In partnership with:

VOLT ACTIVE DATA

STL PARTNERS

EDGE IOT PLATFORMS WEBINAR 1: VERTICALS AND USE CASES DRIVING REAL GROWTH IN 2023

Webinar: Questions and answers

EDGE IOT PLATFORMS WEBINAR 1: VERTICALS AND USE CASES DRIVING REAL GROWTH IN 2023

This document outlines the questions and answers received from the STL Partners and Volt Active Data webinar, 'Edge IoT platforms webinar 1: verticals and use cases driving real growth in 2023' which was hosted on Tuesday 21st February 2023.

You can watch the recording of the session, and also access the slides, using the link here. *We have included the following timestamps for the webinar recording:*

- **02:53** for the introduction to our presenters and panellists
    - **David Rolfe,** Head of Product Marketing, Volt Active Data
    - **Joe Hurman,** Senior Consultant, STL Partners
    - **Dalia Adib,** Director, STL Partners
- **03:30** for STL's presentation on "Verticals and use cases combining IoT and edge"
- **30:11** for Volt Active Data's presentation on "Automated guided vehicles"
- 38:24 for Q&A session

If you have any questions not addressed in the webinar or this Q&A document, or want to hear more about our research findings or from our speakers, please contact:

• David Rolfe, Head of Product Marketing, Volt Active Data, drolfe@voltactivedata.com

• Joe Hurman, Senior Consultant, STL Partners, joe.hurman@stlpartners.com

• Dalia Adib, Director, STL Partners, dalia.adib@stlpartners.com

Post-Webinar questions and answers

How close does the edge need to be hosted to the end premise?

Joe Hurman, STL Partners: This completely depends on the practicalities of the use case and the benefits that an enterprise is seeking to unlock by utilising an edge location. For example, when running video analytics for manufacturing, large amounts of data need to be collected and processed at a single location to ensure speed and reliability, and reduce data backhaul costs.

However, a network edge deployment may be more appealing – potentially to suit different regulations (e.g. in financial services where specific regulation dictates certain personal data may not leave the country), in situations where the low latencies offered by an on premise deployment are not required, and in scenarios where the data collection or processing is distributed and there is no centralised 'premise' (e.g. content delivery, several smart city use cases).

How do you see different verticals adopting different types of edge (device vs on-prem vs network edge)? Are there key factors leading/influencing this adoption?

JH: Much of this question is also answered by the question above. In addition, the popularity of different use cases, as well as broader network requirements, will drive different verticals to have a preference for device vs on-premises vs network edge computing.

Healthcare, for example, would likely benefit from on-premise edge computing due to the critical nature of their use cases, requiring high network reliability and the necessity of maintaining patient data with a high-level of security – best enabled by the data localisation offered by an on premise deployment.

In comparison, verticals that are more distributed in nature, such as smart cities, would likely benefit from network edge deployments which can serve devices across a geographically disparate platform.

With respect to verticals that would benefit from on-device edge computing, those with use cases that require ultra-low latency and reliability, such as autonomous vehicles, would likely require significant on-device edge compute to ensure that if network coverage was lost then they could still operate safely, even if this was combined with some network edge deployments.

Question for David. It would seem that many edge applications would need an orchestration platform. Does that imply AI is going to play a big role?

David Rolfe, Volt Active Data: That would depend on what you see AI being used for. I'm still not sold that AI will be a revolution. Maybe. Maybe not. I can see lots of use cases involving optimization and planning of activities - that's fairly clear cut. But given the very high failure rate technologies like GPT have in real world situations would you really put it in charge of say, a china shop?

What type of SLAs are relevant and viable for IoT applications across the disaggregated ecosystem of partners required in vertical applications?

JH: Downtime is always a key metric for any business-critical use cases, measuring the availability of a specific device or system of devices. Enterprises need to ensure that the touted reliability benefits of edge are being combined with reliable IoT devices to ensure this benefit is being realised in the deployment. In particular, any automation enabled IoT platforms may have strong dependencies across the business process they enable, assuming no workarounds are implemented to cover the edge case where a technical fault within an IoT device disables the associated business process. For these more tightly coupled platforms, availability is a key factor, with downtime the associated SLA used to monitor this.

In addition, the other element to overall reliability, other than the availability referenced previously, is latency. Many OT applications would require significantly lower latencies than IT (around 5ms versus 100ms respectively), so ensuring IoT devices are compatible with these requirements is also paramount, and tracking latency both across an automated workflow and within its components is a must.

How mature are the leading edge computing platforms? And which are those major platforms?

How are companies monitoring Edge platform / computing devices? What are the popular solutions current in the market?

JH: The edge market is still in its relative infancy, and therefore we are still seeing a fairly fragmented market with no clear market leaders across all verticals. The variety of use cases for edge mandate different hardware, middleware and software across different verticals and use cases so it's very difficult to give an answer on who the overall market leaders are in each of these areas. For more information on some of the companies who are driving innovation at the edge, STL's recent article on 100 edge computing companies to watch in 2023 goes into more detail on this topic.

Does Volt's platform use traditional on-premise edge computing hardware or something manufactured by Volt?

DR: We run on x86 Linux, or containers, on anything from part of a core to dozens. We have a Gravitron port coming out later this year. For various obscure architectural reasons we're really, really efficient.

What are the current challenges for mature edge computing platforms?

JH: I'm going to take this question from the point of view of a lead vendor of an edge platform integrating with IoT, as the challenges to enterprises will vary significantly across industry and from deployment to deployment:

There are several key challenges facing even the most mature edge and IoT platforms. The challenge which we've heard the most discussion of in the market currently is the management and orchestration of such a distributed system, in terms of variety of vendors, geographical location of nodes and the range of technologies/ devices required within the platform.

Managing a variety of vendors requires strong commercial and relationship management capabilities to ensure a vendor ecosystem is commercially architected in a way which best suits the lead vendor. Once this has been set up in a way which best facilitates the delivery of the platform, the orchestration capabilities needed to manage a plethora of heterogenous devices both across geographies and technologies is significant, and strong orchestration capabilities are key to the success of any edge platform for this reason.

In addition, the difficulty translating business requirements from the enterprise into technical requirements which an ecosystem of partners can collaborate to deliver against is another key challenge, let alone then having to translate these technical outcomes back into business outcomes to ensure the enterprise can measure the success (or otherwise) of deploying such a platform.

Would edge computing influence how IoT solutions are sold today?

JH: Most current IoT platforms which are sold as a packaged solution currently are closed systems, which use connectivity such as Bluetooth or WiFi to transfer data being collected across the platform (likely through sensors or video cameras) to a single point, before this data is then transferred to a compute-enabled location (either edge, cloud or legacy on prem) for processing. This single point takes the form of an IoT gateway, where data is aggregated and standardised, before being collected across the platform.

With the dawn of edge computing, the role of this traditionally simple IoT gateway can be expanded to include some analysis at the aggregation point, either removing the need to transfer data onto a further location for processing or reducing the quantity of data which needs to be transferred to this cloud or edge location, in turn reducing costs.

Vendors such as Dell are starting to combine the aggregation and standardisation capabilities of a traditional IoT gateway with the analytics capabilities of edge compute in an 'Edge Gateway for IoT', facilitating this analysis of collected data at the edge.

How important are regional edge data centres for businesses striving toward edge - what are the benefits you see for using edge data centres over larger hyperscalers?

JH: Using regional data centres as a form of 'regional edge' can produce similar benefits to a network edge deployment (albeit without the same latency and reliability guarantees associated with storing edge compute within an operator's network). However, the key blocker to network edge deployments is the low maturity of operators' network edge offerings. Regional datacentres exist today and have mature offerings which in some instances meet the flexibility requirements and cloud-like commercial models which allow them to meet STL's definition of edge.

Regional edge provided by data centres is also likely to be a more cost-effective solution than a network edge deployment, and for use cases such as data localisation for regulatory compliance (as discussed earlier in relation to financial services), a regional data centre is likely to be a more cost-effective solution that meets the requirements of the enterprise. Another consideration is the ability for a regional edge deployment to connect to all operators' networks, as opposed to being constrained to one operator's network with a network edge deployment.

There are various different reasons why a hyperscaler deployment may not be appropriate for a given enterprise. However, even if a hyperscaler deployment would be preferable, it is not always an option. The hyperscalers have strong datacentre density in certain regions, however they do still have significant gaps in their datacentre coverage. For example, AWS, Azure and GCP each only have one region in Africa (all in South Africa). Even in areas where it appears on first glance like they have strong local coverage; they may only offer certain services which reduce the appeal of a hyperscaler cloud deployment, but can be solved with ISV partnerships in a regional edge datacentre.

Can go a little deeper into Government? The public sector is enormous and has multiple segments that are very different, even after separating out defence

JH: The government sector is likely to vary significantly by country and could include IoT applications for smart cities, emergency services, public transport, education and of course, defence. Examples of IoT applications that could benefit from edge in each of these sub-sectors are broad and would rely on the varying capabilities offered by edge.

There are some use cases that could be leveraged across most of these sub-sectors e.g. smart buildings. Edge enabled IoT could be used in schools, hospitals and government office buildings to enhance security (e.g. using video analytics), health and safety (e.g. condition monitoring) or energy efficiency (e.g. using sensors and edge compute to analyse data with the objective of reducing heating in rooms with no occupants).

Additionally, many of these sub sectors have high security and reliability requirements which favour edge solutions. Defence, emergency services and public transport could all benefit from using edge to increase data locality and reduce the risk of connectivity outages caused by inconsistent network traffic and services.

However, for other IoT applications, edge requirements would vary widely. Network edge may be necessary for distributed services such as public transport and emergency services, whereas on-premises (or even on-device) edge may be better suited to sub-sectors with reduced geographical distribution such as hospitals and military bases, which may prefer the data locality and ultra-low latencies offered by on-premises deployments. These comparisons can be extended across most sub-sectors with differing requirements, as edge capabilities vary significantly according to edge location.

Is data latency really an issue, any longer? Unless we are talking about urgent, emergent situations (emergency management, combat, surgical procedures), is the latency issue not insignificant?

JH: Latency is definitely driving a disproportionate amount of the hype compared to the proportion of edge IoT demand we expect to be driven by this benefit. Use cases which leverage benefits such as security, cost savings and reduction of external dependencies (eg on connectivity) are on balance probably further down the line in terms of relative maturity. It's also worth noting that the majority of use cases which are expected to drive the demand for edge will benefit from more than just one of the touted benefits, so it's likely that low latency will be a secondary benefit even if security or reliability are the primary drivers of deploying an edge alongside IoT.

However, even considering use cases for which the leading benefit is low latency as not significant would be short-sighted in my opinion. If we take the use cases which fall under the bucket of smart cities, there are some more futuristic, yet realistic, use cases such as autonomous vehicles, which would rely on ultra-low latency to satisfy the safety requirements of such a vehicle. Where safety is the key driver as it is here, it is likely that both the public and private sector would be willing to pay a premium for the combination of reduced latency and increased reliability.

How does Volt support in this effort to reduce latency or facilitate this use cases?

DR: Volt is an unusual platform. Instead of a stack of different components, which adds latency and complexity, Volt combines logic, data storage and streaming into a single platform that scales on a core by core basis.

Again, geographical distance is not relevant any longer, at least not on our planet. However, if we are talking about Space Exploration, where infrastructure is limited, or deepsea…

JH: This is an interesting point – of course currently we're focusing on edge deployments on earth, but who knows where the edge market will be in 10 years' time! Latency would become a huge issue with any deployments in space so it's entirely possible this could be a potential long-term area of growth.

Can you share some use cases related with autonomous smart buildings

JH: Definitely, the following use cases can all be considered under the 'smart buildings' banner:

- Smart building management (utilities/ lighting)
- AGVs for cleaning/ concierge
- Advanced security solutions (video ingest and analytics)
- Smart energy management (expanding utilities to include new energy sources and vehicle to grid EV charging)

Get in touch:

If you have any questions intended for Volt Active Data or would like to learn more about their offerings, please contact:

David Rolfe, Head of Product Marketing, Volt Active Data, drolfe@voltactivedata.com