

Fujitsu develops high-speed thin client technology for 10-fold improvement in responsiveness

May 15 2012

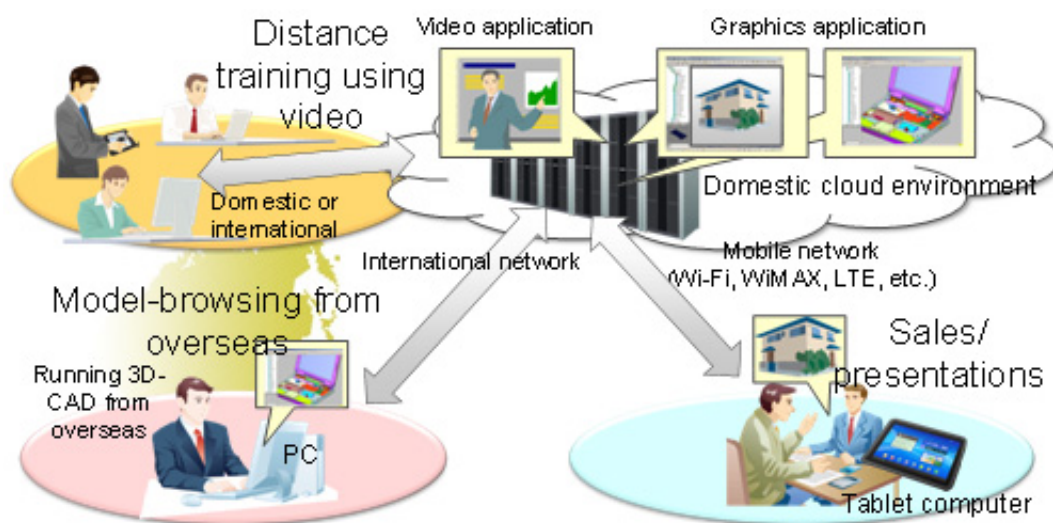


Figure 1: Various uses for this technology in thin-client scenarios

Fujitsu Laboratories Limited today announced that it has developed technology for high-speed thin clients that will respond as much as ten times faster to user operations, even in low-quality network environments.

In cloud-based virtual desktop environments, when client terminals access the virtual desktop environment for video or graphics-heavy applications, system responsiveness has conventionally been poor,

making it difficult to use such applications. Particularly when the terminal using the virtual desktop is in another country or in a mobile setting, latency in data transmissions and data loss (packet loss) make these problems all the more glaring, posing a challenge in using cloud-based applications for companies with global operations.

Fujitsu Laboratories has developed high-speed thin-client technology that slashes response times by up to 90%, even in low-quality network environments. For companies with global operations, this technology will make it possible for users to efficiently operate applications on the virtual desktop, whether at home or abroad. With the design process rapidly becoming more global, this will enable dispersed design teams to collaborate securely. This technology also paves the way for new thin client applications, such as international video-based employee training, remote operation of internal applications using tablets in mobile settings, and interactive product demonstrations at customer locations (Figure 1).

This technology is being exhibited at Fujitsu Forum 2012, which runs May 17–18 at the Tokyo International Forum.

"Desktop as a service" (DaaS) is a cloud-based service that virtualizes a PC's desktop environment and makes it remotely accessible from a thin client. This concept has been gaining interest for the benefits it offers in terms of business continuity planning(1) and security for manufacturers, and from the perspective of the operational management costs of PCs and other terminals. As companies in Japan increasingly operate on a global level, manufacturing is rapidly coming to involve sites in multiple countries. In this context, thin clients are viewed as an attractive way to allow remote work without sending design data abroad, thereby lessening the risk of information leaks. Furthermore, as the performance of smartphones and tablets has increased, there is an increasing need to use virtual desktop services from a variety of mobile environments.

In today's virtual desktop environments, visual data on the virtual desktop is rendered on the server and then transmitted to the client terminal, enabling the desktop to be used remotely. Up until now, however, this has presented the following problems:

1. Latency in data transmission

The networks that tie together Japan with other countries in Asia or elsewhere typically take at least 100 milliseconds for data to make a round trip, so when using graphics applications such as 3D CAD software, where the client is in one country and the server is in another, response times are slow. For example, the round-trip time between Singapore and Japan is about 100 milliseconds, and the round trip between North America and Japan is about 200 milliseconds.

2. Packet loss

Running visually intense software on a virtual desktop, such as 3D-CAD or video, requires tens of megabits per second of bandwidth capacity per client. But with overseas transmissions or mobile clients, packet loss will be frequent, making impossible more than a few megabits per second of bandwidth.

3. Graphics processing

Using video or graphics in a virtual desktop environment requires a special graphics-processing unit (GPU)(2) to draw the screen images. But allocating one GPU per virtual machine would require an amount of installed space that would make server consolidation impossible, raising operating costs.




Figure 2: Sample image-quality and frame-rate adjustments using image-data reduction technology

Fujitsu Laboratories has previously developed a technology called RVEC (announced May 2011) that enables applications that use video and high-resolution graphics in a domestic virtual-desktop environment to run smoothly. To address the issues listed above, now Fujitsu Laboratories has developed high-speed thin client technology that extends RVEC technology, allowing it to work in international and low-quality network situations, such as with mobile devices. It also dramatically improves responsiveness, even in network environments subject to high latency and high packet loss. Some of the features of this technology are as follows.

1. Reduced visual data

In the virtual desktop environment, image quality and frame rate are adjusted when operating an application, in response to the available transmission bandwidth, before sending to the client. This image-data reduction technology ensures that the volume of data being transmitted and the operational latency are both kept in check (Figure 2). Compared to the previous technology, RDP(3), responsiveness is as much as ten times faster.

2. Network latency compensation

To maintain responsiveness even in network situations with relatively high packet loss and latency, such as international connections or mobile networks, Fujitsu Laboratories developed a new protocol. This protocol allows for visual data to be transferred at speeds six times greater than the conventionally-used protocol, TCP(4).

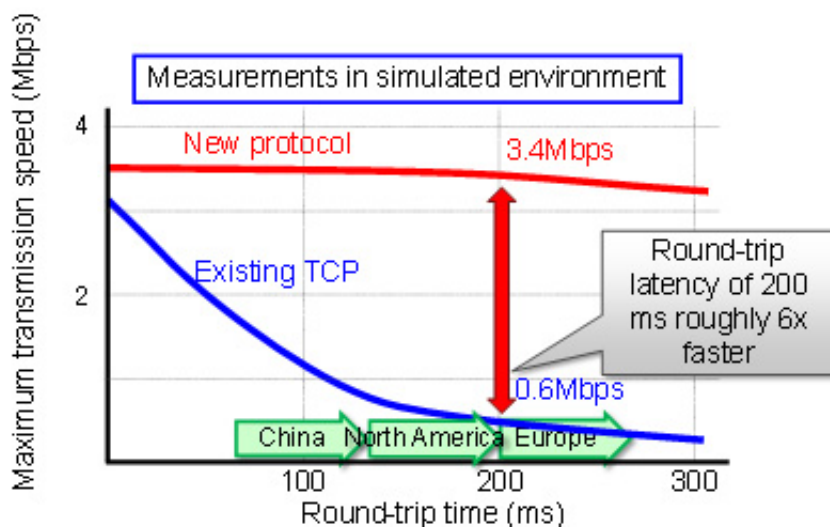


Figure 3: Performance assessment of network-latency compensation

3. GPU virtualization

When running graphically intensive applications in a virtual desktop environment, the GPU that accelerates the drawing of graphics, such as 3D-CAD, is shared out among applications running on multiple virtual machines (Windows, etc.) simultaneously for improved responsiveness. In the Fujitsu Laboratories test setup, one GPU could be shared among four virtual machines (Figure 4). With GPU virtualization technology, a

CPU/GPU switching mechanism allows screen-update decisions and static image compression to be handled by either the CPU or GPU, depending on their relative loads. In addition, a GPU sharing mechanism delivers draw instructions from multiple virtual machines to the GPU in parallel. This makes it possible to increase consolidation of virtual machines in a virtual desktop environment and to deliver predictable, fast responsiveness.

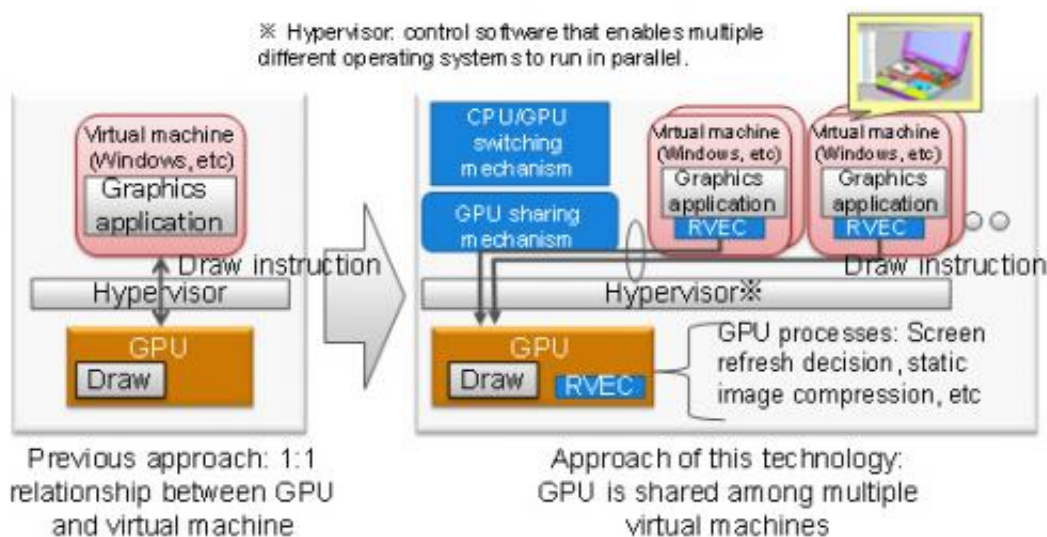


Figure 4: GPU virtualization (sharing) technology

In addition to making it possible for globally scattered design teams to work together more effectively, this technology promises to open up new applications for thin clients. It can be useful not only in 3D-CAD but also in CAE(5) and other graphically intensive applications. Furthermore, it can be useful in training staff at foreign locations using educational video materials. And because internal applications can be used from remote locations more smoothly than before, even in situations with unpredictable bandwidth, such as mobile settings, it can be beneficial for interactive product demonstrations at customer

premises.

This technology is slated to appear in commercial products during fiscal 2012, in the form of the Engineering Cloud™, the next-generation manufacturing environment announced June 2011.

Provided by Fujitsu

Citation: Fujitsu develops high-speed thin client technology for 10-fold improvement in responsiveness (2012, May 15) retrieved 8 April 2024 from

<https://phys.org/news/2012-05-fujitsu-high-speed-thin-client-technology.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--