AUTODESK UNIVERSITY

TR502979

Automotive Virtual Design Review on AWS

David Randle
Amazon Web Services

Learning Objectives

- Understand the Virtual Design Review Use-Case
- Learn about the AWS Quick Start for enabling Virtual Design Reviews
- Understand the Amazon EC2 G5 benefits
- Be aware of an estimated TCO (Total Cost of Ownership)

Description

Autodesk VRED 3D visualization software helps designers and engineers create product presentations, design reviews, and virtual prototypes using interactive GPU ray tracing and both analytic and cloud-rendering modes. This session highlights the AWS Quick Start designed for IT infrastructure architects, administrators, and DevOps professionals who are planning to implement or extend their Autodesk VRED workloads on the AWS Cloud to enable realtime, virtual, design reviews.

Speaker(s)

David Randle is the Global Head of GTM for Spatial Computing at Amazon Web Services, helping people create immersive experiences at scale. He brings 16 years of experience marketing 3D as a high value communication medium. Having democratized real-time rendering at a start-up called Bunkspeed and subsequently 3D design after an exit to Dassault Systemes SolidWorks, he is intimately familiar with the value chain of immersive experiences.

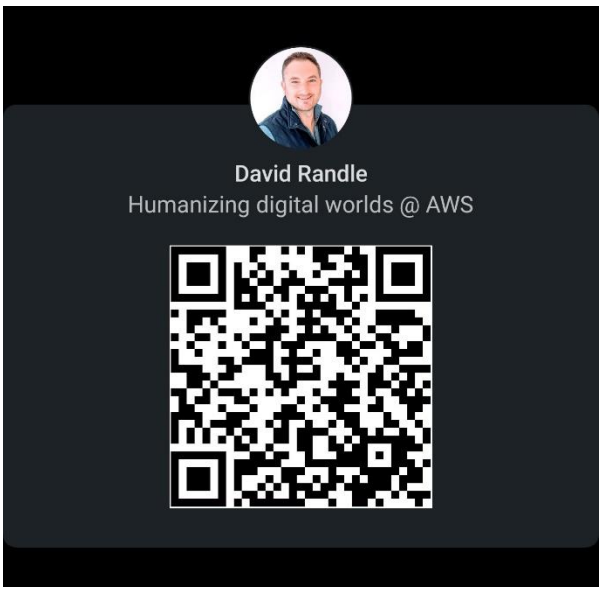

Understand the Virtual Design Review Use-Case

Within automotive customers, automotive design departments (who typically reside within the R&D org) have been augmenting physical processes (2D drawing/sketching/rendering, clay modeling, and buck production) with digital processes (digital sketching/rendering, 3D modeling, and visualization) for more than 16 years in a quest to improve product quality, competitive differentiation, and product market fit. These processes have been used in practice to increase stakeholder confidence during the decision-making stages of product development (i.e. the cutline between the front and rear door of a car design should be moved forward by 3cm to restore visual balance). The COVID pandemic created a dramatic work dynamic shift impacting the design review process, since design reviews would occur in-person at design studios. High end graphics workstations would power real-time visualization of 3D versions of upcoming vehicle designs. Stakeholders (design management, executives, designers, and marketers) would congregate and review the vehicle designs presented on large format screens (powerwalls) and VR headsets tethered to the workstations. Recently (within the last 4 years), accessibility to VR (Virtual Reality) and AR (Augmented Reality) devices (i.e HTC Vive VR headsets and tablets) has been on the rise (thousands of VR devices and tens of thousands of mobile devices being deployed and managed within automotive OEMs) which is leading to new capabilities within the automotive design process. This paired with the need to facilitate reviews among stakeholders who are located remotely has created the perfect accelerant for virtual design reviews.

Learn about the AWS Quick Start for enabling Virtual Design Reviews

In order to meet the desires of automotive OEMs, three components are required: 1/ Design software such as Autodesk VRED that has the capability to render a 3D vehicle in context, at scale, and to multiple stakeholder points of view from an on-premise server or cloud server (VRED Core). 2/ A streaming protocol such as Nvidia CloudXR that compresses server-side rendering and decompresses client-side images at low latency, while transposing 6 DoF (6 degrees of freedom) data from devices, back to the server to render the next frame from the adjusted point-of-view in near real-time. 3/ Cloud or on-prem infrastructure supporting scalable, on-demand real-time graphics workloads with low-latency edge delivery capability. Automotive designers prepare the vehicles for evaluation in Autodesk VRED on desktop workstations, then publish the content to VRED Core (on a Windows G4dn instance). Nvidia CloudXR, on the same instance compresses real-time rendered images from VRED and streams them to Nvidia CloudXR clients (VR headsets or tablets) operated by design review stakeholders from independent locations and perspectives, simultaneously.AWS has created a Quick Start (Cloud Formation Template and documentation) to simplify the provisioning of this solution for all automotive OEM design studios and other Autodesk VRED customers.

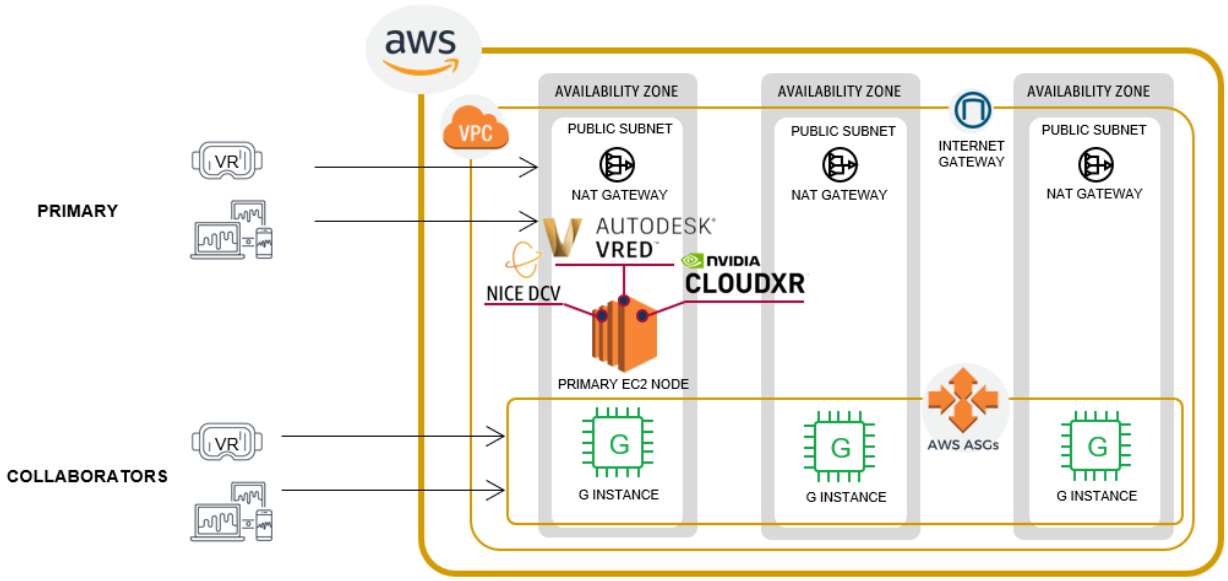

https://aws-quickstart.github.io/quickstart-nvidia-cloudxr/

This Quick Start deploys VRED and CloudXR on the AWS Cloud. If you are unfamiliar with AWS Quick Starts, refer to the AWS Quick Start General Information Guide.

Autodesk VRED 3D visualization software **helps designers and engineers create product presentations, design reviews, and virtual prototypes** using interactive GPU ray tracing and both analytic and cloud-rendering modes.

This Quick Start is for IT infrastructure architects, administrators, and DevOps professionals who are planning to implement or extend their Autodesk VRED workloads on the AWS Cloud. There is no cost to use this Quick Start, but you will be billed for any AWS services or resources that this Quick Start deploys.

Understand the Amazon EC2 G5 benefits

High performance for graphics-intensive applications

- G5 instances deliver up to 3x higher graphics performance and up to 40% better price performance than G4dn instances. They have more ray tracing cores than any other GPU-based EC2 instance, feature 24 GB of memory per GPU, and support NVIDIA RTX technology. This makes them ideal for rendering realistic scenes faster, running powerful virtual workstations, and supporting graphics heavy applications at higher fidelity.

AWS NVIDIA A10G Tensor Core GPU

- G5 instances are the first in the cloud to feature NVIDIA A10G Tensor Core GPUs that deliver high performance for graphics-intensive and machine learning applications. Each instance features up to 8 A10G Tensor Core GPUs that come with 80 ray tracing cores

and 24 GB of memory per GPU. They also offer 320 third-generation NVIDIA Tensor Cores delivering up to 250 TOPS resulting in high performance for ML workloads.

NVIDIA drivers

- G5 instances offer NVIDIA RTX Enterprise and gaming drivers to customers at no additional cost. NVIDIA RTX Enterprise drivers can be used to provide high quality virtual workstations for a wide range of graphics-intensive workloads. NVIDIA gaming drivers provide unparalleled graphics and compute support for game development. G5 instances also support CUDA, cuDNN, NVENC, TensorRT, cuBLAS, OpenCL, DirectX 11/12, Vulkan 1.1, and OpenGL 4.5 libraries.

High performance networking and storage

- G5 instances come with up to 100 Gbps of networking throughput enabling them to support the low latency needs of machine learning inference and graphics-intensive applications. 24 GB of memory per GPU along with support for up to 7.6 TB of local NVMe SSD storage enable local storage of large models and datasets for high performance machine learning training and inference. G5 instances can also store large video files locally resulting in increased graphics performance and the ability to render larger and more complex video files.

Built on AWS Nitro System

- G5 instances are built on the AWS Nitro System, which is a rich collection of building blocks that offloads many of the traditional virtualization functions to dedicated hardware and software to deliver high performance, high availability, and high security while also reducing virtualization overhead.

Be aware of an estimated TCO (Total Cost of Ownership)

This use-case is prevalent across most automotive (consumer and commercial), industrial equipment and aerospace organizations. Each design program is a candidate for this workload. We predict that each brand/customer would leverage this solution on a regular basis, supporting multiple, parallel design programs for new and updated products. For example, for each design program, review sessions might occur 3 time per week and are typically 1-3-hour blocks with 5-8 participants. Projected costs would be ($4.5k - $6k+ per year per customer based on ~$3/hr for G5 x 5 participants x 2 hours x 156 sessions per year).