

## Data Persistence Project Group (Inaugural) Meeting – 8/28/18

**Attendees:** Jim, Trevor, Chandra (Dell), Michael (LF), Andre (Redis), Bruce, Ed (IoTech), Fede (Cavium), Jeron, Markus (ObjectBox). Attendees that may have joined after the start of the meeting may not have been captured and listed.

Discussion and action items as a result of meeting in **RED**

### Old Business

- Requirements finalized / ratified (see previous meeting notes for details)
- Evaluation
  - Current list of candidates
    - MongoDB (use as our baseline)
    - CouchDB
    - Couchbase Lite
    - ObjectBox
    - Key/Value DB
      - Redis (but support more data structures)
      - BoltDB (can run on a gateway; written in Go; lightweight)
    - At least on SQL (may be more useful for data like Metadata)
      - Postgres
      - MySQL
      - SQLite
  - Database comparison matrix
    - See <https://wiki.edgexfoundry.org/download/attachments/18907176/DatabaseComparisonMatrix.xlsx?version=1&modificationDate=1534994948000&api=v2>
  - Database performance evaluation testing
    - See <https://wiki.edgexfoundry.org/download/attachments/18907176/EdgeX%20Data%20Persistence%20Performance%20Testing.pdf?version=1&modificationDate=1535128848000&api=v2>
    - **We would need more benchmark numbers before starting**
    - **Still a good idea to have EdgeX benchmarks – just to be able to evaluate in the future**
  - Discussion
    - What can we glean from the matrix? What types of products should we really be exploring? Embedded?
    - What type of testing is realistic? How much work do we want to do? Can we use an existing benchmark?

### New Business

- Alternate plan for consideration
- **The Edinburgh abstraction seems really far off – could this be accel**
  - **Maybe an unofficial release like we did with Go preview**
  - **Need benchmarks sooner as well**

- For Edinburgh release, the core dev team (that would be Trevor's) works to implement a layer/abstraction system for databases. Allowing any database to be used under the covers more easily.
  - In order to have abstraction, we need to get rid of BSON references (or any reference to DB specific identifier, qualifier, etc.)
- For Edinburgh release, we implement 2 options for the EdgeX reference implementation: a MongoDB and a Redis set of core services (Core Data and Metadata). (Trevor's team)
- For Edinburgh release, we will document how to replace the database for others (Trevor's team).
- For Edinburgh release, we have a certification process for core services – insuring any new core service adheres to the API sets and work appropriately in an EdgeX environment (Jason's team with help from our teams to take care of the technical work)
- For Edinburgh release, we have a marketplace for core services (Jason's team and LF marketing)
- For Edinburgh release, we have a performance test harness for core services that provide statistics on memory usage, CPU usage, speed of queries, writes, etc. (Andy's team)
  - Could impact system management work
  - Deployment footprint measures could involve CI/CD
- For the Fuji release, we make a determination of which of the reference implementations we want to keep long term (may be one or both or even take one out of the marketplace if it is shown to be better)
- No objections to taking this plan to the core working group and others for reaction and possible approach to this effort. Jim will work with Trevor to get it on the upcoming agenda.