Accenture Responses

- What types of services would you want to see ==certified== if EdgeX started a certification program? Would Accenture be a potential contributor or user of certified EgdeX services?

==AJH> Do you mean services in terms of App services/Device services, or at a higher level ? I assume a== ==higher level== – as in services that add value to the monitoring/provisioning/fault finding of deployed IOT gateways and associated IOT device networks?

If the latter, I see there being the following areas of interest (this is my first pass – I will let the questions percolate further, and add more details with inputs from Laura)

We are working to make our Edge Gateway capable of working not only with Accenture's Adaptive cloud (formerly called AIP+), but also with GCP, Azure and AWS cloud platforms. As such, we are re-engineering the cloud tools that we have for registering, authenticating, provisioning, monitoring and debugging our deployed estate of IOT gateways, plus the tools for creating the downloadable apps/analytics capabilities.

I think we are quite comfortable with the engineering that we need to do on the Edge to make this happen – we are migrating to use the App Services for Northbound connectivity, and have Azure IOTHub working well (with some limitations on receiving MQTT comms from the cloud that we are working around)

The major engineering parts are within the cloud- where we want to provide a generic set of tools that work with the Adaptive Cloud infrastructure, but which can also map onto the underlying raw cloud platform functionality so that the Edge can be used without necessarily having to license the Adaptive Cloud platform if it is not required. Abstracting this, we can see that it is relatively easy to get the basic device registration, authentication, security certs, and software load served to the gateway. ==The tricky bit comes with the operational monitoring of both the hardware/software== ==health, and the health of the services being provided by the deployed equipment. We would like to== ==have generic cloud based portals for both of these aspects, and are investigating what is available in== ==the market, and what comes with each of the cloud platforms.== I think the trickiest aspect will be the ability to remotely debug the iOT gateways and IOT networks/devices attached. There are some pre-packed solutions such as Azure and Azure Sphere which we may have to interwork with, but to have an effective solution for all cloud platforms, we need to steer a good, value add line where we can advocate cloud agnostic implementations at the edge to capitalise on   the benefits of the EdgeX Foundry implementation, and extend this to allow us to debug faults in the IOT network and with the more intelligent IOT devices attached.

This is an Area that is of great interest to us if we can benefit from community thinking and collaborate with IOTech or the community in a beneficial way

- Given the Ireland release will be version 2 (a major release), ==would you look to EdgeX to provide== ==example tools/scripts needed for EdgeX 1 to EdgeX 2 conversion== (for things like configuration, databases and profile migration)?

AJH> ==Yes== - having ready made templates is a really good starting point. We are discussing some further abstraction so that the templates are hidden from non Edge development engineers so that

the Edge can be provisioned by Data Scientists - engineers who do not understand the actual physical device and its constraints. We are thinking that we can map their model requirements onto gateway resources and flows via an abstraction tool so that we never expose then to the internals of the gateway. Is this something you are thinking of ?

- Does Accenture have needs / requirements to incorporate a <mark>timeseries database</mark> in EdgeX? Should this replace or augment existing persistence?

AJH> <mark>Yes</mark> – we intend to include our Chronicle TSDB into the device services layer of EdgeXpert 9we have some issues to iron out with our ARM 64 version). We don't see a conflict with the Redis DB at the moment – in our implementation, Chronicle and Redis perform different roles. <mark>Redis will always be present, and Chronicle is provisionable depending on the use case.</mark>

- Do you have <mark>additional language support needs</mark>? Is having EdgeX (documentation, examples, etc) in English prohibiting or slowing adoption from your perspective?

AJH> <mark>We use English</mark> as the generic language within Accenture, so we have not encountered any issues.

- EdgeX has never <mark>removed Docker images from our Docker Hub</mark> account since the inception of the project. While the world is free to use any of the old Docker images, the EdgeX community's support policy (as laid out here) makes clear we do offer any support changes or updates to images beyond the current release or any LTS release. What is your expectation about how long Docker images should be made available?

AJH> we pull the IOTech images when we are ready to integrate them. After that, they are contained within our internal build and test system. We have an open discussion item with Keith about how long images will be maintained , but <mark>in general we will aim to be a maximum of two, and realistically one release behind the latest major IOTech release.</mark> This is of course dependent on feature content. If there is no benefit from Accenture moving to the latest release, we will be unlikely to migrate quickly. The associated question relates to vulnerabilities that IOTech/ EdgeX Foundry have detected and release updates for. This forces an integration of new code, and so the issue is that if the vulnerability is found on the latest current release, will the fix be pushed out to the previous two releases ?

o We would only pull the latest images or any LTS release image and so this does not affect us.

o We would expect EdgeX Docker container images to be available for _____ years and understand the limits of support for old images

o We would expect EdgeX Docker container images to be available indefinitely and understand the limits of support for old images

o (Provide an answer in your own words) Our expectation is: _____.

- Today, sensor data is picked up by device services and put into an event/reading, forward to core services for local persistence, and then processed through application services to export the data (or make it available to local analytics). However, due to the asynchronous nature of the EdgeX micro services, there is no guarantee that messages flowing from south to north (that is from device services to application services and beyond) will arrive in any order. In other words, there is no order assumed

with regard to the arrival and processing of events and associated reading in EdgeX micro services. Would you ever have a need for the order that event messages are processed by EdgeX micro services be preserved (from device service collection to export outside of EdgeX)?

o Yes - We would like the option to preserve event/reading message order through EdgeX and understand that could result in slower processing times

o Yes - We would like the option to preserve event/reading message order and this should be the default means of dealing with sensor data – even at the expense of slower processing times

o No - We understand the nature of asynchronous micro services and messaging systems and understand order is not guaranteed. This works for our use case(s).

(Provide an answer in your own words) Our expectation is:

AJH> this is dependent on use case. Where time stamping/ordering is important, we have the option of using the Chronicle TSDB. In other use cases using e.g. video analytics (such a s computer vision) we will time stamp events when the analytics process has detected them.   I will discuss this point further with Kari and Laura. _____.

- Regarding your device data simulator – could you provide more information on what you have done? Would this be something you would be willing to contribute? Is the "playback" mechanism the most important part of your simulator and can you provide some scenarios of use? How is the playback stored?

AJH> yes – the Device Data Simulator sits within the device services layer and presents data to the higher level entities as if it has popped out of one of the device services associated with an IOT protocol stack.   The DDS can generate random data, or read data from a 'test file' in a number of formats including jpeg and video. We are in the process of wrapping the provisioning of the DDS to make it more friendly for engineers . the intention is that our cloud tool suite for creating and downloading models to the edge also has the ability to associate simulated data streams with the models so that they can be created downloaded in a user friendly manner. Currently, the implementation is suitable for Edge engineers who understand the mechanics of the model and DDS. In its current form, DDS requires a .yaml file to define the simulated device profile data formats and identities, a .csv file that contains the simulated data characteristics. DDS also has a random value generator that can be used to continuously inject data for long duration runs.

- With regard to your losing data from battery powered devices (as they wakeup and report) – can you provide more details on how data is lost and can you provide thoughts on how you think EdgeX can help in this area? Are EdgeX store/forward capability and device service hooks a means to deal with these losses?

AJH> I interpret your question as the battery powered devices being e.g. zigbee sensors that trigger on a state change (door open) .   In this case, the battery powered device will be in deep sleep mode to conserve power until the state change is indicated. When a state change happens (door close to door open), the sensor activates, sends the state change, and then goes back to deep sleep…. Issues are if there is network congestion/interference, or the gateway is experiencing more iOT traffic than it can process, the notification may be lost, so the state change not detected. Possible solution (adopted on

a previous project is to have the concept of prioritized device data flows so that even if the gw rx port is overwhelmed with normal messages, high priority ones will always be processed)

Battery powered devices that go into deep sleep mode between state change notifications are always problematic. They appear to disappear from the network because they are not 'pingable' . However (at least with ZigBee and Z-wave) they have a configurable 'wake up' interval at which point they notify the gw that they are still functioning, and give health indications such as battery state.

A REAL issue that we have and will have with IoT deployments is that it is very difficult to debug faults in the network or with the connected devices. This is an area that we can brainstorm/collaborate on. Debugging requirements will differ depending on device types and network types. Mesh networks are a pain because nodes act as transceivers, and depending on how the network is deployed, you can very quickly get into non optimal network configs because new devices are added and/or existing devices are moved. On my last project, we had to introduce a network healing mechanism so that we could re-configure Z-Wave mesh to its optimal setting if network performance started to decline.

We also had to introduce the capability to force a disconnect of devices that had not been paired correctly before a re-pairing attempt was made. Initially, we had scenarios where a pairing attempt failed part way through, and we were left with Zombie devices in the network

Other factors that we should, consider are the provisioning/monitoring of intelligent IOT devices. This might be a good intersect with XRT because as we have more sophisticated automation, we may have intelligent satellite nodes connecting to our gateway that require monitoring and provisioning also. I am really interested in the concept for using XRT MCUs as satellite nodes to further distribute intelligent processing in the IOT domain

ThunderSoft responses

We also had a few questions (below) that came up as a result of reviewing your presentation. We were hoping you could help answer based on ThunderSoft's use of EdgeX. The answers to these questions will be critical to our scoping efforts that will be part of our Ireland planning meeting in November.

·   What types of services would you want to see certified if EdgeX started a certification program? Would ThunderSoft be a contributor or user of certified EgdeX services?

below are certification services demanded:

 1. sensor/protocol certification - south-bound device/sensor are ported and tested on EdgeX framework

 2. edge device certification - edge devices are to be certified as "Powered by EdgeX"

 3. software module certification - it's like application/micro-service/AI store (TBD)

 4. cloud certification - north-board cloud (public & cloud) certification

ThunderSoft can act as both a certification service provider as well as a user.

·   Given the Ireland release will be version 2 (a major release), would you look to EdgeX to provide example tools/scripts needed for EdgeX 1 to EdgeX 2 conversion (for things like configuration, databases and profile migration)?

yes.

·   Does ThunderSoft have needs / requirements to incorporate a timeseries database in EdgeX? Should this replace or augment existing persistence?

yes, we are engaging with some timeseries DB partner on this already. this will be presented as an alternative to the existing persistence.

·   Do you have additional language support needs? Is having EdgeX (documentation, examples, etc) in English prohibiting or slowing adoption from your perspective?

English is fine.

·   EdgeX has never removed Docker images from our Docker Hub account since the inception of the project. While the world is free to use any of the old Docker images, the EdgeX community's support policy (as laid out here) makes clear we do offer any support changes or updates to images beyond the current release or any LTS release. What is your expectation about how long Docker images should be made available?

   o   We would only pull the latest images or any LTS release image and so this does not affect us.

   o   We would expect EdgeX Docker container images to be available for _____ years and understand the limits of support for old images

   o   We would expect EdgeX Docker container images to be available indefinitely and understand the limits of support for old images

o   (Provide an answer in your own words) Our expectation is: _____.

<mark>We would expect EdgeX Docker container images to be available indefinitely and understand the limits of support for old images</mark>

·   Today, sensor data is picked up by device services and put into an event/reading, forward to core services for local persistence, and then processed through application services to export the data (or make it available to local analytics). However, due to the asynchronous nature of the EdgeX micro services, there is no guarantee that messages flowing from south to north (that is from device services to application services and beyond) will arrive in any order. In other words, there is no order assumed with regard to the arrival and processing of events and associated reading in EdgeX micro services. Would you ever have a <mark>need for the order that event messages are processed</mark> by EdgeX micro services be preserved (from device service collection to export outside of EdgeX)?

o   Yes - We would like the option to preserve event/reading message order through EdgeX and understand that could result in slower processing times

o   Yes - We would like the option to preserve event/reading message order and this should be the default means of dealing with sensor data – even at the expense of slower processing times

o   No - We understand the nature of asynchronous micro services and messaging systems and understand order is not guaranteed. This works for our use case(s).

o   (Provide an answer in your own words) Our expectation is: _____.

<mark>Yes - We would like the option to preserve event/reading message order through EdgeX and understand that could result in slower processing times</mark>

·   With regard to EdgeX providing <mark>more data caching</mark>, can you help us understand this request? What data is being cached and when? What is the need for caching – where are the current performance bottlenecks? How/when would you expect the cache to be reset?

<mark>not sure if data caching is something must-have for EdgeX. Here are some common feedback about caching, for example, the old configuration will still be remained after recreating the profile with the same name.</mark>

At present, EdgeX device data is stored in the database in the core layer, and will be cleared by daily tasks as default. Users can increase the cache by simply modifying the scheduler and event. it may not be an economical way to cache a large amount of data on the edge side. The better way is to forward the data through the data bus.

·   You have created your own <mark>device virtualization</mark>. Can you explain more details behind this? What could EdgeX do to offer assistance in this area?

On the one hand, our device virtualization is built on top of EdgeX to enable easier and more distributed communication between different devices, more automated to register devices in EdgeX and connect with device services. On the other hand，we virtualize the device of another edge device into the local device, for example, we use v4l2lookback to virtualize the cameras of other hosts as the local cameras. we will have a more detailed introduction to the coming Ireland planning.

Jiangxing Intelligence Responses

·   What types of services would you want to see ==certified== if EdgeX started a certification program? Would J.I. be a contributor or user of certified EgdeX services?

Currently, we would like to see the ==cetification of device services==. For industrial users, a cetified device service means high reliability used in an application. J.I. would like to be both a contributor and a user of certificated EdgeX services.

·   Given the Ireland release will be version 2 (a major release), would you look to EdgeX to provide ==example tools/scripts needed== for EdgeX 1 to EdgeX 2 conversion (for things like configuration, databases and profile migration)?

==Yes==, we are using version 1, if EdgeX provides tools or scripts to help users fit the conversion from EdgeX 1 to 2, that would be better.

·   Does J.I. have needs / requirements to incorporate a ==timeseries database== in EdgeX? Should this replace or augment existing persistence?

==Yes==, we use influxdb as our timeseries database, because the applications in our scenario need such a database to store the dynamics of battery, CPU, memory status. Does EdgeX have a plan to propose a new timeseries database?

·   Do you have ==additional language support needs==? Is having EdgeX (documentation, examples, etc) in English prohibiting or slowing adoption from your perspective?

 If there is Chinese language support, that ==will help lots of engineers in China== to read the documents and examples and use EdgeX in their projects.

·   EdgeX has never ==removed Docker images from our Docker Hub== account since the inception of the project. While the world is free to use any of the old Docker images, the EdgeX community's support policy (as laid out here) makes clear we do offer any support changes or updates to images beyond the current release or any LTS release. What is your expectation about how long Docker images should be made available?

Our expectation is: we would expect that EdgeX ==adopts Ubuntu-like version release and management styles==.

·   Today, sensor data is picked up by device services and put into an event/reading, forward to core services for local persistence, and then processed through application services to export the data (or make it available to local analytics). However, due to the asynchronous nature of the EdgeX micro services, there is no guarantee that messages flowing from south to north (that is from device services to application services and beyond) will arrive in any order. In other words, there is no order assumed with regard to the arrival and processing of events and associated reading in EdgeX micro services. Would you ever have a need for the order that event messages are processed by EdgeX micro services be preserved (from device service collection to export outside of EdgeX)?

o   ==No - We understand the nature of asynchronous micro services and messaging systems and understand order is not guaranteed. This works for our use case(s).==

· Would J.I. <mark>consider donating some of their application services</mark> that send data to Alibaba, Tencent, Baidu, etc. to our EdgeX examples? We currently lack examples for the Chinese clouds.

 <mark>Sure, we have the plan to share such services, e.g., to Baidu BIE, to EdgeX examples.</mark>

· Do you <mark>like how the Device Service SDK in C works</mark>? Would you prefer a C library to be delivered vs having to build from source?

Now, we <mark>use Golang</mark> to write device service. But for other programmers, we believe that both C library to be delivered and having to build from source are important.

· You have requested EdgeX provide "device registration optimization where currently a unique device ID is required per instance of Edge Box". Can you help provide more detail and clarity to this request?

Here, we would like to assign a unique ID to a sensor, .e.g., a camera, a PLC controller, or others. A unique ID can help device services and applications in northside to clarify the collected data and the status of these sensors. We therefore need some rules or interfaces provided by EdgeX to help users to detect the status of a sensor. If a sensor with a chip can be identified uniquely, we could add the codes in the corresponding device service to generate and assign a unique ID to it. If not, we also can write a proper device service to handle it.

Tibco Responses

What types of services would you want to see certified if EdgeX started a certification program? Would Tibco be a contributor or user of certified EgdeX services?

More than certification I would like to see a mechanism where all the libraries and underlying dependencies are locked for a certain version. It seems to me that with the current deployment method, a lot of the libraries are in flux and there is no way to predict what will happen with building against those libraries.
This is extremely important if the framework needs to be used in a production environment.

- Given the Ireland release will be version 2 (a major release), would you look to EdgeX to provide example tools/scripts needed for EdgeX 1 to EdgeX 2 conversion (for things like configuration, databases and profile migration)?

Again, for the framework to be considered seriously by the community and be used in real projects, propper updated documentation needs to be provided with each release. My experience so far has been that docs are always out of date and pointing to old references.
Tools and scripts to automate an upgrade will be also required.

- Does Tibco have needs / requirements to incorporate a timeseries database in EdgeX? Should this replace or augment existing persistence?

Currently we don't have that requirement but we are envisioning scenarios/clients that might ask for that.

- Do you have additional language support needs? Is having EdgeX (documentation, examples, etc) in English prohibiting or slowing adoption from your perspective?
No

- EdgeX has never removed Docker images from our Docker Hub account since the inception of the project. While the world is free to use any of the old Docker images, the EdgeX community's support policy (as laid out here) makes clear we do offer any support changes or updates to images beyond the current release or any LTS release. What is your expectation about how long Docker images should be made available?
    - We would only pull the latest images or any LTS release image and so this does not affect us.
    - We would expect EdgeX Docker container images to be available for __1_ years and understand the limits of support for old images
    - We would expect EdgeX Docker container images to be available indefinitely and understand the limits of support for old images
    - (Provide an answer in your own words) Our expectation is:
- Today, sensor data is picked up by device services and put into an event/reading, forward to core services for local persistence, and then processed through application services to export the data (or make it available to local analytics). However, due to the asynchronous nature of the EdgeX micro services, there is no guarantee that messages flowing from south to north (that is from device services to application services and beyond) will arrive in any order. In other words, there is no order assumed with regard to the arrival and processing of events and associated reading in EdgeX micro

services. Would you ever have a need for the order that event messages are processed by EdgeX micro services be preserved (from device service collection to export outside of EdgeX)?

- Yes - We would like the option to preserve event/reading message order through EdgeX and understand that could result in slower processing times
- Yes - We would like the option to preserve event/reading message order and this should be the default means of dealing with sensor data – even at the expense of slower processing times
- No - We understand the nature of asynchronous micro services and messaging systems and understand order is not guaranteed. This works for our use case(s).
- (Provide an answer in your own words) Our expectation is: Currently we don't have a requirement to maintain ordering for the events at the edge. We can handle ordering at the cloud as long as the timestamps are provided.