# iRODS Logical Quotas Policy Plugin

-----

University of Colorado Boulder

# **Research Computing**

- Compute clusters
- Data storage
- Networking















#### **Logical Quotas**

- Available as an RPM and DEB
- Operates as a discrete rule engine

```
"rule_engines": [{
    "instance_name": "irods_rule_engine_plugin-logical_quotas-instance",
    "plugin_name": "irods_rule_engine_plugin-logical_quotas",
    "plugin_specific_configuration": {
        "namespace": "irods::logical_quotas",
        "metadata_attribute_names": {
            "maximum_number_of_data_objects": "maximum_number_of_data_objects",
            "maximum_size_in_bytes": "maximum_size_in_bytes",
            "total_number_of_data_objects": "total_number_of_data_objects",
            "total_size_in_bytes": "total_size_in_bytes"
        }
    }, ...]
```



💿 🕒 🖢 🖿 src — janderson@irods1:~ — ssh -A janderson@192.168.1.145 — 80×24	
<pre>Set is set</pre>	45 — 80×24 nderson/file-limi do iput empty-fil imited ed: <b>••••••••••••••••••••••••••••••••••••</b>
	[janderson@irods1 ~]\$

https://github.com/anderbubble/irods\_rule\_engine\_plugin\_logical\_quotas/tree/ilogiq











🔎 😑 🛑 🖿 src — janderson@irods1:~ — ssh -A janderson@192.168.1	1.145 — 80×24
<pre>     src — janderson@irods1:~ — ssh -A janderson@192.168.1 [[janderson@irods1 ~]\$ imkdir size-limited [[janderson@irods1 ~]\$ ilogiq start-monitoring /tempZone/home/jand 35MB [[janderson@irods1 ~]\$ dd if=/dev/urandom of=random-file bs=1M 10+0 records in 10+0 records out 10485760 bytes (10 MB) copied, 0.0652263 s, 161 MB/s [[janderson@irods1 ~]\$ for f in size-limited/random-file.{14 ile \$f; done remote addresses: 192.168.1.145 ERROR: putUtil: put error for derson/size-limited/random-file.4, status = -130000 status = _INPUT_PARAM Level 0: Logical Quotas Policy Violation: Adding object excee e in bytes limit [janderson@irods1 ~]\$ </pre>	<pre>1.145 — 80×24 janderson/size-limi] erson/size-limited</pre>
[janderson@irods1 ~]\$	<pre>attribute: irods::logical_quotas::total_number_of_data_objects value: 3 units: attribute: irods::logical_quotas::total_size_in_bytes value: 31457280 units: [janderson@irods1 ~]\$</pre>





### **Logical Quotas Experience**

- Real-world vs synthetic
- Experienced some drift in synthetic workloads
- Expecting to be able to clean it up in batch
- Maybe locking later?

https://gist.github.com/anderbubble/027de8b80366cd63df34b3838a7bab1d





#### **Design and Implementation**

- Initially implemented using the Native Rule Language
  - Pre-PEPs: Compute deltas and detect violations
  - Post-PEPs: Update totals using previously computed deltas
  - Successful as a proof of concept
  - Limited
- Reimplemented as a Rule Engine Plugin using C++
  - Made it easier to implement
  - Enabled better encapsulation and design
  - Performance

https://github.com/irods/irods\_rule\_engine\_plugin\_logical\_quotas



#### **C++ Rule Engine Plugin Design**

```
1 class pep api data obj copy final
 2 {
 3 public:
       pep api data obj copy() = delete;
 4
 5
       static auto reset() noexcept -> void;
 6
       static auto pre(...) -> irods::error;
 8
 9
10
       static auto post(...) -> irods::error;
11
12 private:
      // Shared data.
13
14 };
15
16 map<string view, handler type> handlers{
       {"pep api data obj copy pre", pep api data obj copy::pre},
17
       {"pep api data obj copy post", pep api data obj copy::post}
18
19 }:
```





#### **Performance over Correctness**

Correctness Pros / Cons

- Pro: Reported numbers perfectly reflect the zone
- Con: Requires zone-wide synchronization
- Con: Decreased responsiveness on deeply nested collections

Performance Pros / Cons

- Pro: Does not require synchronization at all
- Pro: Interaction stays responsive
- Pro: Can always synchronize the logical totals with the real data
- Pro: The size of the data set does not affect performance
- · Con: Reported numbers are relative due to drift



# **Interesting Findings**

- Opening a data object with O\_TRUNC was not reflected in the catalog immediately
  - Some operations require the data object's size for proper operation
  - Use **dstream**'s seek operation
- Helped lead to the development of **istream** 
  - Required a tool to test stream handlers
- The iRODS Filesystem library drastically simplified development
- Lead to the development of the <u>metadata guard</u> rule engine plugin

https://github.com/irods/irods\_rule\_engine\_plugin\_logical\_quotas



#### **Questions and Contact Info**

- github.com/irods/irods\_rule\_engine\_plugin\_logical\_quotas
- irods.org
- <u>colorado.edu/rc</u>
- jonathon.anderson@colorado.edu
- irod-chat@googlegroups.com



