



iRODS Policy Composition

Jason M. Coposky
@jason_coposky
Executive Director, iRODS Consortium

June 8-11, 2021
iRODS User Group Meeting 2021
Virtual Event

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 a 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
- ...

The iRODS Data Management Model



Some Questions

- 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



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 }
```

Assuming there was even a provided policy enforcement point for the desired event...



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_OUT) {  
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_OUT) {  
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_OUT) {  
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-instance",
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-instance",
19             "plugin_name": "irods_rule_engine_plugin-cpp_default_policy",
20             "plugin_specific_configuration": {
21                 }
22             }
23         ]
24     ]
```

Consider Storage Tiering as a collection of policies:

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



Policies composed by a monolithic framework plugins

Policy delegated by naming 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 overridden by another rule engine, or rule base to customize to 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

input

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
```

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
- Register

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

When - event_handler-data_object_modified

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

Note that order still matters if more than one policy needs to be invoked 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.

What - Simple policy implementations

Standardized JSON interface : parameters, and configuration

iRODS Rule Language

```
1 irods_policy_example_policy_implementation(*parameters, *configuration) {  
2  
3 }
```

Python Rule Language

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

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

What - Simple policy implementations

Policy may be invoked using one of three different conventions:

- Direct Invocation : a JSON object
- Query Processor : a JSON array of parameters
- Event Handler : a JSON object

Each invocation convention defines its interface by contract.

Direct Invocation : Parameters passed as a JSON object

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

Parameters may also be configured statically

```
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 }
```

Query Processor Invocation

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

```
1  {
2      "policy_to_invoke" : "irods_policy_enqueue_rule",
3      "parameters" : {
4          "delay_conditions" : "<PLUSET>1s</PLUSET>",
5          "policy_to_invoke" : "irods_policy_execute_rule",
6          "parameters" : {
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/rods%'",
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:

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

Event Handler Invocation

Serializes dataObjInp_t and rsComm_t to a JSON 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", "data_size": "1", "event": "CREATE", "num_threads": "0",
17  "obj_path": "/tempZone/home/rods/test_put_gt_max_sql_rows/junk0083",
18  "offset": "0", "open_flags": "2", "opr_type": "1",
19  "policy_enforcement_point": "pep_api_data_obj_put_post"
20 }
```

Which is also passed in as the parameter object

Configuration

Any additional statically set context passed into the policy

```
1 {
2     "policy_to_invoke" : "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

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, etc.



Matching a logical path for replication policy 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                  "conditional" : {
8                      "logical_path" : "\/tempZone.*"
9                  },
10                 "active_policy_clauses" : [ "post" ],
11                 "events" : [ "put" ],
12                 "policy_to_invoke" : "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  {
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                  "conditional" : {
10                      "metadata_exists" : {
11                          "recursive" : "true",
12                          "attribute" : "irods::indexing::index",
13                          "entity_type" : "collection"
14                      }
15                  },
16                  "policy_to_invoke" : "irods_policy_indexing_full_text_index_elasticsearch",
17                  "configuration" : {
18                      "hosts" : [ "http://localhost:9200/" ],
19                      "bulk_count" : 100,
20                      "read_size" : 1024
21                  }
22              }
23          ]
24      }
```

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

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



How - Asynchronous Execution

iRODS

```
1  {
2
3      "policy_to_invoke" : "irods_policy_enqueue_rule",
4      "parameters" : {
5          "comment" : "Set the PLUSET value to the interval desired to run the rule",
6          "delay_conditions" : "<PLUSET>10s</PLUSET><EF>REPEAT FOR EVER</EF><INST_NAME>irods_rule_engine_plugin-cpp_default_policy-instance</INST_NAME>",
7          "policy_to_invoke" : "irods_policy_execute_rule",
8          "parameters" : {
9              "policy_to_invoke" : "irods_policy_filesystem_usage",
10             "parameters" : {
11                 "source_resource" : "demoResc"
12             }
13         }
14     }
15 }
16 INPUT null
17 OUTPUT ruleExecOut
```

The execute rule policy invokes a policy engine either from the delayed execution queue or as a direct invocation



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

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_to_invoke" : "irods_policy_access_time",
12                  "configuration" : {
13                      }
14              }
15          ]
16      }
17 }
```



Synchronous Replication

iRODS

```
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          {
6              "policies_to_invoke": [
7                  {
8                      "active_policy_clauses": ["post"],
9                      "events": ["create", "write", "registration"],
10                     "policy_to_invoke": "irods_policy_data_replication",
11                     "configuration": {
12                         "source_to_destination_map": {
13                             "source_resource_0": ["destination_resource_0a", "destination_resource_0b"],
14                             "source_resource_1": ["destination_resource_1a"],
15                         }
16                     }
17                 },
18                 {
19                     "active_policy_clauses": ["post"],
20                     "events": ["create", "write", "registration"],
21                     "policy_to_invoke": "irods_policy_data_replication",
22                     "configuration": {
23                         "destination_resource": "destination_resource_3"
24                     }
25                 }
26             ]
27         }
28     }
29 }
```



Asynchronous Replication

iRODS

```
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_to_invoke" : "irods_policy_enqueue_rule",
10                "parameters" : {
11                    "delay_conditions" : "<ET>PLUSET 1</ET>",
12                    "policy_to_invoke" : "irods_policy_execute_rule",
13                    "parameters" : {
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_to_invoke" : "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_to_invoke" : "irods_policy_enqueue_rule",
3          "parameters" : {
4              "delay_conditions" : "<EF>REPEAT FOR EVER</EF>",
5              "policy_to_invoke" : "irods_policy_execute_rule",
6              "parameters" : {
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                  {
7                      "active_policy_clauses" : ["post"],
8                      "events" : ["create", "write", "registration"],
9                      "policy_to_invoke" : "irods_policy_data_verification",
10                     "configuration" : {
11                         "attribute" : "irods::verification::type"
12                     }
13                 }
14             ]
15         }
16     }
```

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_to_invoke": "irods_policy_enqueue_rule",
10                     "parameters": {
11                         "delay_conditions": "<ET>PLUSET 1</ET>",
12                         "policy": "irods_policy_execute_rule",
13                         "parameters": {
14                             "policy": "irods_policy_data_verification",
15                             "configuration": {
16                                 "attribute": "irods::verification::type"
17                             }
18                         }
19                     }
20                 }
21             ]
22         }
23     }
```

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


- catalog
- filesystem
- checksum

Policy Composed Capabilities



Storage Tiering Overview

iRODS



Data Virtualization (Unified Namespace)

Data Discovery (Metadata)

Workflow Automation (Rule Engine)

Secure Collaboration (Federation)

FEDERATE SECURELY

**OTHER
ORGANIZATION**



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



Asynchronous Discovery and Replication

```
1 {
2     "policy_to_invoke" : "irods_policy_execute_rule",
3     "parameters" : {
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_exists" : {
14                        "attribute" : "irods::storage_tiering::group",
15                        "value" : "{0}"
16                    }
17                }
18            },
19            "policies_to_invoke" : [
20                {
21                    "policy_to_invoke" : "irods_policy_query_processor",
22                    "configuration" : {
23                        "query_string" : "SELECT META_RESC_ATTR_VALUE WHERE META_RESC_ATTR_NAME = 'irods::storage_tiering::query' AND RESC_NAME = 'IRODS_TOKEN_S',
24                        "default_results_when_no_rows_found" : ["SELECT USER_NAME, COLL_NAME, DATA_NAME, RESC_NAME WHERE META_DATA_ATTR_NAME = 'irods::access_tiering::group'"],
25                        "query_limit" : 0,
26                        "query_type" : "general",
27                        "number_of_threads" : 8,
28                        "policy_to_invoke" : "irods_policy_query_processor",
29                        "configuration" : {
30                            "lifetime" : "IRODS_TOKEN_QUERY_SUBSTITUTION_END_TOKEN(SELECT META_RESC_ATTR_VALUE WHERE META_RESC_ATTR_NAME = 'irods::storage_tiering::group')",
31                            "query_string" : "{0}",
32                            "query_limit" : 0,
33                            "query_type" : "general",
34                            "number_of_threads" : 8,
35                            "policy_to_invoke" : "irods_policy_data_replication",
36                            "configuration" : {
37                                "comment" : "source_resource, and destination_resource supplied by the resource metadata event generator"
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_to_invoke" : "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_to_invoke" : "irods_policy_data_restage",
18                "configuration" : {
19                }
20            },
21            {
22                "active_policy_clauses" : ["post"],
23                "events" : ["replication"],
24                "policy_to_invoke" : "irods_policy_tier_group_metadata",
25                "configuration" : {
26                }
27            },
28            {
29                "active_policy_clauses" : ["post"],
30                "events" : ["replication"],
31                "policy_to_invoke" : "irods_policy_data_verification",
32                "configuration" : {
33                }
34            },
35            {
36            },
37            {
38                "active_policy_clauses" : ["post"],
39                "events" : ["replication"],
40                "policy_to_invoke" : "irods_policy_data_retention",
41                "configuration" : {
42                    "mode" : "trim_single_replica",
43                    "log_errors" : "true"
44                }
45            }
46        ]
47    }
48 }
```



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_to_invoke" : "irods_policy_data_restage",
12                 "configuration" : {
13                     }
14                 }
15             ]
16         }
17     }
```




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.**