

# Verizon Graphics Edge

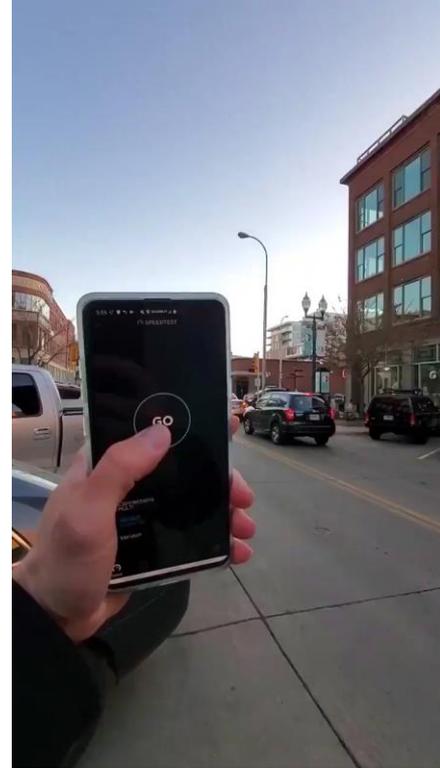
Open Edge Workshop 2020

May 2020

verizon

# 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 are we today with Edge Based Games and 5G?

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

1. Google Stadia, GeForce Now, Microsoft xCloud and other services are trying a fully stream based approach
2. Several companies focused on high performance compute and optimization of GPU services for games
3. Venue based experiences are being piloted across stadiums and arenas
4. Verizon is creating new services and API's



## STADIA



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.

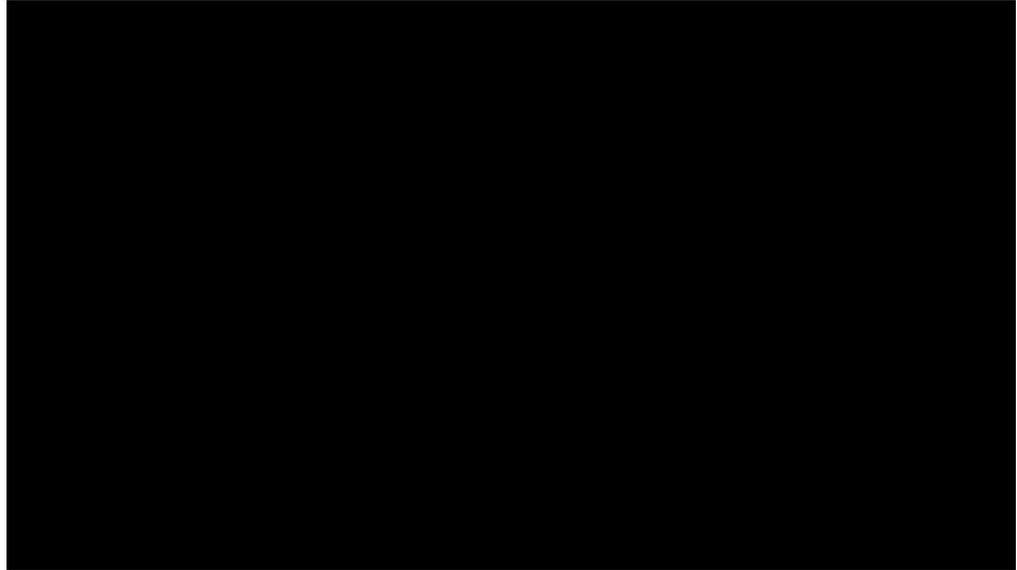
# What Verizon is doing to Engage with Developers

1. Hackathons and Indie outreach
2. Major partnerships: Epic, Snap, Facebook, Niantic, Bethesda, and AAA studios
3. Open platform and services
4. 1<sup>st</sup> Party Game development
5. Localized gaming experiences
6. Deploy of major streaming services with custom built games



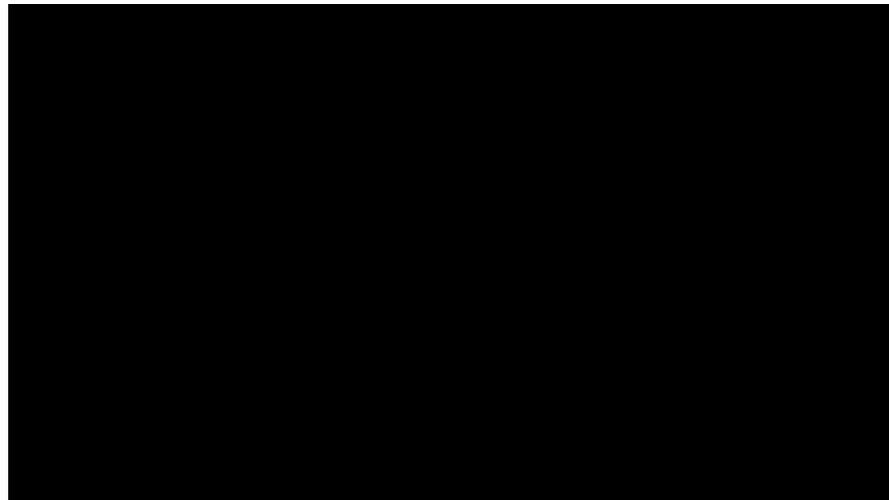
# Consumer Edge Services

1. Computer Vision for AR Headsets, games and retail
2. Immersive communication services for video conferencing, VR meeting spaces and high performance gaming
3. Groundbreaking gaming experiences
4. New next generation retail experiences for consumers



# Enterprise Edge Services

1. Computer Vision for Smart Cities, AR Headsets and enterprise analytics
2. Low latency IOT for commercial, retail, and manufacturing
3. Communication services for remote working
4. Low latency compute clusters for visual analytics and data visualization
5. Operations management in real time using ML



---

## Built for 5G

AI/ML, AR/VR, CV, AAA Gaming, & Real-time Enterprise are **highly dependent on GPU's** for compute

**5G's ultra low-latency properties and bandwidth** effectively created a real-time bus between Edge-GPU and mobile device which **will drive new levels of mobile access to extremely high-powered compute.**

But...

Unlike CPU's, there are no solutions for GPU compute resource management that allow cloud deployments to be affordable/profitable for many of these technologies while guaranteeing network fidelity.

---

# Net Result

## *GPU Based Edge Compute Platform Distributed in the Verizon Network*



- **Revolutionizes mobility for AR/VR** (ultra-lightweight HMDs)
- **New class of affordable mobile cloud services for enterprises** (including real-time computer vision and graphics rendering)
- **Platform for developing <20ms cloud gaming business**
- **Enables the development of dozens of previously cost prohibitive GPU cloud base services** (AI / ML / Encode/Decode, etc.)
- **Accelerates traditional cloud service at the Edge**

# 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



# verizon DevKit



LIANA

We've been at this for ages. When do I finally get to take on a mission for the guild?



DARRIN

Available Now!

# DevKit Highlights

- Verizon is releasing an all-new datacenter grade server for low latency edge compute built with Nvidia and HPE.
- DevKit features two RTX Quadro 8000 GPU's which enables next-generation graphics compute for gaming, computer vision, AI, raytracing, audio and video processing. Bringing 48 GB of high-speed GDDR6 memory and an optimized system for high-performance computing.
- Designed for the Edge: Verizon designed and optimized orchestration platform built on Kubernetes and RedHat Linux combines true GPU multi-tenancy and telco-grade security with lightning fast network delivery.
- Verizon's Graphics Edge portal allows instant load balancing and smart routing across the US and access to best in class edge API's for complex raytracing, deep learning, and visual computing workloads.

# Graphics Edge API's

Enable next generation raytracing,  
computer vision and AI on any  
platform with the Verizon DevKit!



## Computer Vision

- 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.



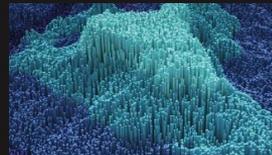
## Raytracing

- 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.



## Spatial Audio API's

- Ability to offload expensive audio processing to free client resources & enable thin clients.
- 15 – 50ms total RTT target, including DSP.
- Supports high-end, desktop-quality immersive & social audio experiences, otherwise impossible on mobile.
- Supports large number of concurrent OTT voice chat & spatial audio users per application.
- Able to do machine learning on voice for captioning, translation, & other data analyses.



## Hybrid Render

- 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.

# Deploy Your Own Containerized Applications!

Take your own applications to the network Edge. Develop and deploy any containerized application and take advantage of highly reliable low latency edge servers distributed across the Verizon network.

**For more information contact:**  
[gpuedge@verizon.com](mailto:gpuedge@verizon.com)

# 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.