CHAPTER 8

The Future of Trusted Computing

The best thing about the future is that it comes one day at a time.

—Abraham Lincoln

This chapter reviews the critical capabilities of a trusted platform and reinforces the benefits of solutions based on Intel® Trusted Execution Technology (Intel® TXT). It discusses the key considerations for implementations and the recommended uses for customers seeking to get started. This chapter also provides a vision of what is to come—explaining why building the infrastructure now will make it easier and quicker to adopt next-generation, trust-based solutions. These new solutions can enhance IT architectures to meet evolving business needs by providing even greater control and visibility into their increasingly virtualized workloads and providing solutions to enable the transition from traditional static datacenter models to embrace new dynamic cloud infrastructures.

Trust Is a Foundation

Increasingly, security is not an option. It is a critical consideration for nearly every business decision: it informs or influences the products that IT buys, the architectures they employ, and the services they subscribe to. With the constantly changing IT architectures, regulatory environment, and potential threats, customers need more tools to provide security in virtualized and cloud environments. The traditional IT security toolbox is just not up to the task to handle all of these changes—with the following being some of the biggest challenges:

- Tools cannot provide coverage to protect assets.
- Tools are no longer architecturally efficient in how they provide protection.

It is already painfully clear that the gap between security capabilities and solutions is a drag on customer adoption of new architectures and use models for the cloud. One can read about it in the trade press on essentially a weekly basis or in discussions with customers or peers, but it is readily evident in nearly every survey that security is a challenge in a global sense.

As shown in Figure 8-1, customers need more protection and controls to make the cloud a viable model for all kinds of workloads. They need protection against emerging threats such as rootkits. Historically, many have viewed these threats as “someone else’s problem” or one that is purely hypothetical. Neither is really true. These classes of stealthy, low-level threats are real and occurring “in the wild.” The recent example of the “Mebromi” BIOS rootkit (Giuliani 2011) was an eye-opener for many. This attack was specifically engineered to target system BIOS code developed by Award for a number of Chinese computer systems, and capable of detecting the presence of several common local antivirus software packages in order to thwart them. Similarly, the discussions driven by Invisible Things Lab with their “Blue Pill” Hypervisor rootkit concept (Rutkowska/Tereshkin 2006) dramatically raised the visibility into security concerns with hypervisor software models. Most were unaware that such esoteric platform components could be attacked and that an attack could be executed in a commercial environment. Unfortunately, as is often the case, it takes a commercial exploit to change the mindset and drive people to take action. And there

are many more IT managers and security professionals that are indeed taking action. In 2012, a growing number of entities, such as the US National Institute of Standards and Technologies (NIST), are designing guidelines for protecting systems, which include recommendations for securing these very basic but highly privileged platform components.

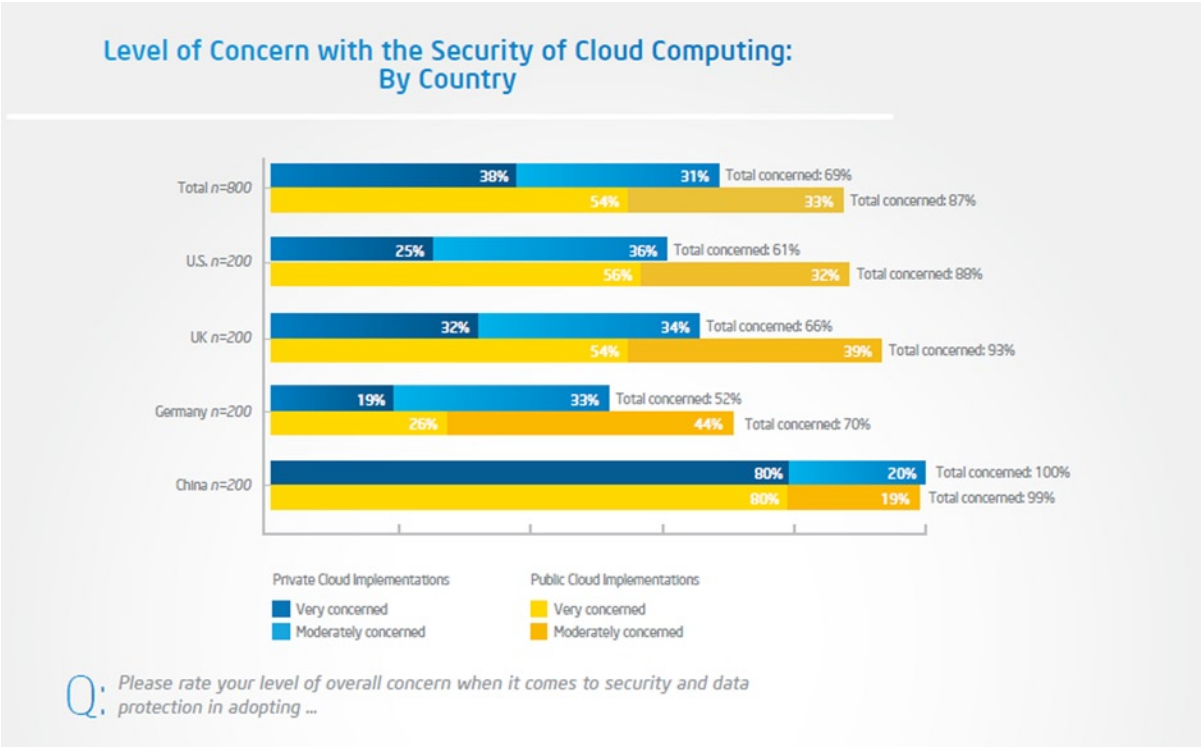

Figure 8-1. *Addressing the need for security in private and public clouds*[1]

More Protections and Assurance

Even as these recommendations from security specialists take form, there is a growing recognition even among the nonsecurity specialists that BIOS is not the only worry. There is the interesting phenomenon that customers and IT managers who implicitly trust the security of their hypervisor in a basic virtualized system grow much more concerned about the security of that same hypervisor when it is deployed in a cloud environment—particularly in a public cloud infrastructure managed by a third party.

This change in perspective is clearly indicative of a customer base that needs new capabilities and controls—more assurance of the integrity of their environment as they virtualize away traditional control and security structures. They need compensating control capabilities to provide assurances for their own concerns, or to assuage the concerns over risk and provide the proofs of protection and control required by others: their information security management professionals or their auditors.

Given the crucial role played by the hypervisor—after all, this essential software is responsible for managing the underlying hardware, allocating resources such as processor, disk, memory, and I/O to the guest virtual machines and arbitrating accesses and privileges among guests—one would want to have the highest levels of assurance that it is

[1]McCann, "*What's Holding the Cloud Back?*" Cloud Security Global IT Survey, sponsored by Intel, May 2012.

indeed uncompromised. This is the rationale for the kind of survey results that appear in Figure 8-2. With this growing awareness and concern has come a corresponding growth in vendors looking to define solutions.

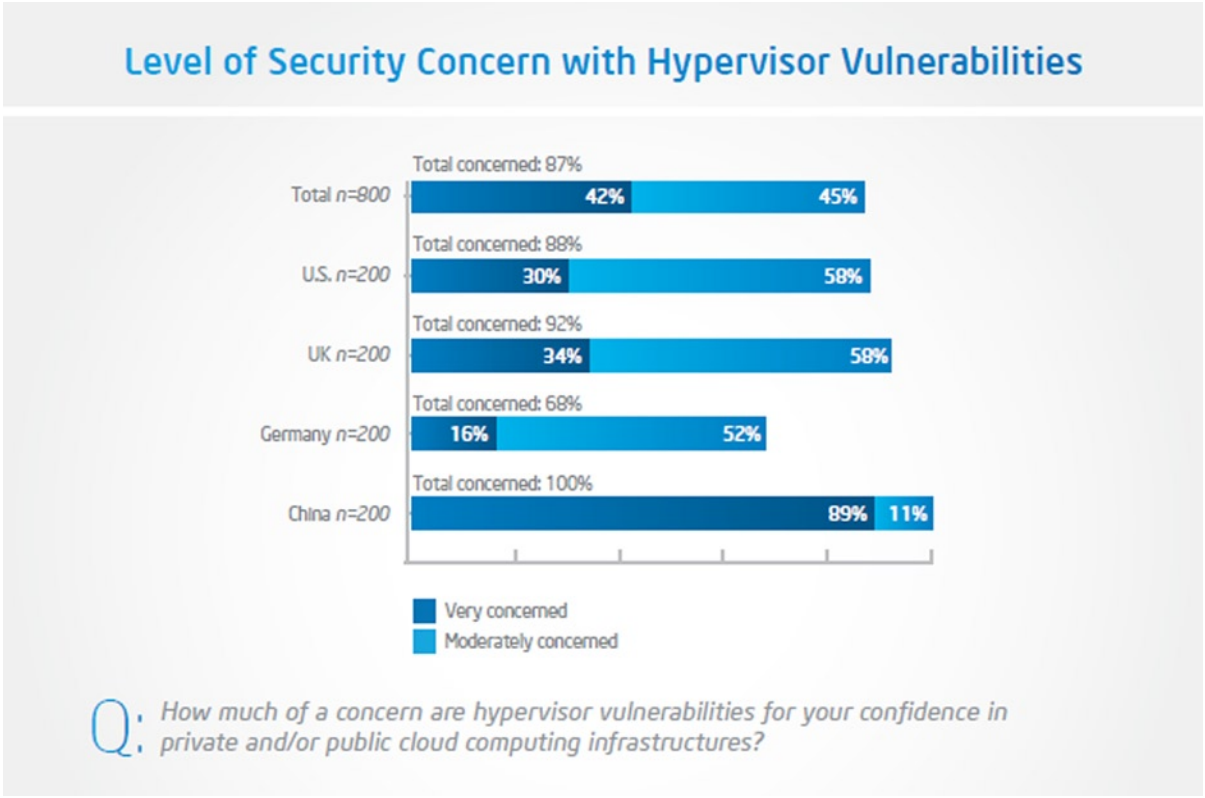

Figure 8-2. *Survey results show concerns over hypervisor integrity across regions*[2]

Along this same line of reasoning, a number of the leading hypervisor platforms have embraced Intel Trusted Execution Technology as a way to allow the hardware and software platforms to work together to provide assurances of integrity. Using cryptographic measurement techniques, applying local whitelist-based policy mechanisms, and creating a tamper-resistant environment, Intel TXT can work with the key platform components to enable a verification mechanism that can help detect alterations in critical BIOS and hypervisor components.

By enforcing this verification on the platform and storing the results in a hardware-based device—the Trusted Platform Module—Intel TXT provides a method for enforcing control for the platform. It also provides visibility into the platform that delivers the assurance benefits needed to meet growing audit and reporting requirements. These requirements are inherent in many corporate IT security policies, as well as rapidly proliferating government and industry security regulations.

Now that Intel TXT is an available, deployable capability on millions of Intel® Xeon® E3, E5, and E7 processor family-based servers from virtually all of the leading OEM[3] and channel providers, a growing ecosystem of software support is emerging to make these visibility and control use models possible. The ecosystem is forming to enable solutions and services to provide these protections and controls. Most of the leading BIOS and hypervisors today

[2]McCann, "*What's Holding the Cloud Back?*" Cloud Security Global IT Survey, sponsored by Intel, May 2012.
[3]For a full list of systems that support Intel TXT, please see www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-server-platforms-matrix.html.

are developed in a manner that will allow them to be measured as an integrity check. Virtualization and cloud management platforms can query the host platforms to verify trust status and differentiate trusted, high-integrity platforms from untrusted platforms—the foundations for trusted pools. Security policy engines are being made "trust aware," such that they can work with virtualization and cloud management platforms to control access and workloads based on the platform's trust status. Other key security management tools—such as governance, risk, and compliance (GRC), and security information and event management (SIEM) systems—can report on platform trust and integrity status and events as part of their monitoring and compliance activities. These software stacks enable platform trust to deliver bigger business value beyond basic anti-malware technology.

With these critical links of trust-based use models, platform-level trust can become a full-fledged aspect of an organizations' IT security management portfolio for traditional and virtualized/cloud models. Taken a step further, and given the changing IT architectural models and threat vectors, one could argue that trust is increasingly likely to be a foundational component to these new security models. Intel TXT provides one of the broadest and more widely adoptable mechanisms to enforce trust on platforms and into the enterprise. Learning how to deploy and use this for your organization today will help provide an advantage to reduce near-term risks and meet tactical compliance challenges. It will also provide a solid basis for leveraging related new trust technologies that expand the benefits and provide more protection, control, and visibility for IT security.

Is There Enough to Trust?

There are no silver bullets for security. The threats are too broad, the adversaries too varied, sophisticated and well-resourced to allow a "one size fits all" or single point solution that can stop all threats. Security is a story of multiple lines of defense (or "defense in depth") and of evolution. As threats evolve, so too must the defenses. Trust is no different: it was created to mitigate threats and meet the needs that have been outlined in this book, and it will need to evolve to provide more protection and value over time. The following section discusses how that evolution may play out. Some of this section is speculative and based on early lessons learned from the process of bringing Intel TXT to market with key customers and hardware and software ecosystem partners. While the final destinations and timing might be unclear, and priority or emphasis may influence one evolutionary path over another, the areas of interest seem to be rather universal and worthy of discussion here.

The trust that is available today is innovative, but has limitations that merit discussion. For instance, Intel TXT

- Only measures at launch time.
- Only measures key system BIOS, firmware, and hypervisor (or operating system in a nonvirtualized use) components.
- Works on a whitelist model.

Some would like to argue the benefits of such a limited approach to trust. But the reality is that these limitations—while real—do little to mitigate the value of Intel TXT, as the use models outlined previously should convince the reader. Moreover, these limitations may be overcome by new use and deployment models, complementary security capabilities, and advances as the technology matures. Let's first address the limitations themselves and discuss how material these limitations are and how these limitations can be reduced.

Measures at Launch Time.

First, there is the "launch-time only" aspect of Intel TXT. While Intel TXT actually does provide some protections that keep secrets in memory safe after a trusted launch, the active measurement component of Intel TXT is only invoked at very limited times—at platform launch or restart or when resuming from a sleep mode. Historically and ideally, these would be very infrequent events: customers would often prefer to set up a server, install their software and workloads, and then never have to power it up or down ever again. Alas, this ideal is seldom the reality because customers have software or BIOS updates or facilities changes that required system restarts. But perhaps more interestingly, customers are looking at new, dynamic, virtualized datacenter models that may lead to *increased* frequency of systems being restarted or powered down. With these highly virtualized datacenter models, customers expect to manage

systems to maximize power savings, powering systems down in off-peak times—and in such models, customers will get an incremental security benefit in addition to power savings as systems can be verified for integrity upon these restart and resume events.

One last consideration that mitigates this perceived limitation is the role of launch-time integrity enforcement in the overall security portfolio. As noted, no single technology solves all security problems, and as such Intel TXT's primary role is that it provides a solid checkpoint that *complements* the (ideally) many runtime protections such as antivirus and intrusion detection systems. If malware such as a rootkit evades these runtime protections, Intel TXT provides a mechanism that can allow for detection at the next restart/resume event. Otherwise, if they can escape detection by runtime protections, these malicious agents can install themselves into locations that can avoid subsequent detection by these traditional tools—which obviously creates troublesome, long-term vulnerability challenges.

What Intel TXT Measures

The second limitation to assess is what Intel TXT measures. Some would say that measuring BIOS and key firmware, hypervisor, or operating system components does not assure complete platform security. Clearly, no one protection does, and as discussed previously, the proper consideration for Intel TXT is as a complement to runtime security tools. Another aspect of the complementary nature is that Intel TXT will be evaluating aspects and components of a system that are generally weakly protected, if at all, with traditional security tools.

It would be easier to assign far more credence to this concern if the threat environment were not showing very real challenges to these components. Adversaries will find a way to exploit these components into more significant attacks if they are left unprotected. Having hardware-assisted integrity assurances provides an additional strong complementary protection for the IT environment. In the highly consolidated and virtualized IT environment, where a single server no longer hosts a single workload and a compromised host can jeopardize multiple workloads, such protections become even more important.

The Whitelist Approach

The final limitation to address is a potential concern over a whitelist approach vs. the well-established blacklisting approach of so many traditional security tools. Again, one can turn here to the consideration of Intel TXT as a complement to existing tools—and in this case, the alternative approach provides a useful contrast, whereby the different approach eliminates the prospects of an "all security eggs in one basket" approach.

Whitelisting is often challenged by the perception of inflexibility, with good reason. The basic principle of a whitelist security model is that it specifies "known good" elements that one wants to allow. Alternately, a blacklist model is based on stopping all "known bad" elements from executing. Each has its challenges in scale and manageability. But one could argue that in some situations one model is advantageous. The case here is that in the tightly defined boundaries of BIOS, firmware, and hypervisor/operating systems, it is relatively easier for an IT manager to exert tight controls and identify a very finite set of these components that they wish to allow rather than having to identify the essentially infinite set of threats they would like to stop.

Whether one considers these aspects of what Intel TXT provides as limited, it is clear that these capabilities provide additional value to the modern IT security toolbox.

The Evolution of Trust

The previous paragraphs put the capabilities into a suitable context, even as prior chapters outlined how to implement the capabilities into an IT infrastructure, and what controls to expect. In doing so, one hopes that the reader has gained a stronger foundation—a foundation of platform trust—to underpin their datacenter and cloud security model. While this itself is helpful and adds value, there is value in starting today and raising the security bar and being better positioned for further advances in technology. In the future, the basic trust capabilities of the platform can evolve to provide more benefits and greater control through unified hardware, software, and services.

Let's consider where the foundational, launch-time-oriented trust capabilities such as Intel TXT may evolve in the future—either through Intel technology innovation or through evolutionary technology from elsewhere in the software, hardware, and services ecosystem, as shown in Figure 8-3. Three of the more interesting and in-demand areas for innovation are related to delivering the following:

- Trusted guests
- End-to-end trust
- Runtime trust

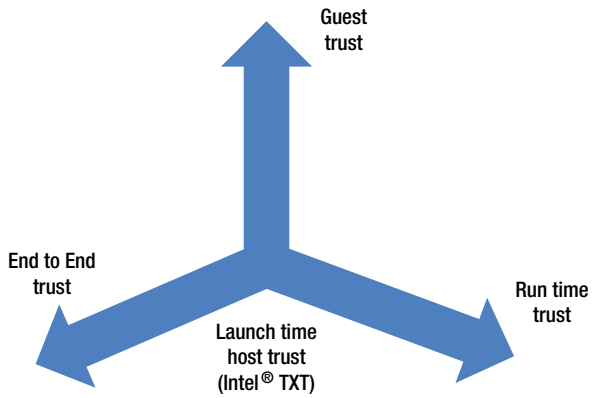

Source: Intel Corporation

Figure 8-3. *The evolution of trust technologies*

Trusted Guest

The first concept is relatively easy to derive from the concept of a trusted host. The premise here is simple: extend the chain of trust established by the hardware at launch, extended through the measured and verified BIOS and hypervisor, and use that trusted base to measure and verify guests as "known good" images. While this cannot be done using Intel TXT hardware solutions today, and there is no significant hardware assistance for such use models today, it is an area of investigation by Intel and the industry in general. But it is not all bad news, for while not rooted in a solid chain of trust, there are currently some promising software-only models that can help verify guest images using whitelist-oriented data management techniques that could be considered quite complementary—especially when run on top of trusted host systems.

End-to-End Trust

The second concept is also easy to grasp if one understands the premise of trusted compute pools already discussed at length in this book. Just as one understands the value in using platform trust attributes to control and manage workloads inside a datacenter or in a cloud environment, one should be able to extend that vision and see a future of trusted clients communicating with trusted clouds. In this manner, platform trust and integrity attributes can benefit both sides of the cloud to client continuum, allowing IT and security managers to establish data and workload access controls based on higher assurances that the client and host have not been compromised. This makes for a very compelling complement to traditional user/role-oriented access control tools. With such a model, security can be based on the users' rights to access data or services, the access device's ability to protect the data or services, and the ability to stop user access to a service or data in the cloud that may have been compromised.

The technologies for such use models are near at hand, and as of mid-2012 proof-of-concept demonstrations of this type of policy enforcement are emerging in the software and service provider ecosystem—including some based on Intel TXT server technologies and McAfee DeepDefender client integrity software on client platforms. Expect a rapid expansion and increased availability and granularity for the types of cloud-to-client trust capabilities that can help enforce data management policies in the coming years.

Runtime Trust

Lastly, trust will also evolve to create better support for runtime-based trust assurances. This chapter has already dedicated space to discussing the need to extend trust and integrity into the runtime space, but has also acknowledged that there is a strong set of products in anti-malware. Intel will continue to evaluate and explore opportunities to use platform hardware enhancements to make trust and integrity assurance more complete, more efficient, and more resistant to attack. As such, one would expect that this evolution will continue over time, with software leading hardware through the development of new use models addressing specific needs and mandates. With the current and expected strong focus for "continuous monitoring" and "continuous compliance" capabilities from entities such as the US government program FedRAMP, the Open Data Center Alliance, and the Cloud Security Alliance, this is a strong area of focus and innovation throughout the industry and across the globe.

The Trust and Integrity "Stack"

Intel has long believed in the need to enhance IT security, and has invested in features such as Intel TXT, Execute Disable (XD) bit, and cryptographic instructions to improve platform and data security. And we've worked with the key security software ecosystem to help assure that these features are enabled and usable to deliver security value to end-user customers. While the collaboration with this broad ecosystem is essential and will continue, in 2010, Intel took a major step and acquired a market leader, McAfee, to move this quest forward. Establishing McAfee as an independent subsidiary allows faster and increasingly deep collaborations between hardware and software capabilities to deliver customers hardened and more efficient, effective security solutions.

Intel and McAfee have established a shared vision for cloud security. Trust and integrity are a critical component of this vision. In this vision, the ability to have a verifiable and auditable high-integrity compute environment can become a reality, with reduced risk from threats, data loss, and downtime, and greater ability to meet compliance requirements.

Increasing the number of enforcement points at multiple layers of the stack will bring higher integrity assurance, especially with additional hardware-enhanced security and software solutions to these areas over time—as discussed in the preceding section. Consider the model from the vision shown in Figure 8-4.

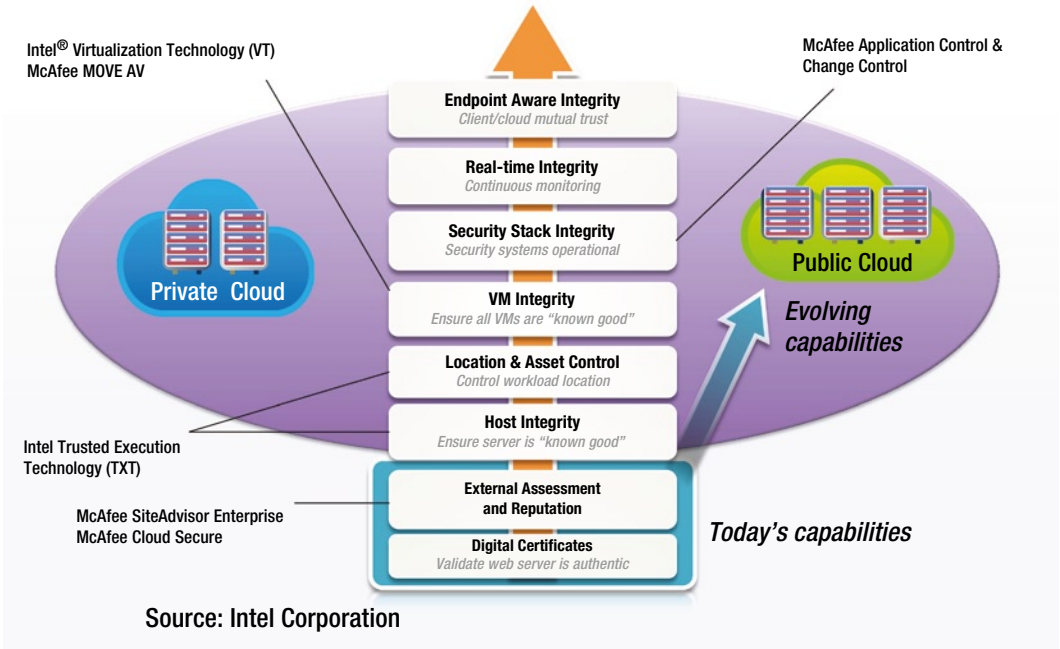

Figure 8-4. *Expanding integrity-based security enforcement points across cloud infrastructures provides more control for IT*

The challenge today is simple: you simply do not know much about any device you may want to connect to—on the Internet or in a cloud. Let's consider an example.

The bottom two items from our diagram are somewhat common today—digital certificates (you see examples when you connect to an online commerce site, for example) or external reputations—if you have a tool like McAfee SiteAdvisor or the like—something that can tell you if a site has been examined for malicious activity. McAfee SiteAdvisor and similar products from other third parties allow you to surf and search the web more safely, avoiding online threats such as spyware, adware, and phishing scams. With the McAfee Cloud Secure program, McAfee enables rigorous security testing, business practice review, compliance certification, and ongoing vulnerability evaluation. Cloud providers and software-as-a-service vendors can demonstrate credible, third-party validated site integrity. *Today, this is essentially all the information you can have when making a connectivity decision*. Unfortunately, that is close to flying blind.

If you want to extend your business to the cloud and make more informed data access and control decisions, you need much more data and control points. Intel and McAfee see a progression of new information that can be provided to allow *much* richer assessments of the security posture of the resources in the cloud. Some of these are being broadly adopted now.

The next step is the host integrity discussion that has been the focus of this book. Using Intel TXT, users are able to verify that the servers they are using have demonstrated integrity. Considering the asset and location control aspect, Intel TXT will also provide a mechanism that will allow a host to store a label in hardware that IT managers or administrators can use to designate the location or other relevant characteristics (customer class, service tier, and so on) of the host server. These attributes can be reported in the same manner as the platform integrity attributes in order to enforce connection or resource allocation policies. One can see how customers can develop use models that are extensions of trusted pools, whereby workloads can be constrained by trust as well as location or asset tag attributes. There are already a number of business and regulatory environments where such extended "boundary controls" based on such physical or logical attributes would be powerful.

Beyond that and further up the stack, McAfee server technologies such as application control and change control work together to reduce overhead on servers and virtual machines while proactively mitigating the risk of data breaches, targeted attacks, and unplanned downtime. These solutions provide proactive security monitoring at the operating system, application, and file level. Additionally, technologies such as the McAfee MOVE Antivirus architecture enhance runtime anti-malware and integrity with an approach that removes the bottlenecks and inefficiencies inherent in the use of scan-based models in virtualized environments. And the previous section discussed how innovative integrators and solution providers can take the Intel and McAfee building blocks to provide end-to-end integrity-based security solutions with a mechanism for mutual verification of trust and integrity between client and cloud to allow for bidirectional control of access to cloud resources.

Although we have discussed the Intel and McAfee joint vision and product synergies in the context of business solutions, rest assured that this is not an Intel/McAfee power play. The preceding text is merely intended to illustrate by concrete example how these two companies see the needs of the market and how we can act together to help address these needs. But the solutions are broader than this—and need to be. Both companies already have leadership positions in datacenter and security markets. Both companies also already have robust ecosystem of third-party partners that provide scalable and interoperable solutions. Both companies are also committed to continue to work with other market and technology leaders and innovators to enhance the robustness, scalability, and completeness of security solutions.

These capabilities promise to deliver greater value to IT and security management professionals as they work to adopt cloud architectures and services. According to data from a McCann Cloud IT security survey sponsored by Intel (Intel 2012), in May 2012:

- Seventy-six percent of IT pros are *very interested* in the ability to measure service providers' security posture in real time
- *Setting and enforcing security policies across clouds* would enhance the confidence of 50 percent of IT pros in adopting public clouds

Starting with a baseline of capabilities enabled in Intel TXT and the use models outlined in this book, and driven by the needs of the market for more and more robust tools for dealing with emerging threats and compliance mandates, Intel and its ecosystem of software and service providers will deliver on-going advancements to hardware and software security for greater control and auditability of cloud and datacenter environments. With these advancements, trust will become a full-fledged complement to today's traditional perimeter and integrity/reputation services for providing security—adding depth and granularity to the controls available to IT and security managers.