Table of Contents

Rucio in Hadoop

Introduction

What's the initial reason this data has been imported? The main developers are Thomas Beermann and Mario Lassnig.

The Data

The data is stored in the `analytix.cern.ch` cluster in the directory:

```
/user/rucio01/
```

Example of access (start from lxplus):

```
$ ssh analytix

$ hadoop fs -ls /user/rucio01/traces
$ hadoop fs -cat /user/rucio01/traces/rucio-server-prod-05.cern.ch.1416407134480 | head
$ hadoop fs -ls /user/rucio01/dq2/traces
$ hadoop fs -cat /user/rucio01/dq2/traces/2014-08 | head
```

Apache Server / Rucio Daemon logs

Stores log files for simple simple cat / grep analysis

- read directly from log file and continuously streamed via Flume to HDFS
- simple text log files
- ~23GB per day

Traces

Contain updates of last access time of files/datasets, will be used for the popularity reports

- update of last access time of files/datasets
- send to ActiveMQ broker and continuously streamed via Flume to HDFS
- text file with one JSON encoded dictionary per trace
- ~5GB per day - 6M entries

There are both traces from DQ2 (historical) and Rucio (current). Traces for DQ2 and Rucio events have the same fields. The only difference that DQ2's traces are stores in a plain-text format, and Rucio's are in JSON.

For DQ2 traces description:

```
$ hadoop fs -cat /user/rucio01/dq2/traces/README.txt
```

Example of loading DQ2 traces data in a Pig script:

```
dq2_traces = LOAD '/user/rucio01/dq2/traces/2014-*' USING PigStorage() AS (uuid:chararray,
eventtype:chararray, eventversion:chararray, remotesite:chararray, localsite:chararray, timestart
timeend:chararray, duid:chararray, version:int, dataset:chararray, clientstate:chararray, protoco
filename:chararray, filesize:long, guid:chararray, timeentry:chararray, usr:chararray, relativest
transferstart:chararray, catstart:chararray, validatestart:chararray, hostname:chararray, ip:char
suspicious:boolean, appid:chararray, usrdn:chararray, rucio_account:chararray, rucio_appid:charar
errmsg:chararray);
```

Example of loading Rucio traces in a Pig script:

```
rucio_traces = LOAD '/user/rucio01/traces/*' USING JsonLoader('uuid:chararray,
eventtype:chararray, eventversion:chararray, remotesite:chararray, localsite:chararray, timestart
timeend:chararray, duid:chararray, version:int, dataset:chararray, clientstate:chararray, protoco
filename:chararray, filesize:long, guid:chararray, tracetimeentryunix:chararray, usr:chararray, r
transferstart:chararray, catstart:chararray, validatestart:chararray, hostname:chararray, ip:char
suspicious:boolean,appid:chararray, usrdn:chararray, rucio_account:chararray, rucio_appid:chararr
```

```
errmsg:chararray');
```

Important fields used from the traces records:

Field name	type	Description
dataset	chararray	dataset or container name
eventtype	chararray	type of an access event; we're interested only in *'get.*'*
usrdn	chararray	DN of user's certificate, used to group events in query
remotesite	chararray	remote site name
localsite	chararray	local site name
tracetimeentryunix (Rucio) / timeentry (DQ2)	chararray	event timestamp
uuid	chararray	job UUID

Oracle Dumps

Contain:

- daily reports for operations / site admins for consistency checks
- file replicas / unique files per storage endpoint
- primary / custodial dataset replicas
- number of replicas per dataset / last access times

Import and sizes:

- daily Sqoop dumps of most important tables to HDFS
- bz2 compressed, tab-separated text files, ~16GB compressed size
    - DIDs: 550.000.00 entries
    - Rules: 7.500.000 entries
    - Replicas: 690.000.000 entries
    - Dataset Locks: 8.000.000 entries
    - RSEs: 700 entries

DIDs and Dataset locks are linked with both *scope* and *name*.

Below there are examples of loading Rucio dumps data from a Pig script. *$CURRENT_DAY* is a parameter having format 'YYYY-MM-DD' to point to the last dumps.

DIDs

```
dids = LOAD '/user/rucio01/dumps/$CURRENT_DAY/dids' USING PigStorage('\t') AS (
  scope: chararray,
  name: chararray,
  account: chararray,
  did_type: chararray,
  hidden: chararray,
  is_open: chararray,
  complete: chararray,
  obsolete: chararray,
  bytes: long,
  length: long,
  events: long,
  project: chararray,
  datatype: chararray,
  run_number: chararray,
```

```
  stream_name: chararray,
  prod_step: chararray,
  version: chararray,
  task_id: chararray,
  panda_id: chararray,
  campaign: chararray,
  lumiblocknr: chararray,
  provenance: chararray,
  phys_group: chararray,
  transient: chararray
);
```

Dataset locks

```
dslocks = LOAD '/user/rucio01/dumps/$CURRENT_DAY/dslocks' USING PigStorage('\t') AS (
  scope: chararray,
  name: chararray,
  rule_id: chararray,
  rse_id: chararray,
  account: chararray,
  state: chararray,
  updated_at: chararray,
  created_at: chararray,
  length: long,
  bytes: long,
  accessed_at: chararray
);
```

RSEs (Rucio Storage Elements)

```
rses = LOAD '/user/rucio01/dumps/$CURRENT_DAY/rses' USING PigStorage('\t') AS (
  id: chararray,
  rse: chararray,
  rse_type: chararray,
  deterministic: int,
  volatile: int
);
```

Rules

```
rules = LOAD '/user/rucio01/dumps/$CURRENT_DAY/rules' USING PigStorage('\t') AS (
  id: chararray,
  subscription_id: chararray,
  account: chararray,
  scope: chararray,
  name: chararray,
  did_type: chararray,
  state: chararray,
  rse_expression: chararray,
  copies: int,
  expires_at: chararray,
  weight: chararray,
  locked: int,
  grouping: chararray,
  error: chararray,
  updated_at: chararray,
  created_at: chararray,
  locks_ok_cnt: int,
  locks_replicating_cnt: int,
  locks_stuck_cnt: int,
  source_replica_expression: chararray,
  activity: chararray,
  notification: chararray,
  stuck_at: chararray
```

```
);
```

Distinguishing *primary* and *secondary* data

- rule expires_at IS NULL and rule locked THEN custodial
- rule expires_at IS NULL and rule not locked THEN primary
- rule expires_at IS NOT NULL THEN secondary
- replica expires_at IS NOT NULL THEN tobedeleted
    - the *tobedeleted* will be changed to *secondary*

Table with column descriptions.

Column name	type	Description
A1	B2	C2
A3	B3	C3

Analysis Code

Where is it?

Table with all the scripts and short descriptions on what they do.

name	type	Description
FilePopularity.pig	pig	Calculates how many times a file has been accessed during...
A3	B3	C3

Table with all the UDFs and their descriptions.

UDF	return type	input parameters	Description
A1	B2	C2	
A3	B3	C3	

Additional information

Talks:

- daily dump data
- more details

Major updates:
-- IlijaVukotic - 2014-11-19
-- SergeyBelov - 2015-02-25

Responsible: IlijaVukotic
Last reviewed by: **Never reviewed**

This topic: AtlasComputing > RucioHadoop
Topic revision: r8 - 2018-06-11 - IlijaVukotic