Technical white paper

Determining the duration of an HP RGS session

Confirmation for HP Remote Graphics Software connectivity

Table of contents

HP's Remote Graphics Software (RGS)[1] is a productivity tool for engineers and creative professionals, and is free to all HP Z Workstation customers. It is simple to install, easy to operate, and enables one or more users to access professional workstation sessions through their networks. There are no complex brokers required, yet it can be used in a brokered environment if desired. HP RGS is robust and versatile to handle many different types of network connections: high-bandwidth, low-latency corporate local area networks; trans-continental networks; and secured VPN networks to home and hotel rooms.

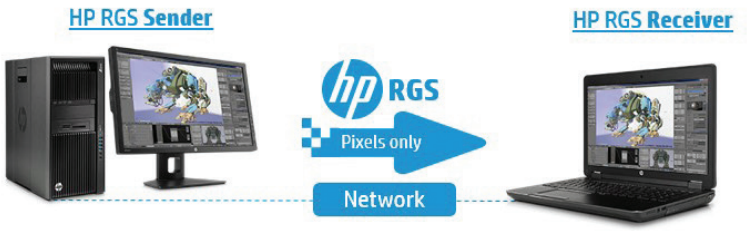

How long will an HP RGS session stay connected?

Customers might ask how long an RGS session will remain connected. To find the answer, we tested and measured a session using HP Performance Advisor (HPPA), an HP productivity tool that is included with all HP Z Workstations. Within HPPA, we used the Workstation Monitor and GPU Utilisation tools to track system loads during heavy RGS connections: 1) On the workstation that is running the user's graphics-intensive professional applications (referred to as the "Sender"), and 2) On the remote system that is connected to the workstation via RGS (referred to as the "Receiver").

This test exemplifies one of the many unique advantages of HP RGS being invented, developed, and supported by the same HP Workstation division that leads the market for high-performance, reliable workstations.

Test set-up

To perform the test, the following was in place:

1. An RGS Sender system running graphics-intensive workload on one or two high-resolution displays.
2. An RGS Receiver system with sufficient CPU and graphics to handle the high-resolution display(s) and remote work.
3. Both systems were connected to an inexpensive Gbit network switch, which was connected to the corporate LAN. This configuration minimised possible network effects to the test. (Note that the network infrastructure effect can be eliminated by connecting the Sender directly to the Receiver system.) Both systems properly negotiated the link to 1000 Mbit (Gbit).

4. Both systems were set to never sleep or turn off the displays.
5. After connecting from the HP RGS Receiver to the HP RGS Sender:
   A. HPPA's Workstation Monitor was started on each system to sample and set at 15-second intervals for multiple days, the RGS application to monitor was identified, and the graphics-intensive application to monitor on the Sender was optionally selected.
   B. A graphics-intensive workload was started on the Sender system and a window size was used that consumed a significant portion of the desktop space.
   C. The HPPA Workstation Monitor was started to track and log the system load.

Test systems

All systems had Windows 7 64-bit

Sender system 1:	Receiver system 1:
• HP Z240 Intel® Xeon® E3-1245v5, 3.5 GHz	• HP Z210 SFF Intel® Xeon® E3-1225, 3.1 GHz
• 64 G system memory	• 8 G system memory
• NVIDIA® Quadro® K2200 GPU	• AMD FirePro™ V3800 GPU
• 2 x 24 in. (1920 x 1200) displays as a single Windows desktop	• 2 x 30 in. (2560 x 1600) displays as a single Windows desktop (larger than Sender)
Sender system 2:	**Receiver system 2:**
• HP Z240 Intel® Xeon® E3-1245v5, 3.5 GHz	• HP ZBook 15 G3 Intel® Core™ i7-6820HQ, 2.70 GHz
• 64 G system memory	• 64 G system memory
• NVIDIA® Quadro® K2200 GPU	• NVIDIA® Quadro® M1000M GPU
• 1 x 24 in. (1920 x 1200) display	• Full HD (1920 x 1080) panel + 24 in. (1920 x 1200) (attached display exactly matched Sender)
Sender system 3:	**Receiver system 3:**
• HP Z820 Intel® Xeon® E5-2697 v2, 2.7 GHz	• HP Z240 Intel® Xeon® E3-1245v5, 3.5 GHz
• 64 G system memory	• 64 G system memory
• NVIDIA® Quadro® M6000 GPU	• NVIDIA® Quadro® K2200 GPU
• 2 x 30 in. (5120 x 1200) displays as a single Windows desktop	• 2 x 30 in. (5120 x 1600) displays as a single Windows desktop (exactly matched Sender)

Test 1 results

The RGS Sender system was fairly heavily loaded the entire time. Because this was a high-resolution, graphics-intensive session, the (default) HP3 codec and default quality setting of 65 was used. NVIDIA®'s Mosaic™ feature was used to present both displays as a single, large display surface to Windows.

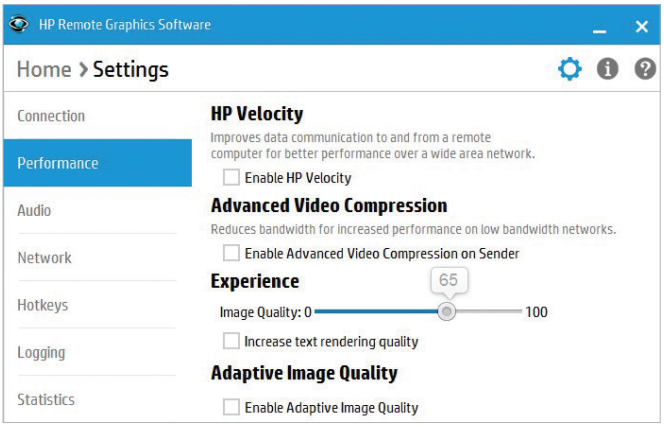

The RGS Sender load was sampled starting at Day-0 1:03:42 p.m., and running through Day-6 3:22:14 p.m. The final sample was taken by HPPA Workstation Monitor at 6 days, 2 hours, and 17 minutes after sampling started. For this particular test, the application graphics load was very intensive for the graphics card, keeping it well above 70 percent utilisation with constant updates of complex geometry and textures. RGS's use of the graphics card was comparatively tiny, but still non-zero to capture the desktop and send updates to the receiver system as fast as possible.

Because of the relatively high combined display resolution, the RGS Sender was using ~8 percent of the overall CPU to capture, encode, encrypt, and send desktop updates to the receiver system. And, due to the high resolution, complex imagery, and interactive update rates, the average network bandwidth was ~50 Mbit/sec, with peak network (not shown) up to 90 Mbit/sec. The variability in network use is due to the variability in desktop updates—size, complexity, and update rate. This sampling shows RGS Sender-Receiver sessions remained connected for more than 6 days, at which point the test was manually stopped.

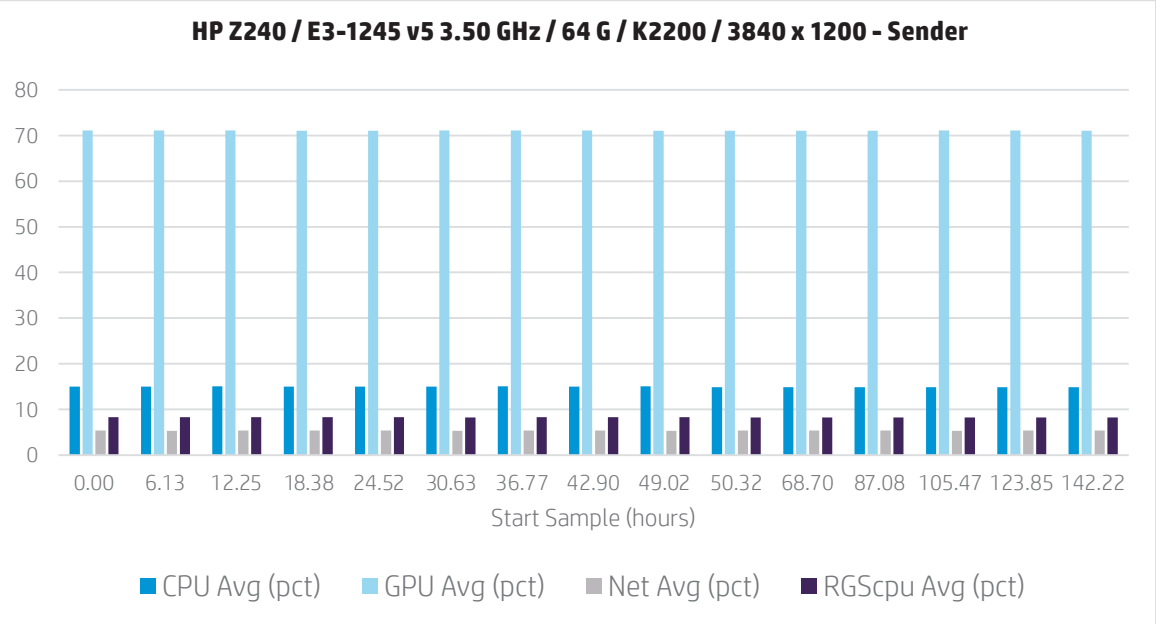

Correspondingly, the HPPA Workstation Monitor sampling results were lower than the RGS Receiver over this same time period. Although the receiver system was driving two 30-inch displays, it was using the AMD FirePro EyeFinity™ feature to combine the 2 displays into a single 5120 x 1600 desktop for Windows. About 14 percent CPU is needed to decrypt and decode the RGS stream, and subsequently display it on the receiver system with this large screen. The local graphics load, even for this high-resolution session, was not too strenuous for the low-end FirePro card. As expected, the average network load matched the Sender system at about 50 Mbit/sec. The Receiver system was running HP RGS and HPPA, but no other applications for this test.

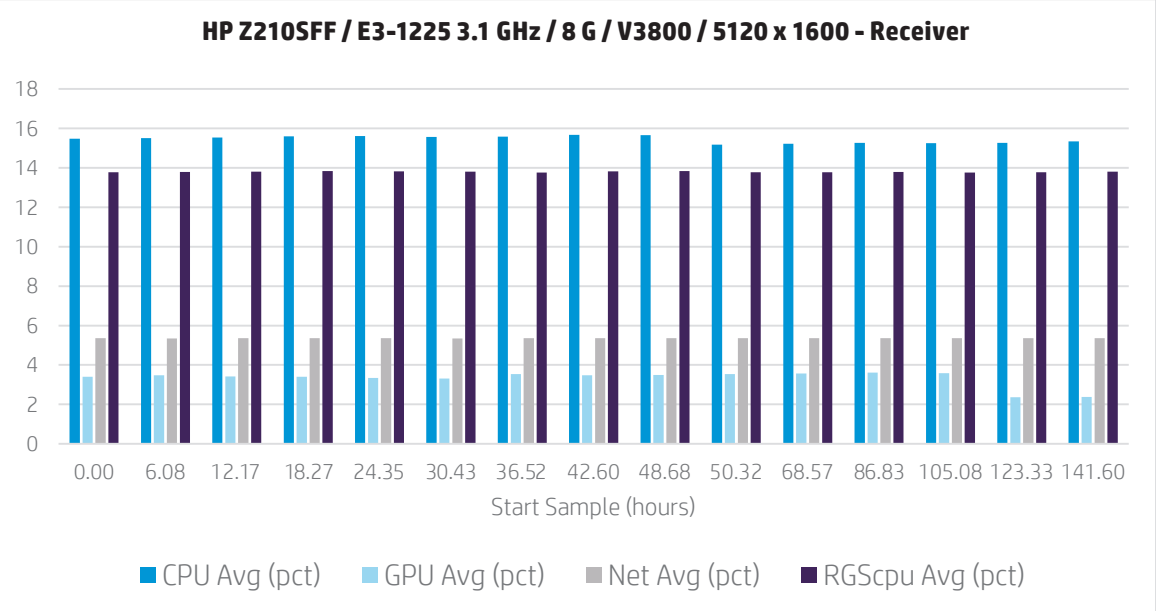

Test 2 results

Next, a simpler 1920 x 1200 display session was used, but changed to HP RGS's Advance Video Compression encoding mode.

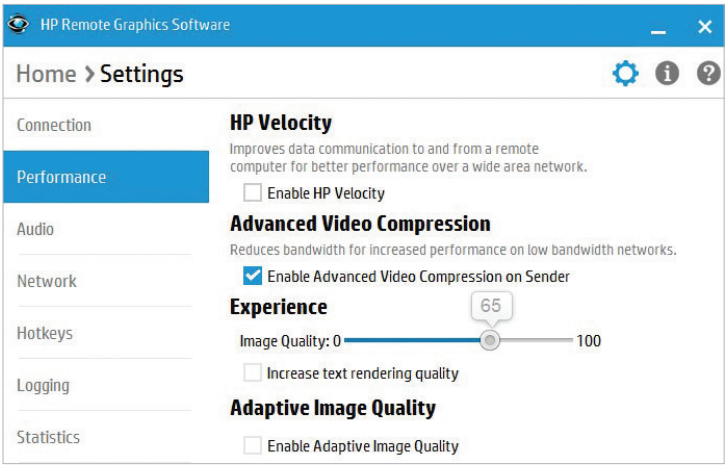

The HP RGS Sender load was sampled starting at Day-0 1:42:05 p.m., and running through Day-4 1:41:38 p.m. The last sample taken by HPPA Workstation Monitor was nearly 4 days after sampling first started. For this particular test, the application graphics load was also very intensive for the graphics card, keeping it well above 95 percent utilisation with constant updates of complex geometry and textures.

In addition to the application's load on the graphics card, Advance Video Compression took advantage of some advanced features in NVIDIA®'s Quadro® graphics cards to handle much of the encoding work. Because the graphics card was doing more of the encoding work in this mode, the RGS Sender used <1 percent of overall CPU to capture, encode, encrypt, and send the desktop updates to the Receiver system. Even so, the overall system used ~17 percent CPU on average, mostly due to the intensive graphics application.

Because the lower-fidelity AVC mode was used, the average network bandwidth was ~2.5 Mbit/sec, with peak network (not shown) up to 2.6 Mbit/sec. The consistency in network is a side-effect of using AVC, where update rates and complexity of the encoding stay fairly constant over time. This HPPA sampling showed that the RGS Sender-Receiver session stayed alive for nearly 4 days, at which point the test was manually stopped.

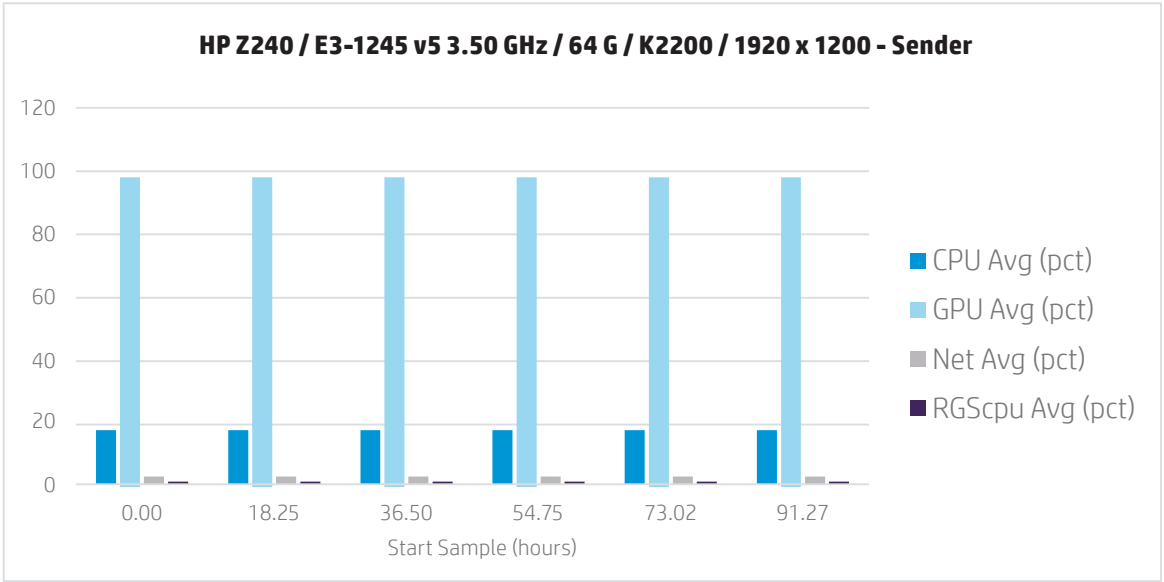

Correspondingly, the HPPA Workstation Monitor sampling results were below the RGS Receiver over the same period of time. The Receiver system was driving a Full HD panel and a 24-inch 1920 x 1200 attached display. About 12 percent CPU was needed to decrypt and decode the RGS stream and subsequently display it on the attached display. The local graphics load for this session was easy for the NVIDIA® M1000M GPU. As expected, the average network load matched the Sender system at about 2.5 Mbit/sec, corresponding to the Sender system's network load. The Receiver system was running HP RGS and HPPA, but no other applications for this test.

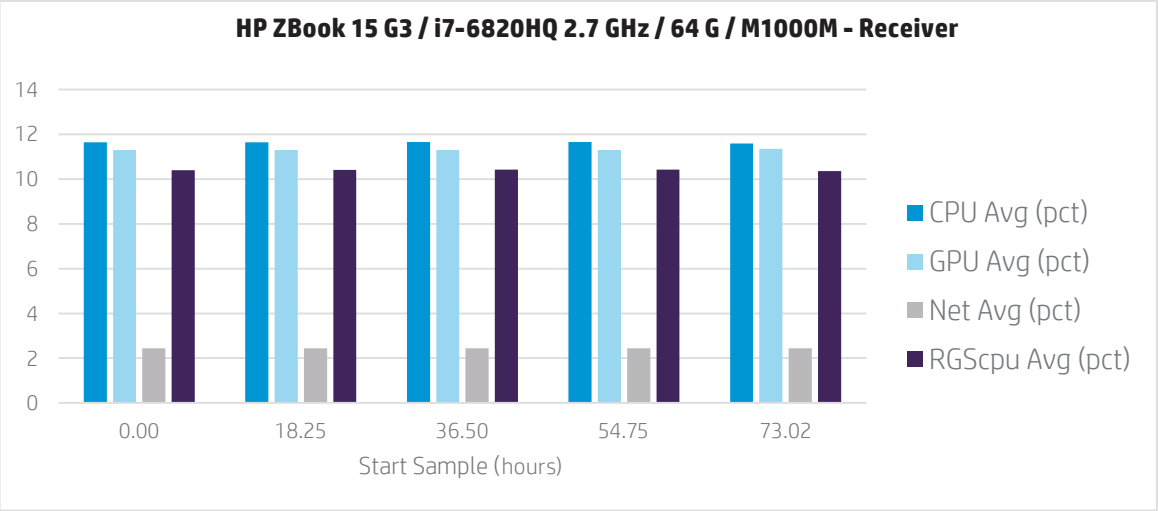

Test 3 results

Again, the RGS Sender was fairly heavily loaded the entire time. Instead of running a single graphics-intensive load on the system in a large window, we ran two graphics loads with each in smaller windows, one using OpenGL, and the other using DirectX. Because this was a high-resolution, graphics-intensive session, we chose to use the (default) HP3 codec and default quality setting of 65. We used NVIDIA®'s Mosaic™ feature to present both displays as a single, large display surface to Windows.

This third test was similar to the first test with a couple of exceptions. The RGS Sender load was sampled starting at Day-0 5:23:13 p.m., and running through Day-5 7:22:17 p.m. The last sample taken by HPPA Workstation Monitor was nearly 5 days and 2 hours after sampling first started. The application graphics load was moderately intensive, keeping it above 45 percent GPU utilisation with constant updates of complex geometry and textures in both application windows. Again, RGS's use of the graphics card was comparatively tiny. Because of the relatively high combined display resolution, we saw that RGS Sender was using almost 6 percent of the overall CPU to capture, encode, encrypt, and send the desktop updates to the Receiver system. Although the screen resolution was quite high, the combined resolution of the graphics windows being updated at 30-60fps was less than the single graphics window resolution in the first test. Hence, the average network bandwidth comes in quite a bit lower, at ~10 Mbit/sec, with peak network (not shown) up to 12.6 Mbit/sec.

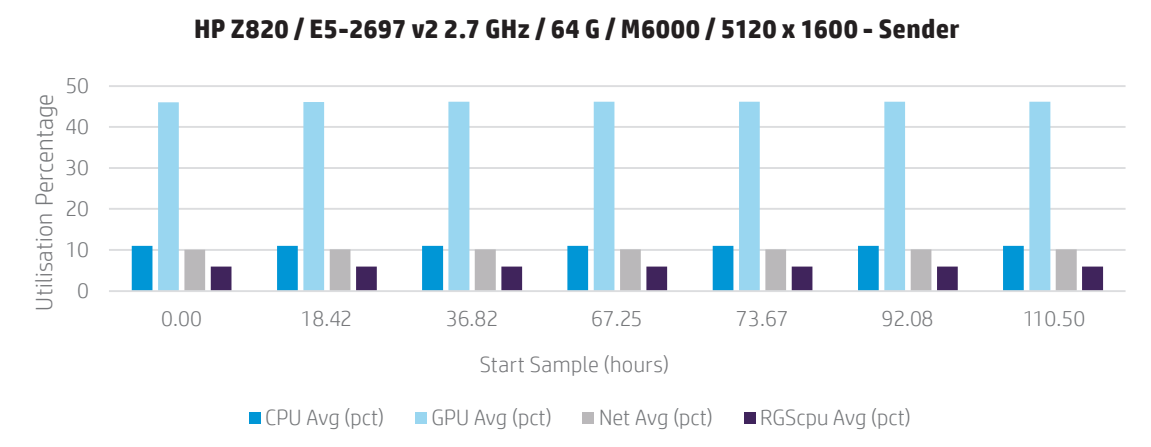

Correspondingly, we see the HPPA Workstation Monitor sampling results below the RGS Receiver over the same period of time. The Receiver system was driving dual 30-inch displays using NVIDIA® Mosaic™ mode to present these as a single desktop to Windows, and the RGS session was set up to remote the entire equivalently sized Sender desktop. About 11.5 percent CPU was needed to decrypt and decode the RGS stream and subsequently display it on the attached displays. The local graphics load for this session was relatively easy for the NVIDIA® K2200 GPU, averaging less than 10 percent GPU utilisation. As expected, the average network load matches the Sender system at about 10 Mbit/sec, corresponding to the Sender system's network load. The Receiver system was running HP RGS and HPPA, but no other applications for this test.

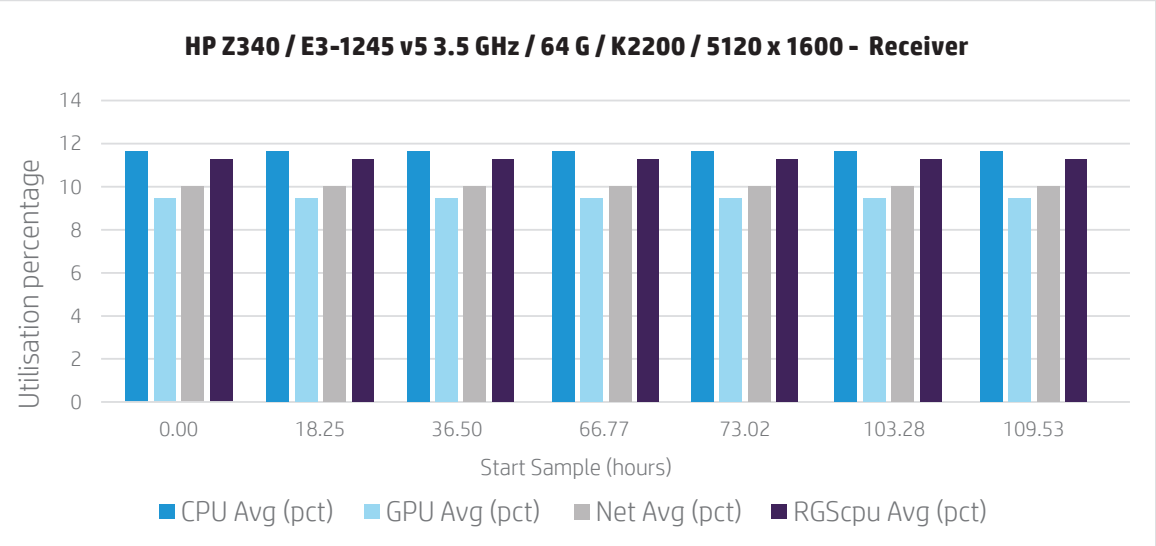

Conclusion

Customers can rely on HP RGS to keep running, even with extremely demanding and graphically intensive remote sessions and applications.

[1] HP RGS requires a Windows, Linux®, or OS X 10.10 and newer operating system and network access.

Sign up for updates
hp.com/go/getupdated

Share with colleagues

4AA6-6237EEW, February 2017