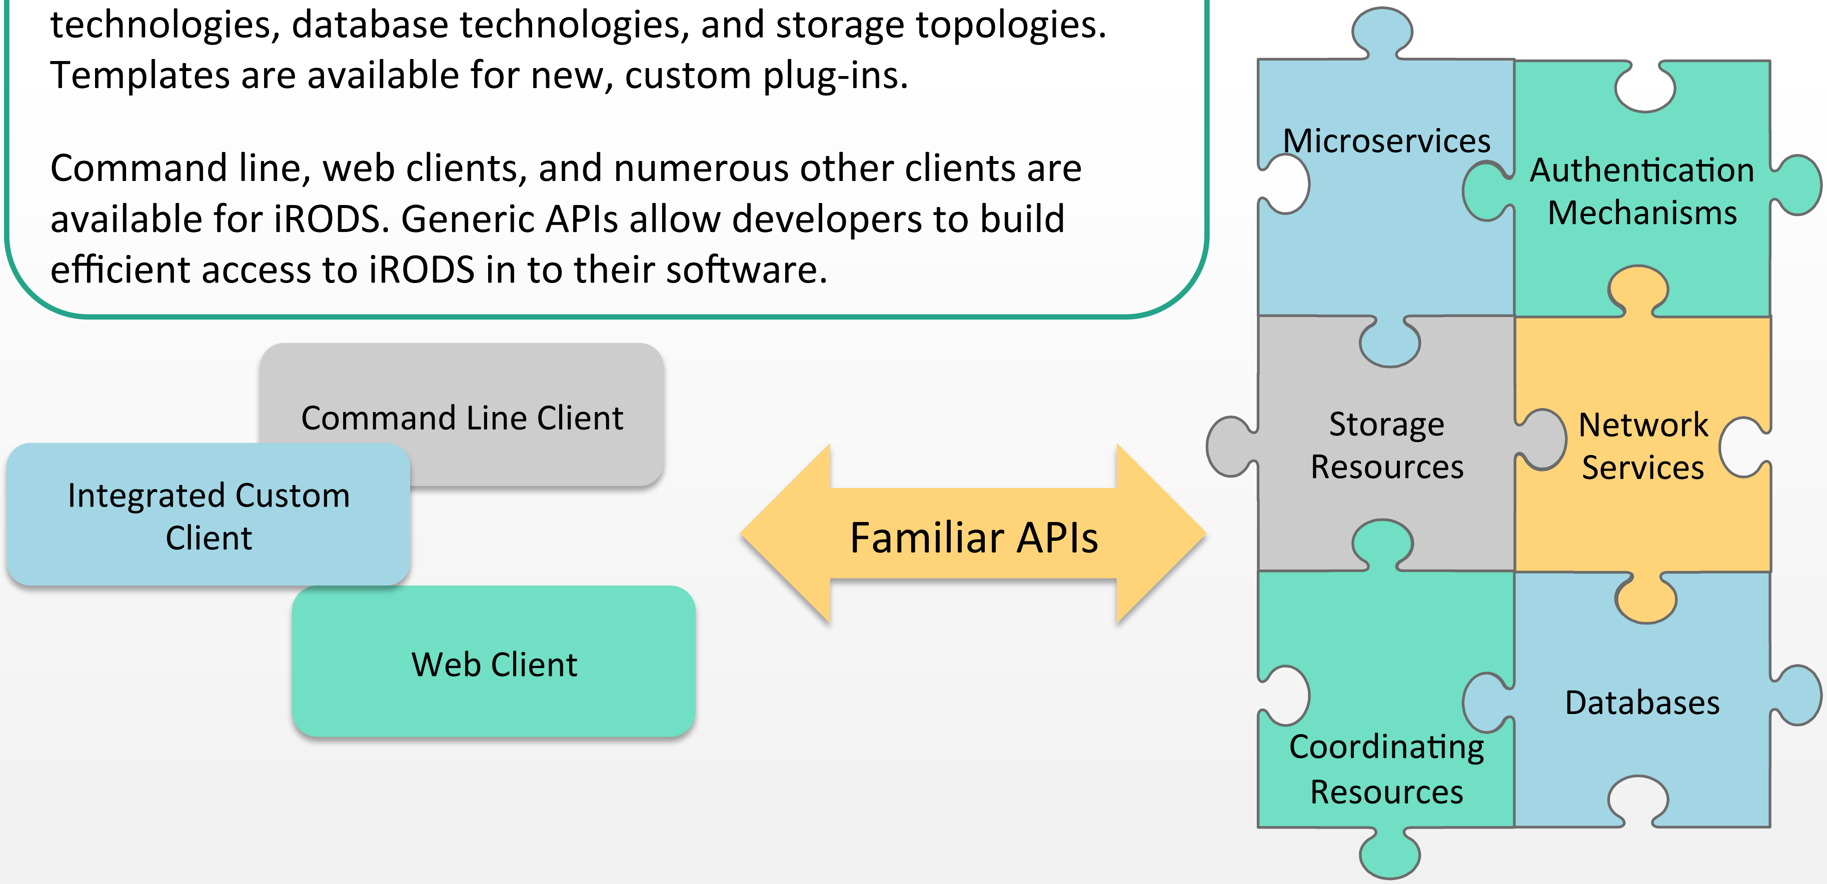

Command Line Client

Integrated Custom Client

Web Client

Familiar APIs

Microservices

Authentication Mechanisms

Storage Resources

Network Services

Coordinating Resources

Databases

WHO USES iRODS?

Who Uses iRODS?

- **Federal Users**
  - National Institutes of Health (NIH)
  - National Aeronautics and Space Administration (NASA)
  - National Oceanic and Atmospheric Administration (NOAA)
  - National Optical Astronomy Observatory (NOAO)
  - US Geological Survey (USGS)
- **Storage Vendors and System Integrators**
  - DataDirect Networks
  - EMC
  - Xyratex
  - Distributed Bio
  - Computer Sciences Corporation (CSC)
- **Commercial Users**
  - DOW Chemical
  - Beijing Genomics Institute
- **Research Programs**
  - The iPlant Collaborative
  - Broad Institute
  - International Neuroinformatics Coordinating Facilities (INCF)
  - Wellcome Trust Sanger Institute
  - Computer Center of the French National Institute of Nuclear and Particle Physics (CC-IN2P3)
  - CineGRID
- **Hundreds of academic institutions** worldwide host thousands of users on their iRODS data grids

iRODS – Proven at Scale

- iPlant: 15,000 users on an iRODS data grid with 100 million files
- IN2P3: over 6 PB of data managed by iRODS
- Sanger Institute: 20+ PB of iRODS data
- NASA Center for Climate Simulations: 300 million metadata attributes
- CineGRID: sites distributed across Japan-US-Europe

WHAT CAN iRODS DO?

What Can iRODS Do?

For Data Center Managers, iRODS simplifies data grid management.

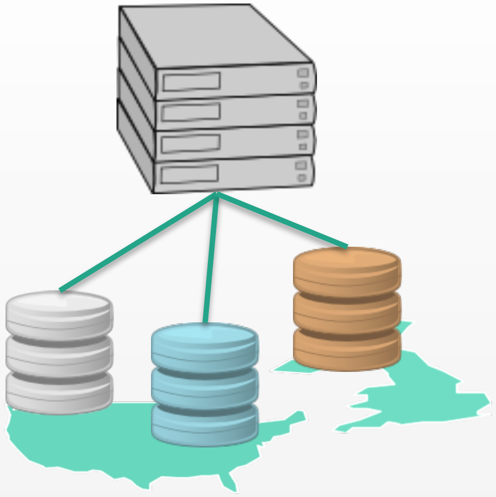

Data on different storage devices at different locations can be centrally managed.

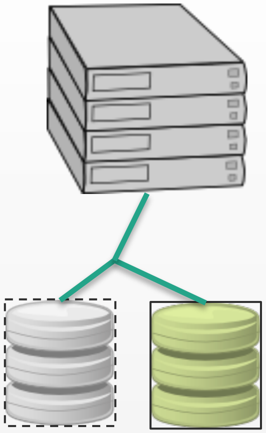

In situ migration to new hardware can be managed by replicating the legacy resource before repurposing or decommissioning it.

BnF (2008)

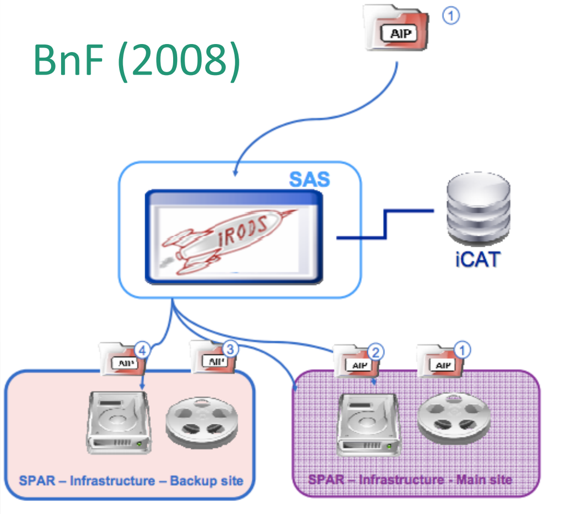
AIP
1
SAS
iRODS
iCAT
AIP 4
AIP 3
AIP 2
AIP 1
SPAR – Infrastructure – Backup site
SPAR – Infrastructure - Main site

Backup and archiving are transparent to the user and highly configurable using the automation and metadata capabilities in iRODS.

What Can iRODS Do?

For Users, iRODS simplifies data discovery, data validation, and data processing.

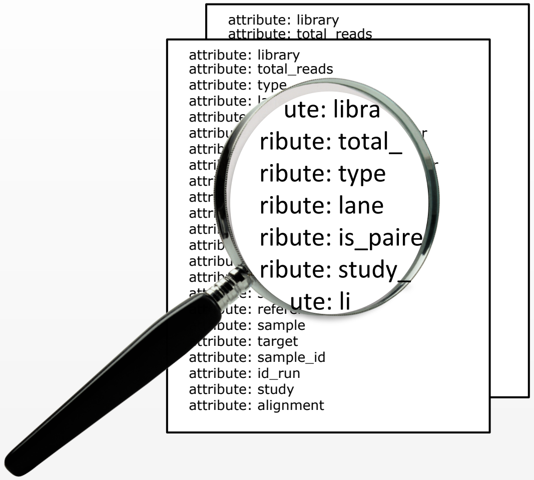

attribute: library
attribute: total_reads

attribute: library
attribute: total_reads
attribute: type
ute: libra
ribute: total_
ribute: type
ribute: lane
ribute: is_paire
ribute: study_
ute: li
ute: refer
ute: sample
tribute: target
attribute: sample_id
attribute: id_run
attribute: study
attribute: alignment

User-defined and intrinsic metadata make stored data searchable.

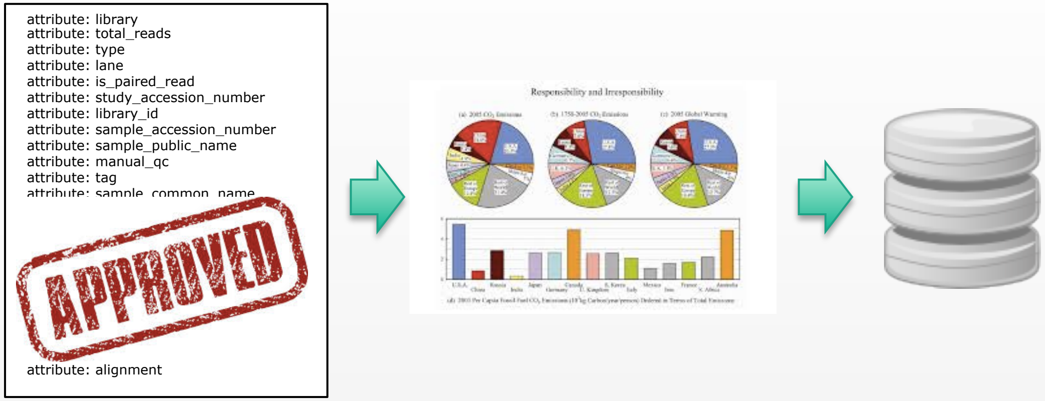

attribute: library
attribute: total_reads
attribute: type
attribute: lane
attribute: is_paired_read
attribute: study_accession_number
attribute: library_id
attribute: sample_accession_number
attribute: sample_public_name
attribute: manual_qc
attribute: tag
attribute: sample_common_name
APPROVED
attribute: alignment

Validation and analytical tools can be automated to process incoming data.

The results and process steps can be stored in the iCAT metadata catalog.

Policy Enforcement and Compliance Verifications

With its metadata catalog and automation capabilities, iRODS presents the infrastructure to enforce mandated data management policies, such as those for records retention and privacy protection.

Audit trails generated by iRODS can be used to verify compliance with policy.

Final Rule:
Retention of R

U.S. Department of Health & Huma

HHS.gov

Improving the health,

HHS Home | HHS News | About HHS

Health Information Privacy

THE FUTURE OF iRODS

The Future of iRODS: Improving Uptake

- Plug-In Bundling: Easier Deployment to Specific Market Segments

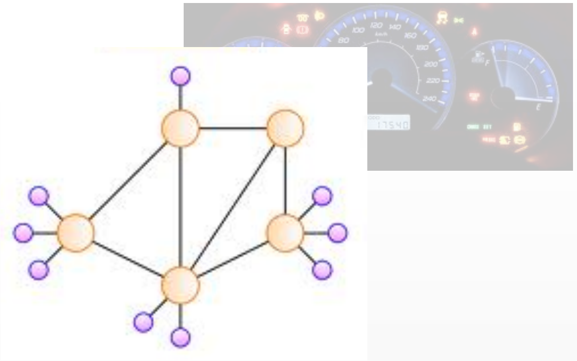

- Registry: Enables Bundling, Easier Upgrades, Dashboard

- Simplifying Connection APIs: Improved Consistency across Programming Languages → More Clients and Plug-Ins

The Future of iRODS: Full Content Indexing

Efforts underway to:

- Perform full content indexing of incoming data
- Connect iRODS to an external database of indexed search terms

Data discovery based on:

- ***Data Content***
- Automatically-generated metadata
- User-generated metadata

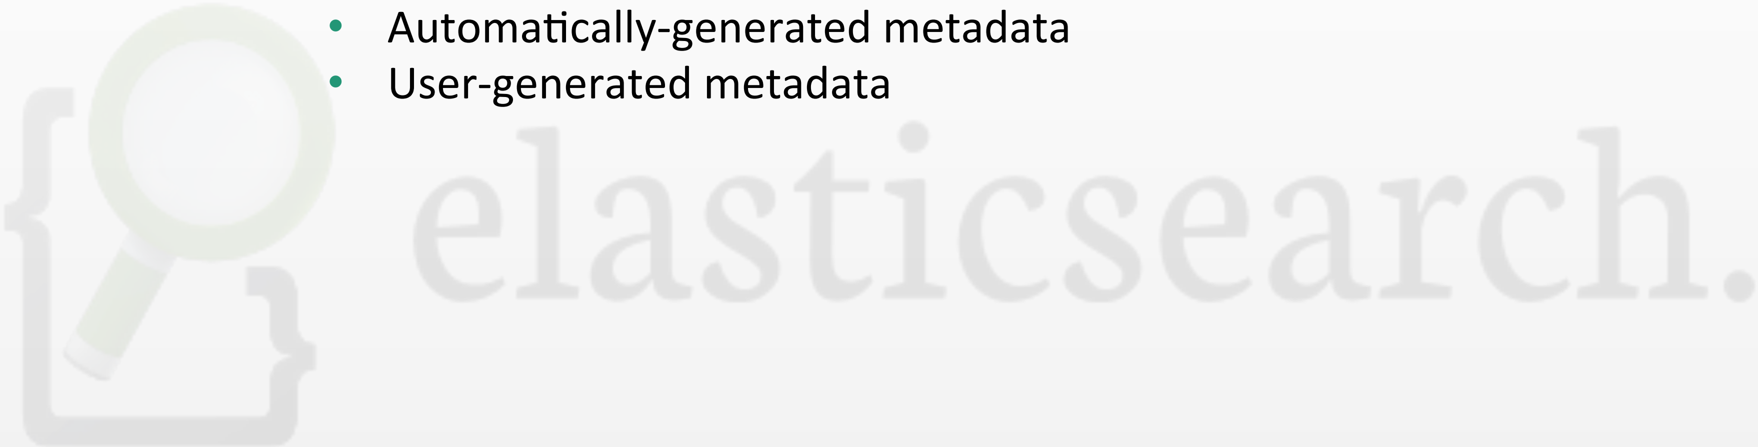

The Future of iRODS: Application-Defined Networking

- iRODS can control a software-defined network (SDN).
- iRODS has been used experimentally with a SDN to maximize bandwidth between storage resources.
- Parallel disjoint network links are dynamically created and torn down in response to network conditions and other parameters.

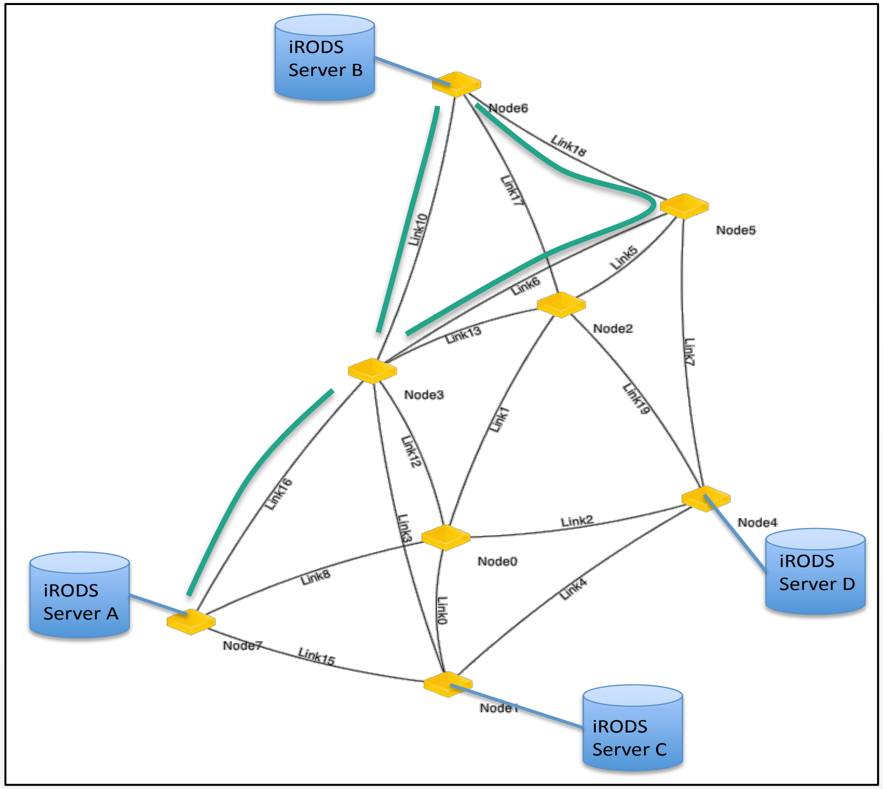

BACKUP SLIDES

Use Case:
University of North Carolina, Chapel Hill (UNC)

Slides courtesy of Charles Schmitt

Genomics Primary Physical Infrastructure

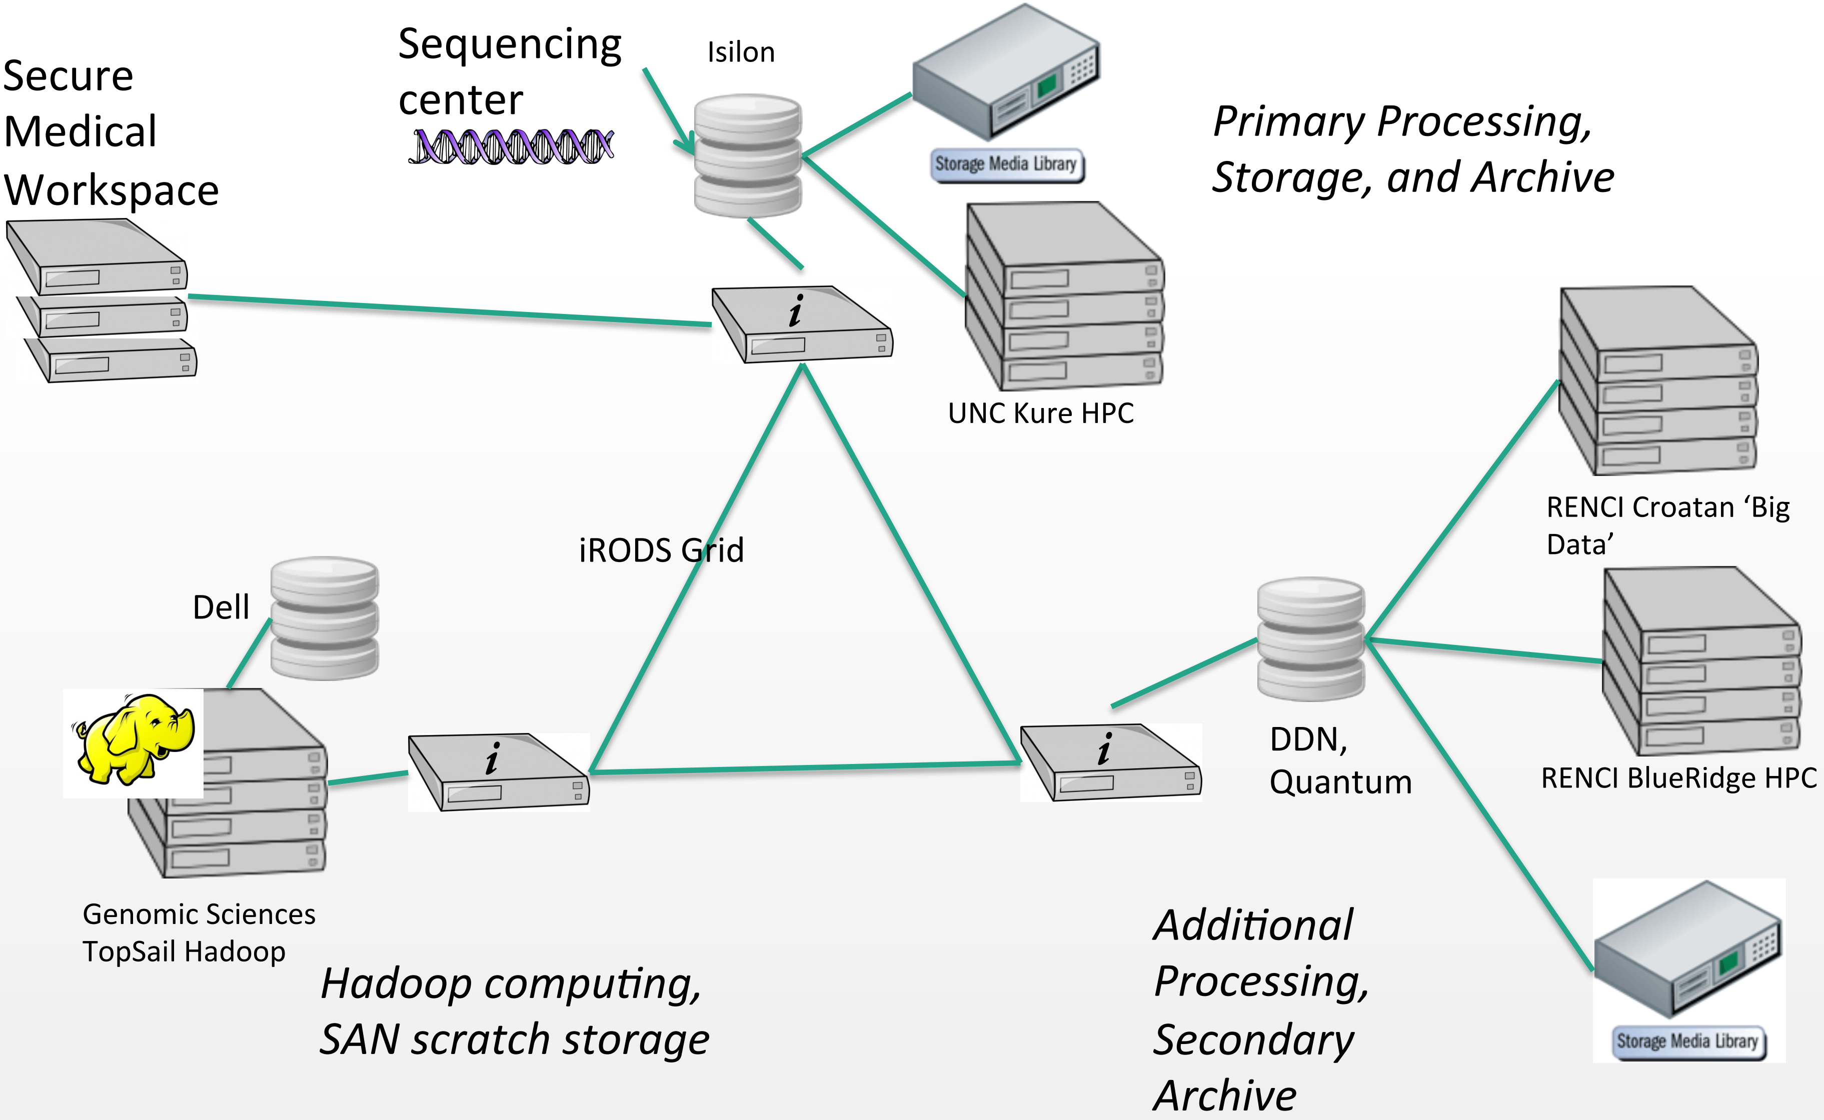

Primary Processing, Storage, and Archive

Hadoop computing, SAN scratch storage

Additional Processing, Secondary Archive

Example: Unified View of Data

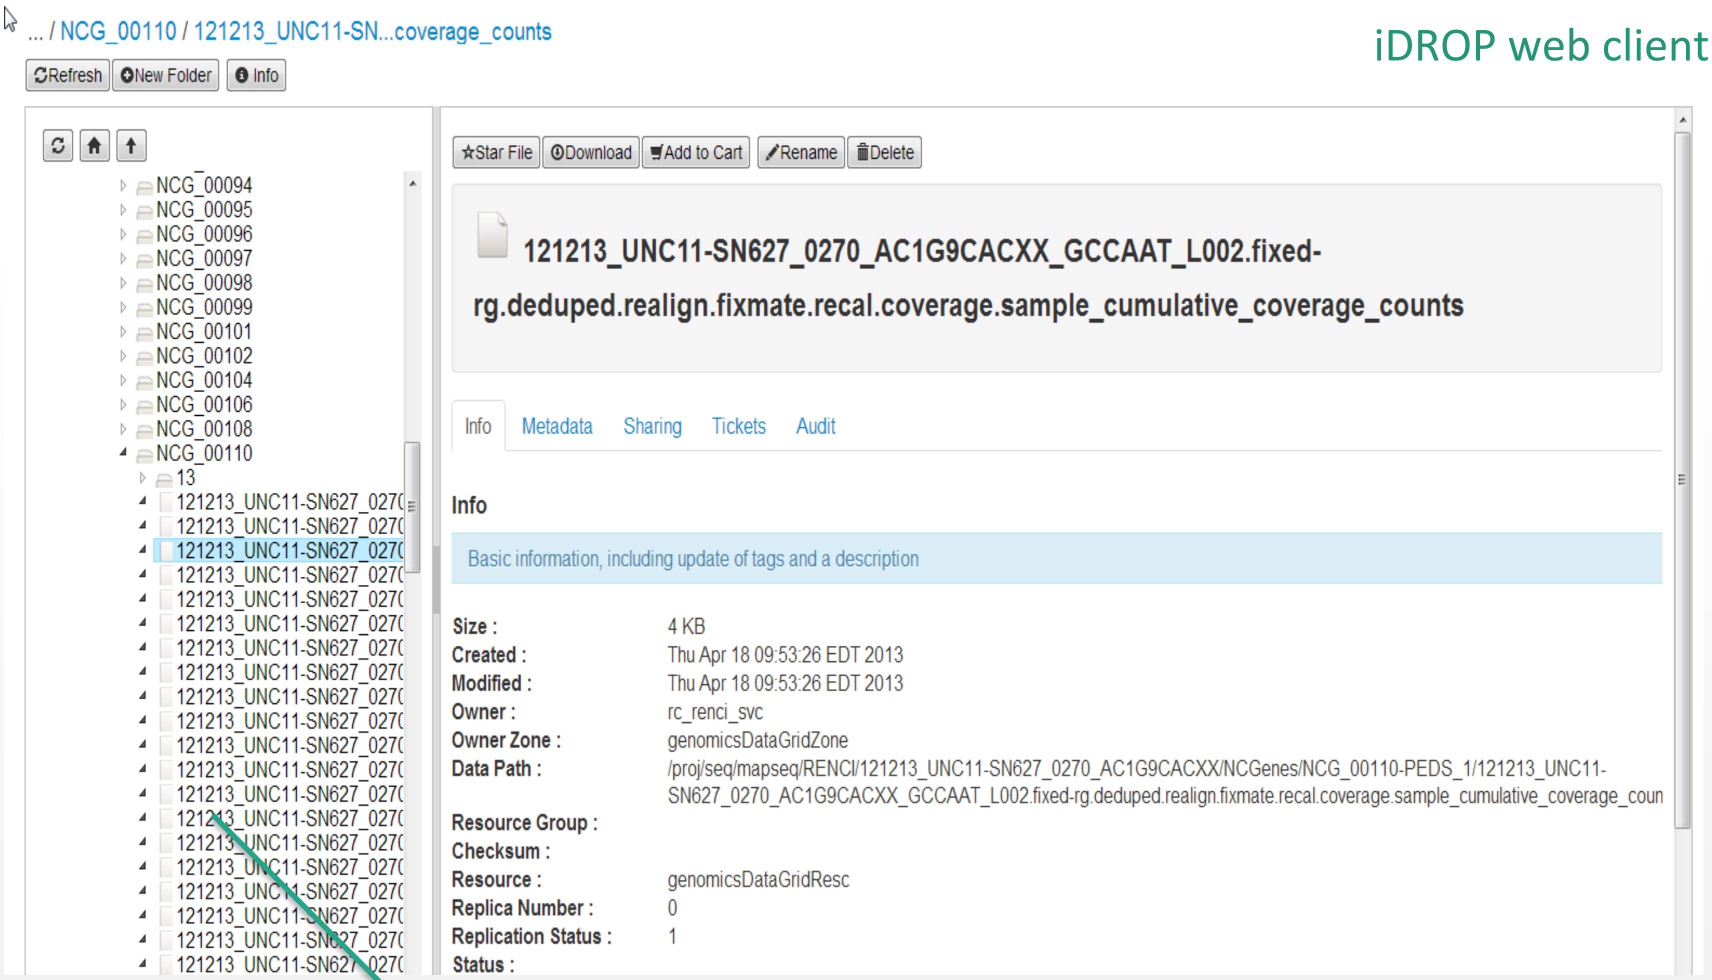

NCG_00094
NCG_00095
NCG_00096
NCG_00097
NCG_00098
NCG_00099
NCG_00101
NCG_00102
NCG_00104
NCG_00106
NCG_00108
NCG_00110
13
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027
121213_UNC11-SN627_027

121213_UNC11-SN627_0270_AC1G9CACXX_GCCAAT_L002.fixed-rg.deduped.realign.fixmate.recal.coverage.sample_cumulative_coverage_counts

Info | Metadata | Sharing | Tickets | Audit

Info

Basic information, including update of tags and a description

Size :	4 KB
Created :	Thu Apr 18 09:53:26 EDT 2013
Modified :	Thu Apr 18 09:53:26 EDT 2013
Owner :	rc_renci_svc
Owner Zone :	genomicsDataGridZone
Data Path :	/proj/seq/mapseq/RENCI/121213_UNC11-SN627_0270_AC1G9CACXX/NCGenes/NCG_00110-PEDS_1/121213_UNC11-SN627_0270_AC1G9CACXX_GCCAAT_L002.fixed-rg.deduped.realign.fixmate.recal.coverage.sample_cumulative_coverage_coun
Resource Group :	
Checksum :	
Resource :	genomicsDataGridResc
Replica Number :	0
Replication Status :	1
Status :	

iDROP web client

Spread across:
1) Disk-storage at UNC, 2) Disk-storage at RENCI, 3) Tape-storage at RENCI

Example: Data Access Policy

- Challenge
  - Millions of files across different projects, growing daily
  - Hundreds of users across different labs, changing frequently
  - *How to control access*
    - *UNIX ACLs became too unwieldy*
    - *Moving data means reproducing permission and group settings*

- Policy: access given if user and data belong to the same groups
  - Tag data with group metadata (e.g., Lab X lung tumor study)
  - Access rule: user's group must match data group
    - E.g. (user y member of Lab X lung tumor study)

Thanks to Sai Balu at LCCC

Example: Data 'Replication' Policy

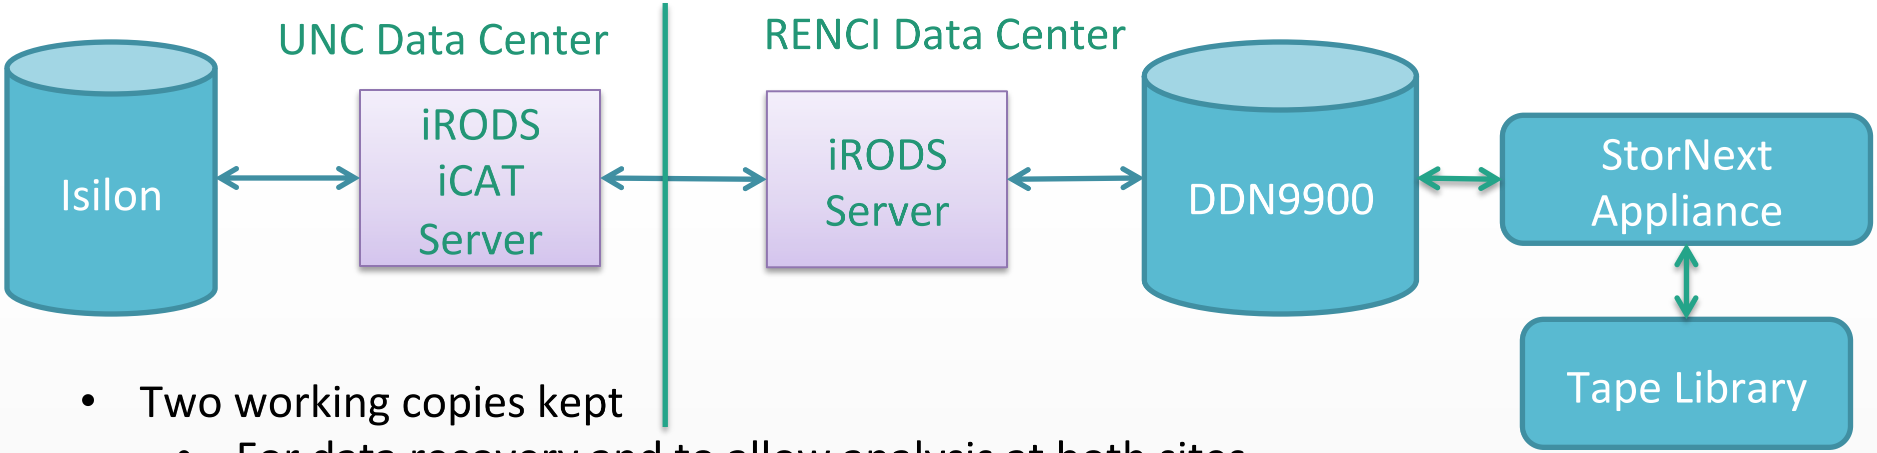

- Two working copies kept
  - For data recovery and to allow analysis at both sites
- 'Copy me' and 'Data copied' metadata control copy process
  - Only on certain files (fastq, 'finished' bam files)
- iRODS rule performs the copy nightly
  - Performs copy, verifies copy successful, resets 'copy me' attribute
- Versioning to allow for re-runs of patient samples

Secure Access to Data on the Clinical Side

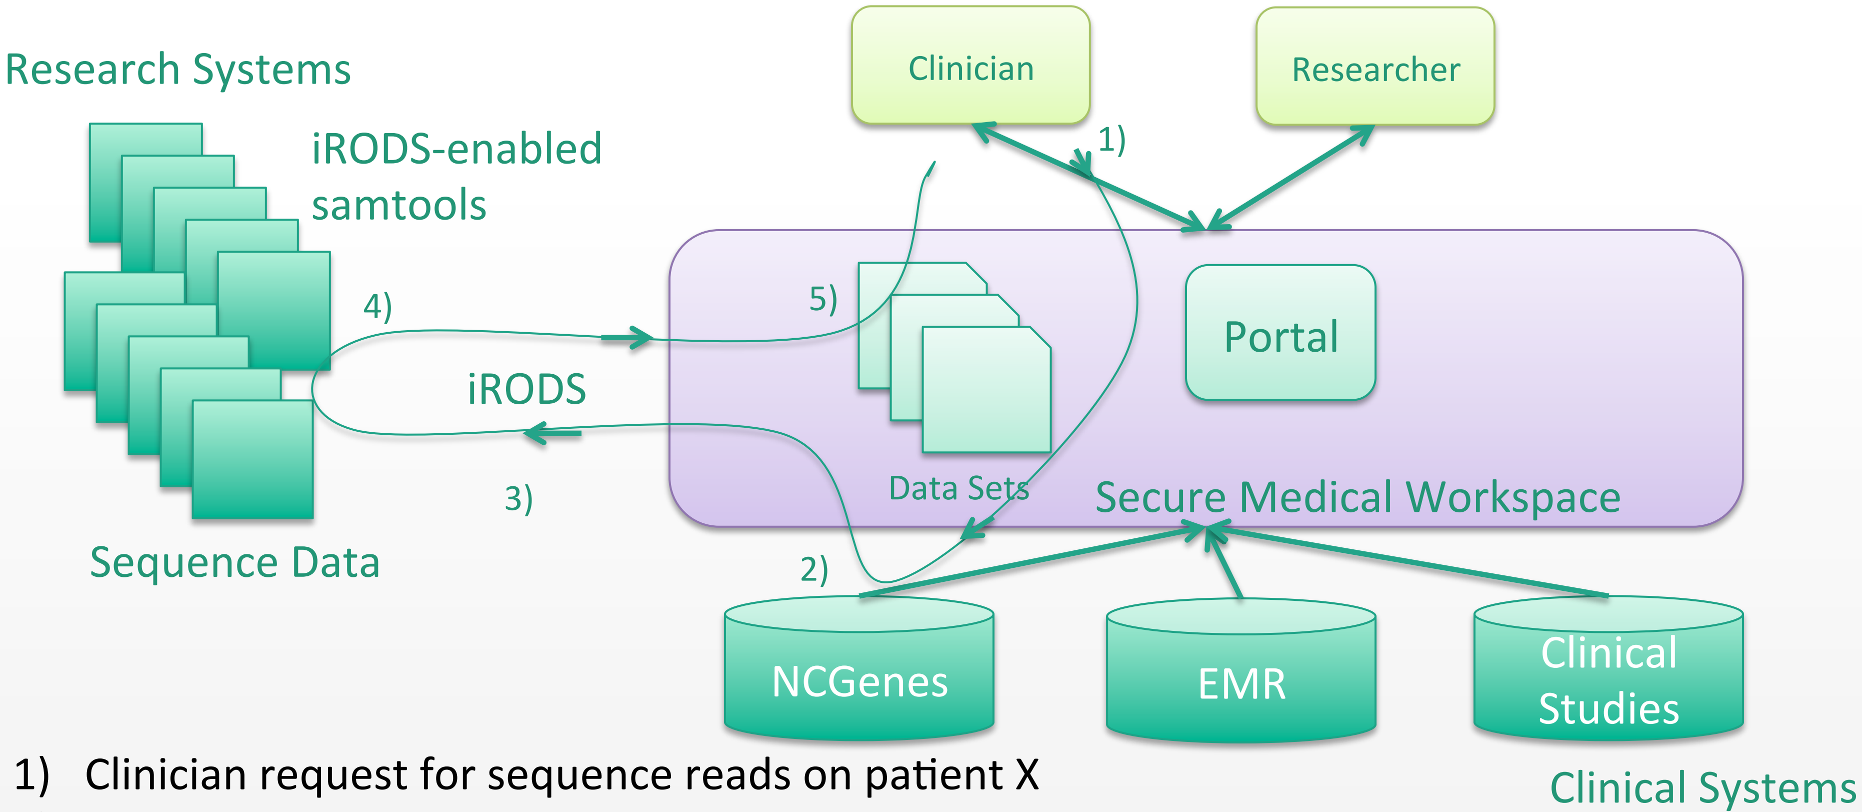

1) Clinician request for sequence reads on patient X
2) Patient id lookup to obtain subject id
3) Subject id lookup in iRODS
4) Data sets packaged in zip file and retrieved
5) Data unzipped and displayed within secure workspace

Use Case:
The Sanger Institute
(Wellcome Trust Sanger Institute)

Slides courtesy of Peter Clapham

iRODS layout

Data lands by preference onto iRES servers in the green data center room

Data is then replicated to red data center room via a resource group rule with checksums added along the way

Both iRES servers are used for r/o access and replication does work either way if bad stuff happens.

Various data and metadata integrity

Checks are made.

Simple, scalable and reliable

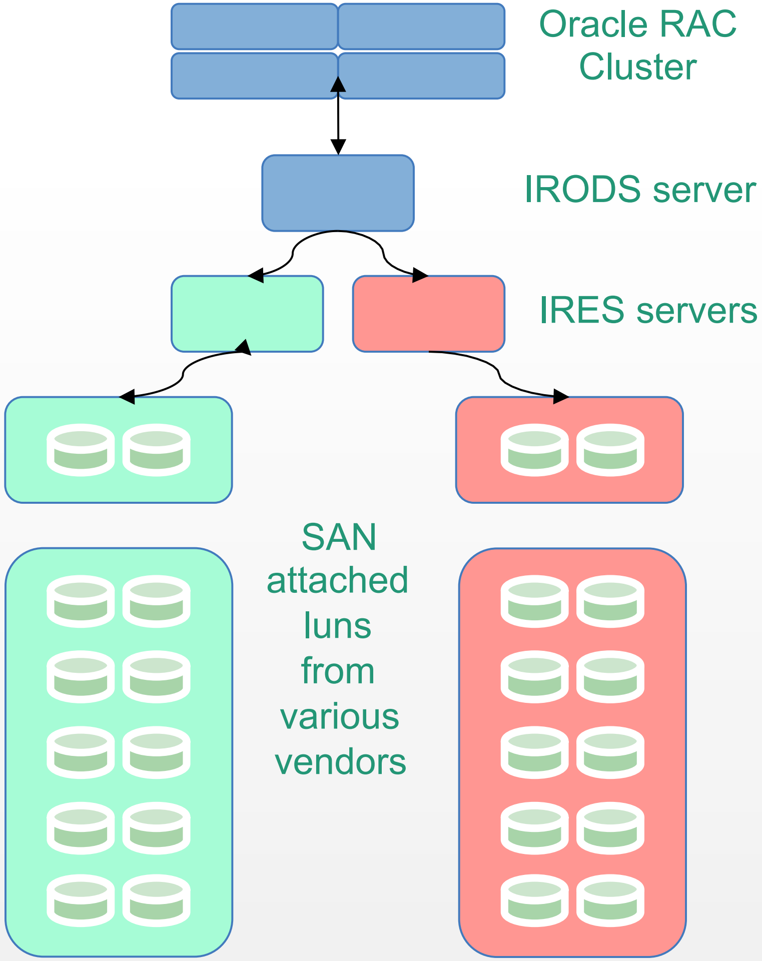

Metadata-Rich

Example attribute fields →

Users query and access data largely from local compute clusters

Users access iRODS locally via the command line interface

```
attribute: library
attribute: total_reads
attribute: type
attribute: lane
attribute: is_paired_read
attribute: study_accession_number
attribute: library_id
attribute: sample_accession_number
attribute: sample_public_name
attribute: manual_qc
attribute: tag
attribute: sample_common_name
attribute: md5
attribute: tag_index
attribute: study_title
attribute: study_id
attribute: reference
attribute: sample
attribute: target
attribute: sample_id
attribute: id_run
attribute: study
attribute: alignment
```

Sanger Zone Arrangement

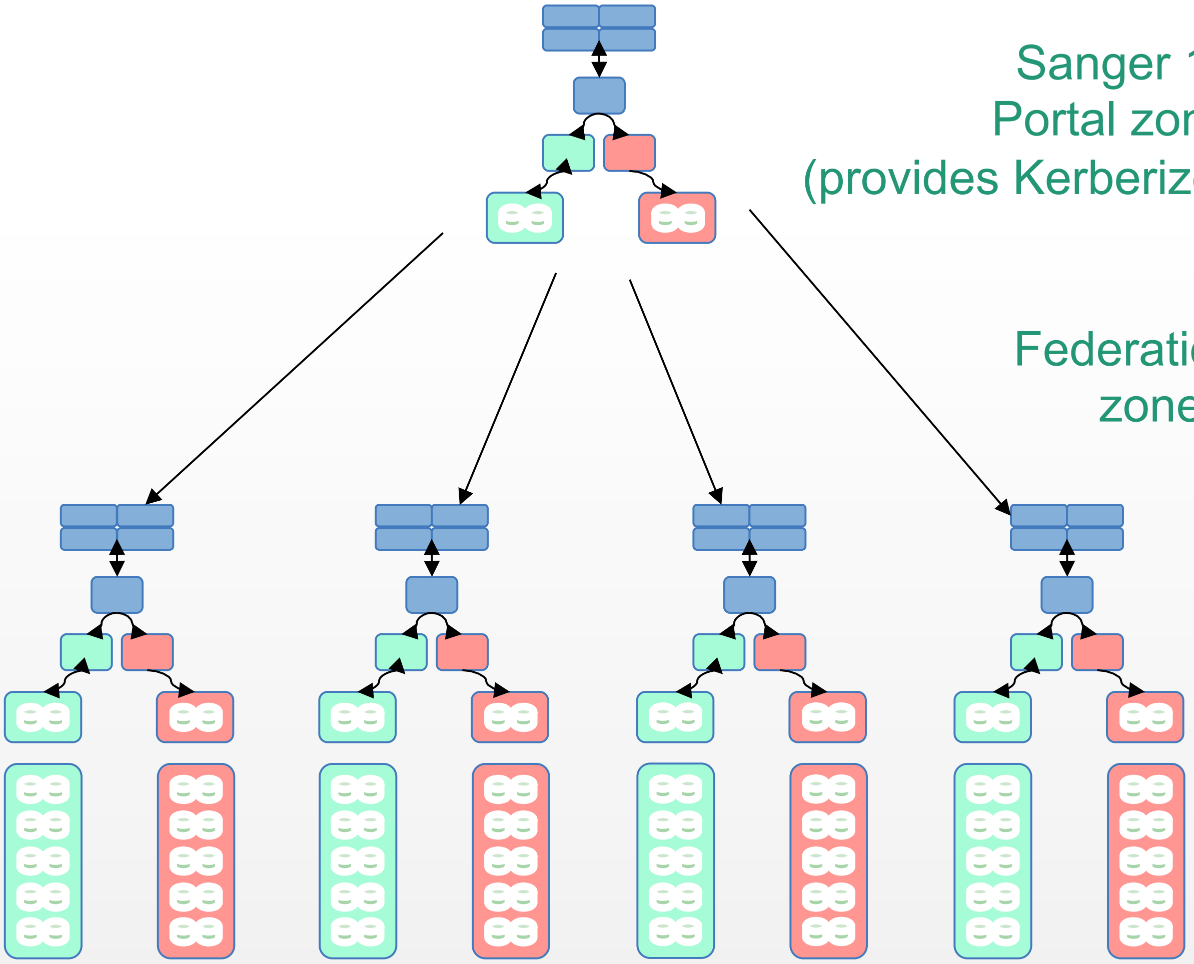

Sanger 1
Portal zone
(provides Kerberized access)

Federation using head zone accounts

/seq /uk10k /humgen /Archive

Baton iRODS "Client"

Thin layer over parts of the iRODS C API

- JSON support
- Connection friendly
- Comprehensive logging
- autoconf build on Linux and OSX

Current state

- Metadata listing
- Metadata queries
- Metadata addition

https://github.com/wtsi-npg/baton.git

Use Case:
Beijing Genomics Institute (BGI)

Slides courtesy of Xing Xin

EasyGenomics BGI

- **The world largest genome sequencing center**
  - Started with Human Genome Project in 1999 with only a few sequencers.
  - Now more than 150 sequencers, 6 TB/day sequencing throughput.

MODEL	ABI 3730XL	Roche 454	ABI SOLiD 4	Solexa GA IIx	Illumina HiSeq 2000
INSTALLATION	16	1	27	6	135

Data Flow - iRODS

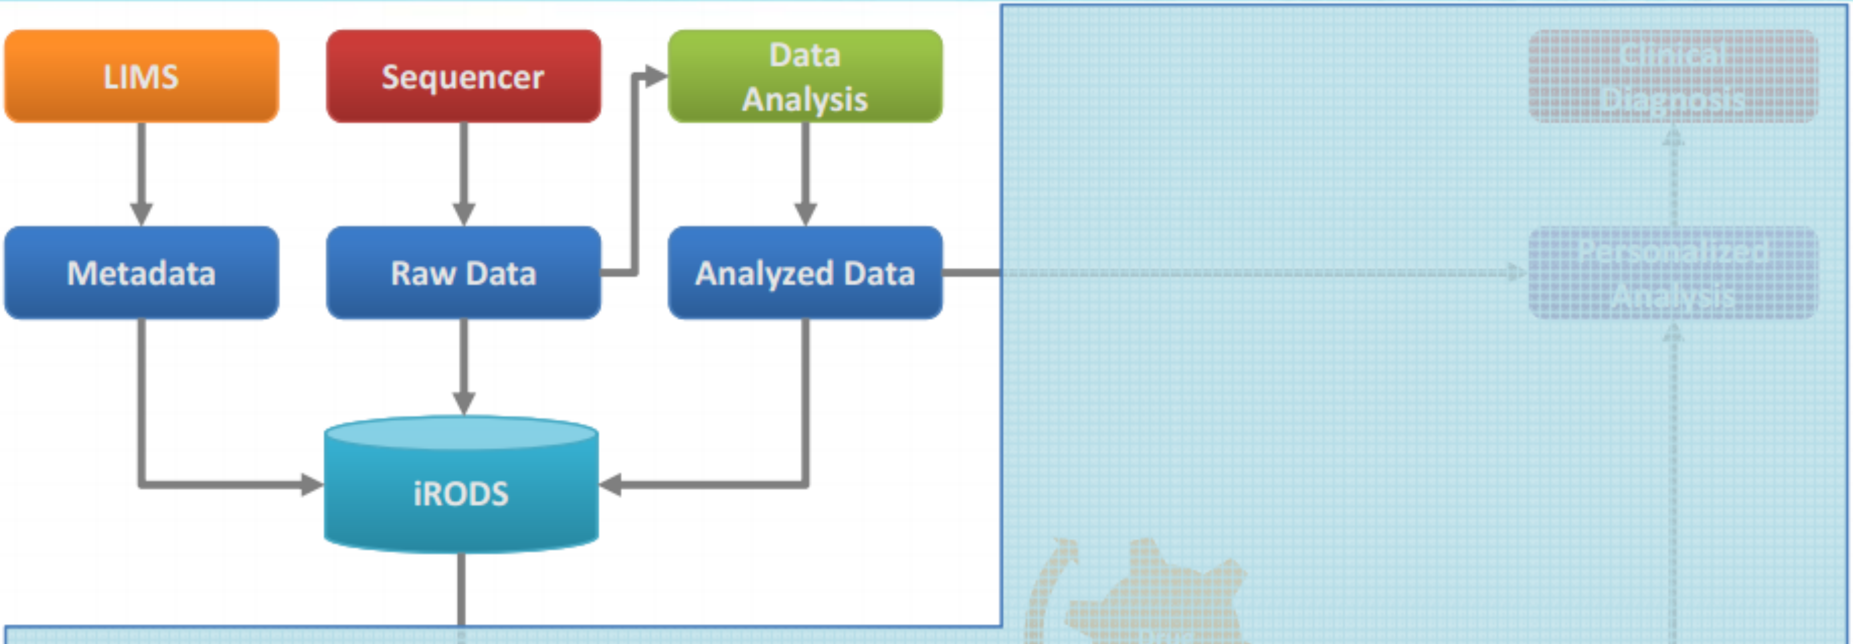

iRODS-based Data Management

- Contents: raw data, analyzed data and related metadata
- Data backup
- Fully integrated with LIMS
- Able to search and access any data according to the metadata from BGI data standard, *e.g.* project, sample, cohort, phenotype, QC, *etc.*
- Federation: integrate separate iRODS zones

Acknowledgements

- Presented work funded in part by grants from NIH, NSF, NARA, DHS. Funding also provided by UNC and the iRODS Consortium.
- Teams involved include:
  - *DICE team at UNC and UCSD*
  - *Networking team at RENCI and Duke*
  - *Data sciences team at RENCI*
- *In collaboration with*
  - *UNC Dept of Genetics, Research Computing, Lineberger Comprehensive Cancer Center, NC TraCS Institute, Center for Bioinformatics, Institute for Pharmacogenetics and Personalized Treatment*
  - *UNC HealthCare*
- *Multiple members of the iRODS community*

Data Policy and iRODS - Definitions

- What is (Digital) Data Management?
  - procedures and operations to assure value of digital assets
    - protection: verification, back-ups & replicas, security, access controls, …
    - maintenance: migrating to tape, integrity checking, …
    - control: format conversion, derived data products, …
    - accessibility and usability: discovery, availability, supporting services, …
  - defined by data proprietors or data administrators
  - enables analytics and operations that pull value from the data
- What is data policy?
  - statement of data management strategy
  - the ensemble of data management requirements and procedures
- Data policy can be *implemented* and *executed* manually or in automated fashion

iRODS History

- SRB: initial product begun by DICE, 1997 at the San Diego Supercomputer Center, UCSD and General Atomics
- iRODS: rewrite of SRB by DICE in 2006; current version: iRODS 3.3.1
- Very close interaction with worldwide user communities who drive development
- Enterprise iRODS (e-iRODS): mission critical distribution co-developed by RENCI and DICE in 2012
- iRODS 4.0: merge of the iRODS and e-iRODS codes by iRODS Consortium to form a common core and full deployment of plug-in architecture