

Software as a Service, and other Cloud Computing approaches

Introduction

In 2009, cloud computing became part of the mainstream IT consciousness. Although many of the concepts embodied by cloud computing have been around for a long time, it is only in the last 18 months that they crystallised into a coherent set of approaches.

This document examines the various approaches, and recommends appropriate uses for each one.

It pays particular attention to Software as a Service (SaaS), which is the most appropriate approach for companies where IT is not their core business.

Approaches

Cloud computing allows IT services to be purchased as a utility, the client pays a monthly fee according to how much IT service they consume. The services are hosted and managed remotely and the client does not need to worry about the complexity of managing those services. Cloud computing allows companies to manage IT services in the same way they manage water and electricity, by identifying a provider and paying for the services they use. The traditional approach to IT, of purchasing and managing all the hardware and software necessary to provide the required services to the company, is equivalent to running a generator non-stop to provide electricity, and having a borehole and water treatment facility to provide water.

There are a number of different approaches to cloud computing. They mainly differ in how much of the “technology stack” is provided by the cloud provider.

Infrastructure as a Service (IaaS)

Infrastructure as a Service gives the client on-demand access to a vast pool of processing or storage resources.

The processing resources normally take the form of virtual machines equivalent to a physical server of a specified processor type, speed and memory (RAM). The client can dynamically scale the computing resources available to them by starting new virtual machines as required. The virtual machines are paid for by the hour while they are running. The leading service of this type is Amazon's Elastic Compute Cloud (EC2).

The storage resources are usually provided as a scalable storage pool where the client pays per gigabyte stored, per month. There may also be data transfer charges for each gigabyte uploaded or downloaded from the cloud provider. The leading service of this type is Amazon's Simple Storage Service (S3).

Infrastructure as a Service gives companies access to highly scalable infrastructure without the cost and complexity of running a data centre. For example, a company offering a new service over the web might not be able to estimate how many people will use the service. If the company overestimates the demand, it will have paid for servers that will be under-utilised. On the other hand, if it underestimates the demand, the infrastructure it has set up may not be able to cope with the load, leading to increased response times, system downtime and unhappy customers. If the company had hosted the web service using an Infrastructure as a Service provider, then as the load increases the customer can start new servers as required. The customer only has to pay for those servers that are running, and no upfront

investment in hardware is required.

This approach is suited to computing tasks that must be able to scale very rapidly. However, it requires skilled IT professionals to design the system so that it can be scaled horizontally (by adding new servers) instead of vertically (by moving the application to the same number of more powerful servers).

If the required resources can be estimated with reasonable accuracy and are stable then it may be better to use either dedicated servers or virtual private servers, as this eliminates the complexity of using cloud resources.

Platform as a Service (PaaS)

Platform as a Service provides a platform with a specific API that can be used to deploy applications written by the customer. Google App Engine and Microsoft Windows Azure are two leading examples of Platform as a Service.

Platform as a Service implementations generally support a limited set of programming languages and a limited or proprietary set of application programming interfaces (APIs). This may limit the portability of applications deployed using Platform as a Service. For example, Google App Engine supports Python and Java Applications and Google's own BigTable non-relational database. This means most applications will not run on App Engine unless specifically designed for it, and it will be hard to move those applications elsewhere.

These types of concerns have made Platform as a Service the least developed approach to cloud computing.

The Platform as a Service market is changing as smaller providers start to offer widely used open source frameworks as a platform, such as Ruby on Rails and Django. Using one of these providers will provide the option of moving to another provider, or to self-hosting, provided that the data can also be moved.

Software as a Service (SaaS)

Software as a Service provides a customer access to a specific application, which is installed and maintained by the provider (normally on the provider's own servers or using an Infrastructure as a Service or Platform as a Service cloud). The customer accesses the application using the Internet, normally via a web browser. Salesforce.com is the archetypal Software as a Service, offering a Customer Relationship Management (CRM) application.

Software as a Service is typically charged per user per month. This approach provides true utility computing to the customer. There is no need for the customer to incur the capital expenditure required to procure licenses for the software and the hardware required to host it, or to incur the higher recurrent costs inherent in maintaining the hardware and software, taking backups, applying security patches, etc.

In general, Software as a Service will provide better disaster recovery and availability as well as lower costs, because of the economies of scale gained from running the same application for multiple customers. The customer is invoiced monthly or annually according to their actual usage.

Software as a Service offerings can be evaluated against those from other providers, and against traditionally licensed software, and selected based the features offered by the application and customer's requirements. It is important that the provider offers the customer the ability to take a full copy of their data, so that they can migrate to a different application in the future should they want to.

Software Technologies Ltd (STL) has traditionally sold licensed software to be installed on the

customer's own hardware and managed by their internal IT department. In 2010 and 2011, STL will offer applications using the SaaS approach, to reduce IT cost and complexity for customers.