iRODS Overview

Kory Draughn
Chief Technologist
iRODS Consortium

April 15-16, 2024
Library of Congress Designing Storage Architectures 2024
Washington, D.C.

iRODS CONSORTIUM

renci

RESEARCH \ ENGAGEMENT \ INNOVATION

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Our Membership

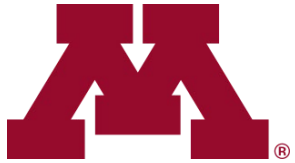

Universiteit Utrecht

Consortium
Member

Consortium
Member

National Institute of
Environmental
Health Sciences

What is iRODS?

Open Source

- C++ client-server architecture
- BSD-3 Licensed, install it today and try before you buy

Distributed

- Runs on a laptop, a cluster, on premises or geographically distributed

Data Centric & Metadata Driven

- Insulate both your users and your data from your infrastructure

iRODS as the Integration Layer

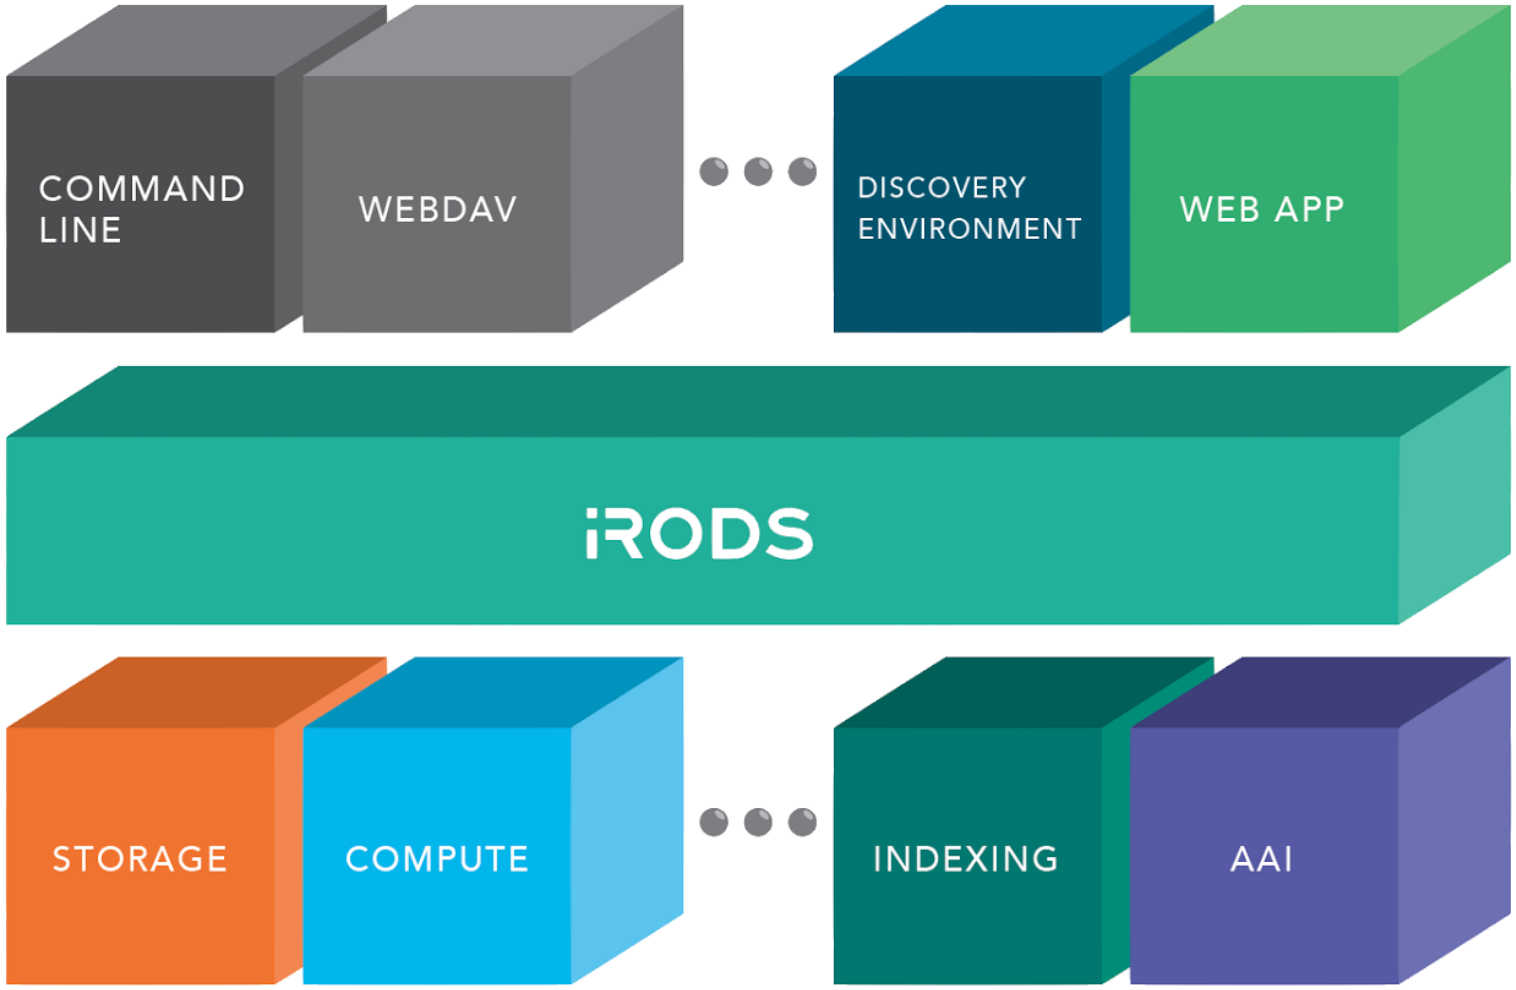

iRODS Clients

iRODS provides a layer of abstraction which integrates with your pre-existing infrastructure.

This flexibility allows your infrastructure to continue to change over time.

Existing Infrastructure

Why use iRODS?

People need a solution for:

- Managing large amounts of data across various storage technologies
- Controlling access to data
- Searching their data quickly and efficiently
- Automation

The larger the organization, the more they need software like iRODS.

The iRODS Data Management Model

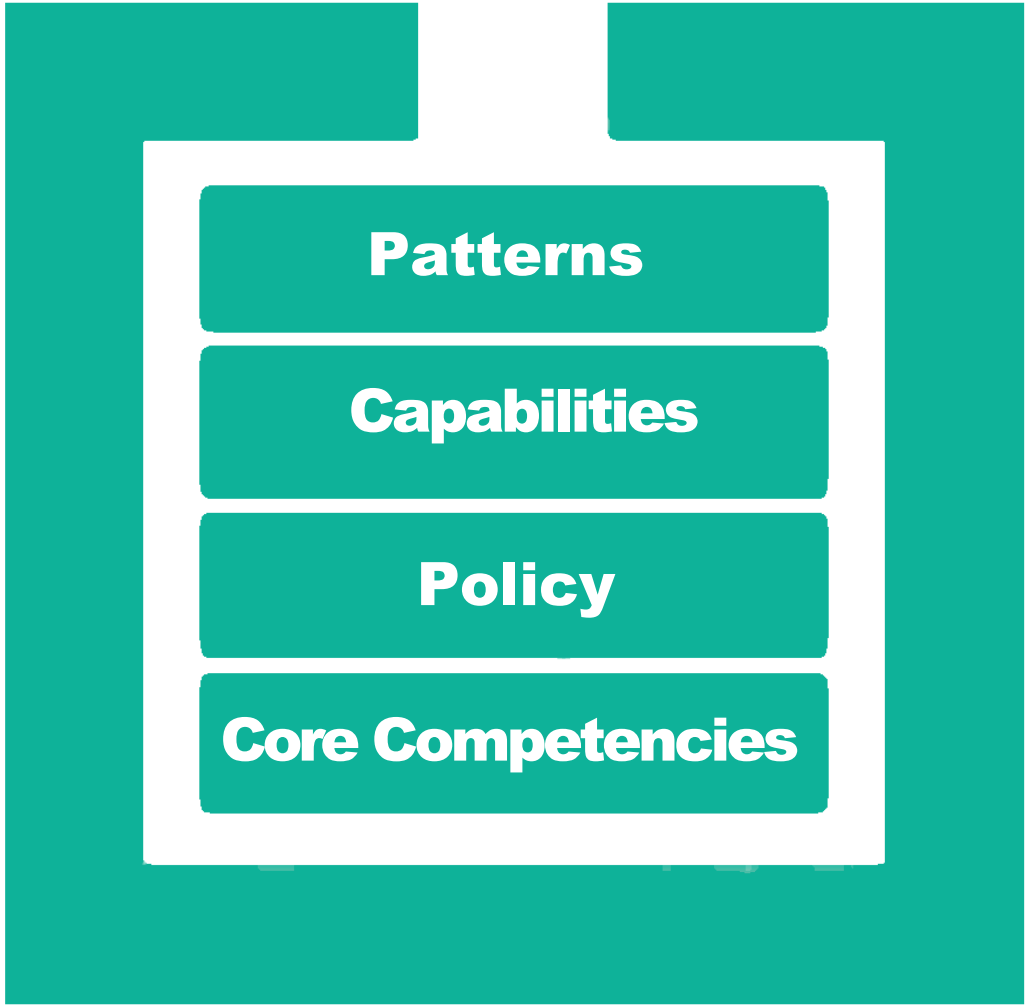

iRODS Core Competencies

The underlying technology categorized into four areas

Data Virtualization

Combine various distributed storage technologies into a Unified Namespace

- Existing file systems
- Cloud storage
- On premises object storage
- Archival storage systems

iRODS provides a logical view into the complex physical representation of your data, distributed geographically, and at scale.

Attach metadata to any first class entity within the iRODS Zone

- Data Objects
- Collections
- Users
- Storage Resources
- The Namespace

iRODS supports automated and user-provided metadata which makes your data and infrastructure more discoverable, operational, and valuable.

Policy Enforcement Points (PEPs) are triggered by every operation within the framework

- Authentication
- Storage Access
- Database Interaction
- Network Activity
- Extensible RPC API

The iRODS rule engine framework provides the ability to capture real world policy as computer actionable rules which may allow, deny, or add context to operations within the system.

Secure Collaboration

iRODS allows for collaboration across administrative boundaries after deployment

- No need for common infrastructure
- No need for shared funding
- Affords temporary collaborations

iRODS provides the ability to federate namespaces across organizations without pre-coordinated funding or effort.

iRODS Capabilities

- Packaged and supported solutions
- Require configuration not code
- Derived from the majority of use cases observed in the user community

Automated Ingest - Landing Zone

iRODS

Data may be automatically ingested from a number of sources which do not speak the iRODS protocol (microscopes, telescopes, sequencers, etc).

These sources could feed a single landing zone or an array of landing zones - this is a design decision for the iRODS administrator.

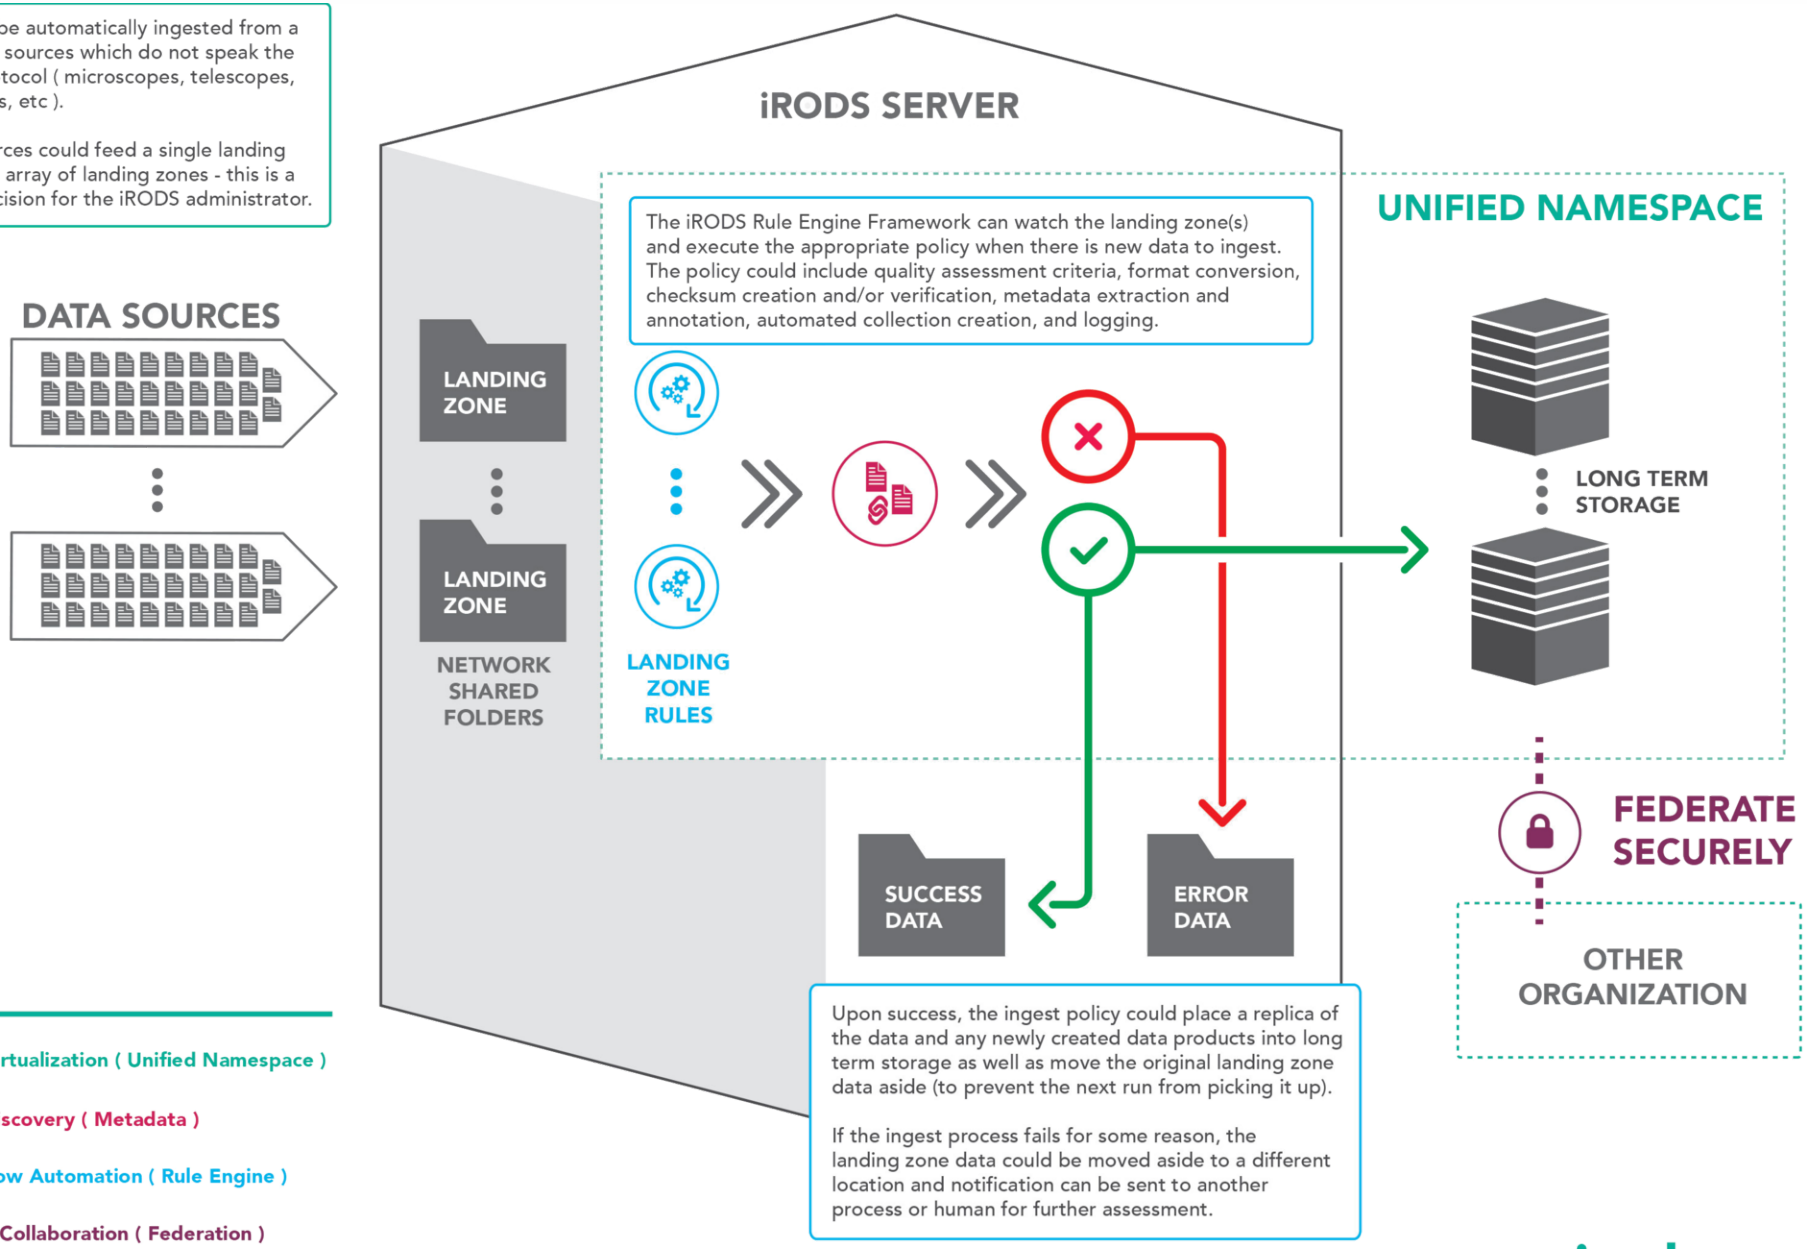

Data Virtualization (Unified Namespace)

Data Discovery (Metadata)

Workflow Automation (Rule Engine)

Secure Collaboration (Federation)

Automated Ingest - Filesystem Scanning

The iRODS filesystem scanner can watch for new or updated files and take appropriate action within iRODS.

This could include registering data in place, ingesting data into an iRODS Vault, or keeping two independent storage technologies in sync. Each of these iRODS actions could trigger policy enforcement points (PEPs) within the Zone to enforce the organizations' data management policy.

A scanning job is periodically added to the queue which generates jobs to register or ingest any new or changed data.

The queue may be scaled horizontally to keep up with arbitrary demand from the data sources (telescopes, simulations, parallel filesystems, streaming devices, etc.).

Metadata can be extracted and applied once data objects are registered in the iRODS catalog.

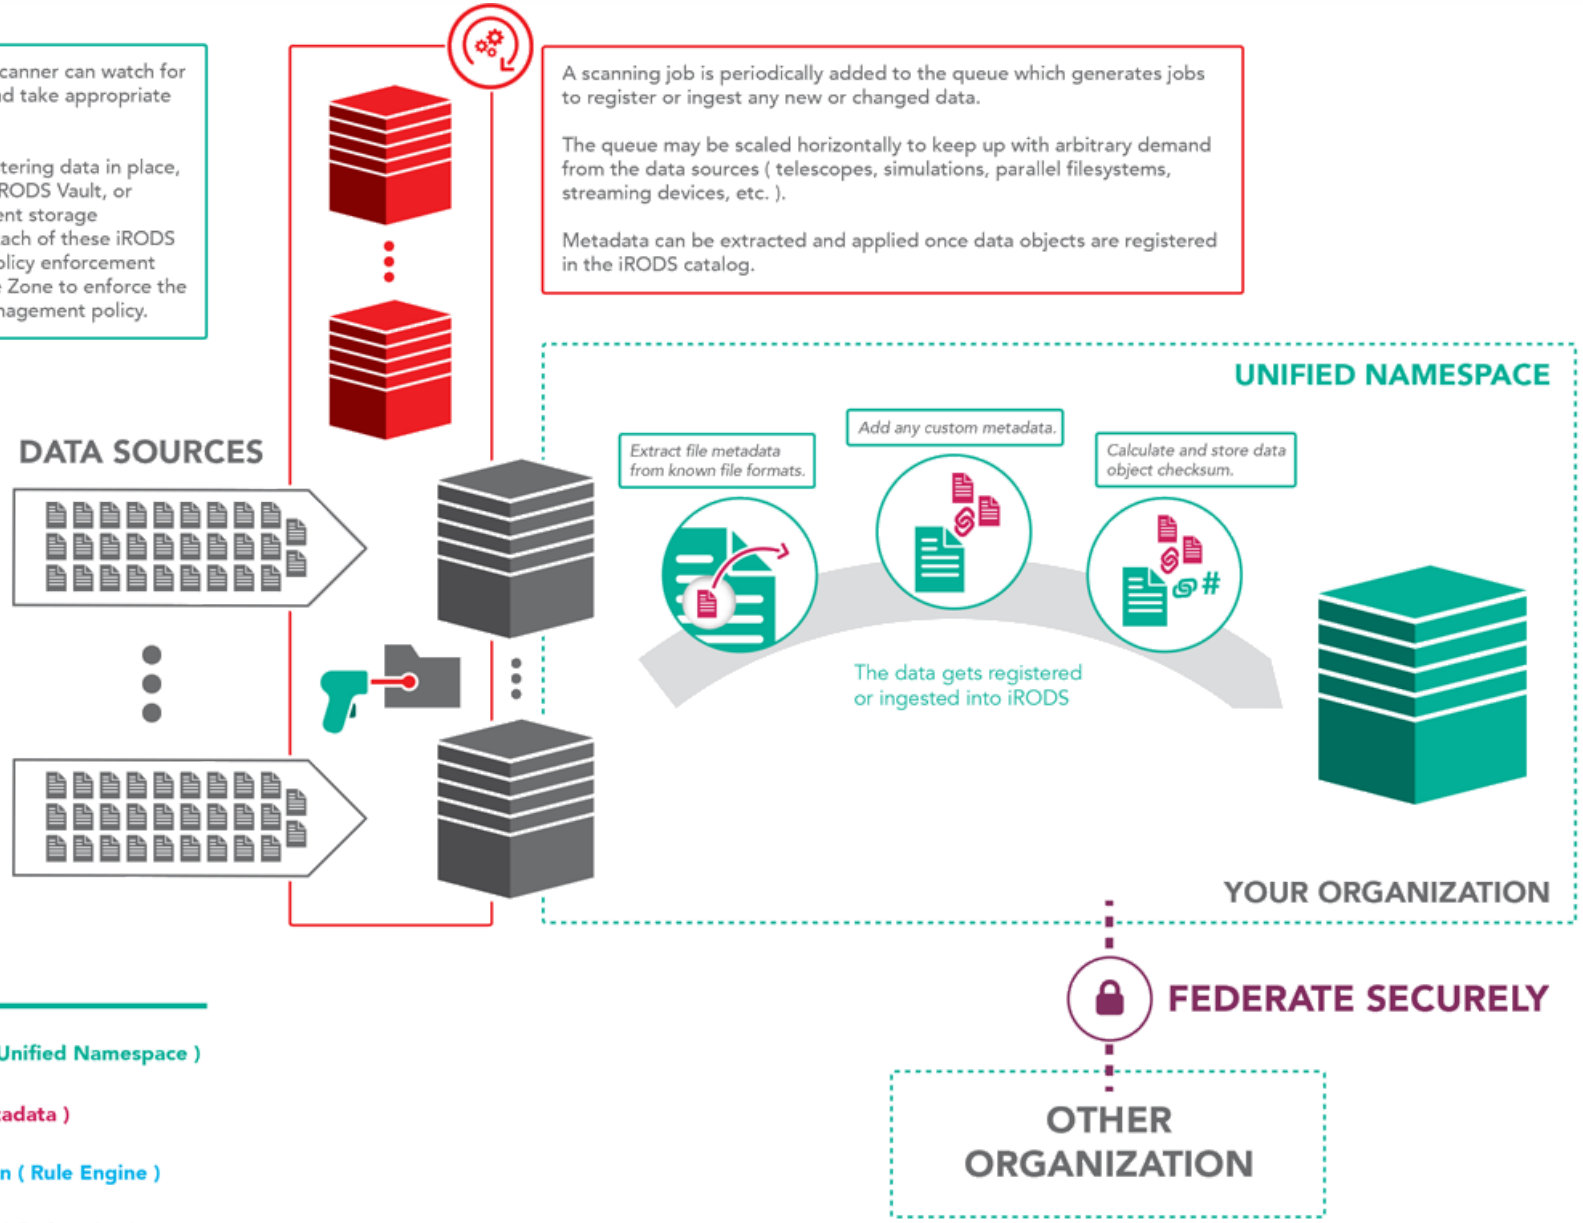

- Data Virtualization (Unified Namespace)
- Data Discovery (Metadata)
- Workflow Automation (Rule Engine)
- Secure Collaboration (Federation)

Storage Tiering

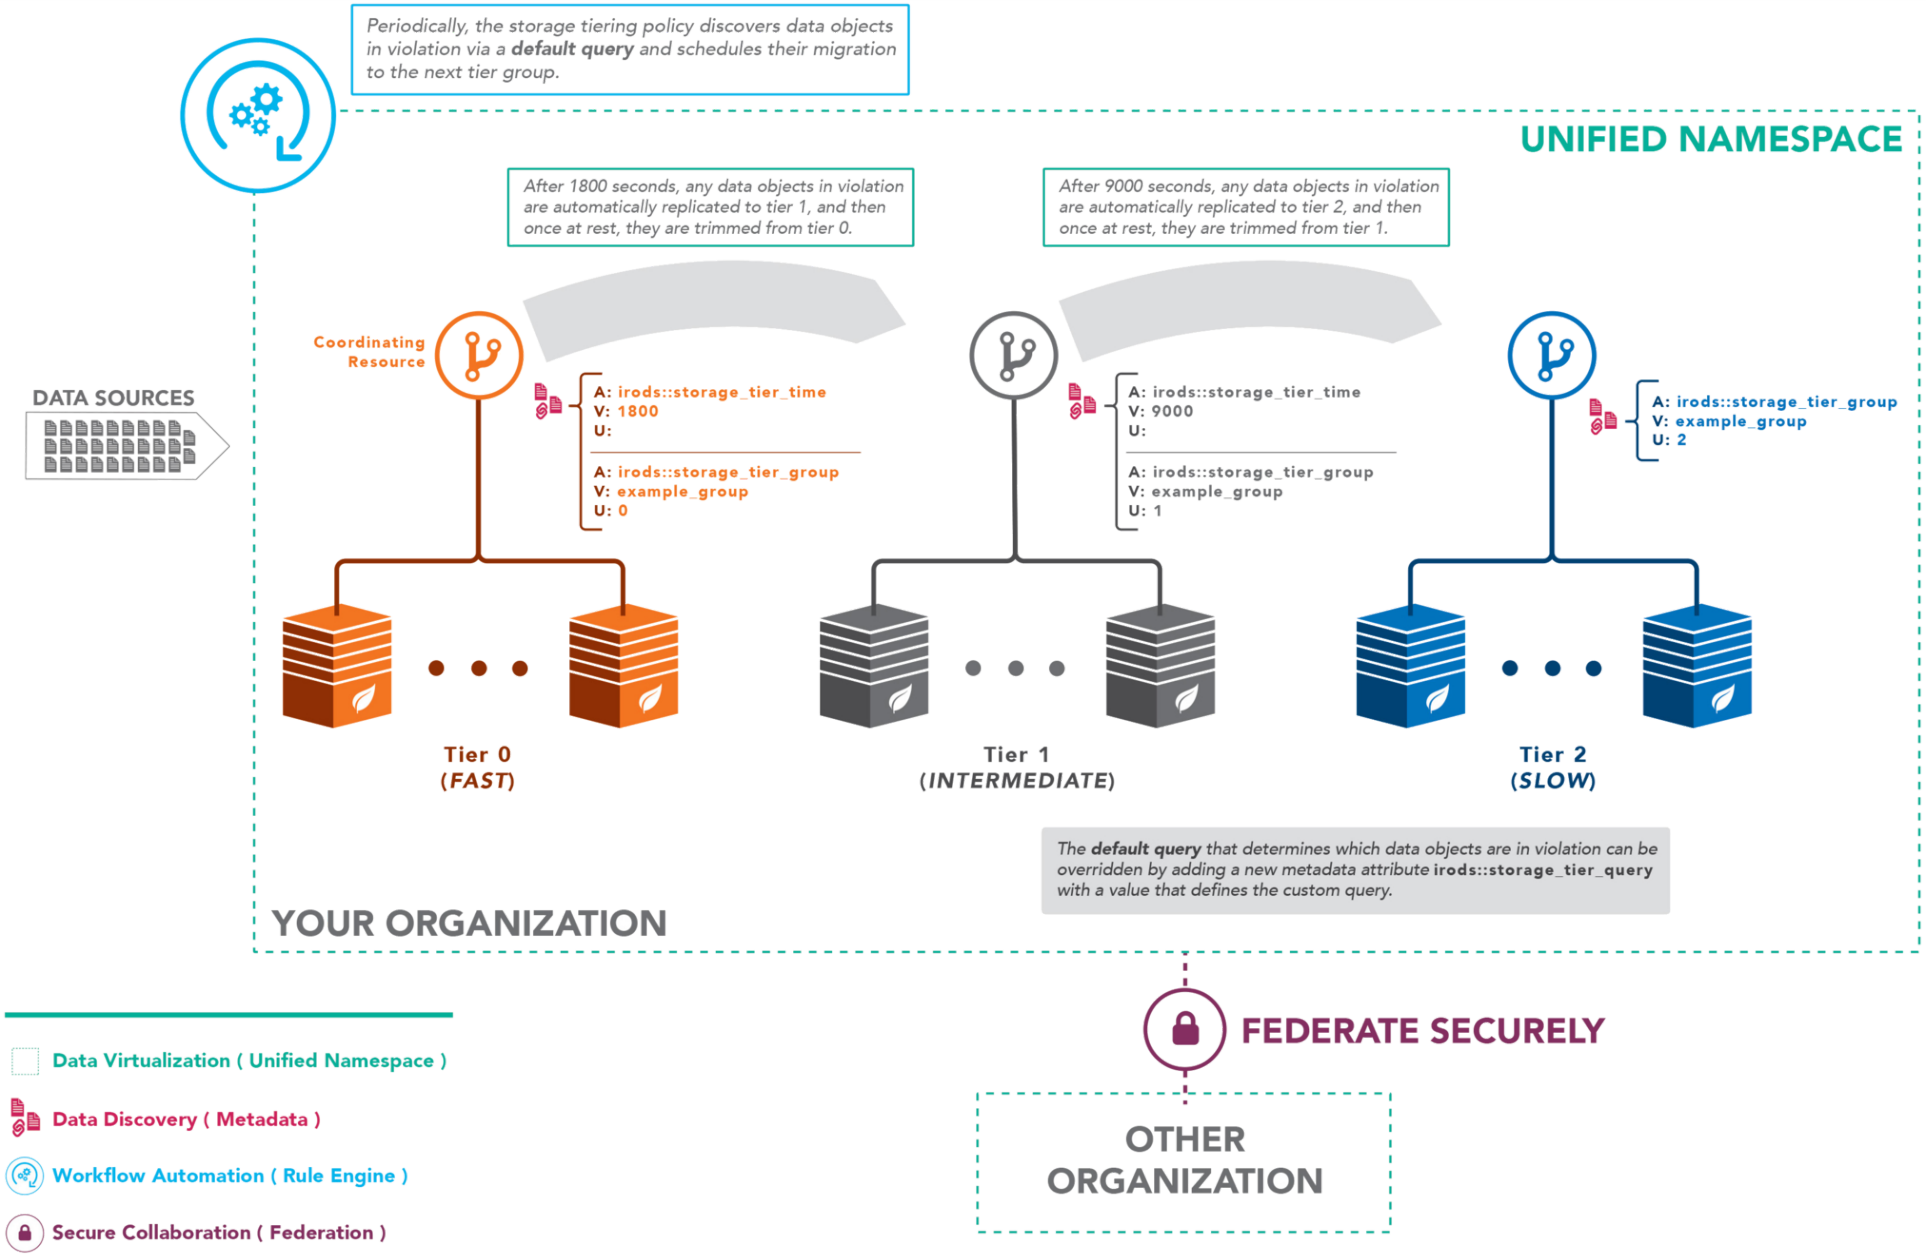

Indexing

The iRODS Indexing Capability provides a policy framework around both full text and metadata indexing for the purposes of enhanced data discovery.

Logical collections are annotated with metadata which indicates that any data objects or nested collections of data objects should be indexed given a particular indexing technology, index type, and index name.

From the configured metadata, the framework composes a rule name and then delegates to the policy implementation through the rule engine.

A new indexing technology can be supported via a rule base or policy engine which provides policy implementations of the form:

- irods_policy_indexing_object_index_<technology>
- irods_policy_indexing_object_purge_<technology>
- irods_policy_indexing_metadata_index_<technology>
- irods_policy_indexing_metadata_purge_<technology>

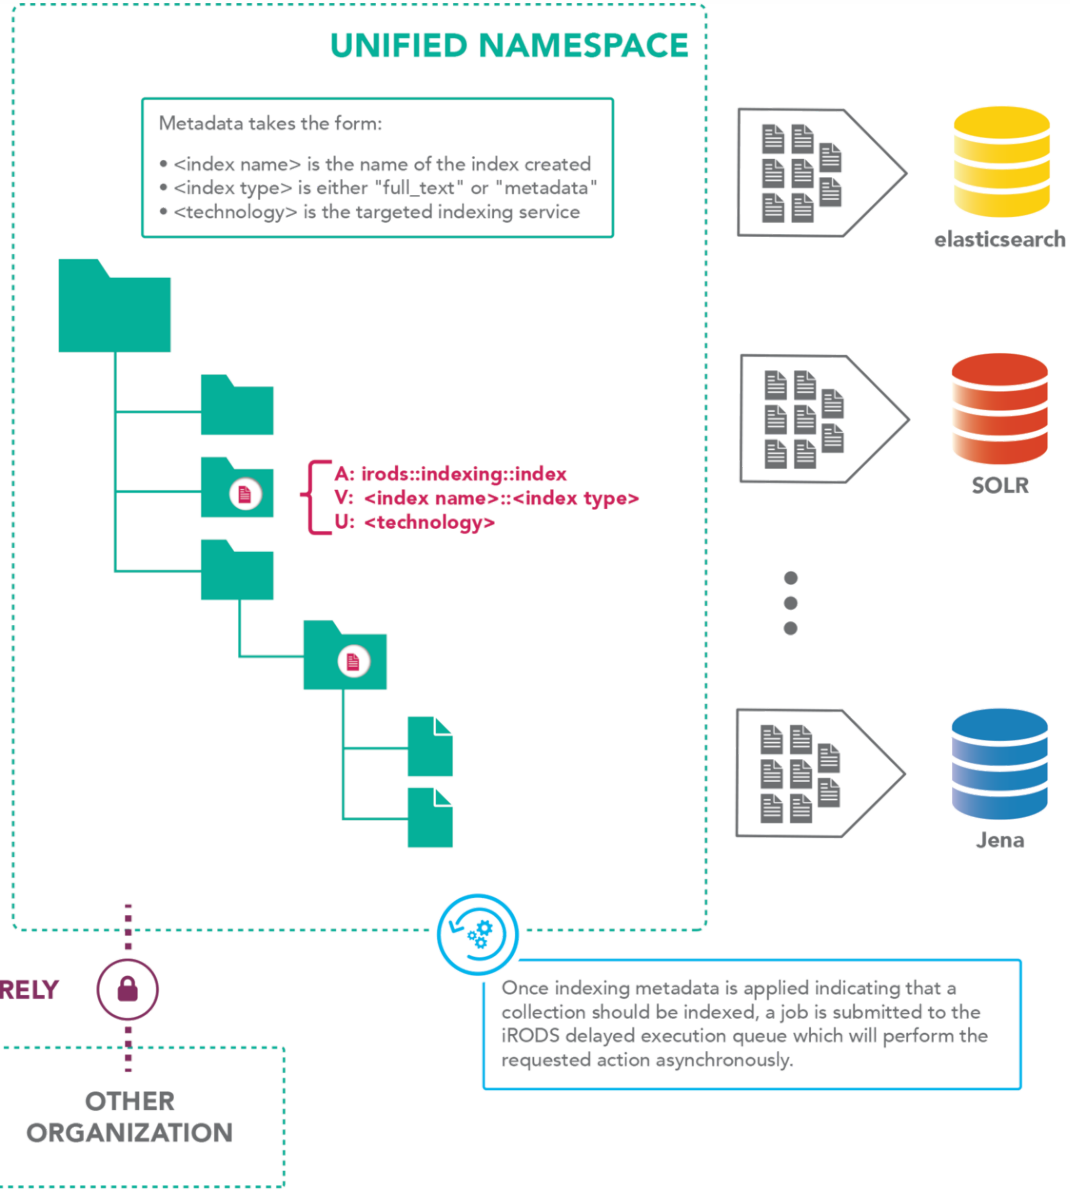

Data Virtualization (Unified Namespace)

Data Discovery (Metadata)

Workflow Automation (Rule Engine)

Secure Collaboration (Federation)

Publishing

The iRODS Publishing Capability provides a metadata driven policy framework for the implementation of data publication to external services.

The policy framework provides:

- Protection for published data, which prevents future modification
- Invocation of secondary policy for the generation and application of persistent identifiers
- Advertisement and possible movement of published data to external catalogs

From the configured metadata, the framework composes a rule name and then delegates to the policy implementation through the rule engine.

A new publishing service can be supported via a rule base or policy engine which provides policy implementations of the form:

- irods_policy_publishing_object_publish_<service>
- irods_policy_publishing_object_purge_<service>
- irods_policy_publishing_collection_publish_<service>
- irods_policy_publishing_collection_purge_<service>

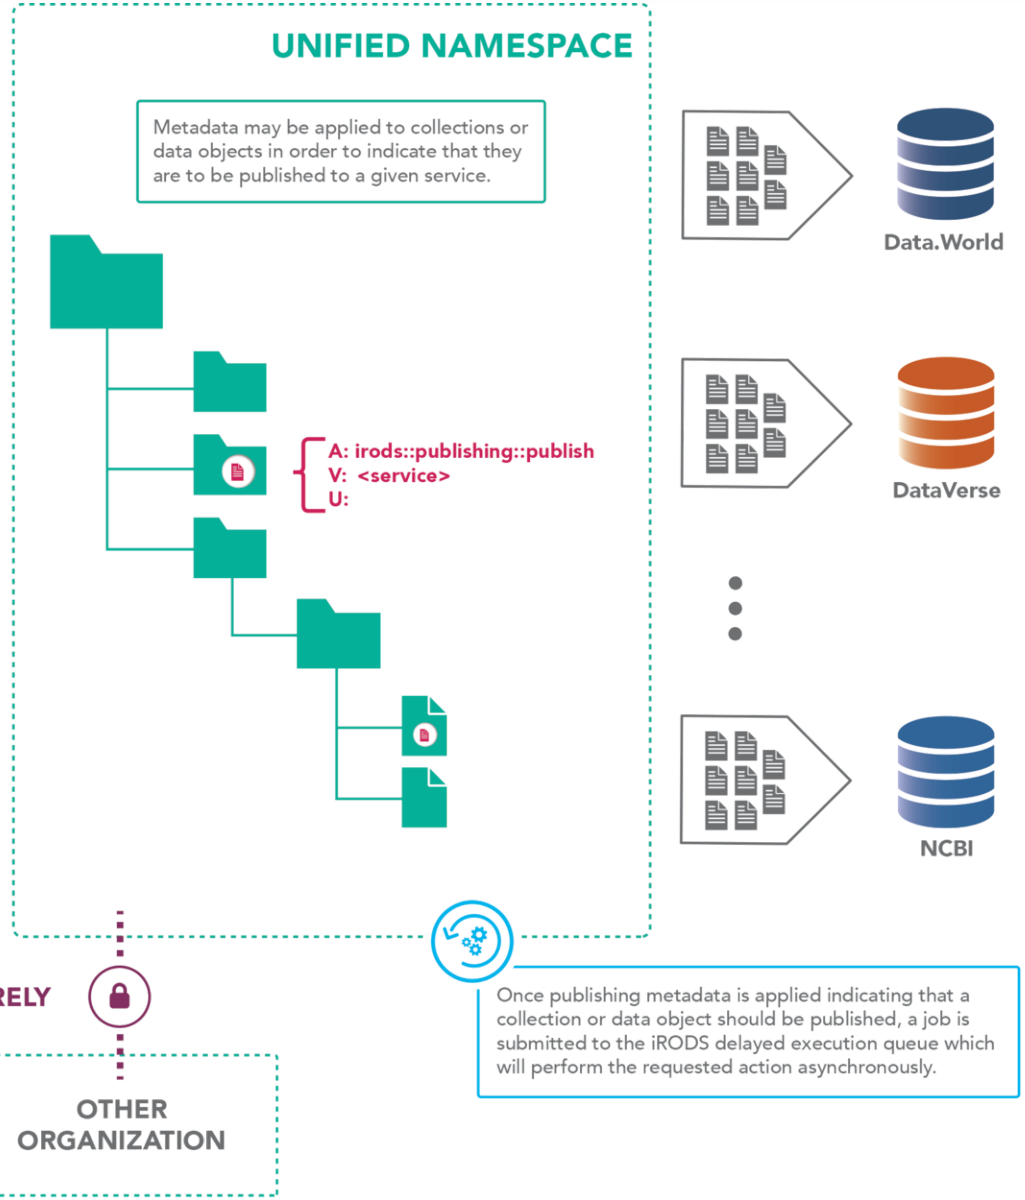

Data Virtualization (Unified Namespace)

Data Discovery (Metadata)

Workflow Automation (Rule Engine)

Secure Collaboration (Federation)

Filesystem Synchronization

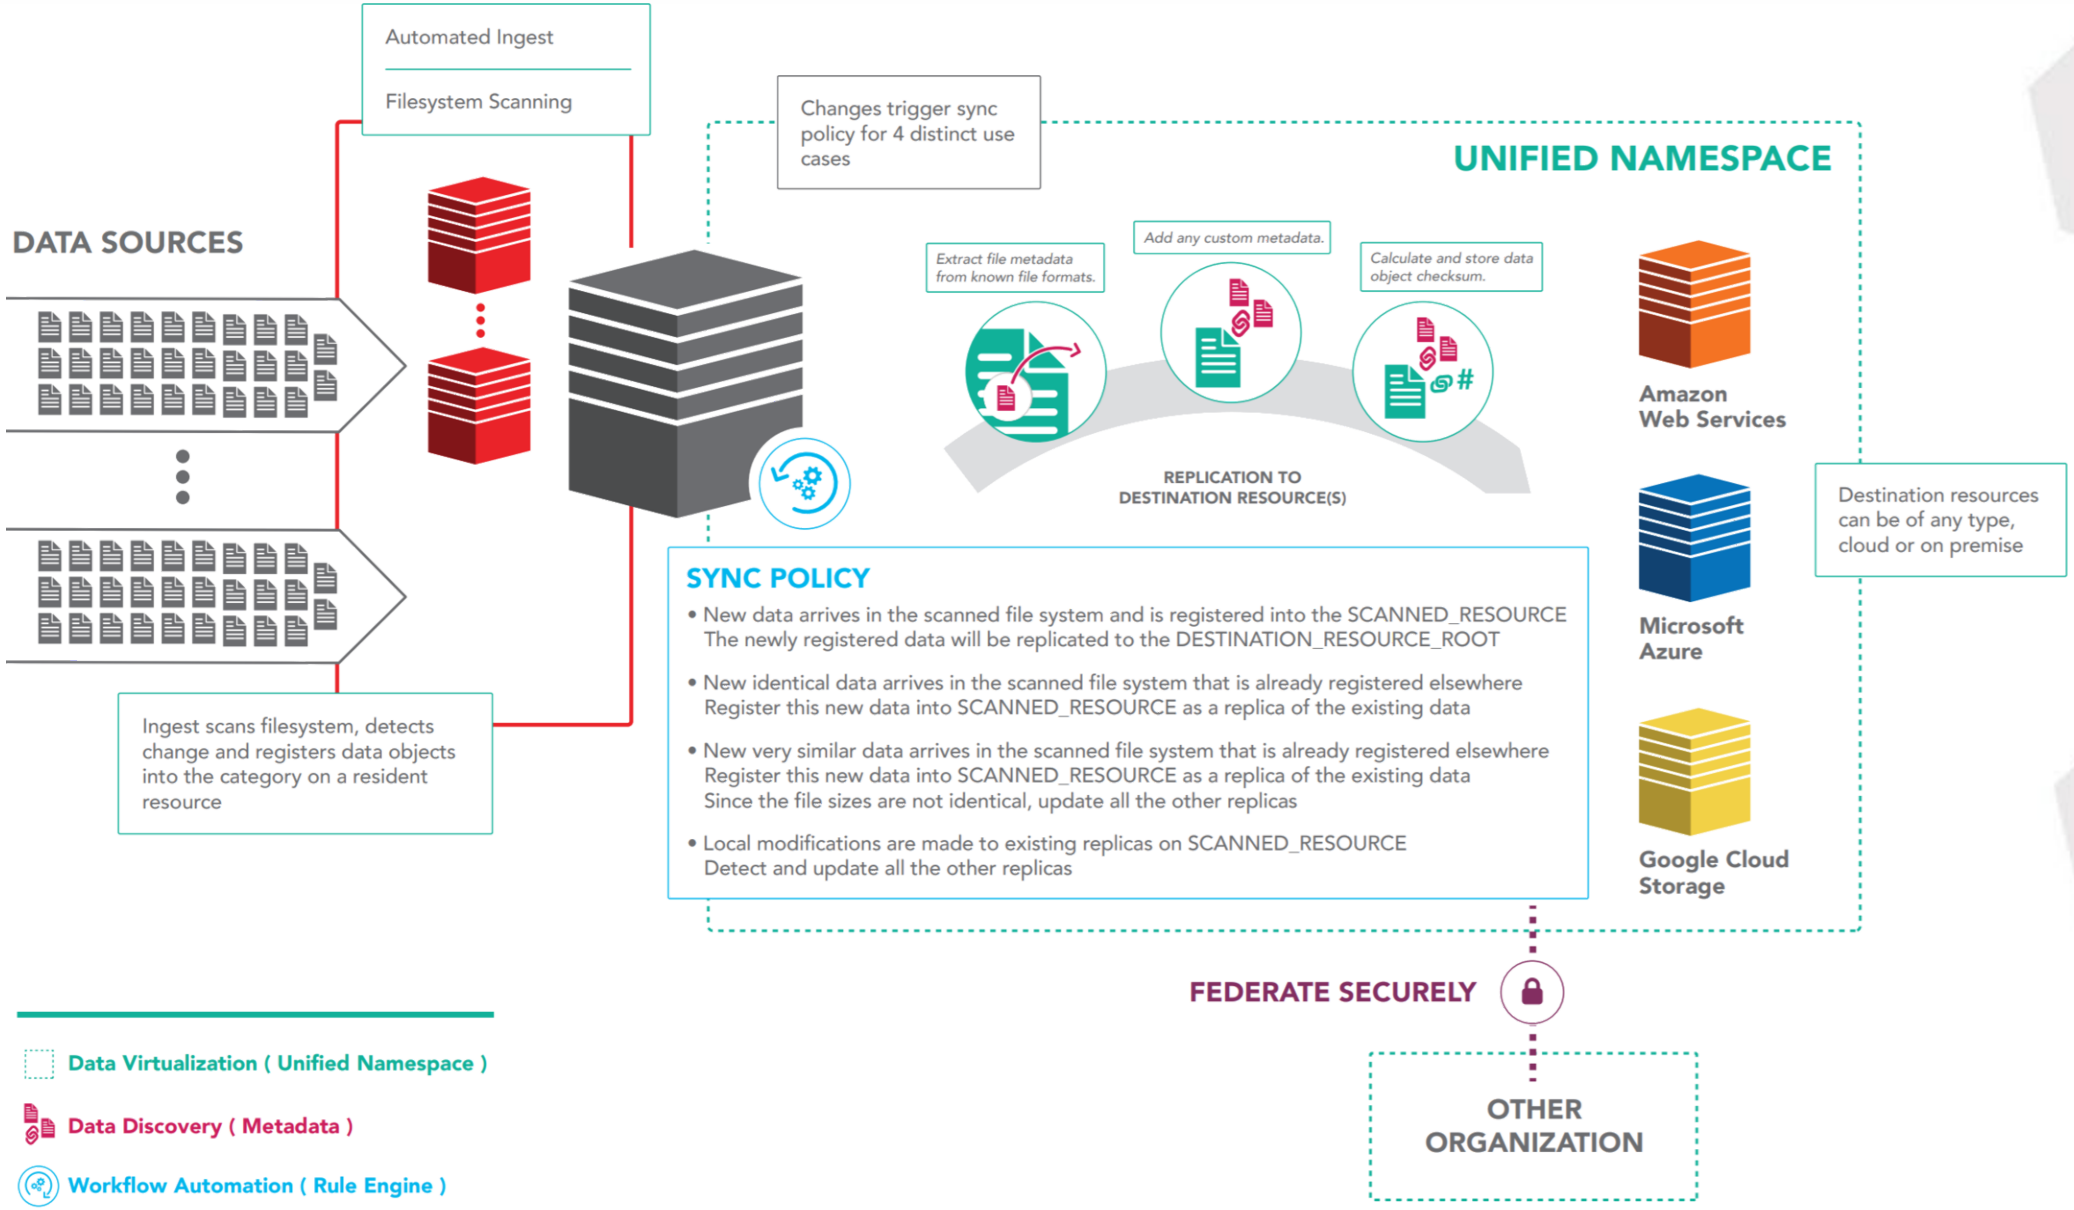

Data to Compute

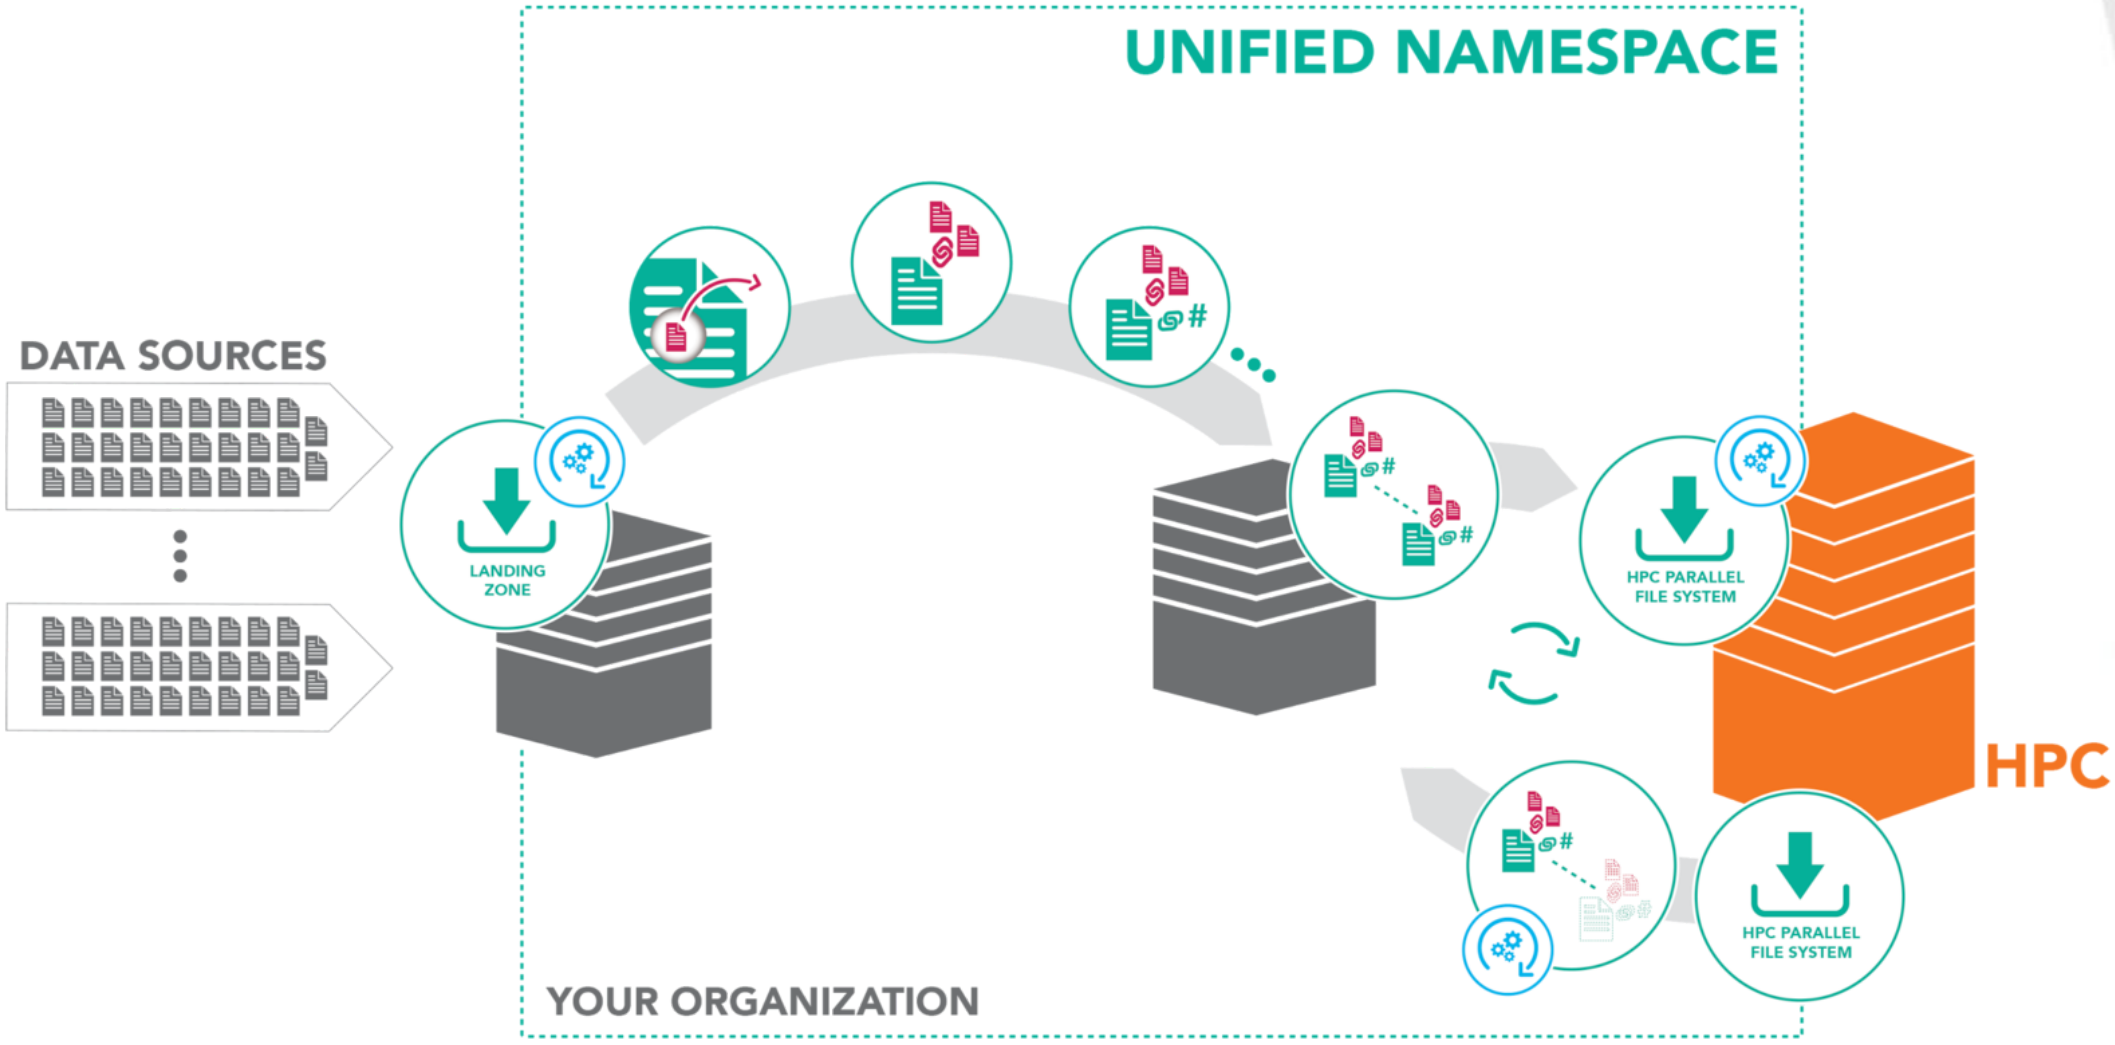

Compute to Data

iRODS

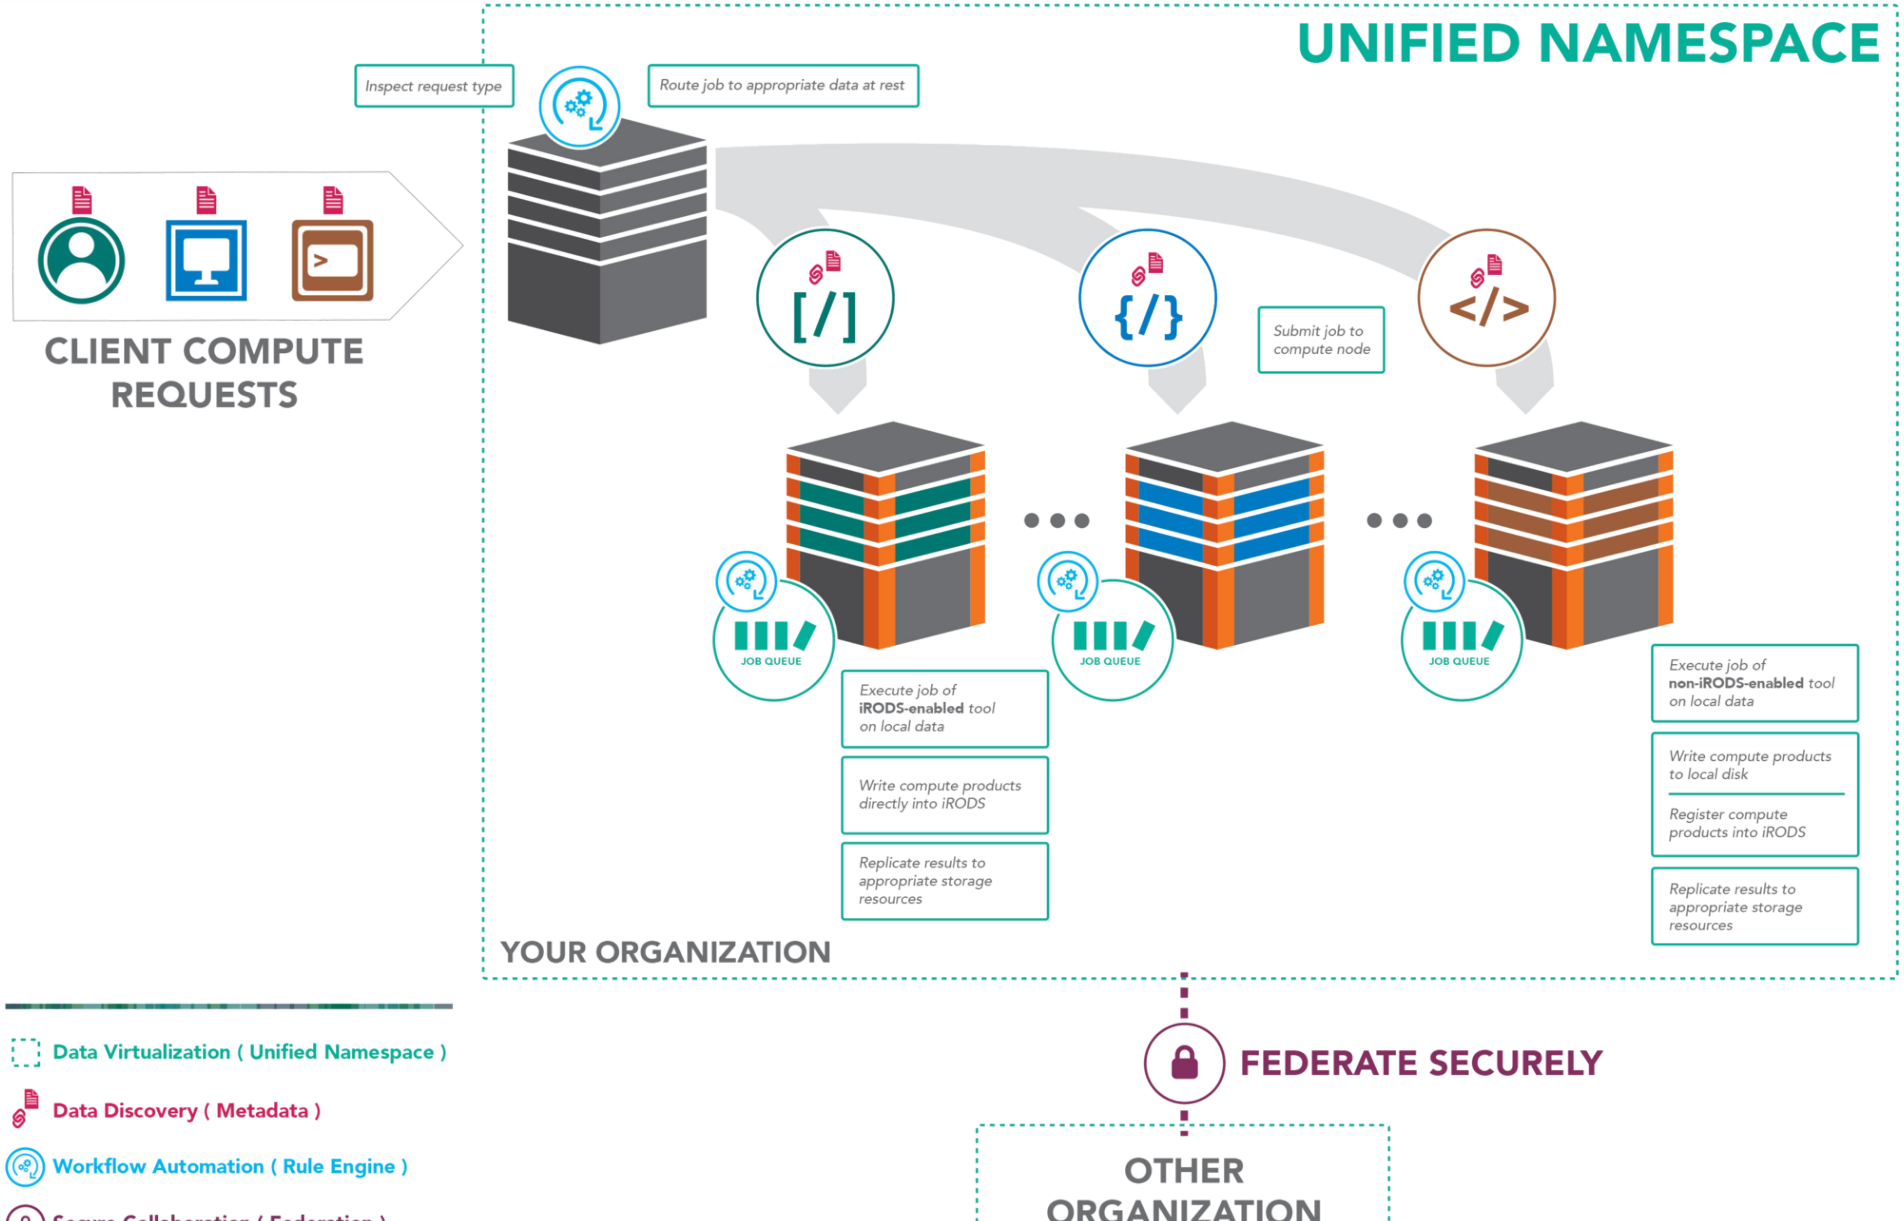

irods.org

Data Transfer Nodes

Moving large datasets across organizational boundaries remains a challenge due to the requirement of exposing high performance hardware to the public network. Data Transfer Nodes (DTNs) provide a secure location for ingress and egress of data while avoiding the performance impact of an organizational firewall.

In the following deployment pattern, iRODS satisfies the requirements of a Science DMZ while also providing automated data management.

> The Science DMZ is a portion of the network, built at or near the campus or laboratory's local network perimeter that is designed such that the equipment, configuration, and security policies are optimized for high-performance scientific applications rather than for general-purpose business systems or 'enterprise' computing.
>
> —ESnet

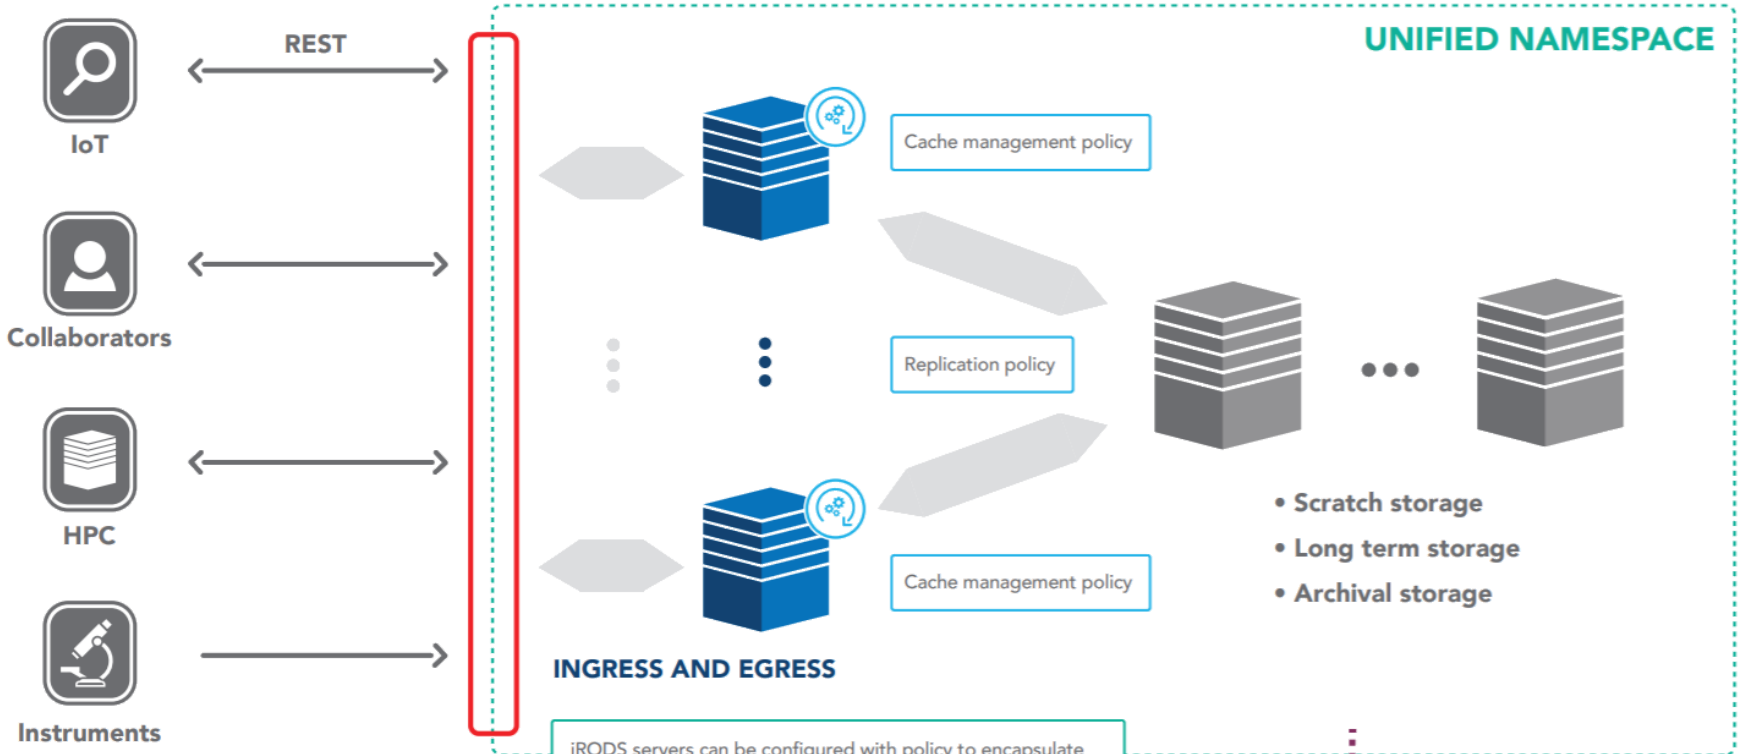

iRODS servers can be configured with policy to encapsulate read and write operations which trigger replication within the Zone. This allows the staging of data for ingress and egress to be separated from scratch storage, long term storage, or archival storage.

Because the iRODS servers act as Data Transfer Nodes (DTNs) the staged data should not be persistent. These servers implement active cache management via iRODS policy which limits the amount of required storage at the edge.

FEDERATE SECURELY

OTHER ORGANIZATION

- Data Virtualization (Unified Namespace)
- Data Discovery (Metadata)
- Workflow Automation (Rule Engine)
- Secure Collaboration (Federation)

irods.org

DATA LIFECYCLE

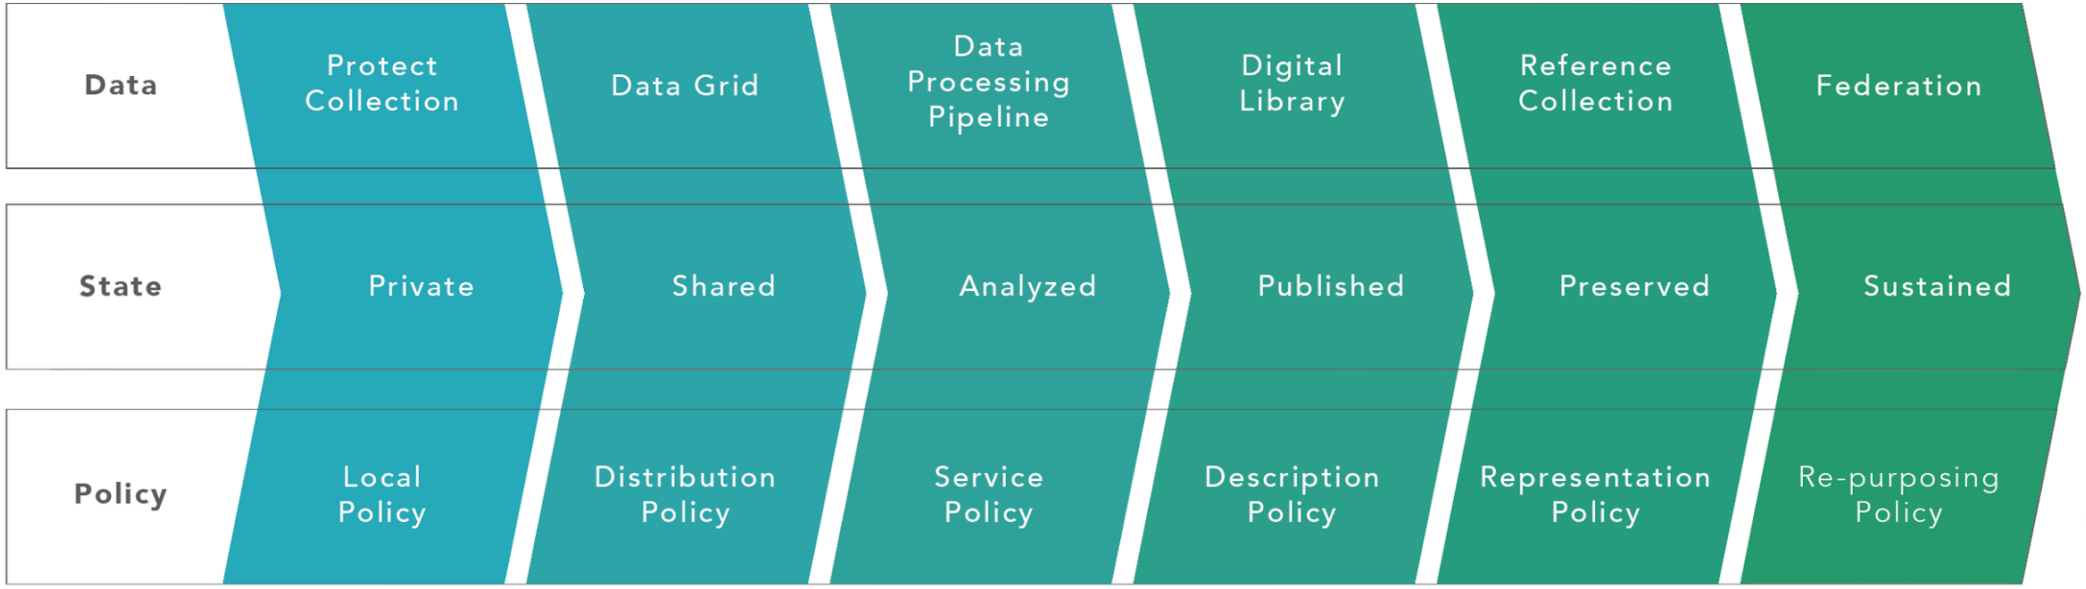

iRODS virtualizes the stages of the data lifecycle through policy evolution

As data matures and reaches a broader community, data management policy must also evolve to meet these additional requirements.

The Data Management Model

iRODS provides eight packaged capabilities, each of which can be selectively deployed and configured.

These capabilities represent the most common use cases as identified by community participation and reporting.

The flexibility provided by this model allows an organization to address its immediate use cases.

Additional capabilities may be deployed as any new requirements arise.

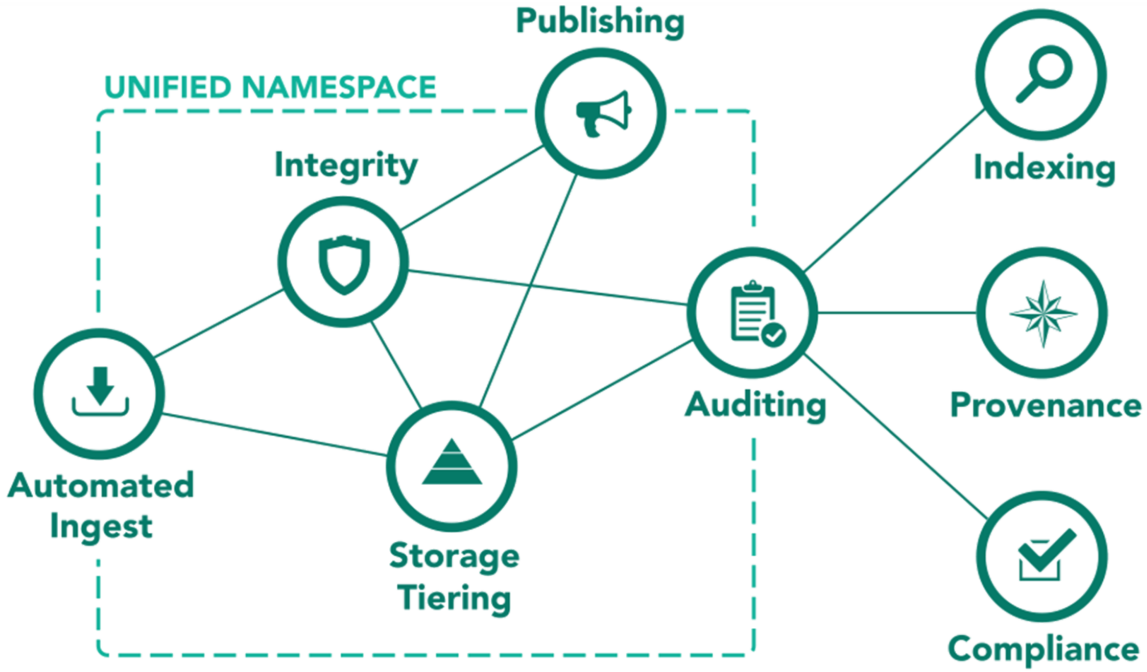

A pattern represents a combination of iRODS capabilities and data management policy consistent across multiple organizations.

Three common patterns of iRODS deployment have been observed within the community:

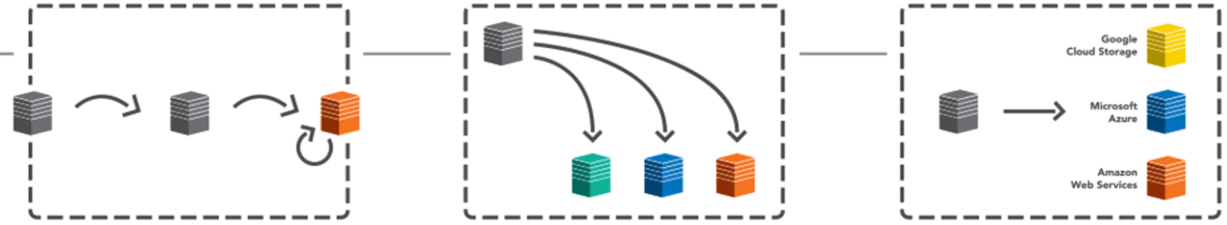

Data to Compute | Compute to Data | Synchronization

Protocol Plumbing - Presenting iRODS as other Protocols

Over the last few years, the ecosystem around the iRODS server has continued to expand.

Integration with other types of systems is a valuable way to increase accessibility without teaching existing tools about the iRODS protocol or introducing new tools to users.

With some plumbing, existing tools get the benefit of visibility into an iRODS deployment.

- WebDAV
- FUSE
- HTTP
- NFS
- SFTP
- K8s CSI
- S3

Thank you.

https://irods.org/ugm2024

May 28-31, 2024