SCTE·ISBE CABLE-TEC EXPO 2019 NEW ORLEANS, LA SEPT. 30-OCT. 3

2019 Fall Technical Forum SCTE·ISBE • NCTA • CABLELABS

Why Gaming Needs An Edge

A Technical Paper prepared for SCTE•ISBE by

Alan Evans
Senior Director Innovation Strategy
EDGE GRAVITY by Ericsson
United Kingdom
+447826952348
alan@edgegravity.ericsson.com

Table of Contents

List of Figures

List of Tables

Introduction

The gaming industry generated almost 140B USD in 2018 (Newzoo Global Games Market Report 2018/2019). It has over taken the entertainment genre of choice for many consumers both young and old, so the need for a high quality of experience has become undeniable; if the experience is poor then the game will not be successful.

The pursuit of high performance is no longer the sole domain of the PC gamer; where thousands of dollars are invested in the latest GPUs and cooling systems to deliver the greatest clock speeds and frames per second.

Rendering performance is only part of the challenge. For many gamer consumers, the performance requirement is focused on the multiplayer experience, where poor performance means poor latency, lost packets and unpredictable jitter, which translates to a laggy experience, poor game play, lost competitive matches, frustration, and abandoned game titles. Our research shows trend where blame for this poor performance is often left with internet service provider.

The concern with quality of experience (QoE) when traversing the internet is not limited to competitive experiences, PvP (Player Vs Player), it also applies to PvE (Player Vs Environment) as well. Failed collaborative experiences where a social network of gamers (aka, Group, Guild, Org, Clan or Faction) are working together on an objective for several hours is just as frustrating as a competitive match, and perhaps for some even more so.

In addition, the Quality of Experience also has an impact on the commercial aspects of the game as well. Many of the most successful titles today fall into the category of “free to play”. Some of the most popular being of a “Battle Royale” game type, are of course demanding on the gameplay performance, but also on the in-game e-commerce, analytics, business intelligence and ultimately a shift to a reoccurring revenue of the traditional one-off purchase. If these less-graphical aspects of the game are affected by poor performance, then the game will not be able to generate revenue from its audience and the game will be a failed investment for the developer.

In this paper we will explore the statement, “why gaming needs an edge”. We will consider the classification of gamers and how they lend themselves to different game architecture. We will also explore the intersection of gaming and media, eSports, and how the broadband service provider edge can play a role in delivering the experience needed. We will explore the role of a Service Provider, gaming performance metrics and our research as it pertains to gaming acceleration at the Edge.

Content

1. Foreword

In 2019 no one would argue that gaming is now a mainstream form of entertainment, and access to games at anytime, anywhere, and on any device is becoming the expected norm. Just as with the TV & Media industry, the gaming industry raises similar questions of how to provide access to content on any device, provide a consistent and high-quality experience, and meet the demands of the various groups of gaming consumer; but are the answers to those question the same for both industries?

2. The Classification of Gamers

To be able to answer this question we should first examine the types of gamer, how the industry tends to classify them and why. NewZoo (*An excellent source of intel for the gaming industry*), and others tend to classify gamers into three categories, Casual, Competitive and Professional. These groupings are good, however do not provide a canonical view or provide any absolutes (but what does), in-fact a gamer that identifies as casual for one game, might identify as competitive in another, and so on. However, the real purpose of this classification is due to the eSports angle. Just as in more traditional sports, let's take football (soccer) as an example; a player might enjoy a casual game with friends, or they might join a competitive five-a-side team, or perhaps they have managed to make a career out of the sport and are paid and/or are sponsored for their skill.

So, let's try to describe or "profile" the three classifications, based on anytime, anywhere, and any device, i.e. time, place and device: -

Table 1 - Cassifying Gamers against time, place and device

	Casual	Competitive	Professional
Time	Opportunistic	Session	Organised
Place		Online	Venue
Device		Expenditure	Sponsored

There is an obvious demarcation between Professional and the other two classifications, and that is that the professional gamers get paid for gaming. But if we look a little deeper we can gain a relevant insight from what is important for professional gamers, and how it relates to both the casual and competitive gamer classifications.

Professional gaming events or "Esports" are organised events, at a specific venue and are nearly always sponsored. They are organised months ahead to allow professional gamers and spectators to a) travel to a venue, b) scope and build for and manage the event, and c) get some hype going to attract as many viewers as possible, both locally and online. The venue will have a dedicated offline network to provide the best possible conditions for the gaming traffic, as well as the best devices with the best technical performance to permit the best human performance possible. Only then can the gamers know for sure that there is no advantage, or limitation based on the technology used.

To a certain extent the "Competitive" classification is very similar to the "Professional" one; games are session based, and typically are considered "online" experiences with other gamers. A session might be a quick 10 min PvP (Player vs Player), or a 5-hour Cooperative PvE (player vs environment) experience. The point here is that a) the players have come together at a specific time, sometimes across multiple time zones, and b) they are all online, and are looking for the best network performance possible to ensure their gaming experience is optimal.

The casual gamer isn't really defined by whether they play online or not, (with other people) its more of an indication of the time they have available to spend gaming, as well as being less likely to buy new hardware or watch gaming streams. One might also choose to imagine that casual gamers are less concerned with their position on a multiplayer leader board, but I have known some serious Tetris players who would disagree, and even friends who have spent hours at the arcades to finally punch their three-letter gamer tag in the number one slot on Pac Man. No, I would hypothesise that the casual gamer is one that just enjoys gaming but doesn't necessarily have the time it takes to make sure their console, PC, game, online store, software is at the latest version before spending 10min on their lunch break playing a

game, nor are they likely to spend $120 on a season pass that provides a deep story line or extensive quests to pursue.

There is a tendency to try to segment or define the classifications by the device, specifically the amount spent on the device. But that can be misleading, most of the time the device can't be used to provide absolute classifications, its more of a reflection of how much a gamer is prepared to spend on their device. That might be a purely income-based decision, one gamer might play several hours a day on a 200$ console, whereas another might spend 30 min a week on a 5000$ PC. Not to mention that many gamers will have multiple devices capable of providing a satisfying gaming experience to meet the goal of anytime, anywhere, any device.

By understanding the criteria for classification of the gamer community we gain an insight into the types of game delivery system as well as supporting architectures types and how they might be improved upon or even reimagined to meet the industry need. We will explore this in more detail later in the paper.

What I find interesting is how the amount spent on a device relates to the mindset of the gamer and their expectation on quality of experience, and the culture of gamers.

3. Gaming Culture – is it possible to innovate around?

Ok, so I am not a phycologist. Neither would I describe myself as a serious student of human behaviour. But being a human being, I am inherently interested.

What was that about I hear you shouting at this text (or however you are reading this, or perhaps it made it to audio book?). Well let me explain, I believe that there is an innate desire to compare, to show off, to understand where we rank against others. I think it gives us drive, makes us to try harder to improve our performance and speaks to our competitive nature no matter how hard we try to conceal it.

So, is there really any difference between Single Player vs Multiplayer games? As I hinted at earlier I have known some very serious Tetris players, and frequenters of traditional arcades, who are incredibly competitive, yes with themselves, and they exhibit a completionist focus, that probably registers somewhere on a scale of obsessive-compulsive behaviour (I may be talking about myself here a little).

So yes, of course there is a difference between single player and multiplayer games. However most of it is focussed on the architectural requirements needed to deliver the compelling gaming experience. The culture of the gaming is such that it really is about the social aspect of gaming, and our need to compare our progress, our level, our score, and often as part of a community or a social network (the traditional sense of the term). In fact I love this quote from Zach Snader of Newzoo – "*Gaming is a way for people to connect with one another on a deeply human level accorss a shared interest. Gaming is a form of communication as much as it is a form of entertainment*"

So, is it possible to innovate around gaming culture? I don't think it is. I believe that there isn't a single gaming delivery technology or architecture that solves all the needs of the gamer community and places it into a single paradigm (not yet anyway, maybe the holodeck?). Gamers will continue to value their social interactions above a specific technology. They will still gravitate to games, that let them demonstrate their skill, via score, level, trophies, vanity accessories, emojis, dance moves (etc etc), to strangers, friends and families, work mates.

They will continue to spend on devices, technology and services then let them achieve this, in the way they want to do that will let their skill shine through. That could mean a high-end gaming rig, with 3ms 144Hz screen, the latest low latency keyboard and mouse, that are wired as to not lose any performance

between a key stroke and the reaction of the game engine, or perhaps a service that optimises the route taken for the gaming command traffic (controller button presses) from the device to the game server, or Broadband service providers that provide the best QoS for online gaming, or game downloads and updates.

It could mean a custom controller for a console gamer that feels less awkward for extended periods of gaming, or a tether free VR headset so there is no fear of tripping up, distracting from giving the best performance. Or quite a simply a leader board that lets one "punch-in" a three-letter gamer tag when a game of Tetris has concluded during a lunch break or daily commute.

Please don't misunderstand, this is not about one-upmanship, this is not about winning a game because the best technology or service was used or paid for. This is about removing performance impeding obstacles (or artifacts) using technology, with the objective to create an as close to a "level playing field" (a term I learnt from one of our technology partners, and fellow gamer) as possible so that high scores and victories are down to player skill, and not how much funding they have.

In our behinds-the-scenes industry of networks, cloud computing and platforms; we need to focus on how we support this culture, how can we enhance it how we can build solutions that enrich the experience of the community rather than exploiting it.

4. Delivering the game expereince

Games are traditionally hosted in the device they are being played on. i.e. the game engine runs on the device, input is taken from a control mechanism, into the game engine, the output is rendered and displayed on the screen. The gamer provides the feedback loop, and the game is experienced.

For general purposes we might draw the interaction between gamer and game like this: -

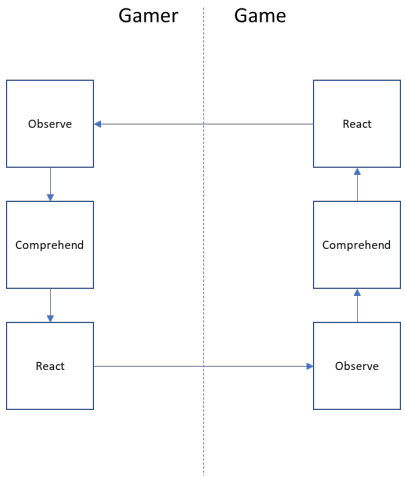

Figure 1 - Interaction between gamer and game

Perhaps an over simplification, but the point is that. Once the game is started the game is presented to the gamer the gamer observes, comprehends, and reacts, the game observes that reaction, decides what to do, and provides a reaction back to the gamer. The purpose for sharing this concept is to illustrate that at each stage there is potential for "friction" to be introduced as a by-product of the technology, mechanics and architecture used to deliver the game experience.

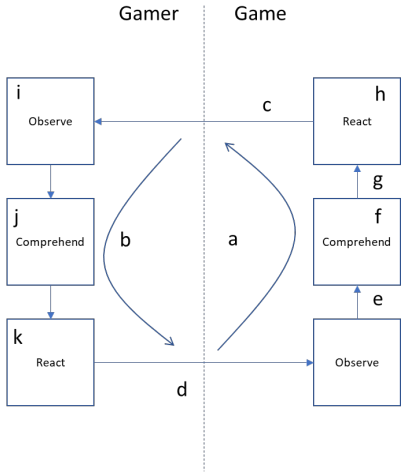

Figure 2 - Potential "friction" on the interaction between gamer and game

Once the propagation delay of the process (a) in figure 2 reaches a certain level, it will have a negative effect on the gaming experience. The gamers overall reaction process (b) will be so affected that the game will become unplayable and likely abandoned, in extreme cases.

If we explore this in more detail, and assume that the game has started, and the gamer is ready and about to react to the game, we can tell that delay introduced in the input (d) or information lost in the transmission would create a false observation in the game, which would affect the games comprehension (f) of what to do with the gaming input, and create an suboptimal output (g) and reaction (h) back to the gamer via (c), if further latency or poor performance affects the presentation of the game (c) to the gamers observation for the game (i), their comprehension will be frustrated (j) and the reaction (k) will attempt to compensate again potentially making the unplayable

Although the described process is quite abstract, the point here is that the interplay between gamer and game is sensitive to the potential short falls of the underlying architecture, and technology used to deliver a given gaming experience. Therefore, game developers state the minimum and optimum hardware (and software) requirements for their games to provide the experience they intended for the game and gamer, at least in the PC gaming market. The console market is slightly different, but once a console starts to "struggle" with the demands of modern game techniques, a new console is traditionaly brought to market. In the case of mobile devices, there as a lot more variation of device capability so principles from both PC and Console can apply.

The interesting part about specifying device requirements to support game experience is that once the game requires connectivity for some or all its game play, it is difficult for the game studio to create a game that works with all connectivity scenarios in possible. Most state "requires broadband connection" and typically that is as much as can be done. Some will perform a network test and might suggest that the speed and ping results are not suitable to play a title. Some even might state that they are detecting poor network performance in real time and terminate a gamers connection to the online servers. Frustrating for the gamer, as this is not something that is always within their control.

5. Gaming Architectures and performance expectations

In this section of the paper we will talk about the gaming architectures. Not the inner workings of the game itself or game mechanics, but the gaming delivery architectures. As I mentioned earlier, to meet the

appetite of the gaming community for gaming, and gaming performance there are various delivery architectures available, and they are needed, as not all architectures work for all gamers, game types, fit all markets or network topologies etc. there are too many variables for a one-size fits all approach, today.

In my opinion we can start of with two main categories:-

1. Device-Side Rendering, and
2. Server-Side Rendering

As the options suggest, we can split the architectures by considering where the rendering occurs. i.e. where the graphics are processed, or more specifically where is the GPU (Graphics Processing Unit) if one is required for a game in question.

The following diagrams are from the Newzoo report on the gamin market. I have borrowed them and have adapted them slightly for my purpose. If you the reader would like to learn more about NewZoo and their market report. There is a reference at the back of this paper.

5.1. Device-Side Rendering

Traditional online gaming architectures perform the game rendering on the client device, PC, Console, Mobile Phone, Tablet, and handhelds. All these devices have varying levels of performance, with PCs clearly out in front with some gamers purchasing nVidia 2080 RTX, or Titan GPUs, that can make their gaming rig cost 10K+ USD, to custom chips made by apple in the iPhone. The purpose of these load specific devices is to reduce the load on the CPU, for a specific task. With the latest iterations of these being capable of providing ray-tracing, (lighting effects) "automatically" in the GPUs, lowering the demands on standard CPU/GPU processes.

For these solutions the game is downloaded to the device (or on physical media e.g. DVD, Blueray, which is largely a legacy approach these days) and rendered locally.

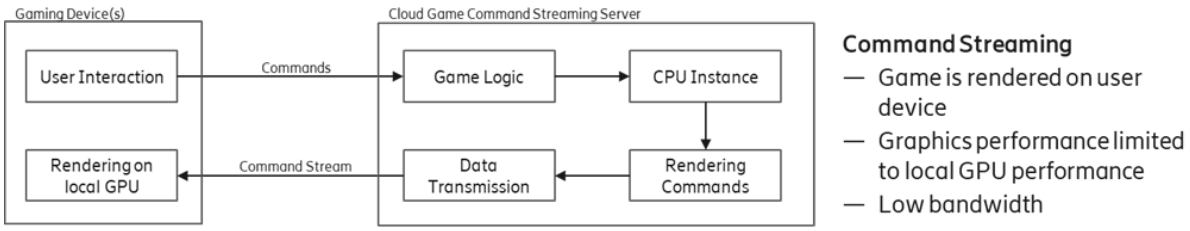

Figure 3 - Online gaming architecture

The challenge with this more traditional architecture; is the potential for introduced lag between the gaming device and the gaming server. Lag is the perceived delay between what the gamers interactions and what they see on screen. Its like a "glitch in the matrix" it can be very off-putting and can potentially make the difference between a positive or negative experience, or in more extreme cases a win or a loss in a competitive multiplayer experience. Therefore, professional eSports games are hosted on a private network to provide a "level playing field". In a Virtual Reality setting, it could result in a gamer being nauseous (I speak from experience). But for player vs player, multiplayer or indeed any online game environment, Round Trip Time (RTT) is a problem. A problem for the gaming experience, and a problem for the ecommerce aspects. The more gamers are in the same online session, the more critical the RTT becomes on preserving the experience.

The great thing about this type of architecture is the game performance. The localised interaction between the gamer and the game is as good as it gets. PC providing the best performance, with graphics rendered natively and delivered to the eyeballs in flicker free 144Hz under 3ms, with all command traffic from the controller being processed locally. The only real lag being a possibility from any online requirements, and the network quality of service (QoS) required is focused on latency and jitter rather than bandwidth.

5.2. Server-side rendering

There is another way to deliver a game to the gamer's device, and that is streaming it as video. By streaming a game as video, the end device doesn't need to be as powerful as it is no longer performing the rendering for the game, it is simply transmitting game control input from the gamer and decoding video set to it from the Cloud. The performance required to create a compelling game experience is now the responsibility of the cloud, the game server, and the networks that deliver the video.

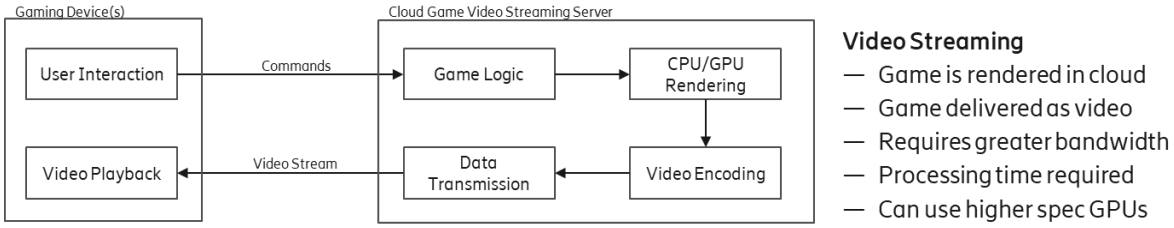

Figure 4 - Cloud gaming architecture

This technology / approach lends itself very nicely to addressing the more casual gamer who want to be able to play a game anywhere any time on any device. This doesn't mean that the more AAA titles can't be delivered in this manor, there are just some considerations that need to be acknowledged.

One benefit is that the game studio only needs to create the game once rather than a version for every platform. As I said earlier, the issue of rendering and delivering the game becomes a challenge for the cloud and the network. There are game service providers offering this type of service, with hundreds of game titles already available, in fact I think this is an important detail, this model lends itself towards the subscription gaming model where gamers pay a monthly fee in an "all-you-can-eat" model, which is a relatively new concept for the gaming community, and we will have to see how it is adopted by the free to play games, that are heavy on in-game purchases.

In addition to using this architecture paradigm for offering a catalogue of games delivered as on-demand video, some companies also offer this as a "gaming-rig-in-the sky" solution. The gamer specs their ultimate machine, and a virtual machine is created in the cloud. It doesn't come with any games, as they need to be bought separately, but the games are delivered as a video stream just the same.

The challenge this architecture faces are like the previous device-side rendering architecture. Quality of Service (QoS) parameters are important as its is critical that all the controller command traffic makes it too the cloud, and quickly, but more so that device-side rendering, because the rendering of the whole game once all input has been gathered, i.e. the player and all other players that might be in a session together, then need to be processed, then rendered, and then delivered back to the device.

The process of delivering the rendered game back to the device can't be done natively. That would require far too much bandwidth. It must be encoded as a video stream. Encoding is not a low latency process, the more time an encoder is provided to encode a picture the better the picture will look. With

some bit rates approaching 15Mbps for HD resolutions at 60fps using H.264, the picture quality, and visuals in general are a sacrifice that is made to accommodate the delivery approach.

Of course, delivering video OTT (over the top) is not a new idea. Network quality issues can be mitigated by using adaptive bit rate techniques and buffering to ensure that the video service is as uninterrupted as possible. Unfortunately, these techniques do not lend themselves well to the gaming use cases, as the gaming experience must be uninterrupted, high quality, and always available.

6. Gaming and the edge

So how can the edge help with these architectures, and does the edge offer additional opportunities for innovation around how games can be designed?

First, let's quickly qualify what we mean by the edge. For the purposes of this paper we are not going to address any single network topology, although of course in a real-world deployment the topology of different network types, 3G, 4G, 5G, HFC, DSL and Fibre (plus wifi) all have their specific considerations.

The broadband service provider edge; is the edge that we are concerned about today, and when we refer to the edge, we are referring to points-of presence (PoPs) with run-time environments that can be used to execute workloads, to gain a performance improvement by some measure.

The figures for round tip times (RTT) are purely illustrative and will vary from network to network.

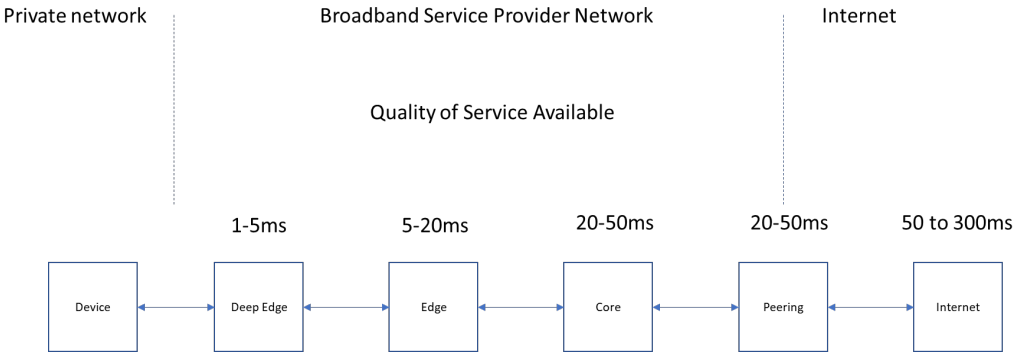

Figure 5 - Latency at the broadband service provider edge

Earlier in this paper we introduced the concept of device-side and server-side rendering along with their pro's and con's, strengths and weaknesses, to highlight that there is potential to provide an improvement to the principle architectures that will result in an improvement of game experience, for casual, competitive and professional gamers alike.

In this paper we are not trying to suggest that there is one true gaming architecture, only that there is a place for them all, but all can benefit from optimisations the edge can provide.

6.1. Ping acceleration and the edge

If a gamer has chosen device-side rendering type of environment then, as we discussed earlier the need for greater performance stems from the competitive eSports culture, and competitive online multiplayer games.

These games are typically PC based, console competitive gaming is also relevant here. Uptime and connectivity are equally relvant on any platform, PC or otherwise, and for any online-connected game. I say PC mainly becase solutions in this space are provided by 3rd parties and may likely require a client to be downloaded to the device. Because latency can ruin a game, gamers turn to companies that can improve the latency, packet loss and jitter of their connection between their chosen gaming device and the gaming server.

These companies act on the command data between the game and the game server. So, this means the controller button pushes from the gamer, where they are in the game, and any information, other players in the game environment need to know about (via the game). They intercept this traffic and find the most optimal route between the gaming device and the game server.

Depending on the location of the gamer's device and the game server they are connected to, latency performance can be improved to ping levels of 10ms and lower, but perhaps just as important packet loss and jitter are reduced to zero due to the ability to perform multipath routing between the device and the game server.

Game Servers are hosted out in the internet somewhere, this can be either in a datacentre or a public cloud. By providing edge locations within the broadband service provider network, within the edge locations of the hosting datacentres and in between, the accelerator technology can make better decisions on how to route game command traffic between the device and the gamer server.

6.2. On-demand Infrastructure at the edge

The ping accelerator technology described above acts on the command traffic between the gaming client, and the gaming server. Although this technology and approach works incredibly well and makes a significant difference for the competitive multilayer experience, the performance increase to the gamer can be variable depending on the gamer's location, the game server's location, and the route between the two.

An emerging approach to improve the gamers multiplayer experience comes from another angle, is the idea of on-demand infrastructure. Rather than accelerating the game data, move the game server closer to where the gamer is.

The challenge here is that this is that for multiplayer gaming, simply moving (instantiating) a game server to "somewhere closer" to the end user is not enough. This is largely because multiplayer gaming goes through a process called "match making".

Match making, is the process of placing a gamer into a multiplayer gaming session with other gamers of a similar skill level and gamer location. It is pretty much an internal game mechanic to make sure that the multiplayer session is a balanced one for all gamers in the session, as well as addressing language issues where communication is important to the experience.

These on-demand gaming infrastructure companies solutions complement this process. During the match making process, these solutions look at which gamers are to be added into the gaming session, and then

instantiating the game servers in the most logical point of presence to provide the best possible experience for the gamers in the multiplayer session. The placement of the server considers the performance of the PoP, both compute and network, as well as its logical and physical location to attempt a level playing field for the multiplayer session.

When the session concludes the instance is destroyed, and the resources are released. By leveraging the service provider edge, game studios would be able to provide an option for multiplayer gamers to take advantage of a multiplayer gaming infrastructure that carried with it the performance of a managed network suitable for an eSports event.

6.3. Server-side rendering (cloud gaming) and the edge

Streaming video is not without its challenges. Delivery of games as video streams increases the complexity further. By bringing the cloud game servers within the broadband service provider network, the round-trip time for the command traffic and the resultant delivery of the video stream could be optimised beyond what can be delivered for this type of service from a public cloud.

The benefit of leveraging the broadband service provider edge should permit the use of higher resolution and the required bitrate of video, whilst reducing the amount of buffering in the gaming device, all of which should contribute to an improved quality of experience for the gamer.

6.4. Device Offload to the edge

Until now we have mostly addressed two main architectures, but the edge might facilitate a third. With edge resources becoming available within the 1ms to 5ms RTT range. It could be possible to offload some of the processing functions of the device, towards the edge. This would allow the devices to perform more complex use cases, without draining their batteries. There would be a need to find the balance between transmit and receiving data, vs processing it on-device, but an interesting study for sure.

Conclusion

Gaming technology is becoming more and more relevant to other industries, just as Sci-fi stories have predicted future technologies, the gaming industry and its technology provide insight into the consumer and enterprise applications we expect to see in the very near future. The performance requirments to create a compelling gaming experience, have even more critical impacts that can dictate the viability of a use case in various IOT scenarios.

At Edge Gravity, we are in the process testing the described architectures in conjunction with leveraging the broadband service provider edge. We are interested in collaboration in all forms, e.g. network type and topology, hardware accelerators and processors, game platforms and game developers and so on.

We expect to see test results that will demonstrate the benefits edge cloud and edge compute for these types of solutions and a show why gaming (and other industries) needs an edge.

Abbreviations

PoPs	point(s)of presence
RTT	round trip time
AAA	triple A rated
HFC	hybrid fiber-coax
PC	personal computed
HD	high definition
QoS	quality of service
QoE	quality of experience

Bibliography & References

Ericsson Consumer Labs 2019

Gameye CEO - Sebastiaan Heijne

Newzoo - Cloud Gaming: The Perfect Storm 2018, https://newzoo.com/insights/trend-reports/free-report-cloud-gaming-the-perfect-storm/

Newzoo - Global Games Market Report 2018 & 2019, https://newzoo.com/solutions/standard/market-forecasts/global-games-market-report/
Ericsson Consumer Labs 2019

