

# The Structure of the New IT Frontier: Market Oriented Computing – Part II

Christian Vecchiola<sup>1</sup>, Dexter Duncan<sup>2</sup>, and Rajkumar Buyya<sup>1,2</sup>

<sup>1</sup> **CLOUDS Lab**, Department of Computer Science and Software Engineering,  
The University of Melbourne, Parkville, Victoria, Australia.  
{csve, raj}@csse.unimelb.edu.au

<sup>2</sup> **Manjrasoft Pty. Ltd.**, Melbourne, Victoria, Australia.  
{dexter, raj}@manjrasoft.com

*Cloud Computing's first outcome is the delivery of an infrastructure model that is approaching commodity like prices and fast implementation times. The main adopters of the public cloud are start-ups, small companies, enterprise test-beds, non-critical applications and delivery of back-up resources. There is a high amount of innovation expected above this commodity layer, namely in platforms (platform as a service) area leading to what we call Market Oriented Cloud Computing (MOCC). Partial adoption of MOCC is happening now inside private clouds with full adoption expected once security, control, trust and "standards" are in place. The cloud qualities enterprises are adopting include reservation aware scheduling, ability to apply Quality of Service (QoS) policies with alternative offers, Virtualization, support of multi-tenant and multi-applications of same resources, ability to apply pricing and charge-back policies and more. This is driving high demand for platforms that help enterprises move into the cloud.*

In the previous article we provided an overview of Cloud Computing. We covered the reference model and provided some real examples of IT companies operating in the field. In this article we dig into the details of the Infrastructure-as-a-Service (IaaS) solutions, which are the most popular outcome of this technology, and introduce a new phenomenon called Market Oriented Cloud Computing.

Enterprise CIOs focus to deliver return on investment and/or provide strategic advantage with applications and services that align to the business (and difficult for competitors to copy). Most CIOs have excess capacity and a huge investment in equipment. With an oversupply of servers and constant requests for more, they are looking to use their internal infrastructure more effectively. The recent global recession helped grow the hype around Cloud Computing as an alternative, and many are adopting the cloud ideas internally before buying whole-heartedly into the public cloud.

The commoditization of infrastructure through virtualization and super efficient management in hosted data centers lowers the cost of delivering the compute and storage needed to a price lower than the costs of purchasing and maintaining current IT infrastructures. A majority of enterprises have successfully installed some form of virtualization or plan to do so in a short time<sup>1</sup>. However, only about 2-5% have built an internal cloud, mostly just based on Virtual Machines (VM's). Only a few large organizations have the man-power and budget to get the same economies of scale that hosting companies provide. Security, control, trust and lock-in concerns are temporary with a new wave of innovation coming that makes it easier to move into the cloud. While some start-ups, small companies and non-strategic compute business areas are moving to the cloud now, the smart ones are developing their IT transformation strategy to focus on innovation to take full advantage of the fast time to market that pay-per-use, elastic services deliver. This new focus on building demand-based private network clouds to work seamlessly with Public Clouds is a natural



<sup>1</sup> From Forrester 2009 survey of 2,600 decision-makers from small and large companies in the US and Europe.  
<http://www.reuters.com/article/pressRelease/idUS126722+04-Mar-2009+BW20090304>

step we call *Market Oriented Computing*. Furthermore, as adoption of Cloud Computing becomes more prevalent a new area called *Market Oriented Cloud Computing*<sup>2</sup> begins to take shape.

*Companies that help organizations align their internal infrastructure to move easily into the cloud are the new heroes and therefore the new winners.*

## Brief Overview of Infrastructure as a Service (IaaS) – offering compute and storage as a commodity

Infrastructure-as-a-Service (IaaS) solutions deliver customizable IT infrastructures on demand. This is the most popular form of delivering IT services introduced by Cloud Computing, and it is also the most consolidated one. Table 1 shows some of the biggest players in the IaaS sector. The range of offerings varies from bare metal servers to entire virtual infrastructures customized for web applications. Some vendors own the underlying physical infrastructure serviced to users while others rely on these providers and provide additional value such as system design and application packaging. The most common pricing strategy is *pay-as-you-go* but many vendors also offer monthly plans.

Vendor	Service Type	Product	Pricing Model
Amazon	Infrastructure	EC2 [C], S3 [S], CloudFront [Cnt], ...	PAYG
Joyent	Infrastructure	Joyent Cloud [C]	Monthly
Rackspace/Mosso	Infrastructure	Cloud Servers [C] Cloud Files [S], Cloud Sites [SD,C,S]	PAYG, Monthly
ElasticHosts	Infrastructure	[C]	PAYG, Plans
FlexiScale	Infrastructure	[C]	PAYG
Rightscale	Infrastructure	[C]	PAYG
Cloud Central (Australia)	Infrastructure	[C,SD]	PAYG
Terremark	Infrastructure	vCloud Express [C], Enterprise Cloud [C]	PAYG, Plan
Melbourne IT (Australia)	Infrastructure	vCloud Express Beta[C]	PAYG
Rejila (Australia)	Infrastructure	[C,SD]	PAYG, Licensing
AT&T	Infrastructure	Synaptic Compute as a Service [C], Synaptic Storage as a Service [S], Synaptic Hosting as a Service [C,S]	PAYG, Monthly
IBM	Infrastructure/Platform	IBM Smart Business Services [C,S,...]	ND
VMWare	Infrastructure	vCloud Express [C]	PAYG
Google	Platform	AppEngine [C]	PAYG
Heroku	Infrastructure/Platform	[C,S]	PAYG
Microsoft	Platform, Infrastructure (M)	Azure [C,S], Dynamic Data Center Toolkit [C,S]	PAYG
Enomaly	Infrastructure (M)	[C]	ND
Eucalyptus	Infrastructure (M)	[C]	Free, Consulting
Elastra	Infrastructure (M)	[C, SD]	ND
Skytap	Infrastructure	[C,S,SD]	Monthly
Zimory	Infrastructure (M)	[C]	ND
Akamai	Infrastructure	[Cnt]	ND
Nirvanix	Infrastructure	[Cnt]	PAYG

**Table 1. Comparison table of the major IaaS vendors.**

Legend:

C	:	Compute
S	:	Storage
Cnt	:	Content (CDN)
H	:	Hosting
SD	:	System Design and Application Packaging
(M)	:	The vendor does not provide any infrastructure support but only the management software

2 R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility, Future Generation Computer Systems, Volume 25, Number 6, Pages: 599-616, ISSN: 0167-739X, Elsevier Science, Amsterdam, The Netherlands, June 2009.

ND : Not Disclosed  
PAYG : Pay As You Go

Amazon is definitely the most popular player in the field of Infrastructure on demand. The company launched Elastic Compute Cloud (EC2) three years ago and it is now the king of offering a pay-as-you-go elastic hosting service and a reference model for many other vendors providing compute on demand services. The entire IaaS offering of Amazon, which goes under the umbrella of Amazon Web Services, also includes Simple Storage Service (S3) providing virtual storage on demand that is combined with EC2 for an integrated solution and maximum savings. The list of Amazon customers ranges from very large enterprises to the end users that simply enter credit card details to instantly provision virtual servers. Another big player in the field is Joyent, which provides hosting, infrastructure, and application services for clients looking to run collaborative applications for their users. Joyent offers compute on demand and web application virtualization and helps to successfully scale social network websites like LinkedIn and Facebook, with nearly have 300 million users. On the same line, GoGrid, ElasticHosts, Rackspace, FlexiScale, RightScale, Melbourne IT and TerraMark provide bare metal virtual servers and storage space on demand.

Some vendors specialize by providing added value on top of the basic IaaS offerings. For example, RightScale eliminates vendor lock-in by letting the user choose the specific virtual infrastructure (Amazon, VMWare, etc) and software stack to compose their virtual environment (SkyTap). Other vendors, such as Cloud Central and Rejila, add specific features for composing your own virtual infrastructure and automating application packaging and deployment. Other solutions are completely specialized in providing a flexible and full featured virtual infrastructure design environment and rely on others to provide bare metal virtual servers or storage (Elastra, CohesiveFT).

With the exception of Google Inc, which offered a Cloud Computing model from inception, the big players of IT such as IBM, Microsoft, Sun, and AT&T made their move quite late and mostly concentrate their efforts in supporting service providers with infrastructure management software and vertically integrated service offerings.

Finally, if you want to build your own Cloud on your premises or elastically manage your datacenter, without leveraging integrated solutions such as those proposed by IBM and VMWare, virtual infrastructure management platforms like Enomaly Enterprise Cloud, Eucalyptus Enterprise Edition, or Zimory are your best choice.

This is just a selection of the plethora of vendors that offer compute and storage services on demand. Such services are the foundation for turning capital expense (CapEx) into an Operating Expense (OpEx) on corporate balance sheets and are called Infrastructure as a Service or IaaS. IaaS is fast approaching pricing models that look like commodity based pricing and the opportunity of exploiting Public Clouds to elastically shape your IT infrastructure is tempting. Rackspace offers compute time for less than 2¢ per hour with additional charges for bandwidth between 8¢ and 22¢ per GB/month. For comparison, Amazon has bandwidth rates between 10¢ and 17¢ per GB/month. Their EC2 compute service starts at 10¢ for a small Linux server with a slightly higher price for Windows. Both charge 15¢ per GB for storage with additional charges of 1¢ per request.

Most CIOs need to ensure that their private data center infrastructure evolves to integrate the new “commoditized” Public Cloud offerings. This allows them to maximize their existing investment and opens up opportunity for a future “unified design” where you easily move your data and share compute resources in a federated model either for purposes of managing peak load (cloud bursting), minimizing time and money (brokering) or as a disaster recovery solution. To make this happen, some key features of the Cloud, virtualization, provisioning on demand & Quality of Service (QoS) based workload management, are being integrated into privately owned infrastructure. This gives the dynamic elastic infrastructure and on-demand, self-serve, self-managed, internet enabled, and consumption based structure needed to take advantage of hybrid or inter-cloud offerings.

## **A Look at the Future: Market Oriented Cloud Computing**

Where is the focus now? What is Cloud 2.0 now that a commodity layer of infrastructure is underway? The window of innovation over the next few years is now clearly focused above this commodity layer. You are not going to make your enterprise hardware cheaper than the big boys operating in the IaaS space. This means that the real place of innovation and growth is in Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS) playground. Expect to see traditional infrastructure and virtualization companies playing aggressively in the Platform area through in-house development or acquisitions. In particular, there is a high amount of innovation beginning to occur above the

infrastructure layer, namely in the PaaS area, leading to what we call *Market Oriented Cloud Computing* (MOCC).

*Market Oriented Cloud Computing encompasses approaches and technologies enabling the trading of IT services, and their seamless integration into the existing infrastructure, according to service level agreements established to ensure the desired Quality of Service.*

<b>Time Frame</b>	<b>Private Infrastructure</b>	<b>Public Infrastructure</b>	<b>Market Situation</b>
2006 – 2010: Cloud Computing Thinking Formation	Legacy Private Networks	Some Cloud Computing Adoption	Confusion, trust, security, control concerns, jurisdiction issues slow down early adoption
2010 – 2012: Partial Market Oriented Cloud Computing Adoption	Cloud-enabled private infrastructures: provisioning on demand and virtualization.	Rapid Cloud Computing Adoption	Market Oriented Platforms (i.e. Manjrasoft) enter Market; Period of Innovation Begins
2013 onward: Full Market Oriented Cloud Computing Adoption	Market Oriented Cloud Computing Blurs line between Private/Public	Cloud Computing is Mainstream	Security, control and trust issues overcome.

**Table 2. Future evolution of the Cloud Computing Market.**

Table 2 shows the expected evolution of the Cloud Computing Market into the adoption of market-based strategies that shape the IT infrastructure of tomorrow. Partial adoption of MOCC is happening now inside Private Clouds with full adoption expected once security, control, trust and “standards” are in place. While we do not expect 100% of the market to move to the Cloud in three years, the mainstream adoption of Cloud Computing is in-line with Gartner Hype Cycle (Figure 1). In the near-term, enterprises are “cloudifying” their infrastructure, which is a partial adoption of the longer term trend to Market Oriented Cloud Computing. The “Cloud Qualities” that enterprises are adding include many areas that allow them to add applications and demand without buying additional hardware for an overall high resource utilization. These qualities include reservation aware scheduling, ability to apply dynamic QoS policies, virtualization, support of multi-tenancy and multiple applications on the same resources, ability to apply pricing and charge-back policies and more.

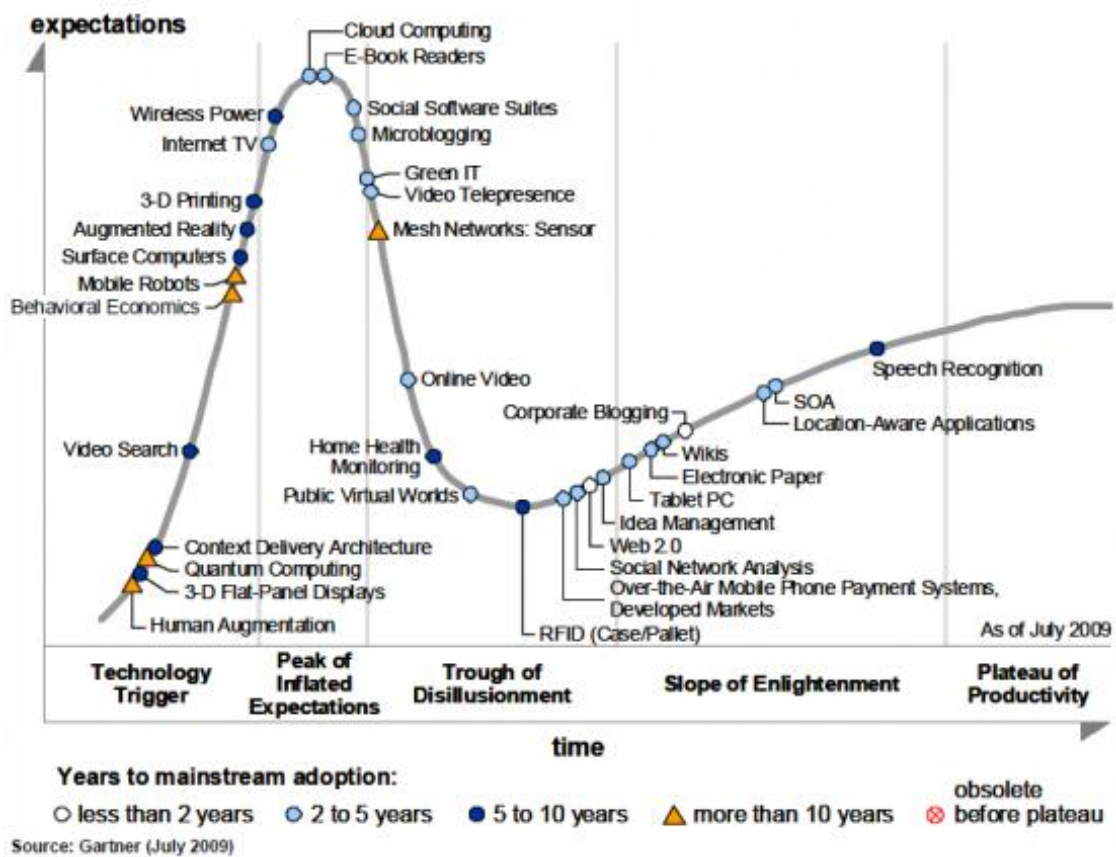


Figure 1. Gartner technology hype cycle.

## Market Oriented Cloud Computing: What More?

How does Market Oriented Cloud Computing differ from Cloud Computing? Is it really a step forward or just another buzzword used to add more confusion?

The Cloud Computing model already introduces on demand provisioning, consumption based subscriptions, and also promises QoS based delivery of IT services. However, at present, service providers have inflexible pricing, generally limited to flat rates or tariffs based on usage thresholds, and consumers are restricted to offerings from a single provider at a time. Also, many providers have proprietary interfaces to their services thus restricting the ability of consumers to swap one provider for another. For Cloud computing to reach a broad and seamless adoption, it is required that the services follow standard interfaces. This enables full commoditization and thus, would pave the way for the creation of a market infrastructure for trading in services.

What characterizes Market Oriented Cloud Computing is the realization of a virtual market place – Cloud Exchange – where IT services are traded and brokered dynamically. This only happens if enterprises seamlessly move from one provider to another, without conversion costs and in a completely dynamic manner. Organizations are building their internal Private Clouds to be similar to Public Clouds. This is the first step that enables the transparent integration of the publicly provisioned IT services with the private premises. What is missing is the availability of a market where desired services are published and then automatically bid by matching the requirements of customers and providers. Some are already moving in this direction: Amazon introduced the concept of “spot instances” that are dynamically offered by the provider according to their availability and bid by the customer. Their effective usage and consumption is then determined by the spot price established by Amazon and the maximum price provided by the customers.

*Market Oriented Computing has the same characteristics as Cloud Computing; therefore it is a dynamically provisioned unified computing resource allowing you to manage software and data storage as on aggregate capacity resulting in “real-time” infrastructure across public*

and private infrastructures. Market Oriented Cloud Computing goes one step further by allowing spread into multiple public and hybrid environments dynamically composed by trading service.

The realization of above Cloud Exchange is technically possible today, but not probable given the lack of standards and overall immaturity of the market. As vendors jockey to put forward their standard, and as concerns on security and trust begin to disappear, enterprises will move to take advantage of the new frontier of commodity like prices and quick time to market. Moreover, the presence of a demand based market place represents an opportunity for enterprises to shape their infrastructure for dynamically reacting to workload spikes and cut maintenance costs. It also allows possibility to temporarily lease some in-house capacity during low usage periods, thus better Return on Investment. The Cloud Exchange is a meeting point where these new opportunities are provided.

## Realizing the Vision: A Platform for Market Oriented Cloud Computing

Figure 2 illustrates the key components that shape the future scenario of Cloud Computing. In this vision, the “market place” plays a fundamental role. The dynamic provisioning of services from different vendors increases the competition in the value offering, resulting in a more consistent effort in providing a high quality and more accessible service.. To achieve this, Cloud providers no longer deploy traditional system-centric resource management architectures that do not provide incentives to share their resources and regard all service requests with equal importance. Instead, market-oriented resource management is necessary to regulate the supply and demand of Cloud resources to achieve market equilibrium (where supply = demand), providing feedback in terms of economic incentives for both Cloud consumers and providers, and promoting QoS-based resource allocation mechanisms that differentiate service requests based on their utility. In addition, clients benefit from the “potential” cost reduction of providers, which leads to a more competitive market and thus lower prices. This is the role of the market place.

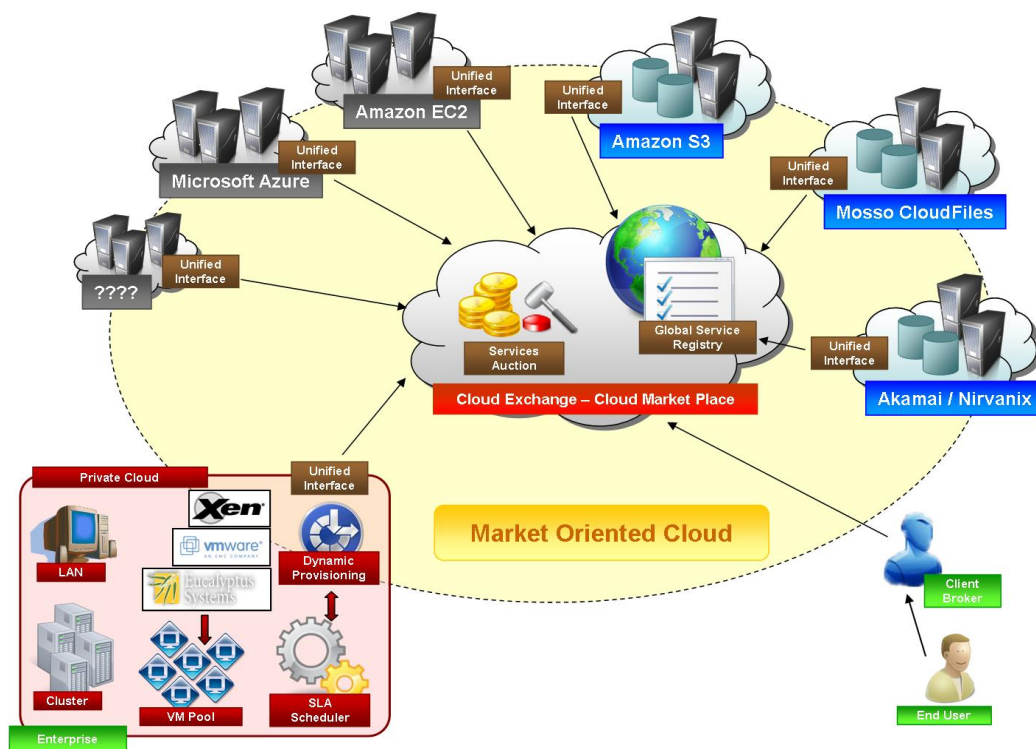


Figure 2: Global Cloud exchange and market infrastructure for trading services.

From the provider point of view a specific infrastructure is put in place to effectively take advantage of the demand based market place. Figure 3 shows a component based view of the infrastructure supporting MOCC.

- *Public Interface Access.* These are a set of standard interfaces for delivering IT services to end users and publishing them in the market place. These interfaces identify the access points of each Cloud Service provider.
- *SLA<sup>3</sup>-based resource allocator and workload manager.* The SLA Resource Allocator and Workload Manager are responsible for ensuring that IT services are delivered with the desired QoS. This layer is a fundamental component and includes different operations such as resource and application metering, billing, accounting and market based scheduling.
- *Virtualization Layer.* Virtual Machine technology allows workload isolation, server consolidation, multi-tenancy, fine grain tuning of allocated resources, and execution sandboxing. Moreover, virtual machine instances are started and stopped on-demand on a single physical machine to meet accepted service requests, hence providing maximum flexibility to configure various partitions of resources on the same physical machine to different specific requirements of service requests.
- *Physical Infrastructure.* It provides the horse power required to execute the Cloud workload. IaaS providers maintain multiple data-center that are maintained and made accessible through virtual infrastructure management software, such as those presented previously (vCloud express, Dynamic Data Center Toolkit, and others).

Such architecture ensures that crucial business operations of companies are performed with confidence by meeting deadlines. Critical QoS parameters that change over time, such as time, cost, reliability and trust/security, characterize the future Cloud Computing market and a flexible and open infrastructure is the way to go for realizing this vision.

---

<sup>3</sup> SLA: Service Level Agreement

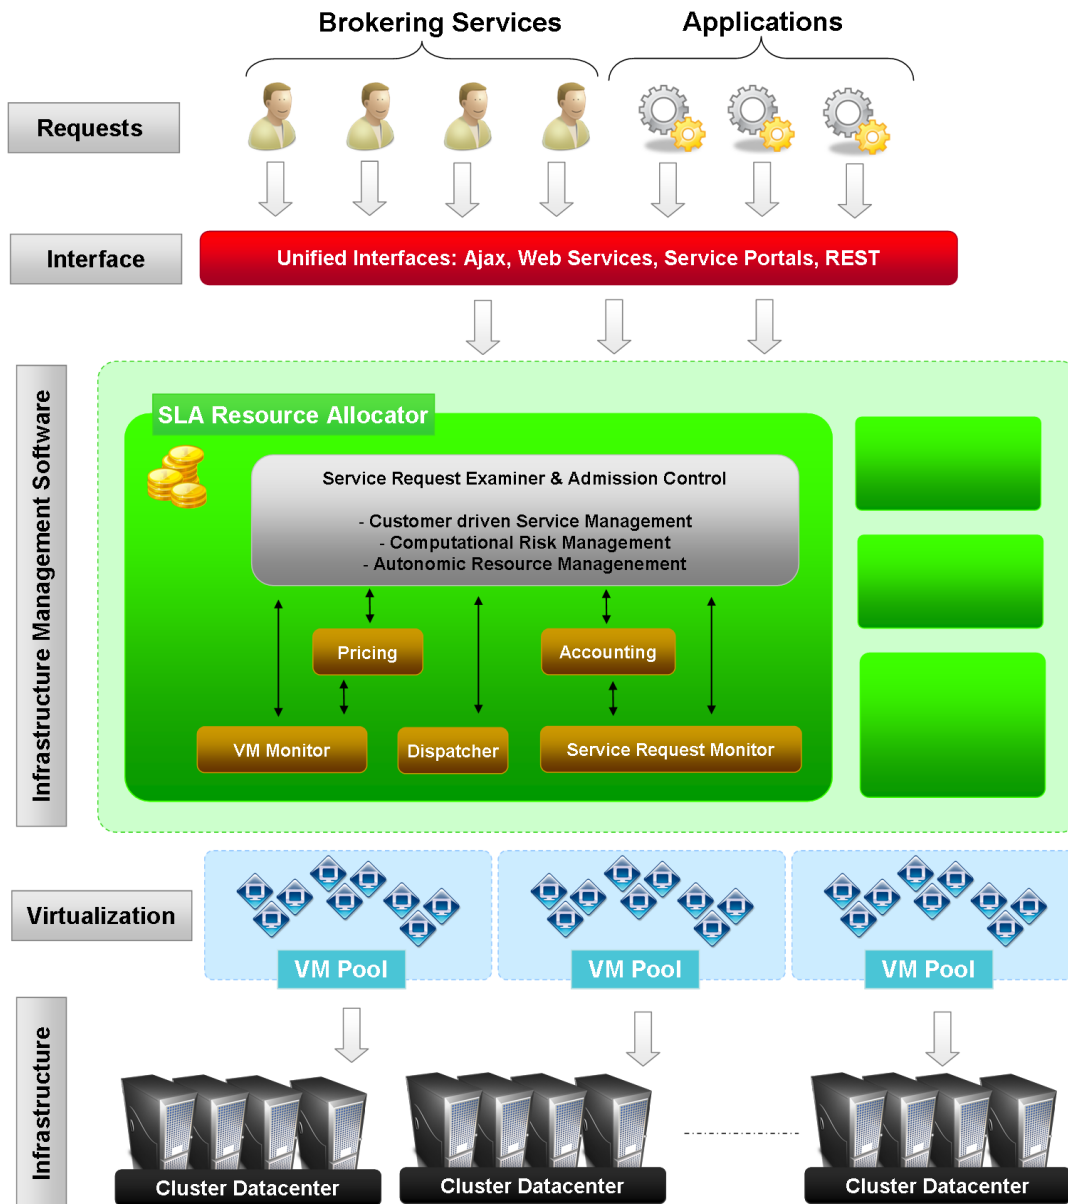


Figure 3. High-level market-oriented Cloud architecture for Cloud services providers.

The similar infrastructure and component stack is expected within the private premises of the enterprise. In order to seamlessly integrate Public Cloud resources, the internal IT organization follows the same resource management model found in public infrastructure, on a different scale. In addition, two distinctive features have to be considered:

- *Heterogeneity of resources.* Private Clouds tend to be composed by different resources, such as clusters, data-centers, workstations, and possibly Computing Grids. These resources are either virtual or physical. This requires uniform access and way of leveraging these resources.
- *Dynamic provisioning.* Public Clouds address the need of integrating additional resources into the existing infrastructure during peak loads. In a true market oriented offering these operations are not performed statically but are triggered by the constraints of the current workload. SLA aware schedulers identify the additional resources (storage, compute, or other) required to meet the QoS requirements and resources are obtained by means of a market mediated negotiation with available providers.



Figure 4 describes the general architecture of a Private Cloud augmented with market oriented capabilities.

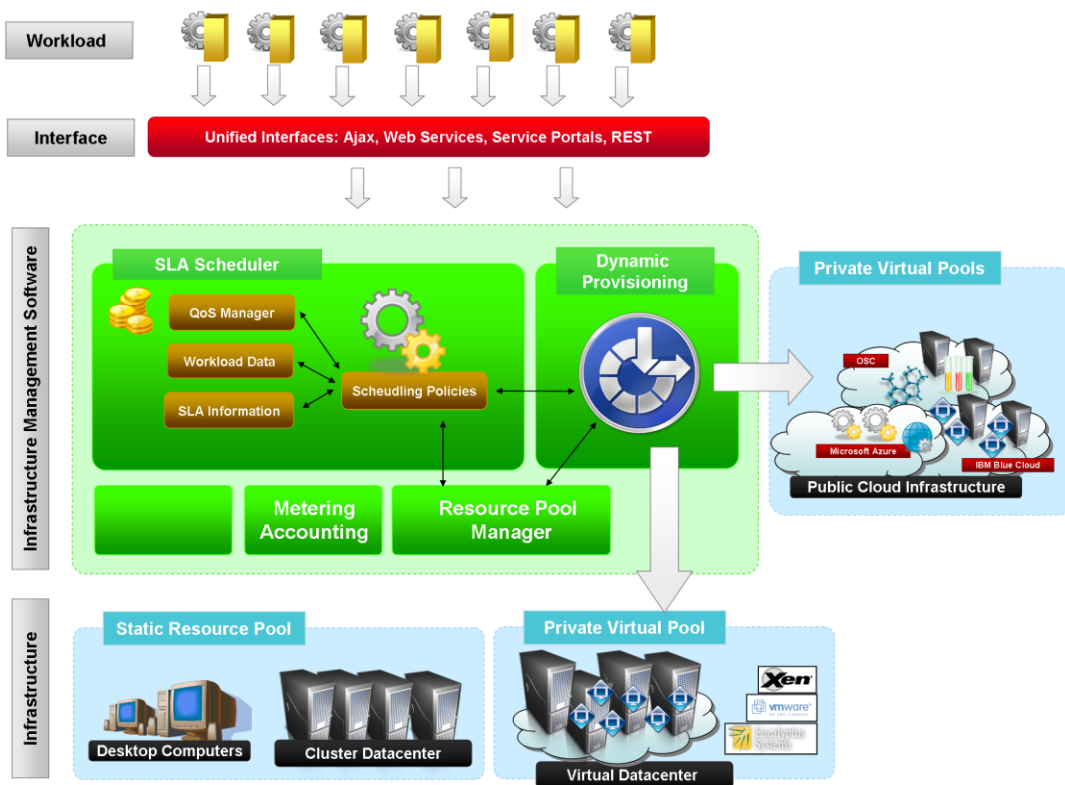


Figure 4. High-level market-oriented Cloud architecture for Cloud services providers.

From a customer point of view, a Cloud based commercial offering for enterprises is expected to:

- Support customer-driven service management based on customer profiles and requested service requirements
- Define computational risk management tactics to identify, assess, and manage risks involved in the execution of applications with regards to service requirements and customer needs
- Derive appropriate market-based resource management strategies that encompass both customer-driven service management and computational risk management to sustain SLA oriented resource allocation
- Incorporate autonomic resource management models that effectively self-manage changes in service requirements to satisfy both new service demands and existing service obligations
- Leverage VM technology to dynamically assign resource shares according to service requirements.

These requirements are most likely addressed by Platform-as-a-Service offerings rather than a pure IaaS. There is now a vacuum in the area of infrastructure and applications platforms that many vendors are rushing to fill. The Platform as a Service (PaaS) market is growing at 160% annual growth rate (according to IDC, Gartner, AMR and William Blair & Co) and expected to reach to just around USD \$9 to 10 Billion in 2012. This is by far the highest growth area in the Cloud.

To understand this better, we put together “Cloud Quality” Characteristics needed for enterprises to Cloud-enable their infrastructure. The Market Oriented Cloud Platform delivers on the qualities by being flexible with ability to easily work in multiple programming and multiple run-time environments. A new framework race is now happening that allows enterprises to manage based on the new demand based commodity infrastructure. These frameworks need to make it easy to manage, provision, monitor, secure and change resources when needed or based on pre-determined service levels that match the applications.

There is a shift from today’s supply-driven “build to peak” to a demand-driven, service-oriented approach. This results in real-time IT infrastructure with little spare capacity, but also gives more freedom to focus on business strategies rather than maintenance and system engineering. A Market Oriented Platform allows easy abstraction of underlying hardware and provides an elastic architecture that provisions and de-provisions on demand based on QoS requirements.

It is important for the platform to be flexible, easily customizable, and extensible while being service oriented to provide unique SLA/QoS & Web Services management, security and dynamic provisioning. This includes the ability to remotely control, monitor and dynamically change single and groups of nodes, which is not only important to simplify the management, but also to identify and remove bottlenecks.

Cloud Quality (CQ) Characteristic	Market Oriented Cloud Platform
Easy to Deploy	Write once and deploy in multiple run-time environments.
Pay as you Go	Pricing/Accounting modules to allow you to price based on resource usage
Demand Based	Dynamic Provisioning based on QoS
Multi-tenant	Resource Sharing: ability to host multiple users and applications across same infrastructure
Scalability	From the small local network to the data-center. Scalable services for handling increased workloads and potentially large infrastructures.
Workload Distribution	Provide custom workload Distribution Algorithms
Failure Management	Strong support for failover and content replication. VMs can be easily migrated from one node to other.
Flexibility	Customizable Services based on needs of applications. Pluggable and extensible architectures for implementing the system infrastructure.
Elasticity	Dynamic Provisioning to VM based on QoS from public and private resources.
Service Oriented Architecture	Service-based approach for infrastructure, workload management, and application submission.
Easy to Manage	Monitoring, Reservation, Meter and Charging features.
Self Managed	Auto plug scheduler as required – dynamic provisioning of scheduler
Quality of Service (QoS)	Extensible set of parameters controlling the scheduling of workloads. Ability of dynamically change the value of these parameters in order to meet the continuously changing business needs of the enterprise.
High Resource Usage	Combination of Multi-tenant and Provisioning based on QoS results in high utilization rate

**Table 3: Expected features for a Market Oriented Cloud Computing Platform.**

Many of the existing middleware platforms are proprietary and tied to their infrastructure or applications’ service such as Google’s App Engine and Salesforce’s Force.com. Most of the non-proprietary ones are Java/Linux based, with a few that rely on .NET framework. Outside those mentioned above, other companies providing Platform as a Service are: 3Tera, Apprenda, DataSynapse, Manjrasoft and Univa UD. For what concerns Market Oriented Cloud Computing, two initiatives are worth mentioning: Aneka and Gridbus Service Broker. Aneka is a middleware platform that provides an easy way to deploy Computing Clouds on top of heterogeneous infrastructure thus allowing the integration of Public and Private Clouds [1]. It supports QoS based resource management and schedules the execution of applications according to their available budget and QoS constraints. The same features are integrated in the Gridbus Service Broker that provides a general brokering service for data intensive applications, which can be scheduled on large variety of middleware: legacy job submission systems, Computing Grids, Public and Private Clouds [2].

## Summary

Infrastructure-as-a-Service solutions constitute the most popular and consolidated approach for dynamically delivering IT services introduced by Cloud Computing. IaaS offerings commoditize infrastructure and hardware and allow enterprise to elastically reshape their IT infrastructure. While the essence of Cloud Computing is provisioning IT services on demand, and IaaS value propositions are considerable step towards this direction, the flexible, efficient, and profitable use of these is still far away. The market segment of infrastructure on demand is now mature and a solid base on top of which build additional services that will lead to the widespread adoption of Cloud Computing.

To make this happen, it is necessary to adopt a market oriented approach where IT services are traded, provisioned, and advertised in competitive and open environment. This is the vision proposed by Market Oriented Cloud Computing. In this article we have characterized the fundamental properties of

such a model and listed a set of requirements that service providers and service consumers (i.e. Enterprises) need to meet in order to realize this vision. Market Oriented Cloud Computing is the natural evolution of Cloud Computing and identifies a scenario where the virtual public infrastructure is seamlessly and dynamically integrated with the private premises.

To realize this vision, Platform-as-a-Service solutions play a fundamental role. In the next article, we illustrate how to leverage these solutions to provide a heterogeneous, dynamic, and market oriented environment for Cloud Computing.

## References

1. Christian Vecchiola, Xingchen Chu, and Rajkumar Buyya, *Aneka: A Software Platform for .NET-based Cloud Computing*, High Speed and Large Scale Scientific Computing, 267-295pp, W. Gentsch, L. Grandinetti, G. Joubert (Eds.), ISBN: 978-1-60750-073-5, IOS Press, Amsterdam, Netherlands, 2009.
2. Rajkumar Buyya, Suraj Pandey, and Christian Vecchiola, *Cloudbus Toolkit for Market-Oriented Cloud Computing*, Proceeding of the 1st International Conference on Cloud Computing (CloudCom 2009, Springer, Germany), Beijing, China, December 1-4, 2009.