

iRODS

Policy Composition


Jason M. Coposky
@jason_coposky
Executive Director, iRODS Consortium

June 9-12, 2020
iRODS User Group Meeting 2020
Virtual Event

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

The iRODS Technology Stack

iRODS



What is Data Management

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

What is Policy

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
- ...



In /etc/irods/core.re ...

```
1 acPostProcForPut() {
2     if($rescName == "demoResc") {
3         # extract and apply metadata
4     }
5     else if($rescName == "cacheResc") {
6         # async replication to archive
7     }
8     else if($objPath like "/tempZone/home/alice/*" &&
9             $rescName == "indexResc") {
10        # launch an indexing job
11    }
12    else if(xyz) {
13        # compute checksums ...
14    }
15
16    # and so on ...
17 }
```



Expanding policy implementation across rule bases

For example: pep_data_obj_put_post(...)

- Metadata extraction and application
- Asynchronous Replication
- Initiate Indexing
- Apply access time metadata
- Asynchronous checksum computation

Rather than one monolithic implementation, separate the implementations into individual rule bases, or plugins, and allow the rule(s) to fall through



Separate the implementation into several rule bases:

/etc/irods/metadata.re

```
1 pep_api_data_obj_put_post(*INSTANCE_NAME, *COMM, *DATAOBJINP, *BUFFER, *PORTAL_OPR_O
2   # metadata extraction and application code
3
4   RULE_ENGINE_CONTINUE
5 }
```

/etc/irods/checksum.re

```
1 pep_api_data_obj_put_post(*INSTANCE_NAME, *COMM, *DATAOBJINP, *BUFFER, *PORTAL_OPR_O
2   # checksum code
3
4   RULE_ENGINE_CONTINUE
5 }
```

/etc/irods/access_time.re

```
1 pep_api_data_obj_put_post(*INSTANCE_NAME, *COMM, *DATAOBJINP, *BUFFER, *PORTAL_OPR_O
2   # access time application code
3
4   RULE_ENGINE_CONTINUE
5 }
```



Within the Rule Engine Plugin Framework, order matters

```
1  "rule_engines": [
2      {
3          "instance_name": "irods_rule_engine_plugin-irods_rule_language-inst",
4          "plugin_name": "irods_rule_engine_plugin-irods_rule_language",
5          "plugin_specific_configuration": {
6              ...
7              "re_rulebase_set": [
8                  "metadata",
9                  "checksum",
10                 "access_time",
11                 "core"
12             ],
13             ...
14         },
15         "shared_memory_instance" : "irods_rule_language_rule_engine"
16     },
17     {
18         "instance_name": "irods_rule_engine_plugin-cpp_default_policy-inst",
19         "plugin_name": "irods_rule_engine_plugin-cpp_default_policy",
20         "plugin_specific_configuration": {
21             ...
22         }
23     }
]
```



Consider Policy as building blocks towards Capabilities

Follow proven software engineering principles:

Favor composition over monolithic implementations

Provide a common interface across policy implementations
to allow transparent configuration



Consider Storage Tiering as a collection of policies:

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



Policies invoked by monolithic framework plugins
and delegated by convention:

- irods_policy_access_time
- irods_policy_data_movement
- irods_policy_data_replication
- irods_policy_data_verification
- irods_policy_data_retention

Each policy may be implemented by any rule engine, or rule base to customize for future use cases or technologies



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.

The When



When - Event Handlers

iRODS

RPC API

```
audit_pep_auth_agent_auth_request_post  
audit_pep_auth_agent_auth_request_pre  
audit_pep_auth_agent_auth_response_post  
audit_pep_auth_agent_auth_response_pre  
audit_pep_auth_agent_start_post  
audit_pep_auth_agent_start_pre  
audit_pep_auth_request_post  
audit_pep_auth_request_pre  
audit_pep_auth_response_post  
audit_pep_auth_response_pre  
audit_pep_data_obj_put_post  
audit_pep_data_obj_put_pre  
audit_pep_database_check_auth_post  
audit_pep_database_check_auth_pre  
audit_pep_database_close_post  
audit_pep_database_close_pre  
audit_pep_database_gen_query_access_control_setup_post  
audit_pep_database_gen_query_access_control_setup_pre  
audit_pep_database_gen_query_post  
audit_pep_database_gen_query_pre  
audit_pep_database_get_rcs_post  
audit_pep_database_get_rcs_pre  
audit_pep_database_mod_data_obj_meta_post  
audit_pep_database_mod_data_obj_meta_pre  
audit_pep_database_open_post  
audit_pep_database_open_pre  
audit_pep_database_reg_data_obj_post  
audit_pep_database_reg_data_obj_pre  
audit_pep_exec_microservice_post  
audit_pep_exec_microservice_pre  
audit_pep_exec_rule_post  
audit_pep_exec_rule_pre  
audit_pep_network_agent_start_post  
audit_pep_network_agent_start_pre  
audit_pep_network_agent_stop_post  
audit_pep_network_agent_stop_pre  
audit_pep_network_read_body_post  
audit_pep_network_read_body_pre  
audit_pep_network_read_header_post  
audit_pep_network_read_header_pre  
audit_pep_network_write_body_post  
audit_pep_network_write_body_pre  
audit_pep_network_write_header_post  
audit_pep_network_write_header_pre  
audit_pep_obj_stat_post  
audit_pep_obj_stat_pre  
audit_pep_resource_close_post  
audit_pep_resource_close_pre  
audit_pep_resource_create_post  
audit_pep_resource_create_pre  
audit_pep_resource_modified_post  
audit_pep_resource_modified_pre  
audit_pep_resource_registered_post  
audit_pep_resource_registered_pre  
audit_pep_resource_resolve_hierarchy_post  
audit_pep_resource_resolve_hierarchy_pre  
audit_pep_resource_stat_post  
audit_pep_resource_stat_pre  
audit_pep_resource_write_post  
audit_pep_resource_write_pre
```

input

EVENT HANDLERS

- Create
- Write
- Read
- Replication
- Unlink
- Rename
- Register

POLICY INVOCATIONS

iRODS_Policy_Example



A Rule Engine Plugin for a specific Class of events

- Data Object
- Collection
- Metadata
- User
- Resource

The Events are specific to the class of the handler

The handler then invokes policy based on its configuration



A Rule Engine Plugin for data creation and modification events

- Create
- Read
- Replication
- Unlink
- Rename
- ...

Policy invocation is configured as an array of json objects for any given combination of events

Unifies the POSIX and Object behaviors into a single place to configure policy



Example : Synchronous Invocation

```
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_access_time",
10                 "configuration" : {
11                     }
12                 },
13                 {
14                     "active_policy_clauses" : ["pre"],
15                     "events" : ["replication"],
16                     "policy" : "irods_policy_example_policy",
17                     "configuration" : {
18                         }
19                     }
20                 ]
21             }
22 }
```

Note that order still matters if more than one policy is configured for a given event

The What



Basic policies that are leveraged across many deployments and capabilities:

- irods_policy_access_time
- irods_policy_query_processor
- irods_policy_data_movement
- irods_policy_data_replication
- irods_policy_data_verification
- irods_policy_data_retention

The library will continue to grow, with a cookbook of usages



Standardized serialized JSON string interface :
parameters, and configuration

iRODS Rule Language

```
1 irods_policy_example_policy_implementation(*parameters, *configuration)
2     writeLine("stdout", "Hello UGM2020!")
3 }
```

Python Rule Language

```
1 def irods_policy_example_policy_implementation(rule_args, callback,
2 # Parameters      rule_args[1]
3 # Configuration   rule_args[2]
```

Policy can also be implemented as fast and light C++ rule engine plugins termed Policy Engines



Policy may be invoked using one of three different conventions:

- Direct Invocation : a JSON object
- Query Processor : array of query results in a JSON object
- Event Handler : a JSON object

Each invocation convention defines its interface by contract



Parameters passed as serialized JSON strings

```
1 my_rule() {
2     irods_policy_access_time( "{\"object_path\" : \"/tempZone/home/rods/file0.txt\"}"
3 }
```

Directly invoked policy via irule

```
1 {
2     "policy" : "irods_policy_execute_rule",
3     "payload" : {
4         "policy_to_invoke" : "irods_policy_storage_tiering",
5         "parameters" : {
6             "object_path" : "/tempZone/home/rods/file0.txt"
7         },
8         "configuration" : {
9             }
10    }
11 }
12 INPUT null
13 OUTPUT ruleExecOut
```

Serializes results to JSON array and passed to the policy via the parameter object as "query_results"

```
1  {
2      "policy" : "irods_policy_enqueue_rule",
3      "delay_conditions" : "<PLUSET>1s</PLUSET>",
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 like '/tempZone/home',
10                 "query_limit" : 10,
11                 "query_type" : "general",
12                 "number_of_threads" : 4,
13                 "policy_to_invoke" : "irods_policy_engine_example"
14             }
15         }
16     }
17 }
```

For example the invoked policy would receive a row like:

"query_results" : ['rods', '/tempZone/home/rods', 'file0.txt', 'demoResc']



What - Event Handler Invocation

iRODS

```
1  {
2      "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-in",
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"],
9                  "policy": "irods_policy_data_replication",
10                 "configuration": {
11                     "source_to_destination_map": {
12                         "demoResc": ["AnotherResc"]
13                     }
14                 }
15             },
16             ...
17         ]
18         ...
19     }
20 }
```

Serializes dataObjInp_t and rsComm_t to the Parameter object

```
1  {
2      "comm": {
3          "auth_scheme": "native", "client_addr": "152.54.8.141", "proxy_auth_info_auth_flag": "5", "proxy_auth_info_auth_scheme": "",
4          "proxy_auth_info_auth_str": "", "proxy_auth_info_flag": "0", "proxy_auth_info_host": "", "proxy_auth_info_ppid": "0",
5          "proxy_rods_zone": "tempZone", "proxy_sys_uid": "0", "proxy_user_name": "rods", "proxy_user_other_info_user_comments": "",
6          "proxy_user_other_info_user_create": "", "proxy_user_other_info_user_info": "", "proxy_user_other_info_user_modify": "",
7          "proxy_user_type": "", "user_auth_info_auth_flag": "5", "user_auth_info_auth_scheme": "", "user_auth_info_auth_str": "",
8          "user_auth_info_flag": "0", "user_auth_info_host": "", "user_auth_info_ppid": "0", "user_rods_zone": "tempZone",
9          "user_sys_uid": "0", "user_user_name": "rods", "user_user_other_info_user_comments": "", "user_user_other_info_user_create": "",
10         "user_user_other_info_user_info": "", "user_user_other_info_user_modify": "", "user_user_type": ""
11     },
12     "cond_input": {
13         "dataIncluded": "", "dataType": "generic", "destRescName": "ufs0", "noOpenFlag": "", "openType": "1",
14         "recursiveOpr": "1", "resc_hier": "ufs0", "selObjType": "dataObj", "translatedPath": ""
15     },
16     "create_mode": "33204",
17     "data_size": "1",
18     "event": "CREATE",
19     "num_threads": "0",
20     "obj_path": "/tempZone/home/rods/test_put_gt_max_sql_rows/junk0083",
21     "offset": "0",
22     "open_flags": "2",
23     "opr_type": "1",
24     "policy_enforcement_point": "pep_api_data_obj_put_post"
25 }
```



Any additional static context passed into the policy

```
1 {
2     "policy" : "irods_policy_access_time",
3     "configuration" : {
4         "attribute" : "irods::access_time"
5     }
6 }
```

May be "plugin_specific_configuration" from a rule engine plugin
or "configuration" from within the event framework

May hold additional policy which to be subsequently invoked,
e.g. the Query Processor

The Why



Each invoked policy may set a conditional around each noun within the system which gates the invocation

- Data Object
- Collection
- Metadata
- User
- Resource

Leverages boost::regex to match any combination of logical_path, metadata, resource name, or user name



Matching a logical path for replication policy invocation

```
1  {
2      "instance_name": "irods_rule_engine_plugin-event_handler-data_object_modified-in",
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                 "active_policy_clauses" : ["post"],
11                 "events" : ["put"],
12                 "policy" : "irods_policy_data_replication",
13                 "configuration" : {
14                     "source_to_destination_map" : {
15                         "demoResc" : ["AnotherResc"]
16                     }
17                 }
18             },
19             ...
20         ]
21         ...
22     }
23 }
```



Matching metadata for indexing policy invocation

```
1 import shutil
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", "write"],
9             "policy" : "irods_policy_event_delegate_collection_metadata",
10            "configuration" : {
11                "policies_to_invoke" : [
12                    {
13                        "conditional" : {
14                            "metadata" : {
15                                "attribute" : "irods::indexing::index",
16                                "entity_type" : "data_object"
17                            },
18                        },
19                        "policy" : "irods_policy_indexing_full_text_index_elasticsearch",
20                        "configuration" : {
21                            "hosts" : ["http://localhost:9200/"],
22                            "bulk_count" : 100,
23                            "read_size" : 1024
24                        }
25                    }
26                ]
27            }
28        ]
29    }
30 }
```

The How



The cpp_default rule engine plugin in 4.2.8 will now support two new policies:

- irods_policy_enqueue_rule
- irods_policy_execute_rule

```
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" : "irods_policy_example",
8             "configuration" : {
9                 }
10            }
11        }
12    }
13 INPUT null
14 OUTPUT ruleExecOut
```

The enqueue rule policy will push a job onto the delayed execution queue. The "payload" object holds the rule which is to be executed.



The execute rule policy will invoke a policy engine either from the delayed execute queue or as a direct invocation

```
1  {
2      "policy" : "irods_policy_execute_rule",
3          "payload" : {
4              "policy_to_invoke" : "irods_policy_example",
5                  "parameters" : {
6                      },
7                  "configuration" : {
8                      }
9          }
10     }
11 }
12 INPUT null
13 OUTPUT ruleExecOut
```



Sample Delayed Rule for Asynchronous Execution by the cpp default rule engine

```
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_example",
8              "parameters" : {
9
10                 },
11                 "configuration" : {
12
13                     }
14                 }
15             }
16 }
17 INPUT null
18 OUTPUT ruleExecOut
```


We no longer need to pay the penalty of instantiating an interpreted language

Policy Composed Capabilities



Storage Tiering Overview

iRODS





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

Asynchronous 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
```

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         {
30             "active_policy_clauses" : ["post"],
31             "events" : ["replication"],
32             "policy" : "irods_policy_data_verification",
33             "configuration" : {
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 }
```



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                  "conditional" : {
8                      "attribute" : "irods::storage_tiering::restage",
9                  },
10                 "active_policy_clauses" : ["post"],
11                 "events" : ["set", "add"],
12                 "policy" : "irods_policy_data_restage",
13                 "configuration" : {
14                     }
15                 }
16             ]
17         }
18     }
```



Data Transfer Nodes Pattern

iRODS


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_NAM
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                 "active_policy_clauses" : [ "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 }
```

Indexing Capability

iRODS

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>`

UNIFIED NAMESPACE

Metadata takes the form:

- `<index name>` is the name of the index created
- `<index type>` is either "full_text" or "metadata"
- `<technology>` is the targeted indexing service



FEDERATE SECURELY



OTHER
ORGANIZATION

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)



Implemented as individual Policy Engines

- 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 }
```



- Capabilities become recipes which are easily configured
- A Policy GUI is now a possibility with the manipulation of server side JSON
- Continue to build a library of supported policy engines, driven by the community
- Data Integrity Capability will now be a collection of policy engines

Questions?

