Clearing up the Cloud: Adoption Strategies for Cloud Computing

by Ed Reynolds and Charlie Bess

Let's get to the point — there is a lot of hype about "all things cloud." Almost anything that can be deployed with a degree of flexibility is claiming to be "cloud." There are several key questions that enterprises should address when considering the impact of cloud computing services:

- How will the business and technology dynamics change in a cloud-connected world?
- How can enterprises use these emergent capabilities to create value and dominate their market?
- What are the key issues for future network-based computing models?
- What are the high-level approaches that enterprises can pursue to improve operational efficiencies and create new business value using the cloud?

In this article, we will focus first on what makes up a cloud solution and then go into a few basic tenets, starting with a business-level discussion of cloud computing and its impact on organizations.

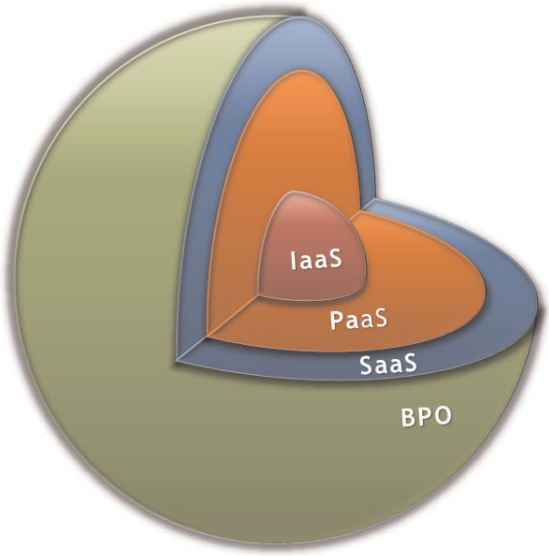

Figure 1 — Layers of the cloud.

LAYERS OF CLOUD COMPUTING

Cloud computing is not just one thing. It is made up of multiple layers, like an onion (see Figure 1). This "onion" has at its core **infrastructure as a service** (IaaS). This is the layer of leveraged hardware resources that can be billed on a consumption basis. Typically, the user has visibility into the underlying operating systems for solutions at this layer. It is usually implemented as a virtualized operating system image that may be moved from hardware device to hardware device without the user's knowledge. It can also be implemented within an organization as a private utility, but since it is a dedicated model, that would preclude it from being called a cloud solution.

The next layer out is **platform as a service** (PaaS). This adds a layer of flexibility and abstraction that usually hides the underlying operating system. These solutions include workflow facilities for solution design, development, testing, and deployment and may also include application services such as team collaboration, Web service, and database integration. These services are provisioned as an integrated solution, usually over the Web with no direct, synchronized contact between the supplier and consumer. A PaaS solution normally hides the complexities of the underlying hardware devices as well as the movement of workload between devices. Software generally has to be written specifically for the PaaS provider and the interface they provide.

The next layer of the cloud is **software as a service** (SaaS). This commonly builds upon the flexibility of the PaaS environment, but since only software interfaces are exposed, it may not be evident how the underlying infrastructure is delivered. SaaS is usually billed on a transaction or user access basis. SaaS's centralized control model enables an ongoing revenue stream for the software creator or supplier, while allowing end consumers the flexibility of paying only for what they consume. The provider performs upgrades, with little control on the part of the end user. If interfaces are implemented correctly, the upgrades should not

affect the production environment. The interfaces to the business model are usually SaaS provider– and service-specific.

Business process outsourcing (BPO) is the outermost layer that conforms to the basic requirements to be a cloud offering. Although not considered "cloud" by some, it is multitenant and billed on a usage basis. This layer could also be called **business process as a service**, since it includes the people, systems, and infrastructure needed to address a business need (such as payroll) for an organization.

When you look at cloud as including all these layers, it can also be viewed from an **everything as a service** (XaaS) perspective. Even though you don't have to have all the underlying layers to be cloud, the most flexible solutions will have them all.

Organizations will need to write software in a different fashion than they do today to take advantage of the appropriate layer of abstraction. Each of these layers is unique in both its benefits and its requirements.

TENETS OF CLOUD

Let's establish some basic tenets to begin clearing up the confusion around cloud computing:

1. Enterprises will use the cloud and cloud computing technologies to add value or reduce the cost of business operations. The cloud is a means to an end, not an end in itself.
2. Cloud computing will not replace private data centers for traditional computing any time soon. Cloud is a selective and additive capability for an enterprise IT portfolio.
3. The future of enterprise IT will be a hybrid environment for the foreseeable future, composed of dedicated and cloud-based resources and services.
4. The cloud is the next stage in the evolution of the Internet, moving beyond transporting information as a service to acting upon it as well.
5. In the long term, the cloud will change the way enterprises use information services.
6. Cloud will accelerate the split of enterprise IT into two classes:
   - Common utilitarian functions
   - Core systems for competitive advantage
7. Enterprise adoption of cloud computing technologies and services is limited by perceived and real barriers stemming from requirements for security, privacy, trust, and quality.
8. Cloud is about spreading resources across as wide a range of consumers as possible.

THE ENTERPRISE CLOUD FORECAST

The manner in which an enterprise employs cloud computing depends on its goals and objectives. Stated in the most basic terms, there are two key questions:

1. Are you trying to reduce cost or add value?
2. Who will benefit from the use of cloud computing — the IT group or business units in the enterprise (including clients or customers)?

We can address these two questions with the familiar four-quadrant chart, which we'll call the Enterprise Cloud Computing Forecast (see Figure 2).

The Enterprise Cloud Computing Forecast shows four broad strategies that enterprises can employ for the adoption of cloud computing:

1. The **replacement strategy** uses the cloud to replace or rehost existing IT systems. The use of cloud resources for workloads such as testing/development, collaboration (Web 2.0), and systems for startup companies is proven and more acceptable. It takes the virtualization strategies currently being deployed by many organizations to a new level of flexibility and reduced cost.
2. The **augmentation strategy** uses the cloud to expand the capacity or reach of existing IT systems. This approach (also known as "cloud bursting") has been employed before, with solutions that have been custom crafted using grid and cluster computing technologies. It involves identifying specific computing models that require significant resources but only for limited periods, making them impractical to own.

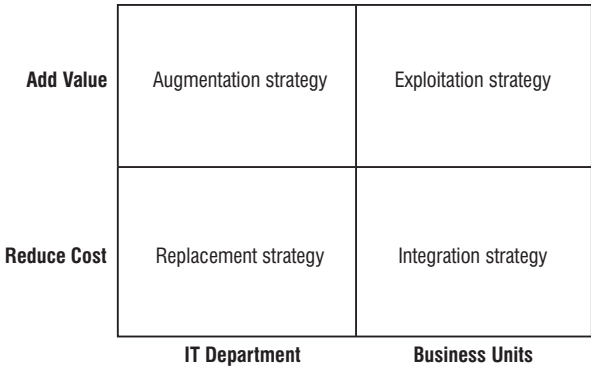

Figure 2 — The Enterprise Cloud Computing Forecast.

3. The **integration strategy** uses the cloud to create more efficient and effective connections within organizations as well as with clients and suppliers. This goes beyond Web 2.0 and social networking. True inter-enterprise integration requires a richer and deeper set of information connections, which the cloud will ultimately enable.
4. The **exploitation strategy** uses the cloud to generate new revenue streams by leveraging existing corporate intellectual property and information. The ability of businesses to generate money from the cloud is a real possibility going forward, but it remains a dream for most organizations at present, since the tools to support them in this kind of effort are immature and custom to the organization today.

The Replacement Strategy

Cloud computing is clearly a hot topic these days, with many of the attributes of a fad. There is an overwhelming amount of information, some of which is conflicting, circulating around the industry. CIOs and other business leaders are trying to keep up with these trends, but it is difficult to separate aspirations from reality.

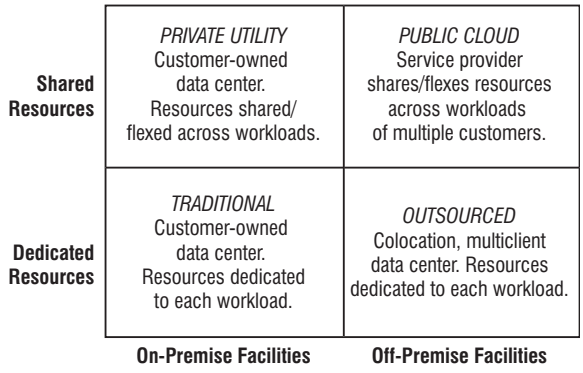

Figure 3 — Infrastructure delivery options.

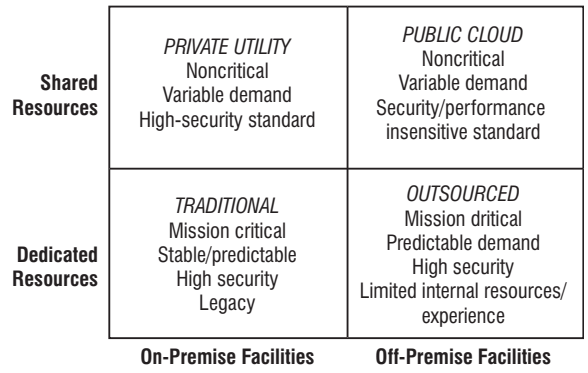

Figure 4 — Workload placement.

The objective of the replacement strategy approach is to reduce current IT costs by replacing existing inhouse services and resources with lower-cost cloud-based resources and services. The combination of the current economic situation along with claims of raw IT services at pennies an hour is driving much of the initial interest in using the cloud as a replacement for private IT infrastructure. What we've seen is that when enterprises inquire about the benefits of cloud computing, they may not realize the alternatives available to deliver similar results. The focus is on the flexibility provided by cloud computing, such as:

- Less investment in physical assets
- Resources that are pooled and shared across workloads
- Scalability and flexibility
- The ability to pay for resources as they are consumed

It is important to take a holistic perspective on all of the various application and infrastructure options available to move toward these goals. Organizations need to recognize that the majority of these benefits can be achieved through inhouse and hosted environments by employing other techniques such as virtualization and automation.

The options for sourcing infrastructure services can be classified along two broad dimensions:

1. Provisioning of resources, either shared or dedicated to applications
2. Facilities that are either on premises or off

Figure 3 depicts these typical choices. The selection and placement of workloads among these styles of infrastructure services is shown in Figure 4.

The **traditional** private data center has been, and still is, well suited for mission-critical applications that require stringent security and have stable and predictable resource requirements. These systems and applications are typically mature legacy applications and COTS packages. As such, there is little to no comparable cloud infrastructure available for these older platforms. Dedicating resources and staff to these workloads makes sense given the stable and predictable nature of the processing and the amount of effort required to port them to a cloud platform. Keeping these systems where they currently reside will likely be more cost efficient than a variable payment approach, although outsourcing or a private utility approach may be an option.

Outsourced facilities provide managed services for applications and systems that are found in the traditional data center model. It is beyond the scope of this article to explore the numerous business drivers for choosing to outsource an enterprise's data center to third-party providers. Generally, the use of outsourcing provides additional resources (hardware, software, and intellectual property) and technical skills/experience. It provides value to the enterprise by leveraging economies of scale and the reach offered by the third-party outsourcing provider. Once again, if the software is licensed or written on specific platforms that require special operations support, it will likely be more cost effective to leave them where they are, although a private utility may still be an option.

The **private utility** model provides infrastructure services for systems and applications that have highly variable demand cycles, such as market-facing Web applications. Some refer to this model as an "internal cloud" or "private cloud," and it is based on the use of technologies such as virtualization and automation within a private data center, to meet the varying business needs of newer applications and systems. A key factor for the effective use of a private utility model is a high degree of standardization within the infrastructure and the application delivery platform. A single operational management standard is necessary to enable the sharing of resources among different applications and workloads. While the use of virtualization has been available for some time, many enterprises are still in the beginning stages of deploying it, choosing to start its use with relatively noncritical applications, such as testing/development environments, internal Web sites, social networking, and collaboration facilities.

The **public cloud** has only recently emerged as a viable option for enterprise IT services, and many still have significant concerns related to issues such as security, privacy, data integrity, lock-in, availability, and so on. The case for public cloud services varies greatly between large enterprises on the one hand and small businesses and startup companies on the other. The latter have little or no embedded infrastructure or applications, and thus the availability of business applications and low-cost, ready-made infrastructure is a compelling proposition.

In contrast, large, multinational enterprises have existing applications and services that have been modified and upgraded over many years to serve the evolving needs of their businesses. They usually have stringent service-level agreements (SLAs) and auditing requirements that are difficult for cloud providers to meet.

Large enterprises need to assess the value of public computing services against the "integration costs" of overcoming these barriers in order to incorporate the services into the enterprise IT service portfolio.

Except for some very basic Web applications, today's enterprise-class business systems cannot be easily moved to a public cloud model. Ultimately, business applications will need to be rearchitected and rewritten to achieve the expected level of performance, scalability, and economic advantage associated with cloud computing. New projects should be evaluated for their ability to take on this more flexible approach.

The Augmentation Strategy

If your enterprise were given a share of a large-scale supercomputer, what would you do with it? Many organizations have not considered this question in the past since it was not economically feasible for them to own or operate a massive supercomputing facility. But now cloud computing services allow an organization to acquire large-scale computing resources on a flexible payment basis, creating new opportunities to improve business performance or gain a competitive advantage.

These opportunities are enabled by a unique attribute of a cloud services–enabled utility computing business model; namely, the notion of "cost associativity." Researchers at the University of California, Berkeley note that using 1,000 machines for one hour has the same cost as using a single machine for 1,000 hours (see Figure 5).[1] The significance of this is twofold: (1) You don't have to buy/build/maintain 1,000 machines, which you would not be able to justify, and (2) you can

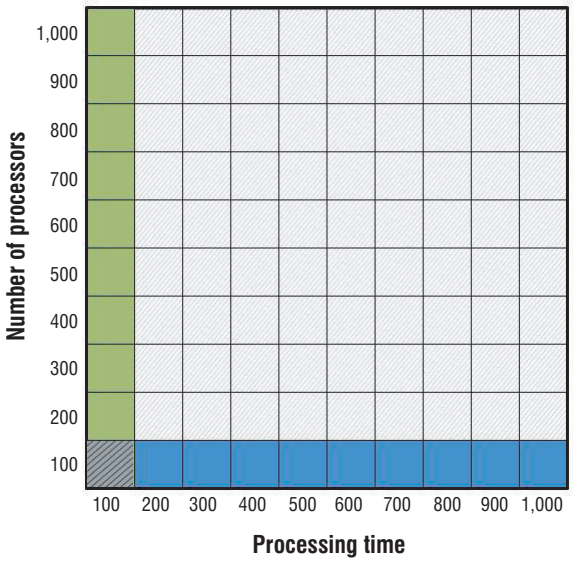

Figure 5 — Cost associativity. (Source: Michael Armburst et. al.)

increase the processing speed 1,000-fold with no increase in total cost (assuming the application or process can take advantage of the added capacity). In this age of rapid decision making, in which the need to reduce latency is acute, cost associativity may be significant to how organizations generate business value.

In effect, cloud services will enable a reengineering of industry ecosystems.

Businesses are collecting more and more information today through the expanded use of IT in their enterprises and also through more dynamic connections to their customers and suppliers. There are many opportunities to process and analyze this information to create a competitive advantage, either in the creation of new products and revenue or through more efficient and timely access to markets. Cloud computing resources can be employed to process customer data and preferences more frequently to drive new revenue from sales and/or to reduce expenses through more efficient processes. Since analyzing this information may require 100 or even 1,000 times the normal load, it would be impractical to retain all these computational resources internally.

As Cutter Senior Consultant Curt Hall observes, "Recent developments have brought together parallel processing and cloud computing technologies in such a way that they are set to change the way organizations look at analyzing massive amounts of data."[2] By augmenting existing IT resources with cloud-based capacity, enterprises can apply this combined power to perform predictive simulations for financial modeling, product design, marketing, and other applications. The net effect is a "more is more" approach of applying massive amounts of computing power to gain competitive advantage and add value to the enterprise.

The Integration Strategy

Today it's no longer a case of you against the world, but instead it is you and the network you can bring to bear against your competitors that will define the winner. At first glance, this approach may seem similar to many former "inter-enterprise integration" proposals, such as supply-chain integration, enterprise information management, business-to-business (B2B) services, and so on. However, we believe that cloud-based services will ultimately enable enterprises to drive improved efficiencies in ongoing operations through the richer and tighter integration of information systems and business processes. In effect, cloud services will enable a reengineering of industry ecosystems. In certain cases, the control point or hub of business activity will be strengthened and amplified by new-cloud based information services. In other cases, these control points will be disintermediated and nullified, as various entrepreneurial approaches seek out new connections and bypass mechanisms. Cloud-based services and SOA will be used to create a series of more thoroughgoing information connections among enterprises — truly tying internal systems together — a sort of "enterprise level mashup" but with more rigor and governance than that term implies.

We can use a "logical data center" analogy to explore this concept further. Think about a model in which an enterprise and all of its key partners, suppliers, distributors, and customers are located in the same "data center." How would you rearchitect the information flows and transaction exchanges if everyone was located in one logically secure facility (i.e., a secure cloud)? How would you use techniques like Lean Six Sigma to remove inefficiencies and interim deliverables? What are the opportunities to remove waste in the process by taking advantage of very tight connections in a secure services architecture?

Many systems today produce "outputs" in some interim summary form that are then passed to other business constituents in a supply chain or ecosystem. Are these interim phases truly needed, or can they be replaced with real-time connections to exchange information and transactions as they occur? In this model you send the "line items," not the summary.

The Exploitation Strategy

One of the traps that many organizations fall into with regard to cloud computing is thinking that it will be used for the same types of solutions as those for which we currently use our dedicated computers. In reality, many of the solutions we have today can be delivered perfectly well on a dedicated platform and continue to do so in the future. There are other techniques, however, that can only be delivered in an abundant computing environment.

There are changes taking place that will drive a different approach to computing. More and more data is being gathered by organizations through the use of

sensors, RFID tags, and other information-gathering techniques at the edge of the enterprise. Digital data is being gathered and fed into organizations at an exponentially increasing rate, with the result that many organizations are, as IT trends analyst and frequent Cutter contributor Brian J. Dooley notes, "strained by the increasing challenge of managing terabytes of data, which are doubling every year."[3] We know the information is out there somewhere — if we could only find it. Cloud storage techniques will allow organizations to store and share this information like never before.

The reason organizations are collecting all this information is that they want greater understanding of what is going on so they can make more rapid decisions. In the last 20 years, the lifespan of products, product lines, and even companies has decreased, and any insight organizations can take advantage of is sorely needed. Anything that can reduce latency must be investigated.

Cloud computing has the potential to usher in an "age of abundance" in computing, providing opportunities to process vast amounts of information in ways that were previously unaffordable or impractical. This will allow for deeper insight and transparency into the operation of the organization, enabling the movement from latency-rich sense-and-response behavior to a much more proactive cause-and-effect. Through modeling, organizations can see where a situation will lead and thus address it early, taking latency out of enterprise response. This will require significant amounts of computing resources for short durations, which may not be practical without resorting to cloud techniques.

We've traditionally designed systems to facilitate normal behavior, viewing change as a disruption. With cloud computing, the analysis — and in some cases the action — for normal situations can be highly automated. The advantage comes to those who can recognize the exceptions and focus their people on turning the anomalies into opportunities. The idea should be to automate the people out of "normal," where they can only add variation.

One of the issues with many of today's solutions is that they are not easily moved into a parallel processing environment, which is the structure of cloud computing. Modeling, simulation, and pattern recognition are ideal for the cloud environment. They can take advantage of both the massive amounts of computing power and data available, allowing organizations to respond more quickly, to be more confident of the eventual outcome of their actions, and to develop a deeper understanding of both the cause and its effect.

There are many ways to use this new computing abundance. We will be limited only by our understanding of how to gain benefit from it. Where the data center is located and who owns the hardware within it will be of less concern than the advantage it provides to the enterprise, but even with all this automation, it will be the people and their understanding of the business that will make the difference.

Anything that can reduce latency must be investigated.

CLOUD CONSIDERATIONS

Since cloud computing is still in the early adoption stage, enterprises must also consider the challenges of moving from a traditional environment to a more service-centric delivery model. These challenges are:

- **Service-level agreements.** If the services are going to provide value, the SLAs need to cover the entire range of business value generation (end-to-end) and not just within the data center.
- **End-to-end management.** Enterprises will use cloud/SaaS to achieve a business objective, requiring end-to-end visibility and control.
- **Security integration.** Currently there is little federated/integrated security support in SaaS or cloud computing. Integration of the enterprise's security with outside services is critical for seamless performance.
- **Enterprise integration.** Today, cloud or SaaS tends to be standalone, with little ability to integrate into existing enterprise IT systems. We believe that cloud services will need to be integrated into a greater holistic service for enterprises, requiring the ability to manage, control, and audit the entire "system" of cloud services.
- **Customization.** Some cloud providers focus on standard solutions and services that are flexible, but not how value is created with that flexibility. Information technology for business needs to be based on generating value for the enterprise, not just cutting costs or making the system easy for the service provider to maintain.

FINAL THOUGHTS

> In bad times, there is a tendency to draw inward, to look out for yourself, and to protect what you and yours (family, company, department) have. But this is also a time to think boldly.
>
> — Ken Orr, Cutter Business Technology Council Fellow[4]

During this economic downturn, many organizations are hunkered down waiting for it to blow over — but there is value to be gained in evaluating cloud computing and its impact on your enterprise. There are bound to be some interesting discoveries made in the process of reviewing your IT portfolio of services and delivery systems. Even if there are no immediate actions to take, a thorough assessment will prepare the organization to think differently about creating new IT solutions and effectively train staff in new emerging capabilities, as these capabilities will most likely mature and be adopted over time.

Perhaps the worst thing an IT organization can do during the current recession is to cut staff and expenses too deeply, thereby eliminating its ability to think strategically. This approach squanders the knowledge and expertise that will be required to exit the downturn with momentum instead of inertia.

Cloud computing is one technology that enables cost cutting (for short-term advantage) as well as increased flexibility (for responding to the recovery). Those who grasp these long-term issues will invest in understanding today what the dividends will be tomorrow.

ENDNOTES

[1]Armburst, Michael, et. al. *Above the Clouds: A Berkeley View of Cloud Computing*, Technical Report No. UCB/EECS-2009-28. Electrical Engineering and Computer Sciences, University of California at Berkeley, 10 February 2009 (www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html).

[2]Hall, Curt. "Hadoop, MapReduce, Cloudera, EC2, and BI." Cutter Consortium Business Intelligence *E-Mail Advisor*, 14 April 2009.

[3]Dooley, Brian J. "Business Intelligence Optimization." Cutter Consortium Business Intelligence *Executive Update*, Vol. 9, No. 5, 2009.

[4]Orr, Ken. "Coming Out the Other Side: Keeping Your Head Up in Bad Times." *Cutter IT Journal E-Mail Advisor*, 15 April 2009.

Ed Reynolds and Charlie Bess are EDS Fellows in EDS's CTO office. The title of EDS Fellow is awarded to the corporation's most innovative thought leaders in recognition of their exceptional achievements. Collectively the authors have over 50 years of experience in IT, and they have each been with EDS for over 20 years. In their current roles, Mr. Reynolds and Mr. Bess are responsible for assessing emerging technologies and developing EDS's future services. They are also regular contributors to EDS's The Next Big Thing *blog. The authors can be reached at ed.reynolds@hp.com, charlie.bess@hp.com.*