



Technology solutions for managing and computing on large- scale biomedical data

Charles Schmitt
CTO & Director of Informatics
RENCI

Brand Fortner
Executive Director, iRODS Consortium

Jason Coposky
Chief Technologist, iRODS Consortium

renci


RESEARCH \ ENGAGEMENT \ INNOVATION

Acknowledgements

- Presented work funded in part by grants from NIH, NSF, NARA, DHS, as well as funding from UNC
- Teams involved include:
 - *DICE team at UNC and UCSD*
 - *Networking team at RENCI and Duke*
 - *Data sciences team at RENCI*
 - *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*


iRODS – an Overview

What is the iRODS Data Grid Middleware?



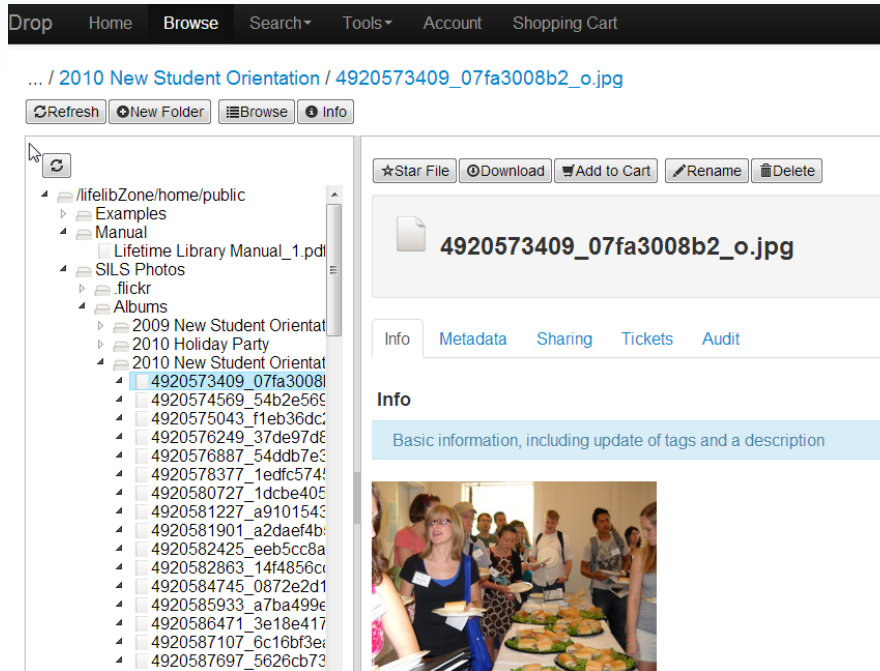
- Applications access data through OS managed file systems
 - RDBMS, web services, office tools, ...
- Standardized interface
 - POSIX operations
 - Limited metadata
 - Hierarchical filesystem
- Alternative interfaces to storage are becoming more common:
 - Object Stores, Cloud Stores

What is the iRODS Data Grid Middleware?

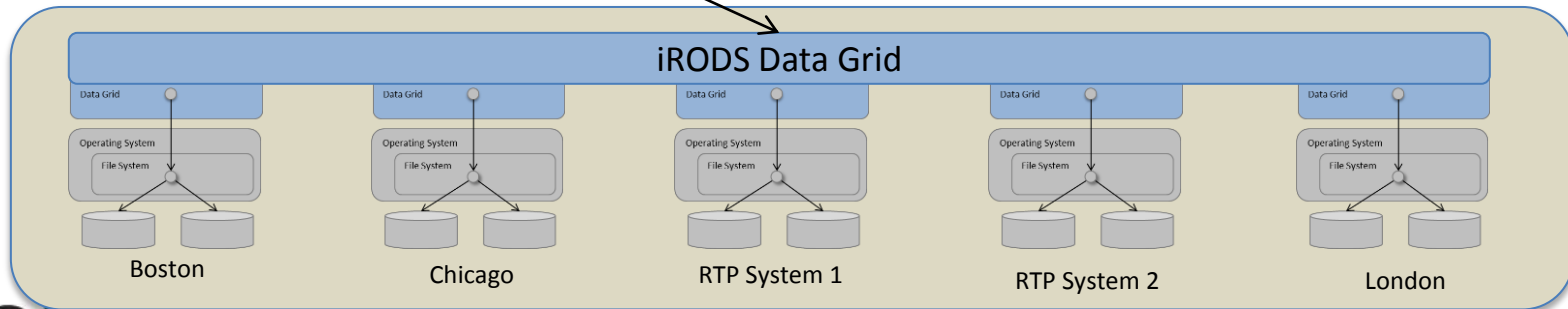


- A middle layer, providing:
 - Uniform interface to POSIX and non-POSIX data systems
 - Running in user space
- Plus, lots of additional capabilities

Middleware for Distributed Data



- Distributed discovery and access
 - By data organization
 - By metadata
- Unified namespace
- Multiple client applications and APIs
- *Breaks dependency between applications/users and where/how data is stored*



Centralized Metadata for Distributed Data

Example Metadata:

Logical Name: /RDDept/LabX/Flow/Study1

Physical Name: /London/var1/proj/labx/stuff

Lab PI: Jane Doe

Date: 12/1/2010

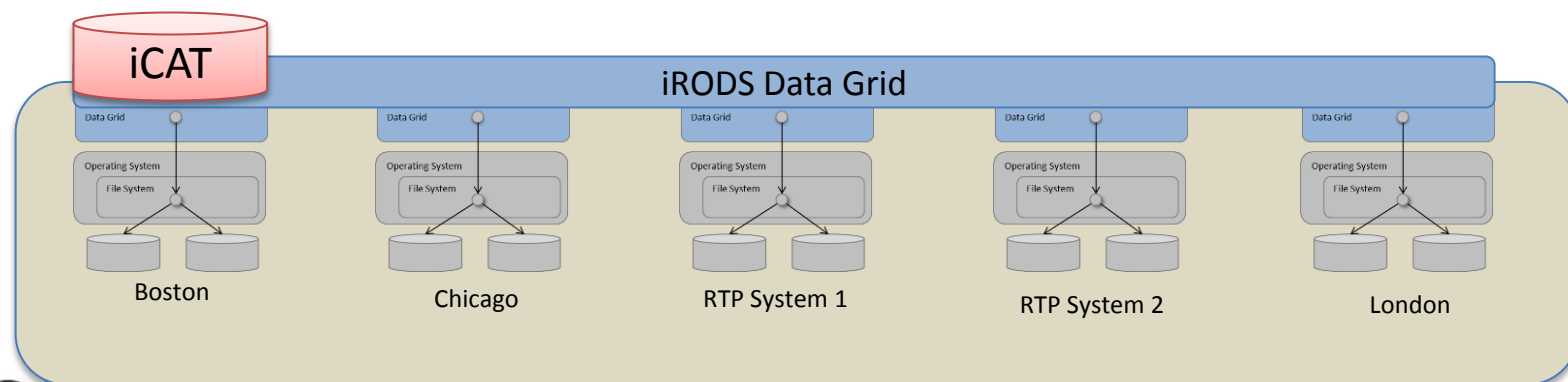
Time: 01:45:12

Study: Proliferation optimization studies

Data Source: Flow Cytometer

Assay Conditions: Data captured under

- Centralized database for metadata
- Access and search by metadata
 - Across data grid
- Open schema, open editing
 - Or restricted by policies
- Driver for more powerful capabilities

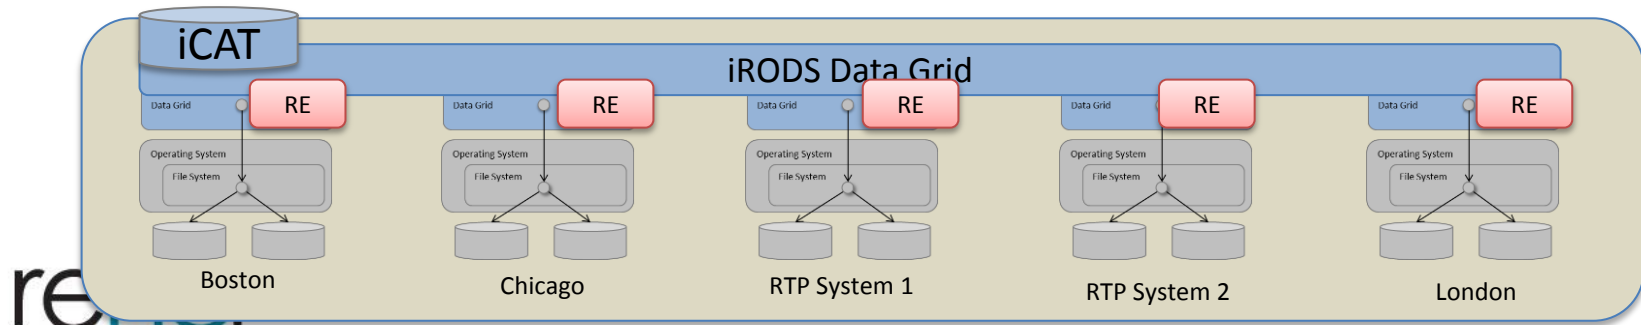


Executable Policies

Manage Organizational and Community Policies

- automate administrative tasks
 - E.g., replicate data to partner labs
- enforce management policies
 - E.g., delete all data older than 1 month
- evaluate assessment criteria
 - E.g., validate checksums on all data files

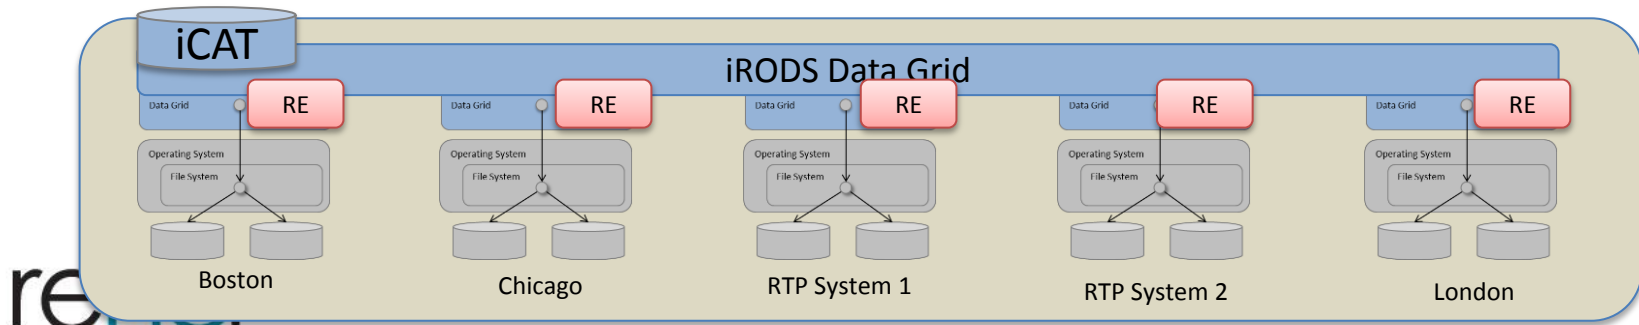
- Policies enforced by Rules:
- Rules executed by Rules Engines (RE) distributed across the grid
- *A way to automate and manage distributed data management tasks*



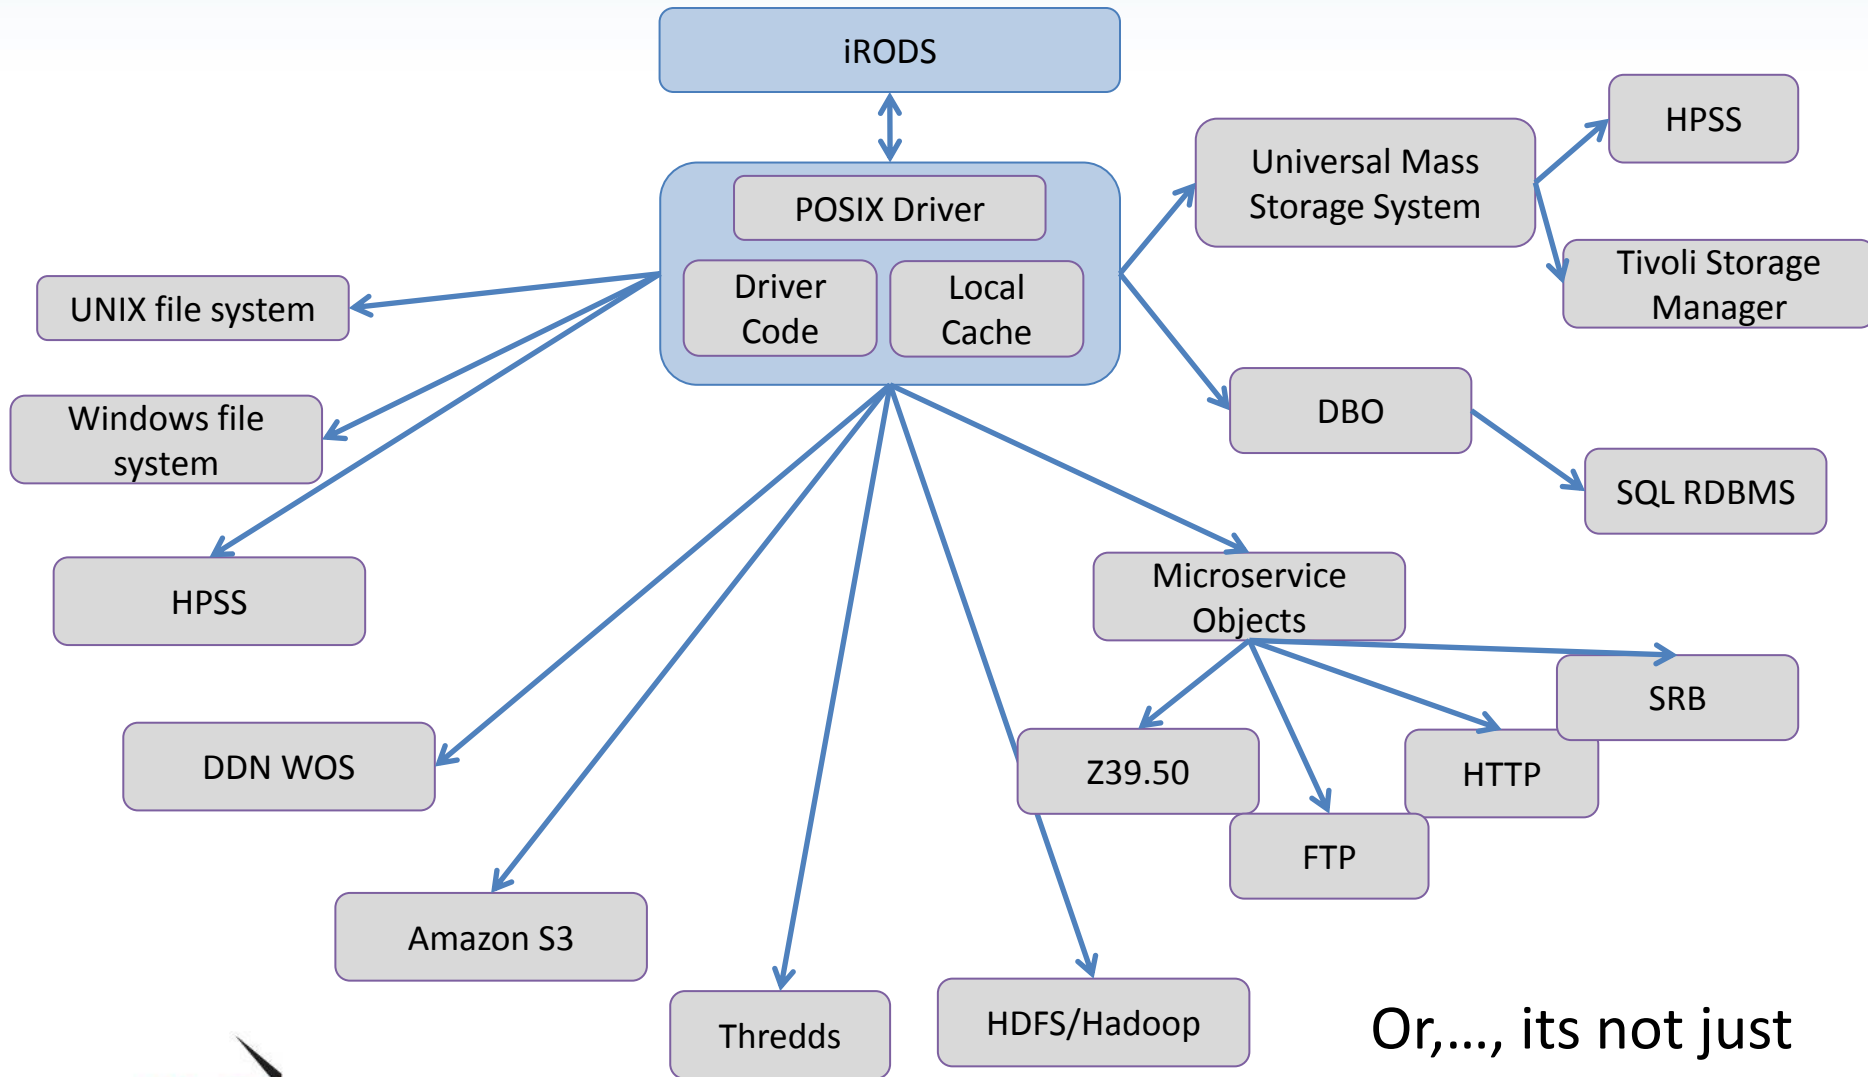
Executable Policies

Complex policies , even data management workflows, can be automated

- Example:
 - On file ingest to iRODS, detect if data is image. If so, tag the file as 'image=yes', then
 - Generate multi-resolution thumbnails, then
 - Replicate thumbnail to storage systems used by web servers
 - On a nightly basis, run image analysis on files tagged as 'image=yes' and tag image as 'embedded barcode=yes/no'
 - On tagging file as 'embedded barcode=yes'
 - Add image file and barcode to product catalog database

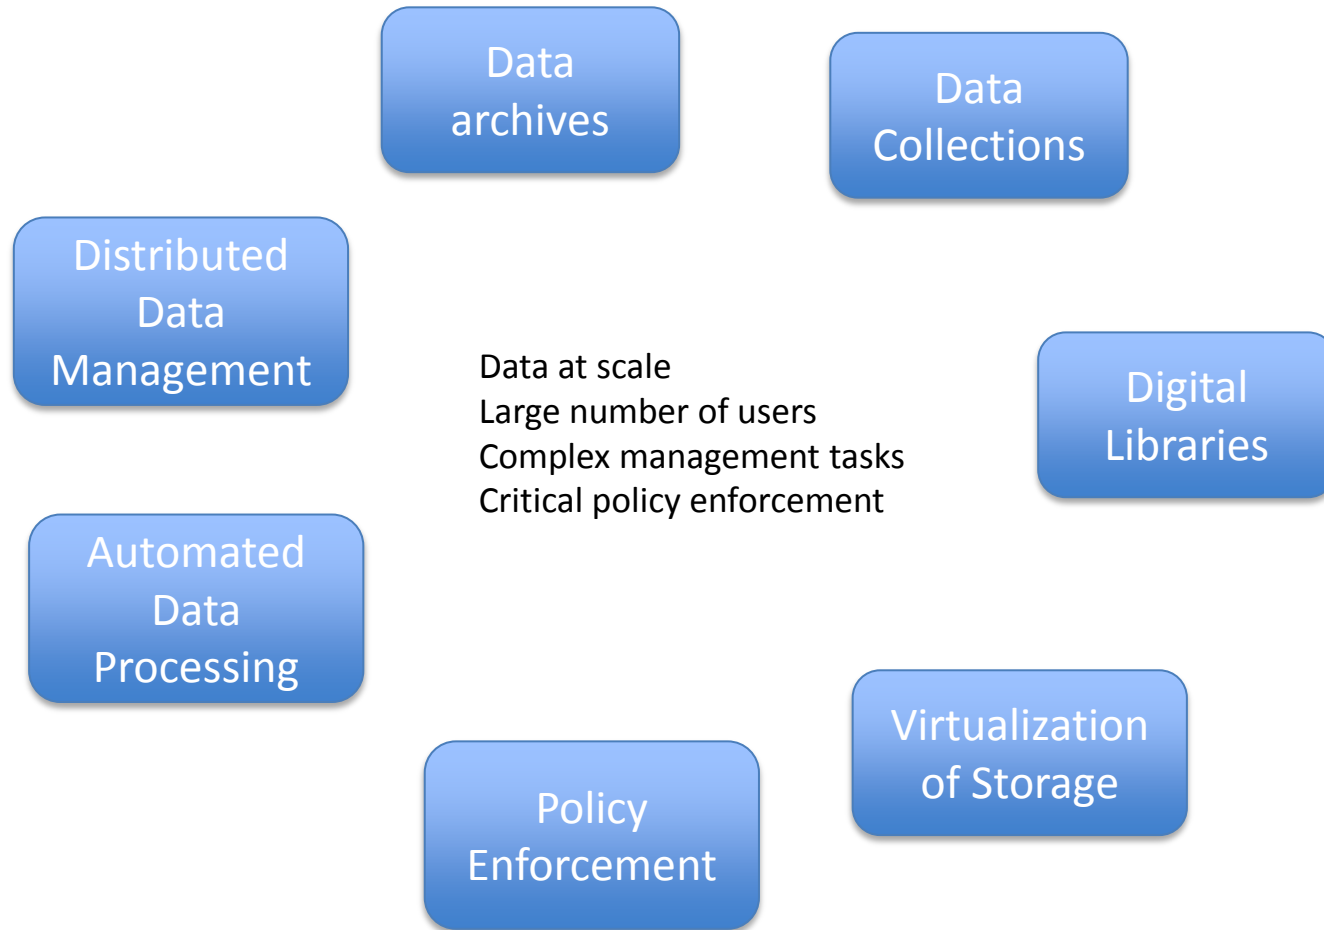


Storage Virtualization



Or,..., its not just files

What is iRODS used for?



iRODS Users - examples

- **Federal Users**
 - National Aeronautics and Space Administration
 - National Oceanic and Atmospheric Administration
 - National Archives and Records Administration
 - National Optical Astronomy Observatory
 - USGS
- **Non-profit/Institutional Users**
 - Broad Institute
 - International Neuroinformatics Coordinating Facilities
 - Wellcome Trust Sanger Institute
 - French national high-performance computing center (CCIN2P3)
 - CineGRID
- **Commercial Users**
 - DOW Chemical
 - Beijing Genome Institute
- **Resellers/Redeployers**
 - Distributed Bio
 - Computer Sciences Corporation (CSC)
 - Cleversafe
- **Academic Users**
 - Too many too list

iRODS Users - examples

- Proven at scale:
 - iPlant - 15k users
 - French National Institute for Nuclear Physics and Plasma Physics – over 6 PB
 - Wellcome Trust Sanger Institute – 20+ Pb
 - Australian Research Collaboration Service - 1024 storage resources
 - NASA Center for Climate Simulations - 300 million attributes
 - Cinegrid – sites across Japan-US-Europe
- Solid foundation:
 - SRB: initial product (developed by DICE Group, owned by General Atomics) in 1997
 - iRODS: rewrite of SRB by DICE Group in 2006; currently on version 3.3
 - Enterprise iRODS: mission critical distribution co-developed by RENCi and DICE in 2012
- Support:
 - Community of developers from groups worldwide
 - Independent groups offering consulting and support and development
 - iRODS Consortium offering formal support, training, involvement, and development help

Policies – the unique aspect of iRODS

Policy-driven data management

- Policies are intentions about data management
- iRODS implements policies as executable rules
- iRODS provides Policy Enforcement Points to enact rules
 - Action: create user, delete user, data delete, ...
 - PreAction: modify metadata, modify access control, ...
 - PostAction: create resource, delete resource, ...
- iRODS provides State information to guide rules
 - Metadata, grid status, user status, ...

iRODS Rules: defining the policies

- Server-side workflows
 - Action | condition | workflow chain | recovery chain
- Action is (possibly parameterized) name for rule
- Condition - test on any attribute:
 - Collection, file name, storage system, file type, user group, elapsed time, descriptive metadata
- Workflow chain:
 - Micro-services / rules that are executed at the storage system
- Recovery chain:
 - Micro-services / rules that are used to recover from errors

iRODS Micro-Services: support the rules

- Functional snippets that wrap a well-defined process
 - Compute checksum
 - Replicate file
 - Integrity check
 - Zoom image
 - Get tiff image cutout
 - Downloads data from NCBI
- Written in C or Python
 - Web services, external applications, can be wrapped as micro-services
- Can be chained to perform complex tasks

iRODS Micro-Services

- Over 300 published microservices
- Pluggable: write, publish, re-use

Examples:

msiRenameLocalZone

msiRemoveKeyValuePairsFromObj

msiTarFileExtract, msiTarFileCreate

msiWriteRodsLog

msiSetACL

msiGetDataObjACL

msiLoadMetadataFromDataObj

msiGetUserInfo

msiGetUserACL

msiMergeDataCopies

....



iRODS Primer: integrated Rule-Oriented Data System (Synthesis Lectures on Information Concepts, Retrieval, and Services) Paperback –

January 19, 2010

by [Arcot Rajasekar](#) (Author) , [Reagan Moore](#) (Author) , [Chien-yi Hou](#) (Author) , [Christopher A. Lee](#) (Author) & 9 more

★★★★★ ▾ 1 customer review

Paperback
\$31.50

10 Used from \$39.33

19 New from \$28.54

Rules – some examples

```
Rule_name (*arguments) {  
  on (expression) {  
    microservice1 (...,*arguments);  
    microservice2 (...);  
  }  
}
```

```
acPostProcForPut {  
  on ($objPath like "*.txt") {  
    msiDataObjCopy($objPath, "$objPath.copy");  
  }  
}
```



Existing microservice

Rules – some more examples

```
//  
// Set access control settings on input of data to irods  
//  
acPostProcForPut {  
    on ( $objPath like /imxgam/*/rawdata/* ) {  
        msiSetACL("default ", "write", "admins", $objPath);  
        msiSetACL("default ", "read", "experts", $objPath);  
        msiSetACL("default ", "read", "utilisateurs", $objPath);  
    }  
    on ( $objPath like /imxgam/pixscan2/* ) {  
        msiSetACL("default ", "write", "admins", $objPath);  
        msiSetACL("default ", "write", "experts", $objPath);  
        msiSetACL("default ", "read", "utilisateurs", $objPath);  
        msiSendMail(leesa@renci.org, "ACLs set",  
                    "User $userNameClient set ACLs for $objPath.");  
    }  
}
```

Rules – some more examples

```
//  
// Migration from disk cache to MSS delayed 1 hour after the file has  
// been written in the cache. Retries up to 6 times until success with a  
// time interval doubling between each attempt:  
//  
acPostProcForPut {  
    ON($objPath like "/ams/data/*" || $objPath like "/ams/mc/nt/*" ||  
        $objPath like "/ams/mc/2011B/*") {  
  
        delay("<PLUSET>1h</PLUSET><EF>1h DOUBLE UNTIL  
            SUCCESS OR 6 TIMES</EF>") {  
            msiSysReplDataObj('stageResc2','null');  
        }  
    }  
}
```

Use Cases - BGI

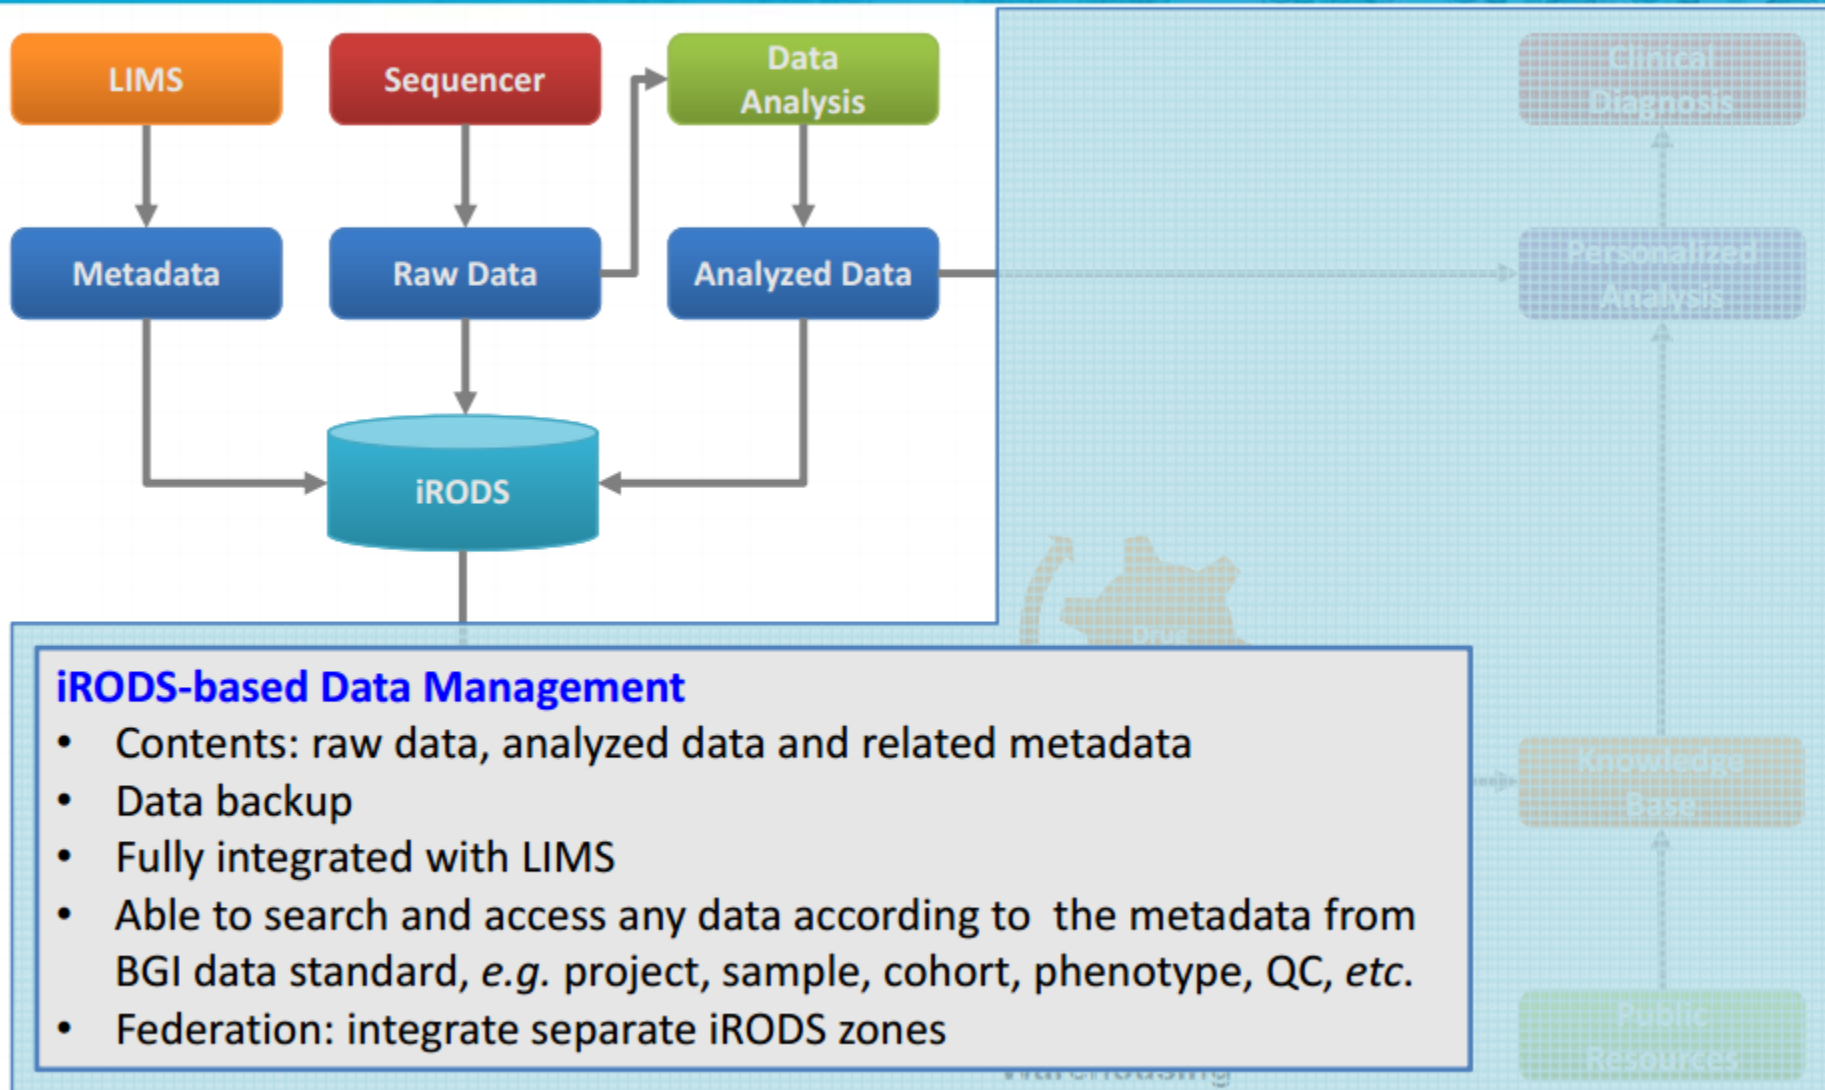
Slides courtesy of Xing Xin

- **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

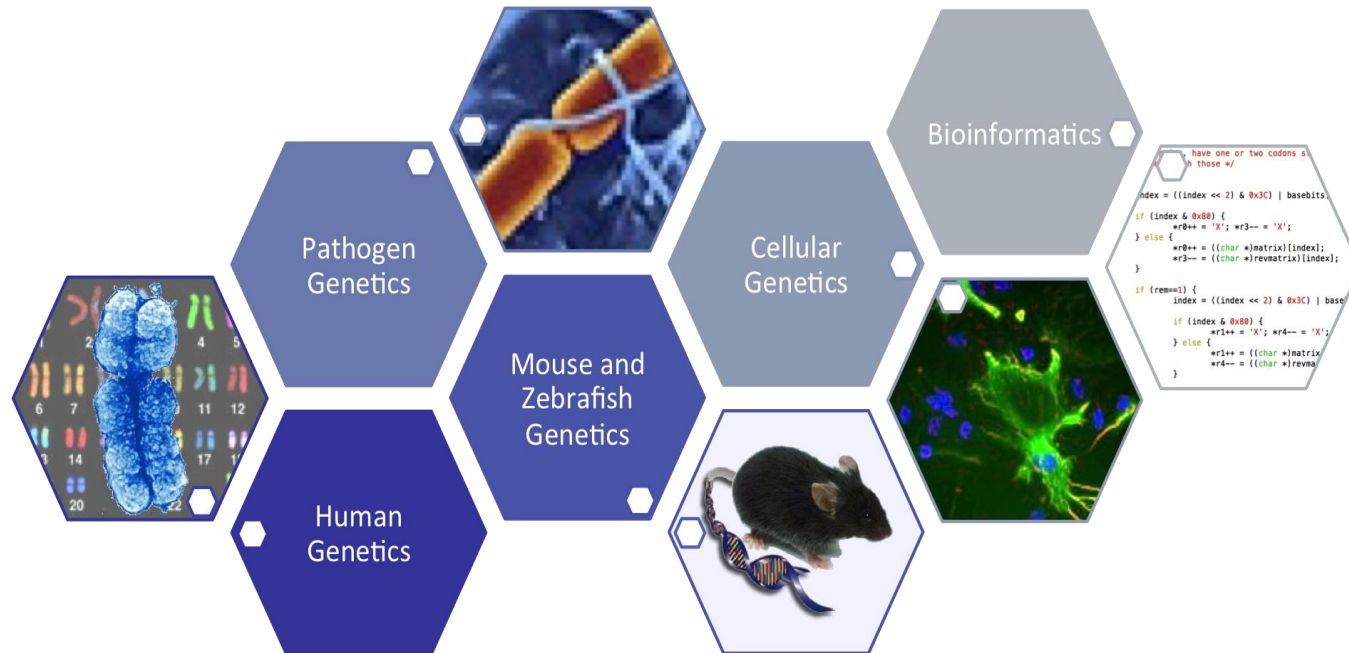





Use Cases - Sanger

Slides courtesy of Peter Clapham

Research Programmes



Sequence till 2011



Beginnings

Sanger was looking for a data management solution that:

- Solid at it's core. 2 am support calls are bad(tm)

- Vendor agnostic.

- Sensibly maintainable.

- Scalable, in terms of capacity and remain relatively performant.

- Extensible

iRODS layout


Data lands by preference onto iRES servers in the green datacenter room

Data is then replicated to Red room datacenter 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 (so far)



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

Sysadmin Perspective

Keep It Simple works.

Reflected by very limited downtime aside from upgrades

The core has remained nicely solid

Upgrades can be twitchy (2.4 → 3.3.1 over the past few year has not been without surprises...)

Some queries need some optimisation. Fortunately we have some very helpful DBA's

End User Perspective

Users are particularly happy with the meta data rich environment.

Now they can find their files and gain access in a reliable fashion.

So far so good. Satisfied users.

So happy they've requested iRODS areas for their specific use purposes

Federating Zones

Top level zone (sanger) acts as a Kerberos enabled portal
Users login here and receive a consistent view of the world.

Allows separation of impact between user groups


- Zone server load

- Different access control requirements.

- Clear separation as groups consider implimenting their own rules within their zone

Each zone has it's own group oversight which is responsible for managing
it's disk utilisation. Separation reduces horse trading and makes the
process much less involved...

Sanger Zone Arrangement



Pipeline Team Perspective

In general stuff is fine BUT some particular pain points have been found.

The good news is that some have been addressed, such as improving client icommand exit codes (svn 3.3 tree) and the ability to now create groups and populate them as an igroupadmin.

Other pain points, data entry into iRODS is not Atomic.

No re-use of connections

Local use of Json formatting, not natively supported by iRODS clients

But iRODS is Extensible

Java API

Python API

C API

Baton

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

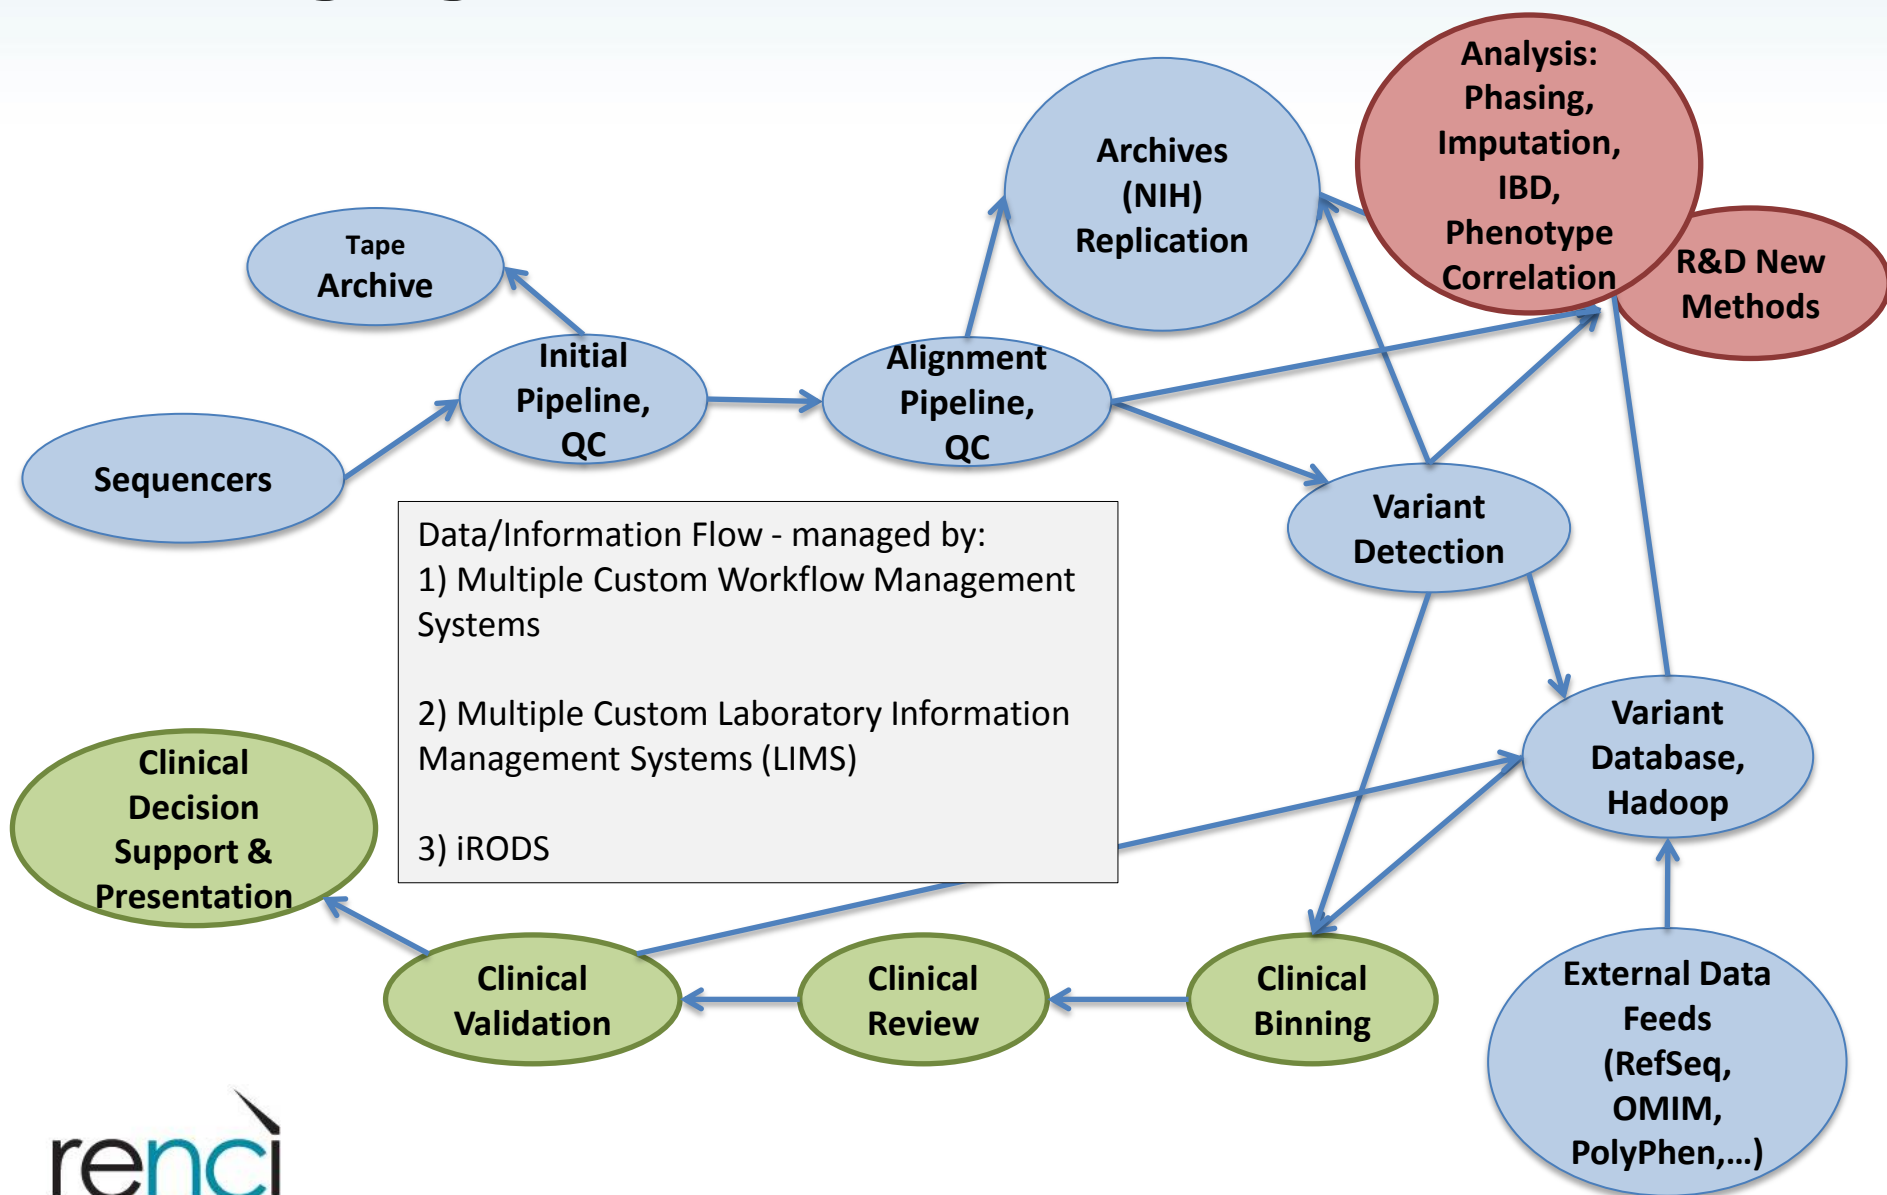
- Metadata queries

- Metadata addition

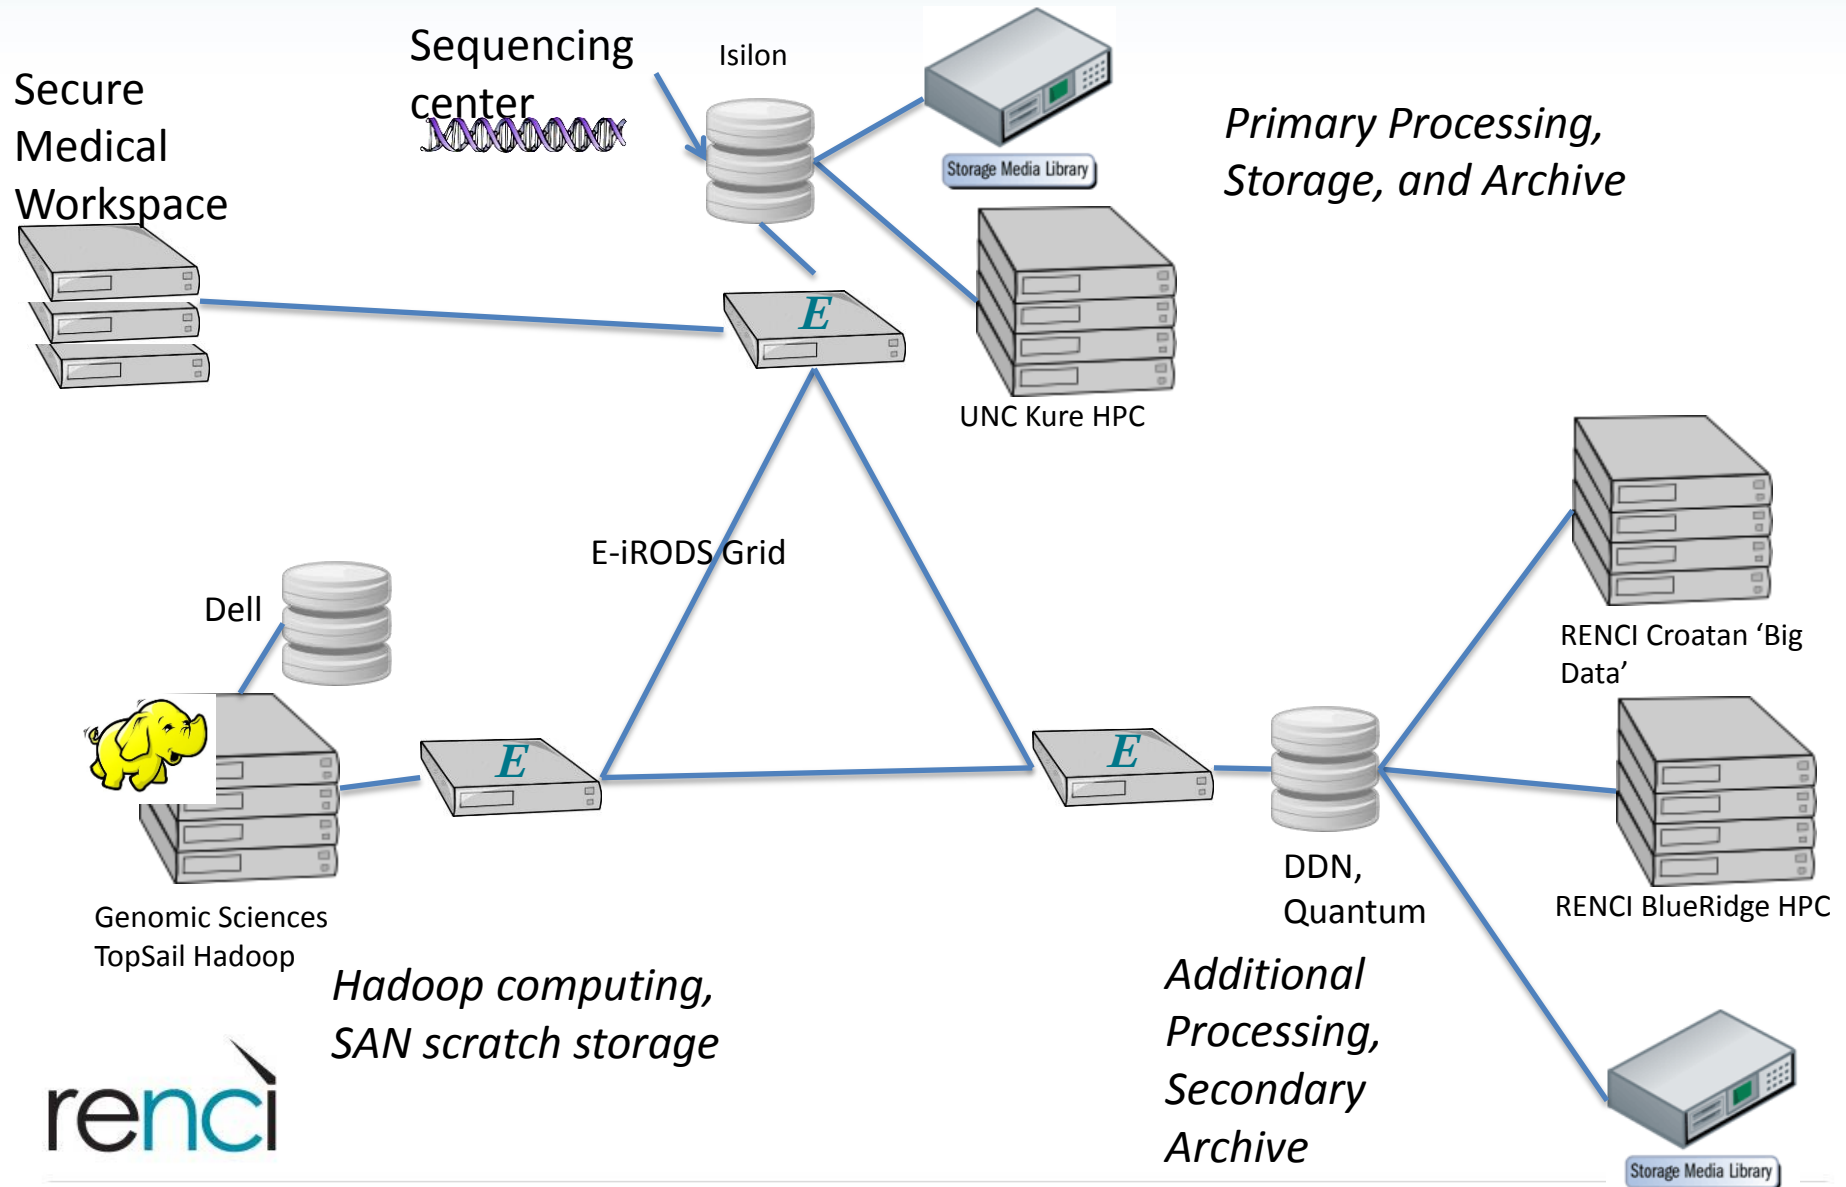
<https://github.com/wtsi-npg/baton.git>

Use Cases - UNC

Managing Research Data: Genomics



Genomics Primary Physical Infrastructure



Example: unified view of data

... / NCG_00110 / 121213_UNC11-SN...coverage_counts

iDROP web client

Refresh New Folder Info

Star File Download Add to Cart Rename Delete

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 :

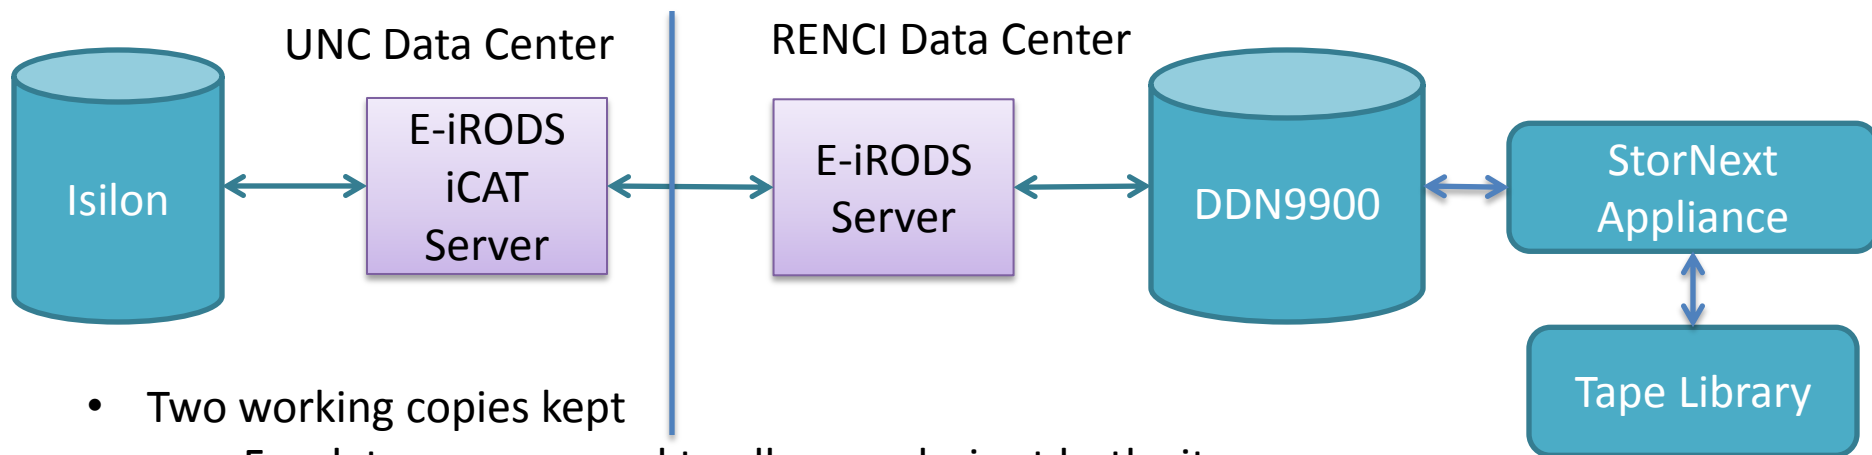
Spread across:

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)

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 run nightly does the copy
 - Performs copy, verifies copy successful, resets 'copy me' attribute
- Versioning to allow for re-runs of patient samples

NC GENES

https://ncgenes.unc.edu/MolecularAnalysis/AnalysisResults.php

How to Change Net...

UNC
SCHOOL OF MEDICINE

Good day, Phil Owen.
Your roles and studies:
NCGENES Administrators, HRC Study, Ophthalmology Study, NCGENES Study

UNC Health Care | UNC School of Medicine | UNC

genetics
NCGENES WorkFlow Manager

HomeWorkflowsAdministrationParticipantsAnalysisCLIAELSIHelpLog out

Variant analysis results for Molecular Analyst

For participant: NCG_00064
Diagnostic type: Myasthenia

Gene filter selection: No filter

Please select an analysis result type: Diagnostic

Load

View coverage data by geneView coverage data by exonDownload Bam files

Are you finished reviewing the diagnostic and incidental results? Select a value..

Class	Calculated Class	Verify	HGNC gene	Type	HGV5 genomic
LB	D	<input checked="" type="checkbox"/>	SCN4A	snp	NC_000017.10:g.6202892

Class	transcr	Loc type	strand	Intron exon dist	Variant Effect
D	NM_000334.4	exon	-	-137	missense

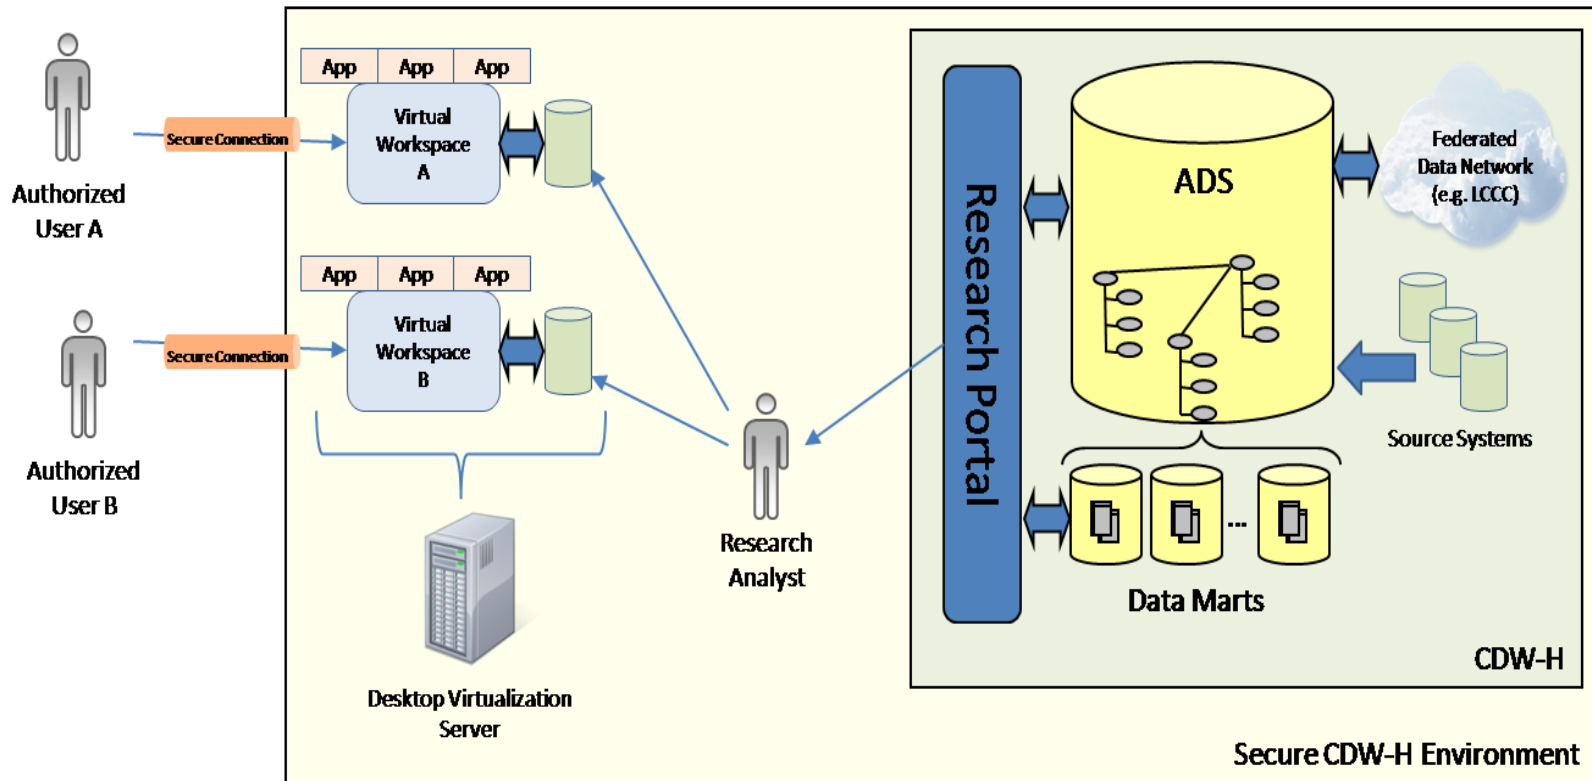
GFPT1delNC_000002.11:g.69565843_69565843delAintron

CHRNA1delNC_000002.11:g.175614909_175614910delAAintron

LB	D	<input checked="" type="checkbox"/>	GFPT1	snp	NC_000002.11:g.69565844A>T	intron	0	36	26.28	3.997	0	0	0	7	29	99	rs116730412	2	AR	0 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	SCN4A	ins	NC_000017.10:g.62045293_62045294insA	intron	0	35	9.92	-1.162		15	2.741	14	16	48.95	rs11418006	1	AR	0 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	CHRNA1	del	NC_000002.11:g.233408221_233408221delG	intron	0	159	7.12	-1.216		7	5.469	104	55	99		1	AR	0 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	CHAT	del	NC_000010.10:g.50856773_50856774delA	intron	0	10	20.96			0	0	2	9	24.07		1	AR	0 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	COLQ	ins	NC_000003.11:g.15539976_15539977insCACACA	intron	0	24	29.87	-0.099		0	0	14	6	99		1	AR	0 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	GFPT1	del	NC_000002.11:g.69565843_69565843delA	intron	0	36	17.39			12	0	8	28	86.65		2	AR	0 notes. [View/Add]
VUS	D	<input checked="" type="checkbox"/>	CHRNA1	del	NC_000002.11:g.175614909_175614910delAA	intron	0	63	11.71	0.897		17	2.14	13	33	93.27		1	CK	1 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	AGRN	snp	NC_000001.10:g.978857T>G	intron	0	151	0.6	2.658	0.01	4	82.866	94	57	99		2	AR	1 notes. [View/Add]
LB	D	<input checked="" type="checkbox"/>	CHRNA1	snp	NC_000002.11:g.233408190A>C	intron	0	176	2.75	-0.632	0	8	135.652	130	45	99	rs142201594	1	AR	0 notes. [View/Add]

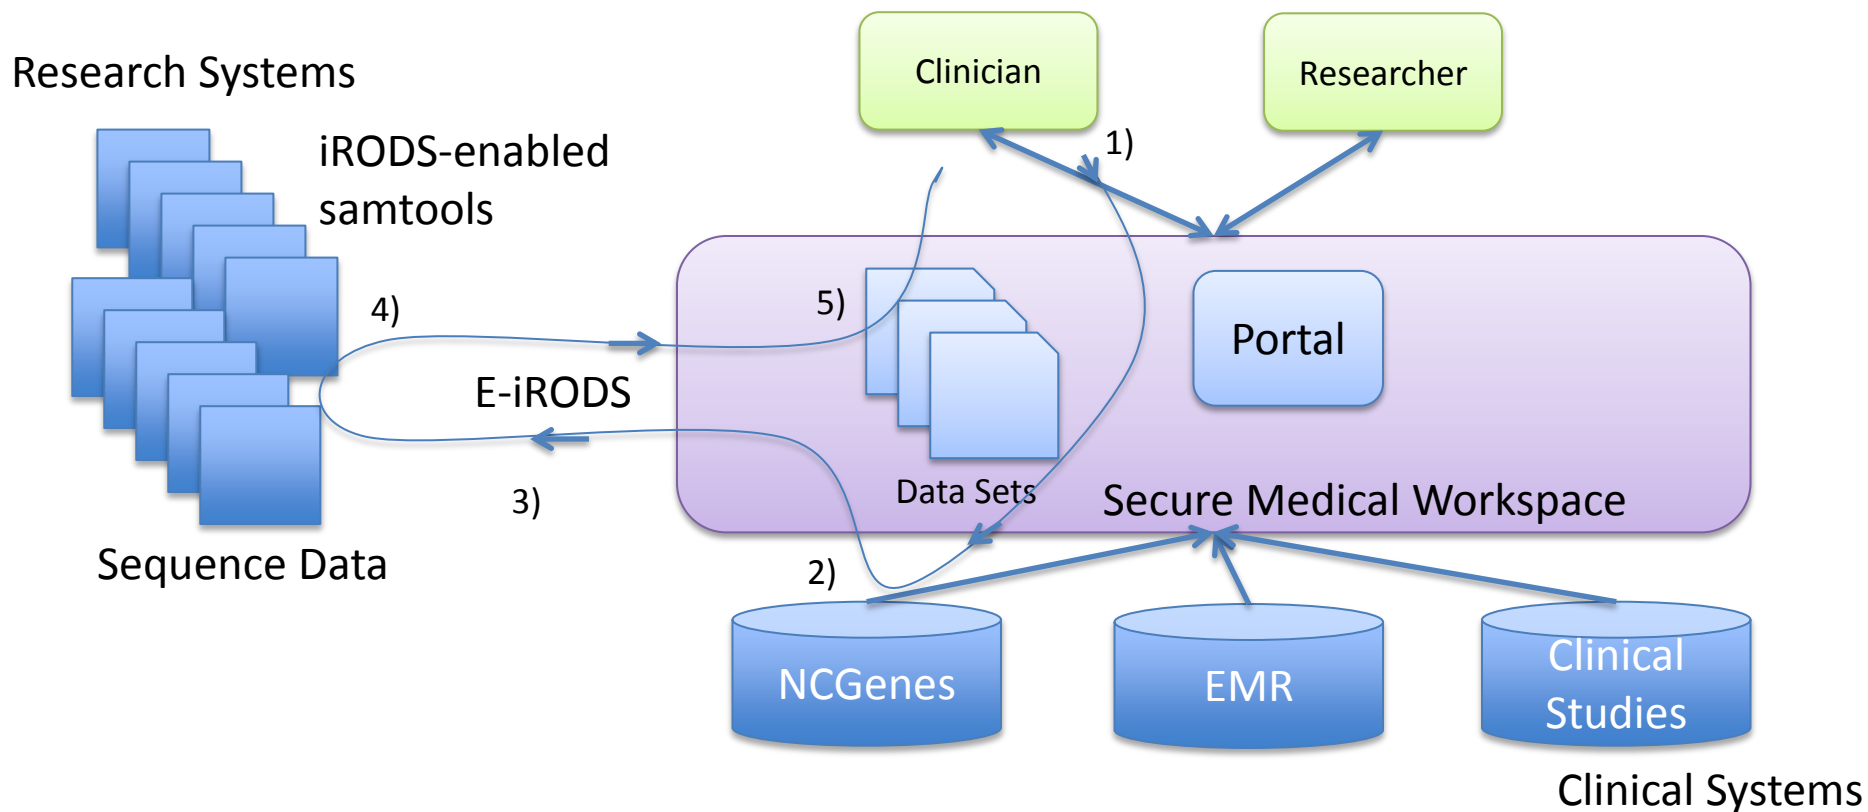
Go to page: 1 Show rows: 10 1-10 of 142

Secure Medical Workspace



- Combines Virtualization, Endpoint Data Leakage Protection (DLP), standard security such as use of VPNs, network sniffing, antivirus, group policies, ...

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 E-iRODS
- 4) Data sets packaged in zip file and retrieved
- 5) Data unzipped and displayed within secure workspace