

Monthly Architect's Meeting, September 20, 2021, 10am PDT

Attendance:



Some attendees may have joined after the call started when this record was captured.

High Priority

- **CORs support decision**
 - Don't set it up in Kong (remove the existing PR)
 - Bryon already had the configuration that needs to be moved over
 - Bryon to remove the work already done
 - Currently have it in edgex-go
 - Move it out into a module (bootstrap or contracts)
 - It should be backward compatible; low risk
 - Submitted PR for device SDK (go) -> look at that as base for code in module
 - Jim to create an issue in bootstrap
 - Impact to C side
 - Impact probably not huge
 - Done in the SDK
 - Iain to create any necessary issues to track
 - Configuration needed (additional elements may come up with implementation)
 - Domain
 - On/off option
 - Jim to reach out to GUI guys on this to see if there is any impact
- ADR reviews
 - Service List
 - Jim to get with Bryon and have him cover his ADR (<https://github.com/edgexfoundry/edgex-docs/pull/389>) in the next Security WG as this may make this ADR moot
 - **All EdgeX architects are asked to review the following** in advance of the next Monthly Architects meeting so this ADR can be dispositioned:
 - Service List ADR
 - Registry ADR (already approved)
 - The getAllServiceEndpoints query on the registry API to be able to query for all services
 - **Architects are requested to put comments in the service list ADR after review** of above

- At the next Monthly Architect's meeting (Oct 18), we'll make a decision as to whether the ADR is still needed, what to update, and create a more detailed design (including secure and non-secure) if it is still needed.
- metrics
 - Architects are in agreement that metrics should be something that is addressed in the Kamakura release (possibly CPE as well and it may be more important – an ADR for it is needed ASAP).
 - The ADR needs to be updated to include **(Jim White task)**
 - Recent changes per Ireland (state of SMA, OMQ, etc.)
 - Change the configuration around message bus to use the common message bus info (message protocol, host, port, etc.)
 - The go-metrics library needs to be explored **(Jim to do)**. This go library would probably be the means to implement this feature. Explore the library for:
 - How it manages the telemetry data (persistence, in memory, database, etc.)
 - Does it offer a query API (in order to easily support the ADR suggested REST API)
 - What does the go-metrics package do so that its features can become requirements for C side
 - Can the metric/telemetry count be reset if needed? Does this happen whenever it posts to the message bus? How would this work for REST?
 - Are metrics collected on a timeframe (“timed”) and can the timeframe be reset? Ex: collect the number of requests on the API and have a way to reset that count back to zero? Are any of the metrics “static” – meaning they are never reset?
 - Global configuration is ruled out (address this in the consequences). However, a global configuration ADR is needed for EdgeX 3.0. **Jim to create the issue.**
 - This ADR to be re-reviewed at the next Monthly Architect's meeting.

The rest of these topics were not covered in this meeting and are deferred until the next meeting

- UoM
- Declarative Kong applicability
 - Allowing us to drop Postgres DB
 - But can you configure groups/users ACL
 - Only support JWT users
 - Lose some functionalities we need: it would make the API Gateway read-only, and users can't create new users and tokens dynamically.

Medium Priority

- Med - Is the Wiki the best place to document project(low level architecture or design) decisions (those outside of or smaller than ADRs). This was our initial take. Should we revisit? (Per Jakarta planning meeting)
 - Jim to research other projects and how they handle this

- Completed
 - K8s: has a architecture SIG, meetings recorded and put in Google Docs
 - Often create PR for issues and debate in the PR
 - <https://github.com/kubernetes/community/blob/master/sig-architecture/README.md>
 - TensorFlow – uses process of RFC (request for Comment) that get submitted by PR to rfcs directory in Github
https://www.tensorflow.org/community/contribute/rfc_process
 - React Native – repository and PR process (but doesn't appear to be very active) <https://github.com/react-native-community/discussions-and-proposals>
 - Most use Github: Visual Studio Code, Google Flutter, Apache Spark, Kubernetes, Pandas, TensorFlow, PyTorch, node.js, vue.js.
 - ADR template is well used but where people keep these is not consistent and at what level an ADR is required is not consistent
- 3rd party lib review process
 - Per planning meeting; relook policies for 3rd party libs and how to evaluate, especially for device services. Also, should we start to fork these “immature” libraries

Low Priority

- Per Core WG meeting of 4/22/21 and Jakarta Planning Meeting (6/21) - how much info should be put in errors and log messages.
 - There is a concern about putting too much information in from those familiar with commercial products
 - Temporary decision was to provide enough so that someone could debug the problem with the information provided and not to be concerned with exposing intellectual property since everything is open source.
 - Per planning meeting: need a survey to address but don't believe we have an issue
 - **Potentially is an issue in defining what is Error/Info/Warning**
- Revisit combine core services at least at all executables in one image
 - Release would be easier but image would be bigger with more complex compose files
 - Per Core WG of 2/18/21 - is it at least worth exploring the combination of Core Metadata and Core Command since the two have to share so much data?
 - Core command is just a proxy service today, but reasons for having a separate service include: additional security to protect actuation; issue multiple device commands with one request (make one request and fire it to all Modbus devices or all devices under the control of one service); provide the means to limit requests down to a device so as not to overwhelm it or wake it up). These needs could also be incorporated into a combined metadata service but there are advantages to separation of concerns.
- Time series database support and applicability
 - Ian Johnson has an example of app service to InfluxDB export (snap in the store)
- Where should tool/script for creating new device and application services be placed?
 - After the architect's meeting of Jan 26, 2021, it was decided that "templates" should be created in all SDKs to allow for the easy creation of new services (removing the old samples in the SDKs). The templates will be a means for users to copy and create a new

service with some instruction on how to rename and replace TODOs with necessary code.

- After the templates are in place, there is a decision to be made about where automation can be placed to use the templates to create new services (versus a manual copy). In the CLI, in a new tool, in a set of simple scripts?

Tabled for now

- V2 API - should we add security foundation added to that (per some of earlier V2 API designs via Dell and Bryon N)?
 - Adding token to authenticate a micro service call (is this in scope for Ireland)
 - May not be needed unless all services are distributed
 - We need to explore alternatives to provide secure / locked out service to service communications
 - **ADR being created and to be reviewed in the Security WG.**
- Digital twin (and LWM2M) applicability
 - **Being worked via liaison with DTC**