Intel's Vision of the Ongoing Shift to Cloud Computing

Executive Overview

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, solutions, and best practices for cloud computing are evolving, and 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 solutions is they 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.

A Paradigm Shift

Rather than a *revolution*, cloud computing is an important transition, a *paradigm shift* in IT delivery – one that has broad impact and important challenges to consider. Cloud computing offers the potential for a transformation in the design, development, and deployment of next-generation technologies – technologies that enable 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. These demands include growing needs to manage business growth and increase IT flexibility. 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 behind a firewall for an organization's internal use).

Public clouds are 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.

Private clouds are 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.

Cloud computing is a step beyond data center virtualization. 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 and 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 optimized use and rapid deployment of resources, improved operational efficiency, and potential for significant cost savings. When fully realized, cloud computing infrastructures can provide competitively significant IT agility, flexibility, and adaptability through systems that are efficient, simplified, and secure.

Though cloud computing can be viewed as an evolutionary step, it is a fundamental shift and there are challenges to consider:

- 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.

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. Several 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**.

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. Public clouds are being driven by explosive growth of Internet data and traffic as the Internet matures and Internet-based services proliferate.

A **private cloud** is a cloud architecture deployed behind a firewall for an organization's internal use. An extension of virtualization deployments, private clouds allow IT to be delivered as a service with significant levels of automation and flexibility.

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. Intel's cloud computing vision calls for a level of federation that enables the movement of workloads and data from one service provider to another; burst implementations between internal private cloud and public cloud providers if additional capacity is needed; and secure and reliable data flow across vendors, partners, and clients.

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. According to IDC's *Data Center Survey* in 2009, virtualization thus far has failed to reduce complexity. The number of server instances that can be managed by the average systems administrator has increased from 37 to only 41 comparing non-virtualized servers to virtualized servers. Moreover, 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. This includes automation of provisioning, resource monitoring, reporting of consumption for bill back, and workload balancing.

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. 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.

Principles to Enable the Vision of Cloud Computing: Efficiency, Simplification, Security, and Open Standards

Evolving the infrastructure to realize the full potential of cloud computing 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:

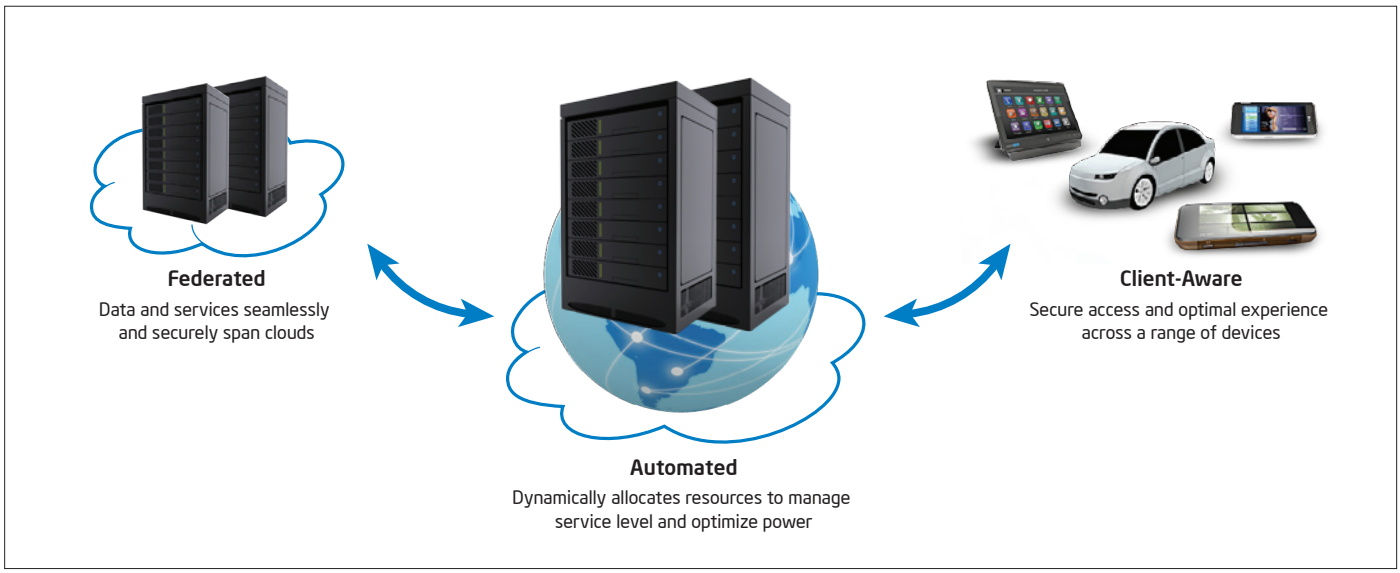

Figure 1. Intel's Cloud Computing Vision

- **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 interoperability.

The companion piece to this brief (see link below) discusses each of these four principles in detail: the drivers behind each 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.

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 the 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.

A prime example is the Intel® Cloud Builder, a program designed to ease the deployment of cloud infrastructure for service providers, Telcos, hosting companies and enterprises. The program, created by Intel in conjunction with leading cloud independent software vendors (including Canonical, Citrix, Enomaly, Microsoft, Parallels, Red Hat, Univa, VMware, Xen, and others), delivers cloud reference architectures that make it easier to deploy cloud infrastructure. These reference architectures combine Intel® Xeon® processor-based servers with leading software solutions to provide deployment recipes that take advantage of Intel technologies and innovations. Intel Cloud Builder (www.intel.com/software/cloudbuilder) also provides best practices, success stories, and information on advanced cloud research. 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.

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.

For more information on how to evolve your infrastructure to be more efficient, secure, and simplified while taking advantage of evolving standards for cloud computing, see the *Cloud Computing: Considerations and Next Steps* white paper at www.intel.com/go/cloud.

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

Printed in USA 0610/SD/OCG/XX/PDF Please Recycle 323941-001US