# IBM SmartCloud: Becoming a Cloud Service Provider

Redguides

for Business Leaders

Joe McIntyre

- Learn about the IBM CCRA Cloud Service Provider adoption pattern
- Understand the key cloud service provider business models
- Gain awareness of cloud service provider starting points

Redbooks

# Executive overview

Cloud computing opens a broad range of business opportunities across the computing industry and enables companies in other industries to provide services to their employees, customers, and partners. Cloud computing provides a compelling approach to addressing this opportunity.

From the competitiveness analysis that was done by IBM® Market Insights[1], the forecast of the public cloud service market opportunity is growing from approximately 30 billion US dollars in 2010, to over 88 billion US dollars by 2015. This forecast shows a 25% compound annual growth rate (CAGR) for cloud services. Of this total global market for public cloud services, the analysis shows that the Software-as-a-Service (SaaS) market and Infrastructure-as-a-Service (IaaS) market opportunities are the largest service segments. This forecast tracks with the focus of companies that are entering the cloud service provider marketplace. In addition, the market shows an emerging trend of providing industry-specific business applications as a service.

IBM identified different types of companies that can use the cloud service provider model to deliver cloud services:

- Service providers, such as communications service providers, internet service providers, and managed service providers.
- Businesses wanting to extend their applications and systems as services.
- Companies participating in ecosystems that are interested in providing applications and services to clients that are adopting cloud computing and accessing cloud services. These companies include independent software vendors, systems integrators, and resellers.

IBM believes that service providers are positioned to support cloud computing offerings. Cloud computing drives significant new revenue streams and strengthens the customer relationship by broadening their portfolio of services. The expanded portfolio can provide capabilities, such as communications, productivity, infrastructure, and business process services. The portfolio is delivered using key cloud computing capabilities, such as:

- Creating extensions to existing offerings using understood technologies
- Providing the ability to serve new markets efficiently
- Offering flexibility so the provider can create new services faster with lower cost (both in capital spending and operational efficiency)

[1] IBM SmartCloud™ Rethink IT. Reinvent Business. Chart 4, Rhonda Weaver, found at: http://www.ptc.org/ptc12/images/papers/upload/PTC12_Data%20Centers%20Wkshop_Rhonda%20Weaver.pdf

Companies rely on their business applications and systems as a part of their business. They can expand the business value of their applications and systems by using cloud computing to enable delivery of these functions as services. Companies have various options when they adopt cloud computing; they can:

- Use existing service providers to operate services on their behalf.
- Implement hybrid solutions that extend existing applications through integration with cloud services.
- Add cloud service hosting capability to their existing facilities.

For ecosystem partners, cloud computing provides compelling capabilities that ease deployment and long term management and maintenance. Equally important, cloud computing facilitates a more flexible business and technical environment. This environment can expand, contract, and adapt as services are added, removed, and evolve. The cloud replaces physical activity that is associated with change and change management by creating a fluid environment that adapts through automation.

Cloud computing looks simple from the outside. This façade understates the complexity of delivering its core characteristics of being reliable, extensible, and managed. To support these core characteristics for cloud service providers, IBM created the IBM Cloud Computing Reference Architecture (CCRA) cloud service provider adoption pattern.

**IBM CCRA:** The IBM CCRA and its adoption patterns are based on the lessons that were learned from thousands of cloud engagements. The adoption pattern embodies the architecture patterns that represent the ways organizations are implementing cloud computing solutions. An adoption pattern guides the definition of your cloud computing solution.

IBM offerings follow the IBM CCRA, including:

- Strategy and business case development services that are offered through IBM Global Business Services®, to define the roadmap for initial entry into the marketplace and plan for growth
- Ecosystem support through IBM PartnerWorld®, providing a ready source for services to deploy in your cloud computing environment or to access from an IBM SmartCloud facility
- The IBM SmartCloud for Service Providers portfolio that dramatically lowers the business and technical barriers of entry to cloud computing.

**IBM SmartCloud for Service Providers:** This portfolio of solutions provides software, hardware, and professional services that enable service providers to deploy pre-integrated cloud computing systems. The solutions include:

- IBM SmartCloud Integrated Infrastructure for Service Providers, which provides a ready to deploy commercial cloud computing environment
- IBM SmartCloud Aggregator for Service Providers, which enables cloud service providers to integrate and offer both internally hosted and third-party hosted cloud services

For more information about the IBM SmartCloud for Service Providers, see *IBM SmartCloud for Service Providers - Launch new cloud services simply and quickly to increase revenue*, found at:

`http://public.dhe.ibm.com/common/ssi/ecm/en/tls03047usen/TLS03047USEN.PDF`

The service provider adoption pattern provides a comprehensive and prescriptive architecture that defines the IBM view of commercial cloud deployments using standards and preferred practices. Companies embarking on cloud services or broadening the scope of cloud services can employ the adoption pattern as an accelerator for their architecture. By using this adoption pattern, cloud service providers can cover current and future cloud computing capabilities using a phased approach.

This IBM Redguide™ publication describes the business and technology choices companies make when they enter the cloud service provider space. It introduces various cloud service provider business models and shows how to apply them to your business. This guide introduces the IBM CCRA cloud service provider adoption pattern, providing guidance about the definition, architecture, and deployment of cloud computing environments. Two cloud service provider deployment scenarios are highlighted throughout the guide, and they reflect the two most common starting points for service providers that enter the cloud computing marketplace. The guide culminates with details about these deployment scenarios and showing how they can be deployed today.

# The cloud service provider business

Are you already a service provider? Do you want to add the capability to deliver services through the cloud? There are many compelling reasons to have a strong business and technical foundation when you enter the cloud service provider business. A comprehensive foundation helps you establish and expand a cloud services business.

## Cloud computing is important to service providers

Cloud computing is important to a range of existing service providers, including communications service provider (CSP), Internet service provider (ISP), and managed service provider (MSP) businesses. It is essential for service providers to understand how to effectively integrate cloud services to maintain and grow their business.

A CSP or ISP focused only on connectivity can include cloud computing as an additional set of services that operate over their network. Often, cloud services are general purpose and best effort services that use bandwidth from a fixed charge connection or consumption based allocation. Cloud services can also be more specialized services with discrete charges based on functions that are used or additional quality of service features.

Most CSP or ISP interests extend beyond just connectivity and include offering services. For these businesses and MSPs, cloud services are a new competitive category of services that use connectivity networks transparently. An example of this transparency is the delivery of the same cloud services that are operated by the CSP, ISP, or MSP to many different devices using the CSP or ISP only for transport to each device.

**IBM Institute of Business Value report:** The IBM Institute of Business Value executive report *The natural fit of Cloud with Telecommunications* is written for the CSP marketplace. This information is also relevant to ISP and MSP. The report provides a comprehensive business view of the marketplace and the current offerings available. For access to this report, go to:

`http://www.ibm.com/services/us/gbs/thoughtleadership/ibv-telecom-cloud.html`

The competitive business considerations include branding, differentiation, trustworthiness, efficiency, and experience/expertise:

- Brand value is a key element in the overall competitive profile of a provider.

  The objective of brand is to retain and strengthen the service providers position even when other cloud service providers emerge. The provider's brand is bolstered by the availability of both the correct quantity and quality of the cloud services offered. Exclusivity, breadth of recognized services, and timely availability of new services enhance the perception of a market leading brand.

- Offering differentiation is important.

  Many cloud service providers have similar cloud service offerings. These offerings include services such as webmail and maps. Differentiation for the common offerings is based on customer experience in activities such as ease of product selection, ordering, and usage. Beyond the commonly available services, differentiation is enhanced by inclusion of interesting niche products and vertical market enablement. Identifying market needs and demographics that match the provider goals for initial deployment and growth are key elements for successfully realizing differentiation.

- Efficiency of delivery.

  Cost and operational efficiency are key considerations in both initial and long-term competitiveness. Customers evaluate their return on investment (ROI) for the use of cloud services relative to other solutions, including on-premise solutions. Automation is a significant factor in delivering consistency in the cost to deliver cloud services and to realize efficiencies of scale.

- Customer trust is a key element of overall competitiveness.

  Cloud services are available from many venues. Customers consider trust a significant factor in selecting a provider, especially when the customers are generating financial transactions. Cloud computing deployments might be a complex environment, with third-party payment providers, and layered service providers. The cloud service providers that manage this environment in a trustworthy manner earn the trust of their customers while they enable the flexibility of operating an aggregate business.

- Experience and expertise in cloud computing sets a provider apart from others.

  Cloud computing is a business of scale with most initial deployments as targeted offerings. The typical business model envisions service provider scale deployment and transaction levels. To be positioned well a service, providers should provide these capabilities:

  - Ensure access control (secure user access).
  - Offer a breadth of compelling services in each market segment that is selected to compete in.
  - Provide an end-to-end delivery experience for users of cloud services.

  The management infrastructure of the deployment environment and support systems are key elements of supporting this scale of deployment.

These considerations become even more important when new entrants into the marketplace introduce additional competitive pressures. The competitive pressures include vertical market specializations and adjacencies to business related to the cloud services that are provided.

## Importance of cloud computing to the services ecosystem

Cloud service providers are not the primary creators of cloud services. An ecosystem of businesses and individuals are the source for cloud services. The ecosystem extends the reach of cloud service providers by expanding the breadth of services that are offered and addressing niche or specialty markets and geographies.

The business models that are developed on these ecosystems rely on a combination of content providers, such as:

- Application developers that create applications and services
- Content creators that generate items that range from ringtones to movies

These ecosystems support various intermediary businesses, such as integrators, resellers, hosting companies, and merchants (online or physical).

Independent software vendors (ISVs) are entering the cloud marketplace using virtualization and cloud deployment models to enable product offerings as a service. To enable this new channel, many ISVs spend considerable time and resources developing services on several platforms. This approach can be cumbersome. What the ISVs want and need is a consistent architectural approach for offering their product as cloud services.

The business case to develop and maintain a cloud offering drives key business decisions that must be answered. For independent software vendors (ISVs), the related questions include:

- What are the revenue expectations, and over what term? Are there opportunities to cross-sell additional services or up-sell from basic to advanced services?
- What is the development platform and APIs needed to support deployment within an aggregated client environment?
- Which business models should you support? Which providers should you partner with to support those business models?
- Which cloud service providers should you work with to ensure market leading exposure while also having low costs to enable an optimal return on investment?
- What is the cost of maintaining and upgrading the software extensions, including multi-channel approaches, which may mix on-premise and cloud delivered versions of products?

Systems integrators, resellers, and other companies that participate in the channel share perspectives from both the ISV and the cloud service provider. In many cases, these companies are the enablers for ISVs to deliver services to specialized marketplaces using business models crafted to individual market segments.

## Cloud computing is essential to many businesses

Cloud services are not limited to enterprise data centers and public Internet company scenarios. Many businesses already provide interactive services as part of their primary business and have a web presence that enables interactions with their employees, customers, and partners. These businesses can realize benefits by adopting cloud computing for the service providing portion of their business.

Internally, companies with affiliated businesses or multiple divisions can use cloud services to enhance their business-to-business (B2B) services. Flexibility in scale, operational efficiency, and consistency of deployment can improve the cost structure of operations.

Externally, businesses with services as adjacent capabilities to their core business can take advantage of cloud computing. Cloud computing offers a rapid market entry capability either as an extension of the current business or through a partnership with an existing cloud operator. For example, some consumer electronics companies provide content as additional services in an adjacent capability. The key business drivers for providing these adjacent services include:

- Providing a consistent end-to-end experience to the customer. This approach combines both the traditional product delivery along with supplemental services enabled by cloud.
- Enabling rapid introduction and responsiveness to market opportunities and competitive threats through streamlined introduction and updates to services.
- Providing the ability to expand and contract the resources that are allocated for services. This approach includes the option to use external resources, enabling efficient delivery without a lengthy skills and facilities build. In the case where external cloud resources are used, you have the option to not build at all.

When the business considers the spectrum of interactions with its employees, customers, and partners, the breadth of opportunities to use cloud capabilities as part of the overall business strategy emerges.

## Monetary drivers

For businesses with an interest in providing cloud services, monetary and non-monetary factors exist that influence the decision. Monetary motivation requires determining the value of cloud services to the business. This value may be determined according to various factors:

- A ROI calculation that determines the revenue gain, cost reduction, or profitability improvement of using cloud computing to deliver a service versus its alternatives.
- A capital investment versus operational expense comparison that is based on immediate investment requirements and ongoing expense. The expense is based on modeling deployment approaches to different cloud deployment options versus their alternatives.
- Opportunity cost determinations that are based on adopting cloud computing, retaining existing capabilities, or comparison to alternatives.

When you consider these factors, the types of services must also be considered:

- Infrastructure-as-a-Service (IaaS)

  IaaS business models typically have a commodity model (for example, metered resource consumption) and a value add model (for example, choice of hypervisor, operating system, and base system stack). IaaS delivers compute, storage, and network resources from the cloud service provider to customers or partners to operate applications and access those applications as services. The pricing model is typically based on a cost per unit of resources used.

- Software-as-a-service (SaaS)

  SaaS application charging is often based on user value (for example, monthly per seat charge). SaaS applications run in the cloud and are commonly accessed by users through a web interface. The pricing model is typically based on a subscription model and can also include charges for functions used.

- Platform-as-a-Service (PaaS)

  Platform-as-a-Service is usually an extension of the IaaS model, where charges are applied based on the resources used. Because PaaS includes additional services (development tools and services) and platform support (for example, application server or database), the cost per resource unit that is used is typically higher. This higher charge per unit that is used allows coverage of acquisition, license, and management costs.

## Market opportunity for cloud service providers

Each cloud service provider has a perspective on the market space that fits them and their existing and potential customers. Key characteristics that determine the market opportunity space include:

- Does the provider have an interest in providing applications as services? If so, do they want to provide the cloud service directly or in partnership with third parties?
- Does the provider want to offer cloud services complementary to other types of services (such as connectivity)? If so, are the services offered exclusively or bundled?
- Are the services limited (by choice or by regulation) to specific geographies?
- What are the market choices? Do the markets include consumers, small/medium businesses, enterprises, or vertical markets?
- Who are the competitors? What markets do the competitors choose to compete in? What specific segments should you enter? Who are the competitors already in or entering those segments?

Enabling customers to purchase cloud resources to run their own applications is a common IaaS opportunity. Common cloud resources include compute, storage, and network. IaaS has a high volume of potential customers, global applicability, and a mix of large and medium providers that offer general purpose offerings.

## Competitive landscape

Cloud computing opens new deployment models for businesses to become cloud service providers. The types of businesses and their motivations define a competitive landscape. Companies entering the cloud service provider business can compete in various categories, including:

- Pure play cloud providers that created their businesses and deployments purely on cloud approaches and technologies.
- Companies that built cloud “like” capabilities within a different business context and entered the cloud provider business as an additional business.
- Companies with a hosting and managed services business, extending the business to incorporate cloud offerings.
- Communications Service Providers adding cloud to their services offerings.

Some companies are entering the cloud business as an extension of their existing business to maintain competitiveness with their peers. Other companies are creating exclusive offerings that enable differentiation from their peers or to create distinct marketplace niches.

Cloud services and solutions are a primary source of new services that can be used by the general marketplace. All service providers require a strong and comprehensive cloud services offering to maintain relevance and competitiveness. IBM delivers the business and technical foundation, along with turnkey cloud solutions, to enable rapid entry and sustained competitiveness to this emerging marketplace. To support time-to-market considerations, IBM offerings include:

- On-premise deployment of solutions to enable companies to deploy and manage their own cloud computing environments. The IBM SmartCloud for Service Providers portfolio delivers turnkey and custom configurations.
- White label cloud services from IBM SmartCloud Enterprise allow rapid market entry with cloud services that are already operational with your branding.
- Aggregation solutions, combining access to on-premise and off-premise cloud services, enabling rapid delivery of services from many sources through a single service provider.
- Build and operate and build, operate, and transfer capabilities for service providers that want to get to market quickly, but have choices about where to operate over time.

## Entering the cloud computing marketplace

With the business and operational foundation that is provided by the cloud service provider adoption pattern, the next step is selection of service models. There are many existing service models, with IaaS, SaaS, and PaaS recognized as the three base service models. An adoption pattern supports each of these service models. Adoption patterns are flexible enough to support new service models that emerge in the future.

IBM, through its CCRA adoption pattern, observed consistency across cloud service providers in the two most common service models that are selected as the entry points into the cloud marketplace:

- Entry as an IaaS provider to establish a base offering that is familiar to customers and extends existing skills and capabilities of the provider.
- Entry as a SaaS provider to establish a market position that is focused on services in a position favorable for being the primary source for cloud services.

In addition, there is significant consistency in the evolution of these initial deployments. The consistency is seen in the scale and function of the initial deployment and in the broadening of business models that are addressed after the initial deployment. Evolution to include PaaS capabilities is accelerating as cloud service providers look to accelerate the availability of services. Evolution to include PaaS capabilities is accelerating as cloud service providers look to accelerate the availability of services. The IBM CCRA architecture and related solutions reflect these entry points and evolutionary paths that enable service providers to deploy and grow on a proven foundation.

Figure 1 shows the entry points and evolution over time to incorporate more service models, along with the capabilities that enhance the overall solution.

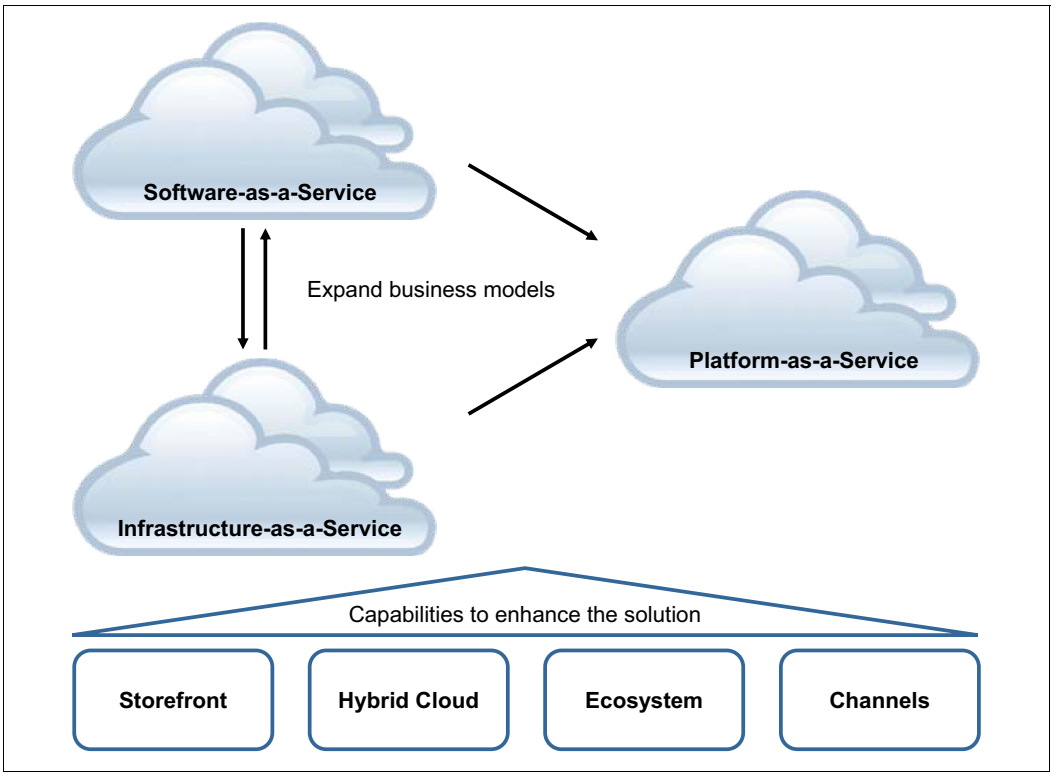


*Figure 1 Cloud computing entry points and evolutionary capabilities*

It is common for the two entry points IaaS and SaaS (Figure 1) to follow one another, with the order decided by the selection of initial marketplace. It is also common to establish a roadmap that supports PaaS, complementing the entry point chosen.

In addition to the types of services, the roadmap contains functional capabilities, including:

- A store front for self service and support for many services.
- Hybrid cloud capability that extends the reach of the provider to encompass a broader range of business customers.
- Ecosystem enablement to support third parties that populate the provider cloud with cloud services.
- Channel enablement to support the sale of cloud services through retail and direct sales channels.

## Two cloud service provider scenarios

Two independent cloud service provider scenarios are used in this guide to explain how IBM applied the CCRA service provider adoption pattern to cloud computing projects:

- IaaS entry point scenario

  In this scenario, the service provider serves a metropolitan city area, with an emphasis on small businesses. The cloud services that are offered enable businesses to run their applications in the cloud, have automatic backup of their data, and are charged only for resources they use.

- SaaS entry point scenario

  In this scenario, the service provider is expanding their mobile business to include cloud services. The distinguishing characteristics of the business include:

  - A high percentage of customers uses only a mobile device (smartphone or tablet) to use cloud services and manage their account.
  - The primary interactive cloud services are productivity and communications oriented.
  - Device management and personal storage (photos and messages) are delivered as cloud services.

This guide uses these scenarios to illustrate the usage of the concepts.

# General cloud service provider business models

The cloud computing marketplace is a heterogeneous environment in which many companies participate in collaborative ways to deliver cloud services. This collaboration often includes the cloud service customers themselves. This flexibility provides many opportunities for business model innovation. This collaboration also creates some complexity in understanding the roles of each participant according to the specific business model that is realized. To ease understanding, the roles are introduced as follows:

- A *customer* is the person or organization that purchases the cloud services. The customer has a business relationship with the cloud service provider defined in the business contract terms and conditions. Users that are authorized by the customer use the cloud service.
- A *cloud service provider* is the organization that makes cloud services available. The provider manages contracts with customers, partners, and other providers that are related to the delivery of the cloud services that are made available.
- A *partner* is a person or organization that works with a cloud service provider to make cloud services available. Partners include:
  - Independent software vendors (ISVs) that create applications that are offered as cloud services
  - Channel partners that assist customers in establishing accounts and ordering services

There are various business models available for cloud service providers, each of which may be applied a stand-alone model or combined. This combination of business models can be built over time. These models enable cloud service providers to start with an initial business model and incorporate additional models over time.

The business models have a range of individual characteristics. Models can contain capabilities that are implemented by the provider themselves and capabilities added through partnerships. A cloud service provider can choose from various models, including:

- Operating their own data centers to offer the services.
- Integrating with other provider clouds, enabling offering services from those clouds to their customers without hosting them directly. Many content aggregators, which are fairly common today, offer their services in this way. This same approach is being applied to cloud services.
- Offering their services through other providers, for example, as a white label provider.
- Creating their own services or acquiring services that they host directly.

- Offering cloud hosting services to third-party developers, enabling the third parties to offer their services through the provider. The provider may optionally provide development platform capability from their cloud.
- Enabling customers to use the provider cloud to host their applications (in whole or in part).

**Choices:** The choices and combination of choices can be realized in any of these business models. The possible business models are extensive, but the list is not meant to be exhaustive. Also, consider that variations of these models can be created by using different combinations of these choices or introducing more choices.

Figure 2 provides a general model that is referenced for the business models that are described in this section.

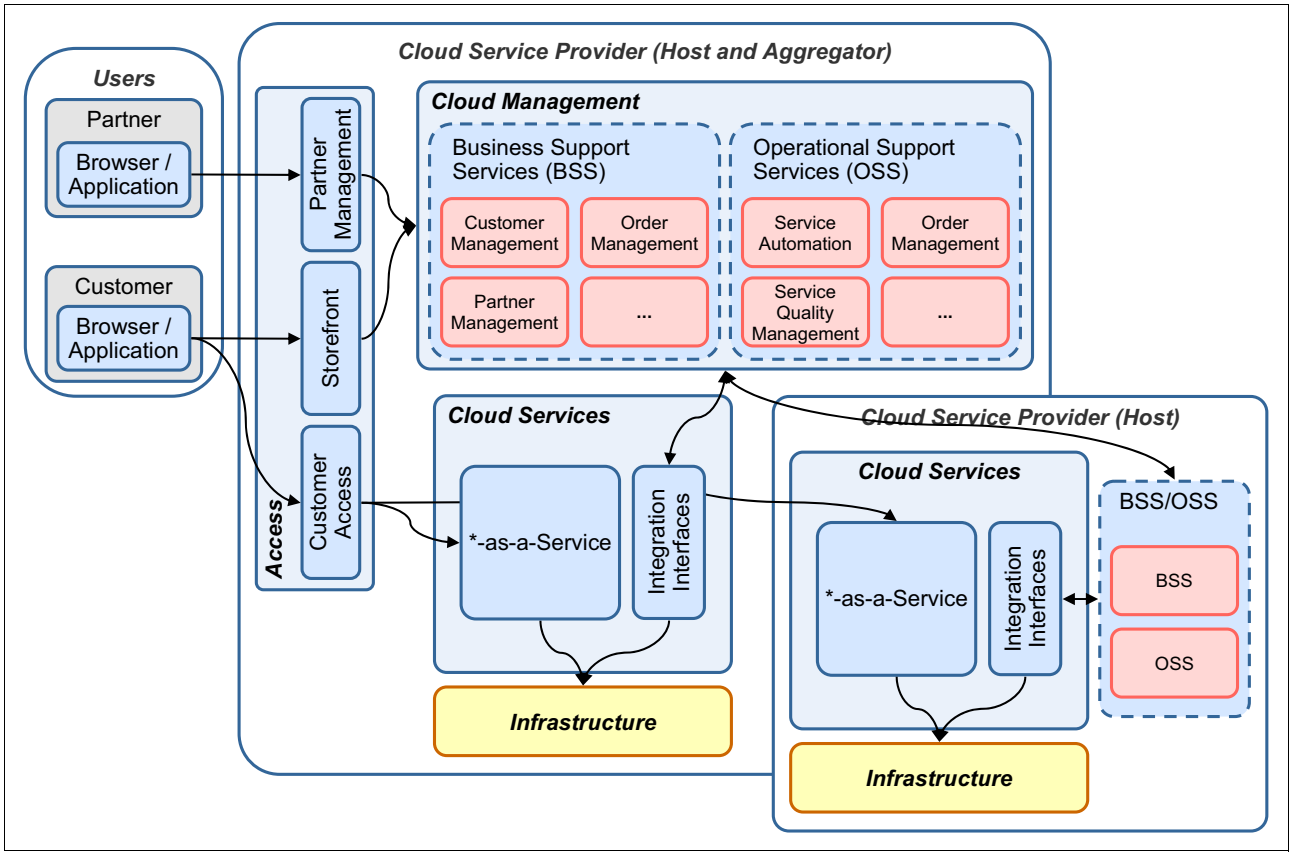


*Figure 2 General model for the business of cloud service providers*

Figure 2 depicts a set of relationships, such as:

- Between the users that access the cloud services from the cloud that is operated by the cloud service provider
- Customers and partners that populate the cloud that is operated by the cloud service provider with cloud services
- Multiple providers in a provider-to-provider relationship that exists for aggregated models

The relationships do not have technology dependencies and various technologies are applicable, including:

- User devices (computers, smartphones, and tablets)
- Machine-to-machine devices (smart meters and home automation)
- Networks (mobile, broadband, public, and private)

The general model (Figure 2 on page 11) identifies:

- The users of the cloud services that are grouped within a domain of users.
- The cloud service provider that interacts directly with the user. The provider acts as the host for the cloud services that are used or aggregates cloud services from one or more other providers.
- The systems that are used by the cloud service provider that enable access to the cloud and its services. The systems also include the support systems that manage the business and operational aspects.

For the IaaS entry point scenario, the customer application hosting is frequently the initial choice. This choice enables the cloud service provider to install and manage the environment. The provider presents customers with simple configuration choices for the initial implementation of products, order processing, and billing. This environment typically can be up and running using an IBM pre-packaged solution configuration in a few months. The deployment includes integration with existing billing functions to provide a single bill capability.

For the SaaS entry point scenario, one of two choices is most often selected.

- Aggregation, including the choice to white label

  This choice allows the provider to enter the market quickly, because the cloud services are already operating and the provider is able to focus on the support systems functions to enable customer management, product definition, ordering, and billing. Either an IBM white label solution or an IBM aggregation solution can be deployed in a few months, typically starting with 4 - 10 services.

- Provider application hosting

  Provider application hosting is usually a stepping stone to supporting a third-party hosting model. This choice is driven by the provider's desire to operate the cloud computing environment and to derive revenues from either reselling or hosting applications for use by their customers. Typically, licensing of the initial applications is done to enable fast time to market, while building an ecosystem capability to onboard new applications. As onboarding capability is deployed, the third-party hosting model emerges. An IBM solution that supports this choice typically includes an initial deployment that is accomplished in a few months and a parallel ecosystem enablement activity that spans 6 - 9 months.

Many business models are a simplified version of Figure 2 on page 11. A specific business model can arise from this context. Specific business models and the associated deployment approaches that are described in this guide are:

- Customer application hosting
- Hybrid clouds
- Provider application hosting
- Third-party models

## Customer application hosting

Customers that are interested in operating their applications in the cloud (in whole or in part) can engage a cloud service provider. The customer specifies their resource requirements to the provider for operating their applications in the cloud. The provider provides the environment and the onboarding facilities to enable the customer to load, configure, maintain, and operate their applications and make them accessible to users authorized by the customer.

IaaS often supports this business model. The cloud service provider offers resources such as compute, storage, and network, for use by an external organization. The customer has an application to be run on the cloud.

A cloud service provider can also offer a PaaS offering that enables the customer to deliver their applications using application servers or similar platforms. The PaaS offering consists of the base stack and configuration to support the customer application.

In this business model (IaaS or PaaS), the customer pays the provider for the use of the resources (whether dedicated or on-demand). The customer is responsible for onboarding the application and initiating its execution.

From the provider perspective, the resources are offered as base configurations with options (IaaS) or with a predefined starting stack (PaaS). Base configurations typically consist of these items:

- Resource profile for compute power (processor count and speed, and memory)
- Storage for use by the virtual machine (VM)
- Internal network between the compute and storage
- Hypervisor and operating system, on which the customer can build their application environment
- External network port to allow customer interaction with the VM (optional)
- Additional options include discrete storage and additional networking options, such as a Virtual Private Network (VPN)

From the customer perspective, they are purchasing resources and predefined stacks that are charged based on their actual use. The customer selects the resources and stacks and creates the configuration for the resources. Next, the customer onboards the application.

**Onboarding:** Onboarding is the process of uploading the image or application content and setting its configuration.

When the customer uses the application, the resource consumption is reflected in the charges generated. These charges are applied to the customer's account using their chosen payment method.

Regarding Figure 2 on page 11, the customer acts as both the developer and the user. The cloud service provider provides the hosting function for the application.

## Hybrid cloud

Hybrid cloud is a linking of a customer application and a cloud service together to provide a complete application. The model has two parts:

- The customer operates the parts of the application that run at the customer premise.
- The cloud service provider operates the parts of the application that run within the cloud.

The separation of the application into parts is specific to each application. This cooperative approach enables applications with data sensitivity to isolate the data that is restricted to use on-premise. Also, the customer can isolate processing functions to those functions that appropriate for each environment and similar considerations.

There are two types of hybrid cloud scenarios:

- Cloud bursting

  Cloud bursting identifies parts of applications that can be isolated and run independently. The location the applications operate in is either in-house or in the cloud. The location is determined at the time the function is to be performed. This approach allows the customer to *burst* out to a cloud service provider in either of these situations:

  - They do not have sufficient capacity.
  - It is economically advantageous to use cloud resources versus in-house resources.

- Cloud integration

  Cloud integration separates an application into multiple elements that operate cooperatively. One part operates in-house and the other part is operated by the cloud service provider.

Film production is an example of using hybrid cloud for cloud bursting. Parts of the film production process require large amounts of computer processing power. This demand occurs only for a relatively short time period over the course of the film production process. The production company can identify these compute intensive tasks within their processes and run these tasks in a cloud. The resulting process is a hybrid, with processing split as follows:

- General processes are performed in-house.
- Specific compute intensive tasks run in the cloud.

In either case, all activity is managed by the in-house processes. This example shows the value of using cloud to shift costs from capital expenditures to operational expenditures. The example shows how to limit capital spending by purchasing only equipment that is consistently in use.

Hybrid cloud can also be used as a staging approach, where portions of applications and processes are ready to move to cloud on different time schedules. This approach enables the workload distribution to evolve over time, enabling the advantages of cloud computing to be realized in stages. The staging approach is preferred to delaying the benefits until all portions of the application or processes are migrated.

Regarding Figure 2 on page 11, the customer acts as both the developer and the user. The provider provides only the hosting function for the parts of the application operating in the cloud. In addition, the customer hosts a portion of their application in their environment. The customer application interacts with the cloud provider through the customer access interface for the cloud portion of their application.

## Provider application hosting

From the provider point of view, the business model covers services that are packaged, offered, and operated by the provider. This model typically uses a usage-based charging model that generates charges for the use of the cloud service. The usage-based charging uses product specific pricing that is set by the provider. This model enables value to be distinguished for each product offered.

PaaS also fits in to this model, where the pricing for the base offering is set by the provider that operates the platform. When the customer builds their specific application using the platform, the application is operated by the provider. Pricing is set by either resource consumption or the platform pricing (decided by the provider).

What distinguishes this business model from the third-party model is the owner of the offering. In this business model, the provider owns these items:

- Packaging
- Product and offering definitions and the application
- Related content (which are either owned or licensed)

Regarding Figure 2 on page 11, the provider creates their own services, or acquires services to host. The customer acts only as a user, and the provider acts only as a host.

## Third-party models

In third-party business models, a third party provides the content (such as the application or platform), which is cataloged and offered by the provider. The customer interacts with the provider to establish an account, interact with the product catalog, create subscriptions, and be charged for usage. These activities are independent of the manner in which the application is delivered.

In the third-party model, the provider acts as the marketing, charging, and relationship management organization. The provider is the first point of contact for care, and accepts payments from the customer. The provider is responsible for:

- Providing onboarding capabilities for both customers and partners
- Managing the product catalog and enabling limited third-party management of products
- Managing third-party contracts and sharing the revenue collected based on those contracts

Today's various application store implementations are examples of third-party models. The cloud computing marketplace shares many of the motivations and operational characteristics of application stores.

The third-party business model has two related submodels, hosting and aggregation, which can be deployed as stand-alone models or together:

- Hosting covers the cloud services that operate in the cloud computing environment that is deployed by the cloud service provider.
- Aggregation presents services from the cloud service provider, but delivers the cloud services from another provider's cloud. Using this model, the provider includes third-party applications in the product catalog, but does not host the application on the provider cloud.

The four third-party models are:

- Pure hosting: The provider hosts all the cloud services offered.
- Pure aggregation: The provider does not host any cloud services; it acts purely in the aggregation role with all cloud services hosted by other cloud service providers.
- White label: A special case of aggregation where the provider rebrands the interfaces and cloud services from another cloud service provider.
- Mixed: A combination of these models where the cloud services provider hosts some cloud services and aggregates other cloud services.

## Pure hosting

In the pure hosting model, the provider hosts all the cloud services offered.

In this model, the provider operates the cloud from which all the services are offered. The third party enters a contract with the provider that specifies the terms and conditions under which the services are made available. The provider also specifies the manner in which payments are calculated. This contract can include fixed payments (either direction), revenue sharing, advertising, or other metrics. The key differences for this model and models where the provider owns the content directly are:

- The service and its related content are owned by entities other than the provider.
- A contract exists between the parties that governs the usage and commercial terms.

Any type of service may be covered under this model. The customer follows appropriate onboarding, product management, offerings, and operations that are suited to the type of service.

When a customer subscribes to a third-party service, the provider is responsible for various activities, including:

- Managing the runtime environment
- Provide access to the service
- Monitoring and metering usage
- Generating charge records and transactions

This model is a common starting point. For many cloud service providers, this model is their primary business model, even if they include other models in their overall offerings. This model is selected because it provides the content that is required to attract customers, while it provides an opportunity for differentiation in the marketplace for their core services capabilities.

Regarding Figure 2 on page 11, the partner provides the application to the cloud service provider. The customer acts as the user and the cloud service provider acts only as the host.

## Pure aggregation

In a pure aggregation model, the provider does not host any cloud services. It acts purely in the aggregation role with all cloud services hosted on external clouds.

For cloud providers using this model exclusively, their business often is related to marketing, advertising, or other lines of business or branding. These businesses apply their distinct value by having the ability to extend beyond just the set of services aggregated. Their business may include vertical market integration or other value-add. They can also provide particular pricing and loyalty models for specific market niches.

Unlike the pure hosting model, the delivery of the cloud services is divided between two cloud service providers, where each cloud service provider provides specific capabilities. The cloud services provider that acts as the aggregator provides:

- Customer management
- Product catalog
- Order capture
- Consolidated monitoring, charging, and billing (or equivalent)

The cloud services provider that acts as the host provides:

- Content for the product catalog
- Hosting of the service

An interface between the aggregator and host cloud service providers supports:

- Provisioning of organizations, users, and subscriptions
- Exchange of monitoring and usage information
- Exchange of revenue sharing information

Regarding Figure 2 on page 11, the cloud service is hosted on the external cloud (lower right of Figure 2 on page 11), but offered through the cloud service provider in the center of Figure 2 on page 11. The customer acts as the user and the provider that is adjacent to the customer acts as the aggregator. The service itself is operated by an external provider that acts as the host.

## White label

White label is a special case of aggregation. In this case, the provider rebrands (with their own brand) the interface and services from another cloud service provider, such as IBM SmartCloud Enterprise. The customer interactions (such as self service, marketing materials, and billing statements) all have the brand of the aggregating provider. However, the actual management and delivery of the services is provided by the *hidden* provider.

This model is often used for common consumer-oriented services that are included in bundles, such as productivity and communications services.

Regarding Figure 2 on page 11, the customer acts as the user and the provider (visible to the customer) acts as the aggregator. The cloud service itself is operated by an external provider that acts as the host.

## Mixed hosting and aggregation

In a mixed model, the provider acts as both a host and an aggregator, providing management and access to both local and external cloud services.

There are many reasons and variations for using a mixed approach, including:

- Both marketplace options and technical reasons, such as limited hosting capacity.
- Need to host data in specific locations.
- Requirement for flexibility of service deployment. For example, some SaaS providers do not have an offering that enables local hosting, so they can be aggregated only.

Regarding Figure 2 on page 11, the customer acts as the user and the provider visible to the customer acts as both a host and an aggregator. The cloud service provider selects the location of the cloud service (by service and customer preference). This approach allows the cloud service provider to choose whether to deliver the cloud service from their hosted cloud or from a cloud that is operated by an external provider that acts as the host.

## A solution that enables many business models

A range of business models have been described. The combinations of these business models provide one of the compelling values of cloud computing. Each cloud service provider chooses an entry point that best fits the current business and customer needs. The correct cloud computing architecture supports those needs as they evolve and grow over time. The overall business solution is enhanced by adding new business models and extending the current deployment to encompass the additional functions.

This scaling model, both of business scope and deployment capacity, may be used as a key differentiator for cloud service providers. This situation is particularly true for agile providers that are able to anticipate new market opportunities by having an architectural foundation that is flexible and extensible in response to changing business needs.

The cloud service provider solution anticipates both the typical initial deployment scoping around one or two specific business models. The solution grows to support more business models over time. Growth of the solution includes the addition of functions, systems integration, and new channels and ecosystem elements.

For both the IaaS entry point scenario and the SaaS entry point scenario, the addition of business models and capabilities are incrementally added to the initial solution. The motivations for selection and ordering of business models vary by deployment, including:

- Changing the mix of hosted and aggregated cloud services to meet revenue, cost, and operating margin targets that are based on the profitability of different services.
- For the IaaS entry point, the addition of PaaS capability to increase the base value of the stack that is offered to customers and developers. For example, a web stack can be priced at a higher price per hour than a plain operating system stack.
- For the SaaS entry point, the addition of PaaS capability can be an accelerator for the creation of new SaaS applications. This approach increases the volume and variety of SaaS applications available in the product catalog. Having a local PaaS capability can decrease testing and onboarding time versus externally developed applications.

IBM SmartCloud solutions provide the general architecture to support a wide span of business models. The IBM CCRA service provider adoption pattern provides the guidance to apply the IBM SmartCloud solution incrementally to fit the business models and capabilities you want to deploy.

## IBM cloud computing solution

A strong architectural foundation is required to enable rapid adoption of cloud computing with the ability to be responsive in the marketplace. Supporting the various business models allows rapid expansion of capability to respond to customer needs more quickly than competitors.

The IBM SmartCloud for Service Providers portfolio provides this foundation. This IBM solution consists of these industry leading elements:

- Robust infrastructure for availability and serviceability suitable for commercial service provider deployment
- Complete business and operational architecture consistent with service provider processes and preferred practices

- Advanced management capabilities that provide secure access, high scalability, efficient resource use, and integration of operational management systems

The solution enables cloud services to be offered with the confidence and quality that is expected by cloud service providers and their customers.

## The IBM Cloud Computing Reference Architecture (CCRA)

The IBM CCRA is a reference architecture that describes the functions and processes that underlie a broad spectrum of cloud computing implementations. The IBM CCRA includes functional definitions and addresses cross-cutting aspects, including security, resiliency, performance, consumability, and governance. This architecture provides a consistent base to support a broad set of business and technical goals that are realized through different cloud computing deployments.

The IBM CCRA is applied to develop prescriptive cloud solution patterns that are known as *cloud adoption patterns*. These adoption patterns represent the collective experience and preferred practices to enable prescriptive guidance on implementing cloud solutions. IBM identifies four key adoption patterns:

- Cloud service provider (described in this guide)
- Infrastructure-as-a-Service (IaaS)
- Software-as-a-Service (SaaS)
- Platform-as-a-Service (PaaS)

The IBM CCRA cloud service provider adoption pattern shows how to provide cloud services through a service provider approach. This approach enables organizations to formally define processes and relationships for making cloud services available in an organized manner. This structured approach provides efficiencies in both managing the environment and in ease of use.

In addition to the cloud service provider adoption pattern, the CCRA supports three more adoption patterns:

- IaaS

  The IaaS adoption pattern defines the core infrastructure for cloud resource management and the operation of IaaS services. The cloud service provider adoption pattern builds on this architecture, adding essential service provider capabilities.

- SaaS

  The SaaS adoption pattern defines the architecture for the definition and operation of SaaS applications. The cloud service provider adoption pattern defines the architecture for the environment into which SaaS applications are onboarded, managed, and offered. SaaS deployment is not dependent on IaaS, but uses the same underlying cloud computing management functions.

- PaaS

  The PaaS adoption pattern relies on the cloud service provider adoption pattern to define the environment for the PaaS ecosystem. This ecosystem supports the ability to build, test, and deploy cloud services. There are two common PaaS environments:

  – PaaS development operations (DevOps) environments provide a development environment to create SaaS applications using development tools and processes that are offered as cloud services.

– PaaS pre-configured runtime environments provide ready to run stacks onto which applications can be deployed. For example, IBM WebSphere® Application Server is available today on different clouds as a pre-packaged platform. Developers are able to select this PaaS base and deploy applications directly to the application server instead of creating full virtual machine images.

- Business-Process-as-a-Service (BPaaS)

  Many businesses develop business processes that interact with many applications to complete their tasks. For example, a payroll processing process interacts with timesheet, tax, benefits, and human resources systems to produce paychecks for employees. BPaaS provides the means to implement these business processes as services. The cloud service provider adoption pattern manages BPaaS processes in the same manner as SaaS applications are managed.

Figure 3 shows the IBM CCRA as applied for the cloud service provider adoption pattern.

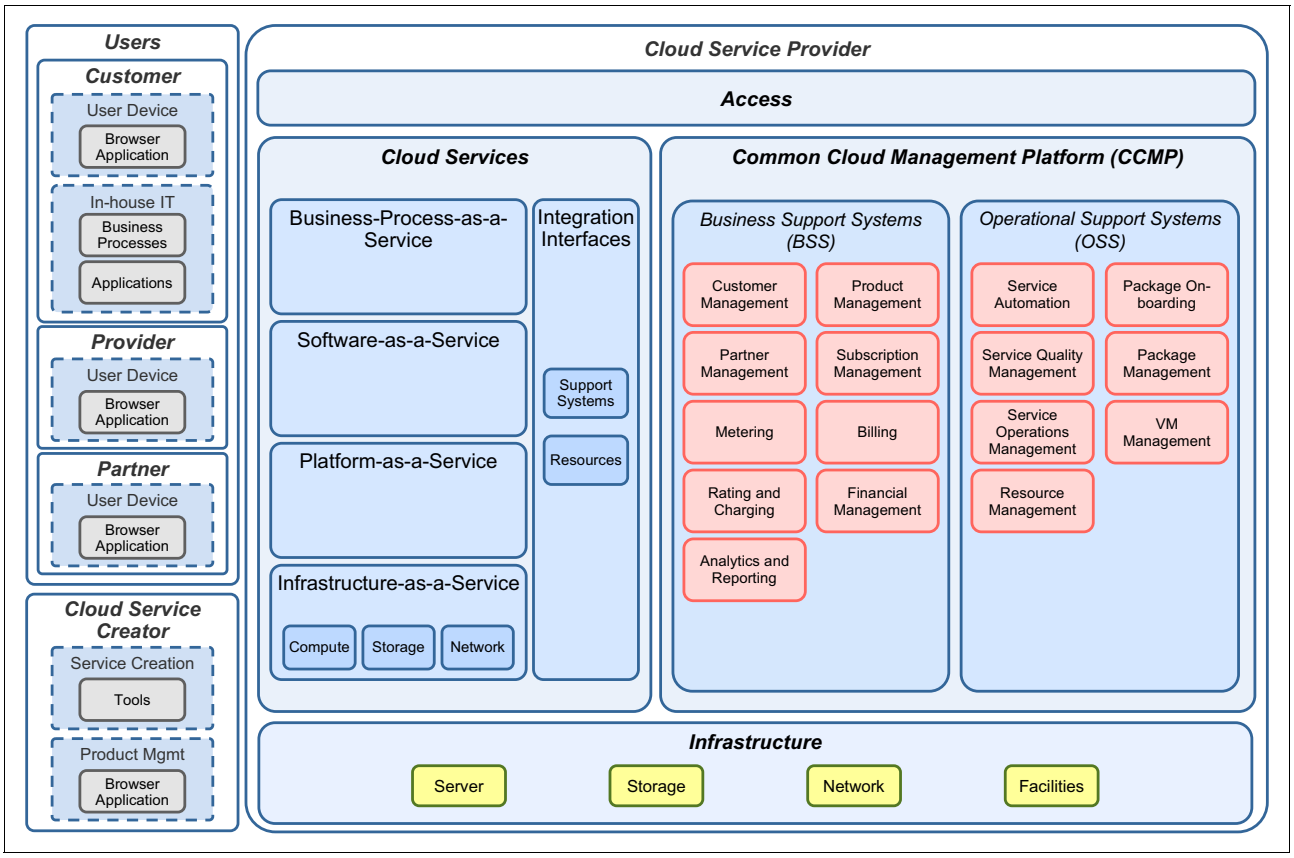


*Figure 3 Component view for the IBM CCRA cloud service provider adoption pattern*

The architecture (Figure 3) enables building solutions that extend the capabilities of cloud to serve new market opportunities. The expansion of capabilities is accomplished by adding component implementations or extending the capability within functions to meet the additional market and technical needs.

For service providers, the implementations of components for management include the integration of external interfaces for customer access, and appropriate functions for management systems.

## Key technical capabilities

Service providers maintain relationships with multiple constituencies, often enabling an ecosystem of business to integrate in the delivery of services to customers. This capability is reflected in the approach to the architecture and the realization of the functions deployed. Capabilities are classified as functional and non-functional and are described in this section.

### Access capabilities

The architecture defines secure access for all interaction points. This solution defines five interfaces. Each interface has appropriate access control functions defined to ensure proper authorization to functions and services. Also, the solution provides network management to enable scalability and availability implementations to meet these goals.

The five interfaces for the solution (presented in the Access component in Figure 3 on page 20) are:

- Customer management: Interface for the customer account manager.
- Storefront: Interface for the customer account manager and customer users that have buyer authority.
- Customer access: Interface for users to interact with cloud services that they subscribed to.
- Partner management: Interface for the partner account manager to interact with the cloud service provider. For developer partners, also includes onboarding interfaces.
- Provider management: Interface for employees and agents that work on behalf of the cloud service provider. For example, retail store employees, customer service representatives, and salespeople.

Each of these interfaces may be visual and non-visual (enabling automation), and have corresponding functional capabilities that are provided by core functional capabilities.

Both the IaaS entry point and SaaS entry point scenarios implement customer management, customer access, and provider management interfaces. The storefront and partner management interfaces are usually part of a second deployment phase, enabling faster time to market for the initial deployment.

### Support systems capabilities

Support system capabilities include the Business Support Systems and Operational Support Systems (see Figure 3 on page 20). These capabilities provide functions that are critical to managing the business and operational systems in the cloud computing environment.

- Customer management

  Customer management supports customer accounts, contracts, and assigned administrators and users. Users are managed, enabling subscription, entitlement, and financial systems to support the various business models available. Customer management maintains an accurate record of the customer activities and accounting.

- Partner management

  Partner management enables those ecosystems that include partners, such as application vendors and retail partners. Partner management essential functions include managing partner accounts, contracts, and disbursements.

- Provider management

  Provider management cloud services, like other services, are managed and supported through administration and care functions. These functions are accessed by employees and agents of the provider. A common provider requirement is facilitating these functions, ensuring proper authority assignments, and enabling integration with related systems for other services.

- Product management

  Product management customers purchase products, often as part of bundles with promotions or discounts, from a selection of product offerings that match their areas of interest. The product catalog and pricing components enable flexible definition and market selection for customer product selection, configuration, and ordering.

- Order management

  Order management translates a customer selection of products into the set of activities necessary to enable access to the services selected. This activity can include initiating actions to create subscriptions, user identities, entitlements, and provisioning actions to the clouds where the services are accessed.

- Service automation

  Service automation provisions resources, service configurations, and user information as part of the processes to make cloud services available to customers. The provisioning processes interact with the service catalog and orchestration functions to determine the steps that are required to fulfill the provisioning request. Orchestration functions have an understanding of the resources to be interacted with. These functions know the order of steps that are required and how to process errors. Error processing includes applying corrections or rolling back partially completed provisioning when needed.

- Metering or charging-to-payment processing

  Metering or charging-to-payment processing resource-based charging uses metering to record usage. A service can also create discrete charges for specific functions. These charges are processed by the charging or rating component. This component applies customer pricing and selects the payment processing method according to customer preference. The payment processor is then invoked, generating a bill entry or interacting with a payment system to collect the charge. Two examples of a payment system are a credit card processor and an online charging system.

- Financial management

  Financial management cloud includes some functions that might rely on integration with existing systems, such as Accounts Receivable and Accounts Payable to complement the cloud functions. These capabilities may include balance management for prepaid customers and revenue sharing for partner distributions from service consumption.

- Service operations management

  Service operations management provides monitoring and event management, cooperatively working with problem and incident management processes. It provides the constant operational view of the cloud and the ability to communicate and respond to any issues that arise.

- Service quality management

  Service quality management recognizes operations in the context of the service quality. It is part of the ongoing processes that determines whether the cloud services and the cloud computing environment are meeting their objectives both in the context of the provider and the customer. The functions in this area provide visibility to the aspects that impact the quality of service delivered.

- Package onboarding

  A package onboarding process to manage the entry of cloud services and content into the cloud computing environment is a key capability. Security concerns that are related to the acceptance of outside content dictate a well-managed process for content acceptance. However, the process must also be efficient to enable timely availability of the services. This function defines the processes, facilities, and automation that is associated with the onboarding process.

- Package management

  Package management cloud services can come in many forms and can be deployed dynamically using automated processes. To enable this automation, an approach using packages is used to assemble the application content, data content, configuration, and service information. To accept these packages and process them, a package management component is defined. The package management function stores the content that is related to each cloud service (such as an image for IaaS or an application package for SaaS), its configuration, and catalog metadata. By storing this information, the package can easily be managed, found, retrieved, and deployed.

- Virtual machine and resource management

  Virtual machine and resource management provides management of the infrastructure, and operational functions (for example, hypervisors) are key functions in the cloud execution environment. This function includes capabilities such as resource provisioning, expansion/contraction, and hypervisor management (such as start, stop, resume, and terminate).

Many of these capabilities are scalable in function, from manual to self-serve. The individual deployment plan for each project determines the order and level of function for each capability.

Typically for the IaaS entry point scenario, the capabilities for customer management, service automation, package management, and virtual machine and resource management are prioritized. Other functions can be partially automated initially or be integrated with existing systems, enabling automation and self service of function over time.

For SaaS entry point scenarios, customer management, order management, and service automation and payment processing are prioritized, enabling ease of service selection and access for customers. Payment processing has a higher priority for those providers that are providing credit card and similar payment options.

## Shared system capabilities

Shared system capabilities underlie the functions that are shown in Figure 3 on page 20. These shared system capabilities ensure the security, availability, performance, and governance of the solution. Security is a key concern and touches many of the components in a cloud deployment. *Cloud Security Guidance IBM Recommendations for the Implementation of Cloud Security*, REDP-4614 provides a detailed view of the different aspects of security and approaches to addressing them. The shared system capabilities are:

- External access security:

  External access security focuses on ensuring that the environment is secure from outside threats and includes:

  - Access control to the cloud, which includes authentication, authorization and identity management
  - Network security to ensure that the network is secure against threats

  - Secure communications, including Virtual Private Network and usage of secure protocols

- Internal access security

  Internal access security secure zones using firewalls, isolated network segments, and protected resources, which provide layers of protection for data and processes. This security approach enables efficient access while providing proper isolation of resources.

- Service security

  Service security within the cloud. Security controls include isolation of processing, memory, network, and storage, ensuring each cloud service resource is protected from other cloud services. It also ensures that proper resource controls are applied to users of each cloud service.

- Availability

  Deployments support a range of availability choices, including high availability within a data center, failover/recovery that spans data centers, and automation of both failover/recovery and disaster recovery processes.

- Performance

  Monitoring of the cloud itself and the cloud services that operate within the cloud is managed through a set of data collectors that feed runtime performance management systems. These performance management systems provide both operational views and evaluation against key performance indicators (KPIs) to recognize performance state and the need for actions.

- Governance

  A cloud computing environment is a complex entity, with many cooperating elements that need definition and maintenance over time. Governance requires a set of processes and enabling capabilities that follow preferred practices, use automation, and have predictable outcomes.

The shared systems capabilities are not optional, but there are choices that affect the depth of function. The IBM solution, for both the IaaS and SaaS entry point scenarios, provides coverage for all these shared functions at a commercial deployment level of capability. Typically, the availability aspect is provided within a data center for initial deployments. Multi-data center capability can be delivered as the volume of use scales.

## High availability

The service provider adoption pattern supports the deployment of solutions that support high availability functions on both ends of the interaction, including:

- Redirection of traffic
- Failover to other local systems
- Failover to remote locations
- Resumption of load after recovery

The usage of these functions requires appropriate facilities and implementation of the high availability function at each location, including consistency of the failover or recovery functions on each side of the interaction.

For HTTP-based implementations (both visual and non-visual interfaces), the high availability model follows existing web patterns for deployment models and preferred practices. For non-HTTP based implementations, the high availability model is defined by the protocol and management functions, and is implemented in a similar manner.

**High availability:** Using external providers to provide high availability services is also supported. The service provider adoption pattern supports this deployment model using the aggregation business model and deployment architecture.

### Interoperability

Interoperability makes interactions between the user and provider, partner and provider, and between providers predictable when they are deployed independently. All interactions between these parties are expected to be Internet Protocol (IP)-based.

The management interfaces for interactions typically use Hypertext Transport Protocol (HTTP) -based protocols, for both visual Hypertext Markup Language (HTML) and non-visual (web services) interfaces. For the non-visual interfaces, HTTP Representational State Transfer (REST) is commonly used as the underlying protocol for management interfaces. There might be other protocols that are used for specific functions, such as File Transfer Protocol (FTP). Secure versions of these protocols can also be used where wanted.

For application processing, protocol interoperability is defined by the application definition. Applications often use web-based presentation interfaces using HTML/HTTP. These applications can choose to use other approaches such as using a local application and interact with the cloud using non-visual interfaces such as web services. Other protocols can also be used, including general purpose protocols, such as messaging protocols, or application specific protocols.

The service provider adoption pattern defines the visual and non-visual interface components for each interaction interface, and supports both general and application specific protocol usage.

## Implementing cloud computing using IBM SmartCloud for Service Providers offerings

Considering the business models and available technologies, there are many options for implementation. Two key considerations that drive the definition of the cloud computing environment are workload definition and the rollout plan. These two factors drive both architectural decisions and the cost basis on which the business case is built.

### Workload definition

A key consideration for deploying cloud services profitably is evaluating workloads well. Finding good matches of customer need, provider capability, and a deployment environment is the evaluation objective.

The starting point for workload definition is the selection of cloud services and creation of profiles that represent the users using the cloud services. These selections and characteristics inform the architecture definition, and provide the basis for modeling the deployment environment.

Some key characteristics that emerge from this activity are:

- Mix of compute, storage, and network resources that are required, including identifying those resources prone to constraints or becoming bottlenecks.
- Lifecycle expectations for the use of resources. Resources might be fluid with many instances allocated and released frequently or relatively static.

- Number of concurrent connections and duration of connections. Connection lifecycles may impact network design and cloud service design to ensure user responsiveness and ability to scale the cloud services as users are added.

Classifying workloads is about identifying the infrastructure, business, and management capabilities that are required for delivering the workloads and planning the operational deployment environment to optimize their run time. This classification provides both efficiency in delivery and the opportunity to differentiate by establishing a distinct value for each workload.

The solution architecture supports selecting and optimizing specific workloads or managing mixed workloads. The business and operational processes can be configured to support specific workload types. This approach enables each type of workload to have appropriate processes in place to optimize:

- Onboarding processes
- Product definition
- Ordering
- Operating the cloud services within the workload

## Phased approach to deployment

Deployment of a cloud computing solution is typically done in a phased approach, with the starting points varying based on the selection of the initial business models.

For those business models where hosting of cloud services is included, the initial phase includes deploying a cloud infrastructure capability. The cloud infrastructure capability is deployed along with its supporting business and management systems. This base configuration includes support for managing virtual images and for managing the lifecycle of cloud services. The most common phases that follow this initial phase are:

- Supporting additional image sources, enabling images to be provided by customers and third parties, which enables a broader selection of cloud services to be delivered.
- Introducing new types of cloud services, such as SaaS applications that use subscription-based models.
- Adding platform images or pre-configured images with pre-packaged platforms, enabling PaaS offerings. These images can include DevOps capability, enabling developers to develop and test their applications in the cloud environment itself.

Cloud service providers that start with an aggregation business model (including white label) where they are not operating the cloud infrastructure initially have a phased approach that adds:

- Additional external cloud service providers to broaden the product selection offered.
- An initial cloud infrastructure deployment that enables local deployment of some cloud services or offering of local cloud resources to supplement the initial services.
- A PaaS offering through an external cloud service provider. Using an external PaaS provider allows this offering to be added without incurring the cost of a creating a hosting environment. This phase may also include offering DevOps capabilities from the external PaaS provider.

The order or combination of phases is not predicated on prior deployment choices. There is a great deal of flexibility in both the choice of initial capability and the direction of additions. Because of the inherent abstraction of cloud services, changes to models (such as bringing a white label service in-house, or vice versa) can be transparent to the customer.

## Architectural decisions

Architectural decisions flow from the choice of business models and marketplace audiences that the provider is looking to serve:

- Hosting/aggregation business model selection

  When you enter the cloud service provider business, one of the first decisions is whether to deploy and manage a cloud infrastructure as part of the business initially or not at all. This choice is a fundamental one about how the business operates and what relative priorities and activities are part of establishing the cloud.

  If you choose to host a cloud, the next choice is whether to also aggregate cloud services from other cloud service providers. If aggregation of cloud services is chosen, the cloud services can be aggregated as pass-through services or white labeled. The white label aggregation approach enables the provider to associate their own branding with the services.

  Implementing an aggregation capability, with or without hosting, requires the architecture to include an integration component both for cross-provider provisioning and operational and business data exchange.

- Storefront inclusion

  A self-service storefront capability provides a product driven interface that has broad reach. This capability is a common choice for cloud service providers that address consumer and enterprise marketplaces.

  This capability incorporates the storefront interface and the supporting product management capabilities in the overall architecture. It provides a richer user experience for customer management functions and reporting functions. Solutions that do not use storefronts typically provide a more administrator oriented interface.

  The storefront decision is equally applicable to both the IaaS entry point and SaaS entry point scenarios. This choice is driven by the audience and the number of cloud services that are offered, not the type of cloud service. For most cloud service providers, if a storefront is not part of the initial deployment, it is a roadmap item that is based on the growth of cloud services offered.

## Cloud service provider adoption pattern principles

The cloud service provider adoption pattern builds on the experience of developing a wide variety of general purpose architectures, solution specific architectures, and solutions for service providers. The principles that guide this work include:

- Application of understood industry terms, capabilities, and patterns across the Internet and communications businesses. These principles enable cloud computing to be more easily adopted as a delivery capability for service providers.
- A business model-led approach to architecture scope and definition, enabling a broad span of adoption across many types of businesses that enable cloud service provision.
- Defined and practical definitions for interactions between providers, customers, and partners, enabling a cloud to manage the complexity of implementation while retaining simplicity in interactions.

The adoption pattern is flexible and adaptable to evolving standards and application of different standards to distinct businesses that become cloud service providers.

## Component model

The component model shows the implementing components and their categorization in applying the IBM CCRA to create the cloud service provider architecture. A high level overview version of the component model is shown in Figure 3 on page 20.

Components are grouped into these categories:

- User groups

  The user groups consist of customer groups, including users and managers that access services from the cloud service provider. Users also include users that work in partner and provider related organizations:

- Cloud Service Creator

  The Cloud Service Creator category includes the development and deployment aspects of creating cloud services and deploying them to the cloud.

- Cloud service provider

  The cloud service provider category includes three major subgroups:

  - Cloud services

    Cloud services encompass the various services that can be provided by the cloud service provider. The types of services include the commonly referenced IaaS, PaaS, SaaS, and business-process-as-a-service (BPaaS).

  - Common cloud management platform

    Common cloud management platform contains the set of operational and business support systems that manage all aspects of the cloud environment.

  - Infrastructure

    Infrastructure represents the set of resources that are used by the cloud components in operating the cloud environment and execution of cloud services.

In addition to the components, there are aspects that are shared across the components, including security, resiliency, performance, and consumability. These common items represent the implementation of features within components that address these areas. Surrounding the deployment and ongoing management of the cloud is the governance capability.

Providing cloud services requires support systems with robust implementations that provide:

- A function that allows access by external users and systems
- Access control that manages access to systems based on the application of authentication and authorization
- Data and process isolation to ensure that proper visibility is applied to data access and the function of processes
- Scaling that is consistent with the workload generated by the users that access the systems

The depth of implementation in the business support systems is significant in cloud service provider deployments. There are also broad interactions with the operational support systems to support the additional business function. For those providers that choose to implement the third-party aggregation business model or host third-party applications in a cooperative model, the additional partner-related business and operational capabilities are applicable.

Figure 4 shows the next level of detail for the business support systems.

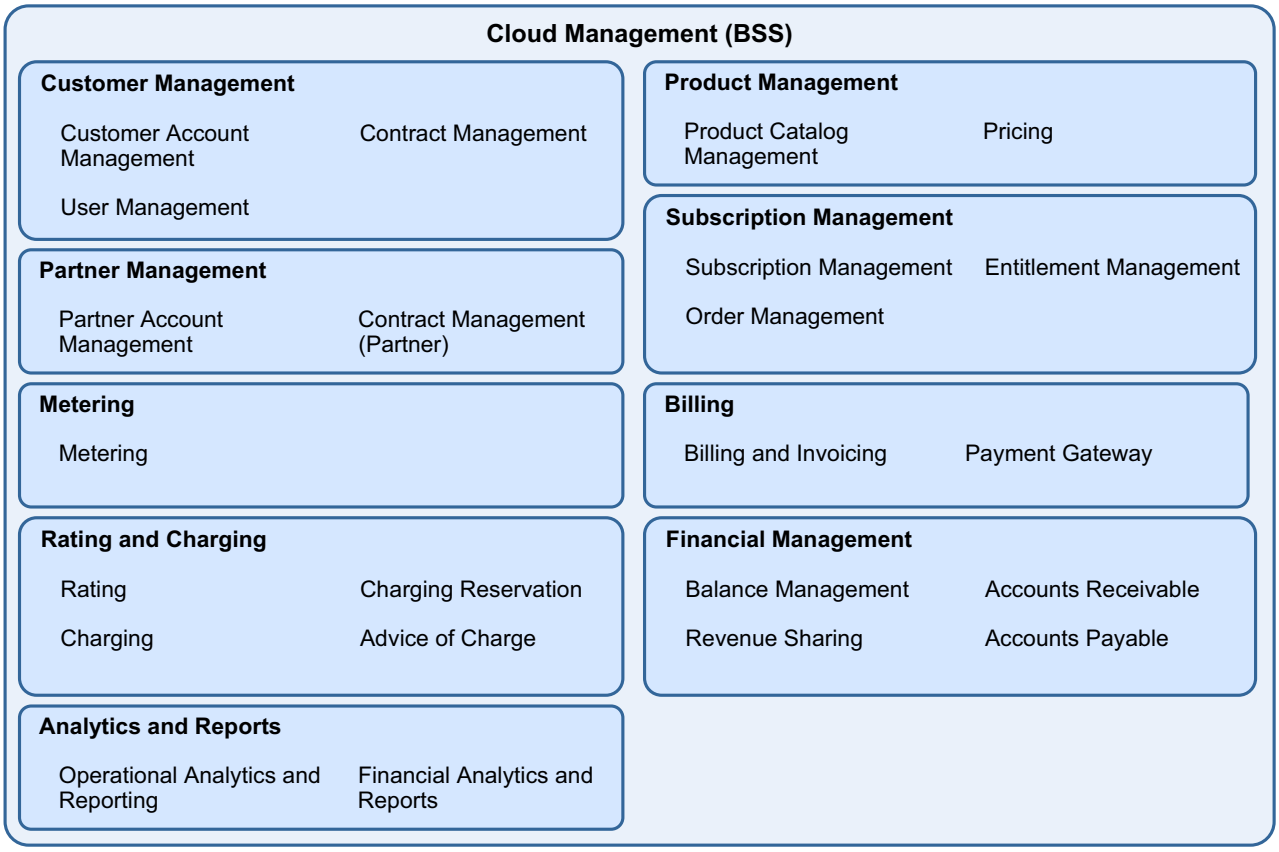


*Figure 4 Component model for business support systems*

The selection of components and their capabilities within the business support system (Figure 4) are driven by the selection of business models:

- Customer management functions are always implemented. User management is more sophisticated for SaaS and aggregation business model support.
- Product management is a key distinction for cloud service providers, enabling user-oriented product definition, selection, and pricing flexibility. Providing customer and partner definition capability for the product catalog provides differentiation and scalability for high volume catalogs.
- Partner management is applicable for third-party service management and for management of channel partners, which are common for cloud service providers.
- Subscription management is another key distinction of a cloud service provider. It enables the provider to provide management of subscriptions and entitlements consistent with the expectations of cloud customers. Subscription lifecycle management and the corresponding robustness of order management provides the flexibility to manage the dynamic product usage that is expected of cloud services.
- Rating and charging brings together the product, pricing, and usage data to create accurate charges. Charging determines how to apply the charges that are based on the customer profile. For business models that support prepaid accounts, reservation management is also provided.

- Financial management for cloud service providers can include revenue sharing for third-party business models, enabling calculation and distribution of payments. Balance management manages balances for providers that offer prepaid offerings. For prepaid systems where the provider does not manage the balance itself, an external balance management system can be integrated to provide the implementation. Some jurisdictions require balance management to be provided by a recognized banking entity.

Figure 5 shows the next level of detail for the operational support systems.

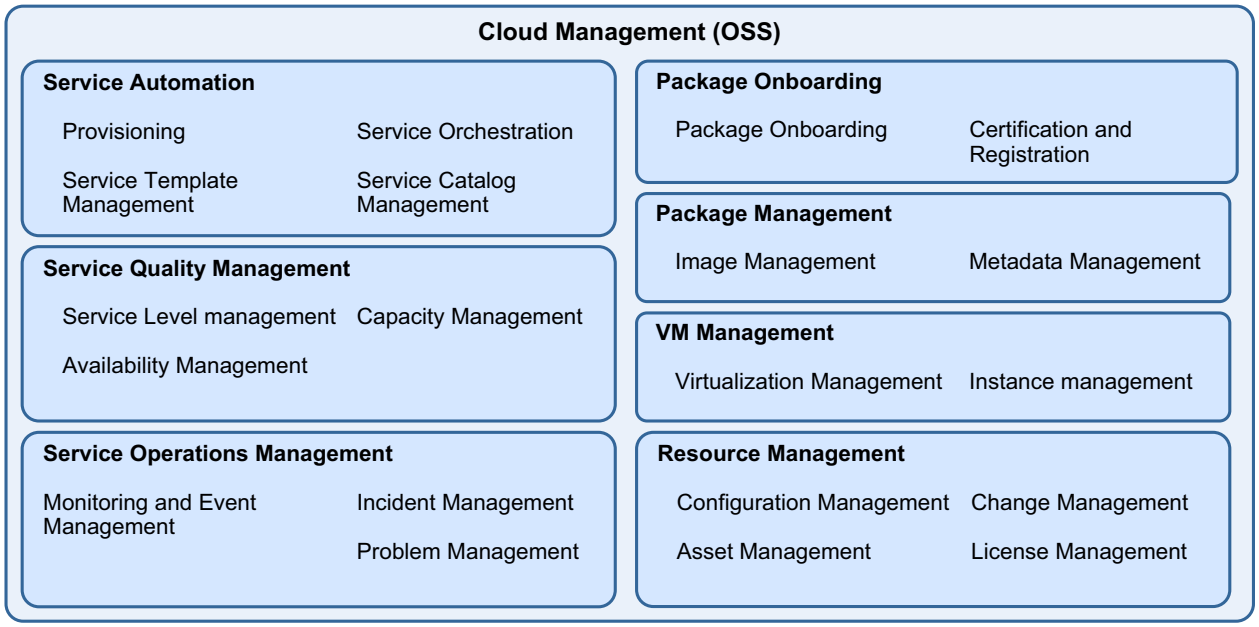


*Figure 5 Component model for operational support systems*

The operational components have a broad subset that is applicable for any cloud deployment (Figure 5). For cloud service providers:

- Monitoring and event management and problem management have external visibility for appropriate content per customer/partner.
- Package onboarding components are typically more robust in order to handle externally sourced content. This component can include appropriate interactions with the related component, such as package management and service catalog management.
- License management can include management of externally sourced licenses.

To support the various business models, the integration between the business and operational support systems includes:

- Service automation capabilities that are coordinated by the provisioning component to interact with the order management function.
- Monitoring generates input for metering that enables the structured capture of resource usage towards payment processing components.

The component model that is defined in the CCRA service provider adoption pattern provides the foundation on which many solution architectures are defined. This foundation enables consistency across many solutions. The component model supports variance in the depth of implementation, supporting staged implementations that expand the function within components as needed.

## Operational views

The different combinations of business models drive different configurations for actual deployments and their respective operational views. To provide a flexible understanding for operational views, multiple views are provided that can be subsets or can be extended to encompass the various business models. These views include a management view and a service use view.

### *Operational environment*

The operational environment for a cloud computing environment is composed of four functional categories (see Figure 6).

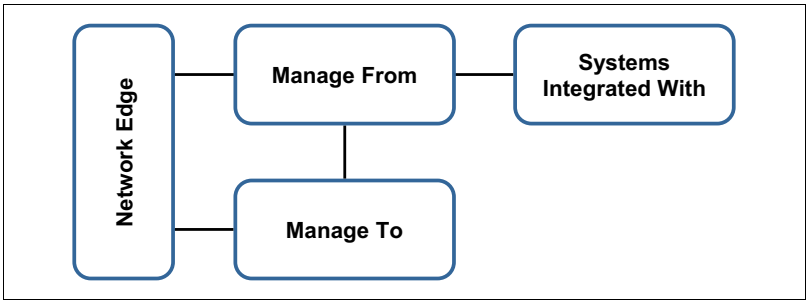


*Figure 6 Categories of operational functions*

Deployments of cloud computing environments implement these categories of functions according to the services offered, support systems that are required, and existing systems to be integrated. The functions of each category (Figure 6) include:

- Network Edge

  Network Edge is an operational area that provides the edge of network functions for routing and security, including routers, appliances such as encryption accelerators, and IP sprayers.

- Manage From

  Manage From contains the support systems and resource management systems that govern and manage the systems that runs the cloud services.

- Manage To contains the systems that the cloud services run on, and the functions, such as hypervisors, probes, and monitors, that interact with the support systems.

- Systems integrated with

  Systems integrated with identifies the systems that are outside the cloud computing environment itself, such as the customer account database, billing systems, and financial systems. These systems are integrated with the support systems, rather than duplicated within the cloud computing environment.

### *Operational view for cloud management*

The cloud management activities in an operational view are shown in Figure 7.

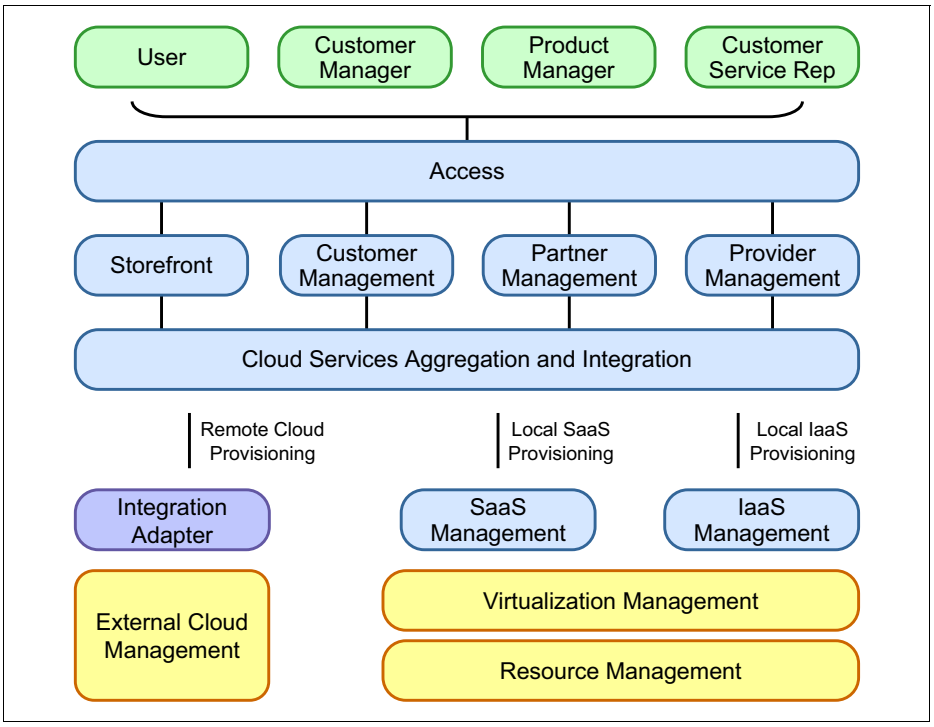


*Figure 7 Operational view of cloud management activities*

At the top of Figure 7, a representative set of users is shown. Below the users is the access interface that all users interact with to reach the cloud management functions. The capabilities available in the access interface include network routing, authentication, authorization, and identity management, and security services. It is common for these capabilities to be shared across many users and protocol endpoints.

The four management interfaces below the access interface serve different audiences (customers, partners, or providers). For each audience, the context of the interaction also selects the interface used (for example, storefront interactions to select and order services for a customer manager).

The cloud services aggregation and integration capabilities implement the business and operational support systems that handle the interactions for each of the management interfaces. The integration capability selects the target management system and location (hosted or external cloud), and initiates actions. These interactions can be sophisticated, including recognizing errors and performing compensating actions.

SaaS and IaaS are shown as representative services. Any other type of service is handled in the same manner architecturally, including the ability to have distinct provisioning interfaces.

### *Operational view for cloud service usage*

Using a cloud service depends on the completion of the management activities that are provisioned for the service and the user authorizations ahead of the time that the service is requested for use. The operational view of cloud service usage (shown in Figure 8) defines the interactions between the user and the service. Figure 8 also shows the interaction of systems that are responding both to the user actions and to the configurations that are established by the management systems.

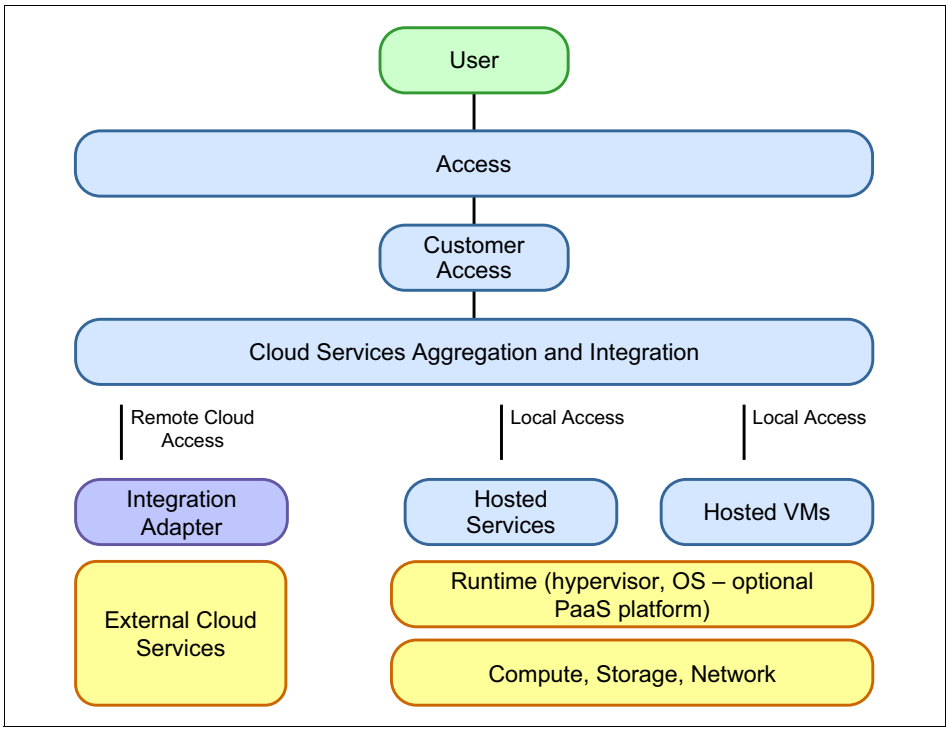


*Figure 8 Operational view of using cloud services*

As in the operational view of cloud management activities shown in Figure 7 on page 32, the access and cloud services aggregation and integration capabilities are also present in Figure 8. However, in Figure 8, they provide the capabilities related to the use of the service, not the management capabilities necessary in setting up and making the service available.

The access capability and the integration adapter provide federated identity management to provide single sign-on capability for the user that accesses the service remotely.

On the right side of Figure 8, the hosted services and hosted VMs are shown relative to their runtime and resource components as deployed by the cloud services provider.

## Solution integration

Companies adding cloud services to their set of products usually have a number of existing systems in place for functions, such as authentication and billing. Extending some of these systems to include cloud services is a frequently chosen and cost-effective alternative to deploying new systems that perform the same function.

The IBM CCRA component architecture approach enables integration to be identified at the component level and for integration to be defined in a manner that does not disrupt the overall solution. For example, integration with an existing billing system does not change the flow between metering, rating, and billing that feeds the data to the existing billing system.

In addition to integration at the component level, systems integration is required across the hardware, hardware systems management, monitoring, hypervisor, and software systems management capabilities. This integration may be pre-packaged in solutions, such as the IBM SmartCloud for Service Providers portfolio, or may be defined as part of the solution definition.

## Common integration points

Common integration points that enable integration with existing systems makes communications between the solution and these systems much easier:

- Authentication

  Service providers usually have authentication, authorization, and accounting (AAA) systems in place, and require services to use this system. Integration using LDAP or similar protocols is common and automatically supported in solutions.

- Authorization

  Similar to authentication, authorization for access is provided through the AAA integration. There are often additional authorization steps, including some that occur in the service itself (such as some privacy checks). Authorization can have additional integration points with a policy management function or privacy function.

- Usage

  Usage, related both to service usage and resource usage, is often collected for usage by charging and network management capabilities. The data collection is usually to a file. A data collector is used as an integration vehicle to retrieve this data on a continuous or periodic basis.

- Billing

  Inputs to the billing system to enable charges to be allocated to customers and invoiced are handled through a custom mediation. This approach uses a file transfer, message queue, or formatted file format to transfer the data into a defined billing process.

- Network and security infrastructure

  Clouds are typically deployed into environments that have existing network infrastructure for the edge of the network (such as routers and edge security appliances). The cloud networking resources are provisioned in accordance with these existing systems as appropriate for the deployment.

- Network management

  Network management systems (including monitoring and event management) often have an existing infrastructure in place. The cloud computing monitoring and event management functions integrate with these systems. This approach provides the service provider with an integrated view that includes the cloud services and infrastructure operations.

- Data management

  Typically, a cloud deployment does not create independent data masters for customer or partner data. Instead, these business support systems can be integrated with the existing master data sources, using either replication or managed access. Replication or managed access are used depending on the expected workload, security, locality, and other considerations. When managed access is used, access uses an approved access interface (either generic or database specific). When replication is used, an approved access interface is used and replication settings are used to manage distributions between the systems. Appropriate security considerations, including placement of the data in secure zones, can be applied.

Each of these integration points has patterns associated with implementing the cloud system with existing systems. In many cases, these systems are consistent with deployed systems that have similar integration points for Internet services and communications services.

## Service development and onboarding

Creating new cloud services is performed by a cloud services developer. A developer is not limited by the organization that they are part of. The developer can be one of the following individuals:

- An employee of a provider, an independent business, or person contracted by a provider. The developer creates a service to be deployed by the provider.
- An employee of a customer, an independent business, or person contracted by a customer. The developer creates a service to be deployed by the customer to a cloud operated by a provider.
- An employee of an independent business or an individual person, creating services to be delivered through one or more cloud service providers.

The creation of the cloud service can be done by using existing development processes and tools or done by using a cloud-based development environment. In both cases, the development steps process is similar, with the result of the process being a set of deliverables that includes:

- Development documents that cover requirements, design, and testing
- Deployment artifacts, including images and runtime components, configuration, and metadata
- Product management information, including product descriptions and pricing

Onboarding includes content and product management functions. Onboarding may also include staging, which is placing content first into a working area before being promoted to the production cloud environment. Deployment processes vary depending on the processes defined by the cloud service provider. The onboarding process can include a combination of manual and automated processes, including:

- Technical processes to deploy and define the service.
- Staging area for uploading the service content to the provider.
- Acceptance processes to validate content, and check for virus or malware.
- Placement of the content in the package management system.
- Creation of the service catalog content that defines the technical aspects of the service and its configuration.
- Creation of product catalog content that defines the products that are orderable and authorizations for who may order the product.
- Business processes to define the licenses and contracts related to the service.
- Customers or partners may need to provide licenses to the provider, covering the dependencies in the services they are deploying.
- Partners may enter revenue sharing or other agreements with the provider that govern the usage of the service, measurements, and payments due based on use.

Delivering services in a service provider model, especially for third-party models, have many considerations beyond just running the service. The onboarding process facilitates ensuring that the processes are completed correctly and efficiently. A robust onboarding process enables the provider to scale the solution, confident in maintaining compliance with provider policies.

Automation of onboarding capability is typically not part of the initial focus capabilities for the IaaS entry point or SaaS entry point scenarios. The initial set of applications is typically configured as part of the initial deployment process, rather than following a post deployment onboarding process. If you are considering an initial deployment that is focused on third-party provided applications, adding this function to the entry point scenario is appropriate, along with starting the business process definitions, such as contract management and revenue sharing.

# Creating a plan to implement your cloud computing solution

From the broad range of service providers that worked with IBM to define their clouds, IBM learned that all cloud implementations are journeys.

This journey includes many factors, such as the services and the supporting human processes and systems to support the business of being a cloud service provider. This journey requires an initial plan and a roadmap to define the capabilities to be enabled over time. This approach supports the practical considerations of entering the marketplace with an initial deployment in a timely manner, and then growing the business over time.

## Influences on the implementation plan

The implementation plan is driven by the selection of business models and cloud services. Combined with a timeline to show the availability for each model and cloud service, the plan provides guidance for the project definition.

The selection of initial cloud services (the set of services to offer and the target markets for the services) helps establish the framework for defining the initial deployment. It is common for some processes to be manual initially. This approach allows services to be made available to a limited audience while you complete the automation activities to scale the solution. There are various influences on the implementation plan, such as:

- Selection of business models

  The selection of business models identifies how to deliver each service and the relationships between the provider and its customers, partners, and peer providers. It is reasonable to both consider initial and future models, including migrating between models as part of the implementation planning. For example, a provider may decide to white label a service from an external provider initially and then choose to host that service later when sufficient hosting capacity is available.

- Roadmap

  The roadmap consists of a path that lays out the systems, human processes, and services involved in the overall solution. By taking the information from the selection activities and combining it with the provider strategy, you can create the timeline and functional requirements against the timeline.

- Capacity determination

  Capacity determination consists of both initial and projected capacity (particularly the user and workloads). It is used to determine the initial configuration and growth configurations. These capacity values may also impact the selection of technologies to use, such as type of storage, network design, and power needs. This information also determines the value of automation investments and the timeliness required for automation to enable the correct mix of human and automated processes.

Using the information collected, the initial implementation plan drives the answers to the architectural decisions. These decisions in turn drive the component selection and component requirements for implementation.

## Usage of existing systems

Cloud computing skills reuse existing skills and preferred practices that were developed in related technology areas. Developing these skills relies on shared experiences developing and deploying clouds computing environments. This history opens the possibility to use existing systems in cloud plans.

When existing systems have sufficient synergy, they may be considered for use with the cloud. Commonality in web and Internet data center technologies and usage patterns allows you to extend their usage to cloud (though capacity must be scaled accordingly). Authentication, routing, dynamic name server (DNS), and similar functions are commonly extended to include cloud services.

Care must be taken when you evaluate suitability for integration. Systems with a similar function, but used for different types of services, might require modifications to their core function. For example, subscription and order management for a content management system can have similar catalog and user definitions, but these systems might have limited capability to address entitlements and service/subscription lifecycle coordination. Evaluating these systems for use with cloud computing systems is part of project specific analysis.

## Multiple data center solutions

The most common reasons to include multiple data centers in the implementation plan are geographic scaling, locality of data, and redundancy.

Placing data centers in multiple geographic locations enables infrastructure to be isolated and placed in proximity to users. A single data center scenario might reach limits on connectivity, power, or space that constrains its ability to meet initial requirements or to grow at the pace to meet demand. Using multiple data centers enables these factors to be managed across multiple locations, where the aggregate of constraints is less limiting than a single constraint.

Locality of data is an issue that reflects both operational considerations (moving data is typically more expensive than moving processes) and regulatory issues (data might be constrained to a jurisdiction). These issues might impact significantly the choices of how services are deployed and how they are provisioned for specific customers that have specific data locality requirements.

Redundancy can be provided within a single data center using isolated network connections, backup power, and standby resource capacity that is allocated dynamically when primary systems experience failures. A multiple data center approach is suitable when:

- Customers have redundancy requirements that exceed single site capacity when site capacity is constrained by failed systems that are unavailable or recovering.
- Disaster recovery for catastrophic failure requires separate physical facilities that have physical distance between them. For example, locations that are in earthquake or hurricane zones.

From a cost and management perspective, most redundancy scenarios are active/standby with data replication to enable the standby systems to continue working with a predictable state.

# IBM cloud service provider solutions

Building a cloud computing environment that follows the service provider adoption pattern starts with the foundation cloud computing capability. IBM has a range of options for cloud service providers, enabling the best fit of business and technology models to your initial and growth needs.

## IBM SmartCloud Integrated Infrastructure for Service Providers

The IBM SmartCloud Integrated Infrastructure for Service Providers solution provides a ready to deploy packaging of hardware, software, and services for installation and configuration of a cloud for service providers. It is delivered in three predefined configurations that enable you to select a starting configuration. These three configurations are extensible beyond the initial configuration by adding capacity in pre-packaged configurations. This packaging enables a ready-to-use option for the most common starting configurations, including the IaaS entry point scenario that can be installed and operational rapidly. Two software products packaged in the solution are:

- IBM SmartCloud Provisioning, which provides the management and administration capabilities
- IBM SmartCloud Monitoring, which provides the operational monitoring capabilities

## Custom service provider solutions

Service providers that build custom cloud computing configurations can select the IBM SmartCloud for Service Provider portfolio, which provides an *a la carte* approach to component selection and implementation. The IBM Service Delivery Manager product provides capabilities that can be customized for specific cloud computing configurations, including:

- Systems management
- Service orchestration
- Monitoring
- Provisioning
- Cost management

The configuration choices enable greater flexibility in the customization of provisioning automation. This flexibility is suitable for those clouds that require customized cloud service configurations and integration with custom cloud hardware.

Delivering custom integration is accelerated by using two IBM products:

- IBM Cloud Services Aggregation and Integration Pack provides pre-built business processes for managing customer, product, and order management functions.
- IBM WebSphere Cast Iron® Integration provides an integration platform and pre-built connectors for IaaS platform (including ISDM) and SaaS application integration.

For more information about IBM Service Delivery Manager, go to:

`http://www.ibm.com/software/tivoli/products/service-delivery-manager/`

## Storefront

The choice of a storefront approach is determined by a combination of the business models selected, target marketplaces, volume of user transactions, and available technology. The storefront is applicable to both the IaaS and SaaS entry point scenarios.

Table 1 shows two instances for the selection of a storefront.

*Table 1 Storefront selection considerations*

| Consideration | Commercial storefront | Extensions to existing customer portal |
|---|---|---|
| Business model | Account management self-service, and automated product selection and ordering | Account management self-service. |
| Target marketplace | Self-service | Self-service. |
| Transaction volume | Medium to high volume | Low transaction volume. |
| Solutions | IBM Business Partners storefront solutions, such as those offered by Parallels and Jamcracker, Inc. Integration to hosted cloud and aggregation with other partners using IBM WebSphere Cast Iron Integration. | Limited extensions to the current customer facing portal. When a function is wanted or a volume increases, consider a commercial storefront. |

Service providers with customer facing portals but without cloud services that want to expand have these choices:

- Add cloud services to the existing portal.
- Migrate existing services to a new portal that spans cloud and non-cloud services.

This situation provides an opportunity for current service providers to move to a complete services portal capability. This approach is a viable one because existing portals are often limited in the services offered and provide limited support for the underlying service and subscription lifecycles.

## IBM ecosystem support

Cloud ecosystems are defining and delivering cloud services. There is an emergence of specialty providers that focus on aggregation, auditing, and other services. Ecosystem development is not limited to one pathway:

- Platform ecosystems create communities of developers that accelerate service creation and delivery.
- Vertical market ecosystems can be created by cloud service providers or customers that address specific communities of customers.
- Cloud service providers can create ecosystems exclusive to their cloud.

IBM PartnerWorld provides the facilities for third-party cloud service creators to certify their cloud services to IBM platforms. The catalog of cloud services maintained by IBM PartnerWorld provides a ready source for cloud service providers to identify and select cloud services to offer. For more information about IBM PartnerWorld Ready for IBM SmartCloud Services solutions, go to:

https://www.ibm.com/partnerworld/wps/servlet/ContentHandler/isv_com_dvm_techval_smartcloud

This flexibility in approach is accommodated by the loose coupling approach taken in the IBM CCRA cloud service adoption pattern. This independence enables the service provider adoption pattern to support new business and deployment configurations for multi-provider scenarios and complex customer provider topologies.

The emergence of new types of services is enabled by the flexibility of the integration of the cloud service execution environment to the surrounding components. As PaaS emerges, many business model and deployment variations are envisioned. An example of such an environment is the IBM SmartCloud Application Services for building and deploying PaaS applications using a cloud hosted service.

For more information about IBM SmartCloud Application Services, go to:

http://www.ibm.com/cloud-computing/us/en/paas.html

The list of cloud services offered continue to vary over time:

- Services commonly expected for each marketplace addressed (for example, webmail for consumers).
- Evolving services that emerge and then subside (such as games).
- New services that are introduced and establish new marketplaces.

The roadmap for services is unpredictable, making the usage of a platform approach to cloud services deployment preferable. Various factors affect predictability:

- Services come and go.
- Subscriber numbers for services peak and stabilize.
- Seasonal/event driven services are introduced.
- Variations in number, type, and usage of each service must be accommodated.

# Cloud service provider deployment scenarios

The report *The natural fit of Cloud with Telecommunications*[2] contains a number of customer scenarios for cloud computing. For this guide, two independent scenarios are shown, one for each of the entry point scenarios. The scenarios show the decisions taken for addressing specific marketplaces and growth paths.

## Scenario one: Vertical market cloud services provider

If you are considering entering the cloud services business in a specific market niche, then this scenario is representative of the decisions and solution definition you can follow. This scenario uses the IaaS entry point.

Niche markets focus on identifying specific interest areas and providing distinct value to the customers within that niche. Niche markets can be defined as:

- Market segments defined by a type of business, such as doctors or retailers.
- Technology specializations, such as high performance computing, or specialized sets of services, such as video encoding.
- Geography based offerings that address specific city, state, or other boundaries. Geographical offerings usually include local content or are affiliated with other business functions that occur in the same geography.

[2] "The natural fit of Cloud with Telecommunications", found at:
http://www.ibm.com/services/us/gbs/thoughtleadership/ibv-telecom-cloud.html

Niche markets have strong synergy with cloud computing because of the relationship between specialization and efficiency of delivery. Niche markets have a high focus on their specific needs. Using shared resources and automation, cloud computing offers the flexibility required of niche applications while being efficient in delivery costs.

For this scenario, the technology/geography niche selected is small enterprises in a city that wants a packaged application management solution. The initial business consists of a customer application hosting along with provider application hosting for related horizontal services. The initial services are VM hosting of customer applications (using IaaS), persistent storage, and backup services.

These customers use the IBM SmartCloud Integrated Infrastructure for Service Providers solution. This solution provides pre-configured hardware and software, including:

- Compute capabilities that support 128 - 1800 virtual machines of capacity in the starting configurations, which is expandable as capacity needs grow.
- Storage and networking delivered in a rack configuration that is extensible.

Figure 9 shows the products used to deliver this cloud computing solution.

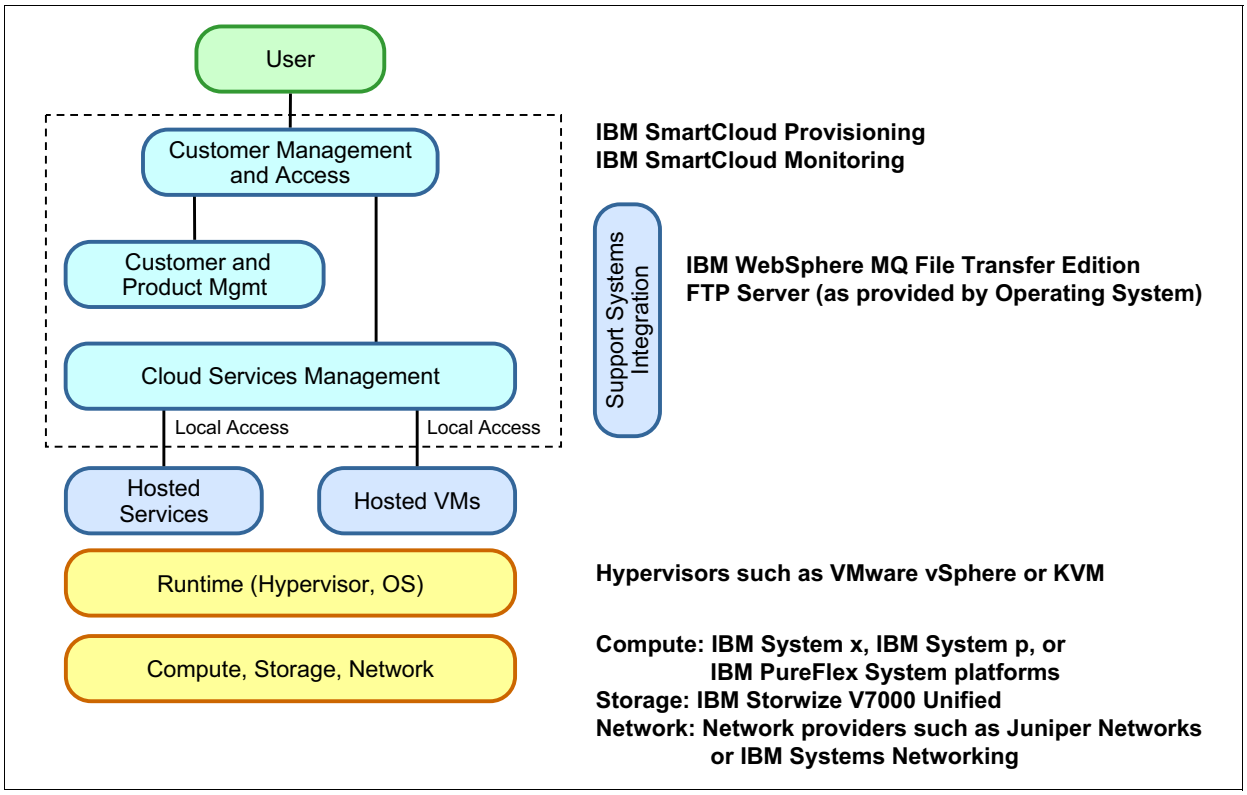


*Figure 9 Vertical market solution deployment configuration and products used*

The functions provided by each of the products are:

- IBM SmartCloud Provisioning

  IBM SmartCloud Provisioning provides a cloud computing environment, including business and operational functions. It provides the interfaces for interaction with the customer for both management and access to cloud services.

- IBM SmartCloud Monitoring

  IBM SmartCloud Monitoring provides execution time monitoring of the operational environment, including recognizing issues and initiating actions in response.

- IBM WebSphere MQ File Transfer Edition or File Transfer Protocol (FTP) Server (as provided by the operating system)

  These solutions enable loose coupling of the cloud services environment with existing systems to ease integration.

- Hypervisors

  Hypervisors (such as VMware vSphere or Kernel-based Virtual Machine (KVM)) provide the virtualization capability in the managed environment and management interfaces that IBM SmartCloud Provisioning uses to perform management functions. Other hypervisors can also be used.

- IBM compute platforms

  IBM compute platforms identify where the solution can be deployed. These platforms include the IBM System x®, IBM System p®, and IBM PureFlex™ System platform offerings. Rack and blade configurations are supported.

- IBM Storage

  In particular, the IBM Storwize® V7000 Unified and IBM PureFlex System V7000 integrated storage solutions provide the ephemeral and persistent storage.

- Network

  Configurations support networking products that provide network capabilities. Various network providers, such as Juniper Networks and IBM Systems Networking, provide these solutions. Both the internal network within the cloud computing environment and external access networks are covered.

The solution also includes backup facilities made available as a service provided by IBM Tivoli® Storage Manager. The cloud service provides automated backup for persistent storage and backup/recover service interfaces for users.

The primary users of the cloud services in this scenario are the people within the small business that are responsible for operating the applications. By using the cloud services offering, they are able to operate the applications without managing the hardware and infrastructure. They also benefit from automatic backup for data, and management of the backups by the cloud service provider.

For the cloud service provider, the management of the systems and applications is handled by the platform. Usage information is collected and correlated for billing. Integration with an existing billing system is used to provide a single bill for the customer that includes their cloud services charges.

As a starting point, the IaaS entry point is often selected because pre-packaged solutions (such as the IBM SmartCloud Integrated Infrastructure for Service Providers solution) provide ready to use configurations. The time to enter the marketplace and offer IaaS services is quick. Expansion of the platform is also easy enabling the platform to grow as usage and needs grow.

## Scenario two: Using cloud to drive mobile applications business

If you are considering entering the cloud services business to enhance a mobile computing capability, this scenario provides business and technical guidance. Mobile computing offers many advantages to customers, but also has many constraints to consider. Cloud computing can minimize the impact of the constraints, and enhance the mobile user experience. This scenario describes the available capabilities, and approaches for addressing mobile-specific capabilities.

Mobile devices, mobile applications, and cloud computing are commonly described together. A broad set of cloud applications are defined by requirements of mobile users. Mobile computing brings a broad range of business considerations, including channels (prepaid accounts being the dominant payment form in many countries), and the expected set of services.

Mobile device considerations also drive customer behaviors and user interface design. You cannot assume that a personal computer is available for activities such as selecting and ordering services. In many countries, a mobile device is the only computing device for the user. Considering these factors provides perspective about the role of the device and channel (retail or call center) in managing cloud services and customer accounts.

For cloud service providers, mobile device service selection drives business model selection. Examples of common starting services include,

- Email, either a webmail or hosted email service
- Device management, with functions such as data backup
- Personal storage, a storage-as-a-service offering for photos and other user content

The cloud services are delivered as an application run on a mobile device with a data plan. These cloud services can be delivered by the same provider that provides the device and data plan, or may be provided by a separate provider. The cloud services can be delivered through the browser on the device, locally installed applications, or a combination of the two.

Software-as-a-service is the dominant service type for this marketplace, as users are selecting services to use rather than deploying or managing their own applications, or acquiring their own applications to place in the cloud. The cloud service provider augments these services with services such as personal storage and data backup. These services are presented as complementary SaaS products, rather than mixing raw infrastructure services with the SaaS services.

Enablement of channels is a key activity, including retail and call center channels. Retail is a key channel for setting up accounts, and establishing balances for prepaid accounts, possibly including cash transactions. Retail also enables *try before you buy* and promotion of services without requiring modification of the user device. Multiple channels may be used for their efficiency with different types of interactions. Retail store and call center interactions for initial account establishment can be more efficient for mobile device only users, and a good store interface on a device is usually preferred for service selection and ordering.

Figure 10 shows the products used to deploy this solution.

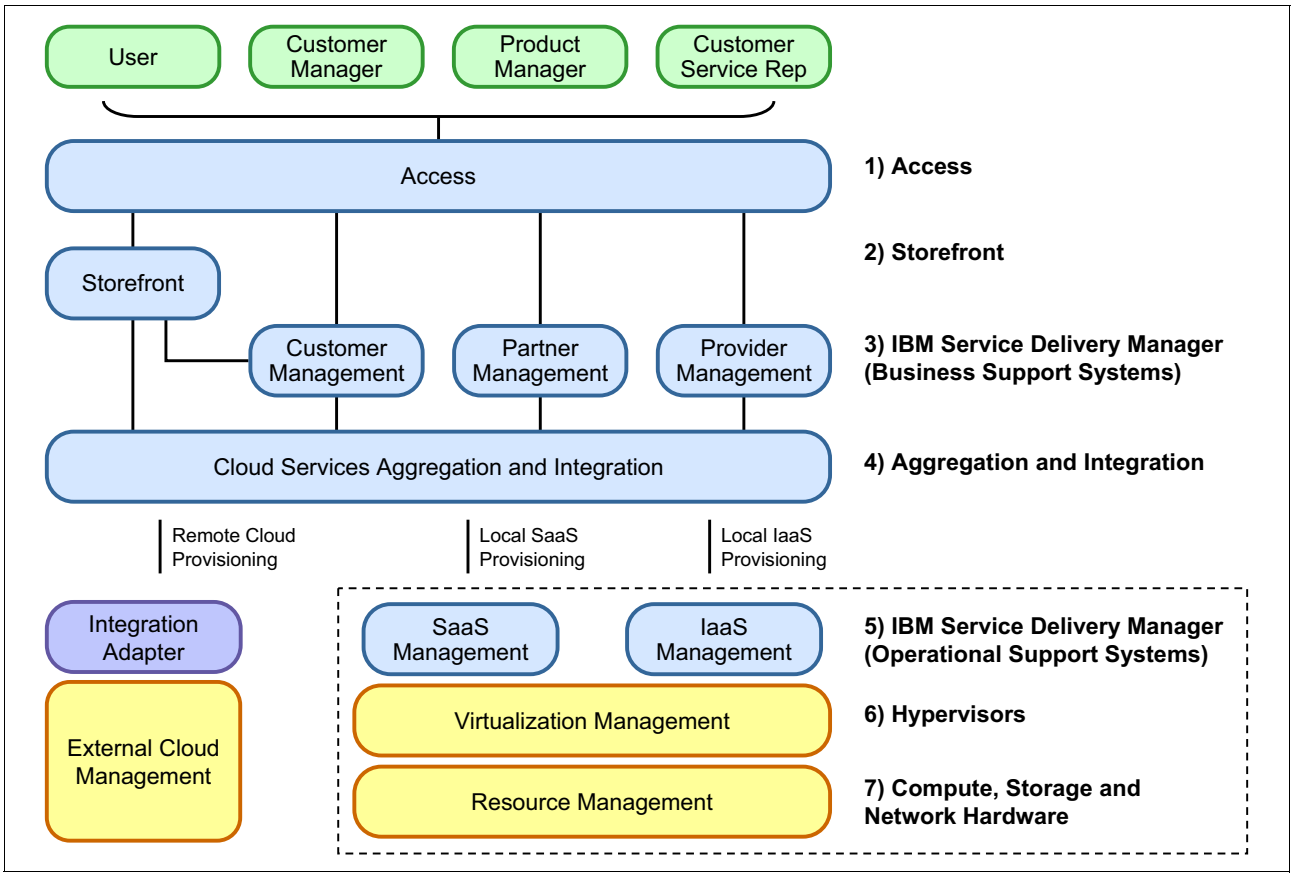


*Figure 10 Solution deployment configuration and products used*

The solution uses products available within the IBM SmartCloud for Service Providers portfolio for the base cloud platform to form the overall solution:

- IBM Service Delivery Manager (Business Support Systems)

  IBM Service Delivery Manager (item 3 in Figure 10) provides the interfaces and manages the data for customer, partner, and provider processes.

- Storefront

  A choice of storefront (item 2 in Figure 10) provides the presentation of the services, service selection, and initiation of the ordering process. IBM Business Partners, such as Parallels and Jamcracker, provide this capability. For service providers that already have a portal, there is an option is to extend the existing portal to add cloud services offerings.

- Aggregation and Integration (item 4 in Figure 10) includes:
  - IBM SmartCloud Aggregator for Service Providers for integration with the cloud computing environments, including aggregation of remote cloud services. This solution is in the IBM SmartCloud for Service Provider portfolio. The solution includes management of the ordering process for cloud services. It uses IBM WebSphere Cast Iron Cloud Integration to provide integration with the remote cloud service providers. It also provides custom integration with an existing customer portal capability to provide account, user, and contract management functions.
  - IBM WebSphere MQ File Transfer Edition or FTP Server (as provided by the operating system) enable loose coupling of the cloud services environment with existing systems to ease integration.

- IBM Service Delivery Manager (Operational Support Systems)

  IBM Service Delivery Manager (item 5 in Figure 10 on page 44) provides the systems management, image, and service orchestration functions for the services offered. It also provides cost management functions for managing financial aspects of service delivery.

- Hypervisors

  Hypervisors (item 6 in Figure 10 on page 44), including solutions such as VMware vSphere and Kernel-based virtual machine (KVM), are commonly supported.

- Compute, storage, and network hardware (item 7 in Figure 10 on page 44) are provided by:
  - IBM compute platforms: The solution can be deployed on System x, System p, and IBM PureFlex System platform offerings. Rack and blade configurations are supported.
  - IBM storage: IBM Storwize V7000 Unified and IBM PureFlex System integrated V7000 storage solutions provides ephemeral and persistent storage for the solution. Additional IBM storage solutions can be used to support specific configuration or cloud services needs.
  - Network: Configurations support networking products such as those offered by Juniper Networks, Cisco Systems, Inc., and IBM Systems Networking to provide network capabilities. Both the internal network within the cloud computing environment and external access networks are covered.

- Access

  Access (item 1 in Figure 10 on page 44) provides the management capabilities and cloud services, and is controlled by functions that include:
  - Authentication and authorization provided by IBM Tivoli Access Manager or an equivalent security service.
  - Identity management provided by IBM Tivoli Federated Identity Manager.
  - Other functions may also be used to provide load balancing, firewall, data caching, and similar capabilities.

In addition to the cloud management functions, cloud services are also included in the solution. These services include:

- The webmail service is configured through aggregation from a third-party provider, using IBM WebSphere Cast Iron Cloud Integration Template Integration Process (TIP) for integration.
- The cloud enabled device management service is installed as a service in the provider cloud, using IBM Service Delivery Manager to provision and manage the service.
- The cloud enabled personal storage service is installed as a service in the provider cloud, using the persistent storage capability and using IBM Service Delivery Manager to provision and manage the service and storage.

This customer scenario, starting with a hosted plus aggregation cloud services capability along with few initial services, is a fairly common first implementation for providers with an early time-to-market priority.

# Summary

IBM provides a comprehensive array of solutions, products, and architecture materials to assist organizations in adopting cloud computing effectively and efficiently.

The IBM Cloud Computing Reference Architecture Cloud service provider adoption pattern provides a comprehensive business and architecture approach for cloud service providers. Using the adoption pattern and the guidance in applying it to your business goals, your organization can build on the experience of cloud computing realized around the world.

The IBM approach to the cloud service provider marketplace is practical and recognizes that the providers must be able to execute nimbly and to establish a marketplace presence quickly. Competitiveness demands a broad architecture approach that enables flexible selection of components to deploy initially and over time. This guide has shown how the adoption pattern reflects these business and technical aspects. It has also shown how to apply these requirements to support initial deployment and growth across various business models.

The architecture has been implemented and deployed for a wide variety of customers. These customers range from small market single purpose cloud service providers to sophisticated (multi-country and multi data center) deployments. Your starting point could be a data center resource offering for infrastructure-as-a-service or enabling a rich set of software and device services to mobile users. In these and other cases, the IBM CCRA adoption pattern provides a sound foundation. As cloud computing becomes a more integral part of the computing fabric, the flexible architectural model provides the robust base platform to grow on.

## Other resources for more information

IBM Redbooks® publications associated with this guide include:

- *Cloud Computing: Save Time, Money, and Resources with a Private Test Cloud*, REDP-4553
- *Cloud Computing and the Value of zEnterprise*, REDP-4763
- *Cloud Security Guidance IBM Recommendations for the Implementation of Cloud Security*, REDP-4614
- *Connect Cloud and On-premise Applications Using IBM WebSphere Cast Iron Integration*, REDP-4674
- *Performance and Capacity Themes for Cloud Computing*, REDP-4876
- *Performance Implications of Cloud Computing*, REDP-4875

For more information about products introduced in this guide, go to these web pages:

- IBM Service Delivery Manager

  http://www.ibm.com/software/tivoli/products/service-delivery-manager/
- IBM SmartCloud Monitoring

  http://www.ibm.com/software/products/us/en/ibmsmarmoni/
- IBM SmartCloud Provisioning

  http://www.ibm.com/software/products/us/en/smartcloud-provisioning/
- *IBM SmartCloud for Service Providers - Launch new cloud services simply and quickly to increase revenue*

  http://public.dhe.ibm.com/common/ssi/ecm/en/tls03047usen/TLS03047USEN.PDF

- IBM WebSphere Cast Iron Integration

  http://www.ibm.com/software/integration/cast-iron-cloud-integration/
- IBM WebSphere MQ File Transfer Edition

  http://www.ibm.com/software/integration/wmq/filetransfer/

## The author who wrote this guide

This guide was produced by a technical specialist working with the International Technical Support Organization (ITSO).

**Joe McIntyre** is a Senior Technical Staff Member (STSM) in IBM Software Group Communications Sector in Austin, Texas, US. He is the architecture team leader for the Communications Sector Industry Frameworks team. He is the architecture lead for the IBM SmartCloud for Service Providers portfolio. Previously, Joe was the lead architect for IBM WebSphere products for the telecommunications industry. He has led standards activities for integration of the Internet and telecommunications networks using web services in 3rd Generation Partnership Project (3GPP), European Telecommunications Standards Institute (ETSI), and Open Mobile Alliance (OMA).

Thanks to the following people for their contributions to this project:

LindaMay Patterson
**International Technical Support Organization, Rochester Center**

Anshu Kak
**IBM Software Group Worldwide Sales, US**

Hamid Khafagy
**IBM Software Group Technical Sales and Services, United Arab Emirates**

Rambabu Parvatina
**IBM GTS Cloud Center of Excellence, India**

Michael Behrendt
**IBM Software Group - Tivoli, Germany**

Tom Recchia
**IBM Software Group - IBM Tivoli US**

Alisa DeMartino
**IBM Sales and Distribution, Software Sales, US**

Bill Lambertson
**IBM Systems and Technology Group, US**

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

**ibm.com**/redbooks/residencies.html

## Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4912-00, was created or updated on December 12, 2012.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at
http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Cast Iron®
Global Business Services®
IBM SmartCloud™
IBM®
PartnerWorld®
PureFlex™
Redbooks®
Redguide™
Redbooks (logo) ®
Storwize®
System p®
System x®
Tivoli®
WebSphere®

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.