

Intergraph™ SmartPlant Review (64 bit)

Citrix XenApp with Azure Virtual Machines

Contents

1. Abstract	3
2. Background	4
2.1 Citrix XenApp.....	4
2.2 Microsoft Azure.....	4
2.3 Intergraph SmartPlant Review.....	4
2.4 CPU Rendering	4
2.5 Threads (Tessellation).....	4
2.6 Vertical Buffer Objects (VBOs).....	5
3. System Configuration and Environment	5
3.1 Machine Specifications.....	5
GPU-Enabled Citrix XenApp Server	5
GPU-Disabled Citrix XenApp Server	5
Citrix XenApp Delivery Controller	5
Citrix Client Workstation.....	5
3.2 SmartPlant Review 64-bit Special Configuration.....	6
3.3 Testing Environment.....	8
4. Functional Multi-User Testing.....	8
4.1 Setup	8
4.2 Procedure	9
4.2.1 Stability Testing.....	9
4.2.2 Performance Testing.....	9
4.3 Results.....	9
4.3.1 Stability Results.....	9
4.3.2 Performance Results	10
4.3.3 Ad-Hoc Testing	11

1. Abstract

This effort investigates different areas of setup, configuration and performance for the cloud-based utilization of Citrix XenApp and Intergraph SmartPlant Review using 64-bit architecture. Virtual machines hosted in Azure are available with a Graphics Processing Unit (GPU). Multiple areas of the product were tested. On premises, physical machines were used to access the Azure cloud system for all tests.

Testing was divided into two distinct sections:

- The first section was to determine maximum stability and user density on a single server. Several configuration options were reviewed and verified in both single and multi-user scenarios. When we determined a configuration that produced the maximum trade-off between stability and user density, we administered further tests.
- The second section focused on performance at user density. We tested performance with and without a GPU-backed virtual machine. Performance was tested using a graphically dense test case roughly 2 GB in size. An API test harness was used to display various views in the model while logging information based on the length of time it took to display each view.

2. Background

2.1 Citrix XenApp

Citrix XenApp allows an application to be hosted on a server machine, and then publishes that application to a client machine. All processing and execution for the application occurs on the Citrix server, while the client only needs to display the graphical data sent over the network from the Citrix server. Since application load is handled by the server, client PC hardware does not have to be as robust.

2.2 Microsoft Azure

Microsoft Azure is a cloud computing platform that offers a variety of cloud-based services. The virtual machines used for this testing were hosted on a regional Azure server.

2.3 Intergraph SmartPlant Review

SmartPlant Review is a visualization environment for reviewing large, complex 3D models. SmartPlant Review takes advantage of graphics hardware and system RAM to load larger models faster. The 64-bit versions of SmartPlant Review can open very large models while remaining responsive to the end user.

2.4 CPU Rendering

If no GPU is available, SmartPlant Review defaults to Software mode or rendering, and offloads all graphics processing to the CPU. This mode is typically less performant than Hardware mode and is not recommended.

2.5 Threads (Tessellation)

SmartPlant Review is configured by default to access multiple tessellation threads when processing data and displaying it to the screen. This is optimized for out-of-box, single-user scenarios. When optimizing 64-bit SmartPlant Review for multi-user and cloud configurations, the maximum number of tessellators should be restricted to 1.

2.6 Vertical Buffer Objects (VBOs)

SmartPlant Review (64-bit) is configured by default to use up to 100% of the available graphics card RAM. When operating in a multi-user and cloud configuration, this number should be divided by the number of planned simultaneous users.

3. System Configuration and Environment

3.1 Machine Specifications

GPU-Enabled Citrix XenApp Server

An Azure NV6 with GPU was configured with the following specifications:

- CPU: Intel Xeon E5-2690 v3 (6 cores at 2.6 GHz)
- GPU: NVidia Tesla M60, Driver version 362.56
- Memory: 56 GB
- Disk: Standard Storage
- OS: Windows Server 2016

GPU-Disabled Citrix XenApp Server

An Azure NV6 without GPU was configured with the following specifications:

- CPU: Intel Xeon E5-2690 v3 (6 cores at 2.6 GHz)
- Memory: 56 GB
- Disk: Standard Storage
- OS: Windows Server 2016

Citrix XenApp Delivery Controller

An Azure DS3 without GPU was configured with the following specifications:

- CPU: Intel Xeon E5-2690 v3 (4 cores at 2.6 GHz)
- Memory: 14 GB
- Disk: Premium Storage – P10
- OS: Windows Server 2016

Citrix Client Workstation

The Citrix Client workstations used during the test were Dell Precision T1600 chassis. There were 12 available with the following specifications:

- CPU: Intel Xeon E31225 (4 cores at 3.1 GHz)
- GPU: NVidia Quadro 2000

- Memory: 16 GB
- Disk: Hard Disk Drive
- OS: Windows 7 Enterprise with SP1

3.2 SmartPlant Review 64-bit Special Configuration

When configuring 64-bit versions of SmartPlant Review in the environment outlined in this document, two changes needed to be made to SPRDisplaySettings.ini file. This file is located in the SmartPlant Review Installation folder, and can be edited with a standard text editor such as Notepad.

- In the section titled [**ViewSettings**] locate the following line:

VBOMemoryPercent = 100

and replace **100** with a value from the chart below. This value is obtained by dividing 100 by the number of simultaneous users.

Number Of Simultaneous Users	VBO Setting
1	100
2	50
3	33
4	25
5	20

For example, if you have five users accessing a single virtual machine at one time ($100/5=20$), the following setting is recommended:

VBOMemoryPercent = 20

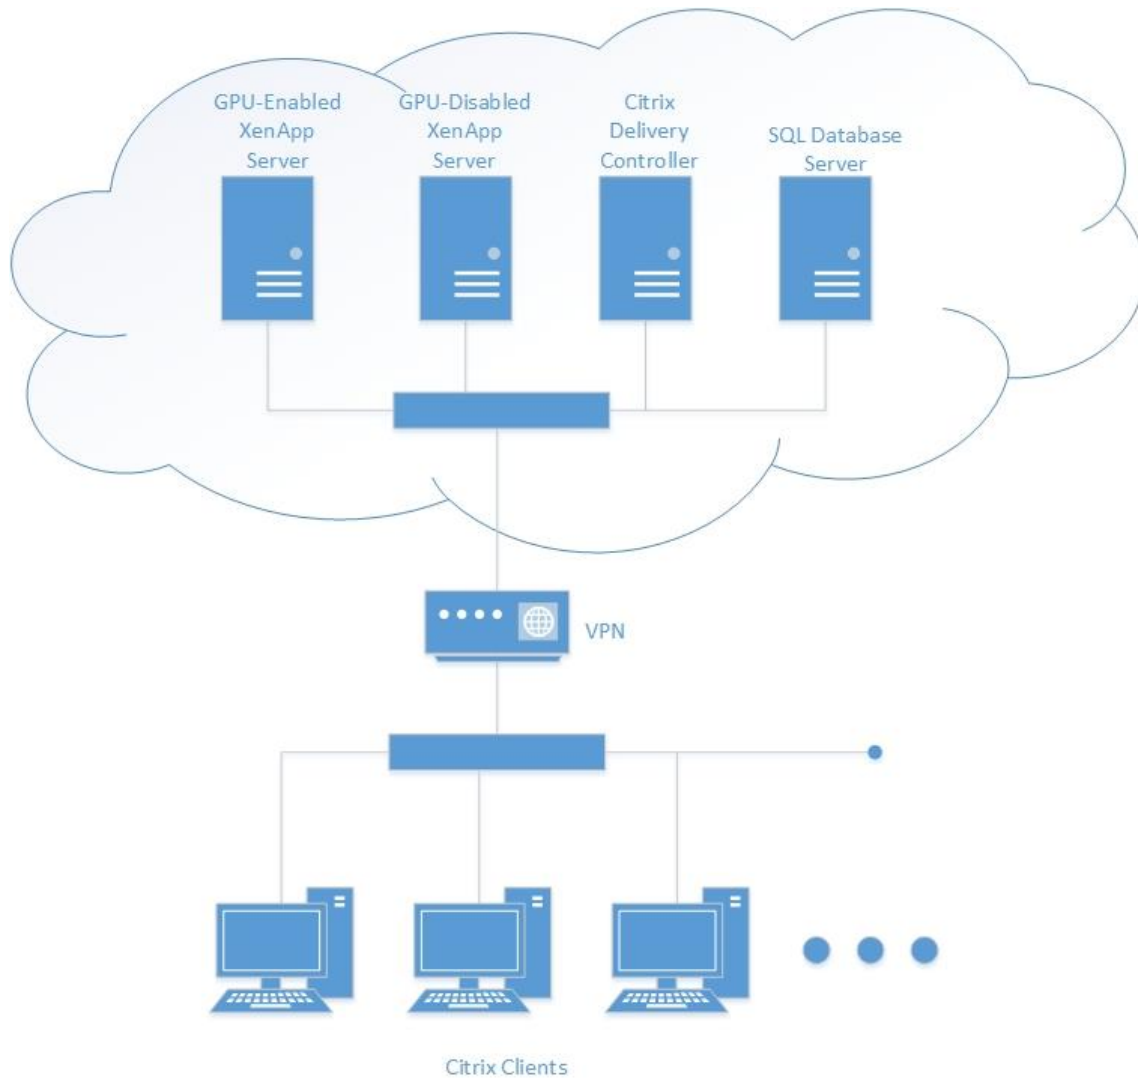
- Directly beneath this line, add the following:

MaxTessellatorCount = 1

```

*C:\Program Files\SmartPlant\Review\SPRDisplaySettings.ini - Notepad++ [Administrator]
File Edit Search View Encoding Language Settings Macro Run Plugins Window ?
SPRDisplaySettings.ini new 2
1  [ViewSettings]
2  BlendedTransparencyCullBackfaceAspectMask = 32
3  BlendedTransparencyCullEndCapAspectMask = 32
4
5  ;VBOMemoryPercent - Set the maximum percentage of available card memory for
6  ;    caching vertex data in VBOs. This number may be greater than 100 to
7  ;    take advantage of the page file for storage of our vertex cache.
8  VBOMemoryPercent = 20
9  MaxTessellatorCount = 1
10
11 [Intel]
12 DoNotUseHardwareAcceleration = 1
13
14 [ATI Technologies Inc.]
15 VBOBufferSize = 524288
16
17 [VMware, Inc.]
18 DoNotUseHardwareAcceleration = 1
19
    
```

3.3 Testing Environment



4. Functional Multi-User Testing

4.1 Setup

Testing used five Citrix client desktops connecting to a single XenApp Server.

The first goal was to find the number of users that could connect to the XenApp server and use SmartPlant Review with reasonable stability.

A unique test case for each login was generated from Smart 3D. The test case was 2.2GB of data comprised of 36 VUE files that were opened as a Streaming VUE File (SVF).

The client had dual-screen monitors set up so that the resolution for each was 1920x1080. The test maximized SmartPlant Review to display across both screens.

The average network latency to the XenApp servers was 54 milliseconds.

4.2 Procedure

4.2.1 Stability Testing

Each round of testing involved modification of the `MaxTessellatorCount` setting and the `VBOMemoryPercent` setting within the `.ini` file. Each pass was done with and without the GPU enabled on the server. Iterations involved one, two, three, four and finally five simultaneous users.



The `MaxTessellatorCount` setting proved to be unstable at values proportional to the number of CPU virtual threads and users. If the value of `MaxTessellatorCount` setting multiplied by the number of users was greater than the number of virtual CPU threads, SmartPlant Review began to show signs of instability. When this setting was verified, the `MaxTessellatorCount` setting was set to the recommended value of 1 for all further tests.

Then focus was turned to the `VBOMemoryPercent` setting. Testing discovered that at this time, the value of this setting should be manually set proportional to the number of users accessing the VM. It was tested at various settings on various levels to come up with the recommendations outlined above.

4.2.2 Performance Testing

When the optimal settings were discovered for managing stability, focus moved to performance of SmartPlant Review and scaling up the number of clients that connect to a single VM. To achieve this, 400 **Fit** commands were called starting in one corner of the model, and moving to the far corner, then returning. Each **Fit** aimed the camera toward model-center. The average time of each Fit operation and the total time were monitored along with server RAM, CPU, and GPU performance.

Testing between **Balanced** and the default performance was compared. Intergraph generally recommends that customers adjust the performance slider to **Balanced** when using SmartPlant Review in multi-user, cloud-based applications for the best performance.



4.3 Results

4.3.1 Stability Results

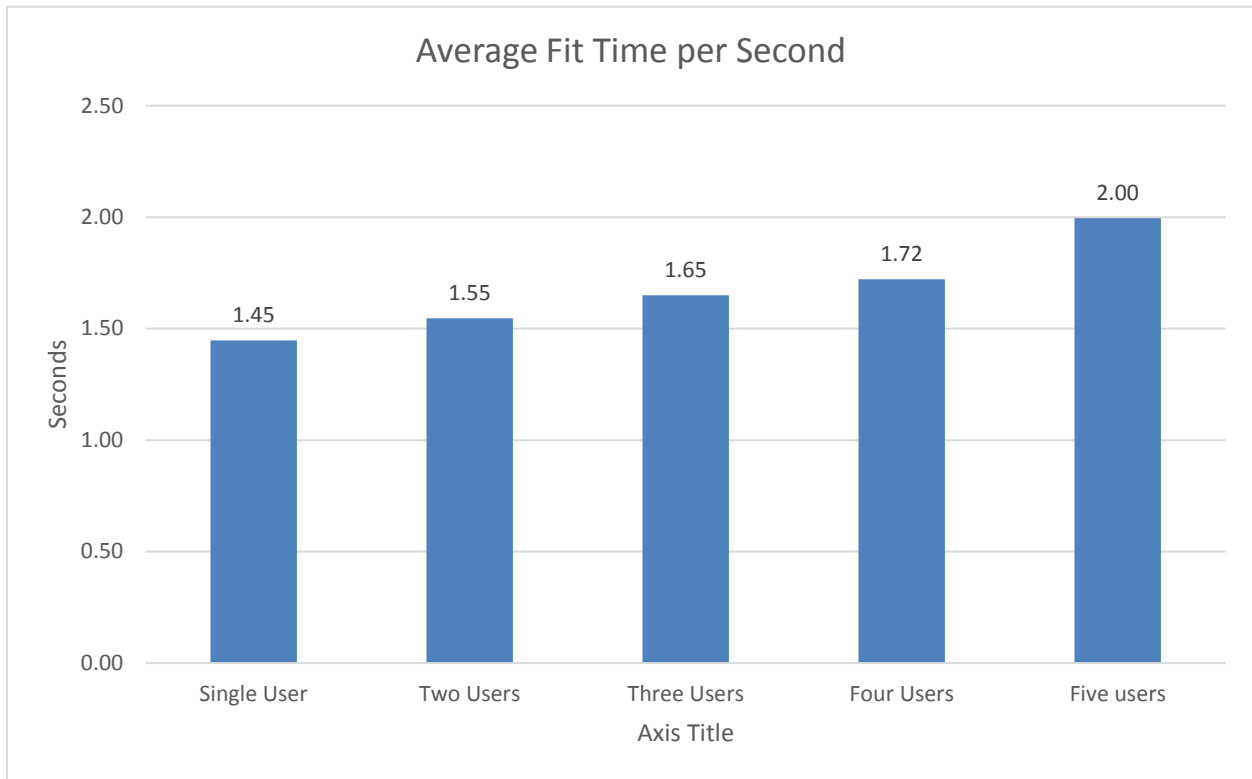
Stability testing eventually proved that the `MaxTessellatorCount` setting `VBOMemoryPercent` setting needed to be managed so that resources were split between each connection. Settings outside the parameters outlined in this paper resulted in resource contention between clients during multi-user testing. This resource contention manifested itself as some clients having completely black screens. Others appeared to momentarily hang. However, in most cases, the machines would eventually return to operation as the operating system managed resources.

Without properly configuring an NV6 system, this resource contention could be observed when three clients were connected. After proper configuration, it was observed that at six simultaneous users, the

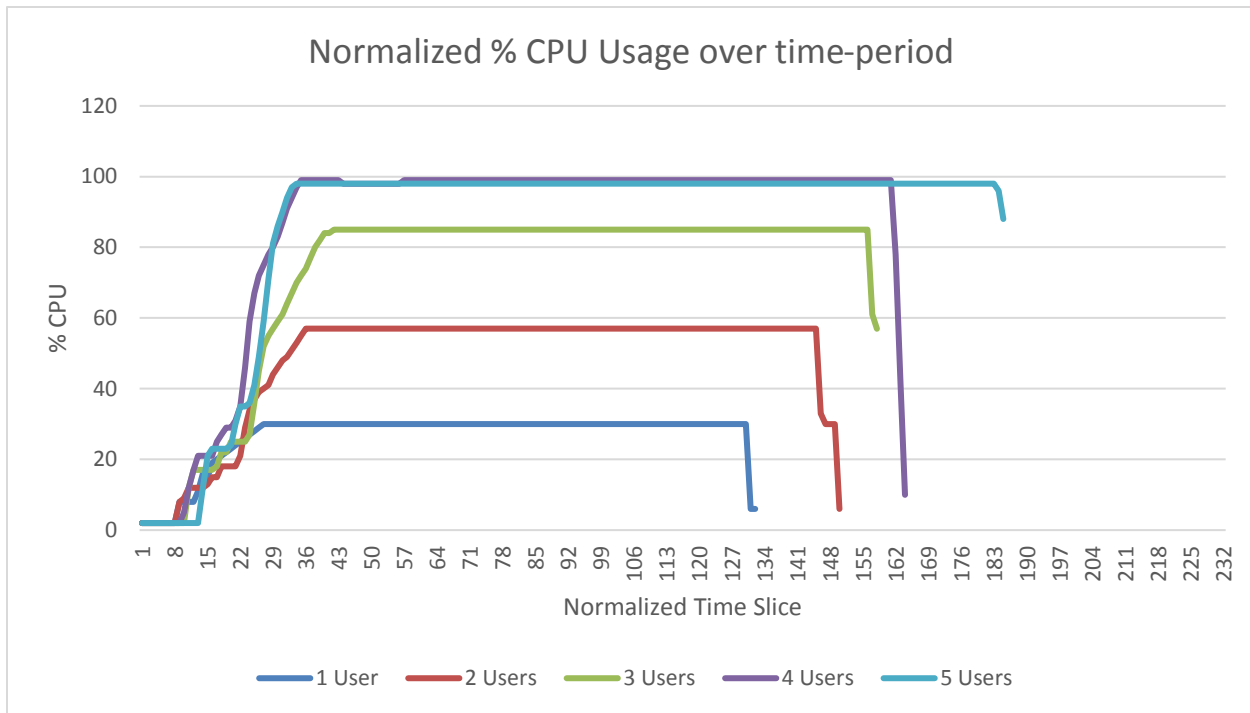
resource contention would manifest again. It was determined that the resources available for testing were enough to support five simultaneous users. This was used as the benchmark for further testing.

4.3.2 Performance Results

When settings were configured as outlined above, the performance slider was set to **Balanced** and multi-user performance tests were performed. Synthetic results were achieved by using a graphically dense and complex test case along with the test harness outlined above to stress each client.



It was observed that for up to four users, the performance decrease was negligible and linear. When a fifth user was added, the linear change in performance skewed. After observing CPU usage at this level, all available resources were being taken. It was concluded that this causes the spike in performance.



4.3.3 Ad-Hoc Testing

The final testing was ad-hoc testing where a user attempted to perform basic navigation and operations in SmartPlant Review while the server was under load. Again, the model was artificially dense, but the user observed that results were what they expected of the model and session. It was observed that many customers do not use models this dense, and that further performance tweaking could be done to improve performance at the cost of graphics quality if a customer did need a comparable model.

For more information about Intergraph, visit our website at www.intergraph.com.

Intergraph, the Intergraph logo, and SmartPlant are registered trademarks of Intergraph Corporation. Microsoft and Excel are registered trademarks of Microsoft Corporation. Other brands and product names are trademarks of their respective owners. Intergraph believes that the information in this publication is accurate as of its publication date. Such information is subject to change without notice. Intergraph is not responsible for inadvertent errors.

© Intergraph Corporation. All Rights Reserved. Intergraph is part of **Hexagon**. Part number information.