JOURNEY TO CLOUD

Volume 1, Issue 2

Bringing Your Cloud Goals into Reach

BUILDING A PRIVATE CLOUD

IS CLOUD COMPUTING THERE YET?

SECURING CLOUD COMPUTING

CLIENT VIRTUALIZATION IN A CLOUD ENVIRONMENT

CLOUD TRANSFORMATION FRAMEWORKS

BUILDING A CLOUD SECURITY PLATFORM

AND MUCH MORE

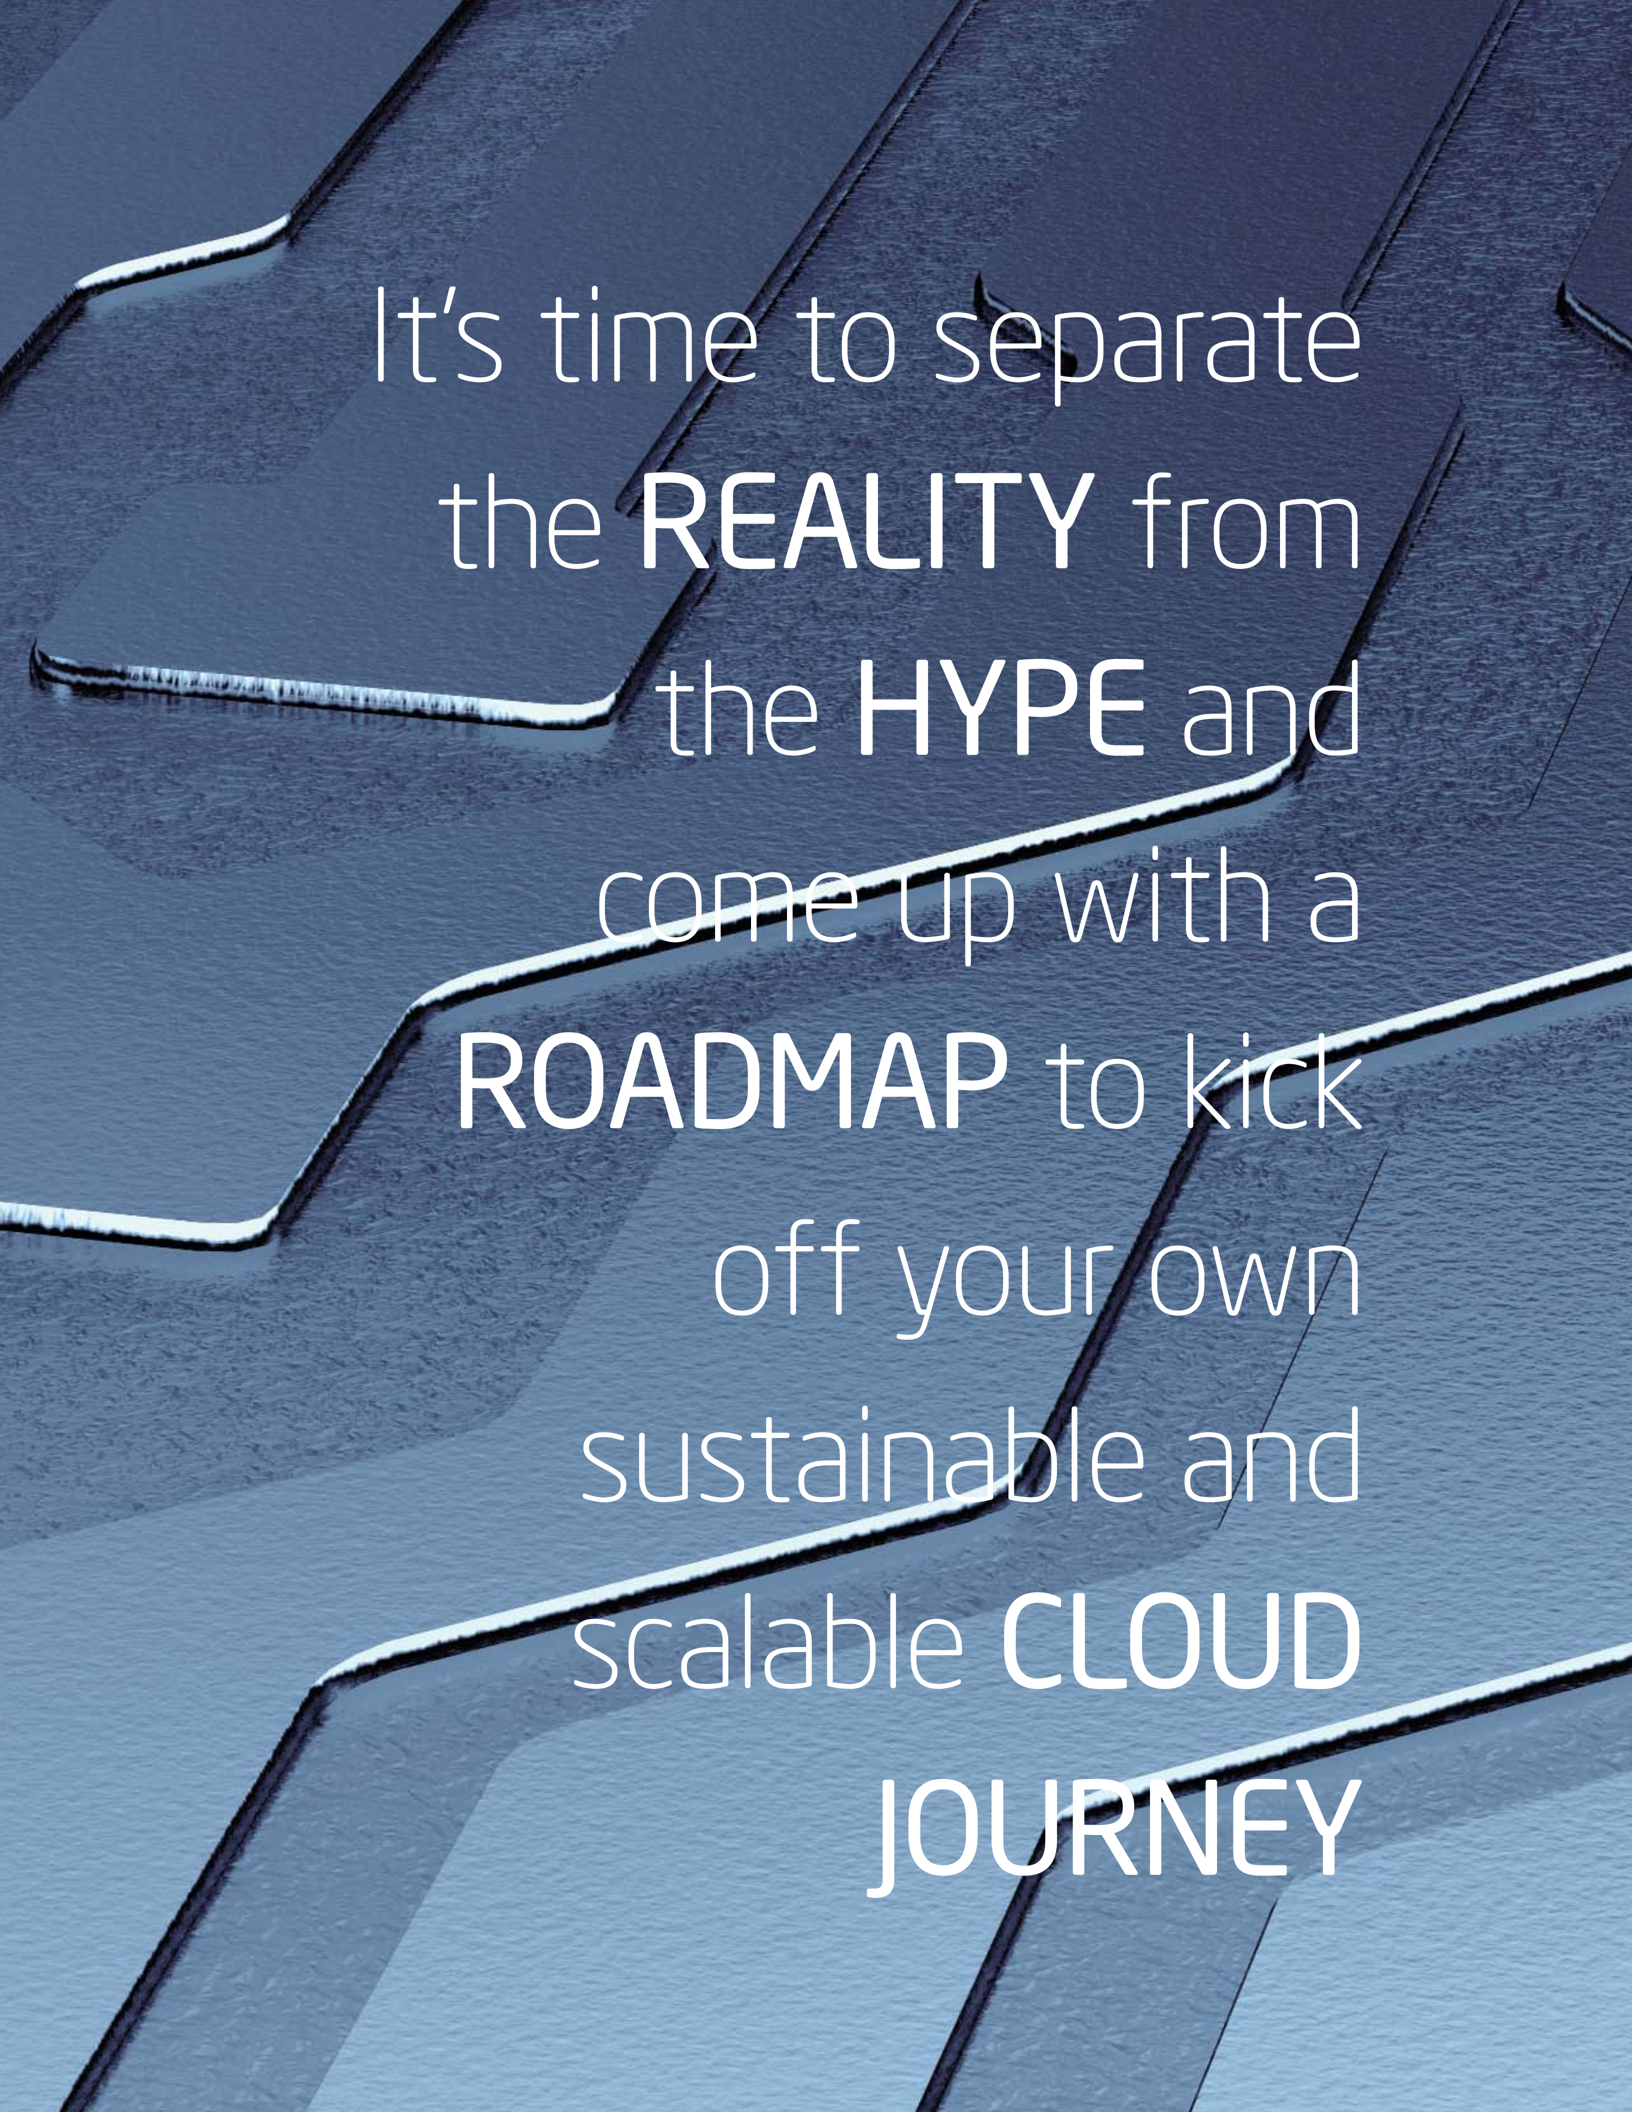

It's time to separate the **REALITY** from the **HYPE** and come up with a **ROADMAP** to kick off your own sustainable and scalable **CLOUD JOURNEY**

CONTENTS

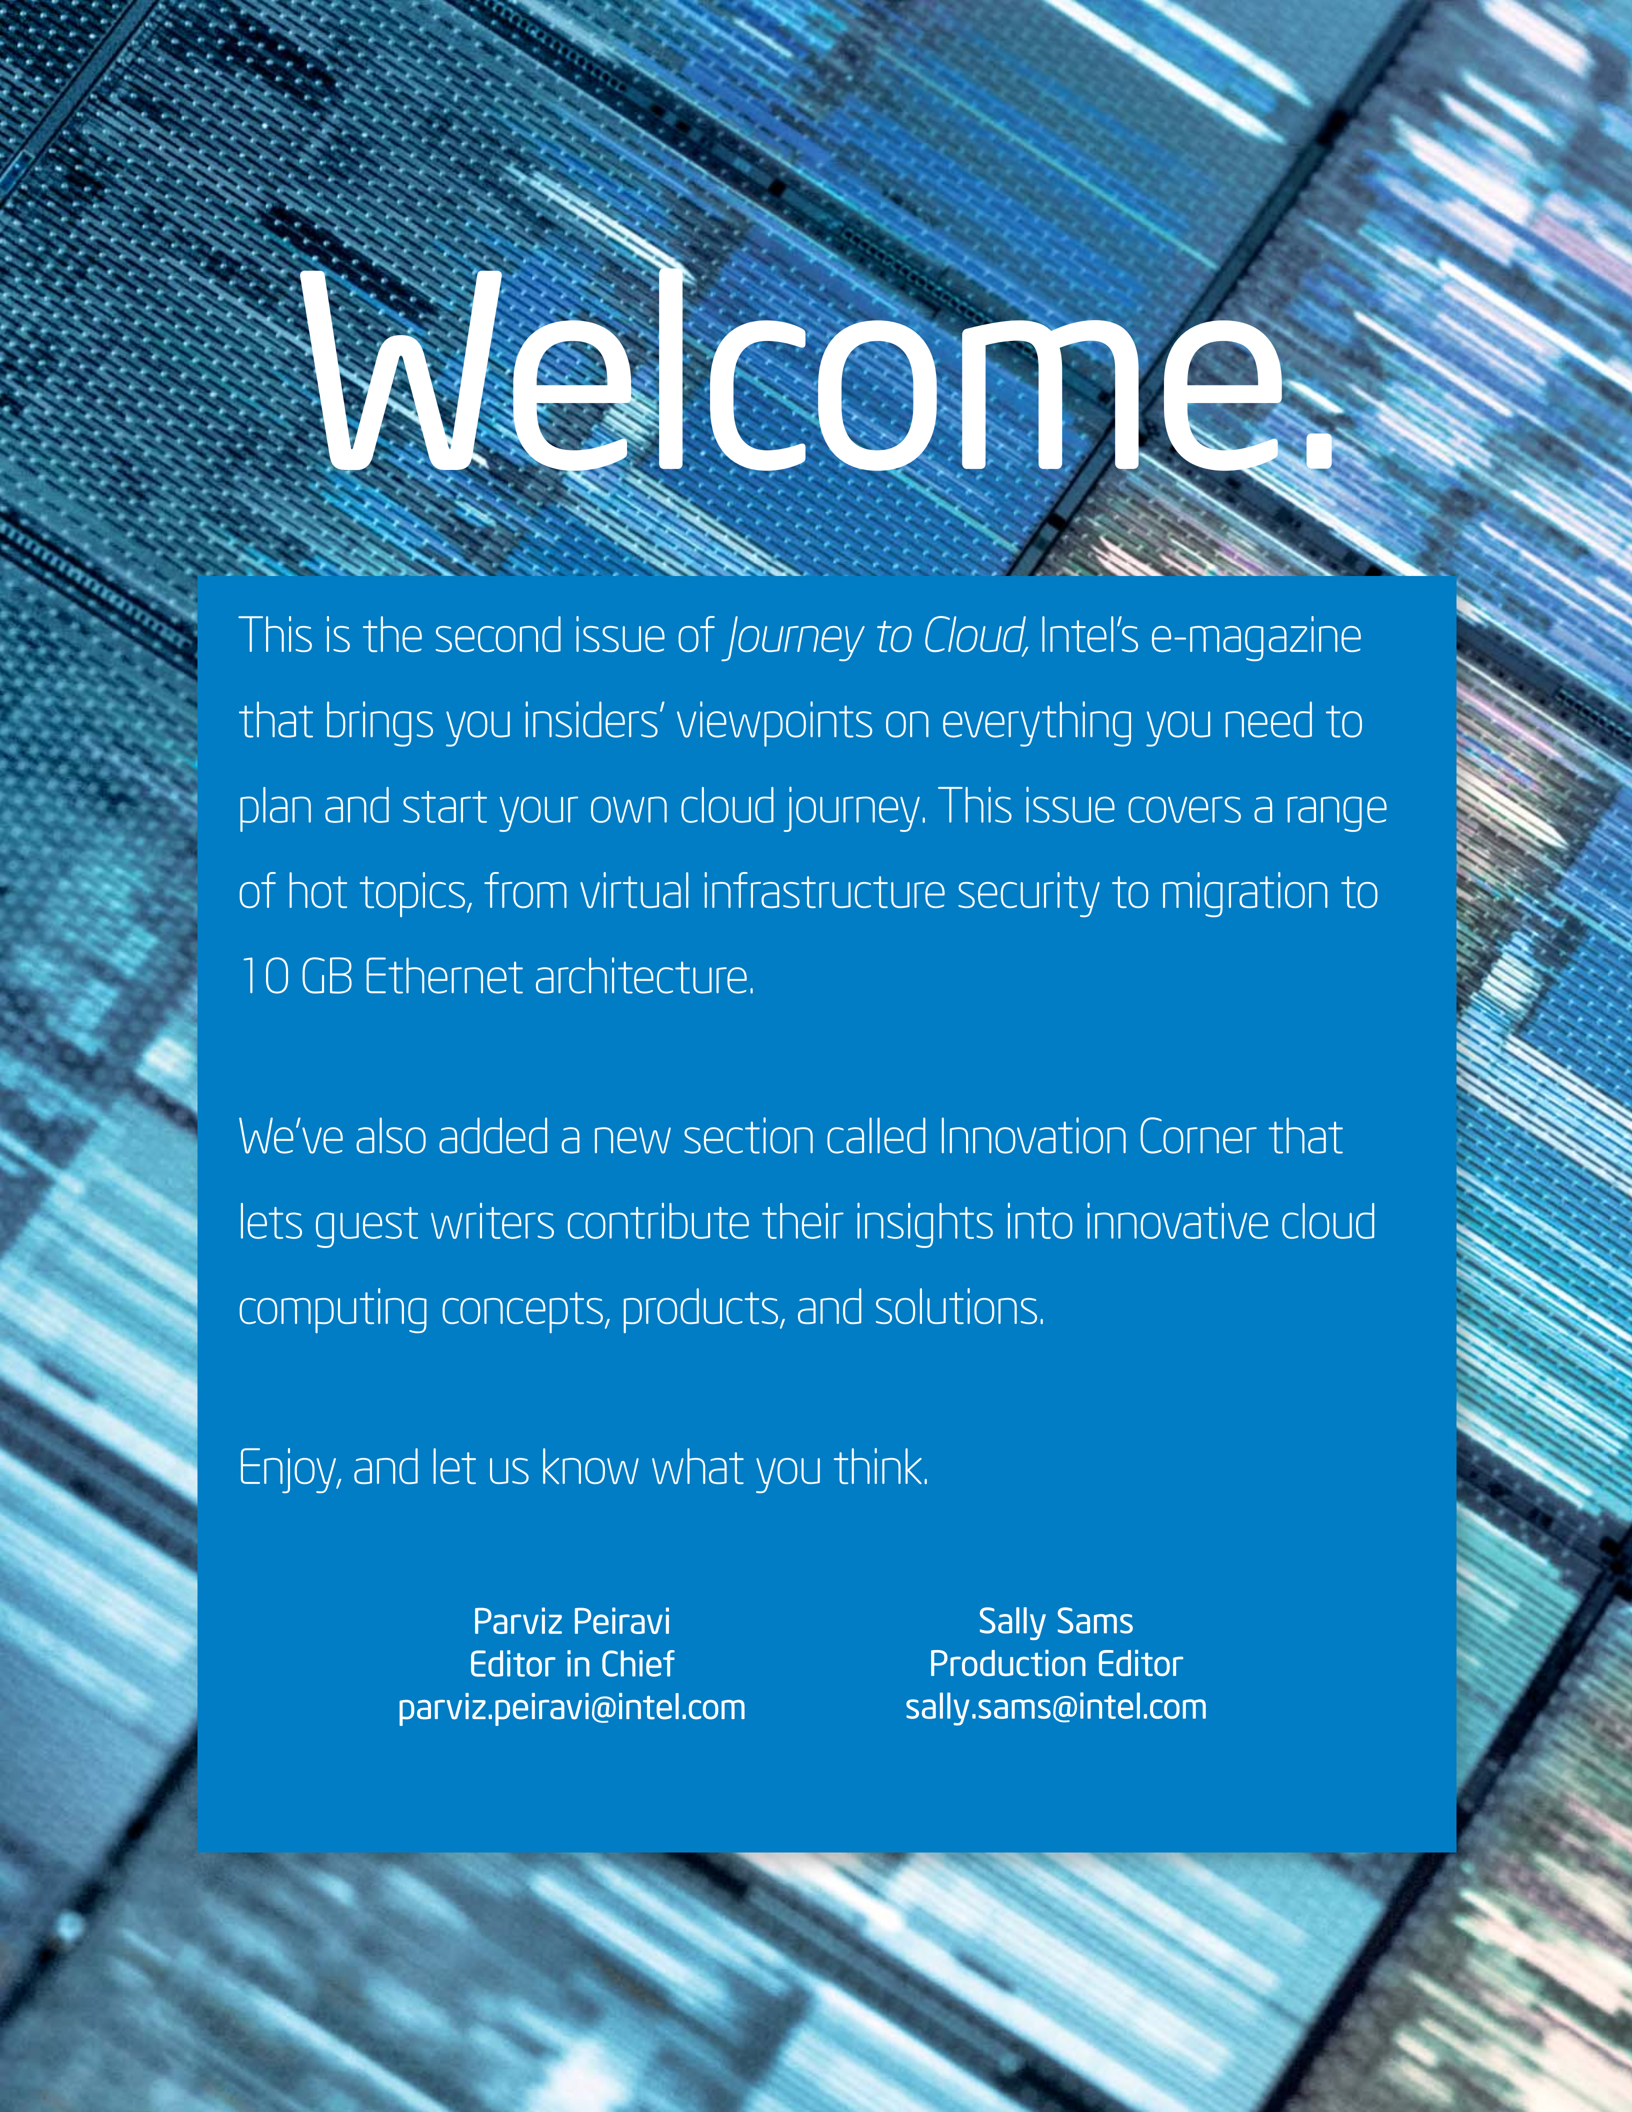

Welcome.

This is the second issue of *Journey to Cloud,* Intel's e-magazine that brings you insiders' viewpoints on everything you need to plan and start your own cloud journey. This issue covers a range of hot topics, from virtual infrastructure security to migration to 10 GB Ethernet architecture.

We've also added a new section called Innovation Corner that lets guest writers contribute their insights into innovative cloud computing concepts, products, and solutions.

Enjoy, and let us know what you think.

Parviz Peiravi
Editor in Chief
parviz.peiravi@intel.com

Sally Sams
Production Editor
sally.sams@intel.com

BUILDING A PRIVATE CLOUD

How Platform Computing's Platform ISF* Can Help

MARK BLACK, CLOUD ARCHITECT, PLATFORM COMPUTING
JAY MUELHOEFER, VP OF CLOUD MARKETING, PLATFORM COMPUTING
PARVIZ PEIRAVI, PRINCIPAL ARCHITECT, INTEL
MARCO RIGHINI, SOLUTION ARCHITECT, INTEL

Cloud computing is a paradigm shift in the way IT is developed, managed, and consumed. It provides infrastructure and computer resources as services. End users request IT services through a Web browser or a computer using an API. These services are provisioned from a pool of shared computing resources according to corporate standards and business policies. Each line of business (LOB)

SUPPORTING FLEXIBILITY AND CHOICE

can have its own self-managed, virtual private cloud to get the benefits of IT consolidation without losing the ability to self-manage and optimize the cloud for their business. However, the private cloud administrator still maintains overall control and defines the top-level constraints for each LOB. The concept is a cloud within a cloud within a cloud. Within each virtual LOB cloud, business policies optimize the cloud according to specific needs such as performance, efficiency, high availability, and scalability.

Enterprises looking to take advantage of the cloud do so for many reasons—chief among them to enhance their agility in response to changing business dynamics. Today's enterprise data center is facing tremendous pressure to both innovate with new cloud architectures and operate legacy applications and heterogeneous systems. With that reality, it's crucial to adopt an open cloud management solution that supports flexibility and choice.

Large enterprises continue to recognize the need for a private cloud to meet regulatory, security, and performance requirements. To achieve agility and cost objectives, enterprises are taking a page from Google and Amazon, using a new cloud data center stack to transform internal IT.

Cloud uses your company's investment in virtualization—such as with VMware vSphere* and other hypervisor providers—and must have a deep integration with virtual machine (VM) technologies. However, the new cloud data center requires more than just virtualization and traditional IT practices. Cloud management is a layer, purpose-built for cloud infrastructure and processes, that co-exists with legacy and bridges to the future. Key capabilities need to include

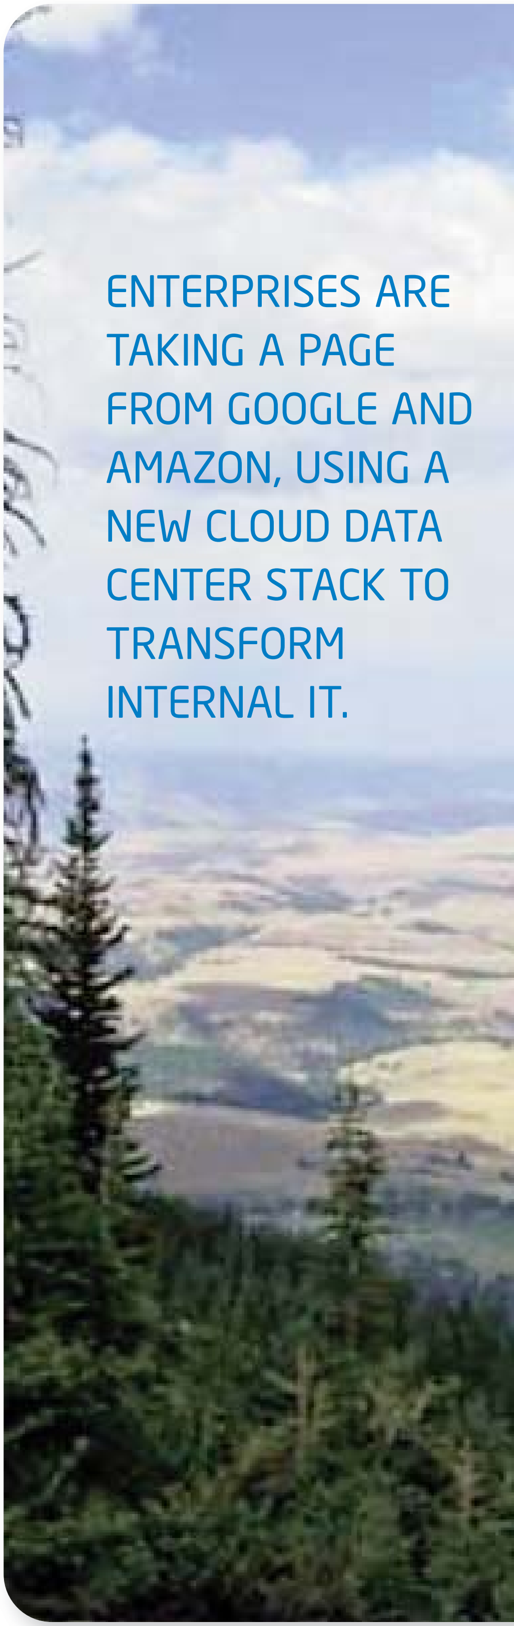

ENTERPRISES ARE TAKING A PAGE FROM GOOGLE AND AMAZON, USING A NEW CLOUD DATA CENTER STACK TO TRANSFORM INTERNAL IT.

MEETING ENTERPRISE DEMANDS

self-service and chargeback, policy-based automated provisioning of applications, dynamic scaling of applications to meet service-level agreements (SLAs), and unification of distributed and mixed-vendor resource pools for sharing.

Enterprises are currently demanding three more capabilities from private cloud management:

- NO VENDOR LOCK-IN: The ability to easily switch hypervisor and provisioning technologies and not be locked into any one major system vendor.
- A COMPREHENSIVE PRODUCT: A streamlined user experience from a single vendor instead of stitching together disparate offerings or multiple, complex tools.
- A PATH TO PRODUCTION APPLICATIONS: Enterprises recognize that infrastructure as a service (IaaS) offers tremendous benefits for development and testing, but that clouds also need to support applications. Cloud management systems must deliver fully operational, multi-tier application environments, or application clouds.

To meet enterprise requirements, Platform Computing offers the Platform ISF* solution. This next-generation cloud management modular software product helps IT organizations build and manage enterprise clouds that span both internal and external resource pools. Platform ISF's application-centric approach automates the self-service assembly and runtime management of the IaaS platform (middleware) as services, up to complex, multi-tier applications on top of shared heterogeneous resource pools. Data centers benefit with a solution that can support the entire application lifecycle, from development and testing to production-ready application clouds, in as little as 30 days.

Platform ISF offers a single platform that delivers:

- LOB SELF-MANAGED VIRTUAL CLOUDS: Enables hierarchical definition of clouds for each LOB to self-manage according to resource quotas and business policies.
- INFRASTRUCTURE TO APPLICATIONS: Supports IaaS, a customer's own platform as a service (PaaS), applications, and instance-specific software as a service (SaaS).
- DEVELOPMENT AND TEST TO PRODUCTION: Defines service templates for simple to complex multi-tier applications that support the entire application lifecycle.
- ALLOCATION AND RUNTIME MANAGEMENT: Manages both the allocation of the environment and the dynamic flexing according to changing workload levels and SLA requirements.

With Platform ISF you can:

- ELIMINATE OVER-PROVISIONING OF INFRASTRUCTURE to meet peak demand, resulting in lower capital and operating expenses and higher utilization.
- AUTOMATICALLY PROVISION APPLICATION ENVIRONMENTS and place workloads on the right systems to meet service levels in a timely and cost-effective way.
- KEEP VMWARE* ENVIRONMENTS OPEN with support for other VMs, physical provisioning, and external cloud providers.

Key capabilities of Platform ISF include:

- SERVICE MANAGEMENT. Self-service portals, account management, unlimited levels of hierarchical administration for ease of management, chargeback, and reservation management.
- ALLOCATION: Service catalogs, allocation engine for policy management, and runtime management.
- RESOURCE MANAGEMENT: Deep VMware vSphere* integration and support for Citrix Xen* and Intel® KVM technology, multiple physical provisioning tool integration, connectors to public clouds such as Amazon Elastic Compute Cloud* (Amazon EC2*), and an open API to other data center systems.
- OPERATIONS MANAGEMENT: A single cloud cockpit monitoring alarms and events.
- ENTERPRISE-CLASS SUPPORT: Global coverage and 24x7 hotline support.

Figure 1 shows how Platform ISF works. End users (at the top) are presented with a self-service portal from which they can request different offerings from the service catalog. Once the request is approved and within the user's allocation limits, a reservation is created for those resources. When it's time to start the request, the allocation engine locates the appropriate resources to run the service and creates the service on those resources. These resources can be obtained from either an internal or a public cloud, or both. Performance metrics are gathered and fed to the reporting and monitoring tools. The data is available to both users and administration staff. A policy engine allows the service to be scaled up or down automatically. The self-service portal allows the user to see the application and monitor the machines in the application. The user can access the services once they are created. The whole process, from the user requesting a service to delivery, can complete in a few minutes with no need for administrators to get involved.

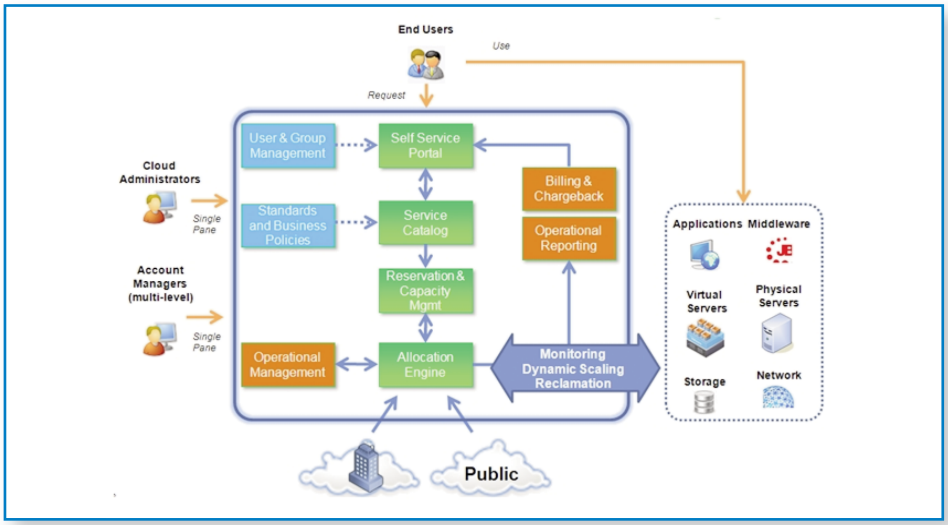

FIGURE 1. PLATFORM ISF USER FLOWS

MANAGEMENT FROM A SINGLE PANE

Cloud administrators can view and manage the cloud resources from a single pane. Different views in the user interface allow the administrator to quickly isolate fault locations. Along with remote log file viewing, remote command execution, and remote consoles, the administrator has all the tools to manage the cloud resources. The cloud administrator is responsible for defining the accounts that will use the cloud and for creating the initial service offerings.

The accounts are hierarchical, so once they are defined their management can be delegated to account owners (at the top level, this could be a business unit). The account owners can create other accounts (such as departments), subject to the limits imposed on their account, and add users to the accounts. They also have access to the billing and chargeback reports. Account owners can also create service offerings. These, in turn, can be offered to different sub-accounts (such as projects within departments within business units).

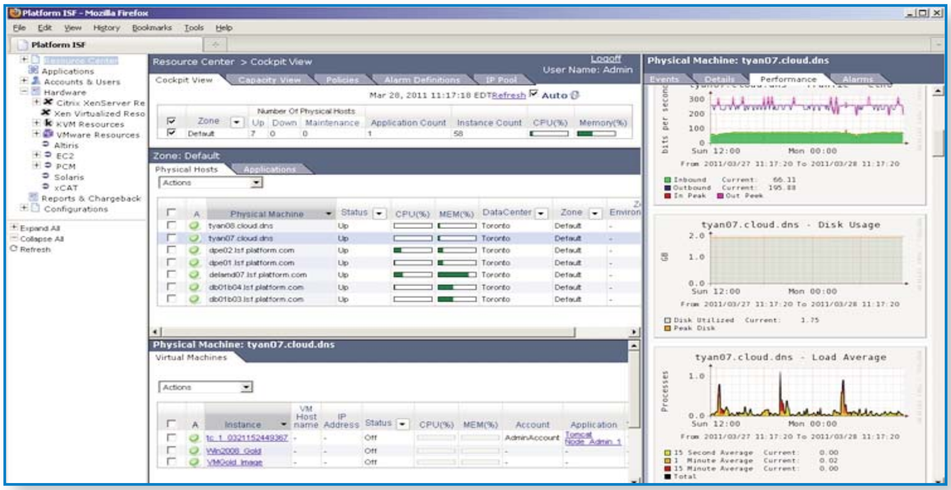

		Number Of Physical Hosts						
☑	Zone	Up	Down	Maintenance	Application Count	Instance Count	CPU(%)	Memory(%)
☑	Default	7	0	0	1	58		

□	A	Physical Machine	Status	CPU(%)	MEM(%)	DataCenter	Zone	Environ
□	✓	tyan08.cloud.dns	Up			Toronto	Default	-
□	✓	tyan07.cloud.dns	Up			Toronto	Default	-
□	✓	dpe02.lsf.platform.com	Up			Toronto	Default	-
□	✓	dpe01.lsf.platform.com	Up			Toronto	Default	-
□	✓	delamd07.lsf.platform.com	Up			Toronto	Default	-
□	✓	db01b04.lsf.platform.com	Up			Toronto	Default	-
□	✓	db01b03.lsf.platform.com	Up			Toronto	Default	-

□	A	Instance	VM Host name	IP Address	Status	CPU(%)	MEM(%)	Account	Application
□	✓	tc_1_0321152449367	-	-	Off			AdminAccount	Tomcat Node_Admin_1
□	✓	Win2008_Gold	-	-	Off			-	-
□	✓	VMGold_Image	-	-	Off			-	-

FIGURE 2. ADMINISTRATOR COCKPIT

Figure 2 shows the cloud administrator's view, with an interface optimized for managing a large-scale cloud without introducing unnecessary management complexity. The left frame shows a variety of virtualization and physical provisioning systems supported by Platform ISF. The top-middle frame summarizes the entire cloud including metrics organized by each data center. The middle frame connects the applications (services) in the cloud to the physical hosts. The bottom frame shows the VMs running on that physical host or as part of that application. This ties together the multiple layers of the stack in one easy-to-browse management view. The right frame shows the performance metrics gathered on the physical or virtual hosts, including any events and alarms.

PLATFORM ISF SOFTWARE ARCHITECTURE

Platform ISF integrates with major hypervisor technologies including VMware vSphere, Citrix Xen, Red Hat Xen*, and Red Hat KVM*. Where a hypervisor presents a central management interface, such as VMware vCenter* or Citrix XenCenter*, Platform ISF executes, controls, and monitors status through that interface. In cases where the hypervisor

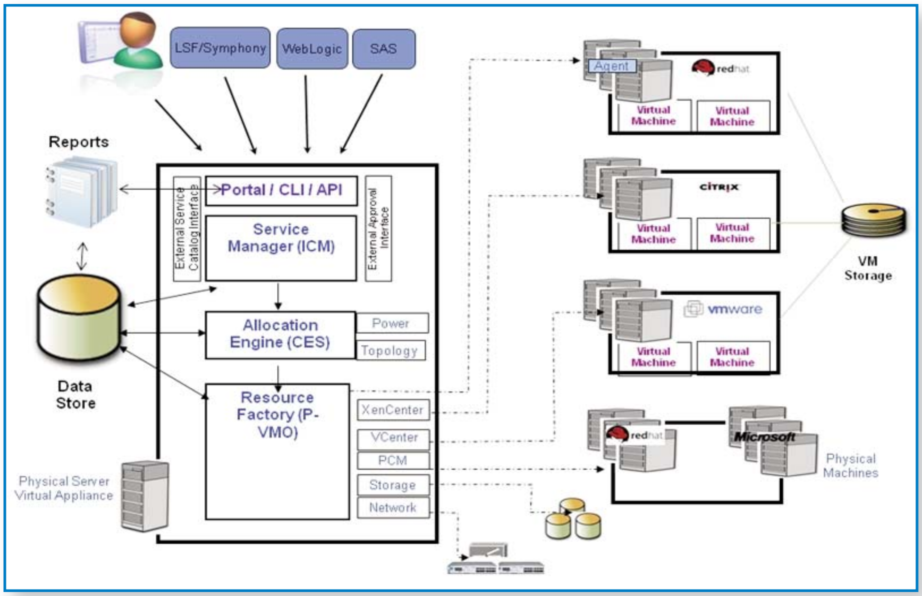

FIGURE 3. PLATFORM ISF SOFTWARE ARCHITECTURE

technology has no central manager, Platform ISF provides its own agent, which is deployed to each hypervisor node.

Platform ISF can manage and provision physical machines from bare metal using its adapter technology (Resource Interface Adapter*) with adapters to best-in-class provisioning tools such as those from Altiris, TPM, BladeLogic, and xCAT.

Platform ISF also provides its own provisioning tool, Platform Cluster Manager* (PCM*), which integrates with Platform ISF and can perform image- and packaged-based installs of the operating system and applications. PCM has also been adapted to provision hypervisor hosts to support demand-based shrink and grow of VM capacity by, say, switching nodes between VMware and Red Hat KVM. Finally, Platform ISF can also manipulate storage and network configuration to suit a particular workload.

Once integrated, Platform ISF can deploy a workload across any resources available (virtual or physical) through its Allocation Engine*, which can create and execute complex policies governing how a workload needs to be initially deployed and how it will behave over the lifecycle of an application. The Allocation Engine also lets the user make current and future reservations on a given request. All requests for a workload are submitted through a Web portal, where users are given a service catalog, published by the cloud administrator, from which to choose. The Service Manager (ICM) manages the application lifecycle (application definitions and instances) as well as users and accounts.

Platform ISF tracks the applications' duration and produces chargeback and billing reports based on the price set for each resource (e.g., CPU, MEM) and the time for which the application ran. It also produces capacity and allocation reports to manage the overall system.

PLATFORM ISF WITH INTEL® TRUSTED EXECUTION TECHNOLOGY

Without a doubt, cloud security is a hot topic and a major concern for enterprises. Cloud security covers a number of important areas including identity management, access control, data protection, data loss prevention, and hypervisor security. Although cloud access and data protection have been at the center of the security discussion, cloud infrastructure security and, specifically, hypervisor security have been getting special attention from enterprise IT. Pervasive use of virtualization technology among enterprises and cloud service providers, with the ability to move virtual machines between internal data centers or external to cloud service providers, raises the question of how enterprises ensure the security of their cloud infrastructures.

It isn't too difficult to imagine that one infected hypervisor or guest VM can spread the infection through virtualized infrastructure when moving from one host to another. While there are solutions to detect infected applications within guest VMs, there are far fewer solutions to detect infected hypervisors today. For this reason, Intel has developed a unique technology called Intel® Trusted Execution Technology (Intel® TXT).

Intel TXT helps prevent software-based attacks on currently unprotected areas, such as attempts to insert a non-trusted VM manager (VMM, or rootkit hypervisor), reset attacks designed to compromise platform secrets in memory, or BIOS and firmware update attacks.
To view it in a different way, Intel TXT enforces control through measurement, memory locking, and sealing secrets. To do this, it also works cooperatively with Intel® Virtualization Technology (Intel® VT).

An Intel TXT-enabled system requires all of the listed components—processor, chipset, TPM, enabled BIOS, and enabled hypervisor (VMM) or operating system. Without a complete set of these components, a trusted launch is not possible (Figure 4).

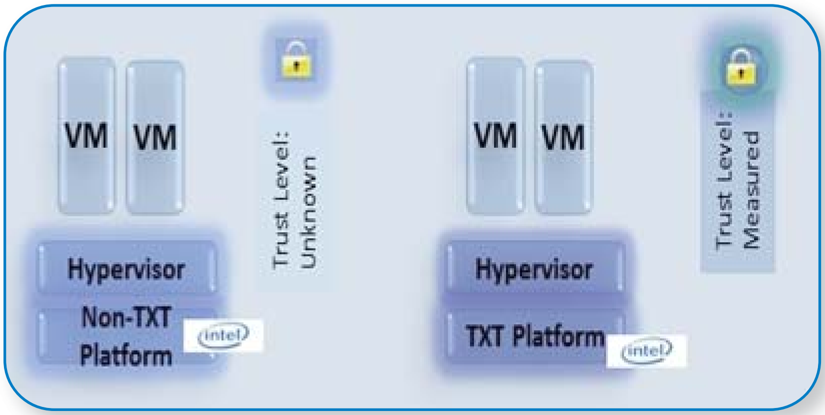

FIGURE 4. PLATFORM ATTESTATION AND SAFER VMM LAUNCH

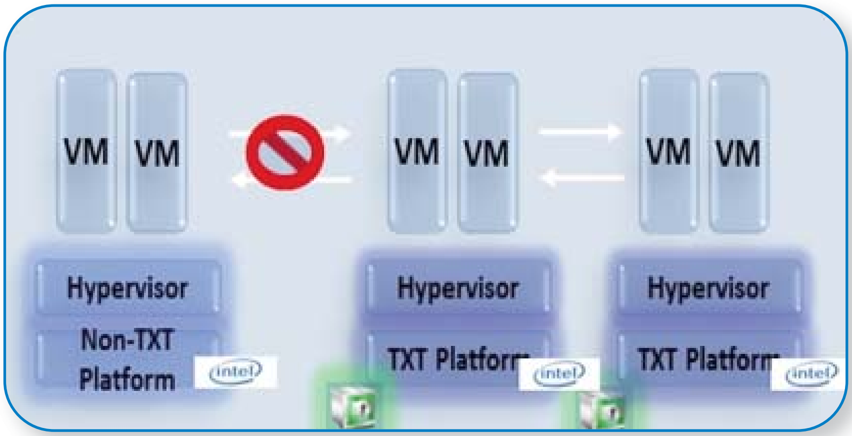

FIGURE 5. TRUSTABLE POOLS AND SECURE MIGRATION

Available on Intel® Xeon® processor 5600 series-based servers, Intel TXT is providing hardware-based protection in the processor, chipset, and third-party trusted platform modules (TPMs) that can better resist software attacks and make platforms more robust (Figure 5).

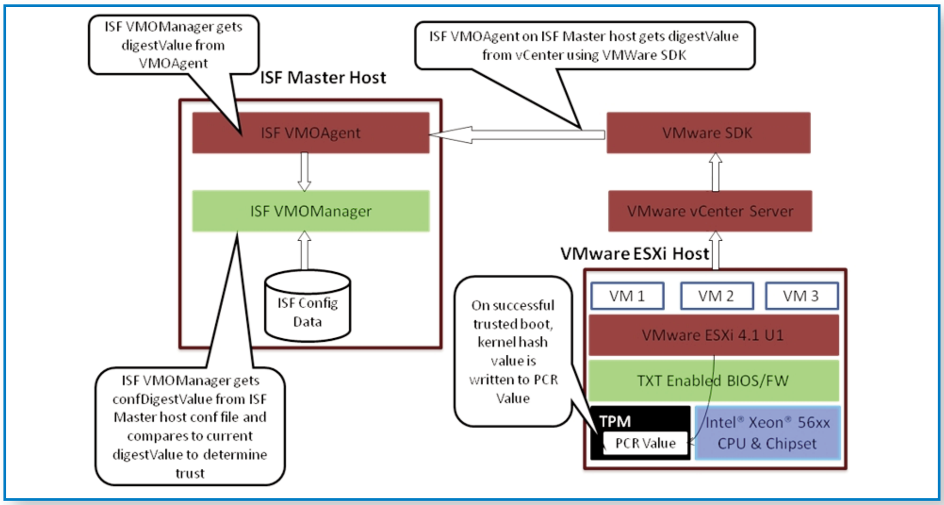
ISF VMOManager gets digestValue from VMOAgent
ISF VMOAgent on ISF Master host gets digestValue from vCenter using VMWare SDK
ISF Master Host
ISF VMOAgent
VMware SDK
ISF VMOManager
VMware vCenter Server
VMware ESXi Host
ISF Config Data
On successful trusted boot, kernel hash value is written to PCR Value
VM 1
VM 2
VM 3
VMware ESXi 4.1 U1
TXT Enabled BIOS/FW
ISF VMOManager gets confDigestValue from ISF Master host conf file and compares to current digestValue to determine trust
TPM
PCR Value
Intel® Xeon® 56xx CPU & Chipset

FIGURE 6. PLATFORM ISF HIGHLY SECURE ENVIRONMENT

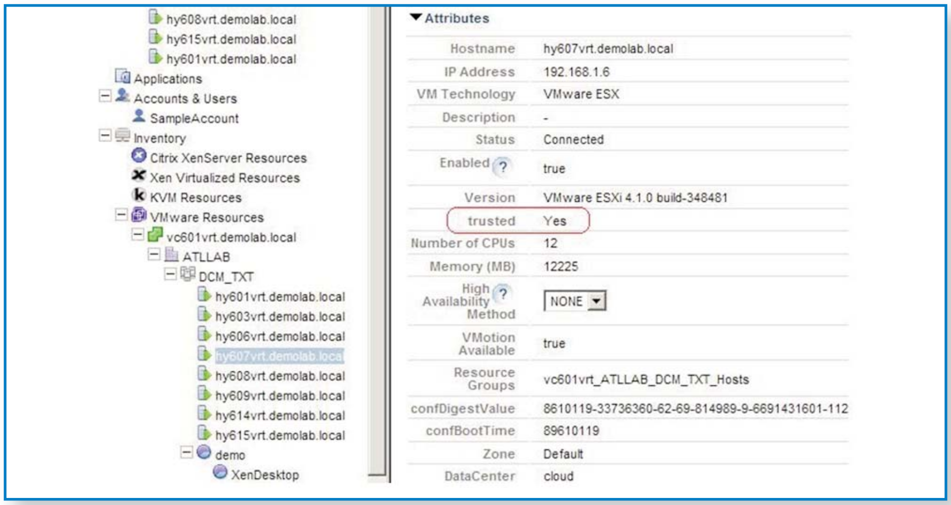

▼ Attributes

Hostname	hy607vrt.demolab.local
IP Address	192.168.1.6
VM Technology	VMware ESX
Description	-
Status	Connected
Enabled ?	true
Version	VMware ESXi 4.1.0 build-348481
trusted	Yes
Number of CPUs	12
Memory (MB)	12225
High Availability Method ?	NONE ▾
VMotion Available	true
Resource Groups	vc601vrt_ATLLAB_DCM_TXT_Hosts
confDigestValue	8610119-33736360-62-69-814989-9-6691431601-112
confBootTime	89610119
Zone	Default
DataCenter	cloud

FIGURE 7. HOST MARKED AS TRUSTED

Platform ISF creates a highly secure environment for running applications in VMs through its integration with Intel TXT, which monitors changes to the BIOS and boot processes of a hypervisor host to ensure it has not been tampered with. When a new host is added to the system, the administrator indicates it's a "trusted" host. Platform ISF then queries Intel TXT for the host's digest value, a cryptographic hash value coded to a number of metrics that measures the unique characteristics of the host's boot sequence. The digest value remains constant as long as the hypervisor installed on the host is not modified. Platform ISF saves this good digest value in its internal database for later comparison (Figure 6).

Each time a host connects to Platform ISF (e.g., after rebooting), the VMOAgent component retrieves the host's current digest value and passes it to the VMOManager component to compare it to the known good value saved. If the values are equal, Platform ISF marks the host as "trusted" (Figure 7).

With this knowledge of trusted and untrusted hosts, Platform ISF enables a number of important policies. Applications can be easily restricted to running only on trusted

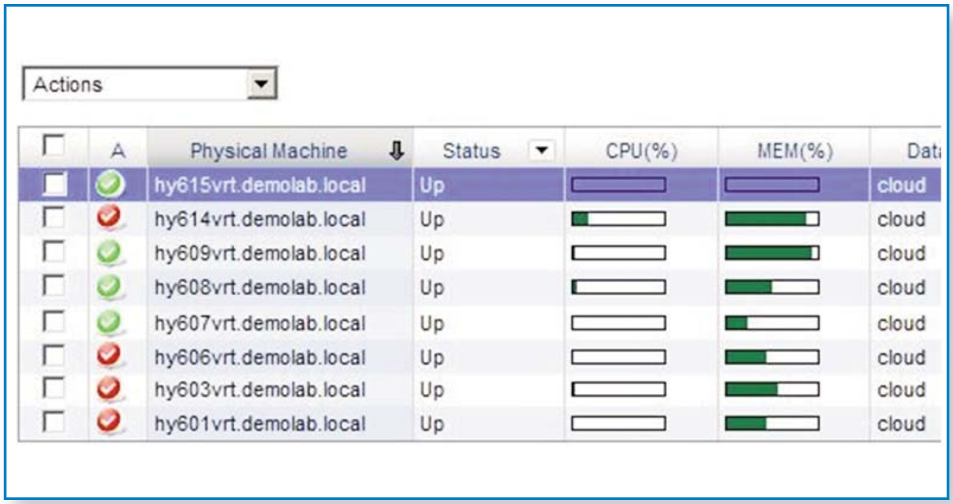

Actions ▾

□	A	Physical Machine ⇩	Status ▾	CPU(%)	MEM(%)	Dat
□	🟢	hy615vrt.demolab.local	Up			cloud
□	🔴	hy614vrt.demolab.local	Up			cloud
□	🟢	hy609vrt.demolab.local	Up			cloud
□	🟢	hy608vrt.demolab.local	Up			cloud
□	🟢	hy607vrt.demolab.local	Up			cloud
□	🔴	hy606vrt.demolab.local	Up			cloud
□	🔴	hy603vrt.demolab.local	Up			cloud
□	🔴	hy601vrt.demolab.local	Up			cloud

FIGURE 8. RED FLAGS INDICATE A PROBLEM

hosts by specifying "trusted==1" in the application definition. This controls not only which hosts are initially chosen to start the application VMs on, but also which hosts VMs are allowed to migrate to. Platform ISF ensures only trusted hosts are used.

Platform ISF can also notify the administrator when a host becomes untrusted by using the configurable alarm feature of ISF. When Digest Values do not match, a red flag next to the host's name indicates a problem (Figure 8). The administrator can then drill down and find the cause of the issue.

It's important to remember that in building secure cloud services, you need to incorporate multi-layer security architecture from the start. The Platform ISF and Intel TXT solution provides a foundation for building secure infrastructure while relying on other solutions to address requirements for data protection and data loss prevention.

BUILDING A PRIVATE CLOUD

In future articles, we'll discuss how to build private cloud using Platform ISF.

For more information, visit the Platform Computing Private Cloud website at www.platform.com/privatecloud or contact the authors: Mark Black (mblack@platform.com), Jay Muelhoefer (muelhoefer@platform.com), Parviz Peiravi (parviz.peiravi@intel.com), or Marco Righini (marco.righini@intel.com).

Back to Contents

IS CLOUD COMPUTING THERE YET?

How Companies are Monetizing the Network and Completing the Cloud

GEOFF BROWN, CEO, MACHINE-TO-MACHINE INTELLIGENCE

Clouds—with their simplified, automated interfaces and dynamic, responsive behavior—have been positioned as an answer to the growing need to support ever-increasing numbers of IT users demanding higher-bandwidth application content on all their mobile and fixed devices simultaneously. That's a tall order for traditional computing architectures.

WHY INCLUDE THE NETWORK IN THE CLOUD?

Many corporate IT departments have begun the foray into cloud computing by using server and storage virtualization technologies to gain efficiencies and explore greater compute scalability and elasticity. These technologies have made a great impact on global data centers—but haven't been able to fully deliver the dynamic, fluid, and adaptable cloud environment that promises even greater resource optimization and, more importantly, new pricing and business models.

Network virtualization is technology that allows layers of security and connectivity services to be deployed over heterogeneous, multi-vendor devices. The term "virtualization" applies since the technology abstracts applications away from propriety holes in and between incompatible devices. Incompatibilities among devices make networks complex and difficult to maintain. Human intervention often leads to mistakes, so it's quite common to see more than 75 percent of traditional IT budgets spent purely on networking. Cloud computing has highlighted the desperate need for advances in computer security, while not a day goes by without yet another high-profile cyber security breach.

WHERE DOES VIRTUALIZATION STOP AND THE CLOUD BEGIN? AT THE NETWORK, OF COURSE

Network virtualization does for data center connectivity and security what server virtualization did for the CPU. Network virtualization dynamically synchronizes access and security with global cloud and application management and business policies. This allows user access control and security to be allocated as global resources, just like compute and storage. The benefits of network virtualization include end-to-end, unified security; fully automated provisioning; dynamic access migration; mobile device integration; automated threat response and mitigation; and standardized, dependable, best-practice network configuration and change.

Why include the network in the cloud? Because:

- ALL END USERS access the cloud application through the network.
- THE NETWORK IS HOME to most security tools and devices.
- THE NETWORK IS COMPLEX and labor-intensive and requires black magic to maintain.

But all these factors are costing us a bundle.

IT experts often refer to "the wilds of the data center." The network is a big reason for the tangled jungle metaphor. The network is complex and heterogeneous, with multiple device types and vendors, plus documentation and controls that are often inadequate. This makes network virtualization a greater challenge than the relatively homogeneous server or storage environments, but it also makes network virtualization crucial to a streamlined, well-behaved cloud. Companies like Machine-to-Machine Intelligence Corporation (m2mi) are here to help by offering network virtualization solutions that knit together the disparate functionalities and translate across vendors and operating systems.

Network virtualization products abstract vendor differences and orchestrate changes across the network to apply and enforce high-level cloud business, access, and security rules in an intelligent "just work" paradigm. For instance, m2mi's Intelligent Network Virtualization* maintains a known, stable network state at all times. Cloud network services first verify the state of the network, auto-discovering the relevant network path as needed, and then ensure the desired change is safe given the network's actual state and make the desired changes. If downstream errors occur, all previous changes are rolled back and an error root cause analysis returns detailed configuration conflicts or automatically fixes common problems. Network Virtualization keeps the cloud traffic flowing through well-defined security layers and perimeter filters.

CLOUD SECURITY: THE BUCK STOPS HERE

The network is a cloud security battlefield—the first and final layer of defense where most attacks start, from distributed denial of service to virus intrusion. For the cloud to be secure, network virtualization must be able to synchronize and coordinate the varied security measures in network and security devices. With m2mi's Network Virtualization coordinates, security and access technologies like the Intel® SOA Expressway and Intel® Expressway Cloud Access 360 turn into sequential layers of dynamic protection, customizing application-specific security settings and augmenting them with automated response capabilities.

SECURITY BEGINS WITH EVERY CLOUD CHANGE, EVERY TIME

At the heart of security and connectivity for cloud computing, traditional IT, and smart grids is the concept of provisioning. Derived from telecommunications providers, the idea of provisioning revolves around the concept of enabling a service—such as voice or data—over a machine network. In network virtualization, provisioning is the process of enabling the correct connectivity, configuration, and security among heterogeneous devices such as mobile platforms, firewalls, load balancers, fabrics, network switches, and virtual machines. Solutions like those offered by m2mi can abstract the complex procedures and configurations of individual security devices, access control tools, and network settings to enforce and automate security best practices. Cloud users set high-level global access and security policies and m2mi's network virtualization tools translate them into configurations, thresholds, and alerts.

From a security point of view, network virtualization must assume all underlying network devices. Not all traffic can be trusted without extensive verification and validation. Once all network elements and components are correctly identified, layered security filters and policies can block all unknown or unauthorized communications among devices, encrypting all critical trusted traffic. m2mi's Network Virtualization uses Intel® Trusted Execution Technology (Intel® TXT) in the cloud to enable encrypted communication between the CPU and the local operating system. Intel TXT and m2mi are used to deliver a trusted boot-up, securing the cloud from start to finish.

The approach of trust through verification and validation, default filtering of unknown traffic, and encryption of all critical communications leads to the highest levels of security possible. This approach can stop common cyber attacks such as distributed denial of services and wasteful irritants such as spam. One reason these cyber threats are so successful is because of IT's overreliance on blacklisting, where known malicious traffic sources are blocked but all others are treated as innocent until proven guilty. Cloud environments can't afford this approach, which wastes resources and bandwidth and leaves the system vulnerable to any previously unknown traffic source. Instead, a cloud need a security-line approach like an airport, where incoming traffic is white-listed (i.e., registered traffic sources are identified and then allowed to move through a reduced security line; average traffic sources are scrutinized and X-rayed; and sources labeled suspect are put through additional measures before being allowed to board the cloud). To accomplish this security differentiation and layering, network virtualization tools must be able coordinate and dynamically adjust traffic routes to sort traffic and sequence policies and filters. To aid in this process, m2mi Network Virtualization uses Intel® Expressway Cloud Access 360 with the network as an identity management and ultimate white-list authority. As traffic is sorted—first by the firewall and Intel Expressway Cloud Access 360, and then directed through serial security features and filters at each level of the network—additional security devices handle unknown or potentially suspect traffic before allowing it to enter the application, including application firewall and packet inspection technologies like Intel SOA Expressway. By integrating with Intel SOA Expressway, m2mi Network Virtualization can offer cloud application owners greater flexibility and customization of security policies specifically designed to protect their applications. Intelligently automating the security line from end to end allows flexibility to deal with threats and customized security policies—and makes cloud security much more robust.

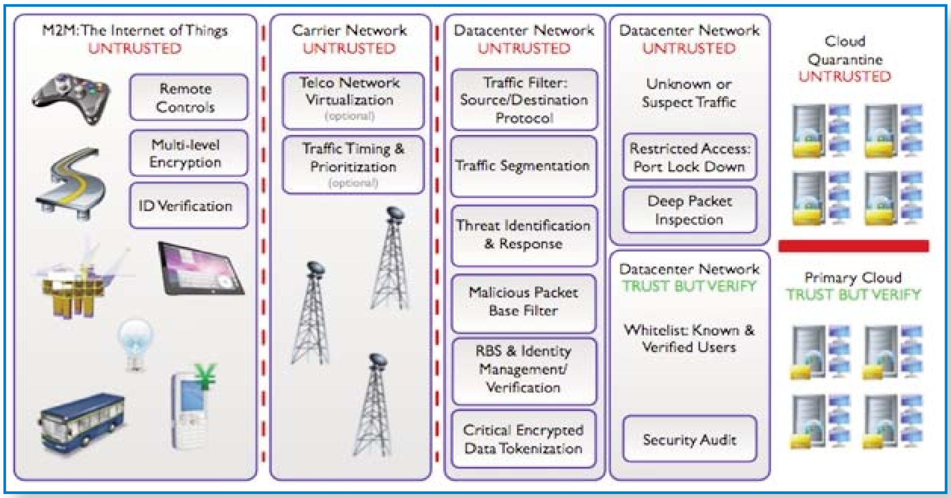
M2M: The Internet of Things
UNTRUSTED
Remote Controls
Multi-level Encryption
ID Verification
Carrier Network
UNTRUSTED
Telco Network Virtualization (optional)
Traffic Timing & Prioritization (optional)
Datacenter Network
UNTRUSTED
Traffic Filter: Source/Destination Protocol
Traffic Segmentation
Threat Identification & Response
Malicious Packet Base Filter
RBS & Identity Management/ Verification
Critical Encrypted Data Tokenization
Datacenter Network
UNTRUSTED
Unknown or Suspect Traffic
Restricted Access: Port Lock Down
Deep Packet Inspection
Datacenter Network
TRUST BUT VERIFY
Whitelist: Known & Verified Users
Security Audit
Cloud Quarantine
UNTRUSTED
Primary Cloud
TRUST BUT VERIFY

FIGURE 1. THE SMART GRID

MONETIZING THE CLOUD: THOSE NETWORK BELLS AND WHISTLES FINALLY PAY OFF

Network and server virtualization are billed differently in the cloud. Compute or server virtualization is sold on a usage model, which quickly resolves to commodity pricing (just ask your cell phone company). Although bandwidth is often sold in this model, the network device manufacturers have filled your network with wonderful features that network virtualization can expose as billable, premium services. These features include secure socket layer (SSL), packet inspection, quality of service (QoS), traffic prioritization, and data encryption. Most IT shops find these advanced network features too complex to configure or maintain across multiple vendors. But in a virtualized network environment, the configuration, metering, and maintenance are automated to follow best practices. This enables premium cloud add-on services to be delivered based on value-add instead of purely usage rate pricing models.

With support for Intel cloud products like Intel Expressway Cloud Access 360 for federated identity management, Intel® Expressway Tokenization Broker for credit card encryption compliance, and Intel SOA Expressway for application traffic inspection, m2mi Network Virtualization exposes network features and functionality as billable premium services.

SMART GRID: A CASE STUDY

M2M, or machine-to-machine, refers to the "Internet of things" or the vast, invisible network that connects oil and gas sensors, mobile phones, transport geolocation, and even the new smart meters being installed by utilities across the globe to better manage the diverse and dynamic resources of the next-generation smart grid. Think of all the environmental sensors being put into smart buildings or the number of mobile phones and tablets being used for computing assets. M2M is big business.

M2M, and particularly smart grid solutions, are natural cloud adopters and users. The cloud offers a simplified, on-demand platform capable of managing the high data and user traffic required. The challenge for M2M solutions in the cloud is how to securely manage across millions of remote mobile devices, through Telco networks, across enterprise environments, and safely through the cloud. The smart grid connects everything from remote smart meters to critical next-generation infrastructure. Simple cloud blacklisting and basic firewall policies just won't cut it. Network virtualization solutions that can segment and encrypt traffic—by security risk, by system permissions, by source, and by intended data use—are the only way to validate and verify all system traffic and comply with security requirements. From a cloud service provider standpoint, these additional security measures and layers are potential value-added options, above and beyond compute usage and bandwidth, for a smart grid or any M2M SaaS offering.

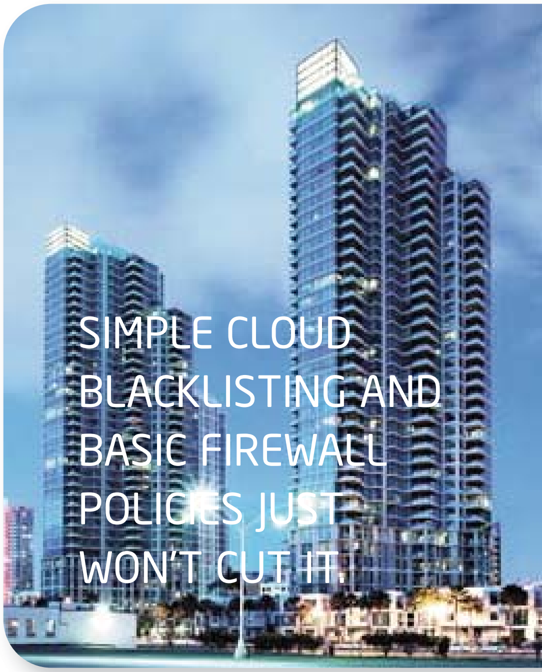
SIMPLE CLOUD BLACKLISTING AND BASIC FIREWALL POLICIES JUST WON'T CUT IT.

Geoff Brown is the CEO and Founder of Machine-To-Machine Intelligence Corporation (m2mi), based in the Silicon Valley at NASA Ames Research Park. Email him at geoff.brown@m2mi.com or visit m2mi at www.m2mi.com.

Back to Contents

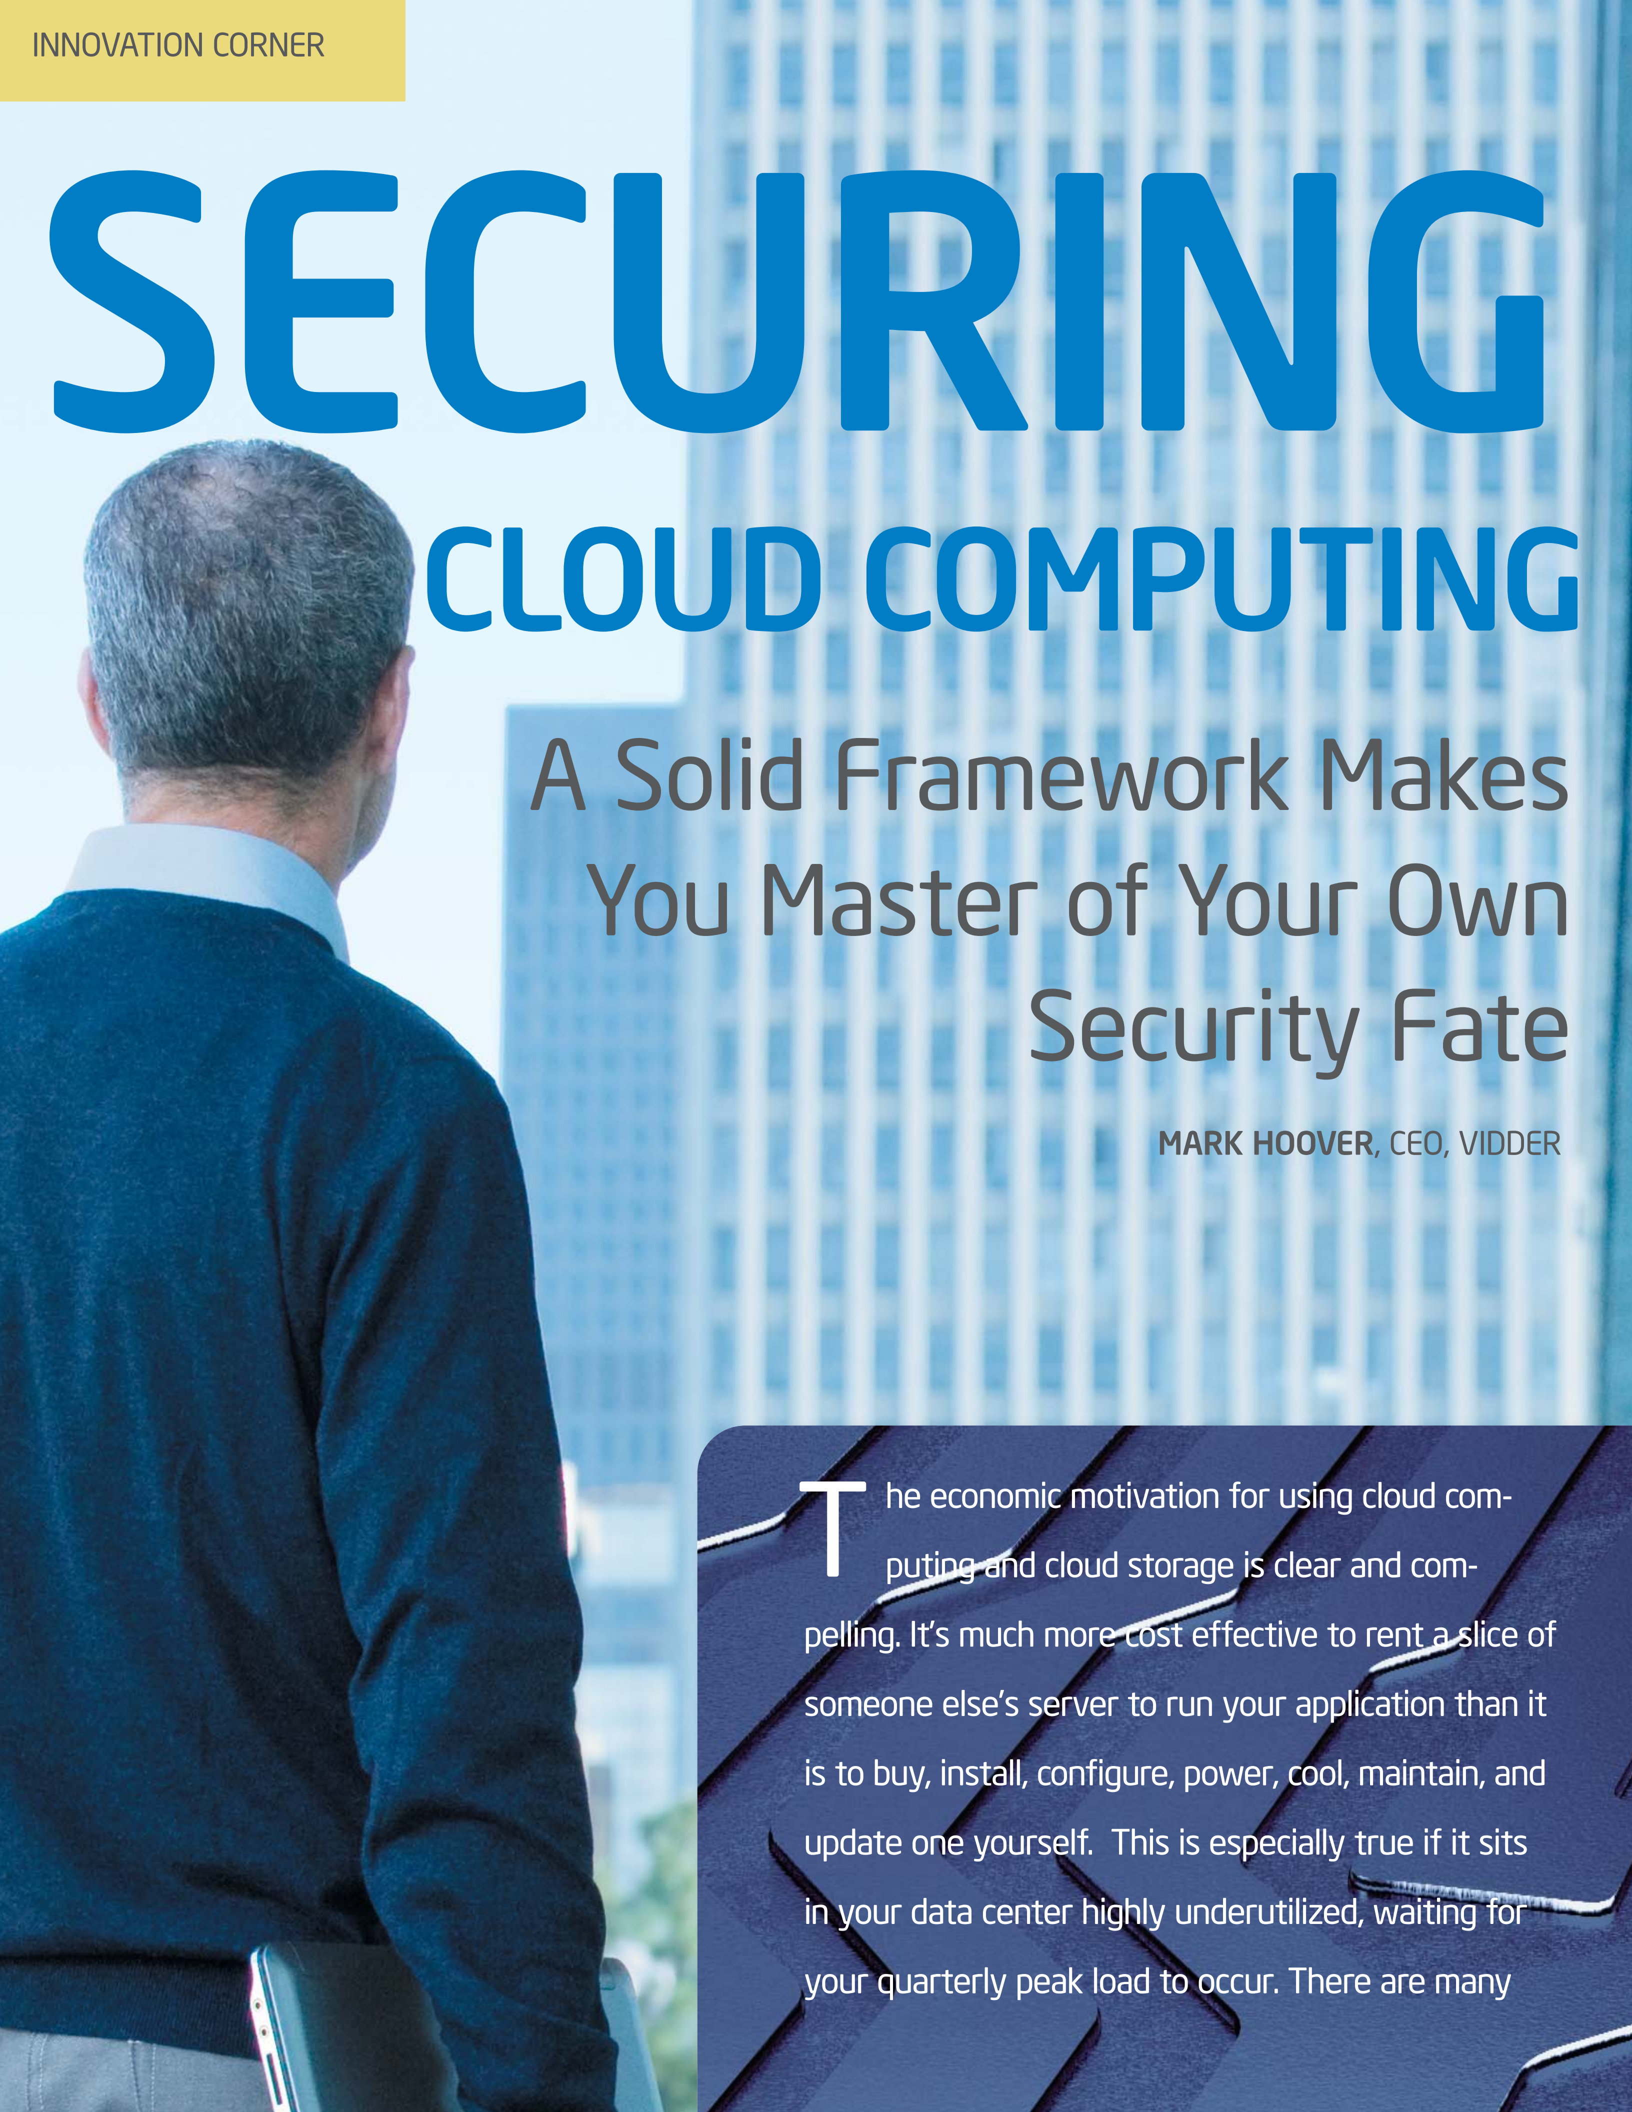

SECURING CLOUD COMPUTING

A Solid Framework Makes You Master of Your Own Security Fate

MARK HOOVER, CEO, VIDDER

The economic motivation for using cloud computing and cloud storage is clear and compelling. It's much more cost effective to rent a slice of someone else's server to run your application than it is to buy, install, configure, power, cool, maintain, and update one yourself. This is especially true if it sits in your data center highly underutilized, waiting for your quarterly peak load to occur. There are many

WHY HASN'T MOST ENTERPRISE COMPUTING MOVED TO THE CLOUD?

public cloud providers in the world today willing to serve you with a variety of different service suites and pricing models to fit your needs.

So why hasn't most enterprise computing and service delivery moved to cloud infrastructure? The answer might be that it's headed that way, but such evolutions take many years and we're only in the early stages.

One major, multi-dimensional issue that's slowed down enterprise use of cloud infrastructure is security. Unless the issues related to security are addressed in the coming years, the migration to cloud infrastructure will slow down and cloud will end up being used only for applications where security is not paramount.

What are the dimensions of the security problem related to cloud infrastructure usage? At a high level, you can put them into two categories: cloud security challenges and cloud outsourcing challenges.

CLOUD SECURITY CHALLENGES

Essentially, all cloud providers use server virtualization as an essential part of their infrastructures. That's what separates cloud from hosting. Virtual services, even when deployed from your own data center, introduce new challenges compared to securing fixed services.

CO-RESIDENCY

Virtualization leads to co-residency—instead of having a single application using the entire resources of a (probably highly underutilized) server platform, you have multiple applications sharing the resources of a common server.

Co-residency results in reduced security partitioning between co-resident applications. A compromise to one application could more easily

ONE MAJOR, MULTI-DIMENSIONAL ISSUE THAT 'S SLOWED DOWN ENTERPRISE USE OF CLOUD INFRASTRUCTURE IS SECURITY.

THE ISSUES: SOFTWARE TRUST AND DATA EXPOSURE

spread to a co-resident application. To the degree that these compromises are introduced by infected client systems, this means that all users on all co-resident applications represent a potential security threat to all applications. Any system-level BIOS, firmware, OS, or hypervisor vulnerability could affect all co-resident applications.

In a public cloud infrastructure, the issues related to co-residency become magnified because you don't know whose applications are running in the same operating environment. You don't know who their users are. You don't know if they're vulnerable to compromises that could leak toward you. You don't know if a co-resident application has been deployed by an attacker just to obtain more direct access to target applications (yours?). You don't know when your neighboring applications start, stop, and switch to new ones, making things like signature analysis and forensics increasingly difficult. And you don't know who ran on the same hardware before you and potentially left behind malware.

Co-residency issues can largely be broken down into two areas: software trust and data exposure.

SOFTWARE TRUST

Cloud providers have been very good at creating portals that allow you to launch servers at their sites based on configuration specifications you provide. It's like magic. All of a sudden you have Microsoft Server* 2008 image running at their site. But where did this software come from? How can you be sure it has all the patches and updates you need? How can you be sure this is a "clean" image like it would be if you launched it from your own image server or disk?

DATA EXPOSURE

Probably even more important than your trust in the software running at the cloud provider is the security around the data operating within your application—because it's your data. How can you be sure it's not compromised when mixed in with data from the cloud provider's other customers? How can you be sure it's not snooped on by other applications running over the same virtual and physical networks? How can you be sure an employee of the cloud provider doesn't access it via a backdoor?

INCREASED VISIBILITY

When you're running applications in your own data center for internal use, you can use internal (private) addressing schemes that are never exposed to the outside world. If you run these applications at a cloud provider, you need to use a pubic IP address to route to it over the Internet. Some cloud providers configure a public IP address on every server you launch at their site, whether communication over the Internet is needed or not. This increased visibility to potential attackers is something you need to defend against.

MAINTAINING AUTHENTICATION AND ACCESS CONTROL

Most enterprises have put a tremendous amount of effort and money into their internal identity management and access control systems. These systems can ensure that users are who they claim to be (authentication) and can access only the enterprise services and data they're supposed to (authorization). Achieving this across different services that are always evolving and being added to, with a user base of employees and authorized consultants and visitors who change on a daily basis, is no easy task. When you're moving services to the cloud, you should use and extend this capability for those services rather than depend on a new system and/or something supported by the cloud provider.

AGILITY AND SUPER AGILITY

Virtual services can be deployed anywhere in a cloud data center and possibly moved for performance or maintenance reasons. Virtual machines that are part of a cluster can also be launched and added to clusters or retired from clusters as needed. As if securing mobile clients weren't enough of a challenge, now you have to secure mobile servers as well.

CLOUD OUTSOURCING CHALLENGES

The act of deploying virtual services on somebody else's infrastructure adds additional business issues to the technological challenges we discussed in the previous section.

MISALIGNMENT OF CONTROL AND RISK

When running your virtual services in your own data center, you can be in total control of the physical access to the environment, the software running, the background of all employees associated with running the environment, and all other aspects of the security of the system. You can put controls in and monitor them in a way that matches cost and effort to risk.

When running your virtual services in a cloud provider's data center, you still have the same compliance, PR, and liability risks, but you have diminished control over the security techniques used to mitigate those risks. You're more dependent on the provider's approach, the ability of its employees to deliver and oversee the approach, and the trustworthiness of those employees as virtual insiders.

Cloud providers are motivated to do their best to ensure security to win and maintain your business. But still, in the end, they don't have quite the same motivation that you do. They don't have to pay the fines or deal with the PR issues that you do if sensitive information is leaked, especially if your customers or collaborators entrusted that information to you.

MIGRATION, PARTIAL MIGRATION, AND DE-MIGRATION

Your use of cloud infrastructure is bound to be dynamic. You may carefully move some services to the cloud, or possibly just part of a particular service delivery capability. You may decide to bring services back into your own data center after deploying for awhile in the cloud.

THE SOLUTION: A SECURE FRAMEWORK

PROVIDER INDEPENDENCE

You probably want to deploy services across multiple cloud providers so that you can benefit from the best pricing and most appropriate features. To achieve this, you don't want your cloud security solution to depend on any one provider's internal security capability.

A FRAMEWORK FOR SECURE CLOUD COMPUTING

Vidder believes the best way for an enterprise to achieve cloud security is to implement a solution itself that address the cloud security challenges, with no particular cooperation from or reliance on the cloud provider.

Such a solution must:

- ADDRESS ALL THE VIRTUALIZATION AND CLOUD SECURITY CHALLENGES we identified in the previous sections
- PROVIDE THE ENTERPRISE WITH A SECURITY FRAMEWORK totally under its control, regardless of which components of the solution are deployed internally and/or at various cloud provider sites

Achieving these goals is a matter of creating end-to-end trust that you can deploy flexibly across any environment, with enterprise control of all security policies and configurations.

Creating end-to-end trust starts with creating endpoint trust. One tool that can help is Intel® Trusted Execution Technology (Intel® TXT). Using cryptographic techniques, Intel TXT acts as the root of trust by ensuring that the foundational capabilities of an end system—the BIOS and the firmware—are uncompromised (known good). With that as a basis, you can examine any additional software launched in a similar way to ensure that only known-good software is running at start-up time for any system.

Such software can include communication agents that are trusted themselves, isolated from other system software components, and that can authenticate to a certified enterprise connection broker to request a connection to a specific service. If the user is properly authenticated and authorized to access the service, the connection broker can download any additional trusted client software required to access the service, as well as provide both the client and the target server the parameters required for them to establish a direct end-to-end encrypted connection.

End-to-end encryption is a crucial characteristic of such a system because it means the chain of trust is carried all the way into the enterprise's virtual machines (VMs) that may be deployed co-resident with other foreign applications in a shared cloud environment.

Using these techniques, the enterprise can establish a virtual security partition where only trusted servers, trusted users, and trusted connections between them reside. Trust can be maintained because the only interaction these trusted components have is with other

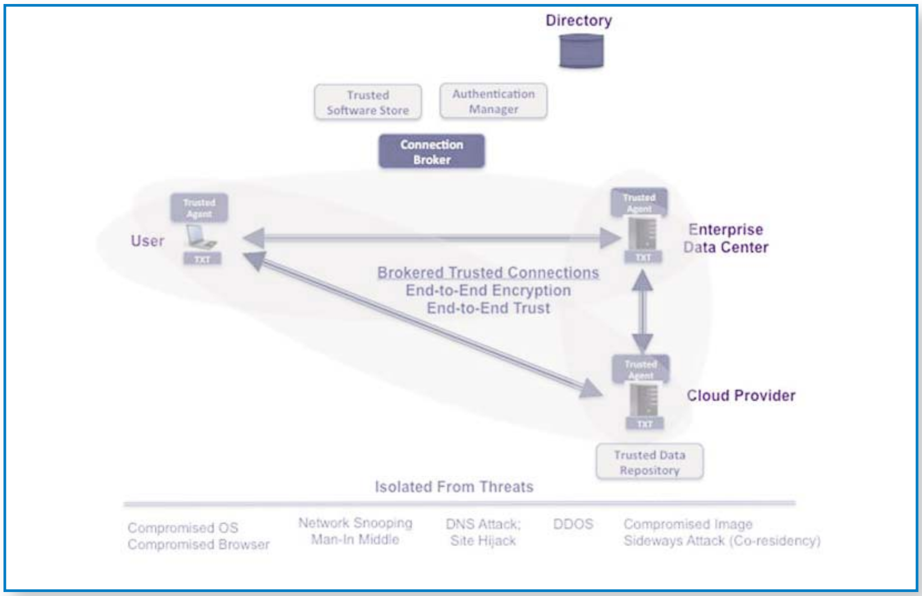

FIGURE 1. CLOUD SECURITY FRAMEWORK

trusted components, which isolates them from network attacks and compromises to co-resident end system software.

This framework for security exhibits many desirable characteristics to enable secure cloud computing utilization by enterprises:

- NETWORK OVERLAY: Operates transparently over any underlying network media, through firewalls, across VPNs and corporate networks, and across clouds.
- ENDPOINT-BASED: Every end system (e.g., user device or server) that participates runs a hardened and secure software agent tied to the endpoint root of trust.
- END-TO-END SECURE: Connected end points, including VMs, communicate with one another only over end-to-end encrypted tunnels. Security is maintained all the way from user-to-VM or VM-to-VM.
- HIDDEN APPLICATIONS: A combination of internal address resolution and network proxies hides the true location of distributed IT resources from attackers.
- INHERENT MOBILITY: Secure connectivity is maintained as devices, including servers and VMs, move. Activities such as the spawning of new VMs are automatically adjusted to.
- ENTERPRISE EMPOWERMENT: Enterprises can deploy the framework themselves, with no cooperation required from cloud providers, integrating them with their own directory, identity management, and authentication systems.

MEETING THE CHALLENGES

There are many challenges to providing the security needed so that enterprises can broadly benefit from the economic advantages of cloud computing. But it's a tenable problem. A solid framework will allow enterprises to be the masters of their own security fate when using cloud computing.

Mark Hoover is the CEO of Vidder (www.vidder.com), a privately-held Silicon Valley start-up focused on working with partners to create new and powerful security solutions for the extended enterprise and other demanding environments.

Back to Contents

CLIENT VIRTUALIZATION IN A CLOUD ENVIRONMENT

Identifying the Issues and Concerns

ENRIQUE CASTRO-LEON,
ENTERPRISE AND DATA CENTER ARCHITECT, INTEL
BERNARD GOLDEN, CEO, NAVICA
MIGUEL GOMEZ, TECHNOLOGY SPECIALIST,
TELEFÓNICA INVESTIGACIÓN Y DESARROLLO

Arguably, computation models seen in the client space are much more diverse than those in the server space proper. For servers, there are essentially two: the earlier model of static consolidation and the more recent, dynamic model where virtual machines are lightly bound to their physical hosts and can be moved around with relative ease. With virtualized clients, there are also two main models, depending on whether the application execution takes place in servers in a data center or on the physical client.

COMPUTATION MODELS

Beyond that, we have identified at least seven distinct variants, each architected to address specific management, security, and TCO needs and with usage models with specific business scenarios in mind. At least for server-based clients, their presence may be an indication of technology convergence between clients and server products in cloud space, a continuation of the trend that started when clients were used as presentation devices for traditional three-tier applications.

Let's look at some of the general issues and concerns regarding client virtualization.

Server-based computation models comprise session virtualization or terminal services (Table 1). Client-based models comprise OS streaming, remote OS boot, application streaming, virtual containers, and rich distributed computing (Table 2).

TABLE 1. SERVER-BASED CLIENT VIRTUALIZATION COMPUTATION MODELS

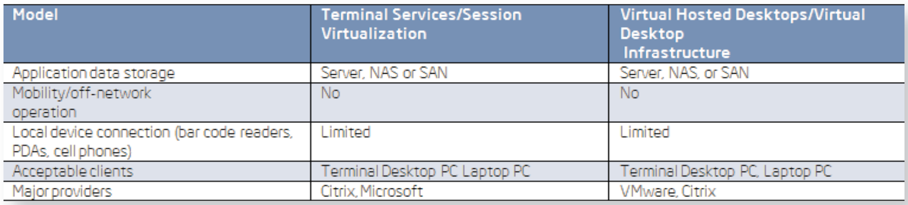

Model	Terminal Services/Session Virtualization	Virtual Hosted Desktops/Virtual Desktop Infrastructure
Application data storage	Server, NAS or SAN	Server, NAS, or SAN
Mobility/off-network operation	No	No
Local device connection (bar code readers, PDAs, cell phones)	Limited	Limited
Acceptable clients	Terminal Desktop PC Laptop PC	Terminal Desktop PC, Laptop PC
Major providers	Citrix, Microsoft	VMware, Citrix

TABLE 2. CLIENT-BASED CLIENT VIRTUALIZATION COMPUTATION MODELS

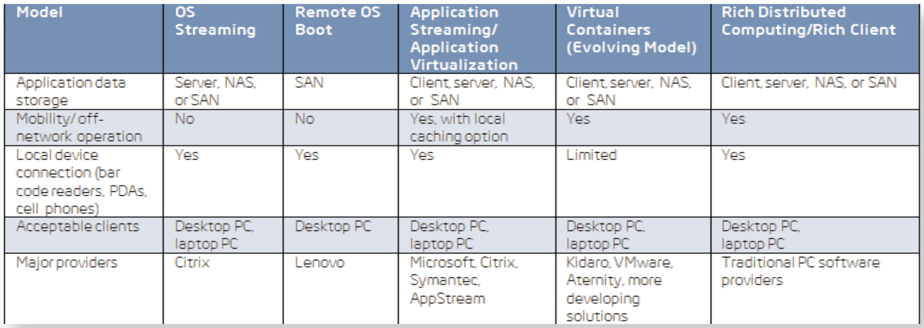

Model	OS Streaming	Remote OS Boot	Application Streaming/ Application Virtualization	Virtual Containers (Evolving Model)	Rich Distributed Computing/Rich Client
Application data storage	Server, NAS, or SAN	SAN	Client, server, NAS, or SAN	Client, server, NAS, or SAN	Client, server, NAS, or SAN
Mobility/ off-network operation	No	No	Yes, with local caching option	Yes	Yes
Local device connection (bar code readers, PDAs, cell phones)	Yes	Yes	Yes	Limited	Yes
Acceptable clients	Desktop PC, laptop PC	Desktop PC	Desktop PC, laptop PC	Desktop PC, laptop PC	Desktop PC, laptop PC
Major providers	Citrix	Lenovo	Microsoft, Citrix, Symantec, AppStream	Kidaro, VMware, Aternity, more developing solutions	Traditional PC software providers

Classifying blade PCs, such as those provided by HP or ClearCube, depends on whether users are assigned to blades in a one-to-one or one-to-many basis. If each user is assigned a single PC blade, the model most closely resembles rich client, except it can only be used in a fixed location and is constantly connected to the network. If an individual PC blade services multiple users simultaneously, the model more closely resembles a virtual hosted desktop.

Each of these computation models has appropriate uses based on the business scenario, user needs, and infrastructure requirements. Intel's position is that the client-side execution models provide the best user experience and can be deployed to meet IT requirements for security and manageability.

Each compute model places unique demands on the enterprise infrastructure.

INFRASTRUCTURE REQUIREMENTS

Moving large amounts of client computation, graphics, memory, and storage into a data center will likely require additional infrastructure build-out, unless the current equipment is grossly underutilized. Infrastructure issues to be considered include:

- SERVER computation capacity
- NETWORK BANDWIDTH, both wired and wireless
- STORAGE of user operating systems, applications, data, and customization profiles
- NEW CONNECTION brokers or remote access gateways
- NEW MANAGEMENT tools
- POWER DELIVERY for additional computation, graphics, memory, and storage now in the data center
- COOLING CAPACITY of the data center
- PHYSICAL SPACE within the data center
- PHYSICAL DISTANCE between the data center and associated clients

CLIENT DEVICES AND COMPUTE MODELS

Conversations around compute models often get intertwined with the device on which they will be deployed. The analysis becomes easier if devices and models are treated separately. For example, the business scenario may dictate server-based computing for a certain application, such as a patient information database. However, this thin client model need not be deployed on a thin terminal. A desktop or laptop PC may actually be a more appropriate device, depending on a user's total application and mobility needs.

MIXED COMPUTE MODELS

In most cases, IT will deploy a mix of computation models, depending on needs for data security, performance, and mobility. Individual users may have a hybrid of models. For example, a construction estimator in the field may use a cellular modem to access the centralized job scheduling tool via a terminal server session, but also have Microsoft Office* locally installed for word processing and spreadsheet work. The complete application and business needs of the user should be carefully parsed to understand which applications and data make sense to centralize versus install locally. Only in certain cases does a 100 percent server-side model make business sense.

SECURITY CONSIDERATIONS

There's no such thing as perfect security. Protection is always a matter of degrees. For simplicity, let's constrain security considerations to software-based attacks (e.g., viruses, worms, and software vulnerability exploits) and remote hacking, assuming that the user isn't a malicious attacker. Let's exclude hardware-based attacks, such as videotaping screen images or attaching purpose-built attack hardware. No compute model is inherently immune to that class of attacks.

BENCHMARKING APPLICATIONS

There are no industry-standard benchmarks for alternative compute models. Under the current state of the art, it's not meaningful to carry out performance comparisons across computation models. You can attempt performance comparisons between models if you measure a common workload, but even under these conditions, issues such as network loading, number of simultaneous users, server and network speed, workload content, and other factors can make simulation results much different from real-world deployments. It's important to evaluate performance claims carefully to understand how they apply to your real-world situations.

STREAMING AND APPLICATION VIRTUALIZATION

Streaming and application virtualization aren't synonyms, even though they're often used interchangeably. Streaming refers to the delivery method of sending the software over the network for execution on the client. Streamed software can be installed locally in the client operating system or, in most cases, it can be virtualized. With application virtualization, streamed software runs on an abstraction layer and doesn't install in the operating system registry or system files. When shut down, a virtualized application can be removed from the client or stored in a special local cache for faster launches or off-network use. The abstraction layer may limit how the virtualized application can interact with other applications. An advantage of application virtualization is that it can limit the continuous accumulation of randomness in the operating system registry and system folders that lead to system instability over time.

APPLICATION VERSUS IMAGE DELIVERY

A helpful way to think of the models and how they fit your requirements is to determine whether the problem needs to be solved at the application or image level (Table 3). In this case, an image is the complete package of the operating system and required applications. Some computation models solve application problems, some solve image problems. It's important to understand the your need in this area.

PUBLIC VERSUS PRIVATE IMAGES

When centrally distributing a complete desktop image with either virtual hosted desktop or operating system streaming, it's important to comprehend the difference between a common public image and a customized private image.

Public images are standardized operating system and application stacks managed, patched, and updated from a single location and distributed to all authorized users. Files and data created by the applications are stored separately. Customization of the image is minimal, but since all users access a single copy of the OS and application, storage requirements are relatively small.

TABLE 3. APPLICATION-LEVEL VERSUS IMAGE-LEVEL MODELS

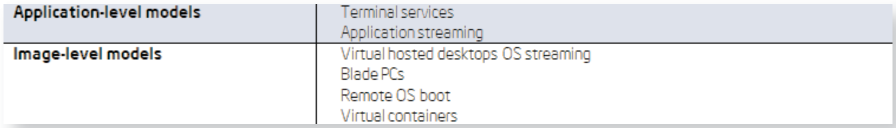

Application-level models	Terminal services Application streaming
Image-level models	Virtual hosted desktops OS streaming Blade PCs Remote OS boot Virtual containers

Private images are operating system and application stacks personalized to each user. Although users enjoy a great deal of customization, each private image must be stored and managed individually, much like managing rich, distributed clients. Current products don't allow private images to be patched or updated in their stored locations. Instead, they require them to be actively loaded and managed in-band on either the server or the client. The storage requirements for private images are much higher, since each user's copy of the operating system and application must be stored.

Find more information in the book *Creating the Infrastructure for Cloud Computing: An Essential Handbook for IT Professionals* by Enrique Castro-Leon, Bernard Golden, and Miguel Gomez.

Back to Contents

CLOUD TRANSFORMATION FRAMEWORKS

It's More than a Discussion about Hardware and Services

BOB DEUTSCHE, PRINCIPAL ARCHITECT, INTEL

By now, you've probably been approached by vendors telling you that your organization can streamline its operations and lower management costs by fully embracing a public or private cloud ecosystem. While this is true at a macro level, it generally ignores the many complex actions you need to have either in progress or in place to realize the promise of cloud computing.

SOLUTION AND FRAMEWORK CONSIDERATIONS

In this article, and others over coming issues, we'll try to separate cloud reality from the hype and help you come up with a roadmap to kick off your own sustainable and scalable cloud journey. The goal will be to give you an enterprise architect's perspective on building a solutions-based framework that can help you assess your ability to embrace the cloud and plan the actions you need to take to make the move.

Before diving into the specifics of the framework you need to implement your cloud ecosystem, let's define the generic considerations for a solutions framework (Table 1). To understand the specifics of constructing a cloud framework, it's important to first see why it's so crucial to focus on solutions if you're considering implementing a cloud ecosystem.

In a typical enterprise, work tends to take place in islands of organizational isolation (IOI). IOIs evolve out of organizational cultures that measure and reward results from a functional, not an enterprise, level. The potential sources of IOIs are many—and mostly a natural result of the way corporations make profits. IOIs can also come from a senior management structure made up of people with narrow skill sets. Regardless of the source, when you have IOIs you need to confront and resolve them because if you don't, they can kill your chances of successfully transitioning to a scalable cloud ecosystem.

TABLE 1. FOUNDATIONS FOR A CLOUD SOLUTIONS FRAMEWORK

HARDWARE	The physical aspects of the framework, servers, end-user devices, network components such as internal and external routers, and the Telco-sourced broadband capacity needed to make it all work.
SOFTWARE	Operating system plus applications and firmware. Considerations include understanding of network traffic generated by the software, open or proprietary sourced code, and data security plan.
IMPLEMENTATION	The "How do I do this successfully?" component. Considerations include balancing capability against the goals and SLAs of your extended ecosystem and time as a factor of expectations.
MANAGEMENT	Quality assurance at the enterprise planning, organization, financial, leadership, and controlling levels.
RESOURCING	How does your baseline project management plan balance against the internal skills required to complete the effort? What external resources will you need? How will you manage them?

Building a vibrant cloud ecosystem demands a level of interaction and cooperation between organizational and technical components that we haven't seen since the days of the mainframe (although, to be honest, the mainframe was more of an IT-driven model, which is the opposite of what cloud offers).

So what's involved in actually constructing a cloud solutions-based framework (CSBF)?

First consider the realities of your organization. By "realities" we mean deadlines, budgets, resources, and skill sets—which often dictate an enterprise's ability to develop an end-to-end solution. When you fully understand what a CSBF is before deploying your cloud environment, you have a basic framework you can build on as your cloud ecosystem matures.

CLOUD SOLUTIONS-BASED FRAMEWORK CONSIDERATIONS

- GEOPOLITICAL REQUIREMENTS. Where is your data stored? How do you access it? This varies by industry type, data type, and the legal requirements of your geography.
- INDUSTRY DIRECTION. Is your industry trending towards shared data or islands of information? Is it well defined or undergoing significant changes that might impact how information is used and shared?
- CORPORATE STRATEGY. Is direction from the top? Or is the cloud initiative driven by a single component of your business?
- ORGANIZATIONAL COHESIVENESS. How well do your business units work together? Do they value your internal IT organization? Are you a black hole into which the corporation continues to dump money?
- ENTERPRISE ECOSYSTEM AWARENESS AND STANDARDIZATION. How well defined is your company's enterprise architecture? Is standardization kind of adhered to, at least in spirit? Or does your environment proudly support one of everything?
- WILLINGNESS TO EMBRACE CHANGE. Implementing a cloud ecosystem, at any level, is disruptive. It forces change down to the business and operational levels. How have these types of actions been received historically in your enterprise? Is your company risk-averse or risk-inclined?
- MANAGEMENT. Cloud is disruptive and requires your organization to consider:
  - Establishing a balanced approach to effective service-level agreements (SLAs) based on business need as a factor of cost
  - Recognizing that in a highly virtualized and geographically diverse operating model, application oversight becomes extremely difficult
  - Evaluating the skill sets required to manage this environment and being open to the suggestion that you may not have the talent you need in your current organization
  - Predicting, proposing, and anticipating a management response based on awareness of these challenges in your company's history and culture

CORPORATE STRATEGY

We can't describe and resolve all the CSBF components you may need because they vary by geography, industry, and the particulars of your organization's management—as well as by management's ability to adapt to change. So let's concentrate on just one key component: corporate strategy. (We'll tackle ecosystem awareness and standardization in future articles.)

FUNDAMENTAL TRUTHS

To cut through some of the hype surrounding the cloud, let's start with a string of fundamental truths we can use as cornerstones for CSBF considerations, statements, and conclusions.

FUNDAMENTAL TRUTH 1: ADOPTING LARGE-SCALE CLOUD COMPUTING IS AN EIGHT-TO-10-YEAR JOURNEY

By its nature, cloud computing is a disruptive business capability. To actually work at an enterprise (not just at an IT) level, it demands repurposing formerly discrete and separate functions to depend on each other and provide utility and value back to the organization.

In a very simple analogy, this means that your combined data center(s), end-user networks (LAN, WAN, and wireless), application portfolios, and related business processes must all be able to co-exist and work together in a way that they were never designed to (at least since mainframe days).

Considering all the difficulties this implies, it's easy to conclude that to gain the benefits of cloud computing you must:

- DEFINE AND MEASURE entirely new success metrics
- IDENTIFY THE NEW NORMAL, which will require radical changes to your business ecosystems (skills, grade levels, responsibilities, etc.)

TO CUT THROUGH SOME OF THE HYPE SURROUNDING THE CLOUD, LET'S START WITH A STRING OF FUNDAMENTAL TRUTHS.

- TAKE AN EVOLUTIONARY VERSUS REVOLUTIONARY APPROACH to change, since revolutions are great for creating martyrs but not so good for surviving in the typical enterprise

FUNDAMENTAL TRUTH 2: CLOUD IS A TOP-DOWN ARCHITECTURAL FRAMEWORK THAT BINDS STRATEGY WITH SOLUTIONS DEVELOPMENT

In the name of simplicity, expedience, and pushing products and services, most vendors associate cloud with your data centers and related technologies. While there's a certain truth in this association, it's also important to remember that the CSBF is part of a much larger ecosystem where any part can impact any other part (see Fundamental Truth 1).

To successfully integrate cloud into your larger CSBF, you need to consider its potential impact on all facets of your enterprise architecture including business, data, applications, and technology (see Fundamental Truth 3).

At Intel, we view architecture at three levels: strategic, reference, and solution or service (Table 2).

In a perfect world, there would be some symmetry in the way one architecture type interacts with and drives its successor or predecessor.

TABLE 2. INTEL'S ARCHITECTURE LEVELS

ARCHITECTURE TYPE	USAGE
STRATEGIC	The longer-term (usually 18 months to five years) articulation of the context, priorities, and plans that set the boundary conditions and roadmap for developing tactical architecture. Conceptual in nature, it understands the high-level roles, capabilities, and processes to achieve the end state. The most abstract of the architectural types, its goal is strategic direction. Provides a template solution for an architecture for a particular domain. Its aim is to stress commonalty in areas such as vocabulary, boundaries, and guidelines to promote integration, standardization, and reuse.
REFERENCE	Characterized by a medium level of abstraction with the goal of identifying and standardizing solutions. Its artifacts form a starting point for solution architecture development. These may range from architectural patterns, mechanisms, and frameworks to complete systems with known characteristics. They may apply to a broad class of systems spanning domains, or have a narrower focus. Describing the scope and design of a change in business functionality, these are normally used to constrain and guide program and project design and implementation.
SOLUTION OR SERVICE	The solution or service architecture is not constrained to changes; in some cases, it can identify a gap and develop a solution to improve the efficiency of existing or ongoing business functionality. A solution architecture must align with the strategic architecture and should reuse or align to the reference architecture as much as possible. Service architecture is an architectural strategy that aims to isolate and separate consumption of business functionality from the provisioning of a function through commonly-defined service contracts.

In the real world, this isn't how things work.

In many organizations, one business unit often initiates an enterprise's need to move to a cloud environment. This is backed up in a recent Forrester research paper, which concludes that end users are driving cloud infrastructure as a service (IaaS) adoption.[1] These informal buyers aren't IT operations staff—they're just looking for quicker and more flexible resources than their enterprise IT organizations are able (or willing) to provide.

If Forrester's research is true, it doesn't bode well for building scalable enterprise cloud ecosystems. In Intel's three-level architecture model (Table 2), business functionality is part of the solution or service architecture. If we follow the logic of the architecture framework, an end-user-driven solution generally doesn't support enterprise-wide business objectives. It's like the tail trying to wag the dog. While it might provide a good solution for one business unit, it may not scale well unless you consider higher-value strategic and reference architectures as part of the roll-out.

FUNDAMENTAL TRUTH 3: YOUR CLOUD ECOSYSTEM IS ONLY AS ROBUST AND ADAPTABLE AS THE SUM OF ITS PARTS

To understand the meaning of this fundamental truth, try describing your business from end to end in just a few words. Having trouble? Fortunately, industry makes it easier through the discipline of enterprise architecture (EA). According to Wikipedia, "EA is a rigorous description of the structure of an enterprise which comprises enterprise components, the externally visible properties of these components, and the relationships between them...This description is comprehensive, including enterprise goals, business processes, roles, organizational structures, and organizational behaviors."

In mature organizations (defined on a 0 to 5 scale), the EA framework is a key component of success, linking the value of the IT organization to the objectives of the enterprise. In practice, though, this is easier to say than to do.

AN END-USER-DRIVEN SOLUTION GENERALLY DOESN'T SUPPORT ENTERPRISE-WIDE BUSINESS OBJECTIVES. IT'S LIKE THE TAIL TRYING TO WAG THE DOG.

EA has four related, but distinct, architectural components (Table 3).

TABLE 3. ARCHITECTURAL COMPONENTS

ARCHITECTURE TYPE	DESCRIPTION
BUSINESS	The business architecture serves as the interface between the needs of the enterprise as reflected in its work and the IT solutions that facilitate that business. The business processes serve as the foundation for a number of important architectural decisions in the balance of the enterprise architectural domains.
DATA	The data architecture is intended to promote information sharing and reuse in support of business processes across the enterprise. This is done via standard description and discovery of common data and the promotion of uniform data management practices.
APPLICATION	Specifies the key elements of information systems Intel uses in executing its business processes. These elements include the services taxonomy and its components. They represent the organization's application portfolio and identify the business systems that enable and support the execution of Intel's business processes (outlined in the business architecture). The application architecture provides a cross-reference of capability or service components to business functions and processes to illustrate application boundaries.
TECHNICAL	Describes current and future technology infrastructure and specific hardware and software technologies that support your corporate information systems. It provides guidance and standards for implementing technologies that are proven to work well with existing and planned technologies.

As we've discussed, successful cloud adoption requires interaction and cooperation between the organizational and technical elements of your company. Simply put, you must understand your enterprise as a factor of its EA elements to be able to scale cloud adoption.

FUNDAMENTAL TRUTH 4: TECHNOLOGY-DRIVEN BUSINESS PRACTICES OFTEN CIRCUMVENT GOVERNMENT REGULATIONS, BUT LEGAL STANDARDS WILL DICTATE CLOUD'S SUCCESS

This is easy to predict in general, but hard to fine-tune for specific cases.

Consider the example of cloud data. The European Union (EU) recently approved data security regulations that establish a cloud code of conduct, founded on EU data protection standards, for all data pertaining to EU citizens. Viviane Reding, vice president of the European Commission and EU justice commissioner, recently proposed four enhancements to these codes of conduct[2]:

1. Companies outside the EU that directly target their activities to EU citizens will need to abide by the new EU data protection rules.
2. A "data protection by design" principle should reinforce existing rules on the security of processing as well as liability of those who control and process data.
3. Data transfer rules need to be revised to streamline and strengthen procedures for international data transfers.
4. The EU needs a way for third-country providers to voluntarily adhere to EU data protection rules. Such a mechanism might include certification and guarantees for auditing and enforcement. Over time, this could become an EU safe harbor system.

In the U.S., we can find more examples in the vagaries of laws from state to state regarding location, privacy, and access to health records stored in the cloud. Every state has its own regulations on patient data. At the federal level, there doesn't yet seem to be any specific guidance regarding Health Insurance Portability and Accountability Act (HIPAA) compliance for patient data in a cloud ecosystem. But it's likely only a matter of time before we see some federal regulation of the storage and access to these records—probably initiated by some privacy breach of a patient's data.

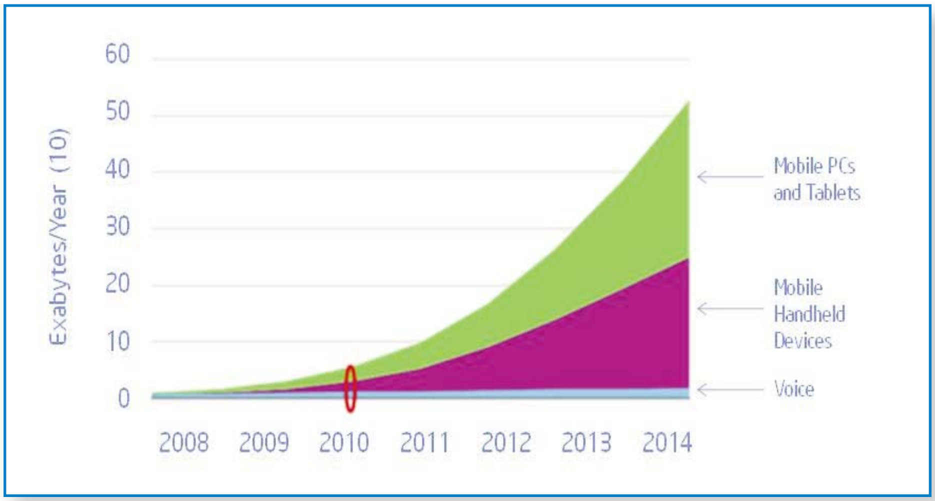

FIGURE 1. GROWING DEMAND FOR BROADBAND SERVICES

We can't really dissect the fine points of any of these examples here. But it's important to recognize that, at best, legal systems tend to be five to 10 years behind technology shifts. Until some fundamental issues are resolved—specifically around data storage and security—the widespread adoption of cloud will be impacted.

FUNDAMENTAL TRUTH 5: BANDWIDTH AND DATA TRANSMISSION MAY NOT ALWAYS BE AS INEXPENSIVE AND UNENCUMBERED AS THEY ARE TODAY.

In the U.S., we seem to believe that we'll be able to simultaneously download a live video stream of our favorite movies and TV series, have instant audio and video chats with friends and family, and keep track of our stock portfolios—all while playing online games. We also fully expect that broadband services used to deliver these services will be inexpensive if not free. It all comes back to our fundamental belief that broadband is an unlimited and unmetered resource.

The premise of free or inexpensive broadband is starting to change. At a recent Telco 2.0 conference, Ericsson predicted that by 2016 there will be around 15 times more mobile voice and data traffic than there is today, primarily driven by video.[3] (This is the traffic on the access side.) In particular, Ericsson stressed that today broadband fees are assessed based on time instead of the network volume used. As the slide was discussed by the conference attendees, there was implied consensus that this model will likely change because of the burgeoning amount of data transmitted over the network. In fact, there's ample evidence that these changes are already happening, at least in the U.S.

In May 2011, AT&T began imposing monthly data limits for its fixed broadband subscribers. With con-

consumers seeing broadband as an unassailable right, this direction is obviously not popular.

Net neutrality, sometimes referred to as data capping, is a topic of great interest and discussion for the U.S.-based Federal Communications Commission. There's every indication that data caps are already being imposed on a limited basis and that with growing business activity in the cloud, they will likely become part of your enterprise cloud business ecosystem—beginning in your data center, extending to how you design cloud-based application portfolios, and including the amount of intelligence in your end-user devices. Time will tell.

BUILDING YOUR CLOUD MODEL

In future articles, we'll use the principles we've discussed here to help you consider enterprise-specific factors for your own cloud adoption:

- BUSINESS PERSPECTIVES AND CHALLENGES including skill set implications and ecosystem considerations
- SERVICE DELIVERY MODELS, both today and tomorrow
- TRANSFORMATION STRATEGIES including maturity assessment and taking a product management office approach

If you want to discuss anything we've covered in this article, email robert.m.deutsche@intel.com. We're looking forward to hearing your point of view.

Back to Contents

[1] Frank E. Gillett. "Navigating the Shifts in Computing Infrastructure Markets" (Forrester Research, March 24, 2011).
[2] Viviane Reding. "The Reform of the EU Data Protection Directive: The Impact on Businesses" (European Business Summit, 2011).
[3] Håkan Eriksson. "The Networked Society" (New Digital Economics Conference, Palo Alto, California, April 2011).

BUILDING A CLOUD SECURITY PLATFORM

Minimizing Threats for a Trusted Compute Cloud

BRUNO DOMINGUES, SOLUTION ARCHITECT, INTEL

> *"Truth, like gold, is to be obtained not by its growth, but by washing away from it all that is not gold."*
>
> *–Leo Tolstoy*

A KEY PUZZLE PIECE: INTEGRITY OF THE CLOUD PLATFORM

Security is a key concern for cloud computing. A recent Forester Research survey on cloud computing adoption backs this up, with 30 percent of respondents agreeing. Data privacy issues were also important for 25 percent of respondents, since security and integrity threats to the infrastructure can expose data and compromise the availability of the whole environment[1]. One of the most important pieces of the cloud security puzzle is the integrity of the platform.

THREATS AGAINST THE INTEGRITY OF THE CLOUD INFRASTRUCTURE

Security threats against operating systems and applications aren't new. But attacks focused on the pre-runtime environment and hypervisor *are* relatively new—and getting more emphasis with virtualization and cloud computing adoption.

Malware such as Blue Pill*, introduced by Joanna Rutkowska, and SubVirt rootkit*, developed as a proof of concept by Microsoft, are virtually undetectable due to their ultra-thin hypervisor approach and ability to virtualize the hypervisor. They are also immune to the most common methods of protecting against rootkit kernel mode (e.g., unlinking EPROCESS blocks from the kernel list of active processes, hooking the page fault handler, and marking some pages invalid).

There are also BIOS-level rootkits that can allow code to persist even if the hard disk is wiped clean, and can hide code, making it possible to execute before any piece of code. BIOS patching can help avoid this threat, with write-protected BIOS the best response.

Low-level attacks are usually hard to detect and can be very difficult to recover from. The best defense is threat security in depth: software- and hardware-based trust.

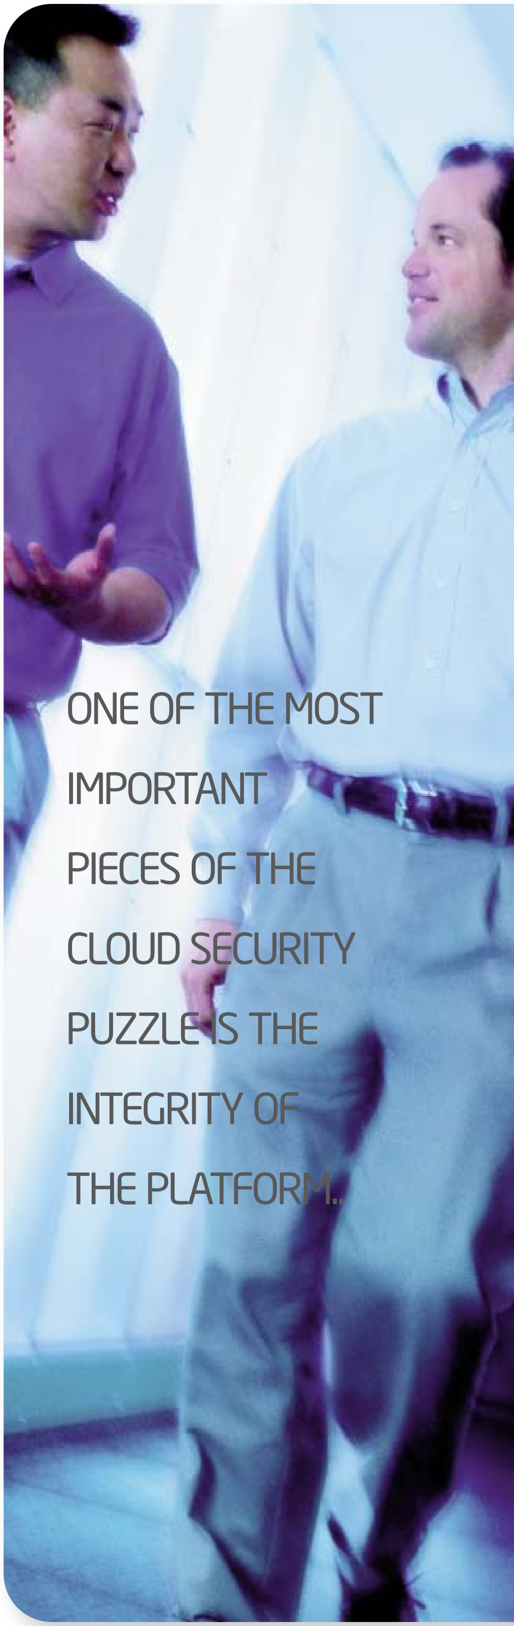

ONE OF THE MOST IMPORTANT PIECES OF THE CLOUD SECURITY PUZZLE IS THE INTEGRITY OF THE PLATFORM..

INTEGRITY MEASUREMENTS

Since the days of the Intel 286 architecture, every instruction executed on the x86 architecture can be executed in four different privilege levels, each defined by two bits. When an instruction is executed with the 00b (0d), it means the highest privilege level, called ring 0 or kernel mode. When these bits are equal to 11b (3d), it means the lowest privilege level called ring 3 or user mode.

In the early 1990s at USENET, Professor Andrew S. Tanembaum, a writer for *Computer Network*, and Linus Torvalds, creator of Linux*, discussed operating system models in terms of the Intel x86 architecture. Tanembaum defended the microkernel model, an elegant proposal from the implementation point of view since the system would be dispersed among the four privilege levels (Figure 1).

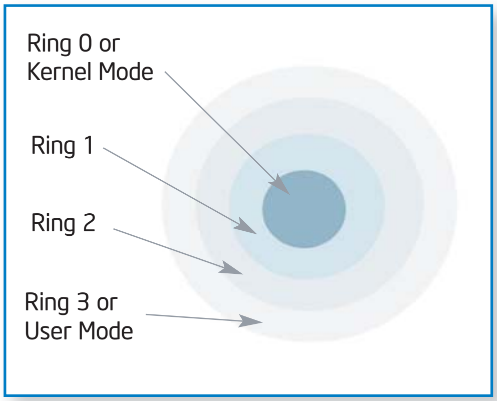

FIGURE 1. OPERATING SYSTEM MODEL

Maybe because of his programming background, Torvalds defended the monolithic kernel, a more pragmatic model where the kernel and device drivers have the same privilege context. From the operating system construction point of view, it becomes simplified because we can avoid context changes due to inter-process calls.

Time showed that Torvalds was right. Linux became popular and MINIX*, developed by Tanembaum, didn't. David Cutler, Microsoft architect and Windows* Internals author, also believed in the monolithic kernel that predominated.

Even so, Tanembaum's concerns were still relevant. From a security standpoint, creating security boarders isolating kernel, drivers, services, and applications is an advance over the Torvalds model. However, a monolithic kernel showed the fastest development cycle due to its simplicity and claimed performance. It's a controversial point, since there are examples of operating systems developed under the microkernel model with excellent performance, such as Cisco's IOS*, which is embedded into Cisco routers.

From a security point of view, in the monolithic kernel any vulnerability or malware loaded by device drivers, or any kind of code running in the ring 0 context, can compromise the entire system. In the microkernel model, where there is a minimum kernel footprint, it is non-extensible and isolated on its own. This trusted computing base (TCB) is resistant against attacks.

In a virtualized environment, hypervisor type 1 (i.e., a bare-metal hypervisor) runs in a ring 0 context more privileged than guest operating systems. What makes it possible is the VTx processor extension. Basically, a single bit in the CPU defines when an instruction will have unrestricted access to CPU registers (i.e., hypervisor or, with restricted access, a guest machine). This means compromising the hypervisor can compromise the whole environment and make detecting malware in this level much harder.

Security should be treated in depth—not only on guest operating systems, network access control lists, etc., but also on hypervisor integrity. There are two approaches to measure validate hypervisor integrity:

- STATIC ROOT trust of measurement (SRTM)
- DYNAMIC ROOT trust of measurement (DRTM)

Both methods rely on a tamper-resistant root trusted device that is used to measure the hypervisor launch. This trusted device is called trusted platform module (TPM), which is the root component of a secure platform. It's a passive I/O device usually located at the low-pin-count (LPC) bus. Nowadays, it can be found as part of the Northbridge chipset. TPM has special registers, called platform configuration registers (PCR), and can do some interesting things such as sealing or unsealing secrets, allowing quoting (remote attestation) and some crypto services such as RSA and PRNG.

TPM is based on PCR extend operations, using the previous PCR value to define the next one:

$$PCR_{N+1} = SHA1\ (PCR_N + Value)$$

A single PCR can be extended multiple times. It's computationally unfeasible to define a specified value to a PCR, so the order where things happen matters: $[(ext(A),ext(B)) \neq (ext(B),ext(A))]$

The secret sealed in a TPM can only be unsealed if the correct PCR values match (Figure 2).

TPM is used also by Microsoft BitLocker*, a full disk encryption technology, where the key to decrypt the disk is located in the TPM chip and the retrieval of this key depends on the integrity of the code that can be executed in memory. This process is known as static root trust of measurement (SRTM) (Figure 3).

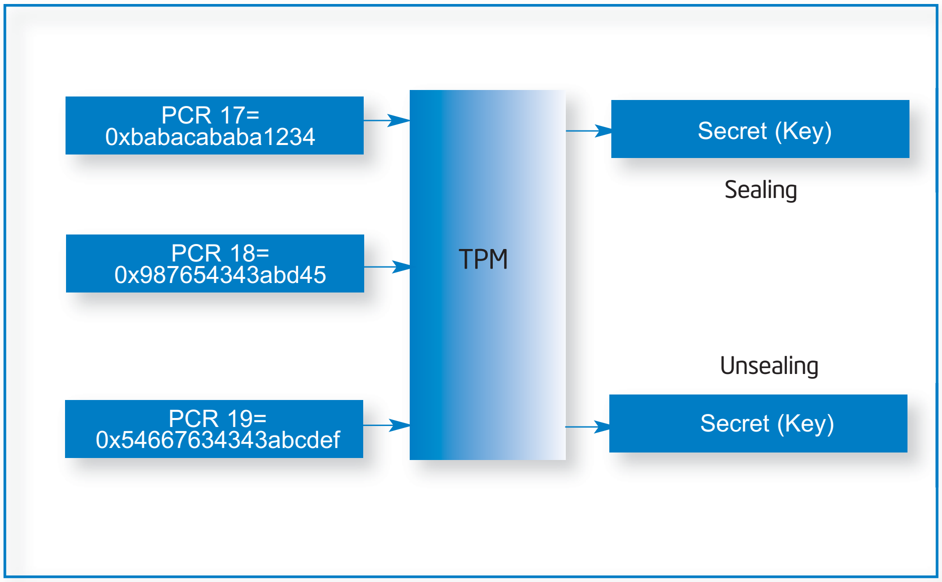

FIGURE 2.
SEALING/UNSEALING TPM OPERATION DUE TO PCR REGISTER MATCHING

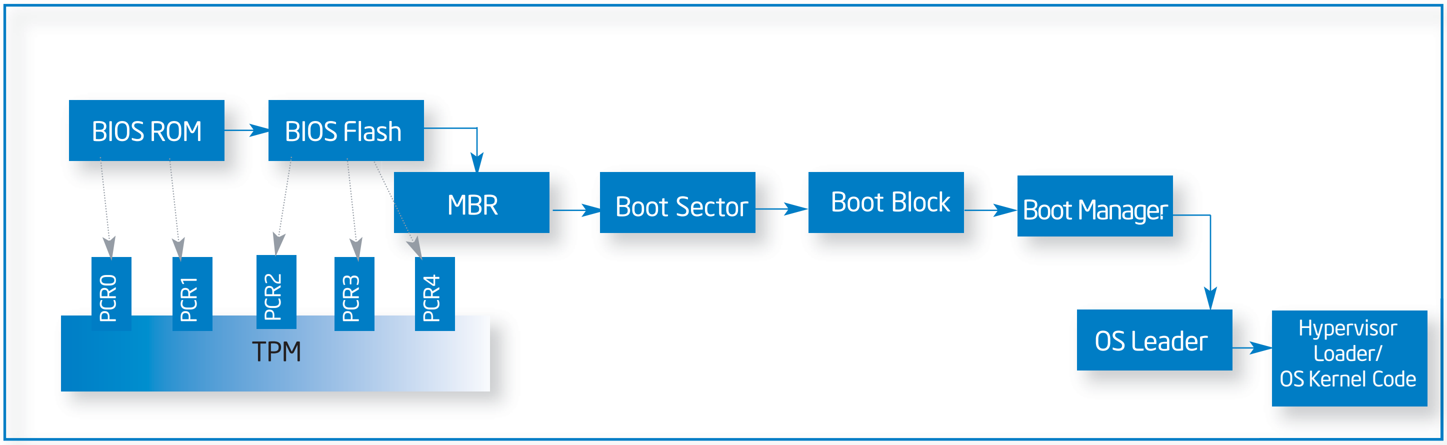

FIGURE 3. STATIC ROOT TRUST OF MEASUREMENT

SRTM produces excellent results and a great level of security—mainly against offline attacks. The problem is that multiple components must be verified in the chain of trust once TPM is initialized. Verifying the integrity of each component in the path of computer initialization (Figure 2) can become hard to manage due to the number of components involved. We need to measure every possible piece of code that might have been executed since the system boot; this imposes scalability issues.

To cope with this limitation, DRTM uses a different approach that is implemented in Intel® Trusted Execution Technology (Intel® TXT). Instead of validating every piece of code, a new instruction called SENTER has the capability to attest to the integrity of the hypervisor loader or operating system kernel code in a process known as measure launch.

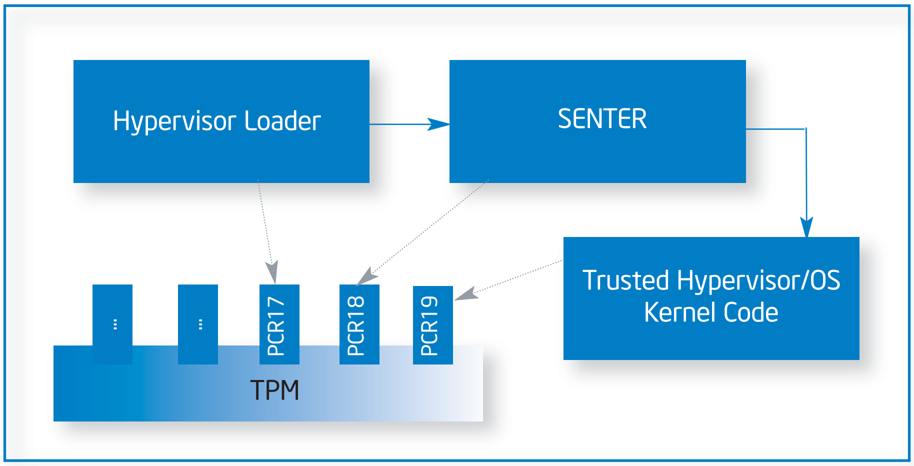

FIGURE 4. DYNAMIC ROOT TRUST OF MEASUREMENT

As Figure 3 shows, the hypervisor loader issues the GETSEC[SENTER] instruction, which essentially performs a soft processor reset and loads a signed, authenticated code module (ACM), which can only be executed if it has a valid digital signature. This module verifies system configurations and BIOS elements by comparing against the known good values protected by sensitive memory areas by using Intel® Virtualization Technology for Directed I/O (Intel® VT-d) and chipset-specific technologies such as Intel® Extended Page Tables (Intel® EPT). Then it verifies and launches the host system (a hypervisor core or an operating system kernel code), which configures low-level systems and protects itself using hardware-assisted paging (HAP).

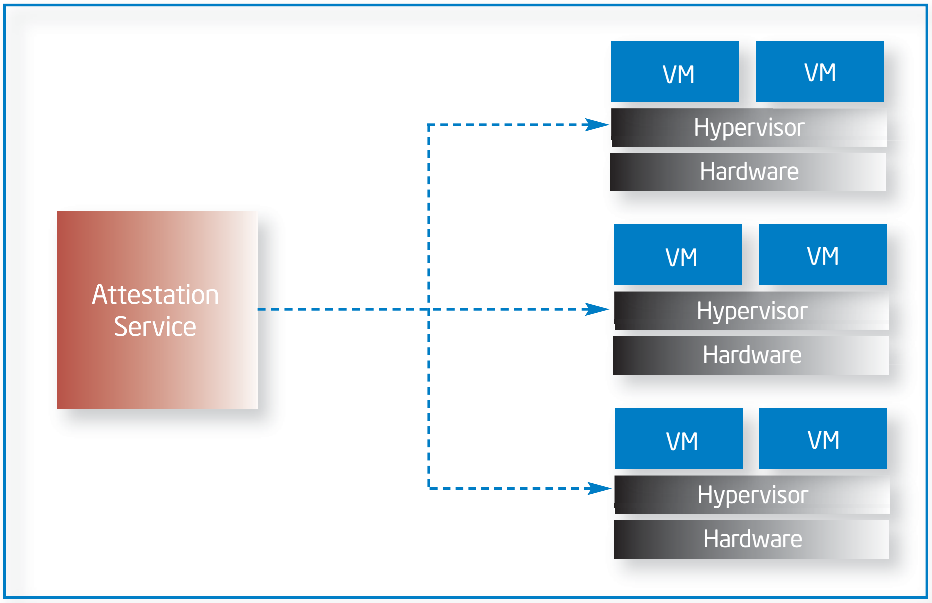

FIGURE 5. TRUSTED COMPUTE POOLS

The beauty of this method is that any instruction to be executed can be measured. That means virtual machines can also be measured before gaining privileged access to execute.

TRUSTED COMPUTING POOLS

Measuring hypervisor and virtual machine (VM) launches is the first step to creating a trusted platform, the brick of trusted computing pools. However, nowadays, with live migration enabled in a virtual environment, and where VMs can run on any host in a cluster, a compromised hypervisor can infect any VM running on top.

The key component in a trusted compute pool is the remote attestation entity, a server that records each hypervisor and VM state in a database. It takes several architecture implementation options to monitor and ensure compliance (they vary by different vendor). However, the basic principle should be the same.

The attestation service sends a challenge to the hypervisor based on the last successful measurement stored in the database, where only the correct PCR value, used as a key, can provide the right response. The secrets checked during remote attestation are known as the manifest.

Beside remote attestation, you can define the upper level of intelligence and rules such as trusted boundaries inside a cluster, and also tag-selected VMs. This lets you confine sensitive VMs to run only on trusted and measured hosts.

Trusted compute pools bring the capability to manage compliance by measuring launch, verifying integrity, providing control to ensure only a trustable hypervisor is run on the platform, protecting the server before virtualization software boot, and enabling compliance support with hardware underpinnings. On top of this trusted layer, you can apply policies to control migration across resource pools and allocate critical workloads to the trusted platform to improve the protection of your most critical applications.

A trusted compute pool is a powerful complement to runtime protection. It's true that on its own, it isn't enough to ensure complete protection. You also need good hiring practices, an established audit process, improved availability measures, and up-to-date security development methodologies such as SDLC and well-known security practices. When you combine them all, you will have a trusted compute cloud.

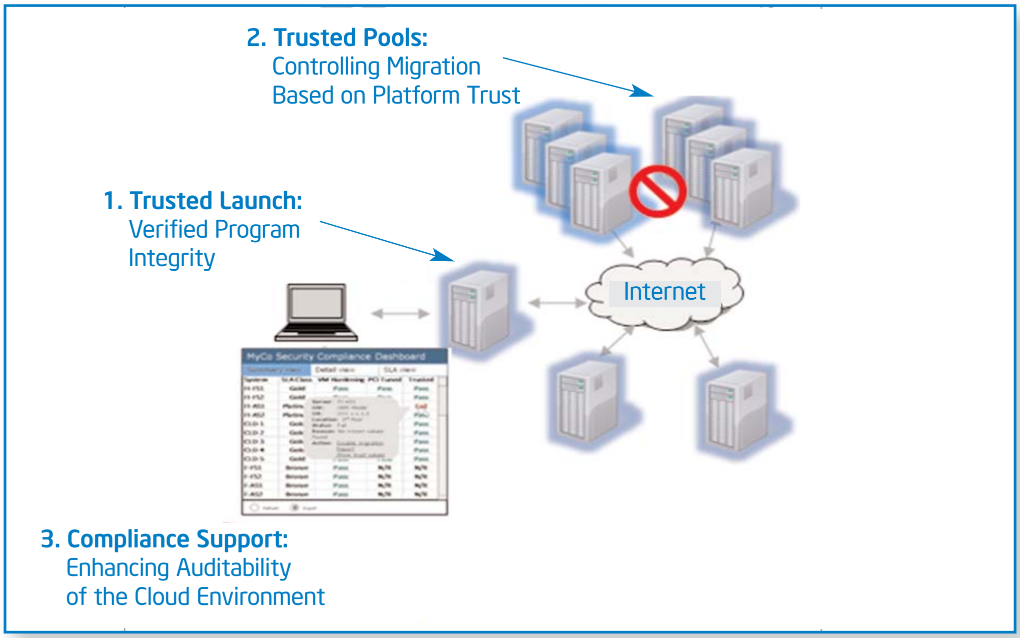

FIGURE 6. TRUSTED COMPUTE POOL USAGE MODEL

Bruno Domingues is a solution architect at Intel. Contact him at bruno.domingues@intel.com.

Back to Contents

[1] Jonathan Penn. "Security and the Cloud," (Forrester Research, October 2010).

CLOUD COMPUTING IN CHINA

A Growing Commitment

PEGGIE ZIH,
SOLUTION ARCHITECT, INTEL

Cloud computing is emerging worldwide as a new paradigm for the enterprise data center. The People's Republic of China is no exception, with a growing commitment to the cloud which it will demonstrate over the next two years with important proofs of concepts in cities throughout the country.

A STRATEGIC TECHNOLOGY

Cloud computing was designated a strategic technology by the Chinese State Council in its twelfth Five-Year Plan in 2010. The Ministry of Industry and Information Technology (MIIT) will be funding research and development for various cloud initiatives including infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS). Orient Securities LLC has predicted that by 2015, cloud computing in China will be a trillion-Yuan market segment, and we are starting to see plans unfolding in 2011.

CITY CLOUD PILOTS

Since the twelfth Five-Year Plan announcement, more than 20 cities in China have announced plans in cloud computing. In October 2010, the National Development and Reform Commission (NDRC) and MIIT jointly issued the Notice on Cloud Computing Service Innovation Development Pilot Demonstration Work, which named five pilot cities (Beijing, Shanghai, Shengzhen, Hangzhou, and Wuxi) for demonstrating cloud computing thought leadership (Figure 1). These pilot cities are focusing on topics such as e-commerce, e-government, small and medium enterprise (SME) enablement, smart city, education, and software. Success of these pilots will depend on a number of factors including data center efficiency, viability, and sustainability of cloud services offered and the ability to scale these pilots into large-scale deployments.

New large-scale data centers are already being built, with more planned to support city cloud projects. For example, Cong Qing is planning for the country's largest data processing center (i.e., an offshore cloud computing base with over one million servers involving billions of Yuan in investment). Industry ecosystem players have been brought together to design the infrastructure, builds the data center, and plan and carry out cloud service offerings.

CLOUD COMPUTING WAS DESIGNATED A STRATEGIC TECHNOLOGY BY THE CHINESE STATE COUNCIL.

SMALL AND MEDIUM ENTERPRISES

Green computing is almost a pre-requisite for these newly-built facilities, especially where we see concentrated efforts in the eastern part of China to meet increasing pressure to lower power consumption. Without efficiently-running data centers, a power shortage may become the by-product of cloud computing before the industry can enjoy the benefits of these city cloud pilots.

These cloud pilots are working to test the viability and sustainability of bringing dynamic and pay-as-you-go capacity to target users including the general public, small and medium enterprises, and government offices. This involves validating the appropriate cloud services and expanding them according to demand. For example, a pilot can start by providing administrative services for e-government and then extend into the healthcare sector, providing integrated service to the general public for handling emergencies. This involves being able to scale the cloud services and extend a single-purpose service to multiple integrated value-add services. It takes flexibility and a high level of resource management automation and security enforcement to develop multi-faceted services.

PUBLIC SAAS AND PAAS CLOUD SERVICES FOR SMALL AND MEDIUM ENTERPRISES

There has also been increasing focus on cloud services to small and medium enterprises. This segment has been growing at exponential rates in China. With more than 8 million small and medium enterprises in 2001, the number grew to 42 million in 2011, accounting for 99 percent of the country's total enterprises.

The pay-as-you-go, high-performance, and cost-effective services of cloud computing are particularly appealing to small and medium enterprises that can't afford large capital outlays in building IT infrastructure. SaaS (e.g., e-commerce, analytics services, enterprise resource planning, finance applications) and PaaS (e.g., platforms for development and testing, business intelligence, database platforms, wil will be particularly appealing to this segment.

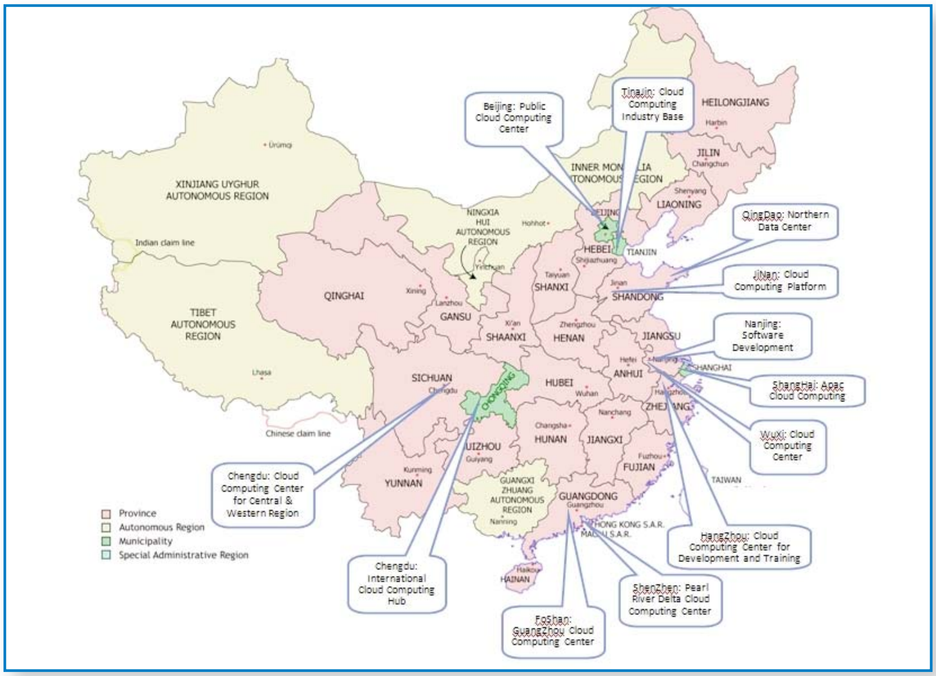
Beijing: Public Cloud Computing Center
TinaJin: Cloud Computing Industry Base
QingDao: Northern Data Center
JiNan: Cloud Computing Platform
Nanjing: Software Development
ShangHai: Apac Cloud Computing
WuXi: Cloud Computing Center
HangZhou: Cloud Computing Center for Development and Training
ShenZhen: Pearl River Delta Cloud Computing Center
FoShan: GuangZhou Cloud Computing Center
Chengdu: International Cloud Computing Hub
Chengdu: Cloud Computing Center for Central & Western Region
XINJIANG UYGHUR AUTONOMOUS REGION
Ürümqi
Indian claim line
TIBET AUTONOMOUS REGION
Lhasa
Chinese claim line
QINGHAI
Xining
GANSU
Lanzhou
NINGXIA HUI AUTONOMOUS REGION
Yinchuan
Hohhot
INNER MONGOLIA AUTONOMOUS REGION
HEILONGJIANG
Harbin
JILIN
Changchun
Shenyang
LIAONING
BEIJING
TIANJIN
HEBEI
Shijiazhuang
Taiyuan
SHANXI
Jinan
SHANDONG
Xi'an
SHAANXI
Zhengzhou
HENAN
JIANGSU
Nanjing
Hefei
ANHUI
SHANGHAI
SICHUAN
Chengdu
CHONGQING
HUBEI
Wuhan
Hangzhou
ZHEJIANG
Nanchang
Changsha
HUNAN
JIANGXI
GUIZHOU
Guiyang
Fuzhou
FUJIAN
TAIWAN
Kunming
YUNNAN
GUANGXI ZHUANG AUTONOMOUS REGION
Nanning
GUANGDONG
Guangzhou
HONG KONG S.A.R.
MACAU S.A.R.
Haikou
HAINAN
□ Province
□ Autonomous Region
□ Municipality
□ Special Administrative Region

FIGURE 1. CHINA'S PROGRESS IN CLOUD COMPUTING IN MAJOR CITIES

China has been unique in its mix of IT investment in hardware infrastructure and software. According to IDC, packaged software and IT services represented just 25.8 percent of the overall IT industry in China, far below the 69.9 percent in the U.S. and 61.3 percent globally in 2007. While China has accelerated its development of intellectual property in the IT industry, cloud computing services can augment this effort. In the small and medium enterprise market segment, the sweet spot will be in bundled IT services (consisting of infrastructure, software, applications, and management of such services). When built and managed properly, these services become economically sound usage models for small and medium enterprises to use to run their businesses, increasing spending in software and IT services. It will be very important for service providers to build services that can scale up and down and are manageable to help lower costs for small and medium enterprises, which will be able to use the new cloud service subscription model instead of having to build their own limited infrastructure and application platforms.

In this segment, potential cloud service providers fall into three key segments:

- TELECOM SERVICE PROVIDERS (e.g., China Telecom, China Unicom, and China Mobile) who have dominated the data center services and independent data center service providers (like 21ViaNet) fulfilling the IaaS and PaaS layer
- INTERNET PORTAL DATA CENTERS (IPDC) (e.g., Alibaba, Tencent, Baidu) who tend to focus on PaaS
- SOFTWARE SOLUTION PROVIDERS (e.g., Ufida, NeuSoft) who tend to focus on SaaS and PaaS

As China enters the cloud era, we expect to see stronger collaboration across tiers of service providers in enabling an integrated service model for small and medium enterprises. For example, e-commerce needs to go beyond offering and managing e-commerce portals to offer business intelligence, scaling services capacity dynamically to align with the fluctuating business needs of small and medium enterprises and helping them expand into the mobile market segment. All these will be value-adds for small and medium

TABLE 1. CHINESE CLOUD SERVICE PROVIDERS

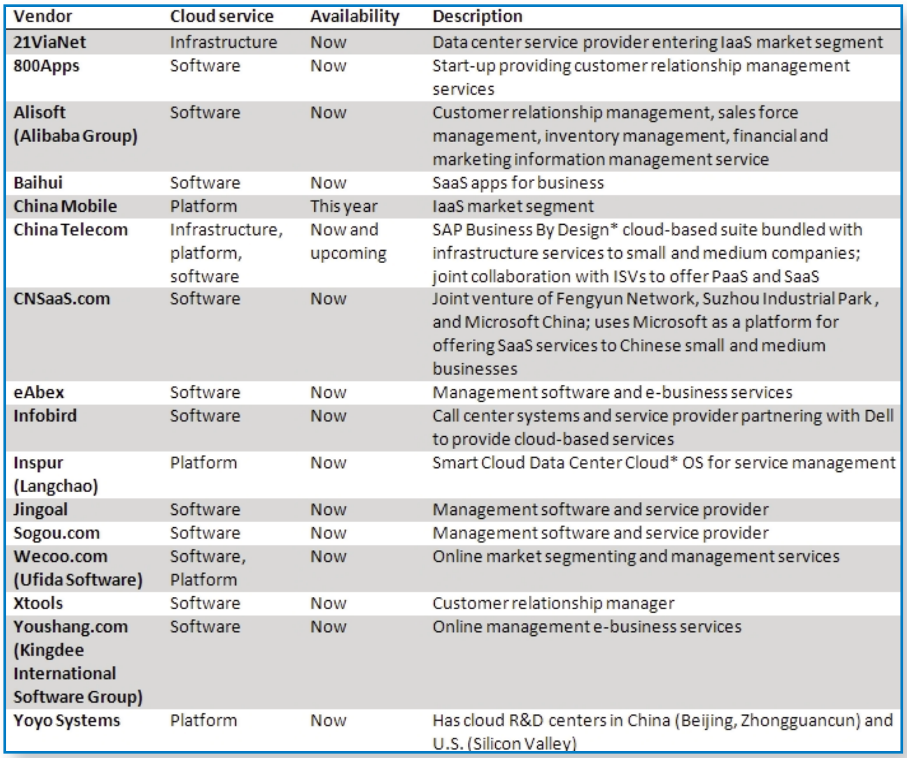

Vendor	Cloud service	Availability	Description
21ViaNet	Infrastructure	Now	Data center service provider entering IaaS market segment
800Apps	Software	Now	Start-up providing customer relationship management services
Alisoft (Alibaba Group)	Software	Now	Customer relationship management, sales force management, inventory management, financial and marketing information management service
Baihui	Software	Now	SaaS apps for business
China Mobile	Platform	This year	IaaS market segment
China Telecom	Infrastructure, platform, software	Now and upcoming	SAP Business By Design* cloud-based suite bundled with infrastructure services to small and medium companies; joint collaboration with ISVs to offer PaaS and SaaS
CNSaaS.com	Software	Now	Joint venture of Fengyun Network, Suzhou Industrial Park , and Microsoft China; uses Microsoft as a platform for offering SaaS services to Chinese small and medium businesses
eAbex	Software	Now	Management software and e-business services
Infobird	Software	Now	Call center systems and service provider partnering with Dell to provide cloud-based services
Inspur (Langchao)	Platform	Now	Smart Cloud Data Center Cloud* OS for service management
Jingoal	Software	Now	Management software and service provider
Sogou.com	Software	Now	Management software and service provider
Wecoo.com (Ufida Software)	Software, Platform	Now	Online market segmenting and management services
Xtools	Software	Now	Customer relationship manager
Youshang.com (Kingdee International Software Group)	Software	Now	Online management e-business services
Yoyo Systems	Platform	Now	Has cloud R&D centers in China (Beijing, Zhongguancun) and U.S. (Silicon Valley)

enterprises, enabling healthy and sustainable growth of cloud services. At the same time, this demands software solution providers to work closely with infrastructure solution providers to meet agility requirements without unnecessary infrastructure investments.

Apart from integration of services, multi-tenancy management will be a critical component, since the small and medium enterprise market segment tends to demand smaller but agile services. Service billing, charging, and auditing capabilities will help to establish trust in the pay-as-you-go services model.

According to IDC, revenue from SaaS in China will grow from USD 88 million in 2009 to USD 3.254 billion in 2014, with a compound annual growth rate of 28.9 percent. Small and medium enterprises will contribute a large portion to this growth, and service providers will be key in meeting the demand for services.

Table 1 shows partial list of Chinese cloud service providers.

LARGE ENTERPRISES ADOPT PRIVATE CLOUD INFRASTRUCTURE CAUTIOUSLY

On the other end of the spectrum, large enterprises will take their time in cloud adoption. This is the segment where private cloud will take off. For the last five years, large enterprises have focused on remodeling data centers and establishing centralized data center operations and management. The attraction of cloud computing to this segment lies in the cost efficiency of infrastructure and agility of service fulfillment, in particular in businesses where there is highly variable demand for services.

IaaS will be most prominent as large enterprises move into the private cloud. With the explosion in data growth, we should expect to see storage taking up a lot of the focus in the IaaS. Maturity in virtualization adoption will help drive the move to private cloud. Over the last couple of years, there have been active pilots and pockets of deployment on server virtualization as large enterprises consolidate and remodel their data centers. Over the next year, there will be broader scale of deployments. The best spot to extend the virtualization deployment into IaaS will be enabling a backup, recovery, and business continuity model. This also fits nicely with the continuing expansion and remodeling of data centers in this segment.

Apart from IaaS, PaaS is also relevant for private cloud as a standard platform for service development and provisioning. It can further reduce provisioning time for application services and/or platforms. To ensure success, it’s essential to assess infrastructure and application requirements and find the right fit with IaaS and PaaS. Cloud is definitely not a one-size-fits-all technology.

Large enterprises outside of China are also exploring public or even hybrid cloud services. However, it’s unlikely we will see public or hybrid cloud services over the next year or two because it will take time for public cloud services to prove their maturity and reliability.

There have been limited trials and deployments today. Large enterprises will take small, cautious steps in building private cloud infrastructures. Internet service providers, IT companies and, to a certain extent, financial companies are probably at the forefront of this segment and actively pursuing pilots in 2011. There will be more pilots over the next months before larger-scale deployments in 2012.

In conclusion, the focus on cloud computing since the twelfth Five-Year Plan announcement has accelerated significantly, mostly in the city cloud pilots and public cloud services targeting small and medium enterprises. The next two years will be crucial for the industry to use this government support and make progresses toward cloud computing in China. This requires careful assessment, evaluation, planning, and management of cloud initiatives. Strong collaboration among telecom, Internet, and software service providers will be needed to build the right cloud service model with intelligent management and cost economies to help China realize the benefits of cloud computing.

Peggie Zih is a solution architect for Intel's Enterprise Solution Sales group. Contact her at peggie.zih@intel.com.

Back to Contents

INTEL POWER MANAGEMENT

Building In Predictable Power Consumption

ENRIQUE CASTRO-LEON,
ENTERPRISE AND DATA CENTER ARCHITECT, INTEL
BERNARD GOLDEN, CEO, NAVICA
MIGUEL GOMEZ, TECHNOLOGY SPECIALIST,
TELEFÓNICA INVESTIGACIÓN Y DESARROLLO

It's worth noting that Intel® Xeon® processor 5500 series-based servers offer the capability to allow the operator to actually set power consumption and trade off the power consumption level against performance. This capability is valuable from two perspectives.

PREDICTABLE POWER CONSUMPTION

First, power capping brings predictable power consumption within the specified power capping range, and second, servers implementing power capping offer actual power readouts as a bonus: their power supplies are PMBus*-enabled and their historical power consumption can be retrieved through standard APIs.

With actual historical power data, it's possible to optimize the loading of power limited racks. Before, the most accurate estimation of power consumption had to be derived from derated nameplate data. The nameplate estimation for power consumption is a static measure that requires allowing a considerable safety margin. This conservative approach to power sizing leads to over-provisioning of power. This was acceptable when energy costs were a second-order consideration, but it's not today.

Integration with Intel® Data Center Manager enables the dialing power to be consumed by groups of over a thousand servers, allowing a power control authority of tens of thousands of watts in data centers. How does power capping work? The foundation for power management resides in CPU voltage and frequency scaling implemented by the Intel® Xeon® processor 5500 series architecture. More likely than not, CPUs represent the most energetic components in a server. If we can regulate the power consumed by the CPUs, we can have an appreciable effect on the power consumed by the server as a whole. Multiply this control over the thousands of servers in a data center. Through this mechanism, we can alter the power consumed in that data center in significant ways.

NESTING TECHNOLOGIES FOR SCALING POWER MANAGEMENT

Figure 1 shows a series of abstractions for a large deployment of servers in a data center, and for cloud storage appliances in particular. Chipsets are used to bind together the CPUs and the memory in a server. A server carries direct-attached (DASD) storage. Servers are organized in racks, and racks are organized in rows. The aggregation of rows encompasses all the servers in a data center.

Today our main lever for power control is throttling the power consumed by the CPUs up and down. In the near future, we can expect to see memory power control added to the mix. The basic mechanism of CPU voltage and frequency scaling allows moving the power consumption of a CPU up or down by a few tens of watts.

TAKING CONTROL OF POWER

Aggregating this capability over the tens of thousands of CPUs in a data center expands the range of attainable power control to tens or even hundreds of kilowatts.

There is also a potential synergistic effect captured by the power usage effectiveness (PUE) factor as defined by The Green Grid industry group. If servers use less power, the power allocated to the cooling equipment can be ratcheted down as well.

Power control for groups of servers is attained by composing power control capabilities of power control of each server. Likewise, power control for a server is attained by composing CPU power control (Figure 1).

Conceptually, power control for thousands of servers in a data center is implemented through a series of coordinated set of nested mechanisms.

MANAGING CPU POWER CONSUMPTION

The lowest level of power consumption is implemented through frequency and voltage scaling: laws of physics dictate that for a given architecture, power consumption is proportional to the CPU's frequency and to the square of the voltage used to power the CPU.

There are mechanisms built into the CPU architecture that allow a certain number of discrete combinations of voltage and frequency. Using the ACPI standard nomenclature, these discrete combinations are called P-states. The highest-performing state is identified as P0, and the lower power consumption states are identified as P1, P2, and so on.

An Intel Xeon processor 5500 series supports more than 10 states, with the actual number depending on the processor model. For the sake of an example, a CPU in P0 may have been assigned a voltage of 1.4 volts and 3.6 GHz, at which point it draws about 100 watts.

As the CPU transitions to lower power states, it may have a state P4 using 1.2 volts running at 2.8 GHz and consuming about 70 watts (Table 1). Actual CPUs may

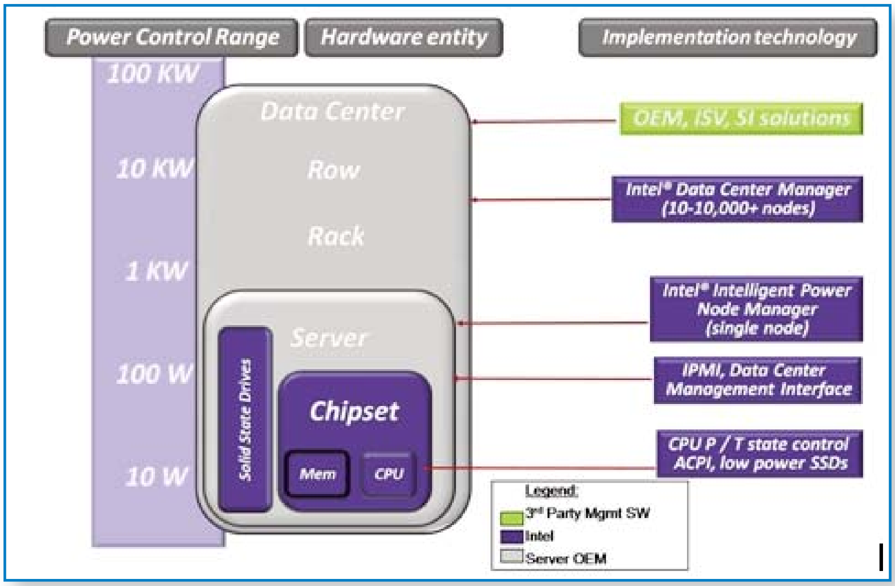
Power Control Range
Hardware entity
Implementation technology
100 KW
10 KW
1 KW
100 W
10 W
Data Center
Row
Rack
Server
Solid State Drives
Chipset
Mem
CPU
OEM, ISV, SI solutions
Intel® Data Center Manager (10-10,000+ nodes)
Intel® Intelligent Power Node Manager (single node)
IPMI, Data Center Management Interface
CPU P / T state control ACPI, low power SSDs
Legend:
🟩 3rd Party Mgmt SW
🟪 Intel
⬜ Server OEM

FIGURE 1. POWER CONTROL HIERARCHY

support more than 10 P-states. (Consult the CPU specifications for the actual values.)

TABLE 1. P-STATE EXAMPLE

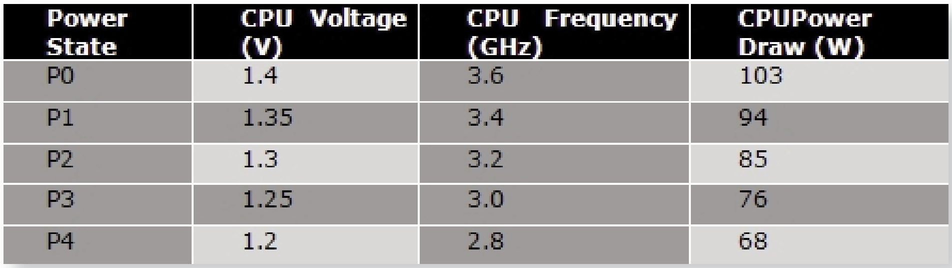

Power State	CPU Voltage (V)	CPU Frequency (GHz)	CPUPower Draw (W)
P0	1.4	3.6	103
P1	1.35	3.4	94
P2	1.3	3.2	85
P3	1.25	3.0	76
P4	1.2	2.8	68

The P-states by themselves can't control the power a server consumes. The CPU has no mechanism to measure the power it consumes. This capability is implemented by firmware running in the Intel Xeon processor 5500 series chipset. This firmware implements Intel® Intelligent Power Node Manager.

MANAGING SERVER POWER CONSUMPTION

If we want to measure the power a server consumes, looking only at CPU consumption doesn't provide the whole picture. That's why the power supplies in Intel Intelligent Power Node Manager-enabled servers are instrumented to provide actual power readings for the whole server through the PMBus* standard.

This process is shown in Figure 2. With these components in place, it's now possible to establish a classic control feedback loop where we compare a target power against the actual power indicated by the power supplies. The Intel Intelligent Power Node Manager code manipulates the P-states up or down until the desired target power is reached.

If the desired power lies between two P-states, the Intel Intelligent Power Node Manager code rapidly switches between the two states until the average power consumption meets the set power. This is an implementation of another classic control scheme, affectionately called bang-bang control (for obvious reasons).

Figure 3 shows more detail on the control loop. Here we can see that the Intel Intelligent Power Node Manager

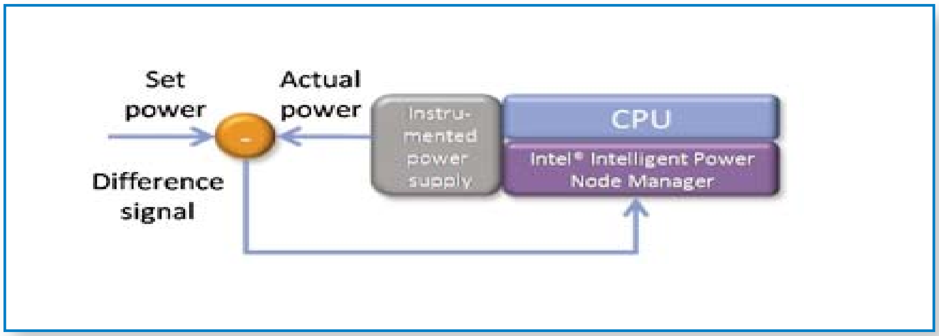
Set power
Actual power
-
Difference signal
Instru-mented power supply
CPU
Intel® Intelligent Power Node Manager

FIGURE 2. POWER CONTROL LOOP IN AN INTEL® XEON® PROCESSOR 5500 SERIES-BASED SERVER

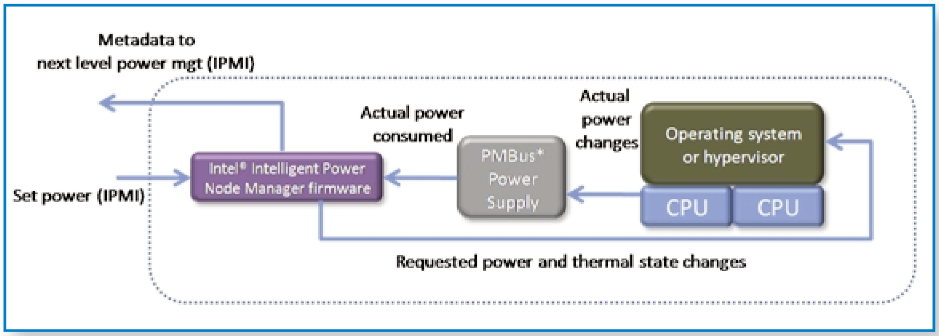
Metadata to next level power mgt (IPMI)
Actual power consumed
Actual power changes
Operating system or hypervisor
Set power (IPMI)
Intel® Intelligent Power Node Manager firmware
PMBus* Power Supply
CPU
CPU
Requested power and thermal state changes

FIGURE 3. EXPANDED VIEW OF THE INTEL® XEON® PROCESSOR 5500 SERIES POWER CONTROL LOOP

firmware implements the difference engine. Intel Intelligent Power Node manager directs the operating system or the hypervisor to change to a target P-state.

Note that the target P-state is set through an in-band mechanism (i.e., through the operating system). Intel Intelligent Power Node Manager does not set P-states directly. It sends requests to the operating system or hypervisor through an API. This is necessary to coordinate power policies with other power policies that the operating system or hypervisor might be carrying out.

Changes in processor P-state induce changes in the level of power consumption registered by the PMBus-enabled power supplies.

MANAGING POWER CONSUMPTION IN SERVER GROUPS

From a data center perspective, the ability to regulate power consumption of just a single server has a small impact and is not intrinsically useful. Harnessing the "power of the masses" represents a key capability. We need a way to control servers as a group, and just as we were able to obtain power supply readouts for one server, we need to monitor the power for the group of servers to allow meeting a global power target for that group of servers. This function is provided by the Intel Data Center Manager software development kit (Figure 4).

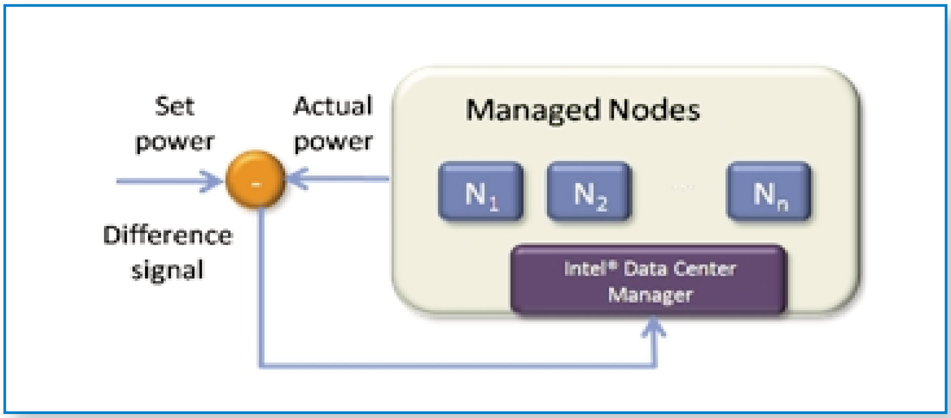

FIGURE 4. POWER CONTROL LOOP IN A GROUP MANAGED BY INTEL® DATA CENTER MANAGER

Note that Intel Data Center Manager implements a feedback control mechanism very similar to the mechanism that regulates power consumption for a single server, but on a much larger scale. Instead of watching one or two power supplies, Intel Data Center Manager oversees the power consumption of multiple servers or "nodes" whose number can range up to thousands. At this level, Intel Data Center Manager is in charge of implementing the difference engine.

Figure 5 shows an expanded view of the Intel Data Center Manager control loop as well as the relationship with the NM control loop underneath. No specific agents need to run in each node. Intel Data Center Manager communicates with the board management controller (BMC) in each node for setting power targets and for doing readouts of the actual power consumed. Intel Intelligent Power Node Manager firmware takes care of ensuring that the individual server meets the assigned power consumption target.

Intel Data Center Manager was purposely architected as an SDK to act as a building block for industry players to build more sophisticated and valuable capabilities for the benefit of data center operators. In one possible application (Figure 6), Intel Data Center Manager is integrated into a cloud storage appliance application. Some Intel Intelligent Power Node Manager-enabled servers come with inlet temperature sensors. This allows the application to monitor the inlet temperature of a group of servers and, if the temperature rises above a certain threshold, take a number of measures, from throttling back power consumption to reducing thermal stresses. At this level, the application code is in charge of the difference engine.

An application interfacing with Intel Data Center manager no longer "sees" individual server nodes; the application code's power policy engine designates a power consumption target to Intel Data Center Manager through the Intel Data Center Manager API. Intel Data Center Manager, in turn, breaks down the power target into power targets to the individual nodes under its command.

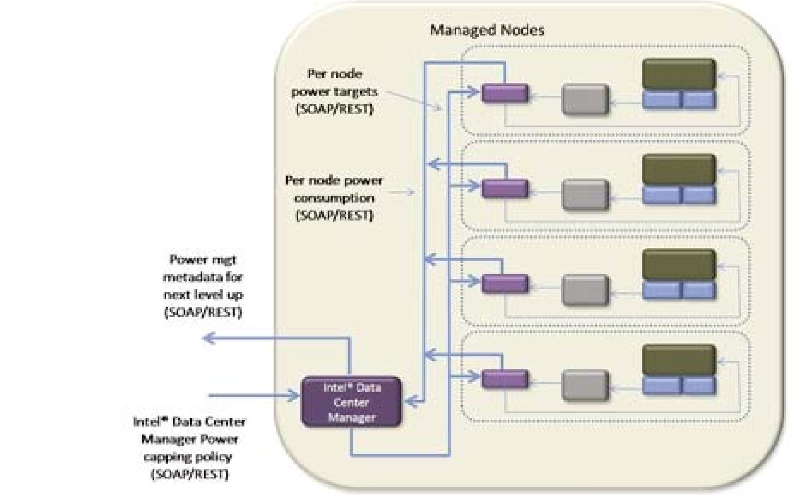
Managed Nodes
Per node power targets (SOAP/REST)
Per node power consumption (SOAP/REST)
Power mgt metadata for next level up (SOAP/REST)
Intel® Data Center Manager
Intel® Data Center Manager Power capping policy (SOAP/REST)

FIGURE 5. INTEL® DATA CENTER MANAGER POWER CONTROL LOOP

CLOUD STORAGE APPLICATION POWER MANAGEMENT

The application code also needs to mind the power consumed by the storage subsystem. Less sophisticated implementations may have a monitor-only capability; the application code reads out the power consumed by the storage subsystem by querying the intelligent PDU that feeds it, and then sets the server power so that the total consumption for the appliance matches the overall set power.

Figure 7 provides a more detailed view of this process, showing the appliance control algorithm implemented by the application's power policy engine. The constituent components are loosely coupled through Web services APIs.

One potential concern with three nested loops operating concurrently is the potential for oscillatory or unstable behaviors. This issue doesn't arise in practice since the time constants across the three levels are spread out. The time constant involved with the Intel Intelligent Power Node Manager loop is in the order of milliseconds, where the Intel Data Center Manager constant is in the order of a few seconds. Finally, the time constant associated with power management for an appliance is in the order of tens of seconds. Because of the time constant spreads, it's possible to optimize each loop individually to ensure that no under-damped, oscillatory behaviors occur.

You can find more information about power management and cloud computing in *Creating the Infrastructure for Cloud Computing: An Essential Handbook for IT Professionals* by Enrique Castro-Leon, Bernard Golden, and Miguel Gomez.

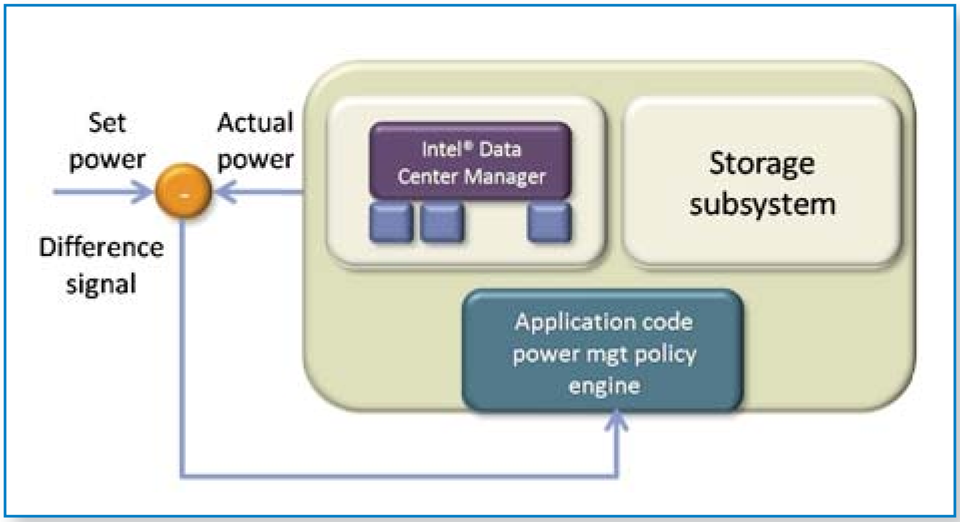
Set power
Actual power
-
Difference signal
Intel® Data Center Manager
Storage subsystem
Application code power mgt policy engine

FIGURE 6. STORAGE APPLIANCE CONTROL LOOP

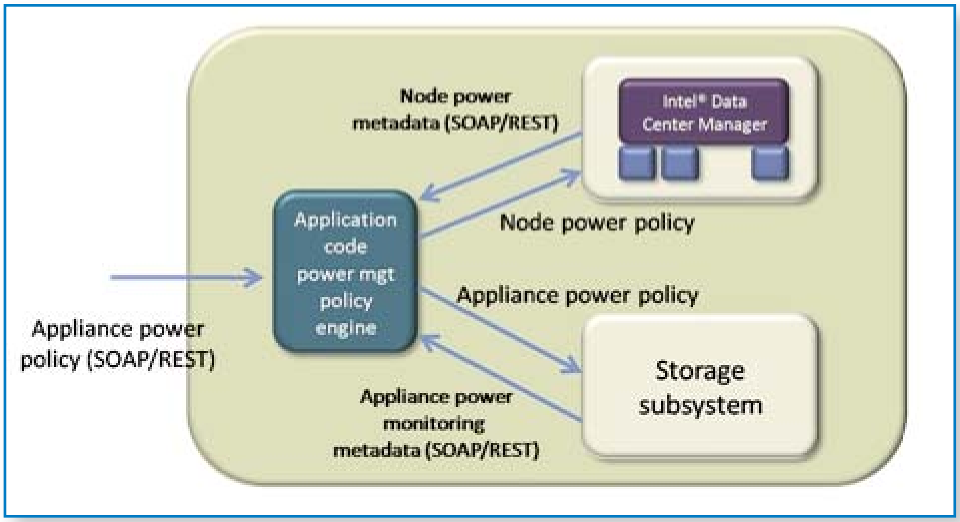
Node power metadata (SOAP/REST)
Intel® Data Center Manager
Application code power mgt policy engine
Node power policy
Appliance power policy
Appliance power policy (SOAP/REST)
Storage subsystem
Appliance power monitoring metadata (SOAP/REST)

FIGURE 7. EXPANDED VIEW OF APPLIANCE POWER CONTROL

To learn more, visit the Intel Press website at http://www.intel.com/intelpress/sum_virtn.htm or see our recommended reading list for similar topics: www.intel.com/technology/rr.

Back to Contents

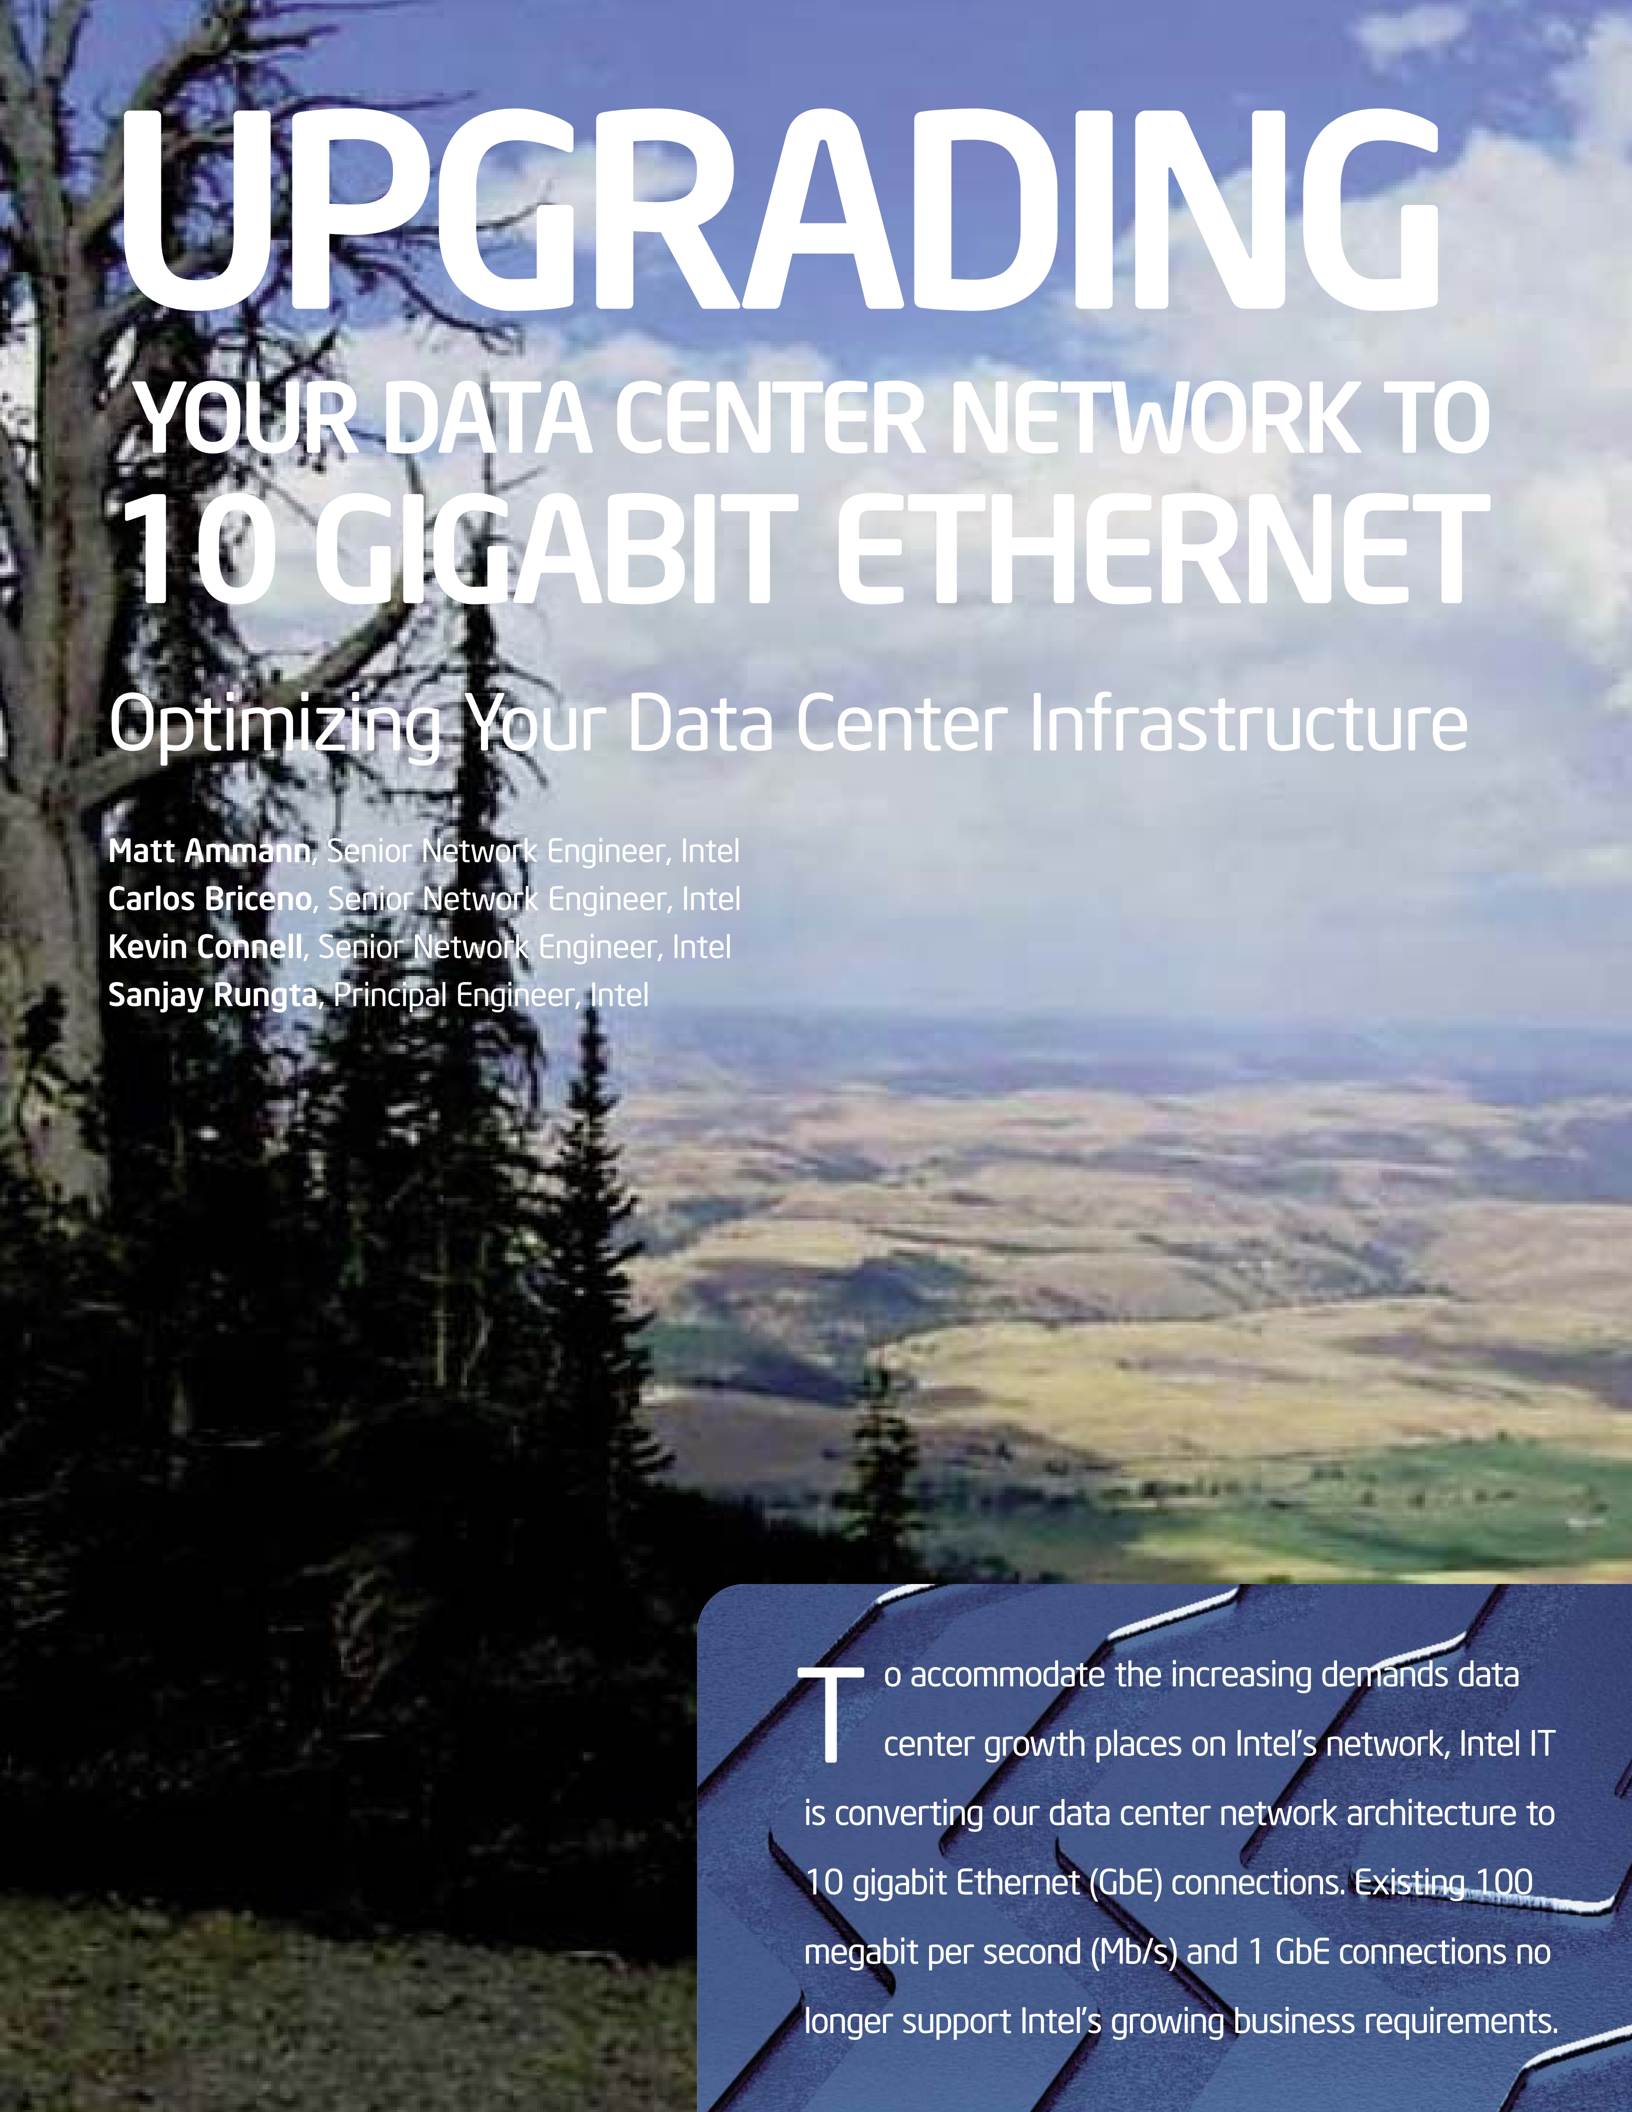

UPGRADING YOUR DATA CENTER NETWORK TO 10 GIGABIT ETHERNET

Optimizing Your Data Center Infrastructure

Matt Ammann, Senior Network Engineer, Intel
Carlos Briceno, Senior Network Engineer, Intel
Kevin Connell, Senior Network Engineer, Intel
Sanjay Rungta, Principal Engineer, Intel

To accommodate the increasing demands data center growth places on Intel's network, Intel IT is converting our data center network architecture to 10 gigabit Ethernet (GbE) connections. Existing 100 megabit per second (Mb/s) and 1 GbE connections no longer support Intel's growing business requirements.

OPTIMIZING DATA CENTER RESOURCES

Intel is engaged in a verticalization strategy that optimizes data center resources to meet specific business requirements in different computing areas. Data center trends in three of these areas drove our decision to upgrade, including server virtualization and consolidation in office and enterprise computing environments and rapid growth in design computing applications and their performance requirements. Also, we experience 40 percent per year growth in our Internet connection requirements.

While designing the new data center fabric, we tested several 10 GbE connection products and chose those that offered the highest-quality performance and reliability. We also balanced ideal design against cost considerations.

The new data center fabric design provides many benefits:

- REDUCED DATA CENTER COMPLEXITY. As virtualization increases, a 10 GbE network allows us to use fewer physical servers and switches.
- REDUCED TOTAL COST OF OWNERSHIP IN A VIRTUALIZED ENVIRONMENT. A 10 GbE fabric has the potential to reduce network cost in our virtualized environment by 18 to 25 percent, mostly due to simplifying the LAN and cable infrastructures. The new system also requires fewer data center space, power, and cooling resources.
- INCREASED THROUGHPUT. Faster connections and reduced network latency provide design engineers with faster workload completion times and improved productivity.
- INCREASED AGILITY. The network can easily adapt to meet changing business needs and will enable us to meet future requirements, such as additional storage capacity.

Upgrading our network architecture will optimize our data center infrastructure to respond faster to business needs while enhancing the services and value IT brings to the business.

BACKGROUND

Intel's 97 data centers are at the center of a massive worldwide computing environment, occupying almost 460,000 square feet and housing approximately 100,000 servers. These servers support five main computing application areas, often referred to as "DOMES": Design engineering, office, manufacturing, enterprise, and services.

Intel's rapidly growing business requirements place increasing demands on data center resources. Intel IT is engaged in a verticalization strategy that examines each application area and provides technology solutions that meet specific business needs. We are also developing an Office and Enterprise private cloud, and we see opportunities to expand cloud computing to support manufacturing.

These strategies, combined with the following significant trends in several computing application areas, compelled us to evaluate whether our existing 1 gigabit Ethernet (GbE) network infrastructure was sufficient to meet network infrastructure requirements:

- LARGE-SCALE VIRTUALIZATION in the office and enterprise computing domains.
- INCREASING COMPUTE DENSITY in the design computing domain.

In addition, high-performance Intel® processors and clustering technologies are rapidly increasing file server performance. This means that the network, not the file servers, is the limiting factor in supporting faster throughput. Our Internet connection requirements are growing 40 percent each year as well, which requires faster connectivity than is possible with a 1 GbE data center fabric.

SOLUTION

To meet these demands, we determined it was necessary to convert our data center network architecture from 100 megabit per second (Mb/s) and 1 GbE connections to 10 GbE connections. Our new 10 GbE data center fabric design will meet current needs while positioning us to incorporate new technology to meet future network requirements. For example, the 10 GbE network will simplify the virtualized host architecture used for office and enterprise computing, and will also reduce network component cost. For design computing, the 10 GbE network will reduce latency and allow faster application response times without the expense associated with alternative low-latency systems such as InfiniBand*. And, although storage area networks (SANs) in the office and enterprise application areas currently use a separate Fibre Channel (FC) fabric, we anticipate that as 10 GbE technology matures, we will be able to consolidate the SANs onto the 10 GbE network—thereby reducing network complexity even further.

SIMPLIFYING VIRTUALIZATION FOR OFFICE AND ENTERPRISE APPLICATIONS

Intel's data center strategy for office and enterprise relies on virtualization and consolidation to reduce data center cost and power consumption, while reducing time to provision servers. Our current consolidation level is 12:1, and we are targeting a 20:1 consolidation level or greater with newer dual-socket servers based on Intel® Xeon® processors.

As virtualization increases, so does the number of server connections. Using a 1 GbE network fabric, a single physical server on the LAN requires eight GbE LAN ports, as shown in Figure 1.

Presently, the SAN will continue to use FC connections and host bus adapters (HBAs).

In addition to simplifying physical infrastructure, a 10 GbE network fabric also has the potential to reduce the overall total cost of ownership (TCO) for LAN components per server by about 18.5 percent compared to the 1 GbE fabric, as shown in Table 1.

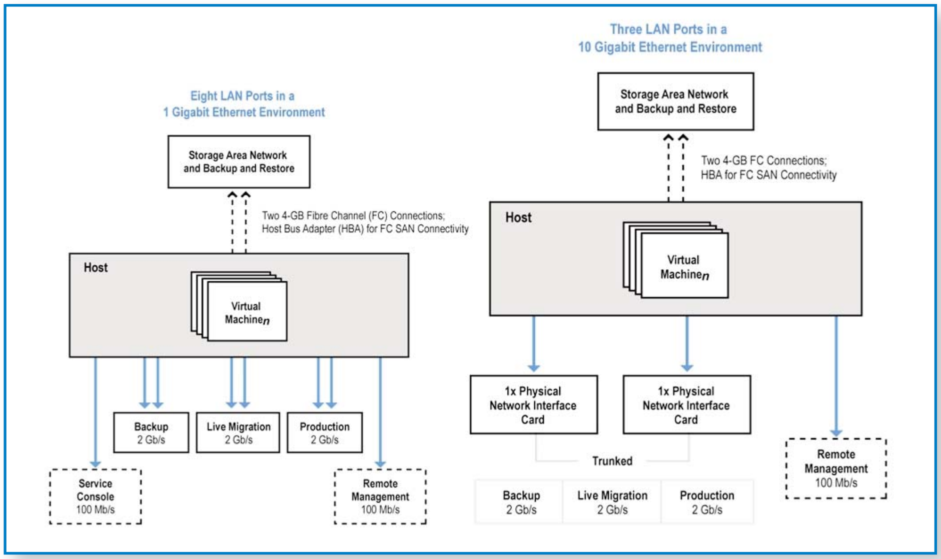
Eight LAN Ports in a
1 Gigabit Ethernet Environment
Storage Area Network and Backup and Restore
Two 4-GB Fibre Channel (FC) Connections; Host Bus Adapter (HBA) for FC SAN Connectivity
Host
Virtual Machine$_n$
Backup 2 Gb/s
Live Migration 2 Gb/s
Production 2 Gb/s
Service Console 100 Mb/s
Remote Management 100 Mb/s
Three LAN Ports in a
10 Gigabit Ethernet Environment
Storage Area Network and Backup and Restore
Two 4-GB FC Connections; HBA for FC SAN Connectivity
Host
Virtual Machine$_n$
1x Physical Network Interface Card
1x Physical Network Interface Card
Trunked
Backup 2 Gb/s
Live Migration 2 Gb/s
Production 2 Gb/s
Remote Management 100 Mb/s

FIGURE 1. IMPLEMENTING A 10 GIGABIT ETHERNET DATA CENTER FABRIC

INCREASING THROUGHPUT AND REDUCING LATENCY FOR DESIGN APPLICATIONS

For some specific silicon design workloads, we needed to build a small, very low-latency cluster between servers. Several parallel applications running on these servers typically carry very large packets. As shown in Table 2, we compared application response times, using several network fabric options. The 10 GbE network provided acceptable performance at an acceptable price point. For messages 16 megabytes (MB) and larger, the 10 GbE performance was about one-quarter of the 1 GbE response time, and was closer to the performance of InfiniBand, a more expensive low-latency switched fabric subsystem. In the table, "multi-hop" is defined as having to use more than one switch to get through the network.

CHOOSING THE RIGHT NETWORK COMPONENTS

We surveyed the market to find products that met our technical requirements at the right price point. We discovered that not all equipment and network cards are identical—performance can vary. With extensive testing, we found that we could reduce the cost for a 10 GbE port by 65 percent by selecting the right architecture and supplier. For example, we had to decide where to place the network switch—on top of the rack, which provides more flexibility but is more expensive, or in the center of the row. We chose to use center-of-the-row switches to reduce cost.

Higher transmission speed requirements have led to new cable tech-

TABLE 1. GIGABIT ETHERNET (GBE) COST SAVINGS

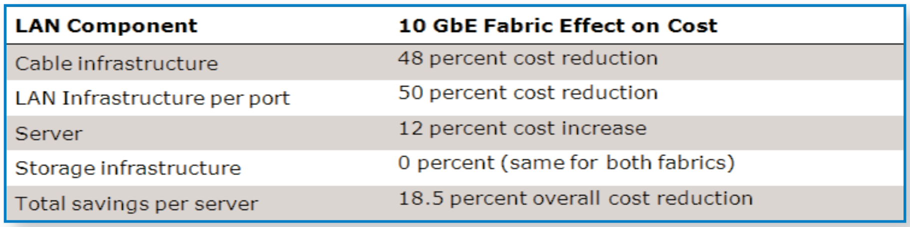

LAN Component	10 GbE Fabric Effect on Cost
Cable infrastructure	48 percent cost reduction
LAN Infrastructure per port	50 percent cost reduction
Server	12 percent cost increase
Storage infrastructure	0 percent (same for both fabrics)
Total savings per server	18.5 percent overall cost reduction

TABLE 2. APPLICATION RESPONSE TIMES FOR VARIOUS PACKET SIZES

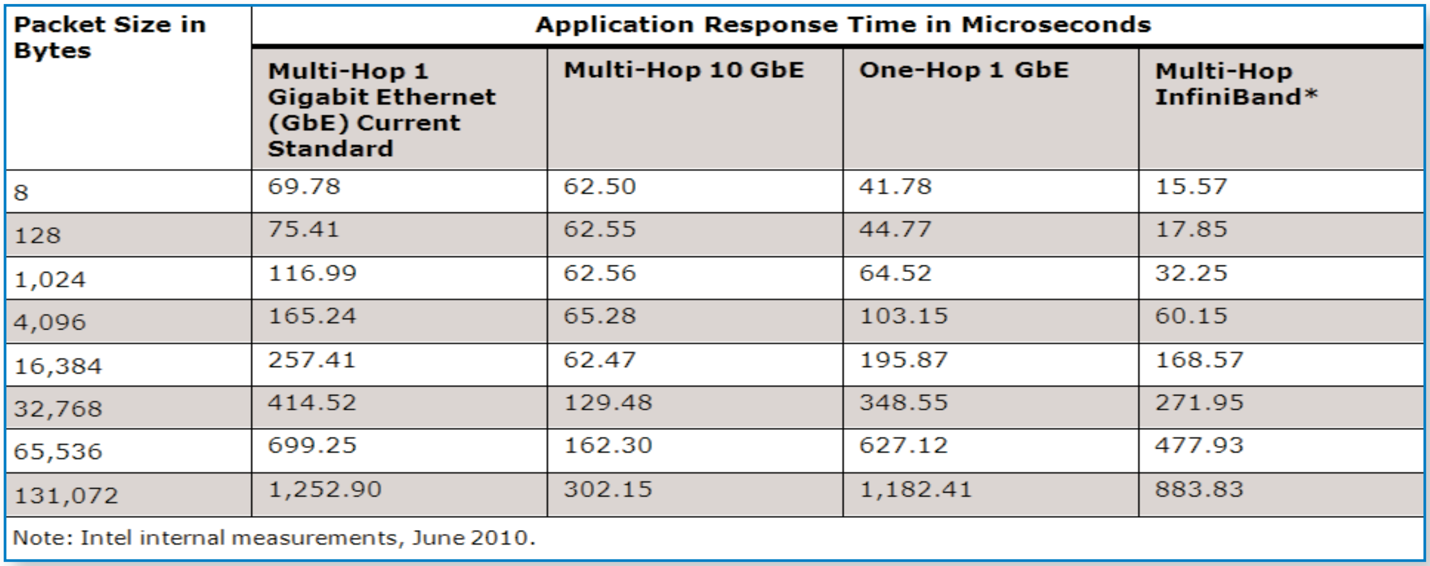

Packet Size in Bytes	Application Response Time in Microseconds			
	Multi-Hop 1 Gigabit Ethernet (GbE) Current Standard	Multi-Hop 10 GbE	One-Hop 1 GbE	Multi-Hop InfiniBand*
8	69.78	62.50	41.78	15.57
128	75.41	62.55	44.77	17.85
1,024	116.99	62.56	64.52	32.25
4,096	165.24	65.28	103.15	60.15
16,384	257.41	62.47	195.87	168.57
32,768	414.52	129.48	348.55	271.95
65,536	699.25	162.30	627.12	477.93
131,072	1,252.90	302.15	1,182.41	883.83

Note: Intel internal measurements, June 2010.

nologies, which we are deploying to optimize our implementation of a 10 GbE infrastructure:

- SMALL-FORM FACTOR PLUGGABLE (SFP+) DIRECT-ATTACH CABLES. These twinaxial cables support 10 GbE connections over short distances of up to 7 meters. Some suppliers are producing a cable with a transmission capability of up to 15 meters.
- CONNECTORIZED CABLING. We are using this technology to simplify cabling and reduce installation cost, because it is supported over SFP+ ports. One trunk cable that we use can support 10 GbE up to 90 meters and provides six individual connections. This reduces the amount of space required to support comparable densities by 66 percent. The trunks terminate on a variety of options, providing for a very flexible system. We also use a Multi-fiber Push-On (MPO) cable, which is a connectorized fiber technology comprised of multi-strand trunk bundles and cassettes. This technology can support 1 GbE and 10 GbE connections and can be upgraded easily to support 40 and 100 GbE parallel-optic connections by simply swapping a cassette. The current range for 10 GbE is 300 meters on Optical Multimode 3 (OM3) multi-mode fiber (MMF) and 10 kilometers on single-mode fiber (SMF).

To maximize the supportable distances for 10 GbE, and 40 GbE/100 GbE when it arrives, we changed Intel's fiber standard to reflect a minimum of OM3 MMF and OM4 where possible, and we try to use more energy-efficient SFP+ ports.

FUTURE PLANS FOR OFFICE AND ENTERPRISE I/O AND STORAGE CONSOLIDATION

Historically, Ethernet's bandwidth limitations have kept it from being the fabric of choice for some application areas, such as I/O, storage, and interprocess communication (IPC). Consequently, we have used other fabrics to meet high-bandwidth, low-latency, no-drop needs, such as FC. The advent of 10 GbE is enabling us to converge all our network needs onto a single, flexible infrastructure.

Several factors are responsible for increasing the I/O demand on our data centers. First, when more servers are added to the data center, it increases input/output operations per second (IOPS), which creates a proportional demand on the network. In addition, as each generation of processors becomes more complex, the amount of data associated with silicon design also increases significantly—again, increasing network demand. Finally, systems with up to 512 gigabytes (GB) of memory are becoming more common, and these systems also need a high-speed network to read, write, and back up large amounts of data.

Our upgrade to 10 GbE products will enable us to consolidate storage for Office and Enterprise applications while reducing our 10 GbE per-port cost. When I/O convergence on Ethernet becomes a reality, multiple traffic types, such as LAN, storage, and IPC, can be consolidated onto a single, easy-to-use network fabric. We have conducted multiple phases of testing, and in the near future, these 10 GbE ports will be carrying multiple traffic types.

SIMPLIFYING VIRTUALIZATION AND REDUCING TCO

A high-performance 10 GbE data center infrastructure can simplify the virtualization of office and enterprise applications and reduce per-server TCO. In addition, 10 GbE's lower network latency and increased throughput performance can support our design teams' high-density computing needs, improving design engineer productivity.

A HIGH-PERFORMANCE 10 GBE DATA CENTER INFRASTRUCTURE CAN SIMPLIFY THE VIRTUALIZATION OF OFFICE AND ENTERPRISE APPLICATIONS AND REDUCE PER-SERVER TCO.

Our analysis shows that for a virtualized environment, a 10 GbE infrastructure can reduce our network TCO by as much as 18 to 25 percent. And, for design applications, where low latency is required, 10 GbE can play a crucial role without requiring expensive low-latency technology. The new fabric will also reduce data center complexity and increase our network's agility to meet Intel's growing data center needs.

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

Back to Contents

Available Soon from Intel Press

To learn more about this cloud computing title and other books from Intel topic and technology experts, visit www.intel.com/intelpress

IT Best Practices Series

Creating the Infrastructure for Cloud Computing

An Essential Handbook for IT Professionals

Enrique Castro-Leon, Bernard Golden, Miguel Gomez, Raghu Yeluri and Charles G. Sheridan

Intel PRESS

Intel® Hyperthreading Technologyequires an Intel® HT Technology enabled system, check with your PC manufacturer. Performance will vary depending on the specific hardware and software used. Not available on Intel® Core™ i5-750. For more information including details on which processors support HT Technology, visit http://www.intel.com/info/hyperthreading.

No system can provide absolute security under all conditions. Intel® Anti-Theft Technology requires an enabled chipset, BIOS, firmware and software and a subscription with a capable Service Provider. Consult your system manufacturer and Service Provider for availability and functionality. Intel assumes no liability for lost or stolen data and/or systems or any other damages resulting thereof. For more information, visit http://www.intel.com/go/anti-theft .

No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology (Intel® TXT) requires a computer system 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 http://www.intel.com/technology/security.

Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM). Functionality, performance or other benefits will vary depending on hardware and software configurations. Software applications may not be compatible with all operating systems. Consult your PC manufacturer. For more information, visit http://www.intel.com/go/virtualization.

Intel® Turbo Boost Technology capability requires a system with Intel Turbo Boost Technology capability. Consult your PC manufacturer. Performance varies depending on hardware, software and system configuration. For more information, visit http://www.intel.com/technology/turboboost

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked “reserved” or “undefined.” Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

Copyright © 2011 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others. Printed in USA   SS/PP/0811 Please Recycle