Chapter 1

Discovering the Fundamentals of Your Computing Environment

In This Chapter

- Deconstructing cloud concepts
- Discovering resource pools/cloud models and services
- Evaluating the role of the data center
- Finding out how the public cloud fits and when the private cloud shines

How quickly things change. Cloud computing has evolved from a risky and confusing concept to a strategy that organizations large and small are beginning to adopt as part of their overall computing strategy. A few years ago when Hurwitz & Associates wrote *Cloud Computing For Dummies,* there was plenty of skepticism. Would businesses really be willing to adopt cloud computing? What exactly is cloud computing, and how does it help businesses be more effective?

The market has come a long way in a short amount of time. Today more companies — large and small — are doing everything from prototyping a new application with public clouds to implementing a complex private cloud as a utility to support customers and partners. Indeed, we have moved out of the first phase of this new market. We are now at the stage when customers are starting to ask not whether they should think about cloud computing but what types of cloud computing are best able to meet their business problems. Companies are realizing that they need many different types of cloud services in order to meet a variety of customer needs. Therefore, the idea of combining all forms of cloud models in conjunction with data center services as an amalgam or hybrid is the direction computing is headed. An organization deploys a hybrid cloud when it utilizes public and private clouds together with its data center(s) and there are touch points between at least one or more of the deployment models. For example, it shares data between a public and private cloud.

In this chapter, we provide an overview of the hybrid cloud computing environment, including the basics you need to understand in order to move forward in the world of hybrid clouds. First, it's important to understand that an ecosystem of participants define the market. This ecosystem consists of three categories of players: consumers of services, providers of services, and designers of services. In addition, a fourth type of player is a combination of the other three: systems integrators can be consumers of cloud services, they can become a service provider themselves, or they can design services. Here are the characteristics of each:

- **Consumers:** There are different types of consumers. A cloud services consumer might be an individual, or a small business team. Departments in large companies can be cloud services consumers. The IT department can use cloud services either to supplement existing data center services, or to provide specific cloud-based applications such as customer relationship management (CRM) to the sales team. Likewise, even a company that provides cloud services to consumers may use third-party cloud services to supplement their capacity.
- **Service providers:** Cloud service providers are companies that offer packaged services to consumers. Many different types of providers range from those who offer services to individuals and those that serve a broad set of constituents. Many service providers focus on certain markets or certain types of workloads so they can optimize their offerings inexpensively. Thousands of cloud service providers provide public cloud services. Other service providers offer private clouds to support specialized services. A service provider can also be the consumer of a service they acquire to support their customers. Some traditional businesses have taken on the role of becoming a service provider to their customers and partners. These companies are discovering that like professional service providers, they can create a private cloud and offer their own set of services to their customers, which is viewed as a new source of revenue.
- **Service designers:** Companies that can create sophisticated services, tools, and applications to support a variety of cloud models have a huge opportunity. These designers typically build everything from a full Software as a Service (SaaS) platform to tools needed for developers or deployers of cloud services. For example, there is an emerging market for companies that design security and governance offerings to support a variety of cloud models.
- **Systems integrators as cloud service providers:** Systems integrators are getting into the act. Systems integrators are helping customers integrate their data center with public cloud services and private cloud environments. These companies are helping to define best practices and implementation road maps. These integrators can provide private clouds that they host and manage for customers.

Deconstructing Cloud Concepts

Cloud computing is a method of providing a set of shared computing resources that includes applications, computing, storage, networking, development, and deployment platforms as well as business processes. Cloud computing turns traditional siloed computing assets into shared pools of resources that are based on an underlying Internet foundation. Cloud computing makes these resources easier to use by providing standardization and automation. *Standardization* is the implementation of services using a consistent approach supported by a set of consistent interfaces. Likewise, a cloud requires that processes are implemented through the use of automation. *Automation* is a process that's triggered based on business rules, resource availability, and security demands. Automation is required to support a self-service provisioning model. To promote efficiency, automation can ensure that after a provisioned service is no longer needed, it can be returned to the resource pool. This type of automation based on rules can help with capacity planning and overall workload management.

Equally important is that the cloud provides a new economic model of computing. Instead of purchasing, managing, and maintaining a self-contained, traditional *data center* — a specialized environment where a variety of different computing resources are managed — a business is able to transform computing into a more streamlined computing environment that better serves changing computing requirements. If a company has already invested in a data center to support a line of business applications, it can transform that data center into a more targeted environment. If the company wants to provide a flexible self-service resource environment on its premises, it can create what's called a private cloud. Likewise, a company can use sophisticated services available from a third-party public cloud provider (either applications or platforms) to extend and enhance the environment.

Most businesses today are already using some kind of cloud service — even if they don't think of it as a cloud. For example, any company that uses ADP for payroll service is using a cloud-based service. A company may use online data backup or storage services that live in a commercial cloud. If employees use Google's Gmail service, they are using a cloud service. Many companies are discovering that having customer relationship management (CRM) available as a service is a better way to support the sales team than the traditional on-premises software options.

You should be getting the idea that cloud computing means that everything — from compute power to computing infrastructure, and from applications and business processes to personal collaboration — can be delivered to you as a service. To be operational in the real world, the cloud must be implemented with common standardized processes and automation.

Clouds come in different versions, depending on your needs. There are two primary deployment models of cloud: public and private. Most organizations

will use a combination of private computing resources (data centers and private clouds) and public services, where some of the services existing in these environments touch each other — which is what we call a *hybrid cloud environment*.

The public cloud

The *public cloud* is a set of hardware, networking, storage, services, applications, and interfaces owned and operated by a third party for use by other companies or individuals. These commercial providers create a highly scalable data center that hides the details of the underlying infrastructure from the consumer. Public clouds are viable because they typically manage relatively repetitive or straightforward workloads. For example, electronic mail is a very simple application. Therefore, a cloud provider can optimize the environment so that it is best suited to support a large number of customers, even if they save many messages. Likewise, public cloud providers offering storage or computing services will optimize their computing hardware and software to support these specific types of workloads. In contrast, the typical data center supports so many different applications and so many different workloads that it cannot be optimized easily.

The private cloud

A *private cloud* is a set of hardware, networking, storage, services, applications, and interfaces owned and operated by an organization for the use of its employees, partners, and customers. A private cloud can be created and managed by a third party for the exclusive use of one enterprise. The private cloud is a highly controlled environment not open for public consumption. Thus, a private cloud sits behind a firewall. The private cloud is highly automated with a focus on governance, security, and compliance. Automation replaces more manual processes of managing IT services to support customers. In this way, business rules and processes can be implemented inside software so that the environment becomes more predictable and manageable.

The hybrid cloud

A *hybrid cloud* is a combination of a private cloud combined with the use of public cloud services where one or several touch points are between the environments. The goal is to create a well-managed hybrid cloud management environment that can combine services and data from a variety of cloud models to create a unified, automated, and well-managed computing environment. In reality, it will be a number of years before the full range of services becomes the norm. Components of this hybrid management approach are available today, but these are the early days. For a good understanding

of this management model, we recommend you read Chapter 4 about hybrid cloud management. Many companies in the market are working on this model because it will solve critical problems. Why is this necessary? A company might use the private cloud as a way to support changing needs of the application development team within various departments where developers need to constantly build new experimental applications or create new value to meet emerging business needs. A private cloud provides a flexible environment with a higher level of security than would be available in a public cloud. The same company might also use public services ranging from compute, storage, platform, and application services to augment and strengthen their changing business needs.

Combining public services with private clouds and the data center as a hybrid is the new definition of *corporate computing.* Not all companies that use some public and some private cloud services have a hybrid cloud. Rather, a hybrid cloud is an environment where the private and public services are used together to create value. In the following circumstances, a cloud is not hybrid:

- If a few developers in a company use a public cloud service to prototype a new application that is completely disconnected from the private cloud or the data center, the company does not have a hybrid environment.
- If a company is using a SaaS application for a project but there is no movement of data from that application into the company's data center, the environment is not hybrid.

A cloud is hybrid in the following situations:

- If a company uses a public development platform that sends data to a private cloud or a data center–based application, the cloud is hybrid.
- When a company leverages a number of SaaS applications and moves data between private or data center resources, the cloud is hybrid.
- When a business process is designed as a service so that it can connect with environments as though they were a single environment, the cloud is hybrid.

Cloud Computing Elements: Resource Pools/Cloud Models and Services

Now that you have a context for the types of cloud environments, it's important to understand the common elements required to make clouds functional. In this section, we give you the basics of what you need to know. Figure 1-1 illustrates the related elements that come together to create clouds. On the bottom of the diagram is a set of *resource pools* that feed a set of cloud delivery services. On the top of the diagram are the common service elements needed to support these delivery models.

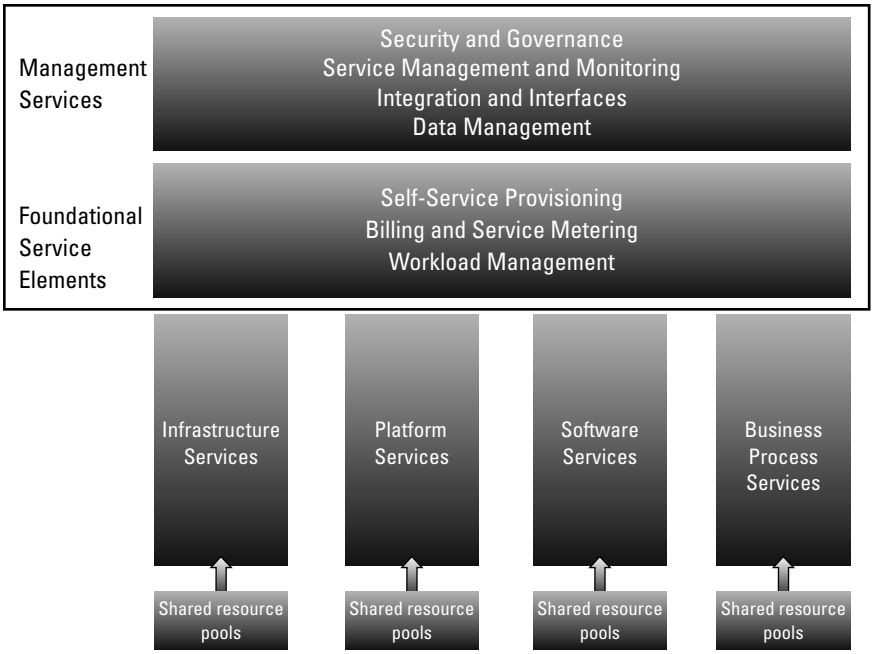

Figure 1-1: How the cloud elements fit together.

One of the fundamental differences between cloud computing and traditional computing is the way a cloud is designed to manage workload resources. Managing workloads is foundational to the cloud. For more details on workload management, you may want to read Chapter 12. Whereas a data center is designed to manage applications, a cloud is intended to manage a pool of resources, which is precisely what it sounds like — a set of shared, configured services that are independent of a physical location. In many situations, cloud service providers create a multi-tenant environment to support the deployment of these resources. Multi-tenancy enables the sharing of a service while keeping the data and configurations of individual customers separate.

For example, say that you are a cloud provider. You don't want your customers to have to select one specific server or one storage system for their workloads or data. Instead, you want the customers to be abstracted from that idea so that they can simply say, "I need some more storage" and those storage resources are pooled together from various physical systems to create a set of resources. In other words, customers don't need to know which storage system they're accessing. These shared resource pools are shown at the bottom of Figure 1-1.

To make resource pooling work, each pooled element needs to be written with service-oriented constructs in mind. Consequently, each resource is written as an independent service without dependencies and with well-defined interfaces. For more details on service orientation, turn to Chapter 3.

Understanding the foundations of cloud computing calls for an understanding of the four different cloud delivery models:

- Software as a Service
- Platform as a Service
- Infrastructure as a Service
- Business Process as a Service

Part II has a chapter on each of these models, but here we cover the basics on the different cloud computing delivery models. These are illustrated as infrastructure services, platform services, and software services in Figure 1-1.

Infrastructure as a Service

Infrastructure as a Service (IaaS) is the delivery of services such as hardware, storage, networking, data center space, and various utility software elements on a request basis. There are both public and private versions of IaaS. In the public IaaS, the user simply needs a credit card to acquire these resources. When that user stops paying, the resource disappears. In a private IaaS service, it is usually the IT organization or an integrator who creates an infrastructure designed to provide resources on demand to internal users and sometimes business partners. Whereas criteria for a public cloud are based primarily on the ability to pay for a service, a private service applies company policy to a service request. IaaS is the fundamental element used by other cloud models. Some customers will bring their own tools and software to create applications. For more details on IaaS, turn to Chapter 5.

Platform as a Service

Platform as a Service (PaaS) is a mechanism for combining IaaS with an abstracted set of middleware services, software development, and deployment tools that allow the organization to have a consistent way to create and deploy applications on a cloud or on-premises environment. A PaaS offers a consistent set of programming and middleware services that ensure developers have a well-tested and well-integrated way to create applications in a cloud environment. A PaaS environment brings together development and deployment together to create a more manageable way to build and deploy applications. A PaaS requires an Infrastructure service. For more on PaaS, see Chapter 7.

Software as a Service

Software as a Service (SaaS) is a business application created and hosted by a provider in a multi-tenant model. Customers pay for the service either

per user on a monthly or yearly contract model. The SaaS application sits on top of both a Platform as a Service and foundational Infrastructure services. Typically, these underlying services aren't visible to end users of a SaaS application. To read more details on SaaS, turn to Chapter 6.

Business Process as a Service

In traditional computing environments, business processes are often built into the application. For example, there may be a payment service or a service that sets up a delivery service request. However, one of the benefits of the movement to cloud computing is that organizations can use cloud-based business process services across the organization. Standardizing on business process services reduces the chance of error and makes it easier for an organization to ensure that the right services are used for the right purpose. If your company is using only a few cloud services, you may wonder why you need to worry about business process services. However, because organizations will use more cloud services over time, business leaders will want the ability to move data among various cloud applications and data center services. So, having a plan for thinking about business process services from a cross-environment perspective will help streamline your organization's movement toward cloud services. These BPaaS can become a revenue source for both cloud consumers and service providers.

Foundational services supporting cloud resource models

No matter what type of resource pool is involved, some core capabilities are essential in the cloud environment, including the following:

- Elasticity and self-service provisioning
- Billing and meters of service usage
- Workload management

These are referred to as foundational service elements in Figure 1-1.

Elasticity and self-service provisioning

One benefit of the cloud is that customers have the potential to access as much of a service as they need when they need it. If they need computing services for a month, they can rent that service for that time period. They don't have to buy a system to handle the need. Although this potential clearly benefits customers, it makes life more complicated for providers of services. In the cloud, the service provider can't know in advance what a customer might need at any given time. So, service providers (whether a public provider or a private

cloud provider) must create an environment that can scale to support a diverse set of customer requirements. A provider of cloud services needs to design a platform architecture that is optimized for the service the company is providing. The platform must be designed so that the users/customers provision resources such as computing or storage resources that they pay for on a per-unit basis. When the user no longer needs that resource and stops paying, the resource is released back into the pool of resources.

There are situations when a service provider can't anticipate the needs of a customer. Therefore, it is common for a service provider to add additional capacity from a third-party service provider. Typically, the consumer is unaware that they are dealing with an additional cloud service provider.

Self-service is one of the most important capabilities of cloud computing. With self-service, the developer of an application, for example, is able to use a browser or portal interface to acquire the resources needed to build an application, which is dramatically different than how a developer works in a data center. In the traditional data center model, a developer must request resources from IT operations, which might require IT to acquire more servers and storage and to purchase the required software. At this point, the system has to be configured to meet the development project requirements. Assuming that the cloud service is designed to meet the required goal, the developer can avoid these complicated steps. We go into more detail about how cloud computing helps with software development in Chapter 7, which focuses on Platform as a Service.

Billing and metering services

A cloud service has to provide a way to measure and meter a service. Consequently, a cloud environment includes a built-in service that tracks how many resources a customer is using. Customers consuming public cloud resources will be charged for a unit of usage — whether that is a CPU hour or an amount of storage. Customers using software as a service will be charged for each registered user on a per-year basis. In a private cloud, employees probably won't be charged directly, but usage will be tracked. In many companies, a department's budget will be tied directly to the amount and type of services it uses.

Managing workloads

As we mention earlier in this chapter, a cloud isn't a single unified environment; rather, it's a combination of resources that could be spread across systems and geographies. The cloud is a federated environment that brings together resources so that they can work together. To make this happen in an organized manner requires an organization of workloads. A *workload* is an independent service or collection of code that can be executed. So, you need to think about the cloud as a group of workloads that are managed as though they were a single cohesive environment. It is important in a cloud environment that workloads be designed to support the right task with the right cloud services. For example, some workloads will need to be placed in a private

cloud because they require fast transaction management and a high level of security. Other workloads may not be so mission critical and can be placed in a public cloud.

When organizations begin thinking about cloud computing as a strategy, they must do more than simply go to a public or private resource and leverage those services. Like any computing environment, cloud computing requires that workloads be balanced and managed, regardless of the resource pool they're based on. Within a well-designed cloud environment, workloads can move among resource pools in different cloud environments. Because, in the real world, you will use a combination of services, it's important to think not just about an individual workload but also about a combination of workloads and how they interact with each other and with collaborators.

Management services

Many other management services are mandatory for ensuring that cloud computing is a well-managed platform that supports customer needs, including ensuring that a variety of workloads are supported in the right way. This is the case no matter what the cloud deployment and delivery model. Some core services are illustrated at the top of Figure 1-1. Security and governance are key services to ensure that your applications and data are protected. Data management in a hybrid environment is also critical since data will be moving between cloud environments. Moreover, the right level of service must be provided to meet consumer requirements. This is the case regardless of whether you're dealing with a public, private, or hybrid model. That's why hybrid cloud management and monitoring is so important. Finally, since various services will need to be integrated in a hybrid cloud, the interfaces between clouds are also important.

Examining the Role of the Data Center

What happens to the data center when companies begin to implement hybrid clouds? First, the data center does not go away. After all, almost all medium-size and large companies run their own data center — which is how many companies operate their systems of record, including accounting systems, payroll, human resources applications, and line of business applications, to name a few. Many data centers have grown in an unplanned manner over many decades. The typical data center supports different hardware architectures, operating systems, applications, and hundreds, if not thousands, of different tools. To make matters worse, a lot of the money spent in supporting a data center is used for maintenance of existing systems, heating, air conditioning, floor space, and labor.

So, it's not surprising that many companies have taken the time to streamline their data centers through technologies such as server virtualization. In

essence, virtualization decouples the software from the hardware. In decoupling, the software is put into a separate container so that it's isolated from the underlying operating system. (See Chapter 16 for more details on how virtualization works.) With the use of virtualization, data center management can more easily and efficiently manage the way applications are placed on servers.

However, even though IT has made the data center more efficient, cloud computing has made it apparent that more can be done to transform computing. Organizations are beginning to take a hard look at what the centralized data center is well suited for and at the changes required to create a computing environment that truly serves the needs of their business.

Companies are beginning to discover that they can have the best of all worlds by finding the tasks that are most appropriate for a highly controlled data center and which workloads are best suited to either a private or a public cloud.

Although a well-designed and well-tuned data center provides essential services to a company, it is often best suited for a complex line of business applications. These are often transaction-intensive applications that need to confirm and track the movement of financial transactions among customers, suppliers, and partners. Additionally, large, often highly customized systems of record are and will continue to be data center–based. These applications are typically tightly monitored for corporate governance and compliance.

The key difference between the traditional role of the data center and the new role of the data center is just beginning to be clear: The traditional data center is changing from a general-purpose repository for all applications to a highly tuned corporate asset for a class of applications. This is explained in Tables 1-1 and 1-2.

Table 1-1 Traditional Data Center

Type of Data Center	*Description*
Traditional Data Center	Data center with all applications and data centrally managed
Consolidated Data Center	Virtualization applied to enable server consolidation supporting existing model

Table 1-2 Next Generation Data Center

Type of Data Center	*Description*
Purpose-Built Data Center	Highly tuned data center for enterprise systems of record
Purpose-Built Private Cloud	Self-service resource available for developers and partners

Traditional data center

One way to think about where we are today with data centers is to think about the car garage in the typical suburban home. When the home is first built, the homeowners use the garage for its intended use — storing two cars. For the first year or so everything works well, but over time the garage is used to store the lawn mower, boxes of books, and random things that are no longer used. Soon there is little room in the garage for even a single car. Figure 1-2 illustrates this concept. The built for purpose garage is no longer efficient and no longer serves its purpose, but the family still needs to store cars and still needs room for everything else. The solution could be to build a shed in the backyard to store yard tools and other stuff. What is the point of this example? It mirrors what has happened to the data center over time.

For the past 30 years, organizations have had traditional data that has been in a centrally managed environment for a huge variety of applications and company requirements. Consequently, the typical data center has become more complicated, expensive, and cumbersome to manage. This situation has resulted in data center consolidation, where IT management has taken a hard look first at what applications are actually needed.

Rethinking the data center

IT management has begun to pare down the number of applications being used, and to use virtualization to consolidate workloads and remove superfluous hardware. Traditionally, it was complicated to determine how much physical space an application would require in the data center. Rather than take the risk that an application would not have enough room to support customer requirements, companies simply bought additional capacity in order to ensure performance. However, this approach to applications management became unsustainable. In essence, through virtualization technology, IT has been able to add a software layer that allows applications to be more easily consolidated onto specific servers. But even these moves have not been enough.

The purpose-built data center

Just as the homeowners reclaimed the garage, IT is beginning to make the transition to purpose-built data centers. These new generation data centers are tuned to support the key systems that manage a business. In some situations, key applications are simply maintained in their current state because of the way they're used for stable, mission critical business processes. In other cases, the IT organization has taken the time to rearchitect these applications

into a set of modular services that can be used to support many different business initiatives. For example, a company might create a service as the authorized way to pay a supplier for a service. By creating a single service that is used across many business units, the data center becomes a more efficient environment. In addition, the emerging purpose-built data center doesn't try to provide all services for all needs. It becomes the source of the business services and data services that apply to the cross-organizational needs.

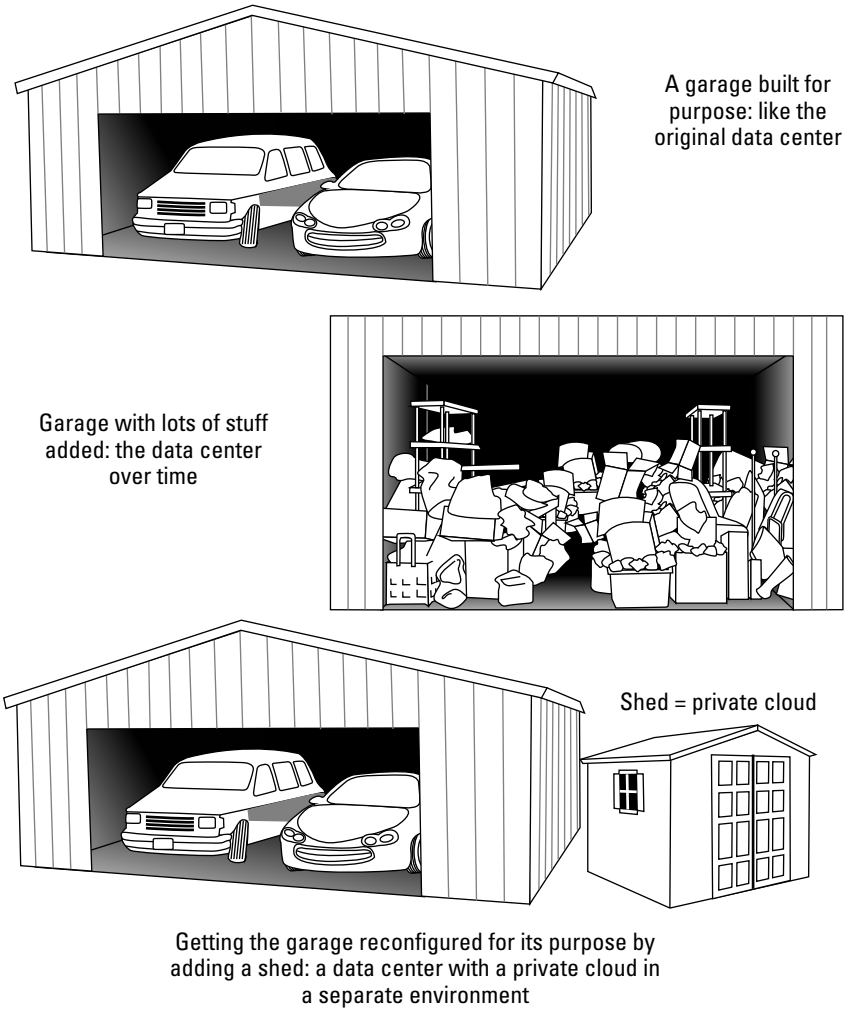

Figure 1-2: The original data center evolved over time just like the typical home garage. The cloud brings back the idea of a fit for purpose environment.

Purpose-built private cloud

IT organizations have discovered that it's much more efficient and effective to create private cloud services for developers to create new applications and services. Therefore, companies are setting up a highly automated computing environment enabled with a self-service portal. This portal is often designed with business process rules that dictate what services a developer or an authorized partner can use. For example, a developer beginning to develop a new application may be permitted to use the Java language, specific types of middleware, and a specified amount of computing capacity and storage. Once the project is completed, there may be a rule that automatically returns the capacity back to the pool of resources. The private cloud service is intended to support an organization's need for speed and agility based on fast-changing company requirements. Having a private cloud available for projects allows a company to easily experiment with new ideas and new applications without having to request funding for a project that might not become a reality.

Seeing how the public cloud fits

Say that your company is a retailer with a vibrant and well-used transactional portal for selling products online. The biggest event for the company is its yearly 50 percent–off sale. During a two-week period, the company has an agreement with a public cloud company to supplement its capacity so that performance is always consistent. Your company has a contract with the public cloud provider to provide capacity services during the sale as well as other times when additional resources are required. Because your company can rely on the public cloud provider, it doesn't have to buy additional servers for incremental usage, thus saving time and money.

Your company has selected two different software applications from service providers to replace existing on-premises applications. One application is a customer relationship management (CRM) that allows the sales force to easily get access to prospective and current customer information from a cloud-based public service. In addition, your company uses a human resources management platform as a service. The company has implemented integration software that allows data to be managed among the data center and the two SaaS environments. Both of these SaaS applications have enabled your company to avoid purchasing additional hardware and software that would require IT management. In addition, because the sales team can access their data much faster from any device they're using, from a business perspective, the sales team's performance is much more effective.

Your company is in a new market where getting a series of new services operational quickly has the potential to leapfrog the competition. Your company hasn't been in business very long. Because the company can use a public service, it can make services available before more established companies can act. At the same time, the company can use its data center to monitor the effectiveness of services, manage data privately, and combine with other services that aren't visible to customers.

Knowing when the private cloud shines

Now, say that your company's products and services are offered in a portal to your business partners. These products are key to a company's revenue. For example, you might be part of a financial services company that offers a key business service that is purchased by various banks around the globe. Because this service is a revenue source and because the company is a well-run and orchestrated computing environment, it makes sense for the company to establish a private cloud to support this business model. An architected private cloud that supports this business initiative is ideal for the financial services company. Using a public service where the company is charged on a time or usage basis wouldn't be cost-effective.

Your company has a large team of developers spread across the world. You decide to create a private cloud that allows any authorized developer to gain access to data, tools, and processes required to create applications. You have set up rules to ensure that developers can access only the tools and data they're supposed to use. The private cloud makes the process of managing a distributed development team more productive and more cost-effective.

What might your cloud deployment look like in the future?

Let's imagine that the year is 2020 and you have implemented a next generation computing environment based on re-envisioning the data center and how you use cloud resources. What might your computing environment look like? Let's say your company is a retailer with stores in 20 different countries and has a healthy online commerce business. In the past you supported three very large data centers in different geographies. You also used a smattering of SaaS applications for sales automation, marketing services, commerce services, human resources, and package tracking. You also used some public cloud storage services to handle transactions during peak holiday seasons.

Inside the data center you implemented server virtualization to reduce the energy and space you required. You also took some time to get rid of applications that weren't used anymore but somehow were still running in the data center.

While all of these changes helped make the computing environment better, none of the changes gave the company the competitive edge that management wanted to see. Management wanted to go further by having the IT organization become a service provider to its customers and business partners. This new business model was intended to help the company become more proactive in responding to opportunities before competitors.

So, the CIO in consultation with the CEO decided to make some radical changes. The first change was to rethink the data center itself. The company transformed the data center into a streamlined environment to manage a line of business applications. Therefore, the data center now included the ERP system, and customized commerce systems that were tailored specifically to the type of retail business. In addition, a specialized system that was used to manage the manufacturing process of

(continued)

(continued)

special-order merchandise was also managed in the data center. Finally, there were four other specialized systems that had been built 15 years ago that were necessary but used only sporadically. With fewer systems, it was possible to optimize the data center to support these workloads much more efficiently. The company was able to purchase new hardware designed specifically for these types of workloads. They were more efficient systems. Networking, storage, and even power, cooling, and floor space were redesigned. More automation was implemented throughout the environment.

To support the needs of the development organization, the company established a private cloud environment. The private cloud was designed specifically to support the development of new applications and prototypes of new innovative ideas that might become products in the future. The private cloud also provided storage, networking resources to support new business services created as a profit center for new business services. Business subsidiaries were able to procure computing services from a self-service portal. IT set up clear business rules that defined what services each business unit was allowed to procure based on the project definitions.

The company also selected some public cloud services to support the business. To support the sales and marketing team, the company continued to use SaaS applications. In addition, the company decided that selecting a public cloud service for e-mail was quite economical, so the company got rid of its internal e-mail system. At the same time, a number of innovative SaaS collaboration services made the process of working with partners a streamlined environment.

The IT team spent the bulk of its time on integration and architectural issues. The team discovered that it was just as easy to create silos with SaaS applications as it was in the data center. This was a difficult lesson to learn. Therefore, the team continues to spend its time on ensuring that the new data center environment includes data and process integration across the data center, the private cloud, and the public cloud services. To manage the environment, the organization invested in service management, security, and governance.

The process of transforming computing was complicated. There were mistakes because either technology was immature or the organization did not consider how all the components of IT needed to be orchestrated. But the IT management team was able to take a step back and create a road map and implementation strategy that took a holistic view of how the next generation of IT would help the company grow in new directions supported by a highly effective and flexible computing environment that was becoming the engine for innovation and business growth.