



Policy Composition: Configuration, Not Code

Terrell Russell, Ph.D.
@terrellrussell
Chief Technologist, iRODS Consortium

January 25-29, 2021
CS3 2021
Virtual

iRODS

— CONSORTIUM —

renci

RESEARCH \ ENGAGEMENT \ INNOVATION



THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Our Membership

iRODS



Western Digital



CLOUDIAN



Universiteit Utrecht

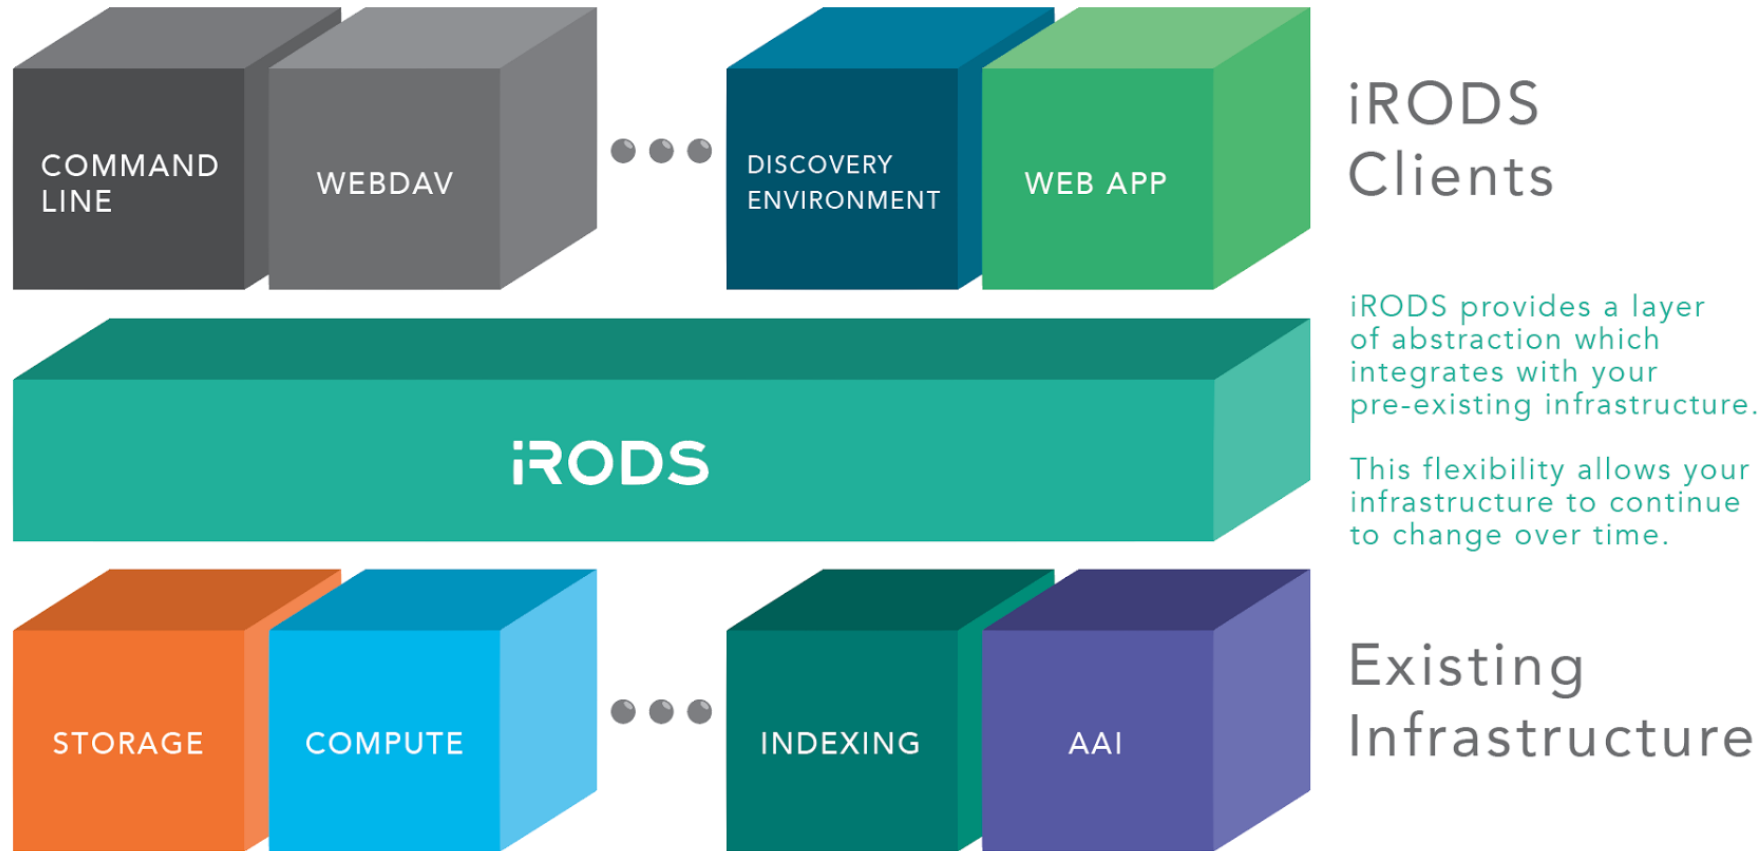


SOFTIRON

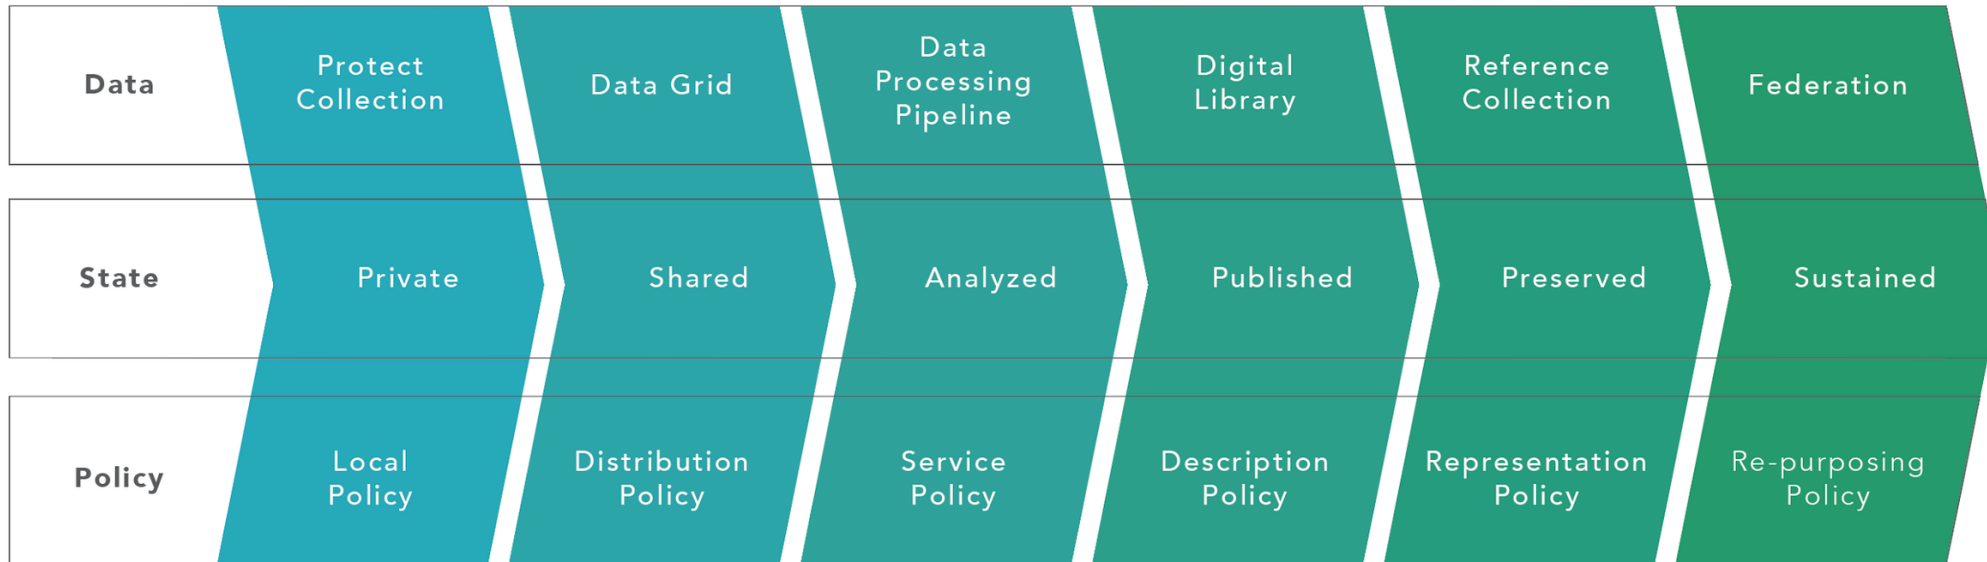


AGRICULTURE VICTORIA





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.

With a twenty-five year history, iRODS open source technology has been used to automate data management across many scientific and business domains. The scale and value of data across these domains drives the necessity for automation. This variety also demands a flexibility in data management policies over time. Organizations have satisfied their own needs by investing in the development of their own policy, but this has led to monolithic, specific policy sets tailored to a particular organization.

As a community, we have observed common themes and duplication of effort across these organizations and now worked to provide generalized implementations that can be deployed across multiple domains. This new approach allows multiple policies to be configured together, or composed, without the need for custom development. For example, our Storage Tiering capability is a composition of several basic policies: Replication, Verification, Retention, and the Violating Object Discovery.

A Definition of Data Management

"The development, execution and supervision of plans, **policies**, programs, and **practices** that control, protect, deliver, and enhance the value of data and information assets."

Organizations need a **future-proof** solution to managing data and its surrounding infrastructure



A Definition of Policy

A set of ideas or a **plan** of what to do in **particular situations** that has been agreed to officially by a group of people...

So how does iRODS do this?



The reflection of real world data management decisions in computer actionable code.


(a plan of what to do in particular situations)





- Data Movement
- Data Verification
- Data Retention
- Data Replication
- Data Placement
- Checksum Validation
- Metadata Extraction
- Metadata Application
- Metadata Conformance
- Replica Verification
- Vault to Catalog Verification
- Catalog to Vault Verification
- ...





- How can we help new users get started?
- How can we make policy reusable?
- How can we simplify policy development?
- How can we provide a cook book of deployments?
- How do we get from Policy to Capabilities?



Consider Policy as building blocks towards Capabilities

Follow proven software engineering principles:

- Favor composition over monolithic implementations

Rules and Dynamic Policy Enforcement Points
can be overloaded and fall through

Implement or configure several rule bases or rule
engine plugins to achieve complex use cases





Consider Storage Tiering as a collection of policies:

- Data Access Time
- Identifying Violating Objects
- Data Replication
- Data Verification
- Data Retention





Continue to separate the concerns:

- When : Which policy enforcement points
- What : The policy to be invoked
- Why : What are the conditions necessary for invocation
- How : Synchronous or Asynchronous

Write simple policy implementations

- Not tied to a Policy Enforcement Point
- Do one thing well
- How it is invoked is of no concern

Each policy may now be reused in a generic fashion,
favoring configuration over code.



Examples





```
1  {
2  "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3  "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4  "plugin_specific_configuration": {
5      "policies_to_invoke" : [
6          {
7              "active_policy_clauses" : ["post"],
8              "events" : ["put", "get", "create", "read", "write", "rename",
9                  "register", "unregister", "replication", "checksum",
10                 "copy", "seek", "truncate"],
11              "policy" : "irods_policy_access_time",
12              "configuration" : {
13              }
14          }
15      ]
16  }
17 }
```



```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "active_policy_clauses" : ["post"],
8         "events" : ["create", "write", "registration"],
9         "policy" : "irods_policy_data_replication",
10        "configuration" : {
11          "source_to_destination_map" : {
12            "source_resource_0" : ["destination_resource_0a", "destination_resource_0b"],
13            "source_resource_1" : ["destination_resource_1a"],
14          }
15        },
16        {
17          "active_policy_clauses" : ["post"],
18          "events" : ["create", "write", "registration"],
19          "policy" : "irods_policy_data_replication",
20          "configuration" : {
21            "destination_resource" : "destination_resource_3"
22          }
23        },
24      ]
25    }
26 }
```



```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "active_policy_clauses" : ["post"],
8         "events" : ["create", "write", "registration"],
9         "policy" : "irods_policy_enqueue_rule",
10        "delay_conditions" : "<ET>PLUSET 1</ET>",
11        "payload" : {
12          "policy" : "irods_policy_execute_rule",
13          "payload" : {
14            "policy_to_invoke" : "irods_policy_data_replication",
15            "configuration" : {
16              "source_to_destination_map" : {
17                "source_resource_0" : ["destination_resource_0a", "destination_resource_0b"],
18                "source_resource_1" : ["destination_resource_1a"],
19              }
20            }
21          }
22        }
23      }
24    ]
25  }
26 }
```




```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "active_policy_clauses" : ["post"],
8         "events" : ["replication"],
9         "policy" : "irods_policy_data_retention",
10        "configuration" : {
11          "mode" : "trim_single_replica",
12          "source_resource_list" : ["source_resource_1", "source_resource_2"]
13        }
14      }
15    ]
16  }
17 }
```

```
1      {
2          "policy" : "irods_policy_enqueue_rule",
3          "delay_conditions" : "<EF>REPEAT FOR EVER</EF>",
4          "payload" : {
5              "policy" : "irods_policy_execute_rule",
6              "payload" : {
7                  "policy_to_invoke" : "irods_policy_query_processor",
8                  "parameters" : {
9                      "query_string" : "SELECT USER_NAME, COLL_NAME, DATA_NAME, RESC_NAME WHERE
10                                     COLL_NAME like '/tempZone/home/rods%' AND
11                                     RESC_NAME IN ('source_resource_1', 'source_resource_2')",
12                  "query_limit" : 10,
13                  "query_type" : "general",
14                  "number_of_threads" : 4,
15                  "policy_to_invoke" : "irods_policy_data_retention",
16                  "configuration" : {
17                      "mode" : "trim_single_replica",
18                      "source_resource_list" : ["source_resource_1", "source_resource_2"]
19                  }
20              }
21          }
22      }
23  }
```



```
1  {
2    "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3    "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4    "plugin_specific_configuration": {
5      "policies_to_invoke" : [
6        { "active_policy_clauses" : ["post"],
7          "events" : ["create", "write", "registration"],
8          "policy" : "irods_policy_data_verification",
9          "configuration" : {
10             "attribute" : "irods::verification::type"
11           }
12        }
13      ]
14    }
15  }
```

The type of verification to perform is stored as metadata on the resource

- catalog
- filesystem
- checksum





```
1  {
2    "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3    "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4    "plugin_specific_configuration": {
5      "policies_to_invoke" : [
6        {
7          "active_policy_clauses" : ["post"],
8          "events" : ["create", "write", "registration"],
9          "policy" : "irods_policy_enqueue_rule",
10         "delay_conditions" : "<ET>PLUSET 1</ET>",
11         "payload" : {
12           "policy" : "irods_policy_execute_rule",
13           "payload" : {
14             "policy" : "irods_policy_data_verification",
15             "configuration" : {
16               "attribute" : "irods::verification::type"
17             }
18           }
19         }
20       ]
21     }
22 }
```

The type of verification to perform is stored as metadata on the resource

- catalog
- filesystem
- checksum



Policy Composed Capabilities





Storage Tiering Overview



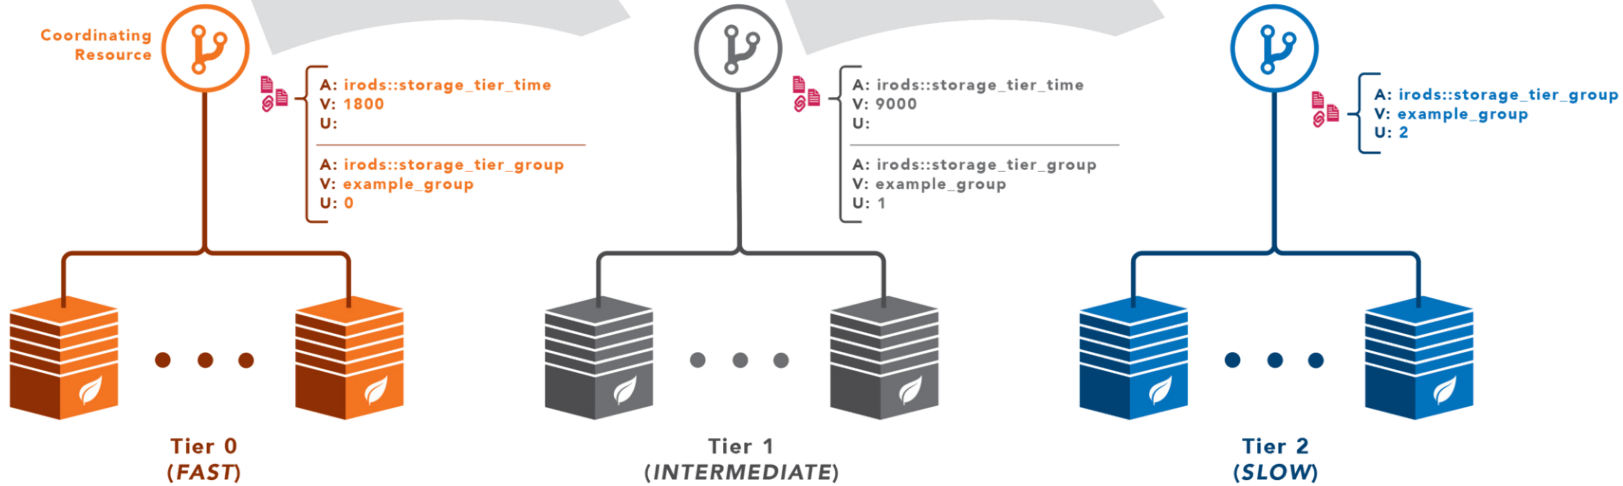
Periodically, the storage tiering policy discovers data objects in violation via a **default query** and schedules their migration to the next tier group.

After 1800 seconds, any data objects in violation are automatically replicated to tier 1, and then once at rest, they are trimmed from tier 0.

After 9000 seconds, any data objects in violation are automatically replicated to tier 2, and then once at rest, they are trimmed from tier 1.

UNIFIED NAMESPACE

DATA SOURCES



YOUR ORGANIZATION

The **default query** that determines which data objects are in violation can be overridden by adding a new metadata attribute `irods::storage_tier_query` with a value that defines the custom query.



FEDERATE SECURELY

OTHER ORGANIZATION

Data Virtualization (Unified Namespace)

Data Discovery (Metadata)

Workflow Automation (Rule Engine)

Secure Collaboration (Federation)

- Asynchronous Discovery
- Asynchronous Replication
- Synchronous Retention
- Resource associated metadata
- Identified by 'tiering groups'





Asynchronous Discovery and Replication

```

1 {
2   "policy" : "irods_policy_execute_rule",
3   "payload" : {
4     "policy_to_invoke" : "irods_policy_query_processor",
5     "configuration" : {
6       "query_string" : "SELECT META_RESC_ATTR_VALUE WHERE META_RESC_ATTR_NAME = 'irods::storage_tiering::group'",
7       "query_limit" : 0,
8       "query_type" : "general",
9       "number_of_threads" : 8,
10      "policy_to_invoke" : "irods_policy_event_generator_resource_metadata",
11      "configuration" : {
12        "conditional" : {
13          "metadata" : {
14            "attribute" : "irods::storage_tiering::group",
15            "value" : "{0}"
16          }
17        },
18        "policies_to_invoke" : [
19          {
20            "policy" : "irods_policy_query_processor",
21            "configuration" : {
22              "query_string" : "SELECT META_RESC_ATTR_VALUE WHERE META_RESC_ATTR_NAME = 'irods::storage_tiering::query' AND RESC_NAME = 'IRODS_TOKEN_S",
23              "default_results_when_no_rows_found" : ["SELECT USER_NAME, COLL_NAME, DATA_NAME, RESC_NAME WHERE META_DATA_ATTR_NAME = 'irods::access_ti",
24              "query_limit" : 0,
25              "query_type" : "general",
26              "number_of_threads" : 8,
27              "policy_to_invoke" : "irods_policy_query_processor",
28              "configuration" : {
29                "lifetime" : "IRODS_TOKEN_QUERY_SUBSTITUTION_END_TOKEN(SELECT META_RESC_ATTR_VALUE WHERE META_RESC_ATTR_NAME = 'irods::storage_tieri",
30                "query_string" : "{0}",
31                "query_limit" : 0,
32                "query_type" : "general",
33                "number_of_threads" : 8,
34                "policy_to_invoke" : "irods_policy_data_replication",
35                "configuration" : {
36                  "comment" : "source_resource, and destination_resource supplied by the resource metadata event generator"
37                }
38              }
39            }
40          }
41        ]
42      }
43    }
44  }
45 }
46 INPUT null
47 OUTPUT ruleExecOut

```





Synchronous Configuration for Storage Tiering

```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "active_policy_clauses" : ["post"],
8         "events" : ["put", "get", "create", "read", "write", "rename", "register", "unregister", "replication", "checksum", "copy", "seek", "truncate"],
9         "policy" : "irods_policy_access_time",
10        "configuration" : {
11          "log_errors" : "true"
12        }
13      },
14      {
15        "active_policy_clauses" : ["post"],
16        "events" : ["read", "write", "get"],
17        "policy" : "irods_policy_data_restage",
18        "configuration" : {
19        }
20      },
21      {
22        "active_policy_clauses" : ["post"],
23        "events" : ["replication"],
24        "policy" : "irods_policy_tier_group_metadata",
25        "configuration" : {
26        }
27      },
28      {
29        "active_policy_clauses" : ["post"],
30        "events" : ["replication"],
31        "policy" : "irods_policy_data_verification",
32        "configuration" : {
33        }
34      }
35    ],
36    {
37      "active_policy_clauses" : ["post"],
38      "events" : ["replication"],
39      "policy" : "irods_policy_data_retention",
40      "configuration" : {
41        "mode" : "trim_single_replica",
42        "log_errors" : "true"
43      }
44    }
45  }
46 }
47 ]
48 }
49 }
```



Possible Metadata Driven Restage for Storage Tiering

```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-metadata_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-metadata_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "active_policy_clauses" : ["post"],
8         "events" : ["set", "add"],
9         "attribute" : "irods::storage_tiering::restage",
10        "value" : "*.*",
11        "policy" : "irods_policy_data_restage",
12        "configuration" : {
13          }
14      }
15    ]
16  }
17 }
```



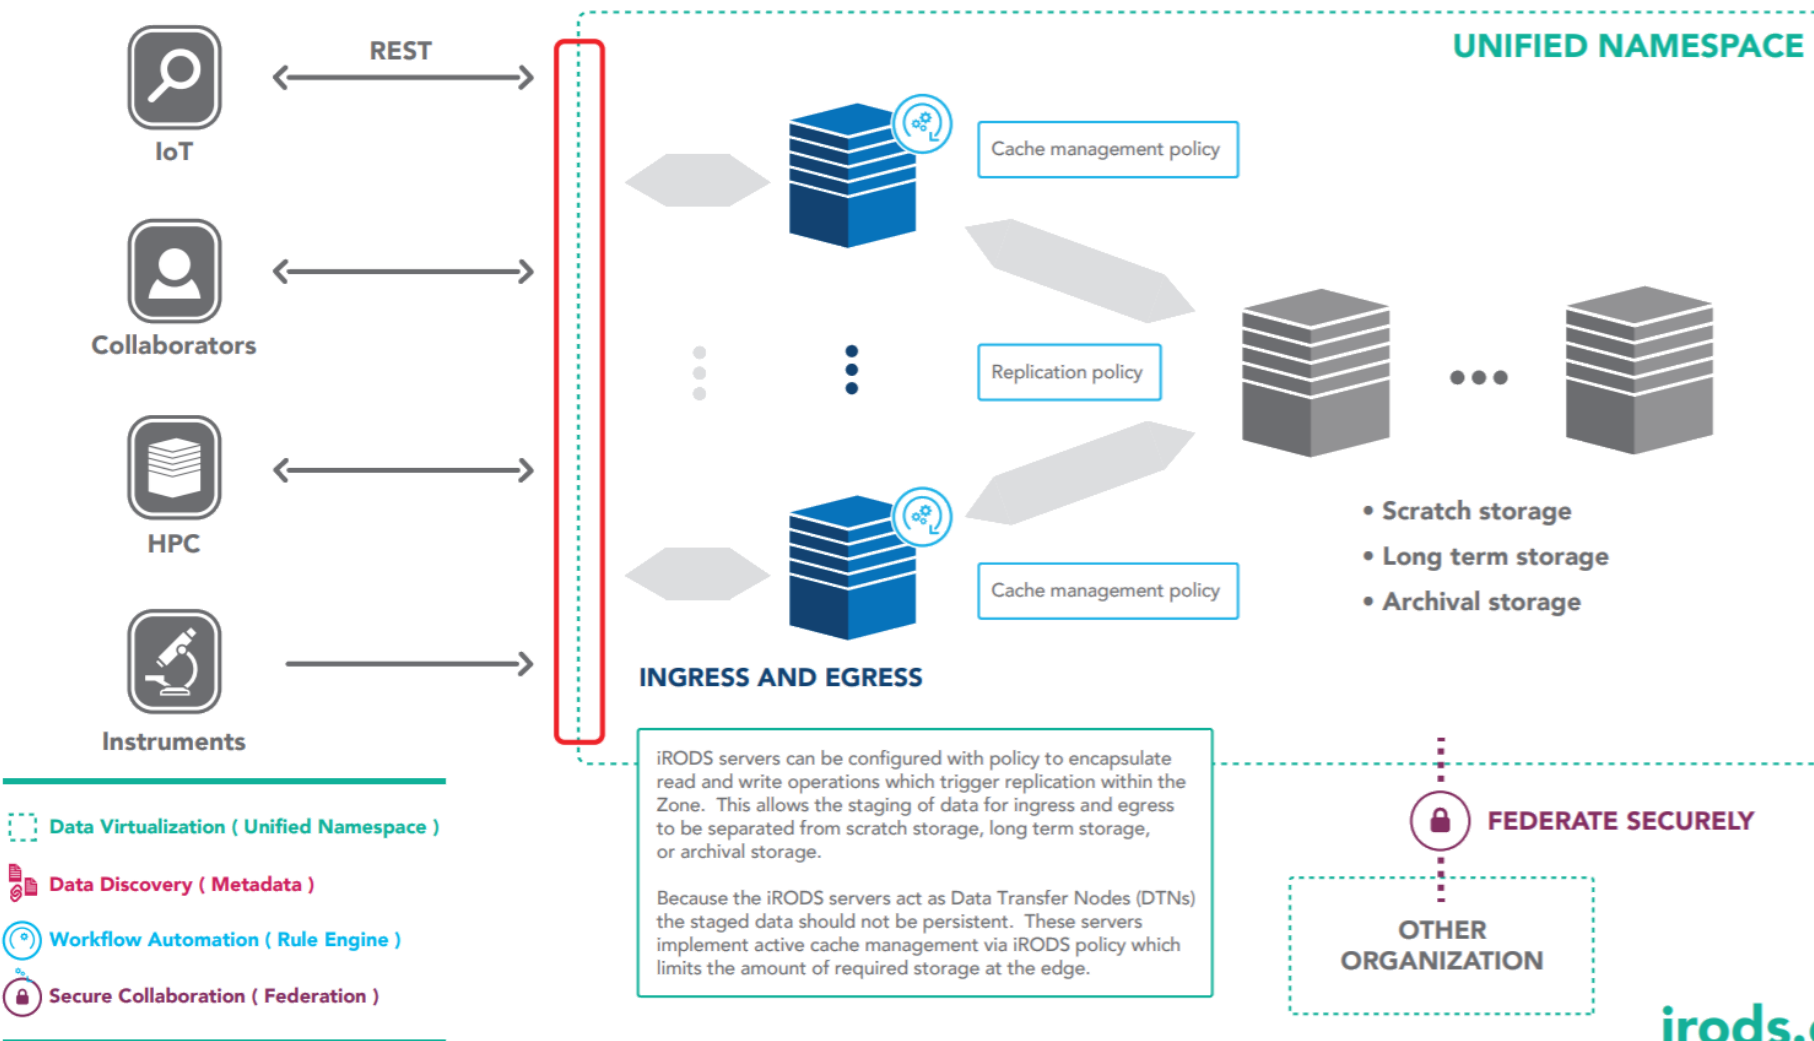

Data Transfer Nodes Pattern

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



- Asynchronous Discovery
- Asynchronous Retention
- Synchronous Replication
- Resource associated metadata
- Identified by 'replication groups'





Asynchronous Retention on Edge Resources

```
1 {
2   "policy" : "irods_policy_enqueue_rule",
3   "delay_conditions" : "<EF>REPEAT FOR EVER</EF>",
4   "payload" : {
5     "policy" : "irods_policy_execute_rule",
6     "payload" : {
7       "policy_to_invoke" : "irods_policy_query_processor",
8       "parameters" : {
9         "query_string" : "SELECT USER_NAME, COLL_NAME, DATA_NAME, RESC_NAME WHERE COLL_NAME",
10        "query_limit" : 10,
11        "query_type" : "general",
12        "number_of_threads" : 4,
13        "policy_to_invoke" : "irods_policy_data_retention",
14        "configuration" : {
15          "mode" : "trim_single_replica",
16          "source_resource_list" : ["edge_resource_1", "edge_resource_2"]
17        }
18      }
19    }
20  }
21 }
```



Synchronous Replication

```
1 {
2   "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
3   "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
4   "plugin_specific_configuration": {
5     "policies_to_invoke" : [
6       {
7         "conditional" : {
8           "logical_path" : "\/tempZone.*"
9         },
10        "pre_or_post_invocation" : ["post"],
11        "events" : ["create", "write", "registration"],
12        "policy" : "irods_policy_data_replication",
13        "configuration" : {
14          "source_to_destination_map" : {
15            "edge_resource_0" : ["long_term_resource_0"],
16            "edge_resource_1" : ["long_term_resource_1"],
17          }
18        }
19      },
20      {
21        "conditional" : {
22          "logical_path" : "\/tempZone.*"
23        },
24        "active_policy_clauses" : ["pre"],
25        "events" : ["get"],
26        "policy" : "irods_policy_data_replication",
27        "configuration" : {
28          "source_to_destination_map" : {
29            "long_term_resource_0" : ["edge_resource_0"],
30            "long_term_resource_1" : ["edge_resource_1"]
31          }
32        }
33      }
34    ]
35  }
36 }
```

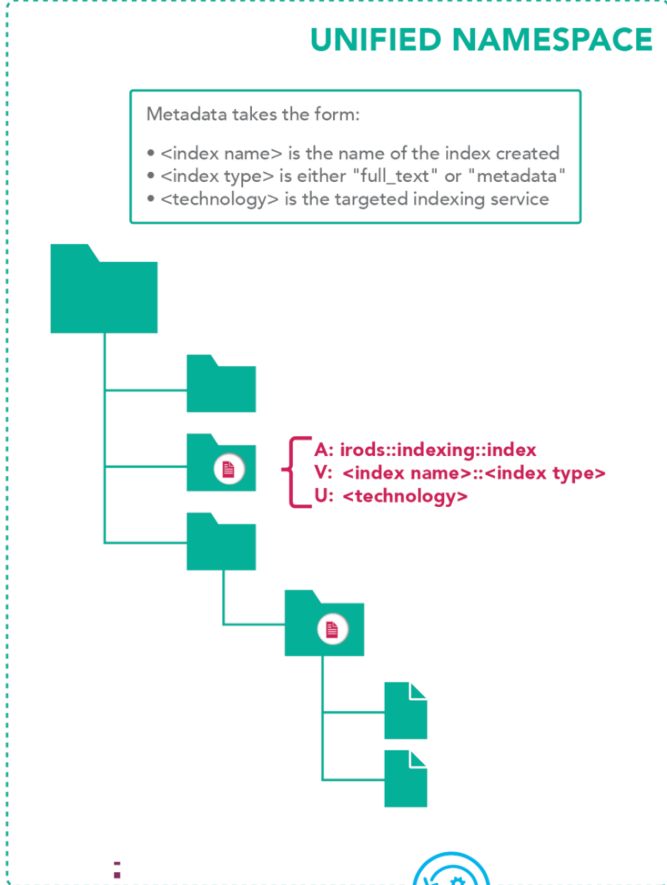
The iRODS Indexing Capability provides a policy framework around both full text and metadata indexing for the purposes of enhanced data discovery.

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>



FEDERATE SECURELY



Once indexing metadata is applied indicating that a collection should be indexed, a job is submitted to the iRODS delayed execution queue which will perform the requested action asynchronously.

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

Policy implemented as separate policy engine plugins

- `irods_policy_indexing_full_text_index_elasticsearch`
- `irods_policy_indexing_full_text_purge_elasticsearch`
- `irods_policy_indexing_metadata_index_elasticsearch`
- `irods_policy_indexing_metadata_purge_elasticsearch`



Synchronously configured full text indexing

```
1 "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-instance",
2 "plugin_name": "irods_rule_engine_plugin-event_handler-data_object_modified",
3 'plugin_specific_configuration': {
4   "policies_to_invoke" : [
5     {
6       "active_policy_clauses" : ["post"],
7       "events" : ["put", "write"],
8       "policy" : "irods_policy_event_delegate_collection_metadata",
9       "configuration" : {
10        "policies_to_invoke" : [
11          {
12            "conditional" : {
13              "metadata" : {
14                "attribute" : "irods::indexing::index",
15                "entity_type" : "data_object"
16              },
17            },
18            "policy" : "irods_policy_indexing_full_text_index_elasticsearch",
19            "configuration" : {
20              "hosts" : ["http://localhost:9200/"],
21              "bulk_count" : 100,
22              "read_size" : 1024
23            }
24          }
25        ]
26      }
27    }
28    ...
```


Synchronously configured full text purge

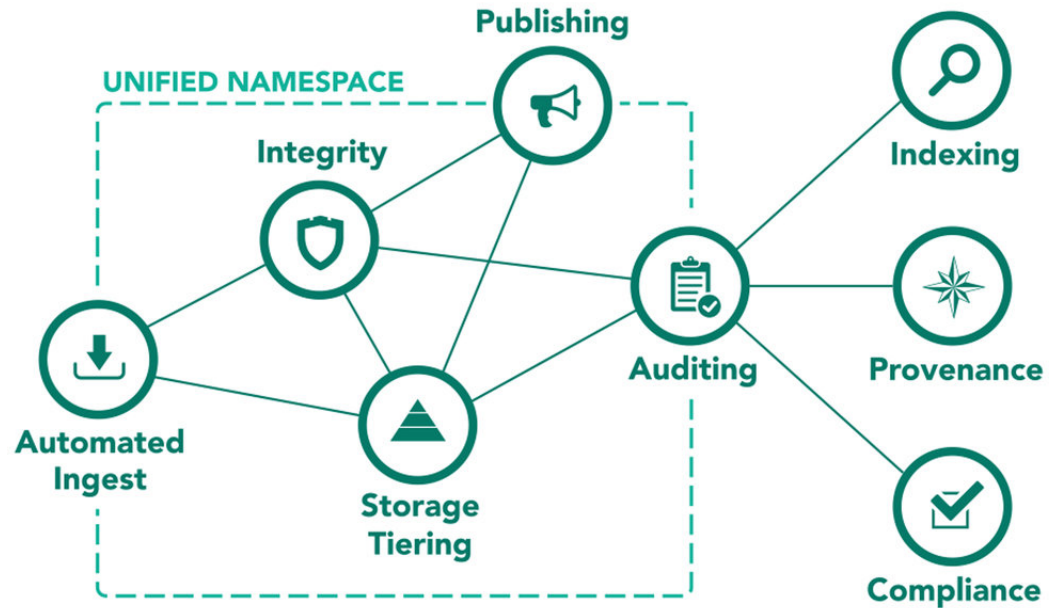
```
1  {
2      "active_policy_clauses" : ["pre"],
3      "events" : ["unlink", "unregister"],
4      "policy" : "irods_policy_event_delegate_collection_metadata",
5      "configuration" : {
6          "policies_to_invoke" : [
7              {
8                  "conditional" : {
9                      "metadata" : {
10                         "attribute" : "irods::indexing::index",
11                         "entity_type" : "data_object"
12                     },
13                 },
14                 "policy" : "irods_policy_indexing_full_text_purge_elasticsearch",
15                 "configuration" : {
16                     "hosts" : ["http://localhost:9200/"],
17                     "bulk_count" : 100,
18                     "read_size" : 1024
19                 }
20             }
21         ]
22     }
23 }
24 ]
25 }
```

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

Capabilities become easily configured recipes.

**A Policy GUI is now a possibility with
simple manipulation of server side JSON.**

**Data management
should be
data-centric and metadata driven.**

**Future-proof automated data management
requires
open formats and open source.**

