L&B and JP usage in job statistics tools

CESNET JRA1 team

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But . . . the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But ... the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But . . . the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But ... the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But ... the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Current situation – our view

- We are proud L&B is used for this purpose
- We are willing to support this further
- But ... the current way is not ideal
- L&B database structure is not public L&B interface
  - We reserve the right to change it with new L&B features (breaking backward compatibility)
- Job state computation is rather complex
  - Naïve L&B event processing may give different results
- OK for now but future time-bomb

Foreseen solution

- Use Job Provenance service, in the way it is designed to be used
  - permanent, well known Primary storage (JPPS) servers
  - L&B and WMS push data into JPPS
  - volatile Index servers (JPIS), created on-demand
  - JPIS are populated with one-time query or continously
- Configure customised index server(s)
  - appropriate sets of queryable and index attributes
  - created and destroyed as required
    (eg. to compute statistics for given month)
  - populated with query for given time interval
  - can register for receiving further updates
- Potential complication
  - L&B uploads information into JP "post-mortem", some time after the job finished

Foreseen solution

- Use Job Provenance service, in the way it is designed to be used
  - permanent, well known Primary storage (JPPS) servers
  - L&B and WMS push data into JPPS
  - volatile Index servers (JPIS), created on-demand
  - JPIS are populated with one-time query or continously
- Configure customised index server(s)
  - appropriate sets of queryable and index attributes
  - created and destroyed as required
    (eg. to compute statistics for given month)
  - populated with query for given time interval
  - can register for receiving further updates
- Potential complication
  - L&B uploads information into JP "post-mortem", some time after the job finished

Foreseen solution

- Use Job Provenance service, in the way it is designed to be used
  - permanent, well known Primary storage (JPPS) servers
  - L&B and WMS push data into JPPS
  - volatile Index servers (JPIS), created on-demand
  - JPIS are populated with one-time query or continously
- Configure customised index server(s)
  - appropriate sets of queryable and index attributes
  - created and destroyed as required
    (eg. to compute statistics for given month)
  - populated with query for given time interval
  - can register for receiving further updates
- Potential complication
  - L&B uploads information into JP "post-mortem", some time after the job finished

Example – custom JPIS setup

Indicators required

- Timing and failure rate of specific application jobs
- Assuming jobs marked in L&B with "AppName = MyApp" tag

JPIS population

- the same query to all (few) relevant JPPS
- query conditions
  - submitted between Jan 1 and March 31, 2005
  - tag AppName = MyApp
- retrieved attributes
  - submitted, matched, started, and done timestamps
  - terminal state
  - job owner
- JPPS's sweep through all jobs in the time interval select only those having the tag, and feed the JPIS

Example (cont'd)

Some possible queries this JPIS can handle

- How long it took to find a matching CE?
- How many jobs failed in January?
- Number of jobs per individual user

Queries this JPIS cannot handle

- Number of jobs per VO
- What is the most frequent failure reason?

Tentative answers to JP questions

- *Do job records contain all information we need?*
  - Yes, as long as the info was logged to L&B.
- *How to manage right accesses?*
  - Trusted-party approach, ie. JPPS knows identities of those who may create index servers
  - Do we need fine-grain access control, ie. JPPS may give any attribute but sensitive JDL to this JPIS?
- *How many primary servers will we need to query? One per RB?*
  - Depends on configuration. We assume one per VO but it should not be enforced

Tentative answers to JP questions

- *Do job records contain all information we need?*
  - Yes, as long as the info was logged to L&B.
- *How to manage right accesses?*
  - Trusted-party approach, ie. JPPS knows identities of those who may create index servers
  - Do we need fine-grain access control, ie. JPPS may give any attribute but sensitive JDL to this JPIS?
- *How many primary servers will we need to query? One per RB?*
  - Depends on configuration. We assume one per VO but it should not be enforced

Tentative answers to JP questions

- *Do job records contain all information we need?*
  - Yes, as long as the info was logged to L&B.
- *How to manage right accesses?*
  - Trusted-party approach, ie. JPPS knows identities of those who may create index servers
  - Do we need fine-grain access control, ie. JPPS may give any attribute but sensitive JDL to this JPIS?
- *How many primary servers will we need to query? One per RB?*
  - Depends on configuration. We assume one per VO but it should not be enforced

Tentative answers . . . (2)

- *Are our criteria queryable attributes?*
  Is the query interface usable for our type of queries?
  - Hopefully yes, if the index server is set up correctly
  - What exactly are the queries?
    (We can still augment JP interface)
- *Would we have to use notifications?*
  - No, notifications are designed to deliver infrequent single job state changes.
- *Would we need to copy all data or would it be possible to query on demand to compute our statistics?*
  - The relevant portion of data is copied from JPPS to JPIS with the initial population.
    JPIS can be queried on-demand, it is pre-configured to handle particular query types effectively.

Tentative answers . . . (2)

- *Are our criteria queryable attributes?*
  Is the query interface usable for our type of queries?
  - Hopefully yes, if the index server is set up correctly
  - What exactly are the queries?
    (We can still augment JP interface)
- *Would we have to use notifications?*
  - No, notifications are designed to deliver infrequent single job state changes.
- *Would we need to copy all data or would it be possible to query on demand to compute our statistics?*
  - The relevant portion of data is copied from JPPS to JPIS with the initial population.
    JPIS can be queried on-demand, it is pre-configured to handle particular query types effectively.

Tentative answers . . . (2)

- *Are our criteria queryable attributes?*
  Is the query interface usable for our type of queries?
  - Hopefully yes, if the index server is set up correctly
  - What exactly are the queries?
    (We can still augment JP interface)
- *Would we have to use notifications?*
  - No, notifications are designed to deliver infrequent single job state changes.
- *Would we need to copy all data or would it be possible to query on demand to compute our statistics?*
  - The relevant portion of data is copied from JPPS to JPIS with the initial population.
    JPIS can be queried on-demand, it is pre-configured to handle particular query types effectively.

Other issues

- *Several jobs characteristics are not in the L&B databases*
  - *CPU time* – yes, if it is logged into L&B by CE provider (SA1 collaboration needed)
  - *abort reason when re-tries* – should be logged by WMS anyway, what is the exact info required?
- *Sara has required a filter to avoid to give us non EGEE data*
  - Doable with advanced JPPS query-based ACL (not in 2.0)
- *LAL only accept to give us statistics*
  - JPIS of their own, statistics inside (not in 2.0)

Other issues

- *Several jobs characteristics are not in the L&B databases*
  - *CPU time* – yes, if it is logged into L&B by CE provider (SA1 collaboration needed)
  - *abort reason when re-tries* – should be logged by WMS anyway, what is the exact info required?
- *Sara has required a filter to avoid to give us non EGEE data*
  - Doable with advanced JPPS query-based ACL (not in 2.0)
- *LAL only accept to give us statistics*
  - JPIS of their own, statistics inside (not in 2.0)

Other issues

- *Several jobs characteristics are not in the L&B databases*
  - *CPU time* – yes, if it is logged into L&B by CE provider (SA1 collaboration needed)
  - *abort reason when re-tries* – should be logged by WMS anyway, what is the exact info required?
- *Sara has required a filter to avoid to give us non EGEE data*
  - Doable with advanced JPPS query-based ACL (not in 2.0)
- *LAL only accept to give us statistics*
  - JPIS of their own, statistics inside (not in 2.0)

Time line

- JP WS interface available now (DJRA1.2)
- JP primary storage almost finished (internal deadline Apr 30)
- index server configuration guidelines, needed to negotiate exact JRA2 requirements (during May)
- whole JP planned for gLite release 2,
  early prototype expected in July