EMC²
where information lives®

Cloud Computing:
A New Taxonomy Based on a Natural Model

EMC Proven Professional Knowledge Sharing 2009

Larry Huetteman,
Technology Business Consultant
EMC Corporation
Huetteman_Larry@emc.com

EMC² PROVEN PROFESSIONAL

Cloud Computing:
A New Taxonomy Based on a Natural Model

2009 EMC Proven Professional Knowledge Sharing

Larry Huetteman
Technology Business Consultant
EMC² Corporation
Huetteman_Larry@emc.com

Table of Contents

Introduction

Is “Cloud Computing” just the latest IT buzz word or a technology that will change the way we think about and deliver IT services in the future? Many of us in the industry are asking ourselves this critical question. The answer will likely fall somewhere in between. This article will not debate the pros and cons of this technology or speculate on its adoption. I assume that the technology is here to stay and since I believe strongly in the promise of the Cloud, am suggesting one way to structure the discussion.

Why call it cloud computing? The very notion of a cloud in an IT context denotes flexibility, abstraction layer, dynamic and ever-changing, volatile at times, and fluid. It can even evoke images ranging from easy, light and airy to powerful and deadly. In nature, clouds take many different forms. Cloud Computing can be almost anything you want or need it to be in the IT segment since it lacks structure or standards. There is a great opportunity to improve the classification of offerings in this burgeoning field.

History of the Cloud symbol in IT

I have seen a great many trends come and go in my 25 year career in IT. Clouds have been used to symbolize a layer of infrastructure for many years. According to Wikipedia, Cloud is a term with a long history in telephony, which has in the past decade been adopted as a metaphor for internet based services, with a common depiction in network diagrams as a cloud outline [1]. I was first exposed to the cloud symbol when IT departments needed more efficient ways to network open systems servers and desktops and used intelligent networks on a local level. This Local Area Network (LAN) was often referred to in architectural drawings as a cloud as illustrated on the next page.

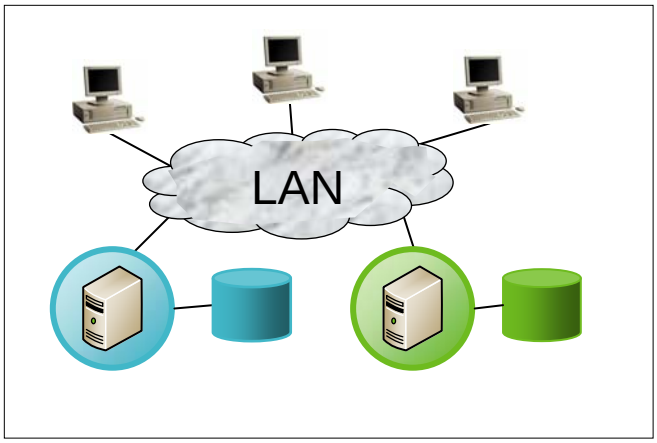

Figure 1 – Local Area Network

The hosts typically had captive storage and tape backup devices internal to the server chassis or directly attached. They communicated to the end user device through a Token Ring or Ethernet connection. We did not confuse the end users, business units, or even IT managers with the gory details under the network 'hood' since this connectivity layer was dynamic and often very complex in large organizations. Some unidentified architect decided that a fluffy *cloud* was the best way to illustrate the entire LAN infrastructure. Either that or he was frustrated with updating his drawings every time he added a node or upgraded a switch (remember AutoCAD?). In fact, the cloud symbol with "LAN" in the middle became a staple in drawing templates in tools such as Visio.

Since this symbol resonated with the IT community, its use and the various components it represented grew over time to include Wide Area Networks (WANs), Storage Area Networks (SANs), links to remote sites and more recently Virtualized Servers and Storage. It became the universal symbol representing the simplification or abstraction of anything in IT that was sufficiently complex, dynamic and difficult to maintain. It almost reached a point where drawings looked more like a weather map than an accurate depiction of the IT infrastructure as illustrated on the next page.

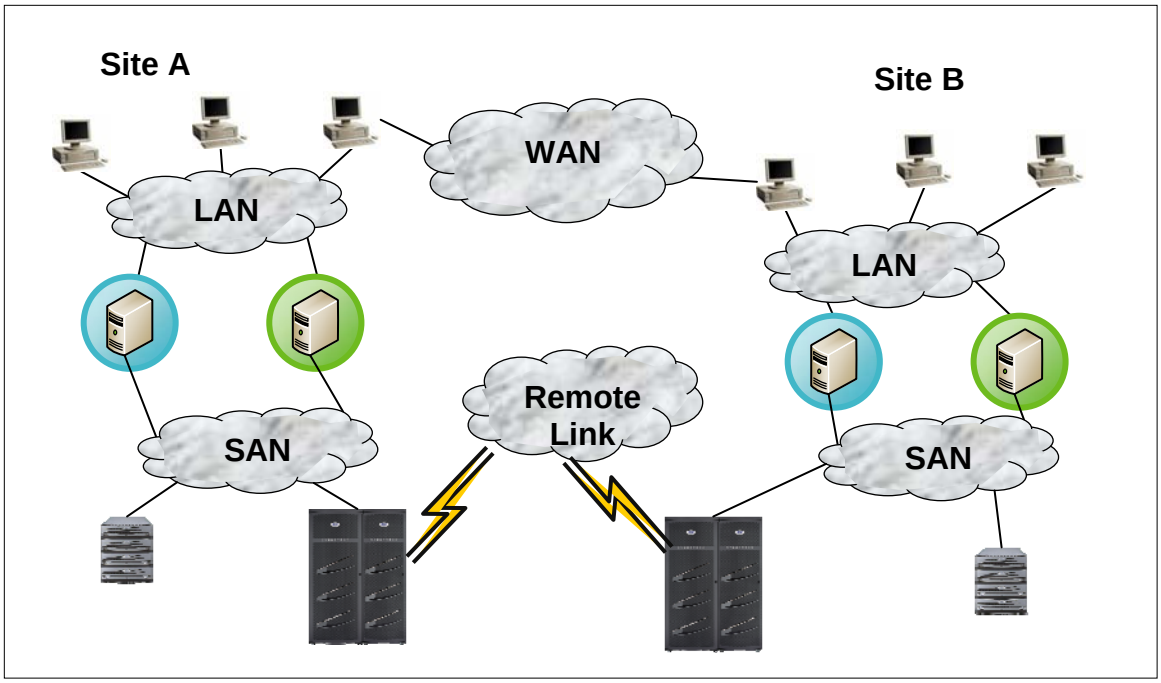

Figure 2 – Dual Site Infrastructure

In recent years, we've evolved this model to its logical conclusion. We surrounded or swallowed up portions or the entire topology into one giant cloud and called it Cloud Computing. This represents the 'hyper-abstraction' of the IT infrastructure and services that use to be owned, managed and maintained by the enterprise or service provider. You could also refer to this approach as the ultimate form of IT outsourcing, where virtually every IT function is delivered back to the business 'as-a-service'.

See illustration on next page.

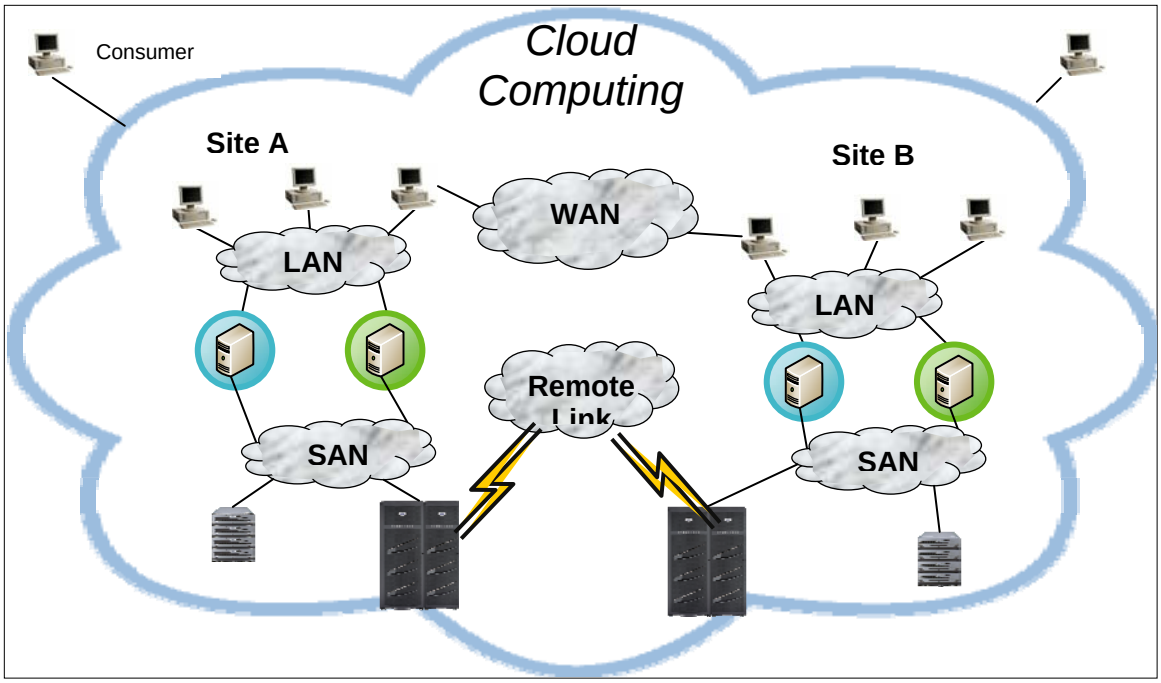

Figure 3 – Evolution of Cloud Computing

This new cloud can be deployed as an on-premise, off-premise, dedicated or shared configuration. The beauty of the model is that you can leverage a single approach or combine multiple approaches to create variations to best suit business needs. For example, a service provider can offer the security and guaranteed performance characteristics of their own on-premise version of a cloud service behind their firewall (private cloud). However, during peak usage periods or period ending batch runs (e.g., month-end, quarter-end), the private cloud can 'flex-out' into a public cloud temporarily to accommodate dynamic workload shifts. This approach offers business users a flexible and cost-effective way to manage seasonal or peak business demands.

No Laughing Matter

I've been the subject of jokes as I've promoted Cloud Computing services. During the last few months, I've heard, "Larry, clearly the stress of this activity has C*louded* your judgment," "It's like a dense fog - get your head out of the *Clouds*!" When it looked like we might close a small cloud storage deal, the Sales Account Manager commented that our portion of the weekly forecast appeared "*partly Cloudy."* We can laugh about it now but must realize that more serious opportunities and challenges lie ahead.

Why a New Taxonomy?

For those of us working in the emerging and expanding field of Cloud Computing, it has become increasingly difficult to have meaningful discussions across organizations without a common vernacular to frame the conversations. If you spend more than a few minutes researching the topic on the Web, you can clearly see that there is very little consistency in the naming conventions in this emerging field. The use of the term 'Cloud' to describe a technology is appealing from the perspective of reducing complexity but also too loosely defined to have any real meaning for IT professionals.

For example, I've seen articles that define cloud computing by using components with terms like Application, Services, Platform, etc. with the word 'Cloud' preceding each component. Others refer to everything 'as-a-service', as in Compute, Application, Storage, Business-Process-as-a-service, while others try to map SOA, Utility Computing and GRID to Cloud Computing Services. In the past, even EMC's suite of Cloud Computing offerings, while technically impressive, seemed to be a loosely coupled group of point solutions (Hulk/Maui/Infiniflex evolved into Atmos, while Mozy and Pi stood alone). This has improved significantly and will continue to do so with the convergence of the practices into a single business unit called Decho (short for "digital echo").

To further illustrate the need for structure, here are quotes from the cloud (Internet) on the status of the Taxonomy (or lack thereof) for the cloud computing industry:

From: David Chappell's Taxonomy of Cloud Platforms and Microsoft [2]
By John Osborn
August 22, 2008

"What's confusing to those of us exploring the substance of this trend for the first time is the sheer variety of open source and proprietary offerings that are "out there." Are any of them truly platforms, suitable for building and supporting applications that will be consumed by thousands or millions of users, and more importantly, how are we meant to knit these services together to build something useful?"

And more recently, from: The Need for a Standard Cloud Taxonomy [3]
By James Urquhart
January 27, 2009 11:15 PM PST

"The Cloud Interoperability meeting prior to Cloud Connect in Mountain View, Calif., last week was a very interesting petri dish of some of the best and brightest in the cloud-computing marketplace… While the whiteboard (filled during the session) may suggest that there was a large amount of agreement on the core concepts and that taxonomy was but a minor player, the reality is exactly the opposite. We couldn't agree on much of anything, except that there is a need for taxonomy and that trust (namely security) was one of the most pressing issues.

Funny enough, this is almost exactly the same conclusion reached in my recent discussions with some Cisco Systems partners, and (apparently) by Reuven Cohen, founder of Enomaly and the Cloud Computing Interoperability Forum (CCIF). Reuven conversed with Cannonical Services Director Simon Wardley--one of the Cloud Interop participants--and reached the joint conclusion that we need a stable, accepted taxonomy for cloud computing to "grease the skids," so to speak, for vendor interoperability discussions."

Vendors, Partners and Clients spend valuable time and energy deciding which model to follow in their discussions, or worst yet, they invent their own. I will attempt to simplify this process and apply some logic to the lack of consistent naming conventions. I propose a new taxonomy for Cloud Computing Service Offerings. Ideally, this convention should be easily understood and applied across a broad spectrum of services. It should also be intuitive so that it can be easily adopted, and based on a logical model. If successful, this model will be flexible and accommodate the anticipated explosive growth of this evolving field of technology.

A Natural Model

I am sensitive to the concern that any taxonomy that categorizes services bears the risk of over-simplification. This can lead to a blurred line of product differentiation and a view that all or most of the products and services placed in that 'bucket' are commoditized. I have made every attempt to provide a multi-layered naming convention that is simple to apply, yet also accommodates the subtle nuances of each product and service. In the end, we're trying to make it easier for consumers of the services to understand the offerings available today and in the future.

The concept of infrastructure and platform tiering or layers of services is one consistent theme in all Cloud Computing discussions. The thought is to use an existing scientific classification and to map the layers to current and potentially future Cloud Computing components or services. The idea occurred to me after one of my many business trips, as I was staring out the window of a plane at the clouds passing by at 40,000 feet. A logical choice is to leverage the familiar and widely accepted term, Cloud, but map the model to common, meteorological cloud terminology. In fact, the atmospheric cloud types are generally defined based on relative height or 'level' that they occur in the troposphere; with others that span multiple levels or layers. Coincidentally, this convention maps very well to the existing landscape of complex and layered Cloud Computing offerings.

After a closer look at the meteorological cloud definitions, we can categorize them into four basic types. There are some inconsistencies in the naming conventions, but generally they follow the guidelines on the next page:

Cloud Type or prefix	Level in the Atmosphere	Description and Examples
Cirrus	High Level	The high cloud group consists of Cirrus, Cirrostratus, and Cirrocumulus clouds. High clouds are made of ice crystals due to the cold air in the upper sky. The base of a high cloud above the surface can range from 6000-18000m in the tropics to 3000-8000m in the polar regions.
Alto	Middle Level	The middle cloud group consists of Altostratus and Altocumulus clouds. Middle clouds are made of ice crystals and water droplets. The base of a middle cloud above the surface can range from 2000-8000m in the tropics to 2000-4000m in the polar regions.
Stratus	Low Level	The low cloud group consists of Stratus, Strato-cumulus, and Nimbostratus clouds. Low clouds consist of water droplets. The base of a low cloud ranges from the ground surface to 2000m.
Cumulus	Vertically developed	Clouds with vertical growth include Cumulus and Cumulonimbus clouds. Vertical growth clouds span all levels of the troposphere and can even shoot into the stratosphere. The clouds develop by warm air rising from the surface. Cumulus and Cumulonimbus clouds provide the most interesting and severe weather.

Table 1 – Meteorological Cloud Definitions

Conveniently, these cloud types are often combined to form hybrids and variations on the original designations as seen in the Descriptions and Examples column. This bodes well for classifying Cloud Computing technologies that are likely to expand and contract many times before the technology matures and consolidates. In fact, there's no reason to limit the bounds of this taxonomy to existing meteorological definitions. For example, you may never hear a meteorologist mention an Altocirrus or Stratocirrus cloud type, but these combinations may make sense as you'll see later in this discussion. The more common cloud types are pictured here with an indication of their relative height in the troposphere:

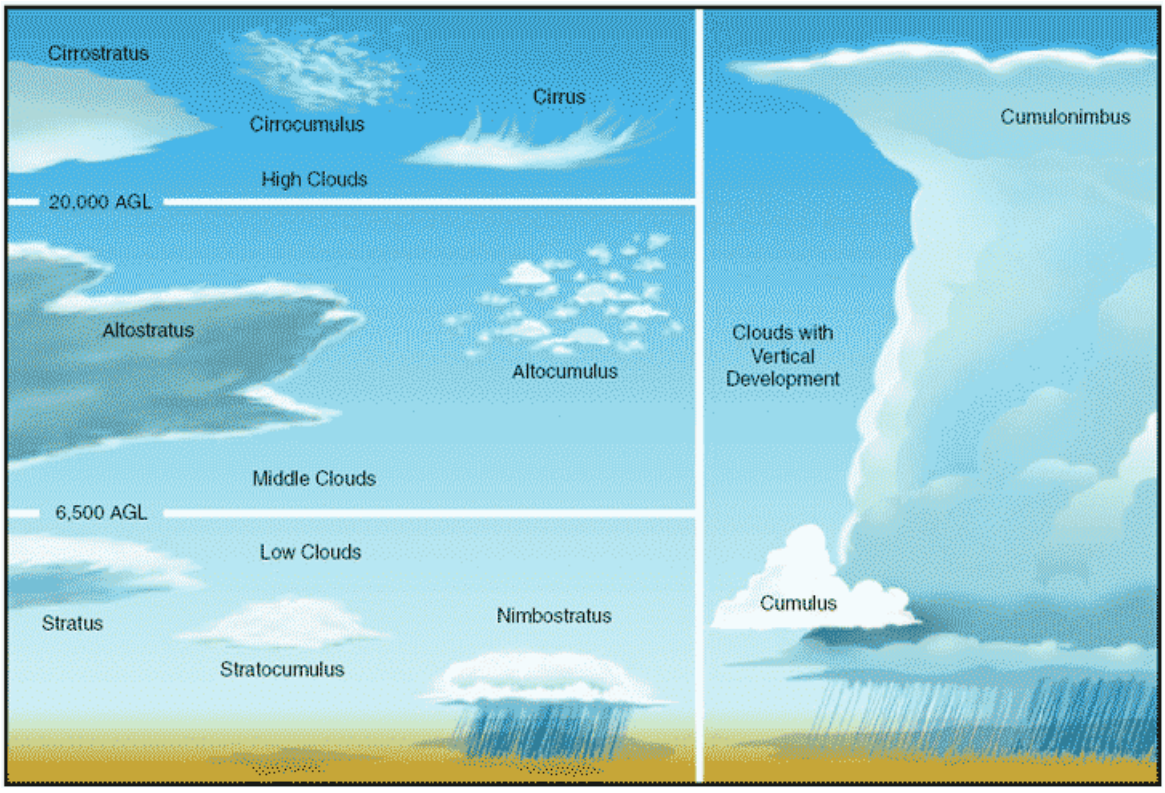

Figure 4 – Weather Clouds and Troposphere Levels

It's interesting to note that the cumulus family of clouds can have vertical development and span all levels of the troposphere. This potentially maps very well to more complex offerings that operate at multiple levels or have an end-to-end scope of service.

Cloud Computing Mapping Process

We can look at some Cloud Computing classification work that has already been done to evolve the discussion and start the mapping process. The following 6 layer stack is one of the more comprehensive models for Cloud Computing [4]:

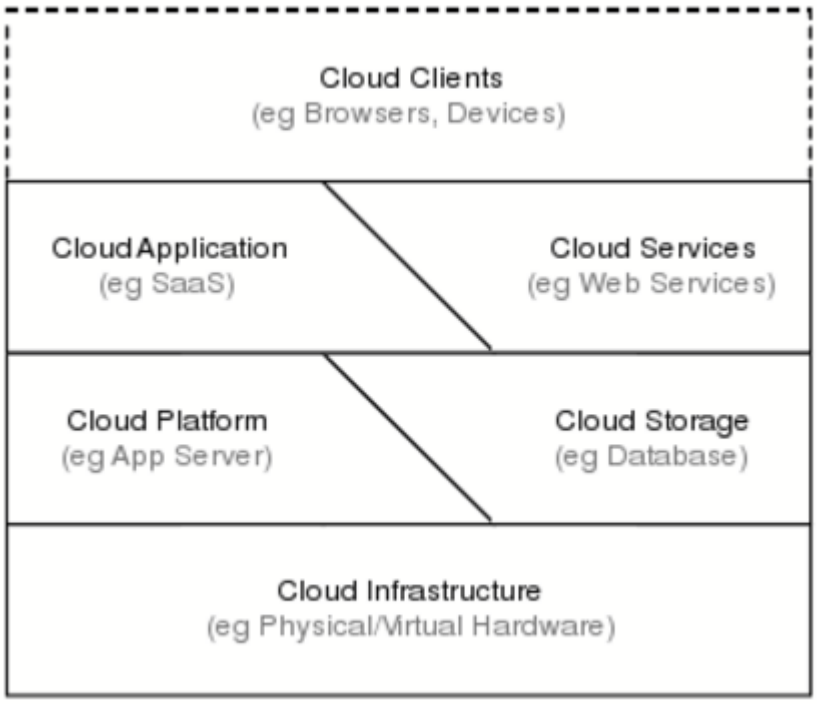

Table 2 – Levels of Cloud Computing

Wikipedia [1] provides a definition and examples for each layer, top to bottom as follows:

Client

A cloud client consists of computer hardware and/or computer software which relies on The Cloud for application delivery, or which is specifically designed for delivery of cloud services, and which in either case is essentially useless without it.

For example:

- Mobile (Android, iPhone, Windows Mobile)
- Thin client (CherryPal, Zonbu, gOS-based systems)
- Thick client / Web browser (Google Chrome, Mozilla Firefox)

Application

A cloud application leverages the Cloud in software architecture, often eliminating the need to install and run the application on the customer's own computer, thus alleviating the burden of software maintenance, ongoing operation, and support.

For example:

- Peer-to-peer / volunteer computing (Bittorrent, BOINC Projects, Skype)
- Web application (Facebook)
- Software as a service (Google Apps, SAP and Salesforce)
- Software plus services (Microsoft Online Services)

Services

A cloud service, such as Web Service, is "software system[s] designed to support interoperable machine-to-machine interaction over a network" which may be accessed by other cloud computing components, software, e.g., Software plus services, or end users directly.

For example:

- Identity (OAuth, OpenID)
- Integration (Amazon Simple Queue Service)
- Payments (Amazon Flexible Payments Service, Google Checkout, PayPal)
- Mapping (Google Maps, Yahoo! Maps)
- Search (Alexa, Google Custom Search, Yahoo! BOSS)
- Others (Amazon Mechanical Turk)

Platform

A cloud platform, such as Platform as a service, the delivery of a computing platform, and/or solution stack as a service, facilitates deployment of applications without the cost and complexity of buying and managing the underlying hardware and software layers.

For example:

- Web application frameworks
    - Ajax (Caspio)
    - Python Django (Google App Engine)
    - Ruby on Rails (Heroku)
- Web hosting (Mosso, Clustered Cloud)
- Proprietary (Azure, Force.com)

Storage

Cloud storage involves the delivery of data storage as a service, including database-like services, often billed on a utility computing basis, e.g., per gigabyte per month.

For example:

- Database (Amazon SimpleDB, Google App Engine's BigTable datastore)
- Network attached storage (EMC Atmos, MobileMe iDisk, Nirvanix CloudNAS)
- Synchronization (Live Mesh Live Desktop component, MobileMe push functions)
- Web service (EMC Atmos, Amazon Simple Storage Service, Nirvanix SDN)

Infrastructure

Cloud infrastructure, such as Infrastructure as a service, is the delivery of computer infrastructure, typically a platform virtualization environment, as a service.

For example:

- Full virtualization (GoGrid, Skytap)
- Grid computing (Sun Grid)
- Management (RightScale)
- Compute (Amazon Elastic Compute Cloud)

I propose the following strategy to attempt the mapping process. We will map each cloud computing service type to a layer of the meteorological model and an associated cloud type. For those services that share a layer, they will share the same suffix (e.g., Alto or Stratus). Starting from the highest level and working our way down the chart:

- Cloud Clients to Cirrus Cloud
- Cloud Application to Altostratus Cloud
- Cloud Services to Altocumulus Cloud
- Cloud Platform to Stratocumulus Cloud
- Cloud Storage to Stratus Cloud
- Cloud Infrastructure to Cumulus Cloud

In my mind, the mapping of Cloud Infrastructure to Cumulus makes sense because of the layered nature of each, thinking of infrastructure services as the bundling of other cloud computing services into one offering.

Cloud Computing Taxonomy

The following chart describes the model and completes the mapping process. This represents the proposed taxonomy for 6 of the more popular cloud computing categories with applications available in the marketplace today. Future offerings will need to be mapped based on their characteristics. We will also need to model a maintenance process to revise the model over time to accommodate as yet undefined offerings:

New Taxonomy Name	Cloud Computing Service	Service Description	Real World Examples
Cirrus	Clients	Computer hardware and/or computer software that relies on The Cloud for application delivery, or which is specifically designed for delivery of cloud services	• Mobile (Android, iPhone, Windows Mobile) • Thin client (CherryPal, Zonbu, gOS-based systems) • Thick client / Web browser (Google Chrome, Mozilla Firefox)
Altostratus	Application	Leverage the Cloud in software architecture, often eliminating the need to install and run the application on the customer's own computer, thus alleviating the burden of software maintenance, ongoing operation, and support	• Peer-to-peer / volunteer computing (Bittorrent, BOINC Projects, Skype) • Web application (Facebook) • Software as a service (Google Apps, SAP and Salesforce) • Software plus services (Microsoft Online Services)
Altocumulus	Services	Software system[s] designed to support interoperable machine-to-machine interaction over a network that may be accessed by other cloud computing components, software	• Identity (OAuth, OpenID) • Integration (Amazon Simple Queue Service) • Payments (Amazon Flexible Payments Service, Google Checkout, PayPal) • Mapping (Google Maps, Yahoo! Maps)

New Taxonomy Name	Cloud Computing Service	Service Description	Real World Examples
			• Search (Alexa, Google Custom Search, Yahoo! BOSS) • Others (Amazon Mechanical Turk)
Stratocumulus	Platform	The delivery of a computing platform, and/or solution stack as a service, facilitates deployment of applications without the cost and complexity of buying and managing the underlying hardware and software layers	• Web application frameworks o Ajax (Caspio) o Python Django (Google App Engine) o Ruby on Rails (Heroku) • Web hosting (Mosso, Clustered Cloud) • Proprietary (Azure, Force.com)
Stratus	Storage	The delivery of data storage as a service, including database-like services, often billed on a utility computing basis	• Database (Amazon SimpleDB, Google App Engine's BigTable datastore) • Network attached storage (EMC Atmos, MobileMe iDisk, Nirvanix CloudNAS) • Synchronization (Live Mesh Live Desktop component, MobileMe push functions) • Web service (EMC Atmos, Amazon Simple Storage Service, Nirvanix SDN)
Cumulus	Infrastructure	The delivery of computer infrastructure, typically a platform virtualization environment, as a service.	• Full virtualization (GoGrid, Skytap) • Grid computing (Sun Grid) • Management (RightScale) • Compute (Amazon Elastic Compute Cloud)

Table 3 – New Cloud Taxonomy Mapping

An Alternative Model

I feel compelled to offer at least one alternative model. It is based on the same principal but with a more simplistic mapping approach. The model above works well from an intuitive perspective and a low to high level mapping, but may not be simple and straightforward enough to encourage rapid adoption.

An alternative model is to abandon the intuitive 'level mapping' and propose a one-to-one mapping with unique cloud names that more closely match the name of the service. This should accelerate memorization of terms and offer an easier method of combining types into bundled offerings. For example Altostratus would be a cloud based application service with a cloud storage component. This alternative model could be represented as:

- Cloud Clients to Contrail (man-made) Cloud
- Cloud Application to Altos Cloud
- Cloud Services to Cirrus Cloud
- Cloud Storage to Stratus Cloud
- Cloud Platform to Nimbus Cloud
- Cloud Infrastructure to Cumulus Cloud

We can get creative and extend the model using different cloud types and weather references, including:

- Water Vapor can map to Cloud Computing startups (not fully formed)
- Cloud Burst can refer to the flexing out of workload to the public cloud
- Funnel Cloud can denote the consolidation of services into a single offering
- Use your imagination for Mushroom Cloud, Toxic Cloud, Storm Cloud etc.

As with any robust classification model, there are numerous cloud types and sub-types so that mapping of current and future Cloud Computing services to their counterparts will be easily accommodated by leveraging existing meteorological terminology. I believe that the Industry has a real need for clarity and consistency in this evolving field and structure can help to accelerate broader industry acceptance and adoption. How well these particular models are received will largely depend on the timing, quality of execution, and intuitive nature of the mappings.

Conclusion

It is as clear as the blue sky above that clouds are building on the horizon. In the ever-changing world of Information Technology, we are constantly challenged to stay ahead of the rapid pace with which new products and services are introduced. For those who support internal or external clients, we need to keep ourselves educated and ensure that our clients are kept abreast of key technologies that can add value to their business. Cloud Computing holds tremendous promise for Enterprises in the areas of cost reduction, streamlined operations, and flexibility to support cyclical workloads and capacities.

This is a critical time for IT practitioners and service providers. The economy continues to exert pressure to do more with less. I am convinced that Cloud Computing can and will deliver on many of its promises as long as it is allowed to age gradually and gracefully. As these solutions mature, many of the current concerns will be successfully addressed. It is my hope that this approach to a *New Taxonomy* is a step toward accelerating industry awareness, providing clarity to the myriad of offerings, and defining a classification process that can stand the test of time.

Appendix A – References

[1] Wikipedia definition of Cloud Computing:
http://en.wikipedia.org/wiki/Cloud_computing#cite_ref-The_Internet_Cloud_6-0

[2] David Chappell's Taxonomy of Cloud Platforms and Microsoft
http://news.oreilly.com/2008/08/david-chappells-taxonomy-of-cl.html

[3] The Need for a Standard Cloud Taxonomy http://news.cnet.com/8301-19413_3-10148806-240.html

[4] Taxonomy: The 6 layer Cloud Computing Stack
http://samj.net/2008/09/taxonomy-6-layer-cloud-computing-stack.html

Author's Biography

Larry Huetteman joined EMC in 1998 and has held various positions including Systems Engineer, Technology Consultant, Client Solutions Lead and most recently Technology Business Consultant in Global Alliances covering the major Outsource Partnership on a global basis. Prior to joining the GA group, Larry was covering large enterprise and Global accounts out of the Chicago field office. He has a consistent and proven track record of technical leadership in the Information Management field. Before joining EMC, he spent 4 years as a Consultant with the Center for Technology Enablement practice at Ernst & Young LLC. He started his career after earning a EE degree from the University of Illinois with 13 years at Exelon Corporation as a software developer, manager and MIS Director.

Laurence's recognition awards include Award of Appreciation, EMC Markets and Channels Reward and most recently, Platinum Player Recognition. His certifications include Engineer in Training (EIT), ITIL Foundations in IT Service Management, and EMC Proven Professional Associate Level certification.