Planning for Cloud Services from Telecommunications Service Providers

By ***Jim Metzler, Ashton Metzler & Associates***
Distinguished Research Fellow and Co-Founder,
Webtorials Editorial/Analyst Division

Abstract

The widespread adoption of cloud computing presents a tremendous opportunity for Telecommunications Service Providers to leverage their networks, data centers, existing customer relationships and other business assets to generate significant new revenue streams derived from delivering public cloud computing services to enterprises of all sizes.

The first step that a service provider should take in planning for a rollout of cloud computing services is to prepare a detailed, multi-year, rolling business plan that outlines both their cloud computing strategy as well as their data center architecture and implementation strategy. The second step is to simultaneously engage in ongoing conversations with potential customers and potential suppliers of key enabling technologies. The third step is to update and adjust their rolling cloud computing business plan in an iterative fashion. Each iteration or phase of the business plan is based in part on the service provider's ongoing conversations with potential customers and suppliers and in part on the desire and the ability of the service provider to leverage wherever possible the technology base and market adoption of the previous phases.

One of the goals of this white paper is to create an outline for a cloud services business plan and to identify and discuss some of the major topics that should be included in the plan. It is expected that service providers will modify this outline for use in their environment. Another goal of this white paper is to identify some of the input that needs to be gathered from both customers and suppliers in order for service providers to be able to update their business plans in such a way that they continually reflect the realities of the burgeoning marketplace for public cloud services.

Introduction

Gartner estimates that between 2010 and 2015 enterprises will spend $112 billion cumulatively on Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS), combined[1]. The great interest that enterprises of all sizes are expressing for public cloud computing services stems from the fact that these services provide a wide range of business and operational benefits including:

- Enhancing business agility
- Linking expenses to resource utilization rather than to resource capacity

[1] http://www.gartner.com/it/page.jsp?id=1389313

- Gaining access to a much richer range of services and applications both at a lower cost and in a shorter time frame than was traditionally possible

Cloud computing delivers these benefits in large part through the pervasive virtualization of data center resources combined with the integrated and automated management of both the infrastructure and the services that it supports. This combination of technologies makes it possible for IT services and applications to leverage pools of resources in order to dynamically reallocate or scale capacity to efficiently meet fluctuations in user demand.

The adoption of cloud computing presents a tremendous opportunity for Telecommunications Service Providers to leverage their networks, data centers, existing customer relationships and other business assets to generate significant new revenue streams. These new revenue streams are based on delivering high value-added cloud-based services to enterprise customers. From the enterprise perspective, the phrase *high value-added* equates to the service provider offering levels of security, availability and performance that closely resemble the service levels that enterprise IT organizations have traditionally provided themselves. Providing stringent levels of security, availability and performance differentiates these services from the best effort services that typically are targeted at small businesses and home users. Service providers can also leverage their experience in operational scaling and assured performance over Next Generation Networks (NGN) to differentiate and add unique added value to their services and they can further differentiate their service offerings by providing meaningful SLAs.

Creating a Cloud Services Deployment Planning Process

Telecommunications Service Providers have historically taken a very detailed, well-planned approach to deploying new services. While that approach is often valuable, it has significant limitations in a market, such as the market for public cloud computing services, that is currently somewhat ill defined and which is likely to evolve significantly over the next several years. An alternative approach is to avoid any attempt at detailed planning and instead implement a series of service trials. There is no doubt that this alternative approach will provide a wealth of valuable input. There is also no doubt that this approach on its own will make it difficult, if not impossible, for service providers to leverage the technology and marketplace synergies of their services offerings.

What Telecommunications Service Providers need is an iterative planning process that balances the benefits of detailed planning with the benefits of continually gathering and leveraging marketplace intelligence. With that in mind, the first step in the service delivery planning process is for service providers to prepare a high level, multi-year business plan that outlines both their cloud strategy and their data center architecture and implementation strategy. This plan should include the definition of the services or applications to be offered, an outline of how the services will evolve over time, the key enabling technologies and products, the target market, the expected adoption rates as well as estimates of the overall ROI and the expected profitability for each scenario. ROI and profitability will depend largely on wide market adoption, the economies of scale derived from delivering essentially the same service to a large number of subscribers and the efficiency of their cloud infrastructure in supporting both a wide range of services and a large number of subscribers from a common resource pool.

The second step in the planning process is for service providers to simultaneously engage in ongoing conversations with potential suppliers and potential customers. The purpose of engaging in ongoing conversations with suppliers is to identify technologies and products that enable the service provider to successfully offer new services. For example, there are several prominent vendors who provide software that would enable a service provider to offer a desktop virtualization service. Service

providers who are considering offering that service should work with these vendors to understand both the strengths and weaknesses of their current solutions as well as their product roadmaps.

The purpose of engaging in ongoing conversations with customers, and possibly leveraging those conversations to create trials of new services, is to identify the specific services, and the primary characteristics of those services, that customers are interested in acquiring from a public cloud computing provider. Staying with the example of a virtual desktop service, service providers who are considering offering this service should work with customers in their target market to identify the key characteristics that these customers want to see in a desktop virtualization solution. For example, do the customers want a single screen and Microsoft software, or do they need a more highly functional solution?

The third step in the planning process is for service providers to update and adjust their rolling cloud computing business plan in an iterative fashion. Each iteration or phase of the business plan is based in part on the service provider's ongoing conversations with potential customers and suppliers. For example, knowing both the roadmap of the key vendors of virtual desktop solutions as well as the detailed needs of their customers enables a service provider to refine their plans for offering a virtual desktop solution. Each phase is also based on the desire and the ability of the service provider to leverage wherever possible the technology base and market adoption of the previous phases.

Market Adoption: The Enterprise's Evolution to the Cloud

At the first stage in their evolution, enterprises are likely to take advantage of basic Infrastructure-as-a-Service (IaaS) solutions such as hosted dedicated servers, virtual machine instances on shared servers and storage services. As previously noted, by offering levels of security, availability and performance that closely resemble the service levels that enterprise IT organizations have traditionally provided themselves, Telecommunications Service Providers can differentiate their basic IaaS solutions from commodity services that are targeted at small businesses and home users.

As early adopters of basic compute and storage services become satisfied with those services, they will likely evolve to become consumers of additional cloud based services. This section of the white paper will describe some of those services. There is no intention in this section to indicate the order in which these services will be adopted, as that will vary based on factors such as the size, type and geographic location of the enterprise customers.

One class of cloud based service that is of interest to enterprise IT organizations is Value-Added Infrastructure as a Service (Value-Added IaaS). The phrase *Value-Added IaaS* refers to infrastructure services that come integrated with additional software or hardware functionality in order to provide a complete solution that addresses a well-defined business need. Examples of Value-added IaaS services include:

- Disaster Recovery and Business Continuity (DR/BC)
- High Performance Computing (HPC)
- Virtual Desktop Integration (VDI)

Many of the early adopters of Basic IaaS and Value-Added IaaS are coming from outside the purview and control of the enterprise's IT organization. This includes research and engineering departments as well as groups that are driving highly visible, special projects. These consumers typically have their own budget for IT services, are self-reliant, technically proficient and expect rapid deployment of services.

IT organizations typically approach cloud services as another form of out-tasking and follow a tightly controlled procurement and sourcing process that is designed to minimize the risk that is associated with adopting any new service. While they can lag behind more nimble groups in their adoption of Value-Added IaaS, IT organizations are beginning to show significant interest. For example, a recent research report[2] presented the results of a survey in which one hundred and sixty four IT professionals were asked to indicate how likely it was over the next year that their company would acquire a Value-Added IaaS. Their responses are shown in Table 1.

Table 1: Interest in Value-Added Iaas

	Likelihood of Acquisition				
	Will Not Happen	**Might Happen**	**50/50 Chance**	**Will Likely Happen**	**Will Happen**
Disaster Recovery	30.8%	23.8%	20.0%	11.5%	13.8%
Virtual Desktops	40.7%	24.4%	18.5%	9.6%	6.7%
High Performance Computing	41.9%	24.8%	16.3%	10.1%	7.0%

The data in Table 1 shows that there is significant interest in Value-Added IaaS. For example, over the next year more than sixteen percent of the survey respondents are either planning to acquire, or are likely to acquire each of the Value-Added infrastructure services listed in the table. This represents the beginning of what could be a fundamental shift in terms of how IT services are provisioned.

At the same time that IT organizations are beginning to adopt cloud based services that are more sophisticated than basic IaaS, most of them are also beginning to implement architectures within their data centers that improve the leverage that they gain from virtualization. These private data centers will continue to be used for at least the foreseeable future for certain business-critical workloads or for workloads that involve data that is subject to strict regulations and restrictions. However, as the evolution to adopt cloud computing continues, many enterprises will decide to transition their on-site private data centers to private cloud data centers that are hosted and/or managed by a service provider. A private cloud data center would typically be implemented using resources dedicated to a single tenant in order to provide security and control equivalent to that of traditional enterprise private data centers.

As enterprise IT organizations continue to implement cloud related technologies within their data centers, these organizations are positioning themselves to be able to leverage the higher scalability and cost effectiveness that are associated with hybrid cloud solutions in which the private data center or the hosted private cloud data center is supplemented by a virtual private cloud data center solution offered by a service provider. A virtual private cloud data center is typically hosted at the service provider's data center and is based on a highly virtualized, multi-tenant, shared architecture. One of the key characteristics of a hybrid cloud is the enterprise's ability to continue to manage the provisioning of resources, as well as to monitor and control application performance across the private/public boundary.

[2] The 2011 Application and Service Delivery Handbook, http://www.webtorials.com/content/2011/07/2011-application-service-delivery-handbook.html

As previously noted, as early adopters of basic compute and storage services become satisfied with those services they will likely evolve to become consumers of additional services. Another class of cloud based service that will be attractive to IT organizations is vertical SaaS applications, such as applications commonly used in the financial services, legal services and healthcare industries. While vertical SaaS applications potentially represent a large revenue stream, providing this class of service will require that Telecommunications Service Providers acquire expertise that they haven't had before. A class of cloud based services that is a closer fit with the skills and expertise of Telecommunications Service Providers is horizontal SaaS applications. This class includes applications such as CRM and ERP and similar to the case with vertical SaaS applications, offering applications such as these will require that Telecommunications Service Providers acquire expertise that they haven't had before. However, there are a number of horizontal SaaS applications that are widely acknowledged to be closely aligned with the skills and expertise of Telecommunications Service Providers. This includes applications such as VoIP, Unified Communications, Network and Application optimization and network management. The previously referenced research report, The 2011 Application and Service Delivery Handbook, reported on the results of a survey in which the survey respondents were asked to indicate how likely it was over the next year that their company would acquire a horizontal SaaS application that was network centric. Their responses are shown in Table 2.

Table 2: Interest in Horizontal SaaS Applications

Likelihood of Acquisition

	Will Not Happen	Might Happen	50/50 Chance	Will Likely Happen	Will Happen
VoIP	34.3%	17.5%	12.6%	15.4%	20.3%
Unified Communications	26.1%	26.8%	16.9%	14.8%	15.5%
Network and Application Optimization	33.8%	22.1%	14.7%	14.0%	15.4%
Network Management	38.8%	26.6%	7.2%	17.3%	10.1%

One of the interesting observations that can be drawn from the data in Table 2 is that over a quarter of the survey respondents indicated that within the next year that their organization either would, or likely would acquire each of the solutions listed in the table.

The Service Provider's Evolution to the Cloud

As described in a preceding section, Telecommunications Service Providers should develop an iterative planning process for a phased rollout of cloud services with each phase designed to leverage wherever possible the technology base and market adoption of the previous phases. Table 3 can serve as a template for specifying how cloud services could be rolled out in a series of successive phases timed to the market[3].

[3] The four phases shown in Table 3 is just an example. The number of phases in the rollout plan will vary by service provider and type of service.

A service provider may choose to offer only a subset of the services listed in Table 3. For example, some service providers may decide to not offer Basic IaaS solutions feeling that this market is too commoditized to provide an acceptable ROI. Also, as was the case with the adoption of cloud based services on the part of the enterprise, the service provider's rollout of services doesn't need to be in any particular sequence. For example, a service provider may decide to offer VPDC services and Horizontal SaaS before rolling out Private Cloud data center services. The most appropriate sequencing of service deployments for a particular service provider will be determined in large part based on the conversations that the service provider has with their suppliers and customers.

Table 3: Phased Introduction of Cloud Services

	Basic IaaS	Value-Added IaaS	Private Cloud DC (dedicated)	Virtual Private Cloud DC	Horizontal SaaS	Vertical Industry SaaS
Phase I						
Phase II						
Phase III						
Phase IV						

Below is a description of each of the service types contained in Table 3.

Basic IaaS Services

These services include basic storage as well as physical and/or virtual server instances combined with storage, load balancers and possibly other physical or virtual appliance-based services. Access is usually provided over the Internet and the typical pricing model is on-demand; e.g., $/hour for servers or VMs and $/month for specified amounts of storage. Additional charges often apply for bandwidth usage and load balancing.

Value-Added IaaS

These services represent a higher level of value that service providers can layer on top of the infrastructure that is needed for Basic IaaS services. These higher value-added services typically require the service provider to provide a complete hosted solution for BC/DR, VDI, HPC or similar services.

Similar to the previously discussed situation with virtual desktops, service providers that are considering implementing HPC solutions need to understand how their target customers prefer to consume this solution. One option is that the customer would provide the application and the cluster computing middleware and the service provider would provide dedicated or virtual servers as well as a high performance cluster interconnect that meets the customer's bandwidth and latency requirements for LAN connectivity amongst the servers. Service providers that are considering rolling out a cloud based, disaster recovery service must work with their target customers to answer questions such as whether the potential customers want a hot stand-by, a warm stand-by, a cold stand-by or some combination of solutions. It is also important to determine if there are any constraints that would impact the user's interest in the service. For example, due to strict data privacy laws, most European companies would have constraints both on where their data can be stored and where it can be moved to in a failure situation.

Pricing for Value-Added IaaS is normally usage-based with a monthly subscription period. Access could be provided either over the Internet or a private WAN service such as MPLS.

Private Cloud Data Center Services

These services are based on outsourcing the enterprise's multi-tier private data center to a service provider. The data center could be located at either a site controlled by the enterprise or at a service provider's site. In most cases service providers will structure these services so that the customers receive the highest levels of support, as well as assurances written into the corresponding SLA for high levels of availability, performance and security. A private WAN service would typically be used to provide access to these services.

Virtual Private Data Center (VPDC)

These services provide an instance of an entire data center hosted on a service provider's infrastructure that is optimized to provide a high level of security and availability for multiple tenants. From the service provider's perspective, the data center architecture for the VPDC would be similar to the architecture used for a private cloud data center except that the resources would be shared among a number of customers rather than being dedicated to a single customer or tenant. The service provider's architecture needs to effectively leverage virtualization in order to maximize the efficient usage of a shared pool of resources. The architecture also needs to allow for a high degree of flexibility in providing a broad range of required network capabilities. This includes WAN optimization, load balancing and firewall services. As noted earlier, service management software must enable the co-management of the VPDC by customers and providers.

The hybrid cloud computing model works best in those instances in which the VPDC and the private cloud data center are based on the same hypervisors, hypervisor management systems and cloud controllers. This maximizes the enterprise's control over the hybrid cloud and allows application and server management to remain the responsibility of the enterprise. Access could be provided either over the Internet or a private WAN service.

Horizontal SaaS

These applications have traditionally been provided directly from the independent software vendor's (ISV's) data center using a version of the application that has been architected specifically for multi-tenancy. Popular horizontal SaaS applications of this type include communications, collaboration, business productivity, sales management, CRM and ERP.

ISVs are expanding the number of times that they host their SaaS applications at data centers provided by IaaS service providers where the multi-tenant separation is provided by infrastructure virtualization, or in some performance-driven cases (e.g., data bases), a combination of virtual and dedicated resources. One option that service providers have is to provide VPDC services to ISVs who want to offer their software in a SaaS model without the burden of managing a multi-tenant data center.

Through partnerships with an ecosystem of ISVs, service providers can also offer SaaS applications hosted at an enterprise's VPDC or an enterprise private cloud data center that is provided by the service provider. The service provider could potentially offer an extensive catalog of horizontal SaaS applications from a variety of ISVs. Additional added value could be based on the service provider offering integration services among popular SaaS applications.

Since SaaS is a relatively new area for many service providers, the introduction of services in this category might span two or more phases of the service provider's rollout plan for cloud services. For telecommunication service providers, the most logical starting point would be to build on prior experience with hosted services such as VoIP, Unified Communications and Collaboration. In a subsequent phase of their rollout of cloud services, more complex SaaS applications, such as personal productivity applications, enterprise-grade CRM, ERP, and data base software could be offered.

Vertical SaaS

These applications provide industry-specific solutions. While industry-specific application software represents a fairly large share of the overall software market, it accounts for only a relatively small share of the existing SaaS market. Vertical SaaS solutions are expected to be more widely adopted when service providers can offer pre-integrated suites of business applications and services that are tightly focused on a vertical industry market, such as health care or education.

Cloud Services Differentiation

The successful introduction of cloud services will depend largely on the service provider's ability to build a portfolio of cloud services that appeal to a wide range of enterprise customers that potentially are in a variety of industry market segments. Since the market for commodity cloud services will be highly competitive, service providers should focus on services that can be differentiated in terms of offering higher added value to the customer. Differentiation can come from the creation of innovative services developed in cooperation with ecosystem partners or from focusing on enterprise concerns regarding manageability, security, availability and application performance.

Service providers who have invested in, or have access to Next Generation Network (NGN) WAN technologies, including Ethernet MAN/WAN transport, MPLS and Layer 2 and Layer 3 MPLS VPNs have a unique opportunity to differentiate their services in terms of application performance and end-to-end SLAs. The NGN offers a wide range of advantages over the Internet as a means of accessing cloud services. NGN capabilities including CoS/QoS, traffic engineering, resiliency/restoration and Layer 2 and Layer 3 VPNs that can be leveraged to support assured network performance. The impact of that assured network performance is reflected not only in network-centric SLAs, but also in application-centric and business process-oriented SLAs. The service provider can further optimize these higher level SLAs by complementing the existing network performance with WOC functionality and by implementing cloud infrastructure capabilities that assure performance within the cloud data center portion of the end-to-end path.

Cloud Services Infrastructure

As previously noted, the successful rollout of cloud services is dependent on the service provider creating a cloud data center architecture and infrastructure implementation that is optimized for the delivery of the planned evolution of cloud services. The data center architecture should be based on pervasive virtualization to provide the flexibility that is required to support a wide range of dynamically provisioned services. Pervasive virtualization includes the virtualization of servers, storage, switches and routers. It also includes the virtualization of appliances such as firewalls, intrusion detection systems (IDS) and intrusion prevention systems (IPS), WAN optimization controllers (WOCs) and application delivery controllers (ADCs).

In addition to pervasive virtualization, the cloud data center architecture should have the following characteristics:

Simplicity: The data center design should standardize on one vendor for each class of data center device and it should only support a limited number of device models and operating system releases within each product set. Reduced complexity generally equates to lower operational expense.

Scalability and Flexibility: A design based on modular building blocks can help provide the scalability required to support rapid growth in the number of subscribers and the breadth of services that can be offered. The infrastructure should be able to support very high virtualization ratios (i.e., VMs per server) and high server densities.

High Availability: Sufficient redundancy and resiliency features should be incorporated into the design to support the availability-related SLAs that are expected of high value added services such as VPDC.

Trusted Multi-tenancy: A robust multi-layer security architecture is required to ensure confidentiality and integrity of the services and the subscriber's data in a multi-tenant environment.

Integrated and Automated Infrastructure and Service Management: Integrated management reduces the number of management interfaces that are involved in implementing administrative workflows. Automation allows services to be dynamically provisioned, modified or scaled without requiring time-consuming manual configuration across the various technology domains of the data center; e.g., compute, network, storage and security. The management suite should also include the application and service level management capabilities that will support the various types of end-to-end SLAs that were discussed in the previous section. From an operational management perspective, the management system should provide additional capabilities, such as cross-domain root cause analysis and service impact analysis, in order to support the highest levels of service reliability.

Support for Enterprise Co-Management: The service management system should provide a web portal supporting the self-service provisioning of new services or the scaling of existing services. The portal should also include dashboards that provide real-time visibility into the performance of applications and services as well as into the consumption of on-demand services.

Compatibility with Enterprise Cloud Implementations: As noted earlier, the efficiency of hybrid clouds is optimized where there is a high degree of consistency across the private and public portions of the solution in terms of the cloud management systems, the hypervisors and the hypervisors' management systems. This consistency facilitates the movement of VMs between enterprise data centers and service provider data centers, and this movement enables the dynamic reallocation of cloud resources. In order to achieve this compatibility, service providers should build their cloud data centers by utilizing the same mainstream technologies that are being used to build the vast majority of enterprise virtualized data centers.

After the basic cloud data center architecture is formulated, the planning for the infrastructure can begin. For each phase of service rollout being planned, the required infrastructure to support these services can be specified, as shown in Table 4. The goal is to maximize the level of integration between technology domains and the degree of management automation across the entire infrastructure in order to control both CAPEX and OPEX. Across the various phases of the rollout, the

goal is to ensure a smooth evolution and scaling of the infrastructure to support new services with the appropriate levels of service quality, including security and performance SLAs.

Table 4: Infrastructure Investment to Support Service Rollout	Servers	Storage	Data Center Network	Virtualization	Mgmt	Appliances & vAppliances
Phase I						
Phase II						
Phase III						
Phase IV						

Building Clouds with VCE Vblock™ Infrastructure Platforms

The Virtual Computing Environment (VCE) company is a joint venture of Cisco, VMware, and EMC. VCE offers packaged, pre-integrated, converged infrastructure solutions that are referred to as Vblock™ Infrastructure Platforms. Vblock platforms are optimized to support data center implementations that are intended to deliver highly secure and scalable cloud services. Vblock platforms come in a range of sizes and are comprised of compute, network, storage, virtualization and management products from the parent companies of VCE.

When a service provider implements Vblock platforms, all of the system integration functions that would normally be required across technology domains that correspond to columns 2 through 7 in Table 4 have already been performed by VCE. In addition, the management solution for Vblock platforms unifies element and service management across the infrastructure while also providing a high degree of management automation. The Vblock converged infrastructure solutions are assembled from the following industry leading product families:

- Compute: The Cisco Unified Computing System blade server platform unites computing, networking, storage access and virtualization into a cohesive system consisting of a single element management domain.

- Network: Cisco Nexus Ethernet switches support full network virtualization, high scalability and continuous systems operation. Cisco MDS 9000 series FibreChannel switches provide FibreChanel SAN connectivity.

- Storage: EMC Symmetrix or EMC Unified Storage (i.e., Celerra, CLARiiON) provide highly reliable and scalable storage with flexible multi-protocol support.

- Virtualization: VMware vSphere, vCenter, and vCloud Director provide a virtualization platform that supports infrastructure and application services within private, public, and hybrid cloud computing environments.

- Infrastructure/Service Management: Unified Infrastructure Management (UIM) Provisioning Center (PC) allows an infrastructure service to be provisioned with a single mouse click. UIM Operations Center (OC) provides operational management of Vblock platforms including in-depth operational analysis across all components of the Vblock platform.

Some of the unique benefits of Vblock Infrastructure Platforms include:

- **Simplicity:** Vblock platform components (e.g., Cisco UCS and EMC Unified Storage) offer a high degree of inherent integration and flexibility, which greatly simplifies the infrastructure. In addition, a Vblock platform is managed as a single element vs. a collection of disparate elements. Simplicity is reflected in low Vblock platform TCO.

- **Rapid Deployment:** Vblocks platforms are pre-integrated, pre-tested and certified. This reduces the service provider's time to deployment and time to market for both new services and new customers.

- **Pervasive Virtualization:** Virtualization capability is built into every aspect of the Vblock platform.

- **Support:** A single source of support is provided by VCE for all aspects of the Vblock platforms. A wide range of support services is available to assist organizations plan and build cloud data centers.

- **Pre-qualification:** Pre-qualified configurations are available for popular applications; e.g., SAP, Oracle RAC, VMware View, Microsoft Exchange and SharePoint.

- **Synergy:** Vblock platforms are highly compatible with the mainstream technologies enterprises are using today to evolve their data centers toward a private cloud computing model. Enterprises that are already using the component technologies of Vblock platforms in their private data centers are prime candidates to adopt VPDC services and managed private cloud data center services from service providers with Vblock platform-based public cloud infrastructures.

Summary and Call to Action

Telecommunications Service Providers have a tremendous opportunity to capitalize on the growing industry-wide adoption of cloud computing. For example, Telecommunications Service Providers are in a unique position to take a leadership role in the public cloud computing market by providing enterprises with high value-added cloud services that also offer the provider the potential for significantly increased revenues and high profit margins.

However, in order to realize the opportunity that cloud computing presents to them, Telecommunications Service Providers must develop a cloud services deployment planning process. An effective plan for the deployment of public cloud services involves three major activities that are designed to balance the benefits of detailed planning with the benefits of continually gathering and leveraging marketplace intelligence. The first activity is for service providers to create a high level, multi-year business plan. This business plan should include the definition of the services or applications to be offered, an outline of how the services will evolve over time, the key enabling technologies and products, the target market, the expected adoption rates as well as estimates of the overall ROI and the expected profitability for each scenario.

Another key component of this business plan is the creation and ongoing modification of the service provider's data center architecture and implementation plan. This architecture and implementation plan must be able to support the range of planned services over multiple phases of evolution. Data center planning should place an emphasis on maximizing service profitability by controlling both CAPEX and OPEX and on supporting value-added service attributes, such as superior SLAs for security, data integrity, availability and application performance. The Vblock platforms from VCE are an example of an enabling technology that has been specifically designed as an optimal solution for supporting a wide range of scalable cloud services.

The second step in the planning process is for service providers to simultaneously engage in ongoing conversations with potential suppliers and potential customers. The purpose of engaging in ongoing conversations with suppliers is to identify technologies and products that enable the service provider to successfully offer new services. The purpose of engaging in ongoing conversations with customers, and possibly leveraging those conversations to create trials of new services, is to identify the specific services, and the primary characteristics of those services, that customers are interested in acquiring from a public cloud computing provider. The third step in the planning process is for service providers to regularly update and adjust their rolling cloud computing business plan based in large part on their ongoing conversations with vendors and customers.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.