

White Paper
Intel Information Technology
Computer Manufacturing
Cloud Computing

Developing an Enterprise Cloud Computing Strategy

Cloud computing is a significant trend with the potential to increase agility and lower costs. Today, however, security risks, immature technology, and other concerns prevent widespread enterprise adoption of external clouds. Intel IT is developing a strategy based on growing the cloud from the inside out. We take advantage of software as a service (SaaS) and niche infrastructure as a service (IaaS) implementations whenever possible, and we are building an internal cloud-computing environment. Our internal environment delivers many of the benefits of clouds and positions us to use external clouds in the future, as supplier offerings mature and barriers to enterprise adoption are overcome.

Hong Li, Jeff Sedayao, Jay Hahn-Steichen, Ed Jimison, Catherine Spence, and Sudip Chahal,
Intel Corporation

January 2009

IT@Intel

A strategy of growing the cloud from the inside out delivers many of the benefits of cloud computing and positions us to utilize external clouds over time.

Executive Summary

Cloud computing is a significant trend with the potential to increase agility and lower costs. Today, however, security risks, the lack of mature technology and standards, and other concerns prevent widespread enterprise adoption of external clouds.

Intel IT is developing a cloud computing strategy based on growing the cloud from the inside out. We already have internal initiatives with many cloud computing characteristics. We will initially grow our internal virtualized computing environment to support an increasing number of cloud-like attributes over time. We plan to aggressively expand and evolve this internal environment.

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

Many applications are not suitable for hosting in external clouds at present. Good candidates may be applications that have low security exposure and are not mission-critical or competitive differentiators for the corporation.

A strategy of growing the cloud from the inside out delivers many of the benefits of cloud computing 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.

Contents

Executive Summary	2
Background	4
Cloud Attributes and Taxonomy	4
Expected Benefits and Risks	5
Benefits	5
Risks	5
Software as a Service and Infrastructure as a Service Experiences	7
Software as a Service	7
Infrastructure as a Service	8
Key IT Architecture Considerations	9
Composite Applications	9
Standards	10
External and Internal Clouds	10
High-Level Cloud Computing Strategy	11
Applications Suitable for External Clouds	11
Software as a Service and Infrastructure as a Service	12
Cloud Computing Adoption	12
Conclusion	14
Authors	15
Acronyms	15

Background

Cloud computing technology is a significant trend with implications for Intel IT. A growing number of suppliers are starting to provide cloud computing offerings, and analysts project that some enterprises will purchase a significant percentage of their applications and infrastructure as cloud computing services within a few years.

We have determined that cloud computing could provide significant benefits to Intel, including increased agility. However, there are also considerable risks. In analyzing cloud computing and developing our strategy, we are attempting to answer two related questions:

- Which services should we move to cloud computing and when?
- How do we map a path to cloud computing from our current environment?

Cloud Attributes and Taxonomy

Because this is an emerging and somewhat confusing area, we have created definitions that provide us with a common basis for discussion and developing our strategy.

We define cloud computing as a computing paradigm where services and data reside in shared resources in scalable data centers, and those services and data are accessible by any authenticated device over the Internet.

We have also identified some key attributes that distinguish cloud computing from conventional computing. Cloud computing offerings are:

- Abstracted and offered as a service.

- Built on a massively scalable infrastructure.
- Easily purchased and billed by consumption.
- Shared and multi-tenant.
- Based on dynamic, elastic, flexibly configurable resources.
- Accessible over the Internet by any device.

Today, we have identified three main categories of external service that fall within our broad cloud computing definition.

- **Software as a service (SaaS).** Software deployed as a hosted service and accessed over the Internet.
- **Platform as a service (PaaS):** Platforms that can be used to deploy applications provided by customers or partners of the PaaS provider.
- **Infrastructure as a service (IaaS):** Computing infrastructure, such as servers, storage, and network, delivered as a cloud service, typically through virtualization.

It is also possible to build an internal IT environment with cloud computing characteristics. We call this an internal cloud, to differentiate it from the external clouds provided by suppliers.

Expected Benefits and Risks

Our decisions about how and when to move services to external clouds are based on the balance between risk and reward. Today cloud computing is relatively immature, and for large enterprises, the risks of wholesale adoption outweigh the potential benefits. As the technology evolves, we expect this balance to shift so that the benefits begin to outweigh the risks for a growing number of applications and services.

Benefits

Potential benefits of cloud computing include:

Agility, Adaptability, and Flexibility

A business group that wants to deploy a new application can do so relatively quickly using cloud computing services, compared with weeks or months it can take with the traditional enterprise model of buying servers, installing them, and then deploying the application to the new servers. In many cases, users can purchase cloud services with a credit card and begin to use them almost immediately.

Because cloud computing is built on a massively scalable shared infrastructure, cloud suppliers can in theory quickly provide the capacity required for very large applications without long lead times. Purchasers of IaaS capacity can run applications on a variety of virtual machines (VMs), with flexibility in how the VMs are configured. Some cloud computing service providers have developed their own ecosystem of services and service providers that can make the development and deployment of services easier and faster. Adding SaaS capacity can be as easy as getting an account on a supplier's host.

Cloud computing is also appealing when we need to quickly add computing capacity to handle a

temporary surge in requirements. Rather than building additional infrastructure, cloud computing could in principle be used to provide on-demand capacity when needed.

Cost Savings

There is a perception that cloud computing can reduce cost. To date, savings have generally been more clearly shown for small to medium-size businesses (SMBs). However, we have achieved cost savings with some of our SaaS deployments, indicating that cost savings can be a factor in propelling enterprise cloud adoption.

The relatively low upfront cost of IaaS and PaaS services, including VMs, storage, and data transmission, can be attractive—especially for addressing tactical, transient requirements such as unanticipated workload spikes. An additional advantage is that businesses pay only for the resources reserved; there is no need for capital expenditure on servers or other hardware.

Risks

The features that make cloud computing so appealing, combined with the fact that services are publicly accessible, can also lead to many potential risks.

Security and Privacy

Today, security and privacy may represent the biggest risks to moving services to external clouds. The advantages of cloud computing—flexibility, easy-to-use service abstractions, and shared infrastructure—also introduce the concern that people may use cloud computing in a way that puts Intel's information and intellectual property at risk.

With cloud computing, data is stored and delivered across the Internet. The owner of the data does not control—and typically does not even know—the location of the data. There is a very real possibility that the owner's data could reside on the same resources as a competitor's application and data. Additionally, in a multi-tenant environment, it may be very difficult for a cloud service provider to provide the level of isolation and associated guarantees that are possible with an environment dedicated to a single customer.

Enterprises cannot rely solely on contractual controls with cloud service providers. In many cases, these controls do not provide adequate protection. It would be difficult or even impossible to use a public cloud for applications that handle controlled technologies, due to the risk of potential compromises and concerns about compliance. For example, external data storage provided by a cloud service supplier might be located in a controlled country to reduce cost.

Standards are lacking for security and for managing service-level agreements (SLAs) that could be used to help ensure compliance with government regulations and Intel standards through independent, third-party audits. Enterprise security policies may stipulate, for example, that all data held externally must be encrypted in transit and at rest. In addition, IT policies may specify that virtual servers supporting certain applications must not share the same physical server. In general, cloud providers do not currently provide capabilities to guarantee compliance with these policies or facilities for auditing compliance. There are also questions about the extent to which cloud

providers would accept legal responsibility, and the damages that could be assessed, in the event of a proven breach of contract involving a security issue.

In many countries, the use and storage of personal data is heavily regulated; without adequate safeguards, this data could be illegally exposed with an external cloud. Applications implemented within our internal environment are relatively easy to audit, and we have well-established techniques for doing this. However, applications implemented outside Intel using cloud computing would be much more difficult to audit. As a result, we might be unable to determine whether applications developed and running outside our firewall had received up-to-date security patches or were otherwise vulnerable to compromise. An additional area of concern is the secure integration of business data generated in an external cloud with existing data stored within Intel's environment.

Enterprise Support and Service Maturity

Cloud computing services may not provide the levels of reliability, manageability, and support required by large enterprises. Today, many services are aimed primarily at SMBs and at consumers, rather than large enterprises.

Uptime SLAs offered by some providers may be inadequate for some enterprise applications. In addition, there may not always be a clearly defined method for validating the SLA.

Cloud computing implementations of some services may lack features we have come to expect in stand-alone enterprise versions and may not integrate well with enterprise applications. We investigated cloud-based e-mail and calendar services, and initially encountered problems synchronizing e-mail, calendars, and address lists with our existing enterprise application. Security was also an issue.

Return on Investment Concerns

General perception is that external cloud computing can reduce costs for large enterprises as well as SMBs. However, the cost advantages for large enterprises may not be as clear as for SMBs, since many large enterprises can reap the benefits of significant economies of scale in their own internal IT operations.

While cloud computing initially appears to be less expensive in terms of upfront costs, the comparison may be much more competitive when total cost of ownership (TCO)—including

recurring costs—and potential risks are taken into account.

There may be other hidden cost impacts. Migration to an external cloud may entail significant changes or additions to the enterprise network in order to provide acceptable performance to corporate users in regions with limited bandwidth. Over the last few years, Intel has been reducing the number of points at which our network accesses the Internet; a cloud model may require more touch points. Increases in bandwidth may be necessary, and in many countries bandwidth is still very expensive.

Software as a Service and Infrastructure as a Service Experiences

We have gained experience with both SaaS and IaaS. We have migrated many individual applications to external SaaS clouds. We use IaaS for some niche applications, and we have also used IaaS platforms to host some experimental and proof-of-concept (PoC) services.

Software as a Service

Intel is already using some SaaS applications. The biggest uptakes have been applications for managing travel, expense reporting, hiring and staffing, and employee benefits. Additionally, Intel has used SaaS for Web conferencing and social media solutions and has explored using SaaS for office applications and customer relationship management (CRM).

The biggest factors in using SaaS have been functionality and project acceleration. Intel used SaaS in order to focus scarce internal resources on the most important tasks within Intel IT and external organizations. Sometimes SaaS was used as a temporary solution while Intel developed in-house solutions. Cost was a factor, but usually a secondary one.

Overall, Intel employees using the SaaS solutions have reported good experiences, particularly with more mature SaaS providers and applications.

Security has been an important underlying consideration, which we have addressed with extensive, up-front due diligence and within our supplier contracts. Intel has high security standards and performs an extensive security audit when considering SaaS.

SaaS delivery models have ranged from subscription-based software to full business process outsourcing. The most successful implementations have been in areas that are self-contained and not core or differentiating to Intel's business. This outsourcing approach has enabled us to take advantage of suppliers' expertise.

The amount of data regularly transmitted between Intel and SaaS providers has been a huge challenge, causing difficulties during initial deployments and upgrades. Testing solutions has also provided challenges, demonstrating the need for full documentation and up-front clarification with suppliers about roles, responsibilities, and process.

Overall, SaaS has been successful in our environment and met Intel's expectations for the intended use of the services.

Infrastructure as a Service

Intel uses IaaS for certain niche applications. For example, some of the content on Intel's Web site is hosted by a cloud service provider. This allows us to take advantage of the supplier's worldwide infrastructure rather than facing the expense and difficulty of building similar infrastructure ourselves.

We also gained experience with IaaS when we built a globally distributed Web-monitoring application. Intel needed a service that would allow us to look at visitors' experiences accessing the Intel Web site from different regions of the globe.

We implemented the service using the distributed systems testbed provided by the PlanetLab global research network. We instantiated VMs on PlanetLab nodes in the geographies of interest and installed our monitoring application on those VMs. From an internal host, we pulled data through Intel's firewalls to an internal display host, which displayed the data graphically on an internal Web page. We automated VM monitoring, instantiation, provisioning, and failover through Web-accessible interfaces available on PlanetLab. Despite the lack of service-level guarantees for PlanetLab nodes, we managed to construct a robust global monitoring service that Intel network personnel continue to be use.

From our deployment of the monitoring application on PlanetLab, we learned that it is possible to create a reliable service from components that do not have SLAs if the service is constructed correctly.

We subsequently conducted a PoC to move our display host from inside Intel's network to a VM on a commercial IaaS provider. We purchased access to the provider's system using a credit card, downloaded the provider's tools, and used them to implement the display host on a VM. We found the cloud computing infrastructure was fairly easy to use: It took only a few hours to install the tools and make the VMs available.

We were initially concerned about the possibility of highly variable and unreliable performance. We found that the VM was highly available and that the file system and network metrics were quite stable. We were pleased to find that the time on the VM was kept reasonably accurate, in contrast to previous experiences with other VM implementations.

Our experience suggests that once security and manageability concerns are addressed, current commercial IaaS implementations may be good for rapid prototyping and compute-intensive batch jobs. After IaaS services prove themselves for these applications, they could be considered for more demanding applications with stringent response-time requirements.

The relative ease with which we set up our IaaS account, particularly the use of credit cards and Web interfaces for VM creation and maintenance, makes it clear that users can obtain cloud computing resources without a central IT organization's knowledge or permission. While such independent implementations can get off to a relatively quick start, they may find that they need to integrate their external cloud application with corporate data on the internal network. Unfortunately, more often than not, this integration requirement may run into insurmountable security concerns and result in abandoning the external cloud approach. The need to avoid these inefficiencies is another impetus for our cloud strategy efforts, and a reason that IT organizations need to stay ahead of users when developing a cloud computing strategy.

Key IT Architecture Considerations

A migration to external cloud computing involves significant changes to the role of an IT organization, supplier relationships, and the way applications are developed and used. From the perspective of an IT organization, there are a number of key areas to consider, including composite applications, standards, and external and internal clouds.

Composite Applications

Ultimately, our cloud computing integration architecture will need to support composite applications built from multiple services from external suppliers and internal IT sources. As shown in Figure 1, an Intel employee might access applications from several sources:

- A directly accessed Internet-located service
- An Internet-located service accessed indirectly through an address managed in an internal cloud
- A directly accessed intranet-located service

The second example illustrates the potential complexity. Using a service may invoke a chain of

sub-operations. These sub-operations may not all be within one supplier's cloud, but the supplier could make this transparent to users by making all the sub-operations appear to be serviced from a single invocation address. Similarly, the location of an external service could be masked by an address owned and managed by Intel.

The IT organization of the future will be responsible for identifying the right location of computing on multiple available platforms—some within the enterprise and some external. IT will evaluate different global cloud suppliers and assemble the right mix of internal and external offerings.




Figure 1. Access to cloud computing services in multiple locations.

Standards

For the foreseeable future, there will be multiple clouds. As a result, standards will be needed to enable these clouds to work as a single entity. Without common specifications for interfaces, protocols, and service announcements, each cloud service will have its own peculiarities of identification and access. For enterprises, this also introduces the risk of getting locked in to specific clouds through the use of proprietary application programming interfaces (APIs).

From an architectural standpoint, a single logical cloud that masks the complexity of different cloud-based offerings is highly desirable in order to minimize application design complexity. This requires developing and adopting foundation cloud computing standards for identity, authentication, federation, and encryption. We would like to see cloud providers converge on a common set of industry standards and differentiate their cloud offerings by the services rendered, including data security, data policy management, multiple levels of service, lower cost, reliability, data transfer bandwidth, low latency, auditing, and compliance.

Clouds must be able to report significant errors, especially faults that affect SLAs or legal obligations. They must be able to report their health through industry-standard interfaces.

External and Internal Clouds

Many technical and legal issues prevent broader enterprise adoption of external clouds. These issues are largely addressed if the cloud operates inside the enterprise, where there is greater control over the cloud. Because of this, an internal cloud is the ideal place to start proving cloud-related technologies and is a logical first step before attempting more widespread migration to an external cloud.

A large enterprise can gain many benefits from the greater abstraction of applications and

infrastructure that accompanies a migration to an internal cloud. Once standard interfaces and protocols exist and technical and legal obstacles have been overcome, IT organizations can start to make greater use of external cloud-based capabilities with minimal disruption to users, while reducing the data center footprint of their internal physical infrastructure.

This progression means that IT organizations need to balance three broad areas of computing while making the transition to the external cloud:

- Current, conventional computing
- Internal cloud
- External cloud

Conventional Computing

Conventional computing will continue to provide the enterprise with capabilities for many years, with a gradual migration of applications to internal and external clouds. Some conventional computing resources are likely to remain in long-term use, including those that need to be physically located on isolated segments or associated with specific hardware.

Internal Cloud

Internal clouds can have most of the features of external clouds. They can use similar technologies to host cloud-aware applications and to provide a dynamic infrastructure that responds to demand and fault signals. IT organizations can try out new chargeback billing methods; these also provide a benchmark for measuring the value of moving a service to external suppliers.

Internal clouds can act as a bridge to a future based on the external cloud. Applications can be developed to standards supported by both internal and external clouds, so that they may be readily migrated to an external cloud as necessary to support business strategy. It should be possible to move an application between

locations within the internal cloud without disruption to users. In the same way, it should ideally be possible to perform live migration of an application from an internal IaaS cloud to an external cloud without disruption to users.

Much of an enterprise's infrastructure could be serviced by a single internal cloud comprised of multiple physical data centers. The internal cloud could be logically and physically subdivided if necessary for business continuity or regulatory purposes.

External Cloud

Ultimately, external clouds will play a significant role in delivering conventional enterprise compute needs, but the internal cloud is expected to remain a critical part of the IT infrastructure for the foreseeable future. Key differentiating applications may never move completely out of the enterprise because of their mission-critical or business-sensitive nature.

High-Level Cloud Computing Strategy

Intel IT's evolving cloud computing strategy is based on growing the cloud from the inside out: building an internal cloud, then migrating to an external cloud as the market matures and security and privacy concerns are addressed. Meanwhile, we are already taking advantage of SaaS for specific applications where there are clear benefits.

Applications Suitable for External Clouds

Not all applications are suitable for external clouds today. Good candidates are applications that do not provide a competitive advantage, are not mission-critical, and are not tightly integrated with other important applications. To minimize security risks, they should not contain sensitive

information. In general, the set of applications deemed to have low security risk should grow over time as more sophisticated techniques to secure cloud-based applications are developed.

Our current view of the key decision-making criteria is summarized in Table 1. We expect to re-evaluate these criteria over time as the market matures.

Table 1. Applications Suitable for Cloud Computing

Typical Attributes of Applications Suitable for External Clouds	Additional Typical Attributes of Applications Suitable for Software as a Service
- Do not deliver competitive advantage - Are not mission-critical - Are not core business applications - Contain less-sensitive data - Are minimally affected by network latency or bandwidth	- Are at a natural re-engineering point in their lifecycles - Have minimal customization - Have industry-standard workflow




Figure 2. A possible scenario for adoption of software as a service (SaaS) and infrastructure as a service (IaaS).

Software as a Service and Infrastructure as a Service

The different characteristics of SaaS and IaaS could lead to significantly different adoption rates. One key factor is whether application re-engineering is required. With SaaS, migration of individual applications may require substantial re-engineering. Migration is therefore likely to occur at specific points in the application lifecycle or when influenced by factors such as mergers and acquisitions or specific business needs. As a result, we expect adoption of SaaS will grow steadily on an application-by-application basis.

In principle, with IaaS there would be little or no need for re-engineering. For example, we might simply move an application from a VM within our internal cloud to a VM in an external cloud. However, before we could do this, we would need to qualify the IaaS offering and position it as the preferred target for hosting a designated class or tier of applications—because of its superior cost profile or scalability, for example. In this scenario, no IaaS adoption could occur until an external cloud offering is qualified as suitable for hosting those applications. After qualification, rapid adoption could proceed as we migrate applications from the designated class or tier to an external IaaS cloud with minimal need for re-engineering.

One possible scenario comparing adoption rates for SaaS and IaaS is shown in Figure 2.

Cloud Computing Adoption

We are developing and beginning to implement a strategy to move from our current environment, through an internal cloud, to external cloud computing. Figure 3 contrasts our current environment with potential interim and future states as we make this transition.

Current: Grow Internal Cloud

The current Intel IT environment consists primarily of conventional computing. However, Intel IT has several ongoing initiatives with many of the characteristics of an internal cloud.

One of our major internal initiatives is data center virtualization (DCV). DCV is a design computing initiative that creates a pool of compute resources located across multiple sites. This lets us apply Intel's global computing resources to individual projects. Design

computing requirements are growing rapidly, and this initiative allows us to provide more compute capacity by increasing utilization of existing resources while reducing the need to add hardware. DCV resulted in significant cost avoidance during 2008.

Data center utility (DCU) is an enterprise computing initiative aimed at building a more agile, dynamic data center environment through virtualization. We are creating flexible pools of compute resources

based on newer, more powerful and energy-efficient servers. Virtualized workloads can be dynamically allocated and migrated between physical servers within these resource pools. DCU is in the early stages of deployment.

Other initiatives include development on demand (DoD), which allows developers to rapidly create virtual development environments hosted on existing servers rather than acquiring new hardware for each project.




Figure 3. Intel IT's high-level cloud computing strategy.

Intel has opportunistically adopted SaaS for some applications that are not mission-critical or core to our business. We have conducted some PoC projects based on IaaS, as well as using IaaS for specific niche applications.

Interim: Focus on Transforming Existing Environment to Internal Cloud

Over the next two years, we plan to focus on expanding our internal cloud environment. We will work with business units to migrate conventional computing services to these environments, while continuing to take advantage of SaaS for specific applications. If standards evolve and barriers such as security, manageability, and reliability are addressed, we may be able to move services to external clouds.

Future: Moving from Internal to External Cloud

As our internal initiatives begin to operate more like a single internal cloud that can scale based on demand, we plan to move a growing number of services to external clouds.

The impetus for this transformation will be standardization. We also envision middleware that allows any client to connect to any service, facilitating the migration of services between internal and external clouds without disruption to users. We will continue to move applications to SaaS and to adopt IaaS offerings such as storage and compute services over time. As external clouds grow in sophistication, they provide segmented services aimed at supporting differing user requirements and client devices.

Conclusion

Cloud computing promises significant benefits, but today there are security, privacy, and other barriers that prevent widespread enterprise adoption of an external cloud. In addition, the cost benefits for large enterprises have not yet been clearly demonstrated.

Intel IT's strategy focuses on growing the cloud from the inside out. As we partner with Intel's business groups to move services to an internal cloud, we see many of the benefits of cloud computing and are positioned to take

advantage of external clouds as supplier offerings mature and other barriers are overcome. We are opportunistically taking advantage of SaaS and gaining experience with IaaS as we prepare for this major transition.

Authors

Hong Li is a principal engineer with Intel IT.

Jeff Sedayao is an enterprise architect with Intel IT.

Jay Hahn-Steichen is a systems engineer with Intel IT.

Ed Jimison is a technology evangelist with Intel IT.


Catherine Spence is an enterprise architect with Intel IT.

Sudip Chahal is a principal engineer with Intel IT.

Acronyms

API	application programming interface
CRM	customer relationship management
DoD	development on demand
DCU	data center utility
DCV	data center virtualization
IaaS	infrastructure as a service
PaaS	platform as a service

PoC	proof of concept
ROI	return on investment
SaaS	software as a service
SLA	service-level agreement
SMB	small and medium-size business
TCO	total cost of ownership
VM	virtual machine



www.intel.com/IT

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

Copyright © 2009 Intel Corporation. All rights reserved.

Printed in USA
0109/KAR/KC/PDF

 Please Recycle
320566-001US