

The Future of Enterprise Computing: Preparing for the Compute Continuum

Our goal is to provide greater productivity and flexibility for Intel employees by enabling them to access information and IT services from multiple devices, whether personal or corporate-owned.

Executive Overview

Enterprise client computing is at an inflection point. A number of trends and pressures are driving a transition from the traditional client computing model toward a future in which employees will use a variety of devices to access information anywhere, at any time.

Intel IT is shifting focus to deliver certain services to any device—and to multiple devices for any employee. By taking advantage of a combination of technologies and trends—such as ubiquitous Internet connectivity, virtualization, and cloud computing—we have an opportunity to meet changing user requirements and redefine the way we provide services.

We call this vision the Compute Continuum.

Our goal is to provide greater productivity and flexibility for Intel employees by enabling them to access information and IT services from multiple devices, whether personal or corporate-owned. These devices will work together seamlessly, while enterprise security is protected.

To deliver these capabilities, Intel IT established a program in late 2010 to chart a path from today's traditional client computing model to the Compute Continuum.

Our strategy includes three overlapping phases:

- **Supporting IT consumerization.** In response to employee productivity needs and technology expectations, Intel IT has begun to provide access to e-mail, calendar, and other business applications from personally owned devices. Our focus in 2010 was on smart phones, and we also plan to deliver services on larger personal devices.

- **Delivering IT as a service.** We are developing and testing capabilities that let us deliver our IT applications and environment as services that run on a variety of personal and corporate devices, rather than as a single integrated hardware and software platform.
- **Delivering the Compute Continuum.** Our goal is that users will be able to securely access services anywhere, at any time, from a growing variety of personal and corporate devices. An increasing number of these services will be delivered from our private cloud and from public clouds.

We are currently developing Compute Continuum use cases in several of the areas that are most important to our users. These include blended use models—such as corporate and personal environments on the same device, and bring-your-own computer—as well as enhanced enterprise collaboration.

By building on our existing support for IT consumerization and exploration of device-independent service delivery models, and taking advantage of new trends and technologies, we anticipate that our Compute Continuum program will result in significant benefits for our users and for Intel.

Dave Buchholz
Principal Engineer, Intel IT

John Dunlop
Client Architect, Intel IT

Contents

Executive Overview.....	1
Business Challenge	2
Toward the Compute Continuum.....	2
Formation of Intel IT Compute Continuum Program.....	3
The Path to the Compute Continuum	3
IT Consumerization.....	3
Delivering IT as a Service	4
Compute Continuum	6
Conclusion.....	8
For More Information.....	8
Acronyms.....	8

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.

BUSINESS CHALLENGE

Enterprise client computing is at an inflection point. A number of trends, pressures, and new technologies are driving a transition from traditional client computing models toward a future in which employees will use a variety of devices to access information anywhere, at any time.

The need to support IT consumerization is a key factor. Like other IT organizations, Intel IT used to control how employees adopted technology. Now, employees are a major influence on IT's adoption of new technology.

Many employees want to use their own devices to access information. The number of handheld devices continues to increase rapidly in our environment: Many employees at Intel already have one or more devices in addition to their mobile business PCs and are looking for IT to deliver information to all of these devices. By responding to this need, we can enable employees to work in more flexible and productive ways.

This requires a significant change in the way we provide IT services to client devices. To date, Intel IT has typically focused on delivering a build—an integrated package comprising an OS, applications, and other software—to a single PC.

As employees use a wider range of devices, we need to shift focus to delivering services to any device—and to multiple devices for any employee.

This makes managing technology and security more complex. It also introduces

issues for Intel Legal and Human Resources (HR) groups since we are providing access to Intel-owned information and applications from devices that are owned by users.

Toward the Compute Continuum

By taking advantage of a combination of technology trends and emerging compute models—such as ubiquitous Internet connectivity, virtualization, and cloud computing—we have an opportunity to proactively address changing user requirements and redefine the way we provide services.

We call this vision the Compute Continuum. We believe this represents the next major change in the way that employees will use technology. Key aspects of the Compute Continuum include:

- Users have access to corporate information and IT services from any device, whether personal or corporate-owned.
- Multiple personal and corporate devices work together seamlessly.
- Corporate information and services are delivered across these devices while the enterprise continues to be protected.

Overall, we anticipate that employees will enjoy a rich, seamless, and more personal experience across multiple devices. They will be able to move from device to device while retaining access to the information and services they need. Their experience will vary depending on the characteristics of the device they are using; services will be context-aware, taking advantage of higher-performing client hardware to deliver an enhanced experience.

Formation of Intel IT Compute Continuum Program

Delivering the Compute Continuum is a priority for Intel IT. In recognition of this, in late 2010 we created a program chartered to determine how to deliver the Compute Continuum to Intel’s users.

We also are documenting IT-specific requirements and use cases to provide input to Intel product groups and to Intel’s suppliers, to help them deliver Compute Continuum capabilities that are ready for enterprise use.

THE PATH TO THE COMPUTE CONTINUUM

We are taking advantage of a range of new technologies and computing trends—including ubiquitous Internet connectivity, cloud computing, and virtualization—to make the transition to the Compute Continuum.

A key aspect of the transition is a shift toward delivering services across multiple devices rather than on managing client hardware. The devices may range from mobile business PCs to smart phones, tablets, in-car systems, wireless displays, and projectors, as shown in Figure 1.

We anticipate that transitioning to the Compute Continuum, which is already underway at Intel, will be completed in three related phases:

- Supporting IT consumerization
- Delivering IT as a service
- Delivering the Compute Continuum, with increasing use of cloud-based services

IT Consumerization

As new and innovative smart phones and other mobile devices rapidly proliferate, an increasing number of Intel employees

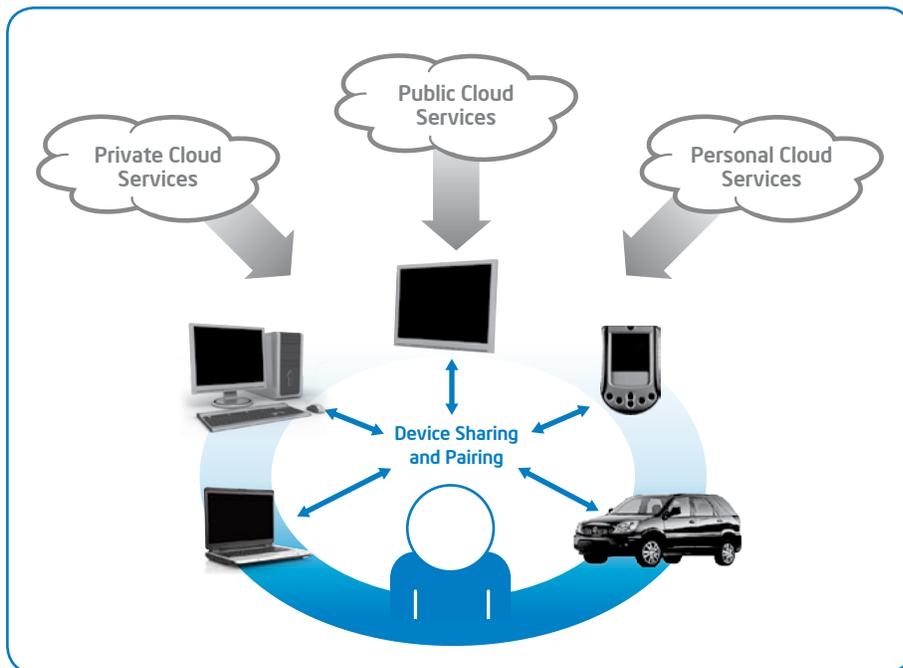


Figure 1. The Compute Continuum: Devices work together to enable a common user experience.

want to use their personally owned devices within the enterprise. They want to choose the platforms, applications, online tools, and services they use to accomplish their jobs as well as manage their lives outside work. This trend is often referred to as IT consumerization.

We recognized several years ago that letting employees use their personal devices could increase employee productivity and satisfaction while helping to attract and retain talented people. By providing employees with access to enterprise information and applications anytime, anywhere, we enable them to work in more flexible and productive ways.

In early 2010, we began letting employees use personal smart phones to access enterprise services such as e-mail and calendar. This represents the first phase of our transition to the Compute Continuum. Previously, we only allowed access from IT-supported corporate-owned devices.

Enabling use of personal devices was challenging because it required a change in mindset not only by Intel IT but also by Intel’s Legal and HR groups. Existing policies assumed that only corporate-owned devices would access Intel information.

Working closely together over an 18-month period, Intel’s IT, Legal, and HR groups formulated new policies that protect Intel while enabling the use of personal smart phones and other devices to access Intel corporate information. The new policies are codified in a connection agreement that defines users’ responsibilities and sets expectations, and is signed by all mobile device users and their direct managers.

Intel employees have rapidly adopted the personal device program, and the program has delivered significant business value. After just over a year, employees are using more than 10,000 personal devices to access corporate information. The groundwork in establishing new policies is enabling us to

Intel IT History of Deploying Consumer Technology

Intel IT has a history of applying technology that enables employees to work in new ways, with the goal of providing greater flexibility and productivity. In some cases, we have been among the first large organizations to deploy technologies that were previously used primarily by consumers.

- During the 1990s, we were one of the first large enterprises to widely deploy mobile business PCs, which are now used by about 80 percent of our employees.
- In 2001, we began implementing wireless LANs (WLANs) at a time when they were primarily employed for home use. We now have WLANs in all campuses worldwide.
- In 2002, we began deploying instant messaging. We delivered a full enterprise production release the following year.
- Starting in 2009, we began deploying a variety of internal social media technologies, including wikis, blogs, and forums.
- In 2010, we began allowing use of employee-owned smart phones and other personal devices in the enterprise.

We are now focused on delivering the Compute Continuum, which we believe represents the next major change in the way that employees use technology.

quickly support additional personal devices, such as tablets, non-standard PCs, and a wider range of smart phones.

Initially, we provided access to corporate e-mail, calendar, and contacts, because those were the capabilities employees most frequently requested. We are now identifying additional enterprise applications and services that we can make accessible to personal devices. For example, we are planning or considering applications that provide conference room booking, network printing, travel booking, and employee benefits.

Delivering IT as a Service

We are beginning the second phase of our transition to the Compute Continuum: Developing capabilities that let us deliver our IT applications and environment as services rather than as a single integrated hardware and software platform.

By developing this device-independent service delivery model, we are building the software foundation for the Compute Continuum—creating a software infrastructure that will make corporate applications and user data available across multiple devices.

Device independence provides the ability to deliver services not only on current devices, but also on new categories of devices that may emerge in the future. We expect that the number and variety of handheld and tablet devices will continue to increase rapidly—and that our employees will want to take advantage of these devices, in addition to the devices they already use.

Depending on a device's capabilities, it may be able to run multiple environments, including separate corporate and personal workspaces.

CLIENT VIRTUALIZATION TECHNOLOGIES AND USES AT INTEL

We are taking advantage of virtualization to provide device independence by abstracting IT services from the underlying hardware. We deliver an application or an entire environment in a virtual container; this enables the service to run on any device that supports the virtualization software. We can also run multiple environments and applications, each isolated within its own container, on the same system.

This allows faster development and introduction of new capabilities at lower cost, because we do not need to certify the OS and each application for every hardware platform. This approach can also reduce IT management cost, because we manage only the virtual containers rather than each hardware platform.

Client virtualization encompasses a range of technologies. These technologies, summarized in Table 1, are grouped into several categories: client-hosted virtualization (CHV), including Type 2 and Type 1 (bare-metal) client-side hypervisors; server-hosted virtualization (SHV); and application virtualization. We anticipate using multiple technologies, depending on the requirements of each use case and the capabilities of the device. Some of these technologies are already in production use at Intel IT; most are at the proof-of-concept (PoC) or evaluation stage.

Type 2 client hypervisors. Type 2 hypervisors run as a software process on a host OS on the client system; they support one or more guest OSs in virtual machines (VMs).

We are currently using Type 2 hypervisors to provide an Intel IT software environment for contractors who develop software for Intel. Previously, to make this environment available to them, we needed to provide

Table 1. Technologies Supporting the Compute Continuum

	Definition	Current or Possible Future Client Computing Use	Enabling Platform Technologies
Type 2 Hypervisor	Hypervisor that runs as a software process on a host OS. Hypervisor can support one or more guest OSs in virtual machines (VMs).	Already in use to provide Intel software development environment on contractors' own PCs.	<ul style="list-style-type: none"> Intel® Virtualization Technology (Intel® VT)
Type 1 Hypervisor	Bare-metal hypervisor that runs directly on the client platform with no host OS. Can support multiple independent OS environments in VMs. Considered to provide greater security and performance than Type 2 hypervisor.	Support personal and corporate workspaces on single mobile business PC; support high-security and standard environments on single PC.	<ul style="list-style-type: none"> Intel VT, including Intel® Virtualization Technology® for Directed I/O (Intel® VT-d) Intel® Trusted Execution Technology (Intel® TXT)
Application Virtualization	Applications are streamed on demand from private cloud (data center) or public cloud to the client, where they execute in a container.	Delivery of specific horizontal and vertical enterprise applications to various client devices, including non-standard and personal PCs.	
Virtual Hosted Desktop	The client environment executes in a virtualized container on a server in a private or public cloud, and is accessed over the network from the user's device.	Delivery of entire client environment to non-standard and personal PCs.	<ul style="list-style-type: none"> Newer protocols can take advantage of current Intel® Core™ processors, including Intel® HD Graphics, on client system to improve performance.
Cloud Computing	A highly available computing environment where services are delivered on demand using a multi-tenant, elastic infrastructure.	Transition to private cloud for office and enterprise computing; use of public clouds for specific applications.	
Context Awareness	Leveraging information about the user's situation to influence the user's experience.	Helping employees find available resources such as conference rooms and printers.	<ul style="list-style-type: none"> Intel Web APIs

them with PCs running the Intel IT software build. Now we simply install a hypervisor on the contractor's own PC and deliver a streamlined development build on top of the hypervisor. To enable this, the PC must meet our minimum specifications, which include Intel® Virtualization Technology (Intel® VT) and a specific OS.

This approach reduces cost and support requirements while also reducing the risk to Intel: We provide the build within a secure, policy-managed virtual container that is fully encrypted, cannot be copied, and will destroy itself if the system does not regularly check in with Intel IT.

Type 1 client hypervisors. These are bare-metal hypervisors that run directly on the client platform without the need for a host OS. They can provide better security and performance than Type 2 hypervisors. However, on client systems, Type 1 hypervisors are less mature than Type 2 hypervisors. Accordingly, we are

researching and testing uses of Type 1 hypervisors but have not yet deployed them for production use. A small number of our users are evaluating the currently available technology to determine whether it is a good fit for the Intel environment.

As the technology matures, we see a number of potential uses within Intel IT. Type 1 hypervisors could be valuable for engineers who work with classified proprietary design information. We could implement two isolated environments on the same client PC: a highly secure environment used for proprietary design information and a standard enterprise environment. This would allow the user greater flexibility and productivity while protecting corporate intellectual property.

We could also use Type 1 hypervisors to implement and isolate personal and corporate environments on the same system.

Server-hosted virtualization. With SHV, software is stored on a server in a data center, and it executes in a container on the

server rather than on the client; the employee interacts with the server-based software over the network. SHV can be used to deliver an entire desktop environment or individual applications to capable client devices.

Traditional SHV approaches do not support mobility, can cause performance problems with compute- and graphics-intensive applications, and increase the load on network and server infrastructure. However, as SHV technologies mature, they are beginning to identify client capabilities and take advantage of them to improve the user experience. For example, newer protocols that are used with SHV can offload some multimedia processing to PCs rather than executing all of it on the server. This can reduce the impact on network traffic and take advantage of higher performing clients to deliver a better user experience.

We are planning a 150-user PoC in 2011 that uses SHV and application virtualization (described below) to deliver our software environment and specific applications to clients

that include personal and non-standard PCs. In the PoC, we plan to use SHV to deliver our corporate OS image.

Application virtualization. With application virtualization, applications are packaged and managed on a central server; the applications are streamed to the client device on demand, where they execute in isolated containers and are sometimes cached to improve performance.

In our PoC, we plan to use application virtualization to deliver core enterprise productivity applications as well as specific line-of-business applications to non-standard PCs and other personal PCs.

CLOUD COMPUTING

Cloud computing is an essential element of our strategy. We are developing a private cloud, built on shared virtualized infrastructure, to

support our enterprise and office computing environment. Our goal is to increase agility and efficiency by using cloud characteristics such as on-demand scalability and provisioning, as well as automated management. We are also selectively using external cloud services such as software as a service (SaaS) applications.

Over time, we anticipate that more of our IT services will be delivered from clouds, facilitating ubiquitous access from multiple types of devices. This will enable us to take advantage of cloud capabilities to broker and manage the connection to the client.

Compute Continuum

Over time, we plan to expand our current strategy to deliver the Compute Continuum to our users. This involves the progressive delivery of new capabilities, as shown in Figure 2, with

the goal of enabling users to securely access services anywhere, at any time.

Our multiyear strategy builds on the capabilities we are already implementing, such as support for IT consumerization, while taking advantage of new technologies—including device pairing for sharing resources among devices and greater levels of context awareness for a highly personalized user experience. We will continue to focus on device-independent service delivery, which will become increasingly important as the number and variety of devices increase. We anticipate a proliferation of new devices including non-traditional devices such as in-car computers and wireless displays connected by Intel® Wireless Display. A growing number of these services will be delivered from our private cloud and from public clouds.

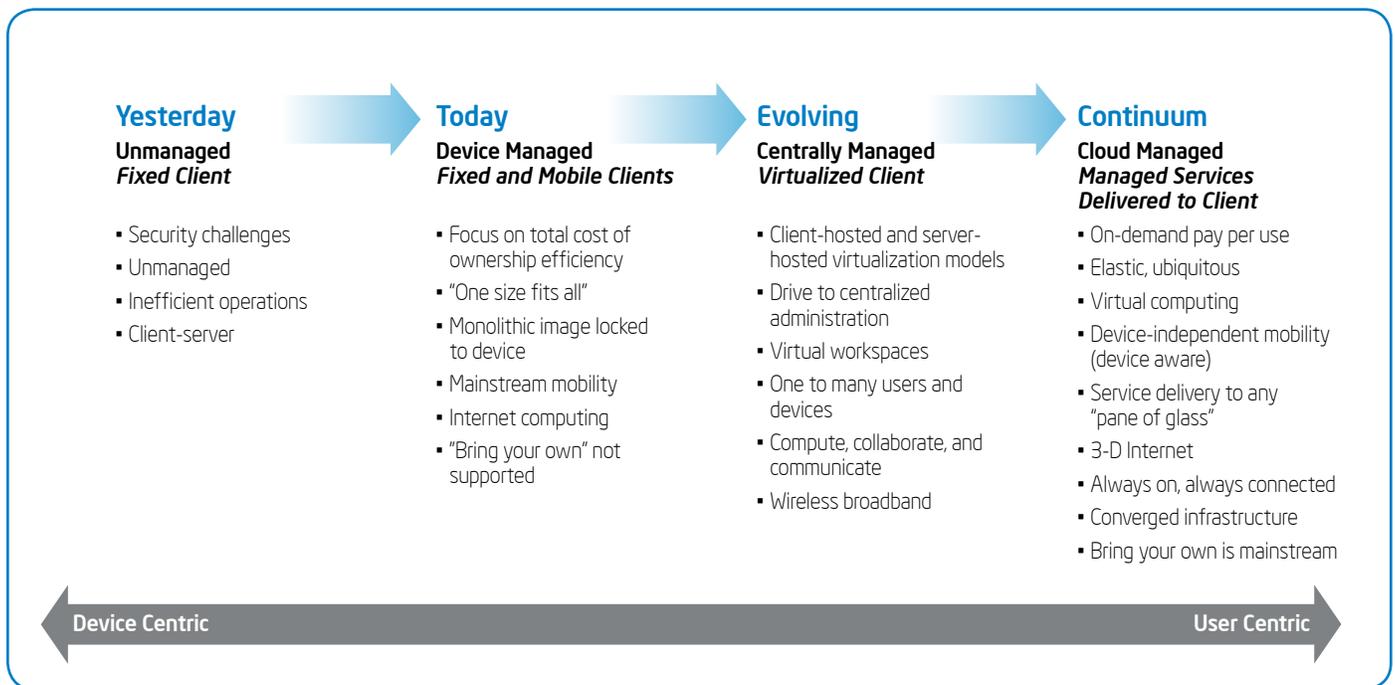


Figure 2. Intel IT path to the Compute Continuum.

We are currently developing Compute Continuum use cases in several of the areas that are most important to our users. These include blended use models—such as corporate and personal environments on the same device, and bring-your-own computer—as well as enhanced enterprise collaboration.

STRATEGY FOCUS AREAS

Delivery of the Compute Continuum affects multiple aspects of our IT strategy. Specific focus areas and challenges include platforms, applications, and security.

Platforms and Applications

Context awareness. We define context-aware computing as services that leverage information about users' situations to influence their experience.

Context-aware services capture information from sensors—determining, for example, user location with a global positioning system (GPS) receiver within the client device. This is combined with other information, such as user preferences, to deliver personalized information or services.

Within Intel IT, we are investigating ways to apply context-aware computing to help employees find available resources such as conference rooms or printers with preferred features. To achieve this, we anticipate that context-aware services will gather information from multiple sensors in client devices, IT infrastructure, and facilities.

Over time, we expect more context-aware sensing capabilities to be built into client platforms. Intel researchers have been working on enabling devices to generate and use contextual information for a greatly

enhanced user experience while helping to ensure the safety and privacy of personal information. This requires services that are also aware of the client hardware specifics, network latency, and other aspects of the environment.

We anticipate the continuing emergence of client-aware applications that can discover the capabilities of client hardware, using mechanisms such as Intel Web APIs, and take advantage of them to deliver the best user experience. For example, a user may see 3-D graphics over the Internet when using a PC based on the latest Intel® processors, but flat 2-D graphics when using an older or lower performing device.

Common development framework. One Intel IT goal is a common development framework that will enable us to quickly develop applications across the spectrum of OSs and devices—without having to focus on the specifics of each platform. Today, we sometimes need to develop multiple versions of an application to provide the same native functionality on different handheld devices.

A related issue is deciding how much effort to invest in customizing software to provide users with an experience that is consistent with other applications on the same device.

Application delivery. To successfully deliver applications across the growing range of user devices, we will need to select application distribution methods that make the most sense for Intel IT and employees. Intel IT, like most IT organizations, typically delivers applications using internal software distribution tools. However, users with personal handheld devices typically obtain their applications from proprietary application

stores that focus on a specific OS. As we develop more applications for handheld and other non-traditional devices, we will identify which of these application distribution methods are most suitable.

Security

IT security policies. Policies need to evolve to comprehend management of non-IT supported devices. This represents a significant challenge for many organizations. At Intel, we have already begun to address this by developing policies that allow use of personal smart phones and other devices.

Multi-level trust model. Intel IT is developing a new security architecture that supports the delivery of IT services and information from our private cloud to a much wider variety of devices and locations, not all of which offer the same level of security. The new architecture moves away from today's binary (all-or-nothing) network access control to a policy-enabled, multi-level trust model. It is designed to dynamically adjust users' access privileges as their risk profile changes, depending on factors that include users' locations and whether they are using a trusted device such as a mobile business PC or an untrusted device such as a personal smart phone.

Streamlining device security. We need mechanisms that adequately secure data on consumer devices without making access cumbersome for users. Some of today's consumer devices do not automatically provide enterprise-class security for data and credentials stored on the devices. Additional protection is needed, requiring the user to perform additional authentication steps.

CONCLUSION

Intel IT has defined a multiyear strategy to deliver enterprise services to employees through the Compute Continuum, which we anticipate will result in significant benefits for our users and for Intel. This strategy builds on our current approach to IT consumerization and our exploration of device-independent service delivery models. We are taking advantage of a number of trends and technologies, including cloud computing and client virtualization.

The Compute Continuum enables greater flexibility and productivity for users by providing them with seamless access to information and IT services across multiple devices, whether personal or corporate-owned, while protecting enterprise security.

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

FOR MORE INFORMATION

Visit www.intel.com/it to find white papers on related topics:

- "Personal Handheld Devices in the Enterprise"
- "Maintaining Information Security While Allowing Personal Hand-Held Devices in the Enterprise"
- "Cloud Computing: How Client Devices Affect the User Experience"
- "Developing an Enterprise Cloud Computing Strategy"
- "Developing an Enterprise Client Virtualization Strategy"
- "Enabling Device-Independent Mobility with Dynamic Virtual Clients"

ACRONYMS

CHV	client-hosted virtualization
GPS	global positioning system
HR	Human Resources
Intel® TXT	Intel® Trusted Execution Technology
Intel® VT	Intel® Virtualization Technology
Intel® VT-d	Intel® Virtualization Technology for Directed I/O
PoC	proof of concept
SaaS	software as a service
SHV	server-hosted virtualization
VM	virtual machine
WLAN	wireless LAN

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel, Intel Core, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 2011 Intel Corporation. All rights reserved.

Printed in USA
0511/JLG/KC/PDF

 Please Recycle
324598-001US

