Defining a framework for cloud adoption

How common ground can help enterprises drive success with cloud computing

IBM

Contents

1 Introduction

3 Looking at cloud computing across two dimensions

3 Cloud delivery models

7 Cloud service types

8 Roles in cloud consumption and delivery

10 Challenges and considerations

11 Summary

Introduction

With the many widely discussed advantages to cloud computing, including elastic scaling, faster service provisioning, greater IT efficiencies and usage-based accounting, businesses are increasingly interested in adopting a cloud computing environment. Yet a swell in media and analyst coverage, plus a proliferation in cloud computing providers and marketing messages, has resulted in a great deal of confusion about what the term "cloud computing" really means.

For example, some providers look at cloud computing as way to provide compute or storage capacity as a service, provisioned from a parallel, on-demand processing platform that leverages economies of scale. Others may equate cloud computing with software as a service, a delivery model for making applications available over the Internet. IT analysts view cloud computing from the perspective of variable pricing without long-term commitments and massive elastic scaling of services. IT leaders look at cloud as an infrastructure architecture alternative that can reduce costs. End users, the media and financial analysts have still other perspectives on what cloud computing represents. Each group is discussing cloud computing, but few are discussing it in the same way.

This multiplicity of interpretations was confirmed in a 2009 IBM study of more than 1,000 IT and line of business decision makers around the world regarding perceptions on cloud computing.[1] The findings revealed that although 73 percent of respondents were familiar with cloud delivery methods prior to participating in the survey, there was little consistency among the terms that respondents associated with internal or external delivery methods. For example, 30 percent of respondents selected "software as a service," while only 24 percent selected "cloud computing." Other choices included "hosting" (19 percent), virtualization/consolidation (16 percent) and utility computing (4 percent).

Without a common vocabulary and a standardized frame of reference, it is difficult if not impossible for organizations to have a cogent discussion about cloud computing—externally with service providers, within the company, between IT and business leaders, or among the professionals within the IT organization. This alone creates a barrier to developing a successful enterprise adoption strategy for cloud computing.

To address this need, IBM has developed a cloud computing adoption framework, which establishes common definitions for cloud computing delivery models and services, illustrates the key capabilities to consider when developing cloud computing strategies, and identifies key aspects required to successfully execute that strategy.

Based on guidelines set by the National Institute of Standards and Technology, this cloud computing adoption framework is also an analytical aid that IBM consultants can use in combination with other IBM intellectual capital to help organizations establish a visual roadmap to cloud adoption. In this context, the framework can help organizations:

- Determine the types of services that can be offered through a cloud
- Identify the service management capabilities required to enable the chosen delivery model and service, such as metering and billing
- Recognize the challenges—both immediate and long-term—that should be considered prior to cloud implementation, such as those related to integration or governance.

The purpose of this paper is to introduce the cloud computing adoption framework from IBM and make it available for use by any organization looking for a standardized frame of reference for cloud computing discussions.

Key features of the cloud computing adoption framework

- Identifies two primary dimensions to be considered when developing a cloud computing strategy
- Identifies key differentiated capabilities along those two dimensions
- Identifies key competencies and considerations to successfully deliver and consume cloud services at the intersections defined by the differentiated points
- Aligns capabilities required with each stage of adoption

Looking at cloud computing across two dimensions

The cloud computing adoption framework defines the two intersecting dimensions that should be considered when developing a cloud computing strategy—the delivery model, and the type of service or services being delivered. When these two primary considerations are combined, a basic adoption framework for examining cloud implementation requirements and options emerges.

Understanding these requirements can help organizations select the optimum combination of delivery model and service type for each workload they want to deploy in a cloud environment. Workloads are the kind of work that an organization needs to accomplish. Each workload has characteristics that make it more appropriate for public or for private cloud delivery, and more appropriate for a specific service type.

Cloud delivery models

The first dimension to consider when developing a cloud strategy is the delivery model to be employed, which is the horizontal dimension of the framework (the x axis). There are two primary cloud delivery models: public and private. A third model, defined below, is a combination of both.

A private cloud is one in which both the consumer of cloud services and the provider of those services exist within the same enterprise. The ownership of the cloud assets resides within the same enterprise providing and consuming cloud services.

A public cloud is one in which the consumer of cloud services and the provider of cloud services exist in separate enterprises. The ownership of the assets used to deliver cloud services remains with the provider.

A hybrid cloud combines multiple elements of public and private cloud, including any combination of providers and consumers, and may also contain multiple service layers.

As shown in Figure 1, the cloud computing adoption framework also defines various subtypes within the private cloud and public cloud delivery models. These segmentations within public and private delivery models allow for a more targeted discussion of roles and responsibilities for both the provider and the consumer of cloud services, based on the specifics of a given scenario. The subtypes also carry an implication of the level of skill and sophistication necessary to successfully deliver and consume cloud services.

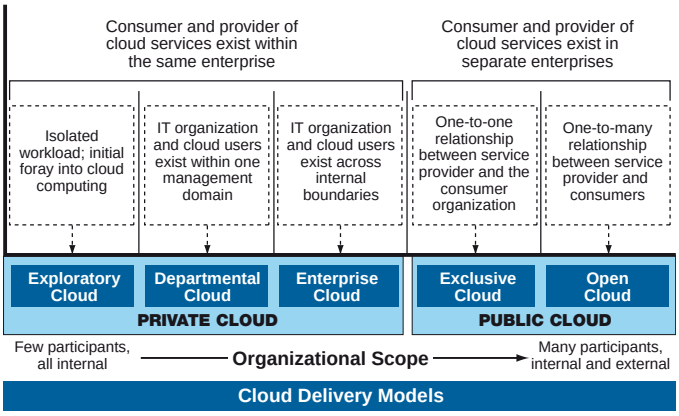

Figure 1: Cloud subtypes. The cloud computing adoption framework defines each delivery model subtype and helps organizations choose the most suitable for successful delivery.

Subtypes of private clouds

Exploratory cloud—This delivery model represents an organization's initial foray into cloud computing, in which the primary objective is to develop cloud delivery skills and experience. The goals in implementing an exploratory cloud—similar to a proof of concept—are to develop consumer and provider competencies and create awareness of unique cloud architectural and management system requirements.

Its value is derived in developing cloud service capabilities and an understanding of the technical nuances required in a cloud environment, as well as a starting point for establishing return on investment via cloud service delivery. An exploratory cloud requires that an organization have documented standards, be well managed and disciplined, have standards-based tools that can be developed or already exist, and have existing tools that are extensible. This type of cloud would likely be chosen to process compute-intensive, non-mission-critical workloads.

Departmental cloud—Both the IT organization supplying cloud services and the functional department consuming cloud services are within the same organization, with the department driving the requirements. In a departmental cloud, the goal is to expand use of cloud computing to consumers who are not familiar with cloud capabilities and to begin developing business support systems and operational support systems capabilities.

Its value is derived from enhanced shared resource competencies and virtualization management skills as well as a deeper understanding of the benefits of alternative delivery models, captured within the business support system. To implement a departmental cloud, the organization must have evolving open

standards as well as capabilities for data virtualization and basic cloud security. Workloads processed within this cloud might include test environments or department-specific applications.

Enterprise cloud—The IT organization supplying the cloud and the organization or organizations consuming cloud services are within the same enterprise, but exist within different management boundaries. The goal in an enterprise cloud is to dynamically solve time-critical business problems with existing, previously idle resources.

This model delivers value by optimizing the organization's investment in IT resource capacity. An enterprise cloud requires that security be both well defined and well implemented, and that the organization have additional service management capabilities, such as monitoring, automated provisioning and user authentication. Metering and billing, which enables IT to charge lines of business for cloud services consumed, is another key capability. E-mail and internal collaboration tools are some of the workloads that might be processed through this cloud.

For each of the three subtypes in a private cloud, both the business and operational support systems are within the same enterprise management boundary as the cloud services.

Subtypes of public clouds

Exclusive cloud—An exclusive cloud is most typically used to provide shared services to the members of a single group or organization, such as a consortium of universities or research institutions, a company and its trusted business partners or remote locations, or an industry or trade association. In an exclusive cloud, the goal is to provide access to trusted participants so they can use mission-critical applications. The known members consume services via their relationship with the umbrella organization, while the umbrella organization has the business relationship with the cloud provider.

There are two critical distinguishing characteristics of the exclusive public cloud. First, the cloud provider and the organization consuming cloud services are known to each other and are able to negotiate service-level parameters. This business relationship can extend beyond that of consumer and supplier—for example, the consumer could have a larger outsourcing agreement with the provider—but it is not required.

Second, any delivery resources used to provide services to the organization are dedicated to this exclusive cloud and are not shared with any other provider delivery environments or any cloud consumers outside of the organization. It is important to note that a "virtual private cloud" is not the same as an exclusive cloud. Rather, a virtual private cloud describes an arrangement where a single consuming organization—for example, a company or a department within a company—has exclusive use of the provider resources that deliver the consumer's services. These resources, however, reside in a data center or even on a server that provides services to many other consuming organizations.

An organization implementing an exclusive cloud can hope to reduce both capital and operating expenses for the consumer and to develop cloud partnering skills, such as securing consensus for changes to the business or operational support system changes, defining security issues and establishing cloud service level agreements. This cloud type requires differing levels of trust, regulations concerning the physical location of data, as well as equitable cost allocations. Providing capacity

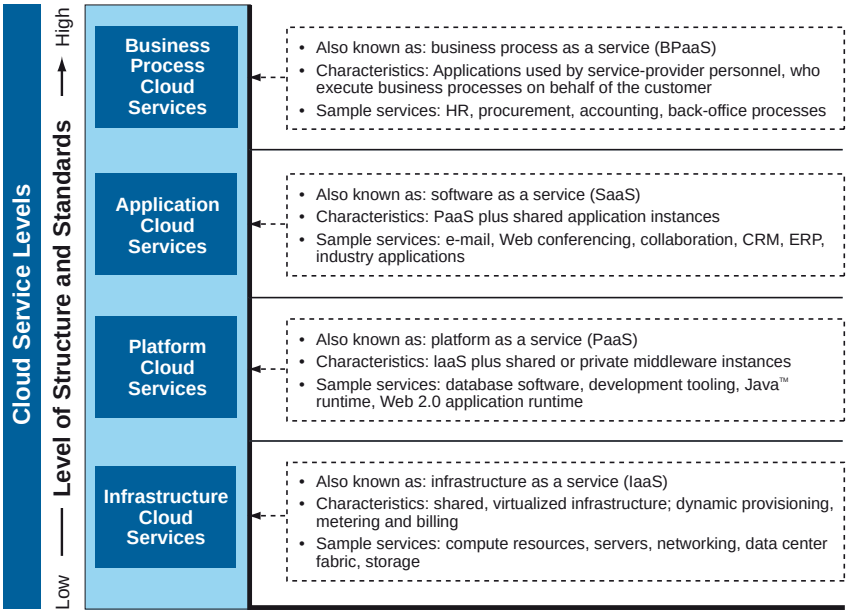

Figure 2: Cloud service types. Each type of service represents an increasing level of structure and standards, with business process cloud services requiring the most.

on demand for intermittent compute intensive workloads such as sophisticated modeling to multiple universities is one example of an appropriate use for an exclusive cloud. Other potential workloads could include industry-specific applications shared by trade association members or analytics applications shared by members of a research consortium.

Open cloud—The consuming and supplying organizations are unknown to each other prior to the presentation of a service request. The primary factor distinguishing an open cloud is that the negotiation of services must be an automated, standards-based event for which governance terms are defined and controlled by the provider. Furthermore, the open cloud delivery model requires automated negotiation and reconciliation of service levels, pricing and policies, as well as automated policy enforcement.

Services are typically ordered, fulfilled and paid for over the Internet, in many cases without human intervention unless the consumer initiates an inquiry or customer service transaction. This cloud type delivers value by eliminating the upfront capital investment that the consumer would have to make to

deliver comparable computing resources in-house and by speeding implementation time and realization of business benefits. While virtually any type of workload could be processed in an open cloud, IBM's analysis of barriers, drivers and consumer preferences[2] identified specific infrastructure, platform and application workloads as the most appropriate entry points for open public cloud delivery. In particular, the top ranking workloads for open cloud included Web conferencing, service help desks, virtual desktop applications such as customer relationship management, test environment infrastructure and storage capacity.[3]

Cloud service types

The second dimension defined by the cloud computing adoption framework is service type. This is the vertical dimension of the framework (the y axis). As shown in Figure 2, the four service types defined in the framework are layered to represent the increasing level of structure and standards. Each service type is built on, and requires the structure and standards of, the one below it. As discussed later in this paper, this progression has significant implications for both the provider and consumer of cloud services. In many cases, organizations will choose one or more service types, adding to the level of complexity in selecting cloud delivery methods for each—and underscoring the need for a visual map from which to examine all aspects of cloud delivery.

The cloud service types defined in the framework are:

Infrastructure cloud services (also known as infrastructure as a service, or IaaS)—These provide on-demand, pay-as-you-use access to infrastructure resources, including servers, storage or network devices, that the consumer configures and controls, running the applications of their choosing. The service can be delivered in a consumption-based business model—for example, by the instance-hour used or gigabyte transferred—or as a fixed fee for a virtual device with a predefined capacity and configuration. Either way, resources are accessed via the network, typically over the Internet. Infrastructure cloud services may also include an operating system with or without basic system management tools.

Platform cloud services (also known as platform as a service, or PaaS)—These services deliver compute capability (infrastructure) plus a predefined middleware stack that is typically structured for developers or advanced IT users. Providers can chose to offer a variety of service products (stacks), configurable to varying degrees by the consumer. Examples include database, Web or application server software. Configuration and management of these middleware resources are the responsibility of the consumer, but the provider may offer to maintain standard images once they are defined.

Application cloud services (also known as software as a service, or SaaS)—The service is a predefined application, such as CRM and ERP, which is typically delivered via a public cloud provider. Consumers from multiple organizations share a single application instance, with virtualization technologies employed to segregate customer data and maintain privacy. Application configuration and management are the responsibility of the service provider. However, an organization may

choose to implement a similar, application-based service within a private cloud to reduce costs licensing fees and other costs.

Business process cloud services (also known as business process as a service, or BPaaS)—This service combines application cloud services and the shared services model in which a single organization delivers business services, such as employee benefits management, help desk or procurement, to multiple internal or external consumers. Although consumers are purchasing business services, not applications, they can benefit from the reduced cost and increased flexibility that the provider realizes by employing a cloud-based service delivery infrastructure.

Once the delivery model and service type have been determined, organizations can begin to identify the service management capabilities required to enable the chosen delivery model and service. An example of this is metering and billing, in which the provider is able to charge the consumer for usage, and the consumer uses the service on a "pay-as-you-use" basis. In a private cloud environment, metering and billing enables the IT department to map IT costs to business units and recover those costs through internal accounting "chargebacks," should they choose to do so. At a minimum, metering should be done so the IT organization gains full understanding of the cost of "manufacturing" a particular IT service.

Roles in cloud consumption and delivery

Together, the two primary dimensions—delivery model and service type—establish a basic adoption framework for examining cloud implementation requirements. The cloud computing adoption framework also identifies the major roles involved in cloud service consumption and delivery and the competencies required for each role. The roles are:

Consumer—In private clouds, the consumers are within the same enterprise management boundary as the provider. In public clouds, the consumers and providers exist in different enterprises. The consumer role is further broken down into two subroles, user and subscriber. The competencies required for each role are:

- *Subscribers*—Must make financial and contractual commitments to consume cloud services and is responsible for making the user population aware of cloud services and entitlements. Subscribers also terminate the subscription.
- *Users*—Are responsible for productively using cloud services, as well as discontinuing them at task completion.

Providers—The provider supplies the cloud service and owns the assets required to produce and deliver cloud services to the consumer. In the private cloud model, internal providers deliver services to internal consumers. In a public cloud model, the provider is an external third party, and may provide services to multiple consumers or to a single consuming organization. Providers can also be part of a hybrid service delivery model, in which they supply both public and private cloud services, from different service types.

Integrators—Integrators provide the necessary level of IT literacy to communicate the consumer's IT requirements effectively to the provider. In a public cloud delivery model, the integrator role resides on the consumer side of the equation and retains accountability to the consuming organization for the cloud services provided by an external third party.

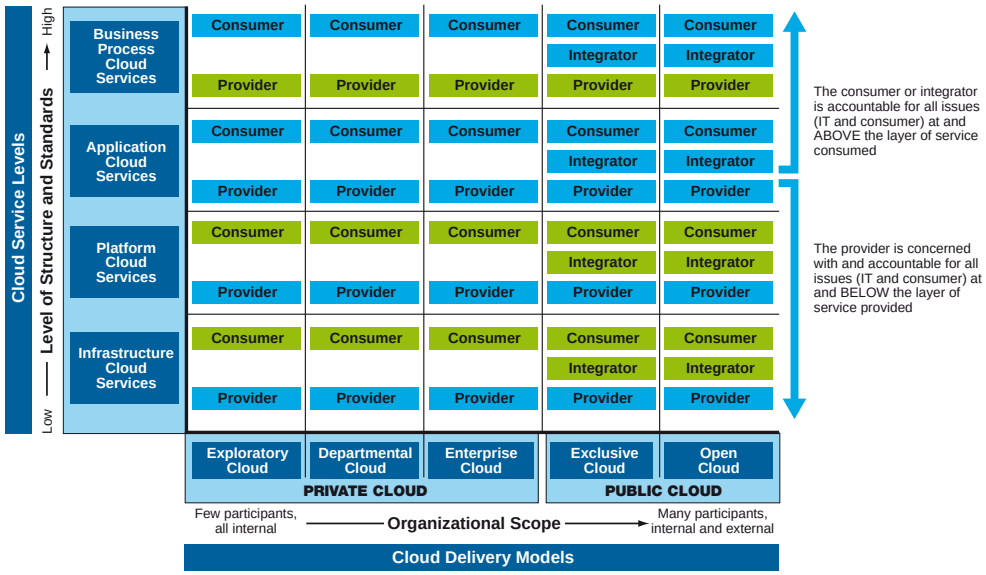

Figure 3: Roles and responsibilities. The cloud computing adoption framework defines the levels of responsibility for the provider, consumer and integrator roles in the public and private cloud models.

The cloud computing adoption framework clearly delineates the responsibilities that each role carries, depending on the chosen cloud service type and delivery model. As shown in Figure 3, collectively the consumer and provider are accountable for all issues. The specific responsibilities of each vary based on service type. Once a specific service layer is identified, consumers have the responsibility to communicate their requirements for that layer to the provider. However, the consumer retains responsibility for all IT-related concerns above that layer. For example, a consumer of platform cloud services is responsible for all of the applications and business processes that the organization requires.

Similarly, the provider is responsible for negotiating consumer requirements for a specific service type. In this case, however, the provider is solely responsible for all IT decisions at and below the selected service layer. Continuing the example, the provider of platform cloud services is responsible for both the middleware layer and the underlying infrastructure layer. An application cloud services provider is responsible for the application layer, the middleware layer and the infrastructure layer.

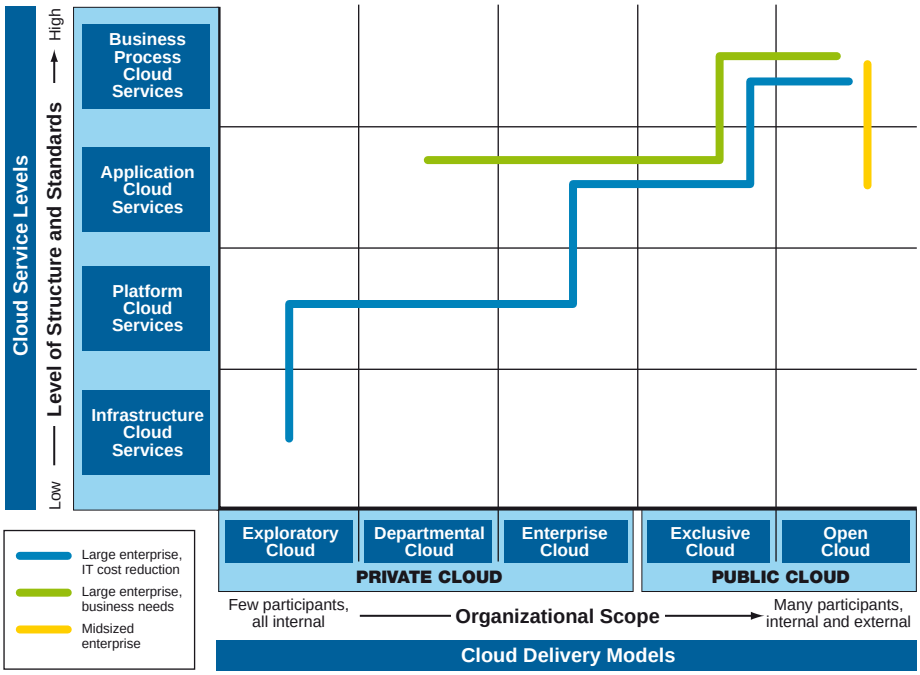

Figure 4: Visual map of potential cloud adoption pathways. Pathways will vary according to the chosen delivery model and service layer, as well whether the migration to cloud is driven by the business or IT.

Challenges and considerations

Once the type of cloud environment has been determined for a specific workload or set of consumer requirements, the cloud computing adoption framework helps organizations identify challenges—both immediate and long-term—that should be considered prior to cloud implementation, such as those related to integration or governance. For instance, in a departmental cloud, organizations must be able to establish cross-functional teams and to share both investments and usage. In addition, a key integration requirement is the ability to ensure data virtualization and data availability. In an open cloud, providers must be able to meet guarantees on service level agreements and resources and, from an integration standpoint, must be able to automate negotiation and reconciliation of service levels, pricing and policies.

Establishing a visual pathway to cloud adoption

Depending on the service type(s) and cloud delivery model subtype(s) that an organization selects, the cloud computing adoption framework provides a tool by which organizations can establish a visual roadmap to cloud adoption. As shown in Figure 4, pathways will vary according to the organization's IT strategy, its chosen delivery model and service or services being delivered, as well as whether the migration to cloud is driven by the business or IT.

Within a large enterprise, when cloud is driven by a strong business sponsor, the pathway to cloud typically begins with a business challenge or a new business opportunity that depends on some kind of enabling technology. The business justification may be built on the lower cost of acquiring a needed resource, but will more likely center on faster time to value and the business advantage that can be gained by implementing new capabilities in days or weeks rather than months. Service types selected will be higher in the stack, typically application or business process cloud services.

Conversely, in a large enterprise where IT is the primary driver, the pathway to cloud will likely be driven by the need for infrastructure optimization in order to provide services more dynamically as well as the need to reduce capital expenses by better leveraging existing infrastructure. In this scenario, implementations will begin with exploratory private cloud projects at the infrastructure level and move up the stack as skills are acquired. Starting with a private cloud implementation would also provide the opportunity for the IT organization to better understand cloud capabilities prior to moving selected workloads—such as Web conferencing or testing environments—to a public cloud.

In a midsized business, the pathway to cloud could be driven by the need to offload basic IT operations to reduce both capital and operating expenses and to decrease the level and sophistication of internal IT skills required. This would allow the retained IT organization to focus on the creation of business solutions while providing elastic scaling capabilities to cope with variable demand. Smaller firms may also utilize business process services from a public provider to reduce requirements for both IT and business process delivery skills.

Summary

As interest in cloud computing continues to gain momentum, there is increasing confusion about what cloud computing represents. Without a common, defined vocabulary and a standardized frame of reference, organizations cannot have a cogent discussion about cloud computing. The cloud computing adoption framework addresses this challenge by providing a context for productive discussion and a structure for planning, both short- and long-term, for a successful implementation.

For more information

To learn more about cloud computing or the cloud computing adoption framework, please contact your IBM representative, or visit the following Web site: **ibm.com**/cloud

IBM Corporation
New Orchard Road
Armonk, NY 10504
U.S.A.

Produced in the United States of America
May 2010

[1] IBM Market Insights, Cloud Computing Strategy Research, July 2009.

[2] Ibid.

[3] "Dispelling the vapor around cloud computing: drivers, barriers and considerations for public and private cloud adoption." IBM, January 2010. **ibm.com**/ibm/cloud/resources.html#5

Please Recycle

CIW03067-USEN-01