Core Working Group Meeting Notes (30-Jan-2020)

Participants (19)	

Type to filter...

TC	Trevor Conn (me)
JW	Jim White
LG	Lenny Goodell (Intel)
	Michael Estrin
a	adamwmit
	Akram Ahmad (Dell)
AS	André Srinivasan	Redis Labs
AB	Anthony Bonafide
	Brandon Forster
BM	Brian McGinn (Intel)
DR	Dave Rensberger
EU	Eno Udoko
JP	Jeremy Phelps
JW	Jim Wang
J	John Luo
MJ	Mike Johanson
RH	Rodney Hess
Tz	Tingyu zeng
	Wolt (Intel)

1.) API V2 Q&A Following from Last Week's Session
   a. Reminder to review https://github.com/edgexfoundry/edgex-go/pull/2285
   b. *Suggestion to leave the PR open into next week b/c of Chinese New Year*

   i. Cloud has done the PR review already, made some comments
   ii. Ah, OK. We're good then.
  c. Ask for merge is based on continuing follow-on workstream
   i. Suggest little risk from the fact that the content of the PR is largely greenfield and isn't mixed with the old implementation
   ii. Lenny said a couple other resources within Intel have looked at it with no adverse comments
   iii. Would like implementation to be somewhat complete before creating an ADR capturing implementation decisions.
   iv. Can we leave open through Monday?
    1. Yes, OK.
    2. Can be left open beyond but recognize that delays the follow-on work.

2.) Bootstrap integration
  a. Add param for name of config file
  b. Filename itself is hardcoded, we already have a cmd line flag for path. Future effort to consolidate these into one?
   i. Consolidation into one param for eventual V2 release
  c. General agreement, yep.

3.) Persistence of data models V1 vs V2
  a. Shared or separate?
   i. Consensus is for "separate"
  b. If separate, what does that entail?
   i. Redis, for example, use namespacing – prefix the data structures with "v2"
   ii. Discussion of contract for credential initialization whereby each service can abstract the knowledge of DB provider and version
    1. Parallel with elimination of config-seed, services to self-seed configuration
    2. Possibly related https://github.com/edgexfoundry/edgex-go/issues/1424
    3. For use case implementation, take same approach whereby persistence V1/V2 are done exclusive to one another. Ideally when V2 is released we just snip the V1 code base. Also leverages V1 as stable.

4.) System-wide service metrics
  a. Following from request from Jim White
   i. "I'd like to raise a question on what metrics would core and device services like to see collected (and offered) via system management for Hanoi and later." (14-Jan)
  b. Here's my list. Are there any others we should consider?
   i. # of API requests/sec
   ii. Avg response time
   iii. Service uptime
   iv. Request success vs. failure vs. invalid (2xx vs 5xx vs 4xx)
   v. Avg request size
   vi. Max request size

vii. Auth/auth failures
   1. Once we get there