Whitepaper

Delivering 3D Virtual Maintenance Training Content: Examining the Deployment Options from the Classroom to the Cloud

Version:

1.1

Released:

19-Jul-2016

Purpose:

This paper focuses on the latest options available for delivering 3D virtual environments to the training consumer along with the pros and cons of each option and key lessons learned for two different types of training consumers; Classroom users and External users.

For Support or Questions:

Name: Scott Ariotti
Email: sariotti@disti.com
Phone: +1 407 206 3390 ext 125

ABSTRACT

Over the past several years, the U.S. Armed Forces have been significantly expanding their adoption of virtualized solutions for use in maintenance training applications. This growing adoption is coupled with a growing interest in expanding the ways in which these training materials are consumed. Unlike traditional operational trainers that teach students on fixed hardware, the virtualized training material for maintainers lends itself well to providing innovative training beyond typical brick and mortar schoolhouses; so long as the content can be effectively delivered. These graphically intense 3D virtual environments often levy hefty requirements on the type of computer capable of delivering an immersive interactive 3D experience (see Figure 1). This paper focuses on the latest options available for delivering these virtual environments to the training consumer along with the pros and cons of each option and key lessons learned for two different types of training consumers; Classroom users and External users.

…virtualized training material for maintainers lends itself well to providing innovative training beyond typical brick and mortar schoolhouses…

The deployment options in this paper compare and contrast traditional desktop use with technologies that incorporate mobile client applications, fixed server rendering solutions, and newly emerging cloud-based application rendering services. The review includes relative cost comparisons, barriers to entry, the consumer access experience, application development considerations, and necessary hardware utilizing real-world examples that encompass both military aircraft and commercial automotive training devices. The paper also introduces discussion topics on information assurance and security considerations for each deployment option.

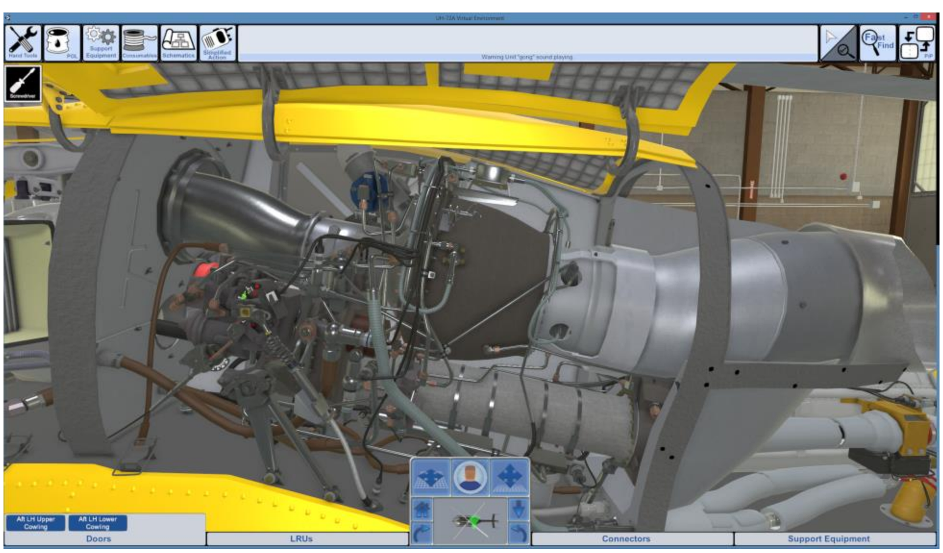

Figure 1. Highly detailed engine bay of an EC145

INTRODUCTION

For nearly a decade, the U.S. Armed Forces have been significantly expanding their adoption of virtualized solutions for use in maintenance training applications. This growing adoption is coupled with a growing interest in expanding the ways in which these training materials are consumed.

Unlike a traditional 3D virtual environment where detail is spread out over a large gaming area, the virtual environments for maintenance trainers tend to be graphically intense 3D environments packed into a single mechanical unit. As a result, a VMT application often levies hefty requirements on the type of computer capable of delivering the immersive interactive 3D experience. This factor relegated VMT instructional material to fixed hardware systems capable of running these applications to the confines of brick and mortar schoolhouses. However, the increasing desire to engage consumers without classrooms, instructors, and training facilities by expanding options for when and how we learn, referred to as "Pervasive Learning" (Freed, et al., 2014), are driving the demand to do the same with VMT applications.

This paper focuses on the latest options available for delivering these virtual environments to the training consumer along with the pros and cons of each option and key lessons learned for two different types of training consumers; Classroom users and External users.

TRAINING CONSUMER TYPES

In this paper we consider two types of VMT training consumers; "Classroom" users, and "External" users. Classroom users are generally brought together at a brick-and-mortar establishment for the purpose of training and/or certification (see Figure 2) (Roganov, 2012). Classroom users train at the facility for a limited time and receive a structured set of training that needs to be completed in the allotted timeframe. Classroom users often receive some type of team training on hardware or virtual assets that mimic the real world environment. An instructor is nearly always available and is likely deeply involved in the student training.

Figure 2. Members of the 31st Test & Evaluation Squadron at Eglin AFB use the F-35 Aircraft Systems Maintenance Trainer

External users, by contrast, are not gathered at a training facility, but instead are distributed at their various workplaces or homes. They may be studying new material, satisfying recurring training requirements, doing just-in-time training, or performing prerequisite training before attending classroom training. They could also be doing free-form exploration of the target equipment using a VMT.

DEPLOYMENT TYPES

When interactive 3D VMTs were first introduced, the only viable deployment option was to provide the training as part of a fixed stand-alone trainer system (Giordano, Jackson, & Blankemeir, 2012). These devices have been implemented as one or more independent trainers, or as multiple desktop devices within an electronic classroom, networked to an Instructor Station via a local area network (LAN) connection. These early VMTs may also have been integrated within courseware applications as Level 3 or 4 courseware assets.

For the purposes of this paper, consider this stand-alone trainer system as a standard classroom environment. For the traditional Classroom user, the traditional classroom environment is a perfect match.

With more need to satisfy External VMT Users, something else is needed. Today there are new options which make delivery to the External users possible. The options available now include:

- Desktop
- Mobile
- Centralized Server
- Cloud Server
- Web

DESKTOP

Desktop delivery of VMTs isn't really new, but is included here as a reference point. Desktop applications are created to run on Microsoft Windows® and sometimes Linux®. They are installed on the local hard drive of the target machine (or possibly a shared LAN drive). This target machine is typically a high-end desktop computer purchased specifically for delivering the 3D training content. The computer is likely to have a high-end graphics card for maximum performance and carry a price tag in the range of US$1,000 to US$2,000 per computer. The installation process for the 3D training application can be time consuming and typically requires administrator privileges.

As a result, Desktop applications are rarely sent to External users because of these high-end system requirements and possible data security concerns. In cases where the application needs to be deployed to External users with lower-end computers, the application would need to be developed at a lower graphical fidelity level, or the classroom and external applications would be generated at two different fidelity levels, resulting in higher cost for development, upgrades, and configuration management.

Desktop VMT deployments are relatively expensive per user because of the required hardware. The issues of getting custom software approved and installed in most environments today is a significant task.

Mobile

Mobile platforms include smart phones, tablets, and phablets. These can be a good target for capturing people's free time at home or other idle times where training can occur. External Users are the primary consumers of mobile applications. However, there is growing interest for using tablets for Classroom Users as a mechanism to provide freedom of movement between the classroom and hardware-based training labs. This is evident by the technical requirements for the U.S Army Gray Eagle VMT and UH-72A Lakota VMT programs (United States Army PEO STRI, 2015).

In addition to the small, portable form factor, for some use-cases, mobile platforms offer the ability to provide simple augmented reality capabilities for the training application. The application can benefit from the built-in camera and accelerometers that most mobile platforms offer. Figure 3 shows an Augmented Reality example running on a tablet. These types of applications can help with on-the-job and just-in-time training, where the mobile device is being used in conjunction with the real equipment on a real task.

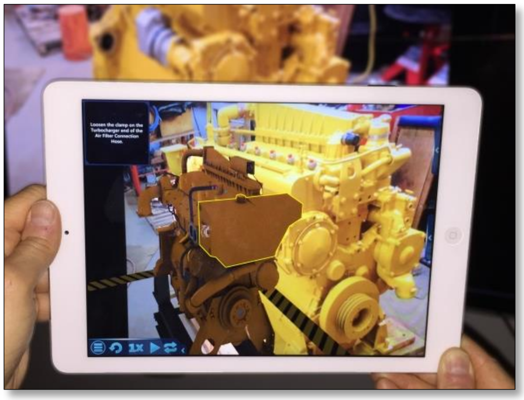

Figure 3. Augmented Reality Engine on Tablet

A big challenge with mobile platforms is to provide an acceptable level of interactivity. The small screens, (especially with phones), limit the types of interactivity that people can reliably perform. It is sometimes necessary to reduce the level of interactivity to keep the engagement level high. Figure 4 shows an example of a maintenance training application running on a tablet. The high level of interactivity would be difficult to support in a phone form-factor.

Figure 4. User Interface Elements on Tablet

The limited graphics rendering performance and memory capacity of some of the mobile platforms, compared to desktop graphics capabilities, constrains the complexity and possible precision of the VMT. This is often acceptable for games, but VMTs often need to maintain a high degree of accuracy and fidelity to avoid distraction and possible negative training.

Centralized Server

One solution to the limited performance of mobile devices and low performance PCs is to move the graphics rendering to a Centralized Server. This allows the Centralized Server to render the high fidelity 3D graphics and then send just the resulting pixels to the target devices. For Classroom users, the

Centralized Server can be located in the classroom or somewhere nearby. The Centralized Server is built to support a fixed maximum number of users.

Sharing Centralized Servers is nothing new for non-graphical applications. However, until a few years ago, no good low-latency solution existed for shared graphics rendering. Now there are a new generation of graphics cards available that facilitate a server rendering solution. NVIDIA offers a product under the GRID™ brand and AMD offers the Radeon™ Sky Series line of graphics cards that help make this possible.

These cards are designed to be installed in rack mounted PCs and work well without a monitor. Instead of outputting the video to an HDMI or DVI port they compress the signal and output it to an H.264 (or other compressed) video stream. This video stream is sent over the network to the target device (NVIDIA Corporation). In this way, many different consumer platforms can be supported by a single application (See Figure 5).

Figure 5. Various Platforms Consume Application Video Stream & Return Input Events

The target device requires an application to decode the video stream. While NVIDIA and AMD provide software tools to accomplish this, software products such as Citrix XenApp™ and VMware® Horizon™ can be used to take advantage of these features.

This type of solution works well for a fixed number of simultaneous users that are co-located, such as the Classroom users on a LAN. The number of simultaneous users (streams) per card and number of cards per server varies based on the model of the card, the type of server, the supporting software products, and the intensity of the 3D application being served. In a 2014 usability test, we configured a Centralized Server with the hardware and software listed in Table 1. In the test, we were able to host up to four (4) simultaneous users on the Centralized Server. Performance degradations with six (6) simultaneous users rendered the system difficult to use. Since conducting this usability testing, Citrix has updated the XenApp software which is reported to improve graphics-core sharing performance amongst multiple users.

Table 1. Centralized Server Usability Test Configuration

Quantity	Item
1	Xeon 1620W 4U Rack-mount Workstation
1	Intel Xeon Six-Core Processor E5-2640 2.5GHz 7.2GT/s 15MB LGA 2011 CPU
4	Samsung DDR3-1600 16GB ECC/REG Samsung Chip Server Memory
2	NVIDIA GRID K520
	Windows Server 2008 R2 SP1
	XenApp 6.5
	Medium Density Part-task Level VMT App

Providing access for External users on this Centralized Server approach over a Wide Area Network (WAN) quickly degrades the viability of this solution. This degradation occurs for a variety of reasons. First, the number of simultaneous users is difficult to predict and manage. Server hardware and software would need to be purchased to handle the estimated peak usage. This can be an expensive proposition, especially since all the servers have to be hosted at a fixed location, be powered-on and available 24/7, configured, networked, maintained, etc. These costs apply even when no External users are connected.

Second, this type of deployment option is subject to video latency issues. The further away an External user is from a Centralized Server the greater the lag they will experience when running the application (Figure 6). As latency exceeds 100ms trainees are likely to observe noticeable delays that could prove to be distracting to the task at hand, depending on the specific training application (Monson, Johnston, & Barnhart, 1997). The only option to combat this latency is to stand up another Centralized Server closer to the External user's geographic location adding to the deployment costs.

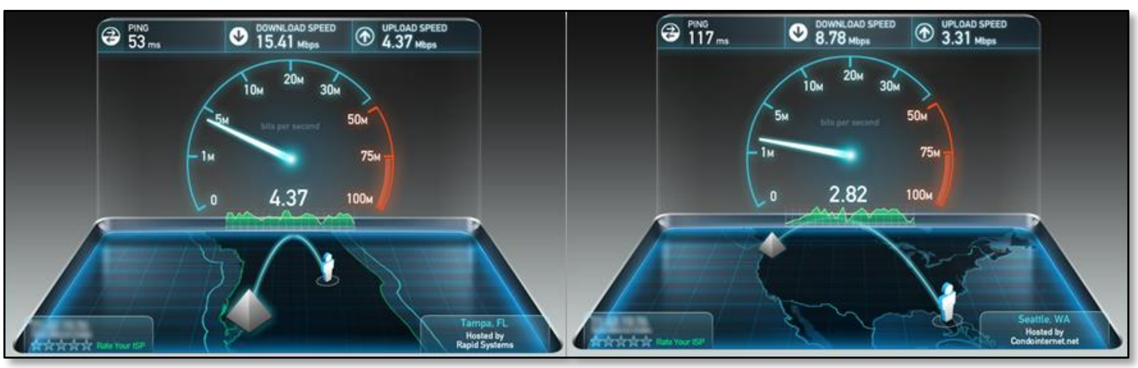

Figure 6. Ping times between Orlando and Tampa (53ms) and Orlando and Seattle (117ms)

Cloud Server

An alternative to the Centralized Server is to put the VMT application in the cloud. Instead of purchasing and hosting hardware on premises, developers upload the application to a Cloud Service Provider (CSP) which hosts the application on their servers. Developers may opt for using a fixed number of always available servers, or by using an automated scaling system that creates new server instances as needs demand. In either case, the Cloud Server or Centralized Server, External users are required to have a high-speed Internet connection. The Internet connection requirement is similar to that for streaming a single 1920x1080 HD movie; 5 Mbps (Netflix Corporation).

Using the option that offers a fixed number of servers is very similar to the Centralized Server solution, but the effort of maintaining hardware servers is transferred to the CSP. Without automated scaling, developers must choose the number of servers to meet the peak demand. Depending on the CSP configuration options, it is possible to mix fixed instances and on-demand instances for a custom solution that minimizes cost.

The need for a software package to decode the application's video stream for the user exists in this method as well. The previously cited examples in the Centralized Server section pertain here. Figure 7 shows a VMware serving a graphics application from an NVIDIA Test Drive site (NVIDIA Corporation).

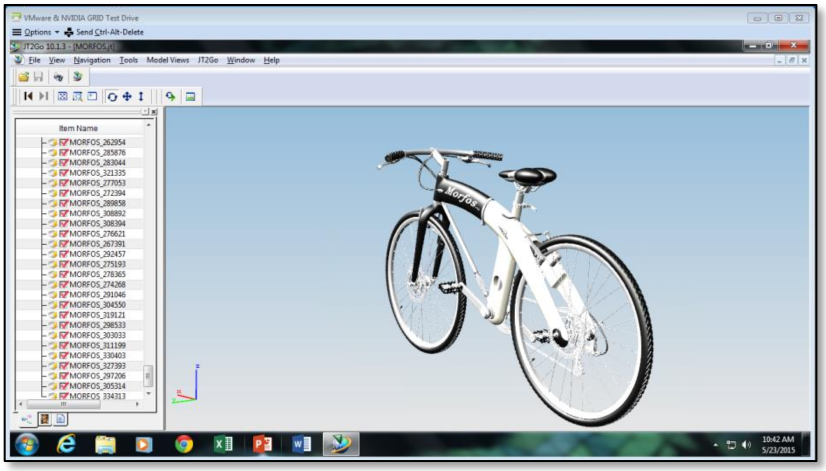

Figure 7. VMware & NVIDIA Test Drive

This means that most of these solutions require a vendor supplied client to be installed on whatever device needs to access the VMT application. While these clients have to be vetted by the appropriate approval authorities prior to their use, they can access any VMT application published to the vendor's cloud service.

The high availability, low cost, and low maintenance make a cloud solution an attractive option for serving External users and possibly even Classroom users. Also, with most of the solutions, updating the VMT software occurs in one place. So deployment of upgrades can be trivial. The new VMT is uploaded and new logins automatically access it. There is no concern for people using outdated training versions. Companies such as Nissan Motors are currently using Amazon's CSP AppStream service to successfully deploy VMT applications to mechanics using low-end laptops in their residences and dealerships (Figure 8) (Barr, 2014).

The tight control over the access and the content, and the scalability also make a Cloud Server solution a good choice to implement Training as a Service (TaaS). This even allows companies to create and deploy content and be compensated only when it is accessed.

Figure 8. The Nissan Virtual Training Solution Deployed Through Cloud Servers

WEB

Over the years, there have been many attempts at a solution for 3D on the Web. Most of these solutions took the form of a browser plugin maintained by a single company. Examples include Flash, Java, Silverlight, and Unity. While these solutions have contributed to expanding the capabilities of 3D on the Web, they have the disadvantage of relying on a single company for updates and at the same time become singular targets for hacking and security holes (Constantin, 2015). Today, the WebGL standard is natively supported on all major browsers so no plug-in is required to run 3D content.

Initially released in 2011, WebGL provides a low level interface to the graphics subsystem (Khronos Group). This allows for high performance 3D graphics on the Web across every major modern desktop and mobile browser. Even the Apple iOS devices support WebGL (See Figure 9). Since the application

Figure 9. WebGL 3D Graphics from a web site on iOS

rendering occurs on the local machine (not the server hosting the WebGL content) there are natural limitations on application size when executing on mobile devices and computers without accelerated graphics. However, these problems can be managed per project as necessary based on the requirements.

With the WebGL standard, it is now possible to deploy VMTs directly to the web, with no plugin reliance or platform dependencies. This can be ideal for smaller VMTs and part-task-trainers. Because the content is just a standard web page, it can be easily included in Learning Management System (LMS) servers without requiring modifications or any plugin approvals (See Figure 10). This WebGL content is available at http://www.disti.com/products/gl-studio/gl-studio-online-demonstrations/.

Figure 10. The same web site as Figure 9, running on a desktop

Deployment Summary

Table 2 summarizes some of the Pros and Cons of each deployment type. Of course, the final choice of deployment type is complex and based on the requirements of the end user.

Table 2 - Deployment Type Pros/Cons

Deployment Type	Pros	Cons
Desktop	• Highest performance • Supports unique interface devices • Traditional approach	• Requires specific hardware • Requires admin privileges to install • Limited accessibility
Mobile	• Train anywhere • Offers Augmented Reality support • Native apps do not require Internet (WAN) connectivity once installed	• Detailed interactions can be difficult • Performance varies across devices • Desktop UI elements may not scale for mobile use, and will need to be customized for the smaller devices
Centralized Server	• Hardware is controlled by the project • No Internet (WAN) connectivity required • Trivial software updates	• Configuration and maintenance • Must be sized for peak access • Not scalable for External users
Cloud Server	• Use virtually any client hardware • Only pay for usage • Full control over current software version • Instant-on • Trivial software updates	• Requires about 5 Mbps Internet connection • Client software install required
Web	• No software install required • WebGL supported on all modern browsers • Transparent LMS integration • Trivial software updates	• Content size limitations • Subject to graphics power on device • Older browsers not supported

INFORMATION ASSURANCE

Cloud-based computing is not new for the DoD. In 2012 the DoD issued its Cloud Computing Strategy from the DoD Chief Information Officer stating that the CIO is committed to accelerating the adoption of cloud computing within the DoD in "achieving IT efficiencies, reliability, interoperability, and improve security and end-to-end performance by using cloud service offerings." (Takai, 2012).

In December 2014 the DoD CIO memo regarding *Updated Guidance on the Acquisition and Use of Commercial Cloud Computing Services* defines DoD Component responsibilities when acquiring cloud services (DoD Defense Information Systems Agency (DISA), 2015). Approved Cloud Service Providers are listed on the Federal Risk and Authorization Management Program (FedRAMP) website. FedRAMP is a Federal Government program focused on enabling secure cloud computing for the Federal Government and provides a standardized approach to security assessment, authorization, and continuous monitoring for cloud services by incorporating the Federal Government RMF processes. The CSPs are categorized as providers of Infrastructure as a Service (IaaS), Software as a Service (SaaS), or Platform as a Service (PaaS).

Coupling a Compliant CSP's offering with the Mission Owner's (MO) system or application paves the way for the DoD Designated Accrediting Authority (DAA) to grant a mission approval to operate (ATO) (See Figure 11) (Bockelman & McDermott, 2015).

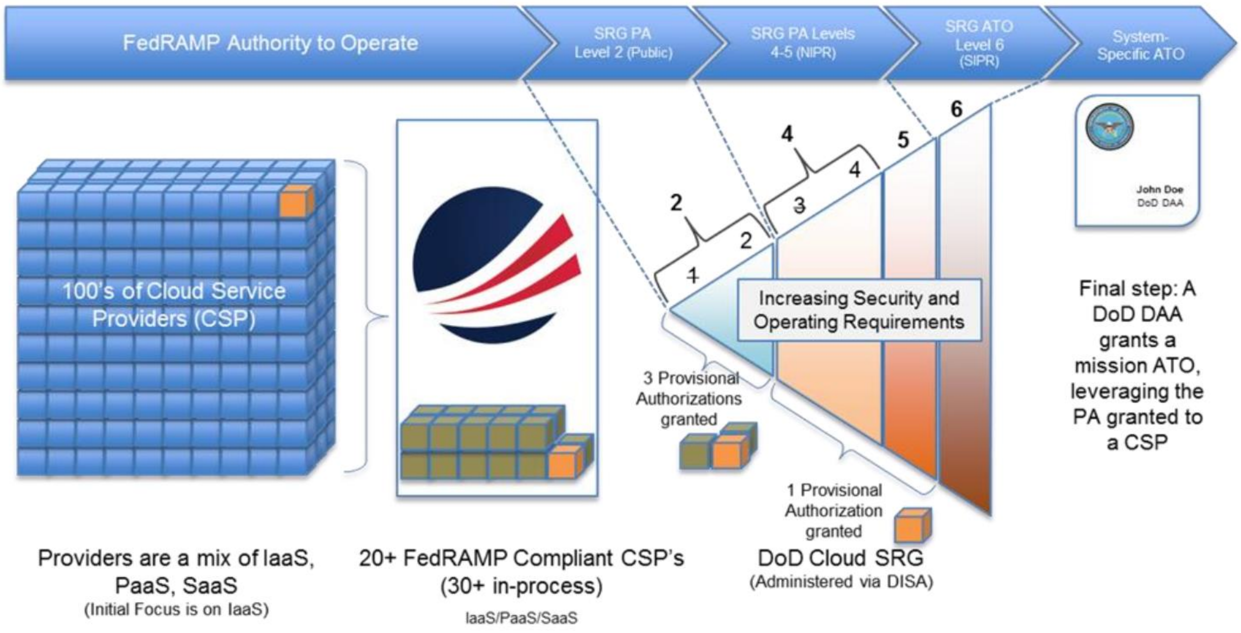

Figure 11. Compliant Cloud Service Provider's Authorization to Proceed Workflow

VMT applications will need an IaaS CSP and generally fall under an information impact level 4 for Controlled Unclassified Information since most VMT training applications are unclassified and are export controlled. Therefore the DoD Cloud Security Requirements Guide (SRG), administered by DISA would serve as the provisional authority.

However, to date, none of the IaaS Compliant CSPs on FedRAMP indicate a video streaming solution in their service descriptions. While Amazon Web Services (AWS GovCloud) is a Compliant CSP on FedRAMP, the AppStream solution currently uses some services that are not contained within the AWS GovCloud umbrella.

CONCLUSION

Deployment obstacles for delivering advanced 3D Virtual Maintenance Training applications must be overcome in order to continue the expansion of the VMT domain. The technologies that have become available in the last two years have significantly expanded the options for delivering the training content.

For Classroom users, Desktop deployment will continue to be a common and reasonable way to use VMT applications within the brick-and-mortar school houses. They may also opt for Centralized Servers for distributing training content on tablets to subsidize the classroom experience.

For serving External users, WebGL will increasingly be seen as a powerful deployment medium while Cloud Servers offer a low cost easy-to-implement strategy for delivery of the applications on arbitrary hardware. Using WebGL and/or Cloud Servers for VMTs will change what people expect from a web-sourced experience.

Lastly, the challenges of information assurance are not insurmountable since the best argument can be made that target devices receive nothing more than a video stream from Cloud Servers. If there is a high enough demand to deliver content in this fashion, then policy makers, industry members, and government personnel will find common ground.

ACKNOWLEDGEMENTS

The contents of this whitepaper were initially authored by DiSTI and published in the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2015 proceedings. I/ITSEC is the world's largest modeling, simulation, and training conference and consists of peer-reviewed paper presentations, tutorials, special events, professional workshops, a commercial exhibit hall, a serious games competition, and STEM events for teachers and secondary students.

Additional information on the conference is available at www.iitsec.org. Copies of papers published by the conference are available at www.iitsecdocs.com.

CITATIONS

Barr, J. (2014, September 30). *Amazon Appstream is Now Available in Japan*. Retrieved from Amazon Web Services: https://aws.amazon.com/blogs/aws/appstream-in-japan/

Bockelman, P., & McDermott, A. (2015, April). DoD-Compliant Implementations in the AWS Cloud. Seatle, WA, USA.

Constantin, L. (2015, February 2). *Flash Player faces its third zero-day flaw in a month*. Retrieved May 18, 2015, from Computer World: http://www.computerworld.com/article/2877986/flash-player-faces-its-third-zero-day-flaw-in-a-month-updates-coming.html

DoD Defense Information Systems Agency (DISA). (2015, January 12). *DEPARTMENT OF DEFENSE (DoD) CLOUD COMPUTING SECURITY REQUIREMENTS GUIDE (SRG).* Retrieved June 8, 2015, from DISA website: http://iase.disa.mil/cloud_security/Documents/u-cloud_computing_srg_v1r1_final.pdf

Freed, M., Yarnall, L., Dinger, J., Gervasio, M., Overholtzer, A., Perez-Sanagustin, M., . . . Spaulding, A. (2014). PERLS: An Approach to Pervasive Personal Assistance in Adult Learning. *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC).* Orlando: NTSA.

Giordano, C., Jackson, S., & Blankemeir, J. (2012). Technology Enhancements & Innovations for Virtual Maintenance Training. *Singapore Aerospace Technology and Engineering Conference (SATEC).* Singapore: Singapore Institute of Aerospace Engineers.

Khronos Group. (n.d.). *WebGL Security*. Retrieved May 13, 2015, from Khronos Group Web site: https://www.khronos.org/webgl/security/

Monson, S. L., Johnston, R. R., & Barnhart, D. J. (1997). Latency - The Adversary of Real-Time Distributed Simulation. *Interservice/Industry Training Simulation and Education Conferece (I/ITSEC).* Orlando, FL: NTSA.

Netflix Corporation. (n.d.). *Internet Conncection Speed Recommendations*. Retrieved 02 2015, from Netflix Corporation Help Center: https://help.netflix.com/en/node/306

NVIDIA Corporation. (n.d.). *The Power of Cloud Gaming.* Retrieved 02 2015, from NVIDIA Corporation Website: http://www.nvidia.com/object/cloud-gaming.html

NVIDIA Corporation. (n.d.). *VMware & NVIDIA GRID Test Drive*. Retrieved May 2015, from NVIDIA Corporation: http://www.nvidia.com/object/vmware-trygrid.html

Roganov, K. (2012, April 11). *F-35 training center begins formal training*. Retrieved May 30, 2015, from Eglin Air Force Base: http://www.eglin.af.mil/news/story.asp?id=123297561

Takai, T. M. (2012). *Cloud Computing Strategy.* Washington D.C.: Department of Defense.

United States Army PEO STRI. (2015). Solicitation Number W900KK-15-R-0004. Orlando, FL, USA.

NOTICE

ALL INFORMATION PROVIDED IN THIS WHITE PAPER, INCLUDING BUT NOT LIMITED TO COMMENTARY, OPINION, DiSTI DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." THE AUTHORS MAKE NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, the authors assume no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of the authors. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. The DiSTI Corporation's products are authorized for use as critical components in life support devices or systems only with the use of the GL Studio Safety Critical runtime libraries and certification kit available from the DiSTI Corporation.

Trademarks

DiSTI, the DiSTI logo, GL Studio, VE Studio are trademarks or registered trademarks of the DiSTI Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.