Executive Briefing Service

In co-operation with:

Hewlett Packard Enterprise

intel

Executive Briefing

EDGE COMPUTING: THE TELCO BUSINESS MODELS

Multi-access edge computing (MEC) has thus far focused on technology-driven use cases, but how could telcos create business opportunities in this quickly evolving ecosystem?

Jonas Strobel, Senior Consultant | jonas.strobel@stlpartners.com | October 2017

Table of Contents

Table of Figures

Introduction

The idea behind Multi-Access Edge Computing (MEC) is to push compute and storage capabilities out to the edge of the telco's network. By that, workloads and applications are closer to telco's customers, enhancing experiences and enabling new services and offers. As we have discussed in a recent report, there is much excitement within telcos around this concept:

- MEC enables a plethora of vertical and horizontal use cases (e.g. leveraging low-latency implying significant commercial opportunities. This is critical as the whole industry is trying to uncover new, exclusive sources of revenue.
- MEC fits in with telcos' 5G and SDN/NFV deployments, which will run certain parts of the virtualised network distributed at the edge. In turn, MEC is dependent on the capabilities of a virtualised network to orchestrate distributed compute.

Figure 1: Defining MEC

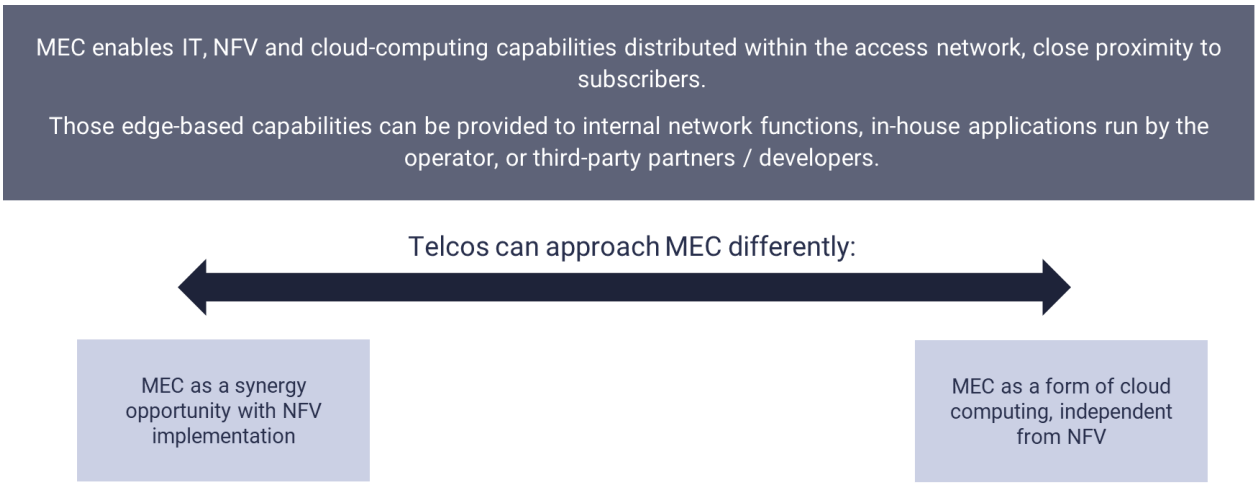

Source: STL Partners

However, despite the excitement around the potentially transformative impact of MEC on telcos, viable commercial models that leverage MEC are still unclear and undefined. As an added complication, a diverse ecosystem around the idea of edge computing is emerging – of which telcos' MEC is only one part.

From this, the following key questions emerge:

- Which business models will allow telcos to realise the various MEC use cases in a commercially viable way?
- What are the right MEC business models for which telco?
- What is needed for success? What are the challenges?

The emerging edge computing ecosystem

Telcos' MEC opportunity

MEC is being discussed in a variety of use case areas. These use cases on the one hand imply advancements for existing applications such as content/application delivery or caching by bringing these closer to the user in a geographically distributed way. On the other hand, MEC is also seen as the key enabler for emerging scenarios, such as:

- Industry 4.0 automation
- Connected/autonomous vehicle functions
- AR/VR gaming
- Drone control

Deployed through MEC, such use case scenarios can benefit from MEC-specific features such as low-latency, network awareness, or optimal allocation of resources. Therefore, MEC will yield benefits for both customers and the operator itself (see Figure 2 in detail).

Figure 2: MEC benefits

Benefit	Internal (operators) Cost saving and/or differentiation	External (end-customers and partners) New (paid-for) services
Low latency	✓	✓
Resiliency	✓	✓
Security		✓
Optimal allocation of resources *Workload or time shifting to optimise costs and performance (e.g. of applications)*		✓
Network (context) awareness *Being closer to VNFs allows telcos (and customers) to analyse network information in real-time to optimise content/applications etc.*	✓	✓
Data sovereignty/localisation		✓
Lower end-device power consumption		✓
Backhaul cost savings	✓	
Core congestion avoidance	✓	

Source: STL Partners

Hence, MEC offers the opportunity for telcos to:

- Improve their network operations to achieve efficiencies and cost savings

- Differentiate their own service offerings through MEC capabilities
- Enable others to make use of distributed compute capabilities

To reap these opportunities, telcos and vendors have been trying to coordinate and integrate their efforts, most notably through the ETSI Industry Specification Group for MEC (founded in 2014). While initially the group's work was focused on features like RAN-awareness in the context of *mobile* edge computing, there has been a reorientation to also encompass fixed-network edge computing (hence, *multi-access* edge computing). A first package of APIs has been released in July 2017 to facilitate interoperability of MEC deployments.

In another move towards interoperability, the ETSI group has recently signed a memorandum of understanding with the OpenFog Consortium, an organisation which promotes the similar concept of *fog computing* (a concept driven by several vendors, which is less focused than MEC on telcos' network access as part of the overall architecture). At first, both parties want to focus on standardising APIs to make it "easier for developers to create common architectures, unify management strategies, and write a single application software modules that run on both OpenFog and MEC architectures".

Hyperscale cloud providers are an added complication for telcos

Recently, it has become evident that the *edge* must not be understood as a concept that is purely related to the telco network. In particular, hyperscale cloud providers such as Amazon Web Services (AWS) or Microsoft Azure have been very clear on their ambition to complement their centralised cloud services with more distributed capabilities.

In a move to "edge-enable" its existing AWS IoT Platform, Amazon launched AWS Greengrass in late 2016. With AWS Greengrass, IoT applications can be run both in the centralised cloud and locally on an IoT device or gateway. AWS Greengrass utilises the same programming model as AWS' cloud, however the capabilities at the edge are restricted to more lightweight functions (such as "serverless" compute models through AWS Lambda). Large parts of an IoT application might still run in the cloud, while there is also application logic deployed locally which performs certain analytical or control functions. In such a model, roundtrip latency and bandwidth cost can be reduced by processing IoT data locally and not sending all information back to the cloud.

While Microsoft's Azure IoT Edge provides similar capabilities as AWS Greengrass, the company has even gone further recently: with the launch of Azure Stack, customers are now able to run the full Azure cloud environment not only in Microsoft's centralised cloud but also on-premise on their own bare metal. This enables enterprises to make use of hybrid cloud environments in which their own premises represents the edge with all the potential benefits that come with localised compute (low latency, compliance etc.).

Both Amazon and Microsoft's moves are facilitated by several trends which shift the economics in favour of more local and distributed compute and storage at the *customer edge* (e.g. gateways or

servers on-premise) or even the *device edge* (e.g. IoT devices performing parts of the overall application logic themselves). These trends include:

- **The proliferation of IoT applications** which create masses of data that do not need to be processed centrally.
- **Chip technologies** which make it possible to perform computational tasks economically on very small space – both for low-performance applications (e.g. through SoCs on mass IoT devices) as well as high-performance applications (e.g. through modern GPUs for AI-dependent use cases such as autonomous cars).
- **Advanced cloud platforms** which now facilitate distributed compute models through capabilities such as hybrid cloud computing (e.g. Azure Stack) or serverless computing (e.g. AWS Lambda). Figure 3 shows how Microsoft as a centralised cloud provider now incorporates the "intelligent" edge (which can be telecoms-agnostic) into their overall cloud vision statement.

Figure 3: Microsoft's new mantra – "Intelligent Cloud, Intelligent Edge"

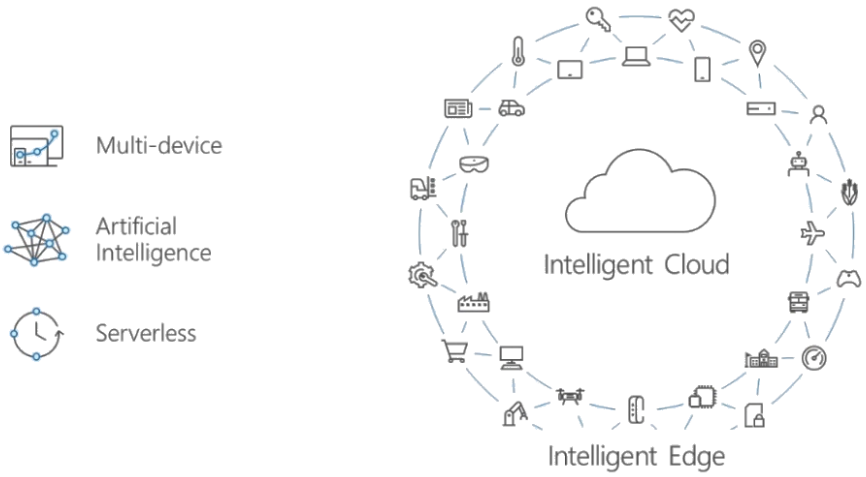

Source: Microsoft

How should telcos position themselves?

STL Partners expects that, by extending their cloud capabilities to the customer edge and/or the device edge, the hyperscale cloud providers will play a strong role in the emerging edge computing market. This represents a significant complication for telcos because these cloud providers will be able to deliver on certain MEC promises without necessarily relying on telco/network-specific distributed compute capabilities.

How should telcos position themselves in this situation? While telcos are rightly trying to come up with common standards and frameworks for MEC, we believe that telcos should not wait for these to mature before trying to tap into emerging commercial opportunities. With strong competition looming,

they should explore what could be viable business models to support the many potential MEC use cases. The next section will provide a starting point for exactly that and present 5 telco business model types for MEC.

5 telco business models for MEC

In the following section, we discuss five potential telco-centric business models for MEC, each of which can support a variety of use cases:

1. Dedicated edge hosting
2. Edge IaaS/PaaS/NaaS
3. Systems integration
4. B2B2x solutions
5. End-to-end consumer retail applications

These involve both business models in which the telco acts as an *enabler* of MEC services for others, as well as business models that imply an *end-to-end* role for telcos (see Figure 4).

Figure 4: STL Partners has identified 5 telco business models for MEC

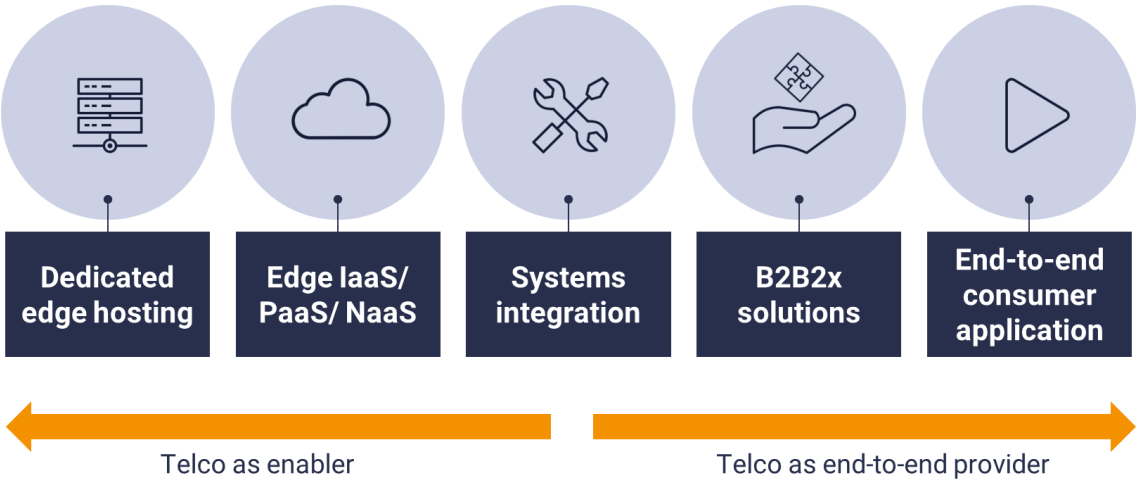

Source: STL Partners

For each business model, we will:

- Provide a detailed description
- Represent the value chain graphically
- Model a quantified, example 3-year scenario.

Business model 1: Dedicated edge hosting

In this first business model, the telco sits relatively low down the MEC value chain taking the role as an enabler of edge computing, rather than a provider. Here, the telco delivers and manages edge-located (bare-metal) compute/storage resources, which are pre-installed and connected to the telco network. The customer/partner would run its software, which could be for example a virtual content delivery network (CDN) or a distributed cloud stack, on top of the telco's edge-enabled dedicated hardware resources. From the partner's perspective, they would have the flexibility of being able to run their MEC-enabled software across multiple telecoms operators to achieve optimal coverage and be able to gain value from the use of network-aware APIs to enrich their offering.

As described in Figure 5, customers for this model could include public cloud (IaaS) providers, such as Microsoft Azure and AWS, looking to expand the coverage of their edge compute service. Others could be SaaS providers, local enterprises or governments and CDN providers. Similarly today, telcos typically offer colocation services to third-party CDN providers whose hardware is installed at the telco's (central) site; this means that the CDN's capabilities are integrated with the telco's network capabilities/data only to a limited degree. In dedicated edge hosting, running the CDN software on top of existing telco edge-based hardware will allow to also leverage data and insights from the network to further improve capabilities and performance of the CDN.

Figure 5: The dedicated edge hosting value chain

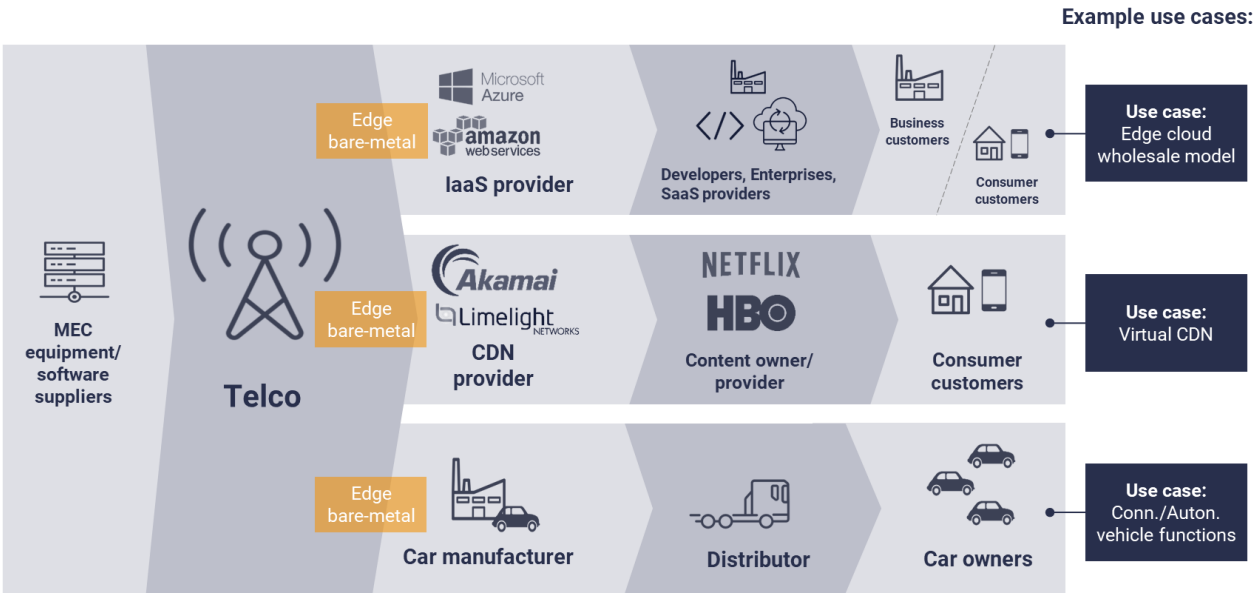

Source: STL Partners

For the operator, this business model is relatively low-risk; it relies on mostly success-based capex with minimum commitment. The telco would only invest in its infrastructure to provide this to partners if there was a customer established. Once set up, each order creates a revenue stream per site from set-up fees, hosting fees and additional capabilities (e.g. APIs, value-added maintenance and monitoring).

This is modelled in Figure 6 as a scenario in which the telco provides dedicated edge hosting to a public cloud provider. The results of this quantified example indicate a relatively high internal rate of return (IRR) after three years. However, the telco most likely would have to enter another investment cycle to replace some of the existing hardware; hence, the terminal value of the investment after three years is limited.

Figure 6: Quantified example – Dedicated edge hosting

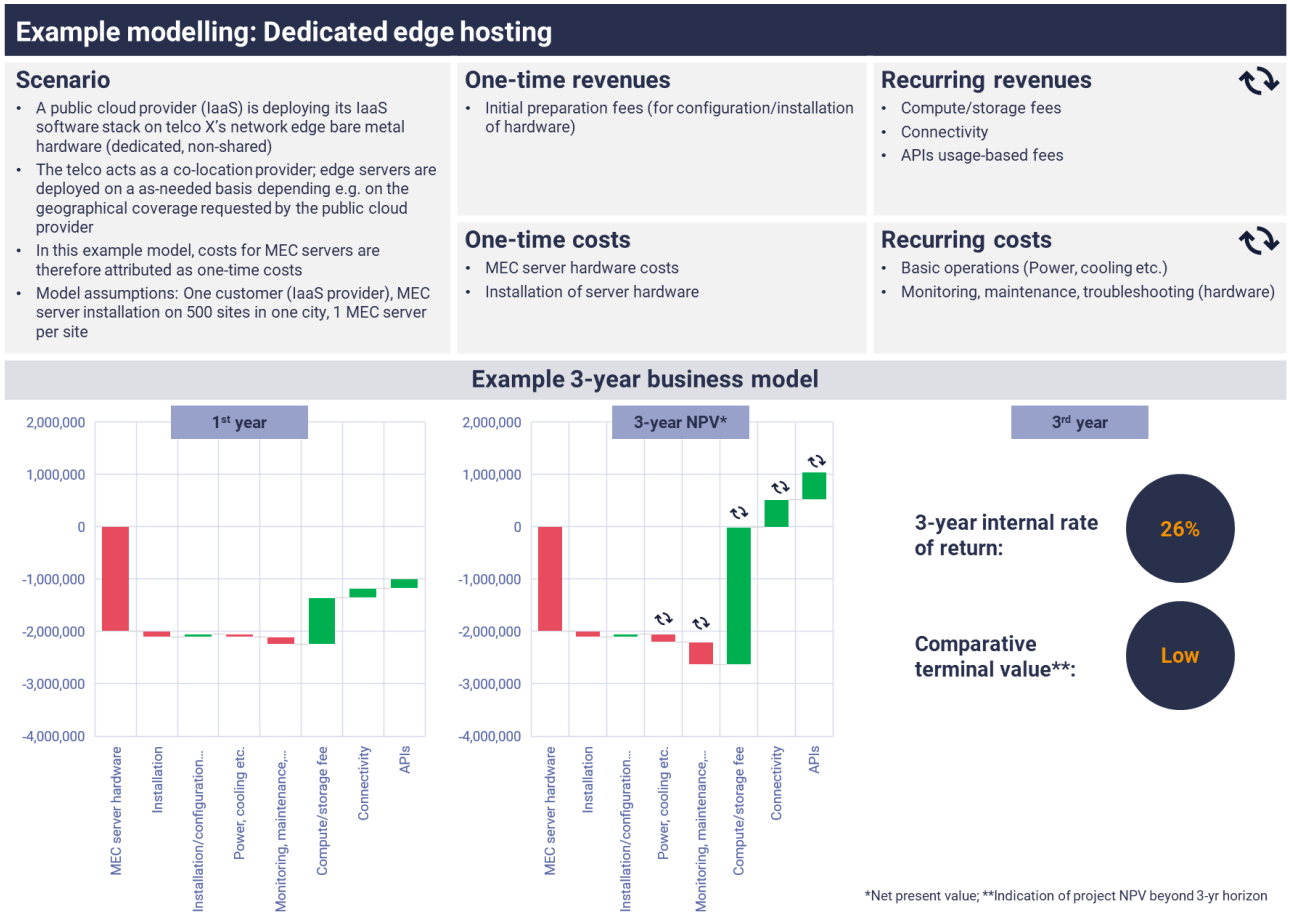

Source: STL Partners

Business model 2: Edge IaaS/PaaS/NaaS

The telco in this business model operates in a similar manner to a cloud provider, providing customers distributed compute and storage capabilities, a platform for developing applications on the edge infrastructure and network services, as well as APIs and virtual network functions (VNFs) in an 'as-a-service' manner through a cloud portal as the customer interface.

Customers would be those looking to deploy applications on the MEC infrastructure and make use of the benefits of the MEC platform capabilities (see Figure 7); for example, IoT application providers who want to optimise applications so that they can analyse the data from devices to trigger actions quickly. Other potential customers include start-ups, large enterprises, systems integrators, CDN providers, content owners and other cloud providers.

In practice, a customer would specify the location of the nodes across the telco's MEC infrastructure and capabilities required and pay for resources (e.g. virtual machines) according to use over time.

Figure 7: The Edge IaaS/PaaS/NaaS value chain

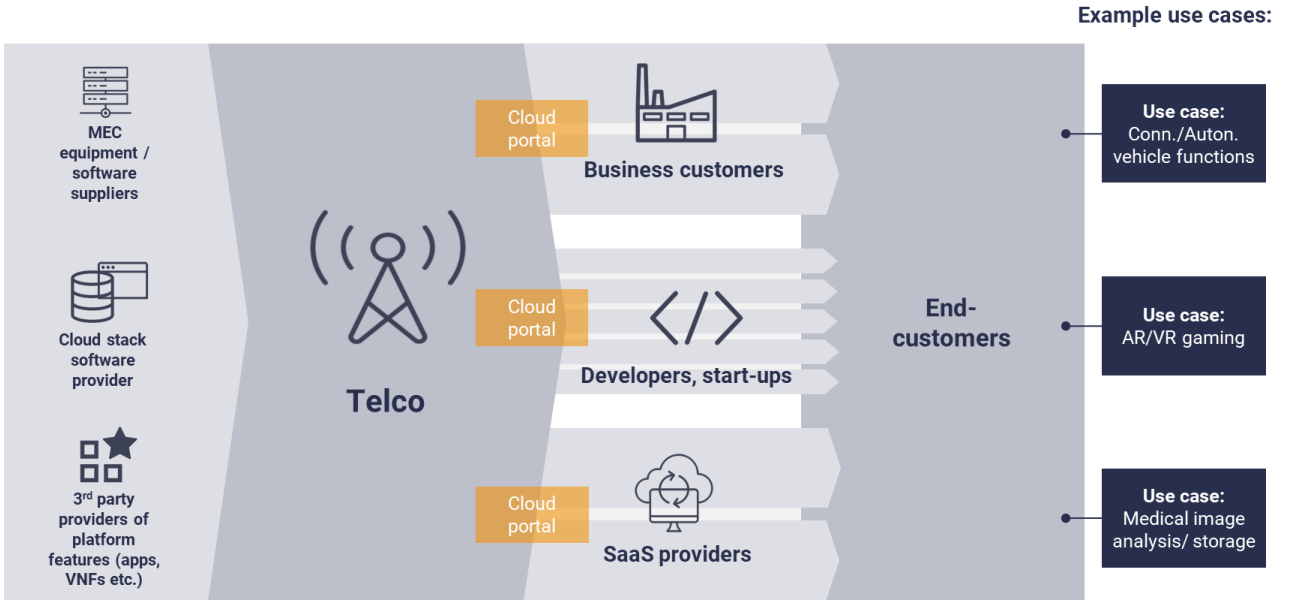

Source: STL Partners

Compared to the dedicated edge hosting business model above, there is higher risk here for the telco, as it needs to invest in MEC coverage and deploying the edge infrastructure (servers, site equipment, etc.) up front before a revenue stream is established.

In Figure 8, this is reflected in high one-time costs involved in rolling out MEC across large parts of the network, and to develop the service. Consequently, the IRR after 3 years is comparatively lower than for example in dedicated edge computing. However, this business model can potentially scale to much larger turnover and after three years the service would have a relatively high terminal value – assuming sufficient adoption among developers.

Figure 8: Quantified example – Edge IaaS/PaaS/NaaS

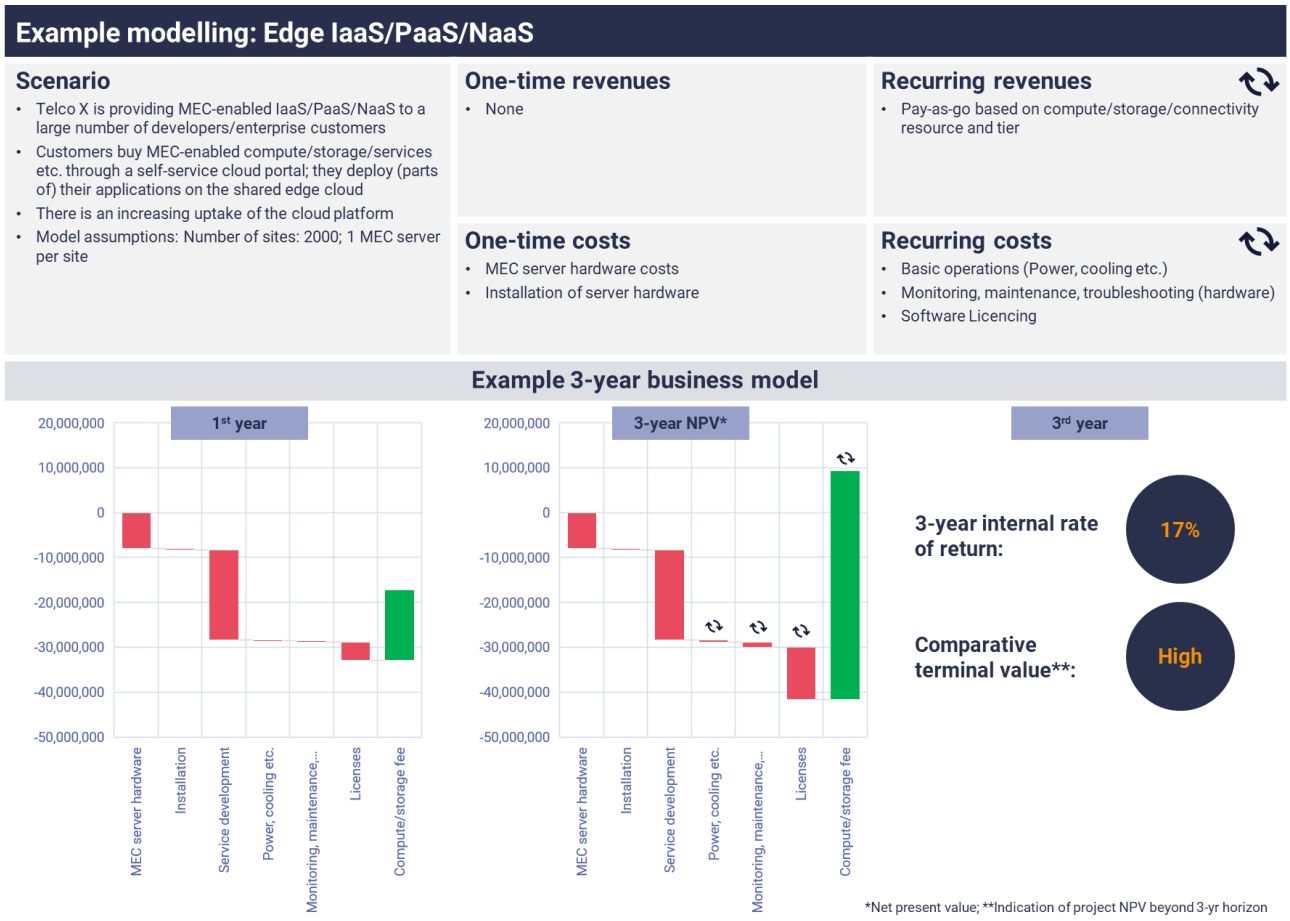

Source: STL Partners

Business model 3: Systems integration

The systems integration (SI) MEC business model will likely build on an existing SI business within a telco, offering custom turn-key solutions for enterprise customers with specific requirements, which are (partially) met by MEC functionality. As seen in Figure 9, MEC functionality is likely to be one of many components that are aggregated through the SI project; other components may include hardware and devices, other telco (not MEC-specific) capabilities, as well as third-party partner capabilities.

For example, a government or local council might invest in a MEC solution for their smart city project, which would require deployment of the MEC infrastructure and any necessary hardware (sensors, actuators and devices), integration of different networks and systems and orchestrating the development of the smart city solutions and applications. Many benefits of MEC are applicable for smart city use cases: security, network resilience and low latency (for real-time applications).

Figure 9: The SI value chain

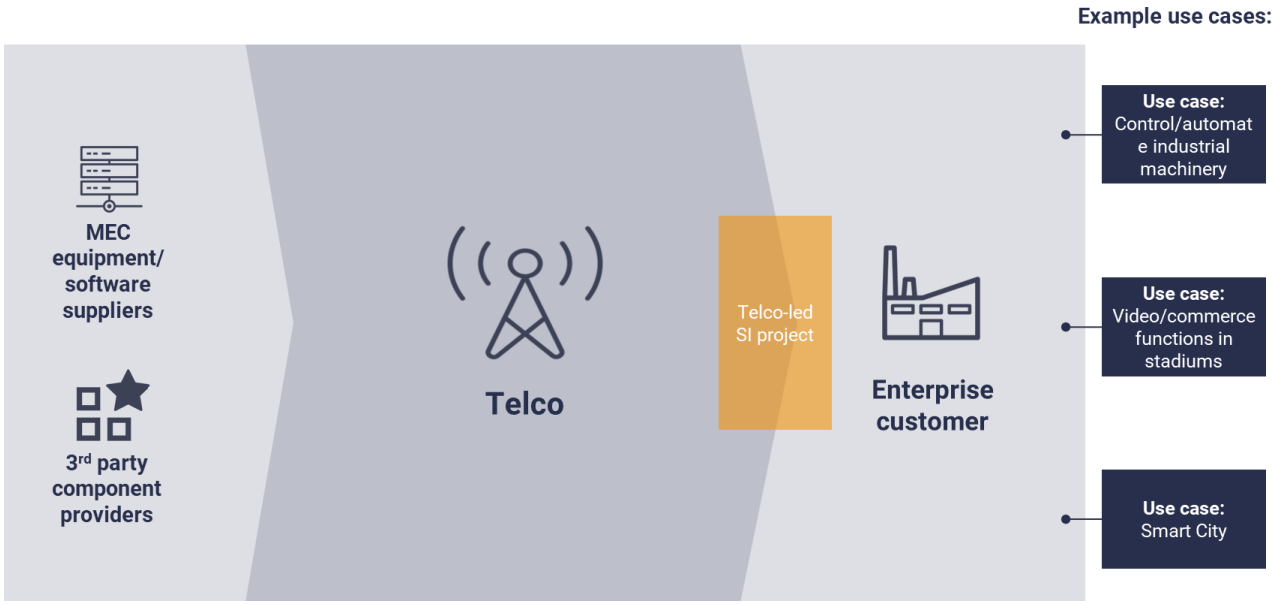

Source: STL Partners analysis

In terms of the revenue model from these potential customers (likely governments or enterprises), the bill looks much the same as a standard SI project – usually broken down according to the level of effort and resources required for each phase of the project (capturing requirements, plan/design, implementation, support etc.). Specific items the customer will pay for will vary by project, including the amount of MEC infrastructure that would need to be deployed.

From a telco perspective, this business model has some reduced risks, as the MEC infrastructure deployed may be explicitly for the customer (i.e. confined to the premises of the customer, for example a manufacturing plant), therefore it does not require significant investment from the telco without an indication of returns on investment.

As with many SI projects, partners are an important part of the offering – from MEC-specific hardware, equipment and software suppliers for servers, operating systems, platforms and applications, to specialist (vertical) project partners. The telco brings value as an SI to orchestrate this ecosystem of different players through their direct channel and provide a reliable, "turn-key" implementation of solutions enabled by MEC.

The quantified example in Figure 10 shows that the telco will recover the biggest part of its costs through the plan and build fees during the actual project. The profitability drivers are then fees for ongoing managed services and additional feature requests which had not been planned.

Figure 10: Quantified example – Systems integration

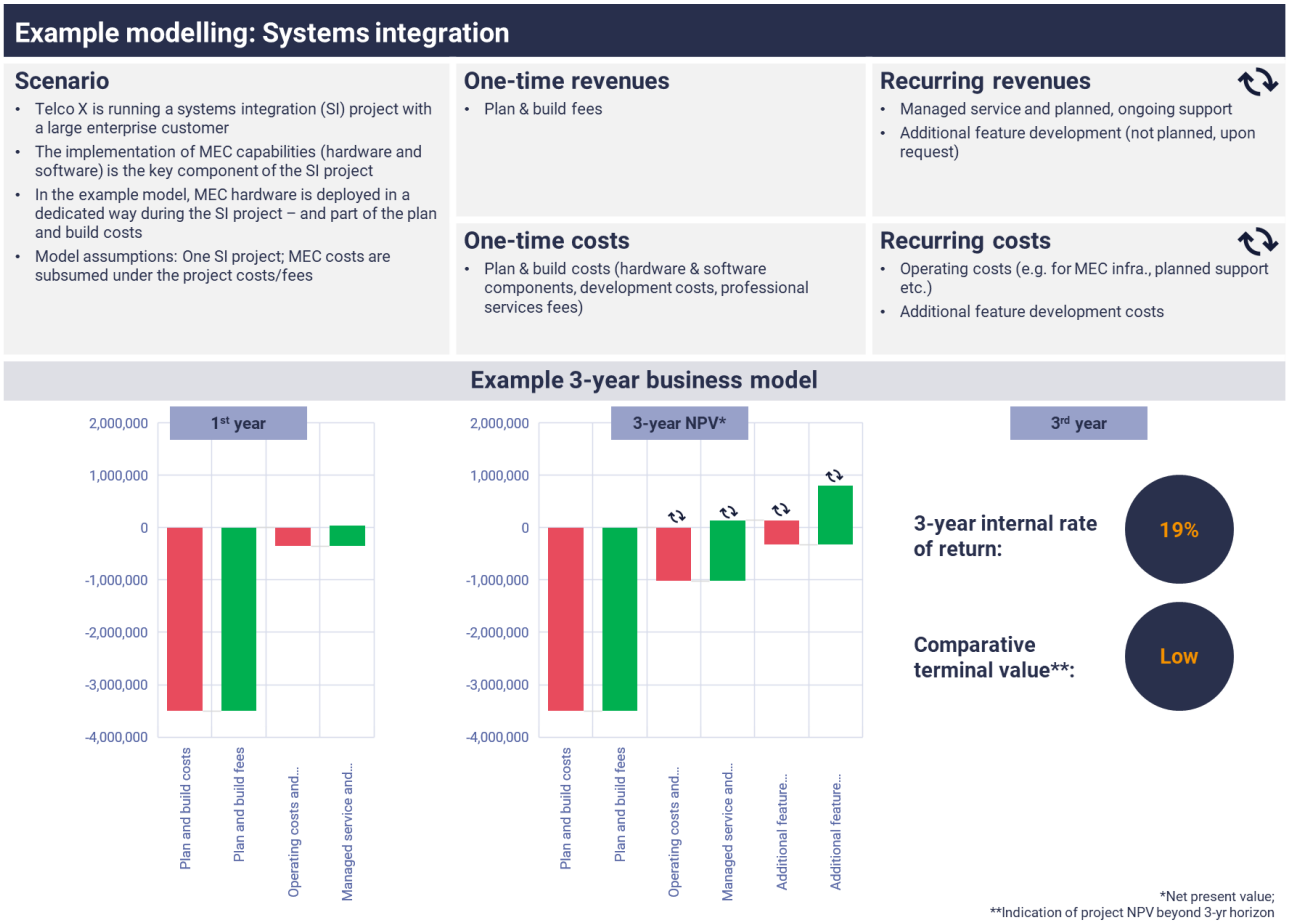

Source: STL Partners

Business model 4: B2B2X solutions

Telcos have an opportunity to create MEC-enabled business solutions for government, enterprise or SME customers. As with existing B2B solutions, these may be for the customer's internal purposes, such as to improve existing processes, or may contribute to an end-customer offering (B2B2X). In general, these solutions will be close to 'off-the-shelf' and will be less bespoke, thus requiring significantly less integration work, compared to SI projects.

A potential solution is indicated in Figure 11, whereby the operator can provide a service for large events with significant network congestion, such as sports stadiums or music concerts, to enhance the attendees' experience by using MEC to provide an immersive video experience, offering a variety of live video streams from different viewpoints. Event organisers can either monetise this and offer it as an additional paid service for attendees, or bundle it in with ticket prices to enhance the customer experience. Another example for this business model is CCTV video surveillance. Transmitting all the video feeds to the cloud would be uneconomical, however if the footage could be analysed at the edge, only events deemed as important would trigger a notification to the relevant emergency services and send the associated feeds to the cloud.

In order to provide these types of solutions, telcos would need to partner with solution technology vendors, to provide specific components and application, as well as MEC suppliers for hardware and software parts. In addition, the telco may not solely sell directly to its customers, but may also partner with specialised resellers who are vertically-aligned for certain enterprise solutions.

Figure 11: The B2B2X solutions value chain

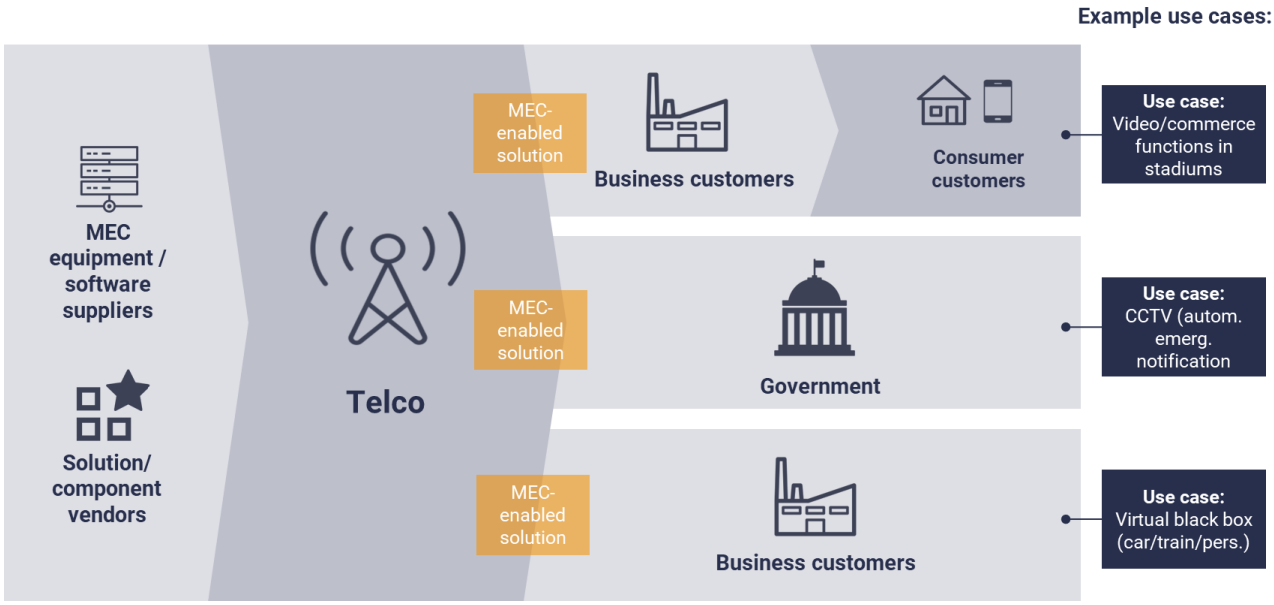

Source: STL Partners analysis

The risks associated with this business model are mainly due to the upfront service development and the uncertain take-up of the solution by enterprise customers. This is also modelled in the quantified

example in Figure 12 which shows that with sufficient adoption, a B2B2x solution (based on recurring fees) might yield significant return and a medium (and potentially high) terminal value as it is a service which can be offered beyond this 3-year horizon.

Figure 12: Quantified example – B2B2x solutions

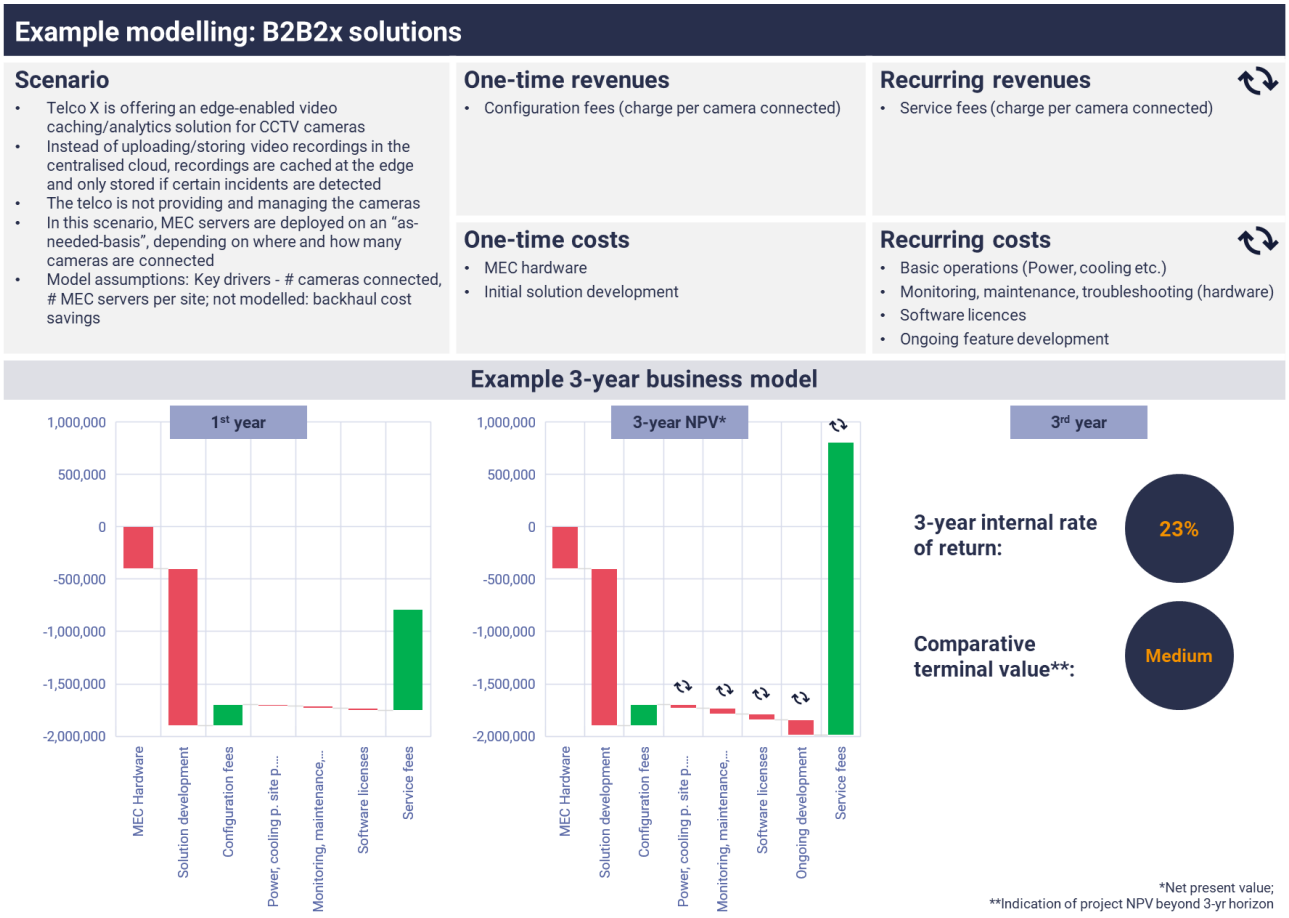

Source: STL Partners

Business model 5: End-to-end consumer retail applications

The last business model is 'end-to-end consumer retail applications'. Here, the operator is playing high up the value chain, acting as a digital service provider for consumer applications. As with B2B2X solutions, telcos would likely need to partner with others to provide these solutions, whether it be MEC suppliers, application technology vendors or channel partners.

MEC-enabled solutions in this category will leverage the benefits of MEC, namely low latency, high throughput and context awareness, to provide consumers with innovative applications. These could include Internet of Things (IoT), augmented reality (AR) or virtual reality (VR) applications that require video transmission in real-time.

Figure 13: Graphical representation of the end-to-end consumer retail applications business model

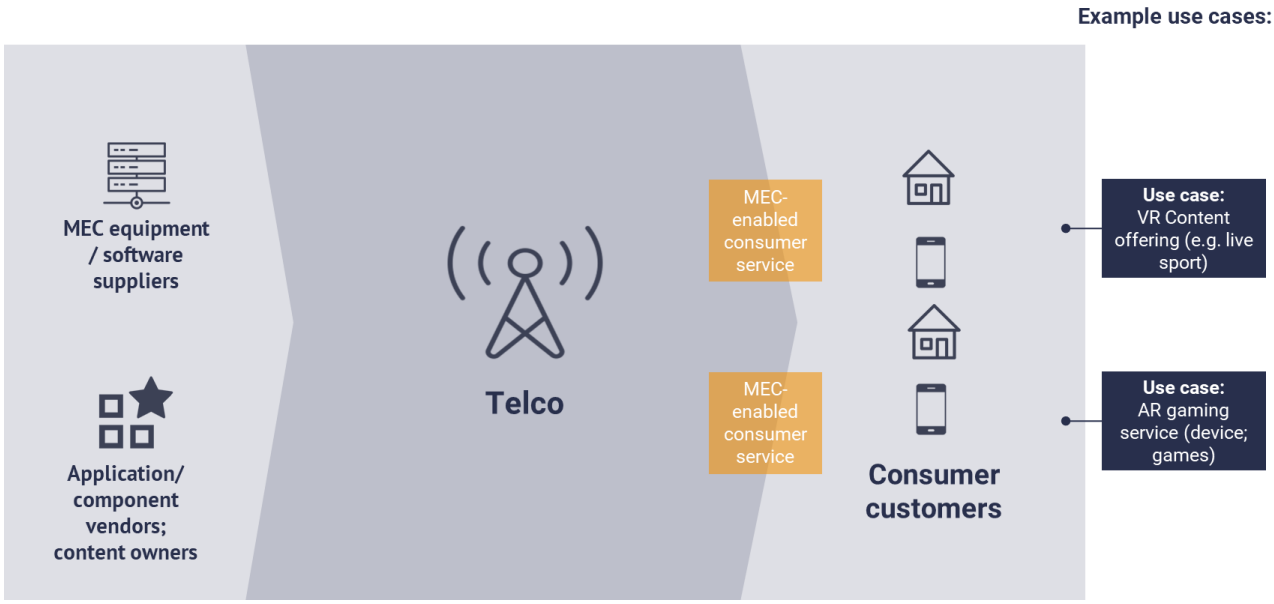

Source: STL Partners analysis

Some operators, such as BT, are already considering extending their existing entertainment service by adding VR to their premium offering to enhance the user experience and differentiate their proposition. Alternatively, operators could offer these services as paid add-ons to their existing services, for example pay-per-view sporting events broadcast in VR.

Such a scenario is also modelled in Figure 14. The quantified example shows that significant investment is needed to augment an existing content offering with MEC capabilities, e.g. for VR live sports. In the model, a break-even is only reached in the 5th year of the service, indicating significant risks (but potentially high revenue) is associated with this end-to-end business model. Note that we have not modelled other positive effects, e.g. on churn, which might be significant.

Figure 14: Quantified example – End-to-end consumer retail applications

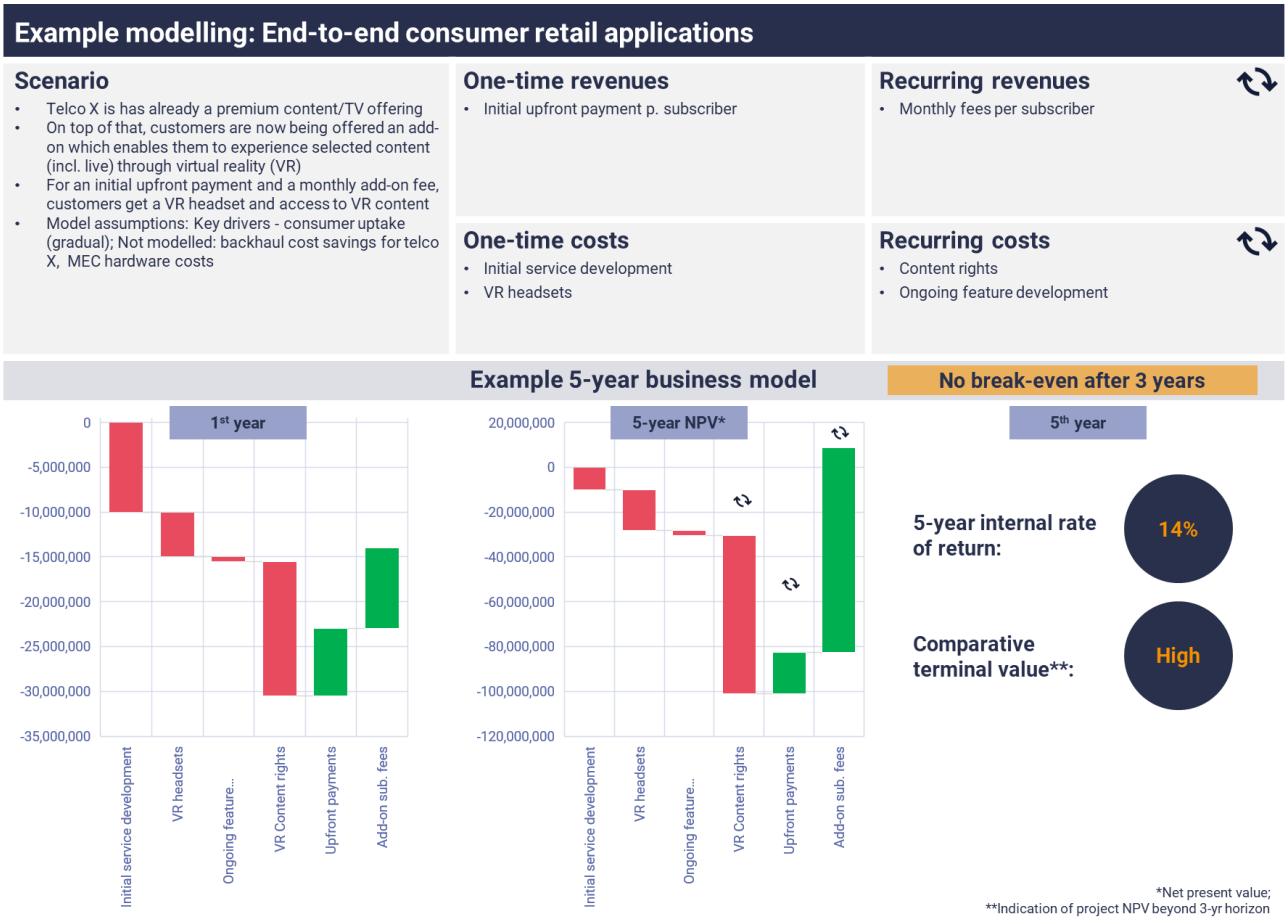

Source: STL Partners

Mapping use cases to business models

Given this broad menu of options (both in terms of use cases and business models for MEC), it is critical for telcos to choose the right business models for the respective application. Figure 15 maps the above business models to possible use cases.

We find that enabler business models tend to be suitable for use cases in which:

- Coverage (geographical, population) is required and application providers therefore need to work with several MEC providers
- There might already be an existing application by a third-party (potentially hosted in the centralised cloud) and e.g. edge IaaS is used to enrich/complement that application
- Net neutrality rules are relevant and telcos need to open up capabilities to third-parties

Whereas end-to-end provider business models are ideal for use cases in which:

- Mission criticality plays a key role and telcos need to provide a SLA-level performance/latency
- Compliance or privacy need to be guaranteed by a single end-to-end provider

Figure 15: Mapping MEC business models to possible use cases

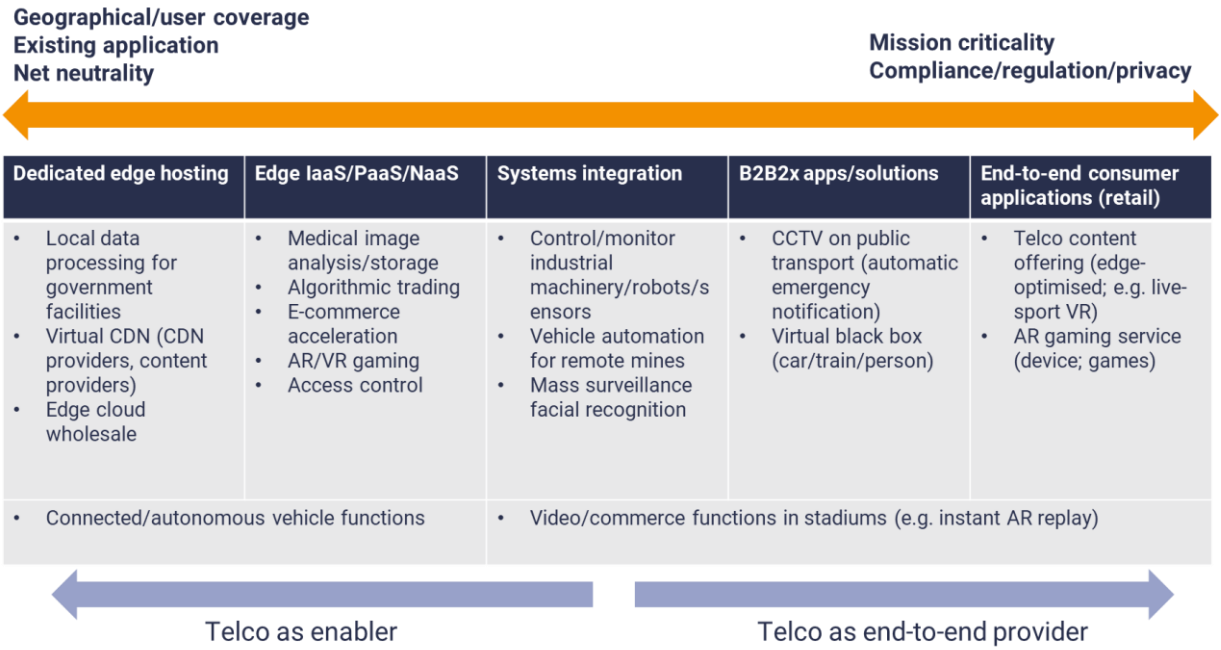

Source: STL Partners

Some business models will require a long-term view on the investment

The quantified models have shown that a higher rate of return after 3 years is typically associated with a lower terminal value of the investment (see Figure 16).

Figure 16: High 3-year IRR correlates with low terminal value

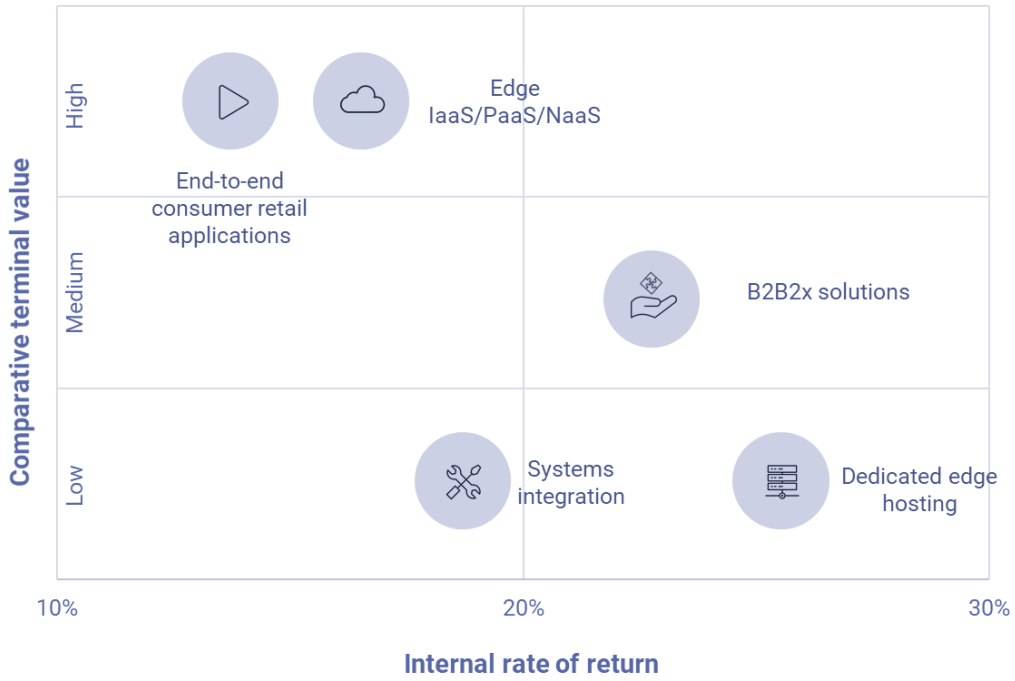

Source: STL Partners

Figure 16 shows, at least from this three-year outlook, that there *are* several business models which can be delivered with considerable profitability (namely "Dedicated edge hosting", "B2B2x solutions", "Systems integration"). However, these tend to be associated with a lower terminal value after this three-year period and a rather limited turnover/scalability overall. Put simply, they are not "cash cows".

In contrast, business models such as "End-to-end consumer retails applications" and "Edge IaaS/NaaS/PaaS" will exhibit a higher terminal value and potential scale after the modelled three-year period. However, the profitability of these business models will be still relatively limited at this point, which means telcos need to adopt a longer-term view on their investment. Figure 17 illustrates this.

Figure 17: Telcos need patience for edge-enabled consumer applications to become profitable (break-even only in year 5)

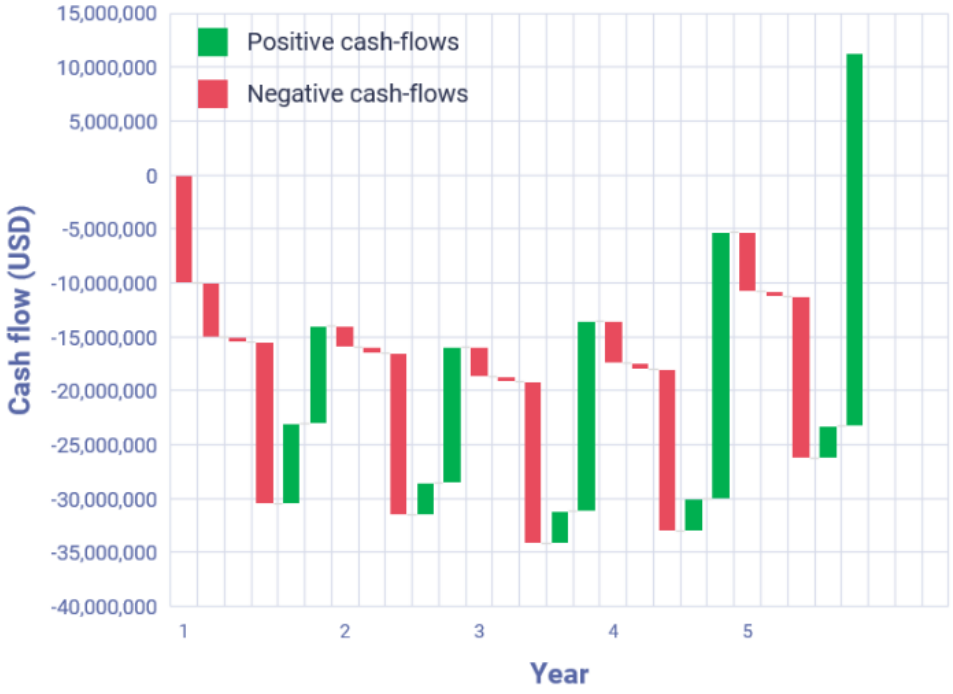

Source: STL Partners

Figure 17 shows how the VR-enabled content offering (as an "End-to-end consumer application") only breaks even in the fifth year. However, with growing adoption the business model then delivers an attractive margin and turnover. A telco therefore might find high value in such a business model, but only if it is prepared to view it as a long-term investment with an associated risk of failure.

Which business models are right for which operator and which operator division?

Adoption of the business models outlined above is not a case of picking just the one with the most favourable business case. In many cases, benefits are uncertain. It is also clear that none of the suggested business models is a cash-cow, and none individually will justify the upfront investment in rolling out MEC technology to the entire network. Operators looking to build value from MEC must therefore endeavour to build propositions in more than one area, where synergy will help drive an incremental benefit.

Even if operators are to explore more than one MEC business model, they must start somewhere. When planning where to begin, there are various considerations, which vary depending on operator characteristics and skills. Some questions to consider include:

- **What are your existing areas of strength, and can MEC help augment these?** For example, operators with a successful existing systems integration business may well look to start by leveraging MEC in the solutions they build for their customers.
- **How can you align your chosen business model with your existing customer base and network infrastructure?** For example, operators with a substantial enterprise customer base and well-developed fixed-line network infrastructure might start with the edge IaaS or B2B2x business models. Those with a large consumer base on cellular or small cell networks (e.g. MNOs) would be more likely to consider end-to-end consumer applications first.

Figure 18: The characteristics and skills required of the MEC operator change as we move along the business models

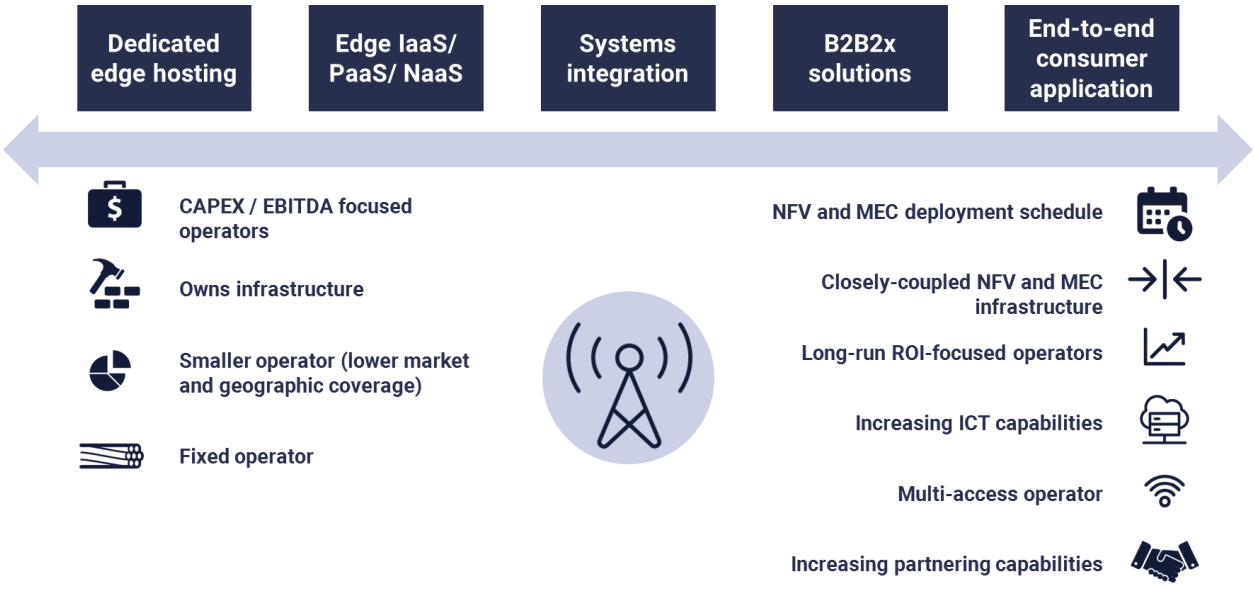

Source: STL Partners

- **Who within the organisation is well-placed to drive efforts?** If we adopt the philosophy that telcos should play to their strengths and current structure, it may be viable to drive MEC service initiatives on a divisional basis. However, if several MEC initiatives are underway (e.g. in for both business and consumer segments), there will be a need to coordinate across the organisation to create synergies; for example, a team developing an end-to-end consumer application might be internal customer of an existing Edge IaaS offering.

- **How easily can you implement a given business model?** Some of the business models will require less up-front investment and effort to bring to market. Enabling dedicated edge hosting, for example, may be as simple as a hardware upgrade on existing co-location sites, as and when your existing CDN partners require MEC functionality. Implementing a new end-to-end consumer application enabled by MEC, such as live VR content, is comparatively more difficult since it requires more investment of time and money in R&D, hardware, product development, and signing up new customers.

- **How does each business model fit with your existing virtualisation or technology upgrade programme?** This might include your organisation's strategy for adopting complementary technologies such as NFV/SDN, the Internet of Things, or 5G. If an existing virtualisation programme is underway, any prospective MEC development must complement it, rather than derail existing initiatives.

- **What is your risk appetite?** As discussed above, some end-to-end business models (such as consumer applications) might deliver considerable revenue – however only in the long run with significant investment needed upfront. Operators with a low-risk profile might therefore focus more on business models in which costs are only incurred when there is a successful customer acquisition (e.g. in dedicated edge hosting or systems integration).

It is likely that many operators will take a modular approach, adopting a complementary business model for an initial deployment, and expanding into further areas as demand becomes clearer and the technology starts to prove its efficacy. Timeline for deployment will vary immensely, and will be heavily dependent on factors including:

- The shape of the emerging 5G standard, of which MEC will form a major component

- Availability of internal resources (including technical skillset)

- Local regulation (currently unclear and likely to unfold over time)

- Competition from others, including hyperscale cloud providers (see the discussion above) and non-telco MEC SPs using small cell/unlicensed spectrum.

Conclusion

- Multi-access edge computing has created significant excitement among telcos and there is a plethora of potential use cases which leverage telcos' distributed compute capabilities.
- However, telcos' MEC efforts form part of a larger edge computing ecosystem which is only now beginning to emerge, most notably through recent moves by hyperscale cloud providers. It is unclear to what extent telcos will be able to secure a share of the opportunity given that there are edge computing concepts which are "telecoms-agnostic".
- To commercially realise the various MEC use cases, we have developed 5 telco-centric business models (Dedicated edge hosting; Edge IaaS/PaaS/NaaS; Systems integration, B2B2x solutions; End-to-end consumer retail applications).
- With these business models, telcos can choose from a menu of options which fit their risk profile and existing capabilities. Indeed, some business models (e.g. Dedicated edge hosting) are a safe bet for telcos, albeit with an overall limited impact on the business in the longer run. Other business models require telcos to invest more upfront and therefore take on more risks; however, such business models might become a significant source of revenue in the longer term.
- To conclude, telcos do have options to counter the emerging competition in the edge computing space. In order to avoid losing any ground, it is critical to start commercialising use cases and business models which are a good fit for an individual telco.

STL PARTNERS

Research

Consulting

Events