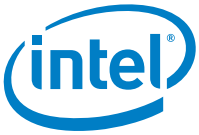

4th Generation Intel® Core™ vPro™ Processors with Intel® VMCS Shadowing

Enhancing the Performance of Citrix XenClient® and McAfee Deep Defender*

CONTENTS

Executive Summary

Responsive and secure desktop virtualization requires tight integration between the virtualization machine monitor (VMM) software that is used to deploy and manage virtual machines and the underlying hardware platform. Intel and Citrix have worked closely together over the past several years to optimize Citrix XenClient® so it takes fulladvantage of Intel® vPro™ Technology[1] to address this integration need.

New usage models are now emerging that can require two or more VMMs to be hosted on the same client system. McAfee Deep Defender* offers a prime example. This advanced security software complements the security safeguards built into XenClient, yet includes—and requires—its own VMM.

This paper is a technology preview that describes a new hardware-based capability known as Intel® Virtual Machine Control Structure (Intel® VMCS) Shadowing, which will be available with 4th generation Intel® Core™ vPro™ processor and describes the hardware-assisted security provided by XenClient, Deep Defender. Intel VMCS Shadowing can enable faster performance for multi-VMM usage models. Both Citrix and McAfee are evaluating this capability for inclusion in future product releases.

Hardware-Assisted Client Security from Citrix and Intel

Citrix XenClient® is a desktop virtualization solution designed to simplify the delivery and management of mobile client environments. Citrix offers two distinct product lines. XenClient Enterprise is built to support the needs of most large businesses. XenClient® XT provides a more targeted solution for federal government agencies and vertical security markets. Both versions support efficient delivery of personalized end-user client environments. They also enable multiple virtual desktops to be run transparently on a single client device, each with its own set of applications and its own instance of the operating system (OS).

XenClient is a bare-metal type-1 hypervisor that provides full system virtualization for x86 client devices. A bare-metal hypervisor, also known as a Type 1 hypervisor, is virtualization software that is installed directly onto the PC's hardware and runs below the PC operating system. IT administrators can use this hypervisor to partition client software components into separate virtual machine packages. The Microsoft Windows* OS, user applications, corporate applications, personalized data, and end-user profiles can all be deployed and managed separately, which helps to reduce network bandwidth and improve serviceability and security (Figure 1).

XenClient supports client security in a variety of ways. A feature known as Citrix Xen hypervisor domain disaggregation isolates I/O traffic from the guest operating system into separate service virtual machines, which helps to protect against various kinds of low-level malware attacks. XenClient also takes advantage of Intel® vPro™ Technology to deliver a number of advanced security capabilities.

Intel vPro Technology is a set of technologies included in select Intel® processors and chipsets. In addition to providing a better foundation for client security, many of these technologies also help to accelerate performance, enhance client management, and virtualization. A few of the capabilities provided by Intel vPro Technology include:

- **Intel® Trusted Execution Technology[2] (Intel® TXT)** measures the boot environment to establish a hardware-assisted dynamic root-of-trust. Intel TXT can be used so that each client system can only boot into a "known good state," which prevents boot-level and firmware-rooted malware.
- **Intel® Virtualization Technology[3] (Intel® VT-x)** provides hardware assists in Intel processors to help improve application performance and to enable the operating system to be completely isolated from the underlying hardware of the client device. It also provides hardware assistance for locking and wiping virtual machines.

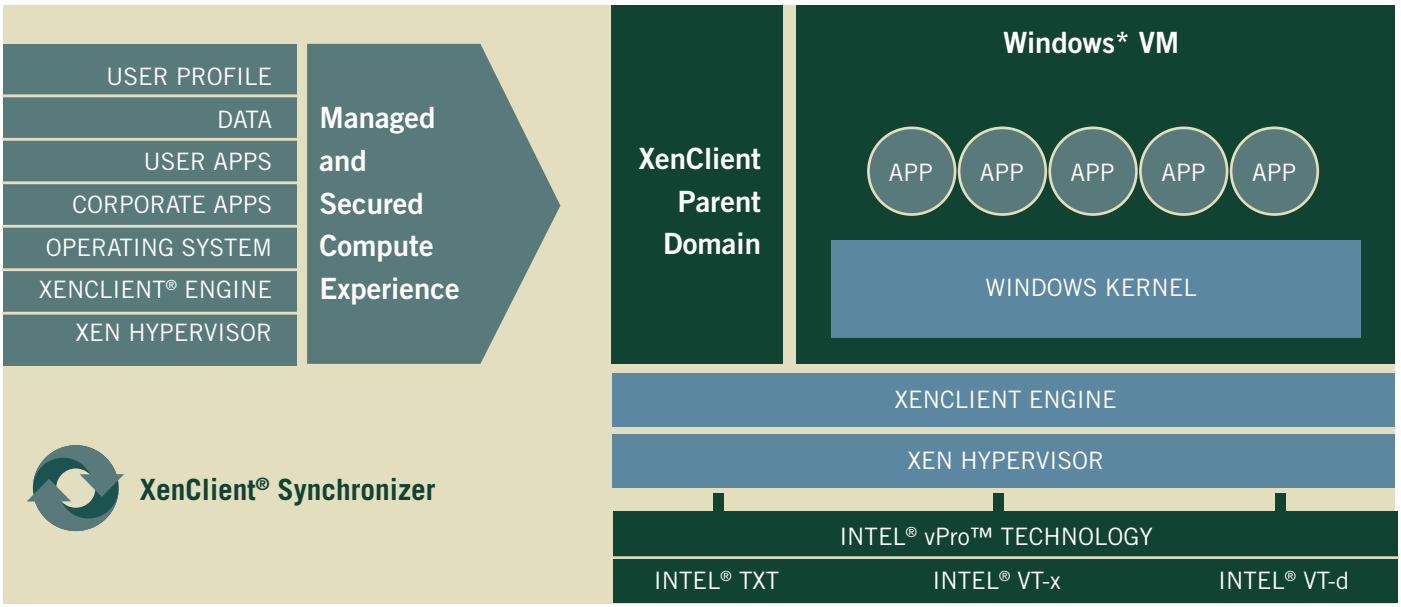

Figure 1. Citrix XenClient® supports flexible desktop virtualization to help IT organizations deliver a managed and secure computing experience to diverse individuals using a wide range of client devices.

- **Intel® Virtualization Technology for Directed I/O (Intel® VT-d)** helps to improve I/O efficiency in virtualized environments. It also enables stronger, hardware-assisted isolation of network, keyboard, disk, and I/O traffic, which helps to protect against a number of attack strategies, such as illegitimate key logging, screen capturing, and covert network channels.

XenClient XT takes advantage of Intel vPro Technology, Intel VT-x, Intel VT-d, and Intel TXT to:

- Provide full system isolation, monitoring, and control over platform hardware, including processors, memory, and I/O resources.
- Support Service VMs, which provide isolated access to corporate networks and storage systems.
- Measure and attest the entire boot system environment—hardware, firmware, and OS loader—to ensure the client system is not tampered with prior to or during launch.

Adding McAfee Deep Defender Technology for Enhanced Security

McAfee Deep Defender provides additional hardware-assisted security capabilities that can be used to supplement and strengthen the safeguards that are built into XenClient and into other, more traditional, security applications. McAfee Deep Defender resides between platform memory and the OS and provides real-time, kernel-level monitoring. It is designed to detect, block, and remediate advanced, hidden attacks, such as kernel-mode rootkits, stealth attacks, and zero-day malware.

McAfee Deep Defender provides this advanced protection using a form of system virtualization, which is furnished by a lightweight hypervisor, or Virtual Machine Monitor (VMM), known as McAfee DeepSAFE* technology (Figure 2). McAfee DeepSAFE technology can be thought of as a kind of micro-hypervisor.

Like XenClient, McAfee Deep Defender takes advantage of Intel vPro Technology and Intel VT-x to provide enhanced, hardware-assisted security and virtualization. However, unlike XenClient, McAfee DeepSAFE technology does not provide full system and I/O virtualization. Instead, it uses hardware-assisted virtualization to monitor and control memory and processor operations, which provides the foundational layer for McAfee Deep Defender security functions.

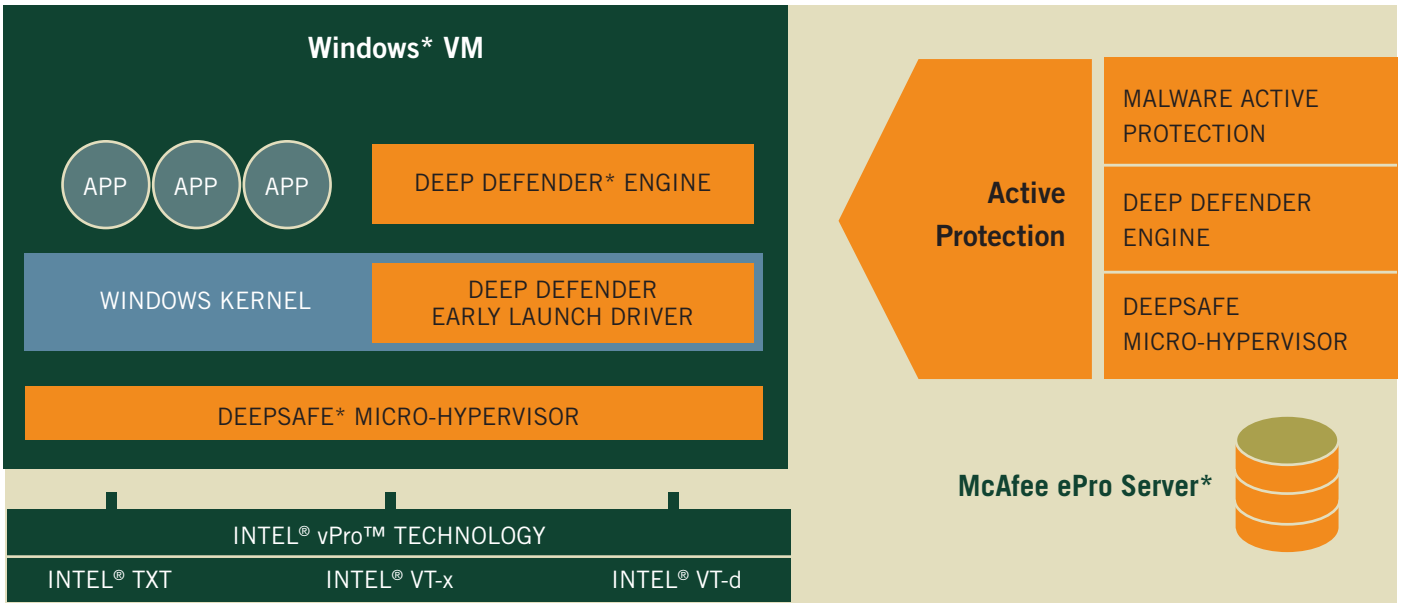

Figure 2. McAfee Deep Defender* uses a kind of micro-hypervisor, called McAfee DeepSAFE* technology, to enable advanced, kernel-level monitoring. It helps to protect against advanced, hidden attacks, such as kernel-mode rootkits, stealth attacks, and zero-day malware.

Together, XenClient and McAfee Deep Defender provide a breadth and depth of security that neither can provide alone. However, running these products together requires installing two VMMs on each client system, the XenClient type-1 hypervisor and the McAfee DeepSAFE technology micro-hypervisor. Both these VMMs require access to the underlying hardware and to the capabilities provided by Intel vPro Technology.

XenClient supports the hosting of two VMMs on a single platform through a feature called VMM nesting. Nesting allows a root VMM to support an Intel VT enabled guest VMM. The XenClient hypervisor functions as the root VMM and provides managed access to key hardware features for the McAfee DeepSAFE technology hypervisor, which functions as the nested VMM. Although this software-based nesting solution provides the necessary functionality for running XenClient and McAfee Deep Defender together on the same hardware platform, it introduces latencies that can degrade virtual machine performance and potentially impair the end-user experience.

The Performance Challenge: Software-Only VMM Nesting

To understand the performance challenges of software-based VMM nesting, it helps to understand how Intel VT can improve virtualization performance in a typical, single-VMM scenario. The VMM manages hardware resources and allocates them as needed to the various VMs running on the platform. Whenever a VM encounters a "privileged" software instruction that impacts shared resources, it hands over control to the VMM to perform the instruction. This helps to ensure there are no conflicts between the guest VMs. The process of transferring control from the guest VM to the VMM is known as a VM exit.

Once the privileged instruction is completed, the VMM hands control back to the guest VM, so the application in the VM can continue to run. This is known as a VM entrance. Intel VT helps to accelerate VM exits and entrances by providing a high-speed, hardware-based memory structure called a Virtual Machine Control Structure (VMCS). The VMCS holds the processor register states of the guest and the host. By providing low-latency access to this information, the VMCS enables VMM exits and entrances to be performed very quickly.

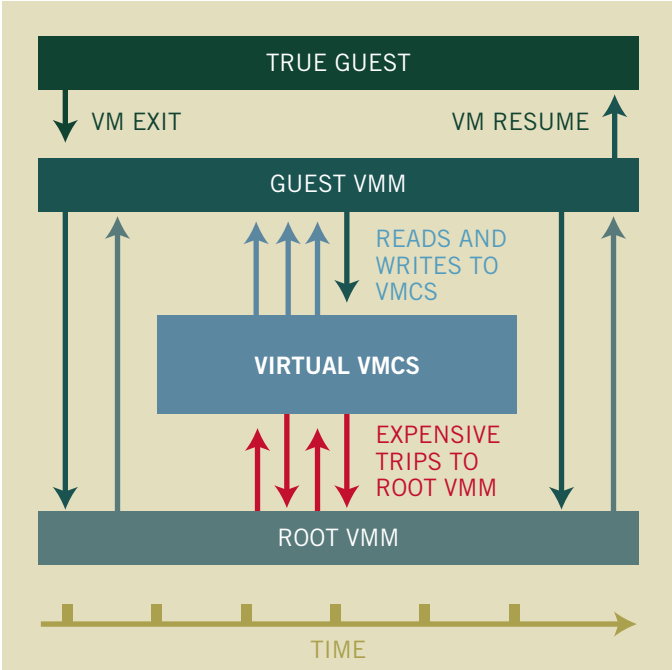

Figure 3. Software-based VMM nesting allows two VMMs to run on the same physical system, but this approach can result in significant performance penalties.

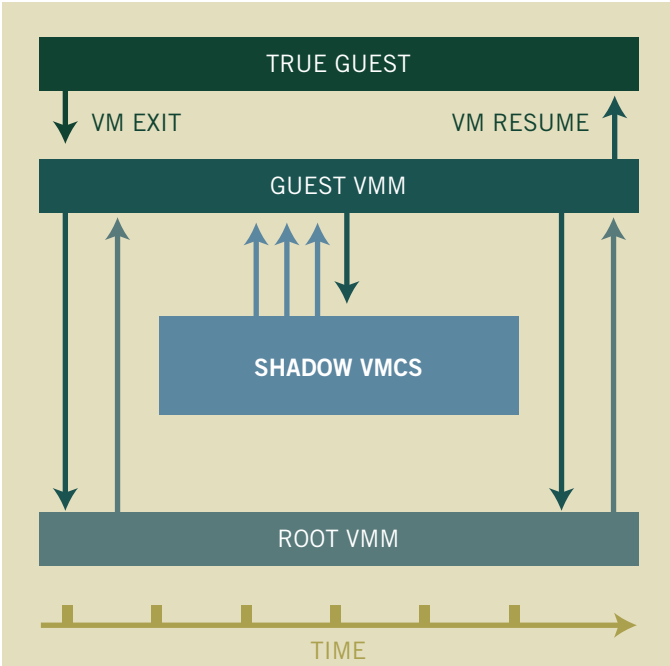

Figure 4. With Intel® VMCS Shadowing, two VMMs can be hosted on the same physical system, without the performance penalties of software-only nested solutions.

When VMM nesting is used, one VMM resides above another on a virtualized platform (Figure 3). The root VMM controls the physical VMCS, and the guest VMM is provided with a software-based "virtual VMCS." Although this approach allows multiple VMMs to operate simultaneously, it can also incur significant performance penalties, because the Nested VMM requires access to the root VMM any time a read/write operation is performed (and also for read and write synchronizations between the physical VMCS and the virtual VMCS). Intel findings show that even a well optimized VMM can have up to 12 reads and 6 writes for every VM exit.[4]

The Solution: Intel® VMCS Shadowing

4th generation Intel Core vPro processors include a capability called Intel VMCS Shadowing that greatly reduces the frequency with which the guest VMM must access the root VMM in a nested environment. With Intel VMCS Shadowing, the root VMM is able to define a shadow VMCS in hardware. A guest VMM can access this shadow VMCS directly, without interrupting the root VMM. This capability can significantly reduce the number of VM entrance and exits. And since the shadow VMCS is implemented in hardware, required accesses can be completed nearly as fast as in a non-nested environment (Figure 4).

Regardless of whether nesting is done using the software-only approach or with Intel VMCS Shadowing, there is additional processing overhead created each time the virtual or shadow VMCS is synchronized with the physical VMCS that is used by the root VMM. When the software-only approach is used, the root VMM knows exactly what needs to be synchronized since it was directly involved in every access. It can therefore perform the synchronization efficiently. However, if Intel VMCS Shadowing is used, the root VMM has no idea which of the more than 130 VMCS fields were accessed, since it was not involved in those accesses. The root VMM must therefore synchronize every field that could have possibly been accessed, even though most of the fields are never touched (Figure 5).

Results from Intel Labs profiling across a wide variety of VMMs, shows that approximately 90 percent of VMCS fields are never read and more than 95 percent are never written. As a result, for most VMMs, a full VMCS synchronization can take approximately 15 times longer than necessary.[4]

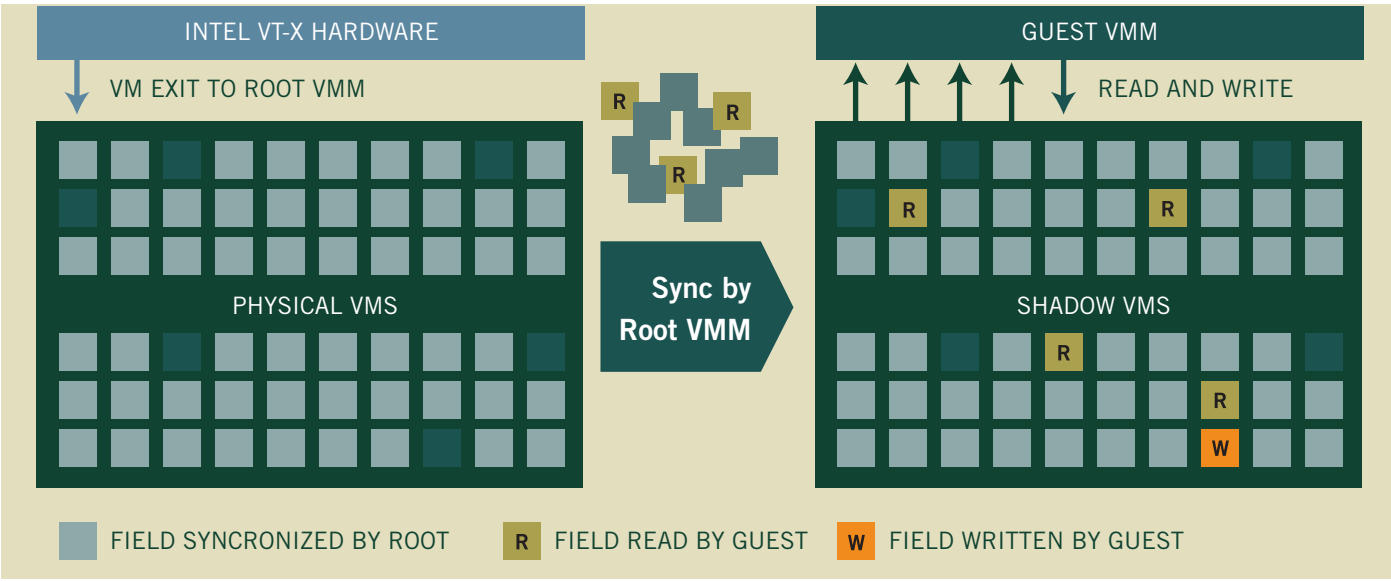

Figure 5. Without specialized hardware support, nearly all 130 plus fields of the physical VMCS and the shadow VMCS must be synchronized, which can result in excessive overhead and slower overall performance of the virtual environment.

To reduce this synchronization overhead, Intel incorporated an additional feature into Intel VMCS Shadowing called VMREAD and VMWRITE bitmaps. These bitmaps allow for selective access to the shadow VMCS. The root VMM can tune the bitmaps so that the 5-10 percent of VMCS fields that are commonly accessed are written directly to the shadow VMCS, while the very rarely accessed fields are synchronized through the slower path that is managed by the root VMM (Figure 6).

By using the VMREAD/VMWRITE bitmaps, the root VMM gets the best of both worlds. Nearly all of the accesses go directly to the fast shadow VMCS and very few extraneous fields need to be synchronized. As a result, Intel VMCS Shadowing enables near-native performance for two VMMs running simultaneously on the same hardware platform.

With the help of Intel VMCS Shadowing, Citrix XenClient and McAfee Deep Defender can be hosted on the same PC, while greatly reducing any potential impact on the user experience. Other multi-VMM usage models can also be accommodated (Figure 7).

Conclusion

As the value and popularity of client virtualization continues to increase, new usage models are emerging that can require two VMMs to run simultaneously on the same physical system. A prime example is McAfee Deep Defender, which runs on top of its own micro-hypervisor to provide advanced, kernel-based monitoring that helps to protect against sophisticated new threats, such as kernel-mode rootkits, stealth attacks, and zero-day malware.

PCs powered by 4th generation Intel Core vPro processors include Intel VMCS Shadowing, which provides hardware assistance for this and other multi-VMM usage models. Both Citrix and McAfee are evaluating this capability for inclusion in future product release.

With this support, customers would have the ability to use Citrix XenClient and McAfee Deep Defender together to enable a more secure client environment, while maintaining fast, responsive client performance. They will also have the flexibility to support additional multi-VMM usage models as they emerge.

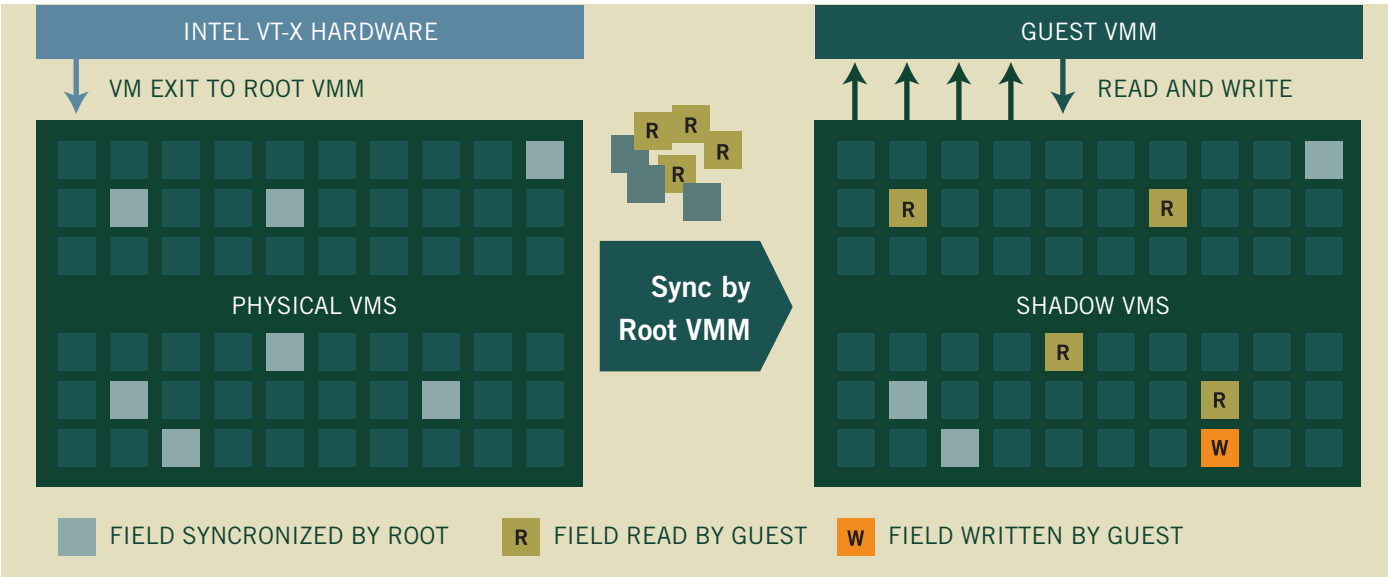

Figure 6. VMREAD/VMWRITE Bitmaps greatly reduce synchronization overhead, while accelerating overall performance by allowing the guest VMM to directly access the shadow VMCS for the majority of fields.

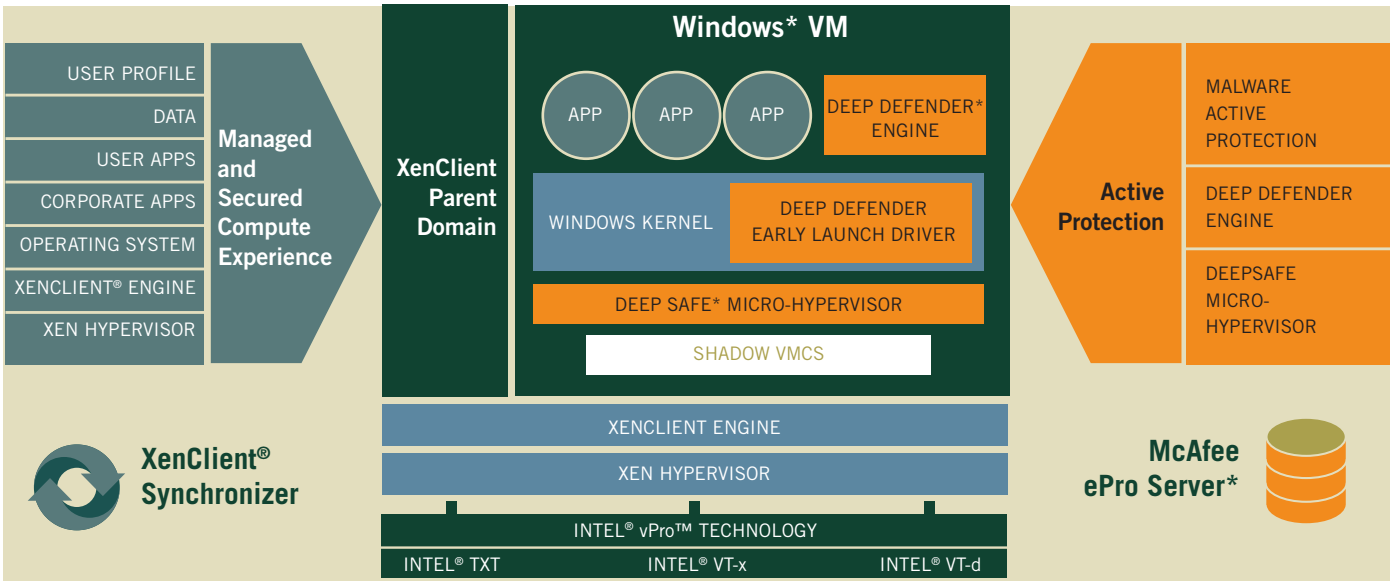

Figure 7. Working together, Citrix XenClient®, McAfee Deep Defender*, and 4th generation Intel® vPro™ Technology (with Intel® VMCS Shadowing) provide strong, deep security protection—without significantly impairing client performance.

About the Authors:

Ahmed Sallam is VP of Product Strategy/ CTO of Client Virtualization at Citrix Systems where he drives technology innovation and product strategy for emerging client virtualization, management, and security solutions. He works closely with software and hardware ecosystem partners as they adopt and integrate new capabilities into Citrix's open, extensible virtualization platforms. Prior to Citrix, Ahmed was CTO of Advanced Technology and Chief Architect at McAfee, which is now part of Intel Corporation. Ahmed was a co-inventor and architect of Intel/ McAfee DeepSAFE and a co-designer of VMware's VMM CPU security technology known as VMsafe. Prior to McAfee, Ahmed was a Senior Architect with Nokia's security division and a Principal Engineer at Symantec. Ahmed is a renowned expert across the industry for pioneering new models in computer system security that help to provide proactive, preventive, predictive and highly-assured safe computing environments to computer networks and devices. Ahmed holds 18 issued patents and has more than 40 published and pending patent applications. He earned a bachelor's degree in Computer Science and Automatic Control from the University of Alexandria.

Greg Boitano is a Senior Product Marketing Engineer with Intel Corporation where he's been part of Intel's Business Client Platform Division for the past several years. Greg has an extensive background working with end customers, ISVs and IT solution providers. Greg is currently responsible for Intel Business Client desktop virtualization marketing. Prior to joining Intel, Greg spent several years in the Telecom industry working with enterprise customers both as a sales manager and an engineering manager. He earned a Bachelor's Degree in Marketing from Western Washington University, and holds an MBA from the University of Oregon.

Ron Talwalkar is a Senior Director of Product Management DeepSAFE at McAfee Labs. In this role, he owns the platform roadmap for DeepSAFE, providing hardware-enhanced capabilities supporting enterprise and consumer security products. Ron has been at McAfee for over eight years working in a variety of roles, from Development Manager to Director of Engineering, and now Product Management. Prior to McAfee, he worked at Intel for 12 years as a Senior Engineering Manager across many organizations, the last of which was the Intel Software Solution Group. Ron holds both bachelor's and master's degrees in Computer Science, as well as a master's degree in engineering management from the Oregon Graduate Institute.

Glossary

Intel® Virtualization Technology (Intel® VT): Comprehensive hardware-based platform capabilities designed todelivernative performance, with best-in-class reliability, security, and scalability. Hardware-based Intel® Virtualization Technology (Intel® VT)1 improves the fundamental flexibility and robustness of traditional software-based virtualization solutions by accelerating key functions of the virtualized platform.

Intel® VT-x: The hardware based implementation of Intel VT in Intel processors, which helps accelerate performance for virtualized workloads along with enabling improved security and reliability.

Intel® Virtualization Technology for Directed I/O (Intel® VT-d): The implementation of Intel VT in Intel chipsets, which helps to improve I/O performance, availability, reliability and security through device isolation and direct assignment of devices.

Intel® Trusted Execution Technology (Intel® TXT): A feature of Intel processors that helps to improve platform security by measuring the boot environment to establish a root of trust and ensure that a computing system can only boot into a known good state.

Intel® VMCS Shadowing: A hardware capability available on PCs powered by 4th generation Intel® Core™ vPro™ processors and designed to improve performance for nested VMMs.

VMM (Virtual Machine Monitor): Computer software that can create and manage multiple virtual machines on a single physical system, each capable of running its own operating system and applications.

Nesting: A software-based strategy for running multiple VMMs on a single physical system, with one VMM acting as the root VMM and others acting as guests.

Root VMM: In a nesting scenario, the VMM that controls access to hardware resources for one or more guest VMMs.

Nested VMM: A guest VMM that depends on a root VMM to provide managed access to platform hardware resources.

VMCS (Virtual Machine Control Structure): A high-speed data structure designed to hold guest and host CPU register states. The VMCS is enhanced and optimized with each successive implementation of Intel VT and is key to reducing virtualization latencies.

[1] Intel® vPro™ Technology is sophisticated and requires setup and activation. Availability of features and results will depend upon the setup and configuration of your hardware, software and IT environment. To learn more visit: http://www.intel.com/technology/vpro/.

[2] No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer with Intel® Virtualization Technology, an Intel TXT-enabled processor, chipset, BIOS, Authenticated Code Modules, and an Intel TXT-compatible measured launched environment (MLE). Intel TXT also requires the system to contain a TPM v1.s. For more information, visit www.intel.com/content/www/us/en/data-security/security-overview-general-technology.html.

[3] Intel® Virtualization Technology (Intel® VT) requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance, or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

[4] Source: Intel Labs.

0513/GPB/HBD/PDF Please Recycle 329012-001US