

PILOT 2 UPDATE

Paul Nilsson

Pilot 2 – Recent developments

- ZIPMAP support
 - Done – for now only updating job parameters (filtering) so should be compatible with Archive tf (to be tested)
- Event service updates
 - AES workflow tested (successfully) by W. Guan (February)
- Full implementation of trace report (i.e. for all copy tools)
- Retriable errors for user jobs
- Transfer time-outs in copytools
- Recent and relevant Pilot 1 changes migrated into Pilot 2 (double checked since May)
- Many new error codes implemented (total: currently at 90 error codes)

Pilot 2 – Examples of new error codes

- Pilot 2 error codes listed here: <https://twiki.cern.ch/twiki/bin/view/PanDA/Pilot2ErrorCodes>

| | | |
|------|-------------------------|--|
| 1313 | NOSTORAGEPROTOCOL | No protocol defined for storage endpoint |
| 1315 | UNKNOWNTRFFAILURE | Unknown TRF failure |
| 1316 | RUCIOSERVICEUNAVAILABLE | Rucio: Service unavailable |
| 1318 | COMMUNICATIONFAILURE | Failed to communicate with server |
| 1320 | LOGFILECREATIONFAILURE | Failed during creation of log file |
| 1321 | RUCIOLOCATIONFAILED | Failed to get client location for Rucio |
| 1322 | RUCIOLISTREPLICASFAILED | Failed to get replicas from Rucio |
| 1323 | UNKNOWNCOPYTOOL | Unknown copy tool |
| 1324 | SERVICENOTAVAILABLE | Service not available at the moment |
| 1325 | SINGULARITYNOTINSTALLED | Singularity: not installed |
| 1327 | UNREACHABLENETWORK | Unable to stage-in file since network is unreachable |
| 1328 | PAYLOADSIGSEGV | SIGSEGV: Invalid memory reference or a segmentation fault |
| 1329 | NONDETERMINISTICDDM | Failed to construct SURL for non-deterministic ddm (update AGIS) |
| 1330 | JSONRETRIEVALTIMEOUT | JSON retrieval timed out |
| 1331 | MISSINGINPUTFILE | Input file is missing in storage element |
| 1332 | BLACKHOLE | Black hole detected in file system (consult Pilot log) |

34 new error types
introduced
(many old removed)

Pilot 2 – Missing pieces / known issues

- Event service updates
 - A few (4) error codes still need to be implemented in the event service code
 - Prefetcher support
- Benchmarking support
 - Will be part of Services API, easily pulled in and used by e.g. new benchmarking transform
 - API now in development
- mv stage-out broken
 - Internal Pilot issue to be looked at asap
- Extension of Pilot workflows (next slide)

Pilot 2 Workflows

- A Pilot “workflow” is an execution path for a given pilot task
- Selected using pilot option(s)
 - E.g. `-w generic_hpc --hpc-resource titan`

| Workflow | Explanation | Status |
|-------------------|---|------------------|
| Grid | Standard/default workflow used on grid sites | Done |
| HPC | Pilot code used on HPCs. Plugins needed for each resource (Titan in production) | Done/in progress |
| Input file stager | Pilot only stages in data (ATLASPANDA-477) | Concept |
| Payload+stage-out | Pilot executes payload and stages out data (ATLASPANDA-477) | Concept |
| Benchmark | Pilot only executes benchmark suite | Concept |

HPC Pilot – Status and Plans

- Developer info with technical details regarding HPC Pilot workflow and plugins
 - Available at <https://github.com/PanDAWMS/pilot2/wiki/HPC-workflow>
- HPC resource plugins
 - Plugin for Titan already developed and used in production since last year (D. Oleynik)
 - Skeleton Pilot 2 plugins exist for ALCF, BNL, NERSC, Summit
- Access will be granted to Summit soon
 - Harvester + Pilot 2 will be setup asap to test e.g. ML payloads
- Yoda/Droid integration to be discussed this (with D. Benjamin)
 - Expect to begin with using Pilot 2 defined error codes in Yoda/Droid asap
- Planning to submit a CHEP 2019 abstract about HPC Pilot

Pilot Source and Version

- Idea: use existing AGIS field 'pilotversion' to select Pilot version for given PQ
 - Will facilitate Pilot 2 commissioning
- Harvester will check this field before downloading the Pilot source tarball
- Status: being implemented in Harvester
- Note: Default field value should be Pilot1
 - Other allowed value: Pilot2

Pilot 2 – Commissioning on the grid (1/2)

- Successful tests on Nordugrid by D. Cameron
 - Two identified issues (long time for pilot to end, JSON from AGIS instead of CVMFS), now resolved
- “Accidental” commissioning on couple of IT queues a few weeks ago
 - Pilot 2 was launched my mistake by wrapper for ~a day
 - Minor issue with Pilot 2 was identified and fixed
 - Bigger issue was later realized by EventIndex team
 - Number of events not read correctly from jobReport -> missing metadata in DB
- Running production tasks on CERN-EXTENSION_HARVESTER_1/2 since last week
 - UCORE queues, rucio copy tool, Pilot 1 running on CERN-EXTENSION_HARVESTER_1, Pilot 2 running on CERN-EXTENSION_HARVESTER_2
 - Help provided by F. Berghaus (will be involved in pilot migration)
 - Minor issues with Pilot 2 identified and fixed

Pilot 2 – Commissioning on the grid (2/2)

- Plan:
 1. Run for a few more days on CERN-EXTENSION_HARVESTER_1/2 with different tasks/job types
 - Compare jobs run by Pilot 1 and 2
 - Task plan setup by Ivan Glushkov (sim, merge, rec, AOD merge, der, NTUP merge)
 - Currently looking at some new DDM related failures
 2. Try a few more production queues for a few days, followed by migration
 3. Migrate most production queues in cloud to use Pilot 2
 - Satellite queues; don't bother with using _1/2 clone queues
 4. Repeat for remaining clouds
- Migrate user analysis queues, gradually
- Avoid queues with non-supported copy tools
 - Some 25 queues still use lcgcp (storm/dCache as well?)