Core Working Group Meeting Notes (11-Apr-2019)

Attendees:

Participants (17)

- Trevor Conn (Host, me)
- Akram Ahmad (Dell)
- André Srinivasan
- Ankur Gargi
- Anthony Bonafide
- Dan Bryant
- ijohnson
- Jim Wang
- Jim White
- Lenny Goodell (Intel)
- Michael Hall (LF)
- Mike
- Rahim
- Rebekah
- Steve Osselton
- Toby Mosby (Intel)
- tonyespy

Old Business

- Call for Fuji F2F topics *(Reminder)*
  - Currently proposed
    - High Availability
      - 12 factor apps
      - Service Orchestration
        - Registration/deregistration
        - When and when not to use SMA
    - Instrumentation
      - We have correlation ID tracking in our logs, but what's next?
      - Reporting capabilities as determined by deployment footprint
    - Architecture Principles
      - Explicit discussion of architectural guidelines to move toward
      - Assembly of resources to support guidelines and educate new team members
    - Persistence implementations
      - Decouple from mono-repo? How to do this?
    - Store and forward capabilities??
      - Collecting events at the edge, connectivity is lost. When restored, push everything stored in the meantime.
- Correlation logging is giving Tony heartburn
  - Should we reorder TRACE and DEBUG and make all correlation logging statements TRACE?
  - Ideally these should be on all the time but the TRACE level would simulate a separate sink for these messages that could be turned on/off.
  - *Correlation logging statements will be logged at TRACE level*
    - *These will not be logged by default*
    - *Folks prefer changing log level for edge deployments in order to not flood the log.*
    - *Edge vs Cloud – Should this be configurable?* *Not right now.*
  - *Re-order TRACE and DEBUG ?? Create our own?*
    - *Correlation logging done at TRACE level*

- *Evaluate and change TBD*
- *Pay attention to where INFO could be logged as DEBUG*
  - *Ping health check logging for example*

New Business

- Core-command PUTs
  - OperatingState / AdminState
    - Why? They're just pass-through to metadata
    - Issues with model validation
    - *Add issues to convert these endpoints to correct REST PUT requests.*
  - Get Commands by DeviceID/Name
    - For read-only operations like this, any reason why Core-command can't just go directly to the database rather than through metadata client?
    - *Discussed, but need to defer until post Edinburgh. Could lead to a wider systemic role for Core-Command. Is there a role for caching here?*
    - *How do performance statistics affect the decision to do this?*
  - IssueCommand doesn't propagate correlation-id currently (issue #1219)
    - Anthony Bonafide is working on this (PR #1240)
- Review of DeviceProfile → ProfileName
  - Core-contracts
    - Issue #27 / PR #63
  - Device-SDK-Go
    - PR #238 (Unbuildable currently)
    - No interlock issue created
  - Edgex-go
    - PR #1236 (Unbuildable currently)
    - No interlock issue created
  - *PRs have been closed. Diana to send email to device-service and core email lists with feature branches. Review will begin there and decision for inclusion in Edinburgh will be made based on subsequent progress.*
    - *Possibly PR against fork's master*
- Invocations of core-command from northside

- In scope for CBOR work?
  - Sounds like it
- *Core-command Request*
  - *Accept (JSON, CBOR)*
- *Response (content-type one of the above)*

Question on extent of core-command documentation. We need to audit and address gaps.