IT@Intel White Paper

Intel IT
IT Best Practices
Cloud Computing and Employee Productivity
August 2011

Applying Client-aware Technologies for Desktop Virtualization and Cloud Services

Chuck Brown
Director, Emerging Compute Lab
Intel Architecture Group

John Dunlop
Enterprise Architect
Intel IT

Ryan Ettl
Systems Integration Engineer
Intel Architecture Group

Tom Mant
Technology Evangelist
Intel IT

Executive Overview

In the context of developing a private cloud, Intel IT is exploring client-aware technologies that take advantage of end-point device capabilities. These technologies can be applied in the areas of desktop virtualization and Web services, and can help maximize productivity and deliver an improved user experience.

Our testing demonstrates that client-aware technologies can offer several benefits to both users and IT.

While Intel employees continue to rely on mobile business PCs as their primary computing devices, they are also using an expanding continuum of companion devices, all of which have widely differing capabilities. Delivering services with a one-size-fits-all approach will not result in an optimal end-user experience and could pose a risk to Intel intellectual property.

Instead, we envision a private cloud that can determine device attributes and user preferences, and tailor services accordingly. An intelligent, client-aware cloud could determine:

- Whether an application would provide the best user experience if executed locally or remotely.
- Which native features, such as a Global Positioning System (GPS) receiver or accelerometer, are available on a device.
- How to apply predefined user and device profiles to customize services to user preferences and the device's level of security access.

Internal testing demonstrates that client-aware technologies such as workload shifting, reverse seamless, and rich Internet applications (RIAs) can offer several benefits to both users and IT:

- **User benefits.** Client-aware services can provide a much richer experience that enhances employee productivity by taking advantage of device-level features such as a GPS receiver. Instead of displaying only a map, a device could provide step-by-step directions to conference rooms, printers, or shuttle bus pick-up locations.
- **IT benefits.** Some of the client-aware technologies we are investigating help resolve performance issues by reducing server and network infrastructure load. Additionally, publishing a single service with varying levels of capabilities based on device features is far more efficient than developing multiple services specific to each device—which would add management and support overhead.

We are continuing to do pathfinding on these technologies through our Intel® IT Innovation Centers and plan to integrate these and future client-aware technologies into Intel's private cloud infrastructure.

Contents

IT@INTEL

The IT@Intel program connects IT professionals around the world with their peers inside our organization – sharing lessons learned, methods and strategies. Our goal is simple: Share Intel IT best practices that create business value and make IT a competitive advantage. Visit us today at www.intel.com/IT or contact your local Intel representative if you'd like to learn more.

BACKGROUND

Intel IT continually investigates ways to enhance productivity, mobility, and security while lowering costs through improved device management. Client-aware desktop virtualization, combined with client-aware Web services delivered using a private cloud, is one approach that could help us meet these goals.

Intel employees want to use a broader range of companion devices with their mobile business PCs, including personally owned smart phones and tablets. We are therefore developing and testing private cloud capabilities to deliver our IT environment and applications as services that run on a variety of personal and corporate devices using a combination of virtualization technologies. These technologies include client-hosted virtualization (CHV), server-hosted virtualization (SHV), and client-aware Web-based services.

This approach allows us to take advantage of the scalability of cloud computing while delivering a better user experience, enabled by ongoing innovation around client devices.

Desktop Virtualization

Of the two types of desktop virtualization solutions available, we primarily use CHV to deliver services to users. For example, we use type-2 hypervisor technology to provide an Intel IT software environment for contractors who develop software for Intel. We are also investigating using SHV, which has the potential to allow centralized, simplified image management.

However, neither CHV nor SHV are ideal solutions for many use cases. While CHV offers the performance and flexibility of local execution, it is most useful only with client devices that have a robust feature set. With SHV, computing occurs in the data center, enabling employees to access applications and data through a variety of client devices, including those with a less robust feature set. However, SHV can provide a poor user experience with graphics- and compute-intensive applications for such devices.

To provide Intel employees with an optimal user experience, we are investigating a solution that includes developing a private cloud that would provide desktop services using both server-hosted and client-side capabilities. This type of client-aware cloud contrasts with a server-centric cloud that simply offers centrally executed services that would not provide an optimal user experience in all use cases. With a client-aware cloud using both CHV and SHV, the robustness of the user experience can scale up or down depending on which client devices employees use.

IMAGE MANAGEMENT

An essential aspect of client-aware desktop virtualization is the need for centralized image and profile management. The virtual client image may consist of one or more components, including the OS, applications, user settings, and user data. Being able to centrally manage the virtual client image provides IT with the benefits of management and security, as well as the flexibility to meet expanding user needs.

Web Services

As we develop Intel's private cloud, we envision services that can determine the type of device accessing the service, as well as various attributes of the device.

Devices differ widely in their attributes, such as screen size, primary input method, processing power, and available memory, and in capabilities such as performance, security, and portability. Device attributes can also be independent of the device type, such as location and other contextual information, or user preferences defined in a profile.

A client-aware private cloud can fine-tune services depending on these attributes, and our testing has shown that these attributes and capabilities can greatly affect the user experience.[1]

- A device may be more capable in one situation—such as when it is connected to a Wi-Fi* network—than in another—such as when it is connected to a weak cellular signal. By detecting the type of connection, a client-aware Web service could deliver content tailored to the available bandwidth and latency.
- A device may have advanced capabilities, such as a GPS receiver, compass, or accelerometer. Client-aware Web services could execute additional code based on which native functions are available on the device.
- An employee with a PC based on the newest Intel® Core™ vPro™ processor might enjoy 3-D graphics, while a tablet user might see a flat 2-D graphic.
- An employee might be presented with a list of available printers, based on preferred features specified in the user profile or based on location.
- A client-aware Web service could have the ability to detect if an employee is located within a corporate facility. If so, the service would allow access to confidential data. But if the employee is located outside of a corporate facility, the service would allow access only to less sensitive data.

By developing client-aware services, we can provide a user experience aligned with the capabilities and context of each client device that accesses the services.

1 See the IT@Intel white papers "Cloud Computing: How Client Devices Affect the User Experience" and "Better Together: Rich Client PCs and Cloud Computing."

NEW CLIENT-AWARE TECHNOLOGIES

We plan to implement several new and evolving client-aware technologies in our private cloud. These technologies will help us deliver enterprise applications to Intel employees in a more efficient manner, enhancing IT efficiency and optimizing the user experience.

The goal of all these client-aware technologies is to balance computing between the client device and the server. As shown in Figure 1, using client-aware technologies enables centralized, layered virtual image management and intelligent decisions about whether local or remote execution is best, given a specific device and context.

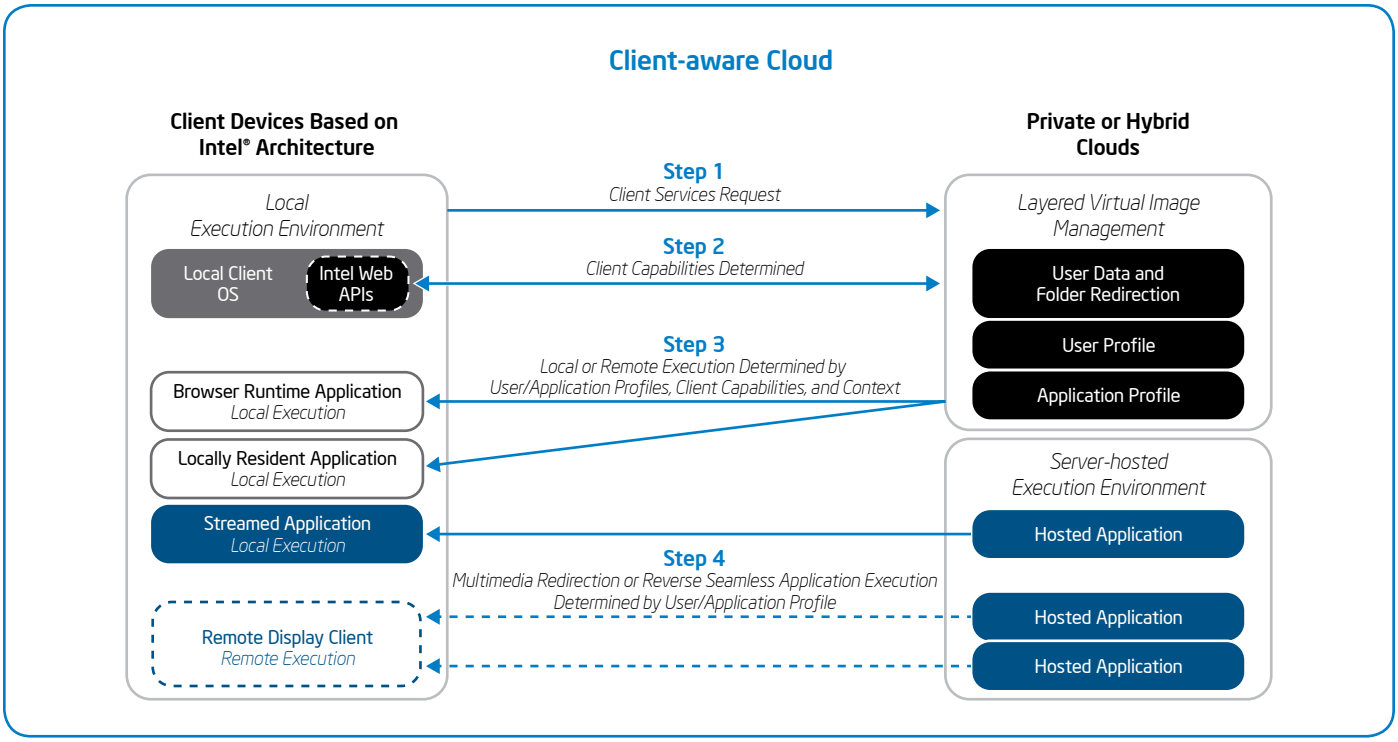

Figure 1. When a server receives a client-side request, client-aware technologies enable the cloud service to decide what should execute on the server and what should execute on the client.

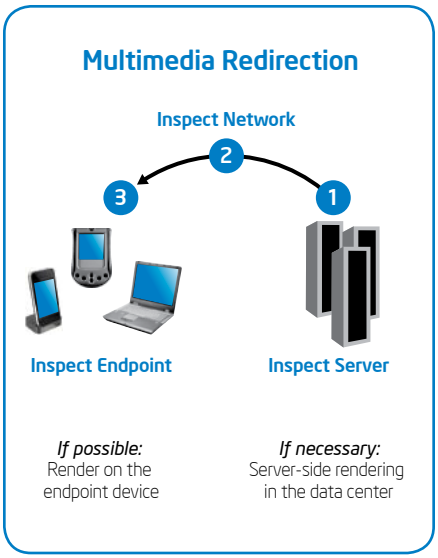

Figure 2. Multimedia redirection inspects the client device to determine if it is capable of rendering an image locally—which speeds application performance.

Improving Client-aware SHV with Workload Shifting

With the addition of client-aware features, SHV could become a more viable alternative compute model for delivering desktop services in some use cases. These features include techniques that shift some of the workload from the server to the client and improve image management.

WORKLOAD SHIFTING

Increasingly, cloud workloads use several approaches to shift graphic- or compute-intensive workloads from the server to the client. For intelligent clients with ample processing power, such as mobile PCs based on the latest generation Intel® processor, workload shifting can result in better application performance—translating to a better user experience.

Multimedia Redirection

For multimedia-intensive workloads such as video and audio, multimedia redirection (MMR) can speed application performance and reduce network traffic and server load. Figure 2 illustrates how, instead of rendering an image on the server and sending uncompressed images to the client, the server can use MMR to send compressed data to a robust client, where the image renders more quickly.

Command Remoting

Similar to MMR, command remoting sends commands to the client for execution. Command remoting can potentially further enhance application performance and decrease network bandwidth usage. For example, graphically intensive Web-based applications with mapping overlays benefit from client-side execution of certain application programming interface (API) commands. We anticipate that future versions of common office applications will require support for a growing set of API commands.

Reverse Seamless Technology

Historically, some end users who become frustrated with slow SHV application performance have resorted to running graphically intensive multimedia applications—such as video streaming, video conferencing, and Voice over IP (VoIP)—locally on their PCs. While this speeds application responsiveness, switching between their local PC desktop and their server-hosted virtual desktop creates discontinuity in the user experience and can be aggravating and inefficient.

Reverse seamless technology addresses this problem by presenting a unified view of the server-hosted virtual desktop and locally executing applications. IT can configure the applications through user device profiles. The profile identifies applications that should execute locally; when the user invokes those applications from the server-hosted desktop, the application is invoked from the client PC and "seamlessly" displays inside the virtual workspace. The application appears to be running in the virtual session when in reality it is running on the local machine.

For example, in Figure 3, the user profile specifies that the text editor should run remotely but that the video conferencing application should run locally. Because the video conferencing application will execute locally, it will be able to leverage the compute power and capabilities of the client system, and the application user interface will not be confined to the boundaries of the virtual desktop window. The photo-editing application in Figure 3 was invoked from the virtual desktop and, based on the user profile, is run on the local desktop. It will therefore behave in the same manner as the video-conferencing application.

Application Code Offloading

Hybrid enterprise applications contain code that can run on either the application server or on the client device. For certain devices, application code offloading can enrich the user experience by using awareness of local device features or device location to provide additional capabilities.

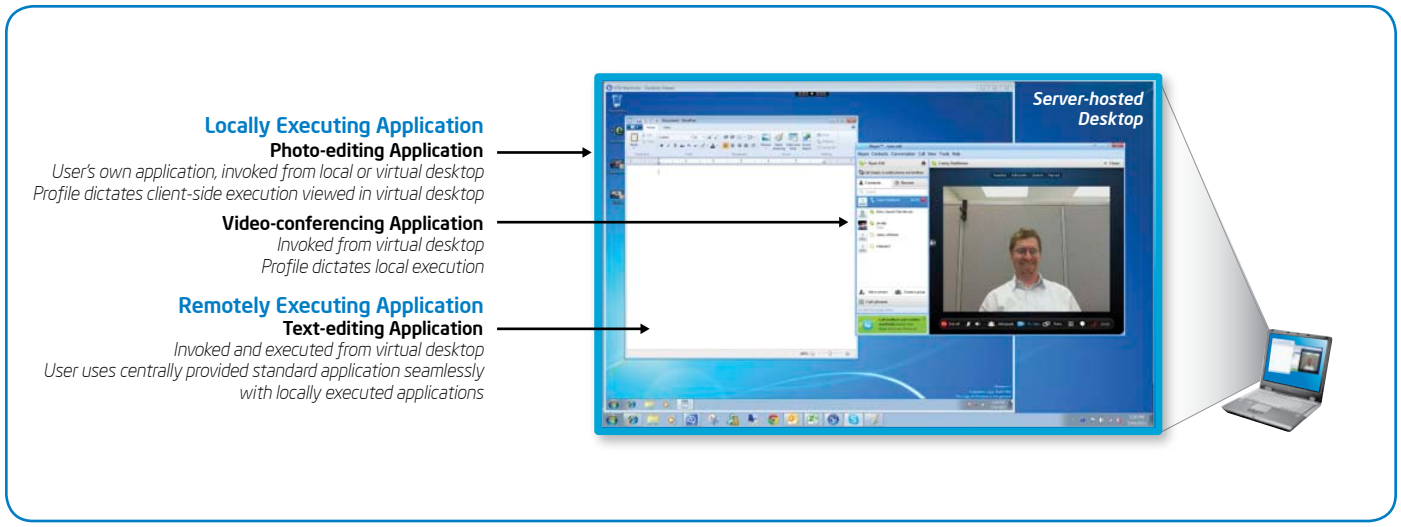

Figure 3. Reverse seamless technology enables users to launch and interact with local applications, without having to switch back and forth between the local desktop and the virtual session.

For example, an enterprise resource planning application could present a different set of functions depending on whether an employee was on the production floor, in the warehouse, or in the office. Using application code offloading allows application developers to write the code only once, but enable implementations for several platforms.

IMAGE MANAGEMENT

Typically, one weakness of SHV is that it uses a monolithic approach to image management, requiring IT to manage a single image of the entire user desktop, including the OS, applications, user data, settings, and profiles. This increases data center storage costs because each client image is stored separately.

Client awareness can improve the efficiency of image management for SHV by enabling layered virtual machine (VM) image management, which supports desktop service delivery across a variety of client platforms. With this approach, the core OS, applications, user data, and user settings and profiles are all separate and managed independently, but are then dynamically reassembled for the user as necessary.

This allows us to use a single stored instance of the core OS, for example, for thousands of users—instead of storing a copy in every user's VM image. By being able to discern the client, IT can use predefined device profiles to decide which parts of the image to deliver to the client in specific contexts.

Improving Application Responsiveness with Rich Internet Applications

Cloud application developers are increasingly looking for a means to offer rich Internet applications (RIAs), which distribute processing between the cloud and the client device to improve application responsiveness. With RIAs, application code downloads to the client device and executes on the client using a RIA software framework, which is typically based on HTML5. We anticipate the continuing emergence of client-aware applications that can discover the capabilities of hardware, using mechanisms such as Intel® Web APIs.

Intel is investing in APIs that allow developers to write RIAs that detect real-time hardware information from the client, such as processor performance, battery life, and network bandwidth. Applications can use this information on a dynamic basis, and the limitation of developing applications for the lowest common denominator is removed.

Products exist today that can use this type of information to make intelligent decisions about the user experience and how applications should run. For example, if employees were taking a certification test in a classroom, the instructor used to have to go around the room and make sure all employees' laptops were plugged into power before taking the test. However, with Intel Web APIs, now a Web application can determine if a laptop is not plugged in. When an employee tries to launch the test application without plugging in the laptop, the application can give an error message, such as "Your laptop must be plugged in to launch this application." Once the student plugs in the laptop and clicks "Retry", the Web application can detect that the laptop is now plugged in, and will start.

Intel® IT Innovation Centers

Intel® IT Innovation Centers provide a place to bring together people, ideas, technologies, and possibilities to generate new concepts and foster Intel's next-generation innovations. The Innovation Centers support the hands-on demonstration of Intel® products and platforms, and foster engagement between internal business groups and with external industry partners. There are six Intel IT Innovation Centers worldwide.

Intel IT's evaluation of client-aware technologies occurs primarily at the Folsom Innovation Center in California. The Folsom Innovation Center features a state-of-the-art executive briefing center equipped with the most up-to-date collaboration technologies, including touch-screen digital signage, smart TVs, electronic whiteboards, video conferencing, and Intel® Wireless Display. The showcase area highlights a range of the latest and emerging Intel® technology demonstrations relating to new tablets and smart phones, video streaming media, context-aware sensors, solid-state storage, in-vehicle infotainment systems, and virtualization involving streaming of virtual desktops, OSs, and applications.

Intel IT Innovation Centers are where dialogue and interaction occur through collaboration with industry peers, community, and academia. Innovation Centers support Intel technology and industry events, such as the Intel Developer Forum (IDF) and the Intel Sales & Marketing Conference (ISMC), through hosting innovation workshops and technology demonstrations. Intel IT Innovation Centers are considered the "place where innovation is put into action."

Here are additional examples of how RIAs can improve user experience and application performance:

- A Web site could deliver a multimedia version to customers on a high-speed Wi-Fi network and a simplified version to customers on a weak cellular connection.
- A Web site could query which graphics applications are installed locally, the type of CPU and screen size, and how much free memory is available. Using that information, the application could decide whether to launch an advanced graphics editor or a simpler application with lower-quality graphics and fewer editing capabilities.

ANTICIPATED BENEFITS

By using technologies such as workload shifting and RIAs to enable a client-aware private cloud, we anticipate benefits for both users and IT.

User Benefits

Each of the client-aware technologies we are exploring for use with our private cloud have the potential to increase the productivity and efficiency of Intel employees by detecting and utilizing client capabilities, and tailoring services to specific devices and contexts. This can result in faster application performance and better responsiveness on client devices with ample processing power and memory. Therefore, users have a better experience—with greater video and audio fidelity and more-responsive application interaction—and increase their productivity.

In addition, technologies such as workload shifting and RIAs balance server workload and network utilization, providing a more consistent and predictable experience for all.

IT Benefits

Client-aware technologies such as workload shifting and RIAs offer multiple benefits to IT:

- **Higher virtualization density in the data center.** These technologies have the potential to increase server density by distributing graphic rendering and computation to the client.
- **Reduced network traffic.** Because they shift workloads to the client, these technologies also reduce the load on the network and on the cloud infrastructure. Some third-party reverse seamless solutions enable IT administrators to determine whether it makes more sense to run an application locally or in the data center, based on the application's network intensity.
- **Improved security.** Some third-party reverse seamless solutions add another layer of security by enabling IT to host more sensitive applications in the data center, while simultaneously allowing users the freedom to run local applications. IT can control which data can and cannot move between the local and remote sessions by performing access control and application execution from the management console.
- **More efficient image management.** With an increasing number of end-user devices, IT must manage a growing set of operating environments, applications, profiles, and data. Use of client-aware image and profile management enables IT to better manage this increased complexity and deploy services faster. Image management is an essential requirement with any client virtualization method. Based on the user device and environment factors, IT administrators can better manage the increased complexity.

INVESTIGATIONS AND PRELIMINARY RESULTS

We are performing several pathfinding explorations in our lab and conducting a proof of concept (PoC) and other studies. The results of these preliminary investigations are promising, and we plan to further evaluate these client-aware technologies and their usefulness for our private cloud and desktop services delivery.

Workload Shifting

We are currently conducting a PoC that tests accessing a virtual desktop hosted on a server from laptop PCs and small form factor devices such as smart phones and tablets. This PoC uses a third-party workload shifting application and command remoting to take advantage of local device capabilities when possible. Although it is too early in the PoC to publish results, we anticipate an improved user experience, increased VM density, and better network bandwidth distribution.

We have also tested MMR in the lab. MMR can increase the number of virtual desktops a server can support by offloading the execution of multimedia content to a capable client. In our tests, we found that enabling MMR can increase server utilization anywhere from 15 to 40 percent, depending on the hypervisor under test.

With reverse seamless technology, our HTML5 testing has shown that this technology, running on PCs powered by Intel Core vPro processors, can merge the benefits of centralized virtual desktop infrastructure (VDI) processing and management with local client computing resources. As shown in Figure 4, reverse seamless significantly affects how fast images render on the client device.[2]

2 See "Intel® Cloud Builders Guide to Cloud Design and Deployment on Intel® Platforms."

Rich Internet Applications

We ran an enterprise customer relationship management application, provided through a public cloud service. We compared the performance of a mobile business PC and an entry-level desktop using both a traditional cloud-based application and an RIA from the same supplier. The mobile business PC executed the script sequence faster in both tests than the entry-level desktop. The RIA provided the highest performance and best user experience—executing the scripted sequence in approximately half the time compared to the traditional cloud application.[3]

Another PoC we are conducting explores how we might provide GPS-based information such as conference room and printer locations to non-Intel-managed smart phones.

3 For more information about this test, refer to "Cloud Computing: How Client Devices Affect the User Experience."

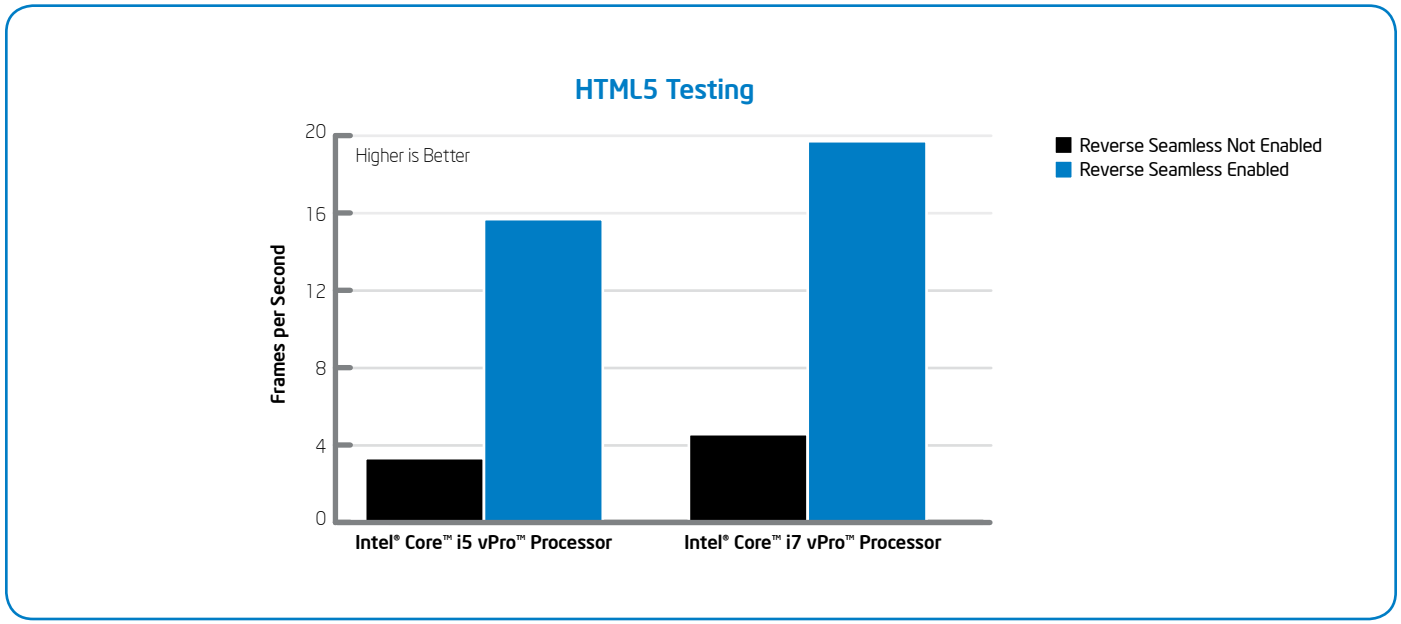

Figure 4. Reverse seamless technology can increase the number of frames rendered per second on the client, thereby increasing user experience and productivity. Intel internal testing July 2011, using a publically available benchmark.

CONCLUSION

To enhance management, mobility, and security while increasing employee productivity, we are exploring client-aware technologies that can support our efforts to develop a private cloud.

Client-aware technologies, such as workload shifting and RIAs, can help boost employee productivity and improve the user experience by taking advantage of client capabilities. For employees using companion devices such as smart phones and tablets, a private cloud that is aware of the device type, user location, user preferences, and other information can deliver services that are optimized and tailored to that user and device.

These same client-aware technologies also offer benefits to IT, potentially enabling us to increase server virtualization, decrease network traffic, enhance the efficiency of image management and application development, and expand our portfolio of SHV solutions without negatively affecting the user experience.

By continuing to monitor evolving client-aware technologies, and planning for a more diverse set of client devices, Intel IT is positioning itself to reap benefits both from the scalability of cloud computing and client device innovation.

FOR MORE INFORMATION

Visit www.intel.com/IT for white papers on related topics:

- "Cloud Computing: How Client Devices Affect the User Experience"
- "Better Together: Rich Client PCs and Cloud Computing"
- "The Future of Enterprise Computing: Preparing for the Compute Continuum"
- "Preparing the Enterprise for the Impact of Alternative Form Factors"

Cloud Builders Web site
www.intel.com/cloudbuilders

ACRONYMS

API	application programming interface
CHV	client-hosted virtualization
GPS	Global Positioning System
IDF	Intel Developer Forum
ISMC	Intel Sales and Marketing Conference
MMR	multimedia redirection
PoC	proof of concept
RIA	rich Internet application
SHV	server-hosted virtualization
VDI	virtual desktop infrastructure
VoIP	voice over IP
VM	virtual machine

For more information on Intel IT best practices, visit www.intel.com/it.

Printed in USA

Please Recycle

0811/JLG/KC/PDF

325585-001US