


Engaging the Game Industry and Edge Updates at Verizon


Dec 2019

5G & The Graphics Edge

- Verizon 5G has demonstrated over 1 Gbps speeds and single millisecond latencies.
- At the same time, Verizon has been building out GPU compute throughout the network in order to create a low latency, graphics edge compute platform.
- By placing computing devices and GPU's right in the last leg of the network we allow for super quick round trip times and rapid options for offload and virtualization.



Where is the Network Edge?



The Network Edge is amorphous and can exist at any or all of these locations.

As 5G and edge computing develop the concept of cloud vs the edge disappears and we begin to utilize resources based on application demands.

...but why?

Benefits Include:

- Processing complex workloads on limited hardware
- Reduce heat on a user endpoint
- Extend battery life on a constrained device
- Smarter, faster, cheaper devices that are more secure!




The State of Edge Based Gaming

Where are we Today with Edge Based Games and 5G?

As carriers are rolling out 5G globally, a number of trends have started to emerge, and a variety of products have launched.

1. Google Stadia, GeForce Now and other services are trying a fully stream based approach
2. Venue based experiences are being piloted across stadiums and arenas
3. Verizon is creating new services and API's



What have we learned so far from streaming?

Virtualized Game services are expensive, low capacity and unreliable

- Gaming is incredibly unresponsive and laggy
- Latencies ranging from 90ms to 9500ms. Far more reliable on mobile networks

This method delivers fantastic results for those who have access to the best available networks.



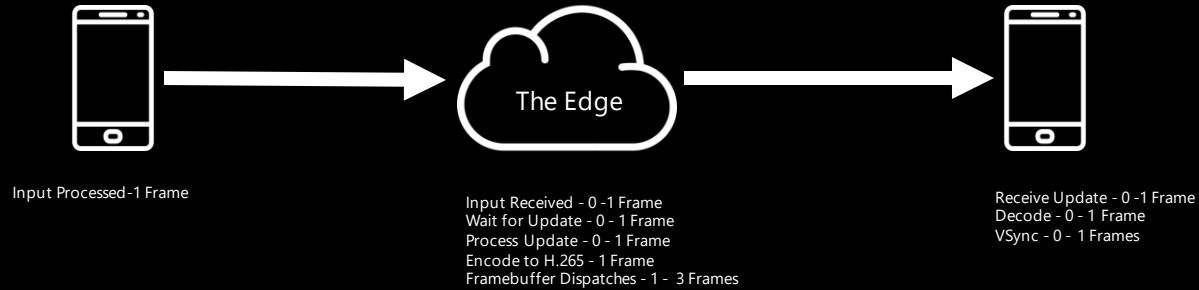
But Why?




Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.

THE PATH OF AN UPDATE

Virtualization – Best Case 3 frames, Worst Case 10+ frames



Microservice based approaches – Best Case 3 frames (Updates and Draws can occur Asynchronously)



What have we learned so far from streaming?

Virtualized Game services are expensive, low capacity and unreliable

- Gaming is incredibly unresponsive and laggy
- Latencies ranging from 90ms to 9500ms. Far more reliable on mobile networks

This method delivers fantastic results for those who have access to the best available networks.

Let's look at what Verizon is doing and some microservice based approaches.



Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.

The Washington Post

Democracy Dies in Darkness

Review

Unplayable at times, magical in others: Stadia's dream is still in the clouds

The Post tested Google Stadia. The input lag is horrendous.

Google Stadia is a great concept for the next generation of gaming, but its input lag — which is variable — makes it nearly unplayable. (Jhaan Elker/The Washington Post)

By **Gene Park**
Reporter

Hands on: Project xCloud review

Microsoft's game-streaming service gives you a choice for how and when you play

☆☆☆☆ Not yet rated By Vic Hood 18 days ago Gaming 40

What is a hands on review?

Image credit: TechRadar

EARLY VERDICT

Project xCloud seems promising with smooth imaging and the ability to run at a lower latency than its competitors, but we're concerned whether this stability will hold up when servers are supporting more than a room full of devices...

+

FOR

Can run at 7-10mb/s
Mostly smooth image and streaming quality
Game progress carries from PC and console
Runs AAA games
Control input is solid

-

AGAINST

Currently only Android has been announced
Struggles with faster paced online titles
Connection occasionally drops out

What Verizon is doing to Engage with Developers

1. Hackathons and Indie outreach
2. Major partnerships: Epic, Snap, Facebook, and AAA studios
3. Verizon is creating new services and API's
4. 1st Party Game development
5. Localized gaming experiences
6. Deploy of major streaming services with custom built games on



Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.

GPU Services on the Edge

Development Roadmap

Create an Edge Native Platform and API Services that allows for new classes of experiences on all devices:

- Develop easy to access API's for Rendering, Computer Vision, XR Lighting, Raytracing, Spatial Audio and more
- Create an easy to use platform that lets anyone create and deploy Edge Services
- Connect Edge Services to client platforms with simple plugins for Unreal Engine and native platforms (Windows, Android, iOS)
- Enable developers to test and build!



Rendering as a microservice

What?


- API Gateway
- Load Balancing
- Centralized analytics
- Configuration Management
- Containerized runtime
- A single automated build / testing platform
- Packaging and deployment

Why?

- Decouples renderer from other subsystems
- Supports many different uses cases
- Scales seamlessly across multiple slices / GPUs / machines

What is Hybrid Edge Rendering?

- An alternate approach to streaming: both server and client render some portion of scene
- Combines benefits of network-accelerated rendering with latency-free experience typical of client-only rendering



Edge XR Lighting API

Benefits of Edge Based Lighting

- Real-time lighting / reflections on mobile devices for XR applications
- Low latency makes real-time reflections in XR possible
- End user setup of live AR portals and subscription could support new types of shopping / tourist / entertainment experiences
- Could also be used for generating in game lighting / reflections, instead of camera, imagery can come from virtual world to provide higher fidelity mobile experience than is available on current hardware



Real-Time Rendering API

- Platform to offload rendering to the edge
- Streams results to client devices
- Distributed render graph framework
- Supports compositing, compression, asynchronous GPU read and streams to client devices
- Client SDK features a mobile friendly forward renderer which composites server data in real-time

Benefits of MEC

- End user can make their own rendering requests
- End user gets dedicated container for custom rendering
- Offloads graphics processing
 - Saves device power
 - Allows you to do things that cannot be done on mobile hardware
 - Supports MEC thin client and path for cheaper phones



Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.




MEC XR Lighting API

MEC Deployment Requirements




- 1GB vGPU per Camera / Environment
- Up to 64 users per container
- Container launched per camera
- 10gbps network total downstream
- 150mbps per user downstream

Benefits of MEC

- Real-time lighting / reflections on mobile devices for XR applications
- Low latency makes real-time reflections in XR possible
- End user setup of live AR portals and subscription could support new types of shopping / tourist / entertainment experiences
- Could also be used for generating in-game lighting / reflections: instead of camera, imagery can come from virtual world to provide higher fidelity mobile experience then is available on current hardware



XRLighting



Hybrid Edge Rendering Vs Streaming

Advantages

- Scales very well in GPU Slices
 - Offload expensive computations to Edge
 - Do not need to run entire game run in a container
- Simplifies build / deployment / testing
- Does not suffer from input lag
- Combined use of mobile hardware
- Saves battery / reduces overheating (but not as much as streaming)

Disadvantages

- Requires tighter integration into existing architecture
- Need to leverage specific use cases

Raytracing API

Benefits of Edge Raytracing

- Supports high end desktop quality graphics on mobile devices
- End user can make their own rendering requests
- End user gets dedicated container for custom rendering
- Vulkan + RTX pipeline
- Lightmap atlas stream combined with local rendering pipeline for hybrid lighting effects



Computer Vision API for XR

Benefits of Edge Based CV

- Real-time horizontal scaling computer vision for 2D and 3D recognition and tracking
- Supports full or hybrid render
- Fully featured webGUI for designing and training targets
- Supports occlusion, classification, and segmentation





Verizon GPU Edge Platform



Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.

GPU Based Orchestration Platform


Verizon XR Team has been building an independent GPU Based Orchestration system specifically aimed at solving the following:

- GPU slicing and management of vGPU/vCPU slices for multi-tenant applications
- Load balancing of distributed clusters of GPU blades (federation)
- Dynamic discovery of available clusters by interfacing with operator network
- Remote provision, deploy for Docker-based GPU Applications
- Rapid MEC development pipeline for mobile device to compute backend using Unreal Engine or Native Applications
- Providing integration and API test bed for Edge Native API's



Network Accelerated API's

- Combined power of 5G and Nvidia Server GPUs
- High bandwidth / low latency
- Powered by Verizon Graphics Orchestration Platform
- Envrmnt Creator to manage asset ingest, content creation, data caching



Platform Stack



Verizon XR GPU
Orchestration Platform:

Shared cluster Paas.
Orchestration and federation
of Kubernetes clusters within
Edge Locations and across the
whole network.



In cluster container
orchestration:

Kubernetes clusters with
custom NVIDIA docker
support.



Virtualization:

Red Hat Enterprise Linux®
kernel, Kernel-based Virtual
Machine (KVM) technology,
and oVirt virtualization man
agement projects.

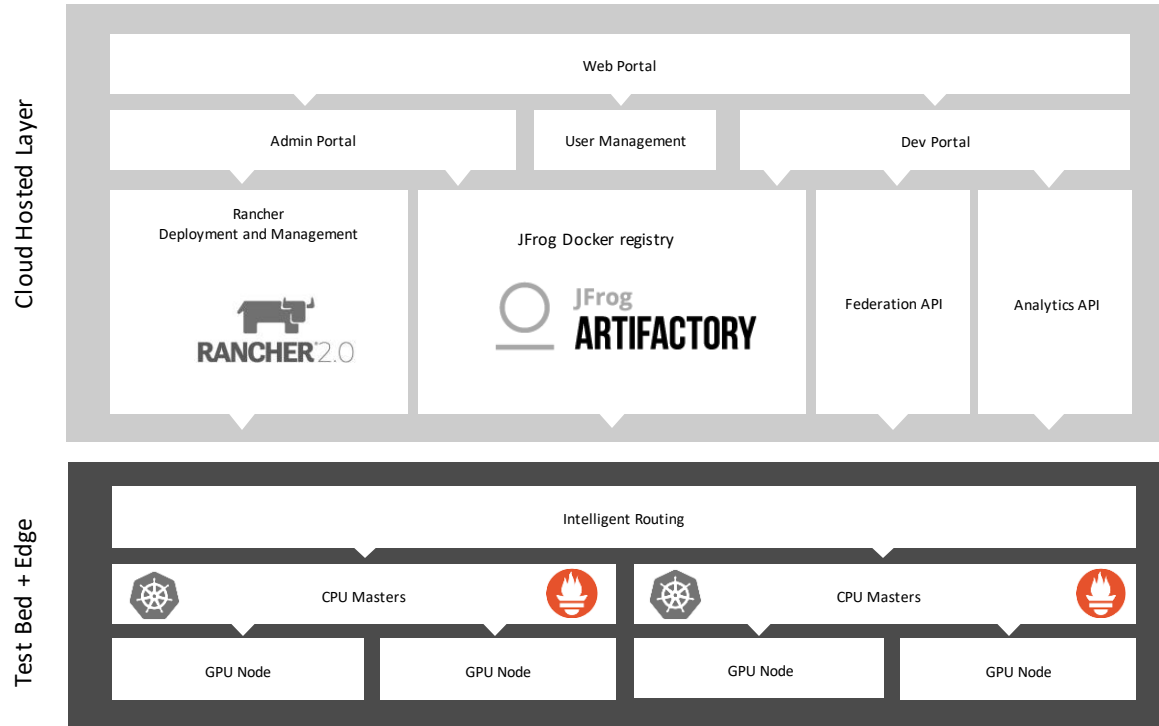
NVIDIA GRID virtualized
GPUs



Hardware:

- Blade Config 1
- 2x Xeon-S 4110, 128GB
RAM, 2x 16GB NVIDIA Tesla
T4
- Blade Config 2
- 2x Xeon-S 4110, 128GB
RAM, 2x 16GB NVIDIA RTX
Quadro RTX8000
- Blade Config 3
- 2x Xeon Gold 6168, 128GB
RAM, 8x 32GB NVIDIA Tesla
v100

Orchestration Platform - Functional Architecture



THANK YOU

Raheel Khalid - @rkhalid890



Confidential and proprietary materials for authorized Verizon personnel and outside agencies only. Use, disclosure or distribution of this material is not permitted to any unauthorized persons or third parties except by written agreement.

EDGE SERVICE CALL FLOW

