White Paper

A Case Study Comparing Data Access Methods Using Remote Graphics Software and Siemens PLM Software, Teamcenter 2007 and NX 5

Prepared for:
Customers of Hewlett Packard and Siemens PLM Software

Authors: Keith Guerin and Pamela Taylor, Hewlett-Packard Co.
Date: May, 2008

Executive Overview

Remote client access to managed data across wide area networks with the characteristic high latency, lower speed and bandwidth as compared to local area networks with low latency and high speed has been problematic. The typical engineering environment using remote clients requires movement of large amounts of data from back end servers across the wide area network to the remote clients for further processing. This data movement compromises the security of data considered as corporate intellectual property. Moving large amounts of data can result in systems that are slow, limit productivity, and are frustrating to the end user. In some cases systems can be rendered unusable because of long delays to move data across slow networks.

A possible solution to these problems may be achieved by reducing the amount of data moved across the wide area network and eliminating the need to move data out of the data center where it is considered to have a higher degree of security. The benefits from not moving the large amounts of data across the WAN are increased system usability, greater productivity, and increased security of corporate intellectual property.

This paper presents hardware and software configuration information and test results obtained from testing Siemens PLM Software, Teamcenter 2007 with NX 5 and HP Remote Graphics Software, RGS. The results compare the data access methods of using RGS and of using Teamcenter 2007 with FCC, the File Management Systems local client cache. The test hardware consisted of client workstations, a blade workstation and Integrity blade servers. The use of Remote Graphics software with an HP workstation, Teamcenter 2007 and NX 5 may provide for greater productivity and greater security for corporate data.

Background

Using a high speed local area network connection with low latency to back end servers, users of Teamcenter 2007 and NX 5 on HP workstations enjoy relatively fast loads and saves of managed modeling data and quick response times for transaction requests. A similarly configured remote workstation client using Teamcenter 2007 and NX 5 and using a WAN connection with high latency to the same backend servers will not enjoy the same type of performance.

Transaction requests as well as loads and saves of large data models across the WAN are considerably slower and are often the source of low productivity and frustration.

HP Remote Graphics Software, RGS, receives keyboard and mouse events and transmits 2D and 3D images from a sender system across standard computer networks to remote users. The remote users interact with a host sender system and its applications as if they were using a local workstation. Because it requires no application modifications, RGS may be used in conjunction with any application software. Benefits of RGS include increased productivity and security for corporate data.

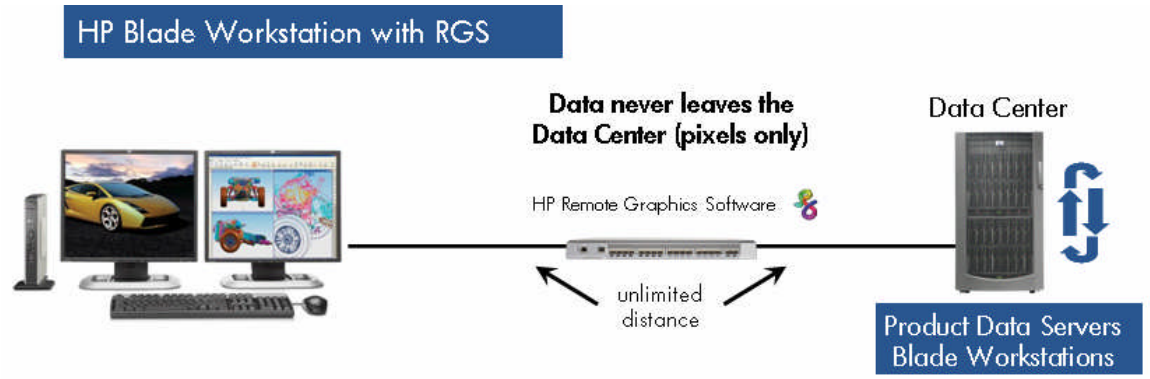

Test Objective

The objective was to measure and compare, under real world conditions, two different methods of client access to Teamcenter data. One method consisted of using a remote client running HP's Remote Graphics Software to access Siemens' PLMS Teamcenter 2007 and NX 5 running on a Blade Workstation located in the data center. The second method was to run Teamcenter 2007 and NX 5 on a remote client located away from the data center.

Hardware and Software Configuration

Teamcenter 2007 was installed in a four tier hardware configuration in the data center located at Siemens PLM Software in St. Louis, Missouri. The Teamcenter 2007 software was running on MS Windows and HP-UX operating systems on a combination of Intel and Integrity Blades in an HP C7000 Blade Enclosure. A component of the Hewlett-Packard Remote Graphics software called a "Sender" was installed on an HP xw460c Blade workstation that served as a local workstation client and was located in the blade enclosure. The workstation used an nVIDIA Quadro FX 1600M graphics card and was running Microsoft Windows XP 32bit.

Remote four tier rich clients were located at Siemens PLM Software in California and were loaded with Siemens' PLMS Teamcenter 2007 rich client with NX 5 and a component of HP's Remote Graphics Software called a "Receiver".

Testing Methodology

Remote clients accessed the Teamcenter data and services by way of various network access points. Latency was measured using the results of a network "ping" from the remote client to servers located in the data center.

Remote clients would access Teamcenter in two different ways. The first method was as a standard Teamcenter login from a remote client running the Teamcenter 2007 rich client and NX 5 software. The second method was to use the same client but running the HP Remote Graphics software and connected to the HP xw460c Blade workstation which was serving as the local client located in the blade enclosure in the datacenter.

In both connection methods the client would perform multiple loads and saves of engineering CAD data and perform a series of graphics operations. Two different size assemblies, one large and one small, were used for the tests. The assemblies are described in more detail later in this document.

Before running the tests various configuration options were set. Screen resolutions of both local and remote clients were set similarly. Graphics operations were performed after the RGS fixed frame rate was disabled and the assembly fully loaded. Timing information for the various operations was recorded.

When the clients connected using the HP Graphics Software they would connect from California to the Workstation Blade in the Blade enclosure located in the datacenter in Missouri by way of different access points and latencies. The latency between the local Workstation Blade in the enclosure and the Teamcenter 2007 servers was less than 1ms. The RGS software would respond to keyboard and mouse events from the remote client and send compressed screen update information back to the remote client. The part or assembly data was never copied to the remote client, only the compressed screen image.

Part Information

The large and small part characteristics are as follows:

Large Part: (250Component Assy, Rev G)
Ugraf Memory Usage: 1,012,332 bytes.
Peak Memory Usage: 1,486,356 bytes.
Total including VM: 1,845,636 bytes.
Note: /3GB setting was required due to the large assembly size.

Small Part: (87Component Assy, Rev B)
Ugraf Memory Usage: 358,500 bytes.
Peak Memory Usage: 362,432 bytes
Total including VM: 654,296 bytes.

Test Actions Measured

Load or Open Assembly
Run Graphics Operations
    Change View to Shaded,
    Change View to Wireframe
    Change View to Shaded
    Rotate
    Change View to 4 View & Work View
Save as New Item Revision
Load or Open Assembly a Second Time
Save Again

Test Results

Load and Saves

The following two charts illustrate the load and save times of the large assembly on different platforms using two different connect methods. The first group of seven vertical bars represents the first load of the large assembly. The second group of seven bars represents the same load performed a second time. The first three vertical bars show the load times without using HP's Remote Graphics, RGS, while the remaining four bars show the load using RGS.

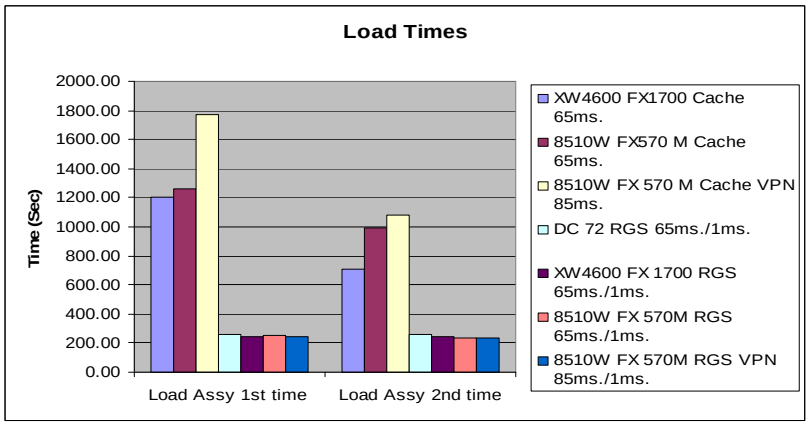

It is clear from reviewing the charts that both the load and save times are dramatically faster when using RGS. This is because when using RGS, loading and saving of the assembly, does not require the assembly to be transported across the network to the remote client. The remote client using RGS is connected to the "sender" on the workstation client in the blade enclosure which is local to the Teamcenter servers also in the same enclosure. This close proximity results in a latency of less than 1ms between the workstation and the servers in the same enclosure and very fast loads and saves as illustrated by the latter four bars of each group of seven in the chart. Only the pixel information from the compressed image is transported across the network to the remote client when using RGS.

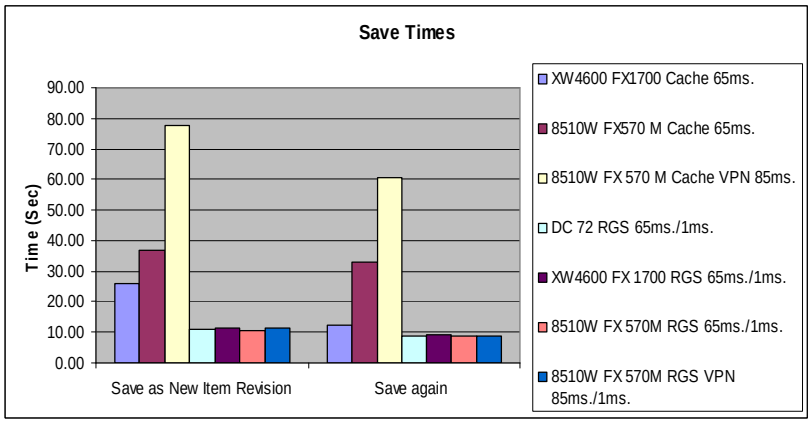

When not using RGS it is evident that the loads and saves suffer from delays of moving the assembly across the network. In this case the network latency was approximately 65 or 85ms contributing to the apparently long times to load or save the assembly. It is also clear that once the assembly is moved, subsequent loads and saves are faster. This is a result of having the assembly loaded or saved from the local Teamcenter 2007 cache, or FCC.

This third chart shows the loads and saves using a small assembly. It is clear once again that there are dramatic savings for load and save times when using RGS. It is also evident that the local Teamcenter 2007 cache, FCC, plays a significant role in speeding up subsequent loads and saves when not using RGS.

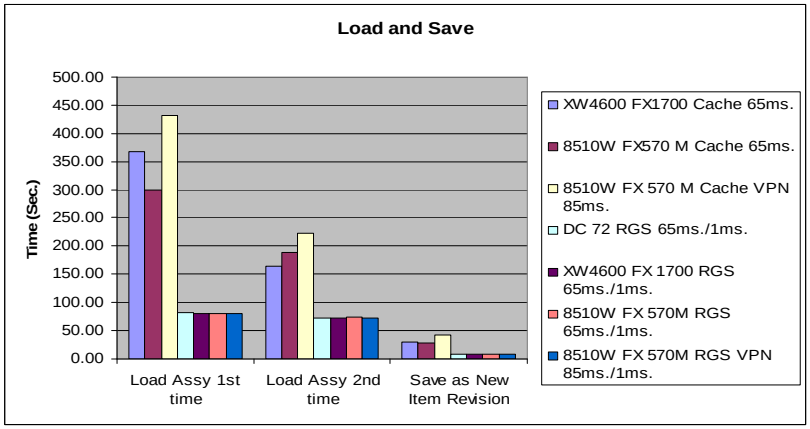

Graphics Operations

Graphics operations were performed after the assembly was loaded into the local cache. Regardless whether the client was located remotely or in the data center the assembly was loaded into the cache closest to the client and took advantage of graphics hardware for the dynamic operations. Operations performed were typical rotations, change of view, shading and wireframe. In this chart you will notice that the latter four bars are all using RGS and the same graphics hardware, nVIDIA Quadro FX 1600M graphics card, on the workstation client located in the datacenter. Therefore the times for the graphics operations were all very similar, approximately 1.6 seconds. The first three bars indicate variations due to different graphics hardware different platforms, clock speeds and some minor variation due to network latency. With the large assembly you can see that the graphics operations times

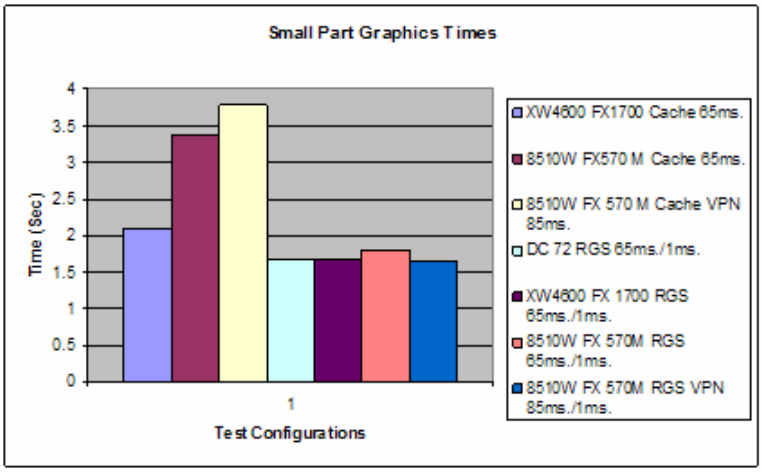

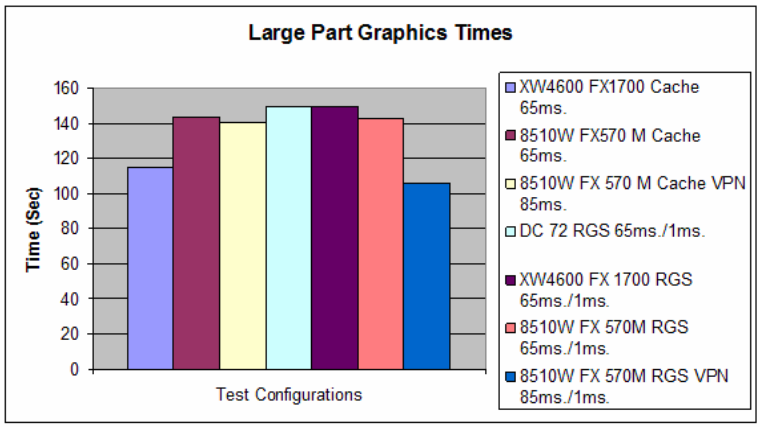

are much closer but take much longer time than the small assembly. This is mostly because of the size of the assembly and the amount of time required by the hardware to effect the required transformations. The benefit of having the assembly reside in the Teamcenter 2007 local cache becomes clear as the total time required for transformations far exceeds that used to access the assembly in the local cache resulting in times that are closer in duration.

The effect of higher latency and slow VPN

This chart is the result of using the small assembly and recording the times for loads and saves with both RGS and non RGS at different latencies.

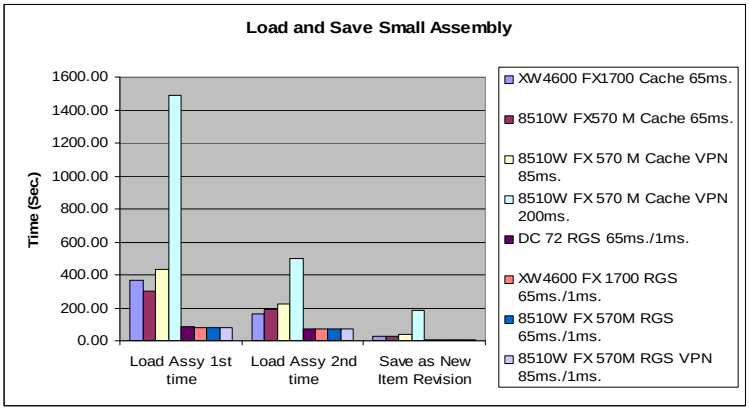

The following tables summarize the chart and indicate the effect of latency on the time to load or save the assembly and its effect on RGS performance.

Using RGS	
Note: RGS was thru the remote network access point to the xw460c workstation client and the local client cache, FCC, in the datacenter. The latency from the xw460c to the Teamcenter servers was less than 1ms.	
Latency (ms)	**Results**
85	The load and save times were relatively fast and the graphics were fast.
200	The load and save times were slower but the interactivity was unacceptable. The warning message "Network Connection Warning" appeared frequently.
>200	Did not attempt tests at higher latency.

Using Non-RGS	
Note: The remote client acts as a Teamcenter 2007 rich client accessing the Teamcenter 2007 servers through the remote network access point with the latency indicated below.	
Latency (ms)	**Results**
85	The load and save times were slower than when using RGS but the graphics were fast
200	Latency of approx 200ms. The load and save times were very slow, 25 min for the load, but graphics performance was good once the download was completed.
>200	Did not attempt tests at higher latency.

Conclusions

Network latency, size of modeling data, and data security are important considerations when deciding the best method for remote clients to access managed data. The two methods to choose between are whether to use Teamcenter 2007 with its local FCC installed on the remote client or to use RGS and connect to an HP xw460c Blade Workstation with Teamcenter 2007 and its local FCC installed and located in the datacenter.

When Teamcenter 2007 is installed on a remote client the data transfers increase in duration and in proportion to the network latency and the amount of data being transferred. When the data transfers have been completed graphics operations are applied to data resident in the local cache and interactivity is very good.

When using RGS the data remains secure in the data center and only the compressed and updated screen images along with keyboard and mouse events are sent across the network to the remote client. Loads and saves of data are very fast but interactivity is reduced as the latency increases and the network bandwidth is consumed by high resolution display settings and frequent changes to display data.

Using RGS with a network latency of 200 milliseconds and using our small assembly the interactivity with the graphics display was unacceptable and made the use of RGS impractical. As the latency and the amount of modeling data increase it may be more effective to use a local installation of Teamcenter 2007 with its local cache on the remote client. This means that it is better to accept the long duration of the data transfers to and from the local Teamcenter 2007 FCC in order to achieve greater interactivity with graphics operations.

For acceptable graphics performance it is recommended to minimize the latency as much as possible by finding alternate access points and by reducing the amount of network bandwidth used. This can be accomplished by lowering the screen resolution and by reducing the amount of modeling data and therefore the amount of network bandwidth used. We experienced acceptable performance using RGS for both the large and small assemblies with latency less than approximately 95 milliseconds.