Growing Your Cloud Infrastructure: Planning, Design and Operation

Abstract

Cloud computing services are expanding and evolving rapidly. But with this fast, large-scale growth comes a complex supply chain of service providers, often involving multiple third parties. In addition, quality standards and security practices in the industry are still unclear, difficult to define, and even harder to verify. To succeed in this uncertain and challenging environment, organizations need a carefully planned, structured approach to the design, implementation, and management of cloud services.

NOTE

Drawing on its Bell Labs heritage, Applied Communication Sciences (ACS) excels at creating innovative technologies and services to solve the most difficult and complex information and communications problems across commercial, carrier and government sectors. Applied Communication Sciences is legally registered as TT Government Solutions, doing business as Applied Communication Sciences. ACS is a wholly-owned subsidiary of The SI Organization, Inc.

Rapid Growth, Lingering Challenges

The cloud is rising. In 2010, IDC estimated the cloud market at $16B and predicted growth to $56B by 2014. Other market researchers are even more optimistic, with Gartner estimating the cloud market at $150B by 2013 and Merrill Lynch predicting growth to $160B in 2011. But no matter where the final number lands in the next few years, one thing is certain: Planning a move to the cloud isn't a breeze. It involves navigating a whole world of service models, providers, vendors, and markets.

Still, organizations are expected to keep migrating to cloud services in droves as they look to the cloud's promise of increased business agility and decreased operational costs. But even as their migration intensifies, debate and uncertainty continue to swirl around cloud computing definitions, use cases, underlying technologies, issues and security risks. What's more, migration typically requires entrusting mission-critical data to a complex supply chain of third-party providers – and relinquishing at least some control over network performance, reliability and security. Even after moving to the cloud, an organization's private information may bounce between multiple providers, creating ongoing security and management challenges.

To ensure the quality of every aspect of the cloud supply chain, organizations need a carefully planned, structured approach to the design, implementation and management of cloud services. From choosing the right service model to developing effective service level agreements (SLAs) to verifying compliance, organizations must seek out and implement cloud computing best practices at every step of the process – so they can succeed in spite of the cloud's ongoing issues and unknowns.

Service and Deployment Architectural Models

The definitions, attributes, and characteristics of cloud computing are expected to evolve and change over time as the technology matures. Currently, the National Institute of Standards and Technology (NIST) provides the following definition and models for cloud computing:

> "Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of three **service models** and four **deployment models**."

Service Models:

1. **Cloud Software as a Service (SaaS)**. The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage or

even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

2. **Cloud Platform as a Service (PaaS)**. The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems or storage, but has control over the deployed applications and possibly application hosting environment configurations.

3. **Cloud Infrastructure as a Service (IaaS)**. The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models:

1. **Private cloud**. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.

2. **Community cloud**. The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy and compliance considerations). It may be managed by the organizations or a third party, and may exist on-premise or off-premise.

3. **Public cloud.** The cloud infrastructure is made available to the general public or a large industry group, and is owned by an organization selling cloud services.

4. **Hybrid cloud**. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

Understanding and Evaluating Network Complexities

Before the introduction of cloud services, all of a user’s data and applications were stored and hosted by a hard drive on a computer or a set of servers in the IT department of an enterprise network. The figure below shows that with cloud computing, there isn’t any need for that hard drive or for those servers. Everything is in the cloud – stored, shared, or managed with “anywhere, anytime” availability. For example, an enterprise cloud user can access the company network from home or from the road, day or night – all that’s needed is an Internet connection.

Based on the models defined by the NIST, individual cloud service could be classified against a cloud service model (e.g., SaaS, PaaS or IaaS) and the deployment model being utilized (e.g., private cloud, community cloud, public cloud, or hybrid cloud). However, there are additional complexities to consider from the perspective of the underlying communications network. These include:

- The access network arrangement and technologies (e.g., fixed and mobile)
- The core network arrangements.

Figure 1 – Delivery of cloud services

There are different implications for performance, reliability and security depending upon how the cloud services are accessed – whether through the managed IP network services of a provider or the open Internet. Therefore, the planning, design and operation of cloud services must involve a complete system engineering view of the overall architecture that includes the relationships with the underlying network.

Setting Effective Criteria for Quality of Service and Security

The main challenge facing the cloud industry is a lack of effective criteria to describe and measure quality of service and security assurance. This absence of clear and precise standards makes it difficult to pinpoint the critical elements to be included in SLAs. More importantly, it's also challenging to determine effective means to verify compliance to the technical and security criteria that *is* stipulated. For mission-critical applications, these issues must be resolved by setting effective cloud criteria for a network's:

- Performance
- Reliability, Availability and Disaster Recovery
- Security

Performance Issues

Cloud planning, design and operations must take end-to-end performance into consideration. Given the variety of service architectural and deployment models, as well as variations of the underlying communications network arrangements, performance in terms of quality of service (QoS) is unpredictable. This volatility in the cloud can lead to network delays, latency, jitter and downtime.

There are no well-known or established criteria to describe the requirements for end-to-end quality of service that takes into consideration both performance at the cloud application level and performance of the underlying communications network. This issue is further complicated by the fact that the communications services provider may be different from the cloud service provider. Therefore, key performance issues must be identified by organizations themselves—and verified with the service providers they use.

Reliability, Availability and Disaster Recovery Issues

Certain critical applications, if moved to the cloud, would demand a guaranteed QoS, high levels of reliability, and continued availability from their computing infrastructure. While SLAs can be structured to meet the demands of various businesses, the reality is that some level of failures will occur when using commodity-based hardware solutions.

It is likely that events resulting in networking failures and outages will occur from time to time. Therefore, cloud deployment and implementation planning must include the operational measures to verify enforcement of diversity and redundancy rules to survive a variety of uncontrollable perils. Operational measures must be implemented, enforced, and audited for noncompliance.

The planning, design, and operation of cloud services should include identifying physical and logical interdependencies in the cloud provider's infrastructure together with the underlying network infrastructure. Understanding how resource democratization occurs within the cloud provider to best predict system availability and performance is critical to the migration and ongoing maintenance process.

Security Issues

There are obvious concerns about the security of transfer and retention of sensitive data, especially as a cloud-based infrastructure seems to imply a less transparent mechanism of storing and processing data. Many companies and government agencies are uncomfortable with their data being located on hardware outside of their direct control. This unease quickly turns to fear when you add the fact that cloud computing services are often multitenant, meaning that other companies, even competitors, are sharing the same hardware resources. In addition, a heightened regulatory environment has driven the need for organizations to be extremely cautious with their information due to serious legal and financial sanctions if personal data is compromised.

Cloud providers are faced with the need to provide proper security measures and are unsure what functionalities are appropriate or adequate. On the other hand, the enterprise customer of cloud services faces difficulty determining the requirements for security and the appropriate measures to check for compliance. Mitigating security risks and threats associated with the use of cloud infrastructure, platforms and applications is a major challenge. When choosing a cloud service provider, the following security issues must be considered:

- Chain of Trust
- Identity Management (IdM)
- Third Party and Supply Chain Management

Establishing a Chain of Trust

The assurance that the risk from using cloud services is at an acceptable level depends on the trust that the organization places in its cloud service provider. In some cases, the level of trust is based on the amount of direct control the organization is able to exert over the provider’s security techniques.

The security planning, design and operations for cloud must establish and maintain appropriate measures to enforce a chain of trust. This includes building confidence that each participating service provider in the cloud supply chain offers adequate data protection and security. The chain of trust can be complicated due to the number of entities participating in the consumer-provider relationship chain and the type of relationship between the parties. It’s also possible that cloud service providers may outsource services to other external entities, making the chain of trust even more complicated and difficult to manage.

Ensuring Effective Identity Management (IdM)

To deploy and maintain cloud services, organizations must establish clearly defined and enforceable policies for identity management among the multiple parties involved, including the cloud customer, the cloud provider and any third parties.

A clear and structured approach to risk assessment should be established and used among all parties to determine, select and implement the necessary security controls. For example, to enable strong authentication, cloud applications could support the capability to delegate authentication to the enterprise network that is consuming the services. In this case, the enterprise can enforce strong authentication using existing infrastructure and authenticate with open standards such as Security Assertion Markup Language (SAML) with the cloud provider. This would require cloud providers to externalize authentications and consider supporting various strong authentication options such as one-time passwords, biometrics, digital certificates and Kerberos.

Managing the Cloud Supply Chain

Supply chain security issues must be carefully considered when selecting cloud providers. To prevent security breaches, organizations must maintain tight control over key business, technical and security issues. Specifically, the cloud provider's supply chain – and the provider's management of those relationships – should be assessed to the extent possible.

Assessment of third party service providers should specifically target the provider's:

- Incident management policies
- Business continuity and disaster recovery procedures
- Co-location and back-up facilities
- Metrics illustrating the effectiveness of controls in these areas.

Considering Operations from Every Perspective

Operations need to be considered from a variety of different perspectives and in the context of several different relationships. Using the diagram below as a guideline, consider the following perspectives when designing, deploying and maintaining cloud services:

1. Business consumer perspective
    a. Relationship to enterprise IT department

2. Enterprise IT department perspective
    a. Relationship to business consumer
    b. Relationship to communications service provider
    c. Relationship to cloud service provider

3. Communications service provider perspective
    a. Relationship to enterprise IT department
    b. Relationship to cloud service provider

4. Cloud service provider perspective
    a. Relationship to communications service provider
    b. Relationship to enterprise IT department

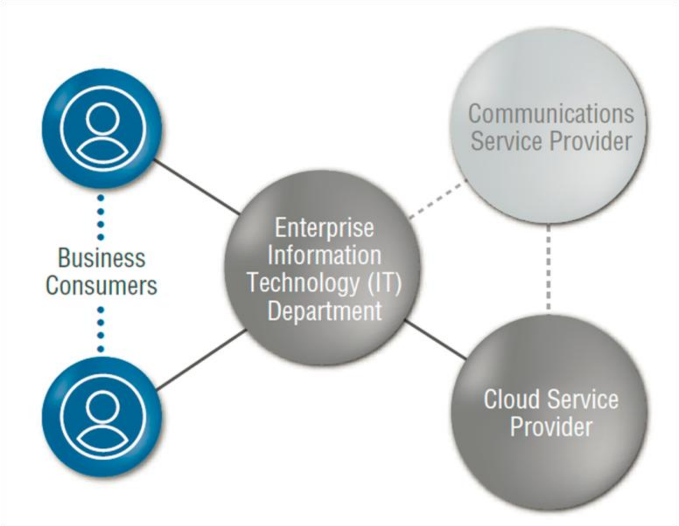
Communications Service Provider
Business Consumers
Enterprise Information Technology (IT) Department
Cloud Service Provider

Figure 2 – Roles and relationships involved in the provision of cloud services

Business consumer perspective: This perspective should not be impacted by the introduction of cloud services.

Enterprise IT department perspective: ITIL provides best-practice advice and guidance on all aspects of managing the day-to-day operation of an organization's information technology (IT) services, including relationships with business consumers. According to ITIL Service Design, an IT service is constructed from a combination of IT assets and externally provided 'underpinning' services. Cloud services, delivered via communications services, provide a new opportunity to exploit the concept of 'underpinning' services. While ITIL Service Design is focused on the design of new services, it may in some cases be appropriate to redesign existing services to exploit the advantages offered by the cloud. In this perspective, the cloud service provider and the communications service provider become *suppliers* to the enterprise IT department. The ITIL Supplier Management process ensures that suppliers and the services they provide are managed to support IT service targets and business expectations. Note, however, that while the ITIL Supplier Management process provides best practice guidance, the application of this process to a specific enterprise IT department adoption of cloud services will require additional customized and detailed process engineering.

Communications service provider perspective: For communications service providers, TM Forum's Business Process Framework, also known as eTOM, serves as the blueprint for standardizing and categorizing business activities (in the form of process elements) that will help set direction and the starting point for development and integration of business and operations support systems (BSS/OSS). Even prior to the cloud, the communications services and IT services worlds had been moving closer together, including the corresponding business and operations processes needed to support such services. This has been well recognized throughout the industry and is discussed in more detail in a TM Forum study on integrating ITIL and eTOM to provide a pragmatic joint solution for business support in the communications sector. The convergence of communications and IT is further accelerated by the introduction of cloud services, particularly for the case when an existing communications service provider wants to enter the cloud service provider market. Note, however, that while the TM Forum study provides guidance on how to “blend” ITIL and eTOM processes, the application of a converged process set within a specific communications service provider environment will require additional customized and detailed process engineering.

Cloud service provider perspective: For the cloud service provider (whether stand-alone or integrated with a communications service provider), the operational processes constitute those disciplines that directly provide IT service delivery (e.g., ITIL). Reflecting on the terminology used by the ITIL framework, this would include the core disciplines of service operations:

- Access management
- Event management
- Incident management
- Problem management
- Request fulfillment

The nature of the operations disciplines do not easily amend to a strict demarcation between the cloud service provider and the consumer’s responsibility across OS’s, virtual machines, servers, databases, middleware, and applications. There will need to be shared responsibilities when data must be accessed or a workflow operates across both entities. Hence, each party must provide data visibility to the other, mutually set priorities and coordinate tasks. At a high level, the operations responsibilities for a cloud service provider are as follows:

- **Monitoring service:** The service provider must monitor the environment (including event, capacity, security and utilization) to ensure SLAs are met. The provider’s monitoring data must be provided over standardized APIs.

- **Management of incident:** Each party must inform the other of incidents which may affect the other. Predefined agreement must be reached on incident prioritization and the level of effort required by the service provider during an incident. Automated and standardized interfaces are to be established to manage incidents.

- **Coordinating the operation of change, configuration, and release/deployment***:* Each party will notify the other when a change in configuration or other operational aspect may affect the service capabilities of the other party. Proactive management is required to ensure a stable environment.

- **Establishing governance:** The service provider must adhere to and permit enforcement of governing frameworks and policies, internal/external audits, minimum standards/certifications and security controls. Penalties and contract termination may be established for instances where requirements are not adhered to.

- **Provisioning of services:** The service provider must have effective automated (wherever possible) mechanisms to request, provision (resources, access keys, etc.), manage and meter usage of services.

Establishing Effective Service Level Agreements

Cloud computing uses the concept of service level agreements to control the use and receipt of computing resources from, and by, third parties. Any SLA management strategy considers two well-differentiated phases: the negotiation of the contract and the monitoring of its fulfillment in real time. Thus, SLA management encompasses:

- SLA contract definition including basic schema with the QoS parameters
- SLA negotiation
- SLA monitoring
- SLA enforcement—according to defined policies

The underlying benefit of cloud computing is shared resources, which is supported by the underlying nature of a shared infrastructure environment. Thus, SLAs span across the cloud and are offered by service providers as a service-based agreement. Measuring, monitoring, and reporting on cloud performance is based upon an end user experience or the end user's ability to consume resources. A major challenge for cloud computing, relative to SLAs, is the difficultly in determining root cause for service interruptions due to the complex nature of the environment.

From the diagram below, it can be seen that the importance of SLAs spans all relationships that exist between a "provider" of a service (any type of service) and the "consumer" of that service.

Historically, service providers have specified the details of what is contained (or not contained) in the SLA based on what service capabilities they are able to deliver according to their deployed technology infrastructure. However, the recent trend is for customers to become more and more involved in the SLA specification process based on their particular business needs. With respect to cloud computing and SLAs in particular, customers would like to see the following attributes addressed:

- Clarity of pricing unit options
- Clarity and completeness of definitions, targets, failure types and remedies, monitoring and reporting
- Service credit policy in the context of a service management tool (not to be confused with liability concerns)
- Ability to change and modify, within parameters, at the customer's discretion
- Year-over-year improvement if a longer term commitment is required
- Clarity of "day one" commitment
- Metrics addressing (at minimum): availability, provisioning, performance, incident and problem management, reporting quality and timeliness, change responsiveness, regulatory driven responsiveness, provision of underlying data and service sets, overall customer satisfaction, sustainability and efficiency

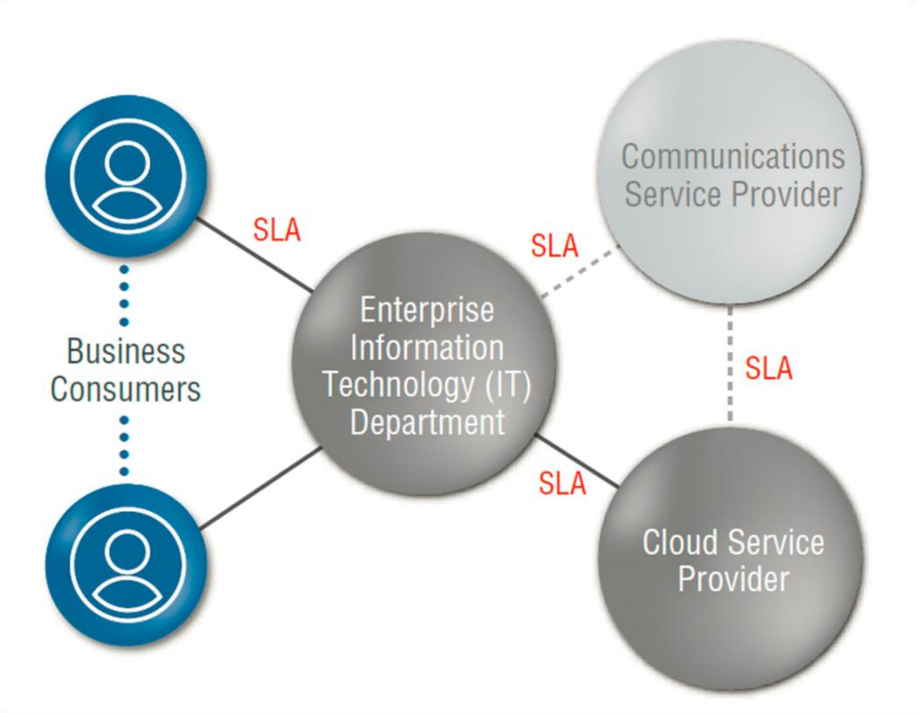
Business Consumers
SLA
Enterprise Information Technology (IT) Department
SLA
Communications Service Provider
SLA
SLA
Cloud Service Provider

Figure 3 – SLAs applicable to the provision of cloud services

Choosing the Right Service Provider

In today's competitive cloud environment, organizations should expect the highest quality and professional standards from their cloud service providers. Because the cloud is a relatively new phenomenon, not all providers have the necessary expertise to solve the complex network challenges that organizations face. To ensure a reliable and secure network in the cloud, potential providers should be evaluated on best practice criteria, which call for experience in:

- Safeguarding sensitive, mission-critical data in the cloud environment
- Managing complex cloud supply chains that involve multiple parties
- Planning effective service level agreements that ensure network performance, reliability, disaster recovery, and security
- Due diligence for compliance and verification of all technical requirements and security policies

Conclusion

The transition to cloud services for critical applications that demand quality of service and security present tough new challenges for organizations to solve. To succeed in this difficult environment, organizations need a structured approach to planning, designing and implementing their cloud structure.

By demanding transparency and accountability from cloud computing service providers, negotiating for more effective SLAs, and ensuring ongoing compliance to technical requirements and security policies, organizations can maximize their cloud infrastructure’s performance, reliability, and security—and meet their core business goals.

FOR MORE INFORMATION ABOUT EFFECTIVELY GROWING CLOUD INFRASTRUCTURE, CONTACT US AT:

Applied Communication Sciences
150 Mount Airy Road
Basking Ridge, NJ 07920
info@appcomsci.com
www.appcomsci.com

APPLIED COMMUNICATION SCIENCES®