

Cloud Computing: Considerations and Next Steps

How to Navigate the Ongoing Shift to Cloud Computing

White Paper
Cloud Computing
Key Considerations



Executive Summary and Purpose of This Paper

Cloud computing is an important transition and a paradigm shift in IT services delivery – one that promises large gains in efficiency and flexibility at a time when demands on data centers are growing exponentially. The tools, building blocks, and best practices for cloud computing are evolving, and the challenges to deploying cloud solutions need to be considered.

The technology and industry leadership that Intel brings to this environment is broader and deeper than most realize. Intel's vision for cloud computing is for solutions that are federated, automated, and client-aware. Moving the industry toward that promise will require a focus on four industry-wide principles of cloud computing that will make this vision a reality: efficiency, simplification, security, and open standards. This white paper defines Intel's cloud computing vision and then, for each of the four principles, explains both Intel's activities and Intel's recommendations for evolving cloud infrastructure over the next several years to be more efficient, secure, simplified and built on open standards.

Introduction: Demands, Challenges, Solutions

A Paradigm Shift

Rather than a *revolution*, cloud computing is an important transition, a *paradigm shift* in IT services delivery – one that has broad impact and can present significant challenges. Cloud computing represents a transformation in the design, development, and deployment of next-generation technologies – technologies based on flexible, pay-as-you-go business models that will alter the future of computing from Mobile Platforms to the Data Center.

The impetus behind cloud computing is the ever-increasing demands placed on data centers that are near capacity and resource-constrained. In response to these challenges, cloud computing is evolving in the forms of both **public clouds** (deployed by Internet companies, Telcos, hosting service providers and others) and **private** or **enterprise clouds** (deployed by enterprises deployed behind a firewall for an organization's internal use).

Cloud computing is being driven by explosive growth of Internet data and traffic as the Internet matures and Internet-based services proliferate. By 2015, over 2.5 billion people with more than 10 billion devices will access the Internet – over 2x today's demand – requiring a cloud infrastructure with one billion virtual servers.¹ The monumental requirements associated with the data center build-outs needed to satisfy this growing demand can only be met with the increased efficiency, performance, and flexibility of cloud architectures.

Table of Contents

Introduction: Demands, Challenges, Solutions	1
A Paradigm Shift.....	1
Intel's Role in Achieving the Cloud Computing Vision	3
The Three Elements of Intel's Cloud Computing Vision	3
Efficiency	5
IT Challenges	5
Intel Advancements to Improve Efficiency	5
Key Considerations for Efficiency.....	7
Simplification	7
IT Challenges	7
Intel Advancements to Simplify Data Center Infrastructure	7
Key Considerations to Simplify your IT Infrastructure	8
Security	9
IT Challenges	9
Intel Advancements in Security	9
Key Considerations for Security	9
Open Standards	10
IT Challenges	10
Intel and Open Standards	10
Key Considerations about Driving Open Standards	11
Conclusion	11
For More Information	11

Within enterprises, cloud computing is being driven by the expanding business demands on enterprise IT. More and more data centers find themselves facing real limits, whether based on lack of power, lack of room, lack of server capacity, or lack of network bandwidth. Expanding traditional infrastructures to meet these challenges quickly uncovers multiple inherent inflexibilities. The resulting complexity breeds cost, deployment risk, and operational risk. For the vast majority of enterprise IT organizations, these challenges are serious today and will be larger tomorrow. Public cloud service providers are experiencing significant growth and have a need to continue to evolve and advance their IT infrastructure while addressing the challenges related to cloud computing services build out.




Figure 1. Virtualization: Evolving Towards the Enterprise Cloud

Cloud computing is a step beyond data center virtualization (see Figure 1). Initially, virtualization technologies allowed data centers to consolidate server infrastructure to save cost. Next, flexible resource management technologies added the ability to more dynamically allocate data center resources. This further reduced costs but also increased data center flexibility and performance, ushering in a new era of technology development and deployment. Software vendors have begun to design robust management features and technology optimizations for enterprise and public clouds based upon virtualization. Hardware vendors have extended their management tools and reliability features to include increased flexibility. The era of cloud computing can be seen as the next natural step, where significant automation and scalability become possible. Cloud computing offers a path to the needed reductions in cost and resource drain. Beyond technology, cloud computing creates new organizational requirements within IT departments, as separate teams responsible for networking, compute, and storage must now work together to manage and deploy a common infrastructure. When fully realized, cloud computing infrastructures can provide competitively significant IT agility, flexibility, and adaptability through systems that are efficient, flexible, and secure.

Though cloud computing can be viewed as an evolutionary step, it is a fundamental shift and presents challenges:

- Maintaining the stability of mission-critical applications as you transition into cloud environments is paramount.
- Intellectual property protection, data security, and privacy all require additional attention and new tools if shared resources in a public cloud are to be used.
- The automation and flexibility of resource pools will be imperfect while cloud computing tools evolve.
- Selection of solutions that provide for flexibility and interoperability.

Intel's Role in Achieving the Cloud Computing Vision

To meet these challenges and reduce the risks that IT architects and managers face, Intel is providing leadership and advancing efforts in each of four principles – efficiency, simplification, security, and open standards. Our efforts extend across the cloud computing ecosystem, from silicon to platforms to software architecture to data center design. This leadership includes technology innovation, standards leadership, and first-hand experience – combining vision (where cloud is going) with know-how (how practical cloud infrastructures are built and deployed). This cloud computing leadership is immediately useful to IT professionals as they consider and implement their own cloud computing infrastructures.

While Intel is widely seen as a semiconductor supplier, our role in the cloud computing architecture ecosystem actually goes beyond supplying silicon building blocks. We engage extensively with enterprise decision makers, software vendors, system integrators, computer manufacturers, and value-added resellers to work on key cloud computing challenges. By bringing engineering resources, a very large network of global relationships, significant market development investments, and specific-industry expertise, Intel also acts as a major catalyst working with the industry to develop effective, open standards.

In addition, as a typical large enterprise managing about 100,000 servers globally, Intel is actively developing a strategy for and deploying cloud computing-based solutions within our own IT department to support our corporate business goals. The first-hand experiences we've gained internally help to inform our engagements with the rest of the ecosystem.

The Three Elements of Intel's Cloud Computing Vision

Cloud computing technology is maturing at a fast pace and many cloud services and vendors are entering the market to enable the development of private clouds for enterprise IT. In several cases, public cloud providers are expanding their services to support enterprises. In Intel's numerous conversations with vendors, analysts and customers, we've identified key themes that emerge as critical to what customers want from cloud computing infrastructures and solutions. Intel's vision for cloud computing over the next five years centers on three themes that are essential to help overcome key challenges and realize the full potential and value of cloud computing solutions: **federated**, **automated**, and **client-aware**.

Federated means communications, data, and services can move easily within and across cloud computing infrastructures. To accomplish truly federated systems, smooth interoperability across many platforms and solutions must be a reality. Today, the industry is just reaching the point that enterprises can move or migrate workloads within and between their own data centers. Data center operators are far from being able to have data and services seamlessly and securely scale beyond their borders to span public and private clouds when desired.

Automated means that cloud computing services and resources can be specified, located, and securely provisioned with very little or zero human interaction. Today, the industry faces many gaps in automation. Virtual machines are generally very statically provisioned vs. automatically responding to user needs. Data center management remains very manual today – patching of servers doesn't scale reliably. Intel's cloud computing vision calls for automation that dynamically allocates resources to agreed-upon service levels and optimizes the data center for maximum resource utilization and power efficiency.

Definition of Cloud Computing

Cloud computing is an evolution in which IT consumption and delivery are made available in a self-service fashion via the Internet or internal network, with a flexible pay-as-you-go business model and requires a highly efficient and scalable architecture. In a cloud computing architecture, services and data reside in shared, dynamically scalable resource pools, often virtualized. Those services and data are accessible by any authenticated device over the Internet. The key attributes that distinguish cloud computing from conventional computing are:

- Compute and storage functions are abstracted and offered as services
- Services are built on a massively scalable infrastructure
- Services are delivered on demand through dynamic, flexibly configurable resources
- Services are easily purchased and billed by consumption
- Resources are shared among multiple users (multi-tenancy)
- Services are accessible over the Internet or internal network by any device

A **public cloud** is a cloud architecture deployed by Internet companies, hosting service providers, and Telco service providers to deliver services to a broad range of consumers and/or businesses via the Internet. A **private cloud** is a cloud architecture deployed behind a firewall for an organization's internal use.




Figure 2. Intel's Cloud Computing Vision

Client-aware means that cloud computing solutions adapt seamlessly to the end user's device and use model regardless of the type of client system they are using. Today, there are certain frameworks that allow for some level of data center intelligence and scaling to support the client being served, but they are neither consistently applied nor ubiquitous. There are many end user examples (ranging from web frameworks to gaming to virtual desktop infrastructure) in which software vendors are defining their own approaches to service delivery – but there is no consistency across providers. Many of today's Internet services support the lowest common denominator device even if the user is accessing the service with a powerful desktop computer. Conversely, other services are difficult to use on a handheld device because they were written for a PC. Intel's cloud computing vision calls for the data center and service provider to enable secure access and optimal experience across a range of devices, by making the cloud knowledgeable about client device attributes. These attributes include the device's capabilities, location, policies, and connectivity. At the same time, client device capabilities can affect the overall performance of cloud solutions: intelligent performance on the client device can deliver better end user experiences through local computing power; security capabilities on the client device can ensure security policies are applied at the device; and pervasive communication (using any choice of connectivity, such as LAN, WAN, Wi-Fi, or personal area network) enables work-from-anywhere flexibility.

Making these three essential characteristics of more mature cloud computing real will not be trivial. It will require cooperative development and specific focus by many providers and customers across the IT landscape. We believe that to move towards this vision of cloud computing, individual organizations and the IT industry as a whole need to focus on **four principles**:

- **Efficiency:** While the need for computing throughput increases exponentially, resources are limited. These resources include space, power, cooling capacity, qualified IT professionals, dollars for infrastructure, and dollars for operations. Doing more with existing or available resources will require increased efficiency from infrastructure and processes.

- **Simplification:** Generally, the growth of a system inherently increases complexity, and this is certainly true of IT infrastructures. Multiple architectures complicate management. Increased server utilization raises network bandwidth requirements. And systems from different vendors typically present integration complications. For cloud computing environments to deliver on their promise, simplification must underlie cloud architectures and practices.
- **Security:** Both business risk and compliance requirements make data security paramount. In an environment with abundant traditional security issues, cloud computing creates new challenges because it moves data in new ways, often outside of traditional physical boundaries. The successful implementation of cloud computing requires new security models to meet new challenges.
- **Open standards:** When multiple providers (of solutions, hardware, software, integration, or processes) act *independently*, poor interoperability and lack of flexibility are the natural results—and are in direct contradiction to the main promises of cloud computing. The evolution of cloud computing requires open standards that are carefully constructed and create greater levels of interoperability.

In the following pages, for each of these principles, we discuss the drivers behind the issue, what Intel is doing in the cloud computing ecosystem, and what actions data center architects and managers should take now to best align their own cloud computing strategies with current and future developments. As a key player in the IT ecosystem, Intel is heavily focused on driving industry-wide developments to continue to advance these four principles to enable the evolution of cloud environments over the next five years and beyond.

Efficiency

IT Challenges

In many data centers today, the cost of power (electricity) is only second to the actual cost of the equipment investments. Today, cloud data centers consume 1-2 percent of the world's energy.² If left unchecked, they will rapidly consume an ever-increasing amount of energy, deplete natural resources and drive increasing costs for power consumption. Moreover, as business demands drive the build-out of data centers, architects are challenged in a "traditional" sense to get more out of existing resources, i.e., increase the efficiency of dollar and employee inputs. But they are also bumping up against newer challenges in the form of practical limits to physical resources: running out of available space, power, and cooling capacity. A number of cloud computing developments are bringing relief now – specifically higher server utilization, better throughput for network and storage traffic, and storage optimized by data type and need. Moreover, IT departments are evaluating how to run their data centers at higher temperatures to reduce cooling costs; improvements in compute density and manageability; deployment of containers; and other advancements to improve efficiency.

Intel Advancements to Improve Efficiency

Together with industry partners, Intel develops leading products, building blocks, and software tools that significantly improve data center performance and efficiency. These efforts extend across the breadth of the data center and involve multiple areas of focus, including silicon power optimization, advanced data center power management, improved virtualization capabilities, data center research and optimization – all focused on accomplishing greater levels of data center efficiency. In addition, Intel continues to work closely with a wide range of systems vendors to help enable a broad range of power-optimized systems. These advances are outlined on the following pages.

Silicon advancements that enable increased server performance per watt:

By using server platforms optimized for performance per watt (as well as for density, power management, and virtualization), data centers lay a foundation for efficiency, maximized workloads, and flexibility.

Intel has a long history of delivering energy-efficient performance that follows Moore's Law. Each generation of Intel® Xeon® processor-based server has historically delivered up to an average of 40 percent increase in power efficiency (measured in performance per watt).³ Server platforms based on the Intel® Xeon® processor 5600 series are built on Intel's industry-leading, extremely energy-efficient 32nm process technology. The Intel Xeon processor 5600 series delivers up to 60 percent performance boost and up to 30 percent lower power at the same performance as the previous generation.⁴ In addition, they include multiple technologies that increase performance and performance per watt:

- **Intel® Hyper-Threading Technology¹:** Delivers thread-level parallelism on each processor, resulting in more efficient use of processor resources (higher processing throughput) and improved performance on multi-threaded software.
- **Intel® Turbo Boost Technology⁵:** When a server's workload demands additional performance, the processor frequency is dynamically increased by 133 MHz on short and regular intervals, until an upper limit is met. Conversely, when the limit is reached or exceeded, the processor frequency will automatically decrease by 133 MHz until the processor is again operating within its limits. This increases the performance of both multi-threaded and single-threaded workloads, providing performance when and where it's needed.
- **Intel® Intelligent Power Technology:** Integrated power gates allow individual idling cores within a CPU to drop to near-zero power, independent of the other cores. This can reduce idle power consumption. Integrated power gates can also be used to automatically manage processor power, or can be controlled manually.
- **Intel® QuickPath Technology:** Provides point-to-point high-speed links to distributed shared memory and compliments the parallel processing performance of Intel® Microarchitecture codenamed Nehalem. Each processor core features an integrated memory controller and high-speed interconnect, linking processors and other components to enable dynamically scalable interconnect bandwidth, and excellent memory performance.

Learn more about the Intel Xeon processor 5600 series at www.intel.com/xeon

Intel Software Optimization

Optimizing performance, reducing power consumption, and enabling software code that takes advantage of future platforms is important to help realize the benefits of cloud computing. Intel offers several software tools to optimize workloads and code for Internet services, such as Intel® Thread Checker, Intel® Parallel Studio, and others. To learn more about these software products, visit software.intel.com/en-us/intel-sdp-home.

Intel also provides customers with software consulting and training, helping customers analyze and eliminate bottlenecks, optimizing their code to take full advantage of Intel platforms and technologies, resulting in performance gains and power-efficiency improvements.

Advanced power management: Typically, server power is over-allocated, as lack of visibility into real and varying power loads means data center managers must over-allocate power to account for worst-case loads. Similarly, rack space is under-populated to account for worst-case thermal loads, exacerbating space constraints. Intel has developed a number of technologies that enable *dynamic* power management to enable more efficient use of power and space as well as more flexibility within demanding operational scenarios. These solutions allow IT managers to manage power at the system, rack and data center level.

- **Server platform instrumentation:** The foundation of intelligent power management in the data center is capable instrumentation at the server level. The Intel Xeon processor 5600 series contains instrumentation for temperature (CPU and fan inlet), idle power, average power, power states, and memory states. This data is used by the higher-level tools discussed in the following section to monitor and manage power at the server, rack, and data center level.
- **Intel® Intelligent Power Node Manager:** The under-population of server racks results in decreased performance per rack footprint and often paying for power capacity that will never be used. Intel® Intelligent Power Node Manager is a system-level technology, implemented through a combination of platform power supplies and Intel's manageability engine (embedded in the chipset and firmware) that delivers power reporting and power capping functionality for individual servers. The application runs as a closed-loop algorithm with periodic power consumption analysis and real-time synchronization of OS power and thermal states. It uses multiple power-throttling mechanisms, including processor package power limiting, memory power limiting, and dynamic CPU core allocation. This allows data center managers to dynamically throttle system and rack power based on expected workloads. This in turn allows the installation of more servers per rack, better utilizing available space and power

Oracle Uses Intel® Intelligent Power Node Manager to Increase Rack Density

Oracle was intent on increasing the computing density of its data centers so it could serve more customers at a lower cost and be in a strong competitive position when the economy improves. But the gating factor was power consumption. The only way to reduce the amount of power each server used would be to pack more processing capability into the same space. Oracle is using the Intel® Xeon® processor 5500 series with Intel® Intelligent Power Node Manager to match the power to the application. Intel Intelligent Power Node Manager allows Oracle to manage power consumption, allowing more servers in each rack and increasing computing density. As a result, Oracle projects energy savings of 35 percent (saving money) and 50 percent more servers per rack (saving data center space and enabling future growth).

For more information, visit software.intel.com/en-us/articles/intel-cloud-builder-success-stories.

capacity. Deployment of Intel Intelligent Power Node Manager in very large data center environments has led to up to 40 percent efficiency improvement.⁵ In one specific example, Baidu, the largest search provider in China, was able to increase rack density by up to 40 percent (see software.intel.com/en-us/articles/intel-cloud-builder-success-stories) and Oracle IT saw a 35 percent energy savings with Intel Node Manager (see sidebar, *Oracle Uses Intel® Intelligent Power Node Manager to Increase Rack Density*).

- **Intel® Data Center Manager:** By scaling Intel Intelligent Power Node Manager to the data center level, Intel® Data Center Manager enables fine-grained control of power for servers, racks and groups of servers in data centers. It allows managers to dynamically migrate workloads to optimal servers based on specific power policies. Real-time monitoring of actual power and inlet temp data is aggregated to rack-, row-, room- and user-defined logical groups. Built-in, policy-based heuristics maintain group power capping while dynamically adapting to changing server loads and minimizing the performance impact of workloads. Actual power consumption data can be used to create specific scenarios – for example, linking the power data in real time to HVAC usage – to enable charge-back models based on actual power used. Intel Data Center Manager supports a range of capabilities to better manage power at the rack level and across multiple racks in the data center, and is available as a Software Development Kit (SDK) to integrate with existing management consoles. (Get more details on Intel Node Manager and Intel Data Center Manager at software.intel.com/sites/datacentermanager.)
- **Intel® Solid State Drives (Intel® SSDs):** Intel offers a portfolio of solid state drives. Intel SSDs provide dramatic performance and power advantages compared to standard hard disk drives, providing up to 6x the read performance of rotating drives and up to 80 percent power savings in typical data center applications – resulting in significant reductions in power consumption.⁶ As data center demands continue to increase and solid state drive costs continue to drop, solid state drive deployments will accelerate. (Learn more about Intel SSDs at www.intel.com/design/flash/nand)

Intel is also engaged in rack- and data center-level optimization to help reduce Total Cost of Ownership (TCO) with a focus on reducing the PUE (Power Usage Effectiveness) closer to PUE 1.0. These efforts include container optimization, power and thermal optimization, and improvements in rack power delivery, among others. For example, we are developing proof points and best practices on running the data center at higher ambient temperatures to reduce cooling costs.

Use of these Intel power technologies, capabilities and optimizations can greatly assist data center architects and managers in increasing performance to meet growing user demands and specific workload requirements while addressing the real physical constraints related to power, space, and cooling.

Improved virtualization capabilities: Many cloud computing deployments use virtualization to achieve greater levels of utilization and workload consolidation with a goal of enabling multi-tenancy, i.e., supporting many customers on the same server, storage and networking infrastructure. Virtualization is a technology that spans hardware and software, so real virtualization performance gains depend on

enhancements that work across platforms and virtualization applications. Using a platform approach, Intel® Virtualization Technology® (Intel® VT) provides hardware assists included in the CPU, the chipset, and networking devices that tune virtualization performance. By off-loading workloads to system hardware, virtualization software can provide more streamlined software stacks. Using Intel Xeon processor 5600 series-based platforms with Intel VT provides the foundation for cloud computing infrastructure with up to a 42 percent performance increase in virtualization performance over previous-generation Intel Xeon processor-based servers⁵ (Learn more about Intel VT at www.intel.com/technology/virtualization/server.)

Key Considerations for Efficiency

Select building blocks for your cloud computing environments that are capable of establishing a highly efficient, flexible infrastructure that makes the most of available resources. Advanced power management capabilities such as Intel Node Manager, Intel Data Center Manager, Intel's virtualization technologies, and leading performance per watt in servers based on the Intel Xeon processor 5600 series provide an ideal foundation to enable a more highly-efficient and cost-effective data center. In addition, Intel SSDs deliver excellent performance and significant improvements in power efficiency compared to standard hard disk drives.

Simplification

IT Challenges

As data centers expand and proliferate, new complexities have arisen – in fact, scaling is a major source of overall complexity. Where server sprawl once reigned, now data center managers must deal with virtual machine sprawl. Increased utilization of physical servers through virtualization has caused large increases in network bandwidth requirements and storage demand. More connectivity means cabling, more domains, and more management. The presence of multiple architectures in the data center adds to complexity and increases operational costs. For data centers to handle rapid rates of growth, infrastructure must be dramatically simplified. Fewer configurations, fewer ports, fewer cables, fewer servers, fewer points of management, and simpler integration between vendors – any of these would help the rising complexity of today's (and tomorrow's) data center.

Intel Advancements to Simplify Data Center Infrastructure

Intel is active on multiple fronts to simplify data center architecture and deployment. Three significant areas of work are virtual machine migration, unified 10GbE networks, and the convergence of servers and storage.

Simplified virtual machine migration: Using Intel® VT FlexMigration enabled by the leading virtual machine software vendors, data center managers and system administrators can set simple rules for VM migration based upon workload, time of day, and memory requirements. These tools, such as VMware's Enhanced VMotion[®], enable a simplified and disaster-resistant solution for today's data center infrastructure. Intel VT FlexMigration works in conjunction with VM vendor solutions (such as VMotion) to protect infrastructure investments in Intel Xeon processor-based platforms through architectural compatibility from one generation to another, enabling live VM migration among platforms. This preserves the value of your investments today while facilitating live migration to future performance and efficiency gains in the future. Over time, Intel will continue to enable further capabilities for Disaster Recovery and Business Continuity into Intel VT, which will allow for better instrumentation and additional application workload optimization.

Intel® Cloud Builder Program

Delivering on the promise of cloud computing starts with the ability of the underlying compute, network and storage infrastructure to act as an efficient, shared resource pool that is dynamically scalable within one or across multiple data centers. With this foundation, critical higher-level capabilities such as secure multi-tenancy, guaranteed quality of service, federation and data center automation are made possible.

While it may be easy to agree with this conceptual framework, building a cloud that conforms to it can be quite a daunting task. Specifically, the "getting started" problem posed by the myriad hardware and software combinations that could be used to realize a given set of architectural tradeoffs can result in "analysis paralysis."

In recognition of this challenge, Intel formed Intel® Cloud Builder – a program designed to ease the deployment of cloud infrastructure for service providers, Telcos, hosting companies and enterprises. The Intel Cloud Builder program delivers cloud reference architectures, combining Intel® Xeon® processor-based servers with leading software solutions to provide deployment recipes, best practices, success stories using the latest Intel® technologies. In addition, Intel Cloud Builder provides information and white papers on advanced cloud research.

In cooperation with leading cloud independent software vendors, including Canonical, Citrix, Enomaly, Microsoft, Novell, Parallels, Red Hat, Univa, VMware, Xen, and others, Intel is delivering cloud deployment guides, or recipes on how to get started building a cloud. Intel Cloud Builder recipes provide an easy starting point to build a cloud environment based on a basic hardware blueprint and using available cloud software management solutions. This helps to eliminate time-consuming back-end work by providing a documented solution to ease the deployment of a cloud.

For more information, see www.intel.com/software/cloudbuilder.

Unified networks based on 10-Gigabit Ethernet (Intel® Ethernet 10GbE):

A number of factors are driving greater demands for network bandwidth and causing I/O to become a bottleneck within the data center. Increases in virtual machine density consolidate more and more I/O on the server bus. Faster processors require more network bandwidth. And more complex applications (such as databases and ERP) are being virtualized.

A unified network fabric based on Intel Ethernet 10GbE can simplify existing infrastructures and lower TCO while positioning the data center for continued growth in bandwidth requirements. This approach has a number of advantages:

- **Higher network performance with reduced complexity:** 10GbE increases the speed of Ethernet networks to 10 Gbps (yielding up to a 10x increase in I/O bandwidth) and reduces power per Gb by 4.5x. It does this within the familiar Ethernet environment, maintaining familiar management tools and drawing on a common skills base. Port counts can be reduced up to 5x⁶. The number of switches and cables can be reduced, and a simpler network yields lower management costs.
- **Improved virtualization:** Intel® Virtualization Technology for Connectivity (Intel® VT-c) enables improved virtualization across Intel 10GbE networks. One of the core technologies of Intel VT-c is Virtual Machine Device Queues (VMDq) which improves network throughput via efficient routing of packets to the correct target virtual machines using hardware queues. Intel VT-c also supports the PCI-SIG specification for Single Root I/O Virtualization (SR-IOV), which improves throughput and lowers CPU utilization by enabling direct I/O connection to virtual machines, by-passing the Virtual Machine layer.
- **Improved storage networking:** Intel Ethernet 10GbE has a number of features that enhance network storage. Data Center Bridging (DCB) improves Quality of Service for data and storage traffic on the same network. 10GbE with DCB provides sufficient bandwidth and support (with DCB) for storage over Ethernet, including Network Attached Storage (NAS) and Storage Area Networks (SAN). Internet Small Computer System Interfaces (iSCSI) enable IP-based storage transmission over LANs or WANs. Fiber Channel over Ethernet (FCoE) allows fiber channels to use Intel Ethernet 10GbE networks while preserving current investments in the Fiber Channel protocol. iSCSI and FCoE are both block storage that are used in SANs.
- **Ease of transition:** Since existing Ethernet infrastructures are easily interoperable with 10GbE, it offers a straightforward migration to higher network performance. 10GbE uses processes, protocols, and management tools already deployed in the management infrastructure.
- **Lower costs:** Because Intel Ethernet 10GbE builds on existing processes and tools, both acquisition and support costs are lower versus alternative technologies. Both the interoperability with existing network infrastructure and the large number of equipment vendors lead to lower adoption cost as well as a high level of flexibility in network design.

The demands of cloud computing architectures and applications require a significantly more capable network infrastructure. We have shown that unified networks based on Intel Ethernet 10GbE and the Intel Xeon processor family deliver a leading combination of capabilities based on their performance, cost, flexibility, and room for growth.

The convergence of servers and storage: The increasing demands on the data center naturally include, and contribute to, increasing demands on storage architectures. Storage capacity and performance requirements are growing rapidly as data center compute volumes and throughputs increase. At the same time, storage uptime is critical and the growing complexity of storage architectures puts greater demands on management systems and capabilities. Fortunately, a convergence of storage and computing is underway, as Intel® architecture building blocks are increasingly the foundation of storage systems that can keep up with growing demands. Industry leaders like EMC have chosen Intel Xeon processors as their architecture of choice (with the announcement of the EMC Symmetrix* solutions.) We expect that by the end of 2010, seven out of every ten external storage systems shipped will be based on Intel architecture.

Key Considerations to Simplify your IT Infrastructure

Consider infrastructure strategies that build simplicity in today and help maintain it going forward. Adoption of flexible Ethernet across the data center infrastructure and conversion to 10GbE helps meet today's bandwidth and complexity challenges and provides headroom for growing cloud computing deployments. In addition, to further improve and streamline virtualization, consider advanced technologies like Intel VMDq.

Security

IT Challenges

In a cloud computing environment, new security issues arise while old security issues continue to evolve. In a public or virtual private cloud, your data is on a server controlled by someone else, so traditional security models aimed at protecting the perimeter of the organization are no longer sufficient. Methods to control data corruption, access, disruption, and loss must be adapted to cloud computing. Cloud computing shares some of these dynamics with today's cross-business and cross-supply chain collaboration models. The growth of cloud computing will simply elevate the problems.

At the same time, traditional security problems unrelated to cloud computing continue to grow. Attacks are changing from hackers working on their own and merely looking for personal fame to organized, sophisticated efforts targeting specific types of data and seeking to gain and retain control of assets. And finally, regulatory environments continue to change, with increased requirements for compliance, audit, reporting, privacy protection, and indemnification. The risks and costs of non-compliance are large and growing.

Intel Advancements in Security

Within this environment, Intel has developed a number of technologies that make robust levels of security easier and more efficient to accomplish for improved protection. These technologies are built into the Intel Xeon processor 5600 series to assist with protecting data both in the data center and when it is on the move.

- **Intel® Virtualization Technology (Intel® VT) for increased isolation of virtual machines:** Traditional software-based virtualization is improved by the hardware-based capabilities of Intel VT. By embedding specialized functions within the processor, chipset, and BIOS, Intel VT can simplify resource management and improve reliability and availability. Intel® Trusted Execution Technology[†] (Intel® TXT), when combined with Intel VT, provides users with a trusted pool of virtual machine resources and IT managers with the security to avoid “rootkit” attacks. Intel® Advanced Encryption Standard-New Instructions (Intel® AES-NI) enables more secure SSL and IPSec transmit capabilities throughout the cloud computing environment, providing an isolated virtual environment to protect against intrusion. Using Intel VT, Intel TXT and Intel AES-NI (combined with VM and/or related security software) in cloud computing provides IT managers with the ability to isolate, enforce and encrypt their cloud computing environments.
- **Intel® VT FlexMigration for safer migration of virtual machines:** By enabling virtualization software vendors to deliver live migration of virtual machines across different generations of Intel Xeon processor-based servers, Intel VT FlexMigration dramatically enlarges the pool of servers that can be deployed in a virtual machine migration compatibility resource pool. This allows IT managers to accelerate adoption cloud computing architectures. Intel VT FlexMigration protects your investment in Intel Xeon processor family servers by providing generational compatibility from this year’s Intel Xeon 5600 and 7500 family of processors to previous Intel Xeon processors (introduced in 2007 and later) and into the future of the Intel Xeon processors family roadmap.⁷ This technology allows IT and data center managers to shorten their refresh cycles, increase consolidation ratios based upon the latest technologies, and provide the secure infrastructure availability that your end users are demanding.
- **Intel® AES-NI hardware acceleration to enable pervasive encryption:** One of the barriers to using encryption more broadly has been the computational tax it puts on hardware performance. By integrating Intel AES-NI in the server processor, Intel has accelerated data encryption significantly – up to 89 percent reduction on encryption/decryption time⁸ – and lowered the overhead required. By removing this key barrier, pervasive encryption of data now becomes feasible. For example, Intel AES-NI can be used to optimize SSL and IPSec to protect data in flight. Additionally, Intel’s NICs also have the ability to offload the IPSec computation; with this combination, line speed encryption can be achieved for data in flight. Within the data center, Intel AES-NI hardware acceleration makes it feasible to deploy full-disk encryption and database encryption widely to protect data at rest.
- **Intel® Trusted Execution Technology (Intel® TXT) to increase security of hypervisors and virtual machines:** In a virtual machine environment, launch-time has been a relatively unprotected phase of operation. Intel TXT works cooperatively with Intel VT to provide hardware-assisted protection against rootkits and other malware

threats at launch – for example, attempts to insert non-trusted VMM (rootkit hypervisor); reset attacks designed to compromise platform secrets in memory; and BIOS and firmware update attacks. Intel TXT locks down the virtualization configuration so that it can only be changed by entering into the TXT-driven secure mode. This complements runtime protections such as anti-virus and intrusion detection.

Key Considerations for Security

As you design cloud computing deployments, a full understanding of your unique security needs and threats is essential. That audit can then serve as a guide to selecting hardware and software building blocks for your infrastructure. Standardizing on Intel architecture platforms allows for advanced security technologies developed by Intel and the industry’s security ecosystem to be available across your infrastructure. Each layer of virtualization security, enterprise wide security, and system level security reduces the “surface area” for malicious attacks across your cloud computing environment. Over time, Intel’s node, rack, and data center instrumentation will allow IT managers to isolate virtual and physical machine attacks across their IT infrastructure. It’s also important to utilize security best practices to proactively and efficiently focus on prevention, detection, and response resources. For an example from IT@Intel, visit www.intel.com/it.

Intel IT’s Inside-Out Cloud Strategy

Managing nearly 100,000 servers, Intel IT is evolving its own cloud computing strategy by starting inside the corporation (private clouds) and positioning the company to use external (public) clouds over time. We have implemented a design computing grid that has many cloud computing characteristics. Our priority is to initially grow our internal enterprise virtualized computing environment to support an increasing number of cloud-like attributes over time. We plan to aggressively expand and evolve this internal environment with on-demand self service, an automated hosting framework, and measured services

Additionally, Intel is already taking advantage of external cloud computing technologies. We have many opportunistic software as a service (SaaS) implementations, including Web conferencing and social media solutions. Our preliminary experiences with infrastructure as a service (IaaS) suggest that it may be suitable for rapid development and some batch applications.

A strategy of growing the cloud from the inside out delivers many of the benefits of cloud computing now and positions us to utilize external clouds over time. We expect to selectively migrate services to external clouds as supplier offerings mature, enterprise adoption barriers are overcome, and opportunities arise for improved flexibility and agility as well as lower costs.

For more information, visit www.intel.com/it.

Open Standards

IT Challenges

The attraction and promise of cloud computing is the ability to rapidly deploy IT services and solutions at lower cost, thus enabling faster innovation by IT users. But to realize this potential, cloud computing requires an open environment with interoperable building blocks and solutions based on multi-vendor innovation. Proprietary solutions quickly drain the cost and flexibility advantages from cloud computing. While few standards exist today, Intel and other hardware and software vendors are deeply involved in creating and evolving a core set of standards designed to keep cloud computing simple, efficient, secure, and open.

Intel and Open Standards

Intel has a long history of open standards leadership within the IT industry and has played a historic role in the growth of non-proprietary solutions across the data center and the enterprise. That leadership is continuing in a number of areas critical to the future of cloud computing, spanning components, platforms, software, and management technologies. The underlying objective is to promote multi-vendor innovations that drive down acquisition, deployment, and operational costs. To the extent the industry succeeds at implementing open standards, data center architects and managers will experience more choices, robust interoperability, ease of integration, workload portability across clouds, simplified management, and greater security.

Intel is active in multiple industry forums to enable key cloud computing standards. For example:

- **Interoperability standards:** Intel is co-chair of the Open Cloud Standards Incubator within the Distributed Management Task Force (DMTF), which is defining standards for interoperable cloud interfaces and portable workloads as well as standards for power management integrated with building management systems. Intel is also actively involved with the Open Grid Forum (OGF) which is defining the Open Cloud Computing Interface (OCCI), a practical solution to interface with cloud infrastructures exposed as a service (IaaS). As a major contributor to the Server System Infrastructure (SSI) standards body, Intel is helping create hardware standards that are well suited to cloud computing such as half-width motherboards, and the emerging low-end, entry-level single-socket server category of micro servers.

- **Manageability standards:** Intel was a leader in defining the Data Center Manageability Interface (DCMI) specification, a subset of the IPMI specification providing low-cost, non-proprietary server manageability tailored to the specific needs of large-scale data centers.
- **Efficiency standards:** As a member of the W3C HTML5 Working Group, Intel is helping to form the HTML5 specification to enable superior web/runtime experiences that can also take advantage of Intel platform features. Intel is a key contributor in two forums on the Green Grid: the Power Management Data Collection Task Force, which is driving industry surveys and data collection for power management feature requirements; and the Container Data Center Task Force, working to standardize energy efficiency metrics. Intel is also investing significant engineering work and proof-of-concept efforts targeted at data center best practices, with a goal of driving data center Power Usage Effectiveness (PUE) towards 1.0. Examples of these efforts include Intel's work on the Green Grid; engagement with standards groups like Energy Star; and proof-of-concept studies such as the Advanced Cooling Environment study with LBNL, IBM, HP and Emerson; and warmer temperature data center studies with ASHRAE.
- **Security standards:** As a member of the Cloud Security Alliance (CSA), Intel is helping to drive best practice recommendations and provide better mapping of CSA recommendations into Intel's silicon capabilities.
- **Storage and networking standards:** As a contributor and reviewer to the IEEE 802.1 and 802.3 standards, Intel is helping to shape a study group focused on cloud computing networking standards. Within the PCI-SIG, Intel was a major contributor to the SR-IOV specification. Regarding storage standards, Intel is a member of the Storage Networking Industry Association (SNIA) and a contributor to the STORM (STORage Maintenance) task force, providing inputs into iSER and iSCSI specifications.
- **Virtualization Standards:** As one of the largest contributors to virtualization standards in the technology industry, Intel participates actively in the Xen Open Source Hypervisor community, SpecVirt Benchmarking committee, and multiple industry forums for the advancement of virtualization technologies. Within the PCI-SIG, Intel was a major contributor to the SR-IOV specification. We continue to work with the industry at large to define new usage models for virtualization and actively contribute to the standards organizations driving these changes.

Key Considerations about Driving Open Standards

Industry-wide standards and solutions are the foundation of interoperability and lower cost. Whenever possible, take advantage of industry standards. You should consider playing an active role within the industry in enabling standards to provide greater levels of interoperability, choice, and flexibility.

Conclusion

Cloud computing is major paradigm shift in IT development with powerful advantages. For any enterprise IT department, standing still is not an option. Public cloud service providers also have a need to continue to evolve and advance their IT infrastructure.

Holding the promise of delivering significantly better automation, scalability, and resource efficiency, cloud computing is a transformation in the design, development, and deployment IT technologies. But to deliver on the promise, suppliers and architects will need to focus on coordinated efforts to evolve cloud computing environments that are federated, automated, and client-aware.

Intel believes the path to that promise is a broad and deep focus on key principles of cloud computing infrastructure: efficiency, simplification, security, and open standards. As a supplier and leading member of the cloud computing ecosystem, Intel is investing heavily in each of the four principles. We believe that in order for cloud computing environments to evolve and meet the requirements of users, it's important that the industry focus on these four principles to achieve a vision that provides for greater levels of federation, automation, efficiency, security, and simplification built around open standards.

For More Information

- Intel® Hyper-Threading Technology, visit www.intel.com/technology/platform-technology/hyper-threading/index.htm.
- Intel® Turbo Boost Technology, visit www.intel.com/technology/turboboost/index.htm.
- Intel® QuickPath Technology, visit www.intel.com/technology/quickpath/index.htm.
- Intel® Intelligent Power Node Manager, visit software.intel.com/sites/datacentermanager.
- Intel® Data Center Manager, visit software.intel.com/sites/datacentermanager.
- Intel® VT FlexMigration, visit www.intel.com/technology/virtualization/server/index.htm.
- Intel® 10-Gigabit Ethernet, visit www.intel.com/network/connectivity/resources/technologies/10_gigabit_ethernet.htm.
- Intel® Virtualization Technology (Intel® VT), visit www.intel.com/technology/virtualization/server/index.htm.
- Intel® AES-NI hardware acceleration, visit software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni.
- Intel® Trusted Execution Technology (Intel® TXT), visit www.intel.com/technology/security/index.htm.
- Intel® SSDs, visit www.intel.com/design/flash/nand.

Summary of the Intel® Xeon® Processor Family

- **Intel Xeon processor 3400 series-based platforms** are one-socket servers that unleash the computing power of Intel Xeon processors for entry-level servers and basic web workloads. For more information visit www.intel.com/xeon.
- **Intel Xeon processor 5600 series** is an ideal foundation for maximizing the efficiency, security, and scalability of cloud computing infrastructures. For more information visit www.intel.com/xeon.
- **Intel Xeon processor 7500 series** and Intel Xeon processor 6500 series are an ideal foundation for cloud computing environments running mission-critical applications requiring the highest reliability. For more information visit www.intel.com/xeon.

^A Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

¹ Hyper-Threading Technology requires a computer system with an Intel processor supporting Hyper-Threading Technology and an HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See <http://www.intel.com/info/hyperthreading/> for more information including details on which processors support HT Technology.

⁵ Intel® Turbo Boost Technology requires a Platform with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your platform manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see www.intel.com/technology/turboboost.

⁶ Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

⁴ 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). The MLE could consist of a virtual machine monitor, an OS or an application. In addition, Intel TXT requires the system to contain a TPM v1.2, as defined by the Trusted Computing Group and specific software for some uses. For more information, see <http://www.intel.com/technology/security>.

¹ Sources: IDC "Server Workloads Forecast" 2009; and IDC "The Internet Reaches Late Adolescence" Dec 2009

² Jonathan Koomey, 6/09. http://www.nytimes.com/2009/06/14/magazine/14search-t.html?_r=3&ref=magazine&pagewanted=all.

³ Source: Intel internal estimates, 2010.

⁴ Source: Internal Intel measurements for Intel® Xeon® processor X5680 vs. Intel® Xeon® processor X5570 on BlackScholes*. Also, Fujitsu Performance measurements comparing Intel® Xeon® processor L5650 vs Intel® Xeon® processor X5570 SKUs using SPECint_rate_base2006. See <http://docs.ts.fujitsu.com/dl.aspx?id=0140b19d-56e3-4b24-a01e-26b8a80cfe53> and <http://docs.ts.fujitsu.com/dl.aspx?id=4af74e10-24b1-4cf8-bb3b-9c4f5f177389>.

⁵ Gain is on VMmark* benchmark vs. a Intel® Xeon® processor 5570 series processor-based server. VMmark* is designed as a tile-based benchmark consisting of a diverse set of workloads commonly found in the data center. The workloads comprising each tile are run simultaneously in separate virtual machines at load levels that are typical of virtualized environments. The performance of each workload is measured and then combined with the other workloads to form the score for the individual tile. Multiple tiles can be run simultaneously to increase the overall score. A tile is a collection of six diverse workloads concurrently executing specific software. Running on one of two separate operating systems (Windows or Linux), each workload runs in its own virtual machine and executes specific applications. Included in a single tile are a web server, file server, mail server, database, Java server, as well as an idle machine. Results are posted at <http://www.vmware.com/products/vmmark/results.html>. Performance 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 <http://www.intel.com/performance/resources/limits.htm>.

⁶ Intel Internal estimates, 2009.

⁷ Backward compatibility for live VM migration also exists with dual-core Intel® Core™ microarchitecture products (Intel® Xeon® processor 5100 and Intel® Xeon® processor 3000) and forward compatibility with future dual and multi-core processors. Contact your preferred VMM vendor for support requirements. All future products, dates, and figures are preliminary and are subject to change without any notice.

⁸ Oracle 11g with TDE, time takes to decrypt a 5.1 million row encrypted table with AES-256 CBC mode on Intel® Xeon® processor X5680 (WSM, 3.33 GHz) optimized with Intel® Performance Primitives crypto library (IPP) vs Intel® Xeon® processor X5560 (NHM, 2.8 GHz) without IPP. Timing measured is per 4K of data.

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.

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 © 2010 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and Xeon inside are trademarks of Intel Corporation in the U.S. and other countries.

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

