INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH & DEVELOPMENT

ISSN 2278 – 0211 (Online)

Study Of Cloud Bus Toolkit For Market-Oriented Business Cloud Computing

SK. Saidhbi
MCA, M.Tech(CSE), (Phd)
University Of Gondar, Ethiopia

Abstract:
This paper: defines the architecture for creating market-oriented Clouds and computing atmosphere by leveraging technologies such as virtual machines; provides thoughts on market-based resource management strategies that encompass both customer-driven service management and computational risk management to sustain SLA-oriented resource allocation; presents the work carried out as part of our new Cloud Computing initiative, called Cloud bus: (i) Aneka, a Platform as a Service software system containing SDK for construction of Cloud applications and deployment on private or public Clouds, in addition to supporting market-oriented resource management; (ii) internetworking of Clouds for dynamic creation of federated computing environments for scaling of elastic applications; (iii) creation of 3rd party Cloud brokering services for building content delivery networks and e-Science applications and their deployment on capabilities of IaaS providers such as Amazon along with Grid mashups; (iv) CloudSim supporting modelling and simulation of Clouds for performance studies; (v) Energy Efficient Resource Allocation Mechanisms and Techniques for creation and management of Green Clouds; and (vi) pathways for future research.

Key words: *Cloud Computing, Cloudbus, Utility Computing*

1.Introduction - Technology Trends

This vision of computing utilities, based on a service provisioning model, anticipated the massive transformation of the entire computing industry in the 21st century whereby computing services will be readily available on demand, like water, electricity, gas, and telephony services available in today's society. Similarly, computing service users (consumers) need to pay providers only when they access computing services, without the need to invest heavily or encounter difficulties in building and maintaining complex IT infrastructure by themselves. They access the services based on their requirements without regard to where the services are hosted.

This model has been referred to as utility computing, or recently as Cloud computing.

Cloud computing delivers infrastructure, platform, and software (application) as services, which are made available as subscription-based services in a pay-as-you-go model to consumers. These services in industry are referred to as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS), respectively. "Cloud computing, the long-held dream of computing as a utility, has the potential to transform a large part of the IT industry, making software even more attractive as a service".

Clouds aim to power the next generation data centers by architecting them as a network of virtual services so that users are able to access and deploy applications from anywhere in the world on demand at competitive costs depending on users Quality of Service (QoS) requirements. It offers significant benefit to IT companies by freeing them from the low level tasks of setting up basic hardware (servers) and software infrastructures and thus enabling them to focus on innovation and creating business value for their services.

Cloud computing has high potential to provide infrastructure, services and capabilities
required for harnessing this business potential. In fact, it has been identified as one of the emerging technologies in IT as noted in "Gartner's IT Hype Cycle" (see Figure 1). A "Hype Cycle" is a way to represent the emergence, adoption, maturity and impact on applications of specific technologies.

Cloud computing is definitely at the top of the technology trend, reaching its peak of expectations in just 3-5 years. This trend is enforced by providers such as Amazon1, Google, SalesForce2, IBM, Microsoft, and Sun Microsystems who have begun to establish new data centers for hosting Cloud computing applications such as social networking, gaming portals, business applications, media content delivery, and scientific workflows. It is predicted that within the next 2-5 years, Cloud computing will become a part of mainstream computing; that is, it enters into the plateau of productivity phase.

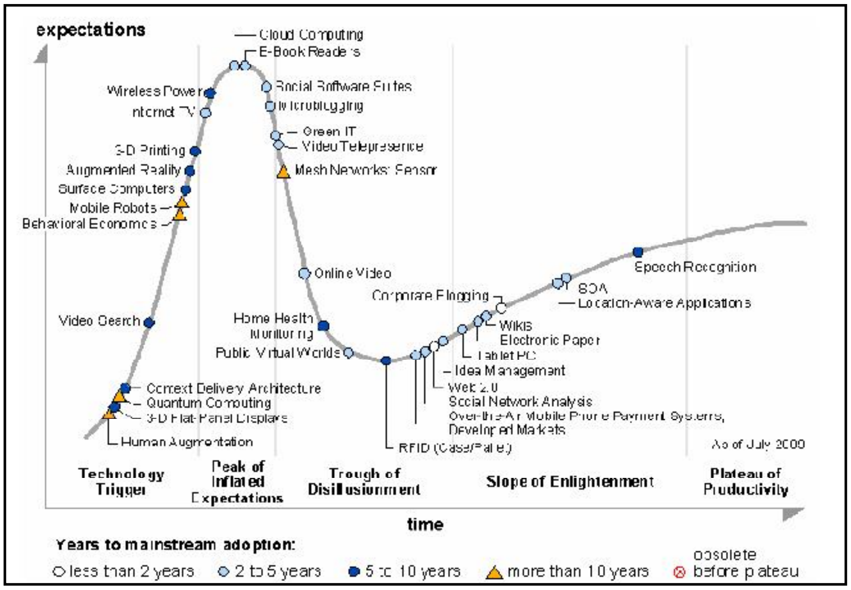

Figure 1: Hype Cycle of Emerging Technologies, 2009
Source: Gartner (August 2009)

The rest of the paper is organized as follows: Section 2 presents a high-level definition of Cloud computing followed by open challenges and a reference model; Section 3 presents Cloudbus vision and architecture in conformance with the high-level definition; Section 4 lists specific technologies of the Cloudbus toolkit that have made the vision a reality; Section 5 talks about integration of the Cloudbus toolkit with other Cloud management Technologies; and finally, Section 6 concludes the paper providing insights into future trends in Cloud computing.

2.Cloud Computing

Cloud computing is an emerging paradigm that aims at delivering hardware infrastructure
and software applications as services, which users can consume on a payper- use-basis. As depicted in Fig. 1, Cloud computing is now at the peak of its hype cycle and there are a lot of expectations from this technology. In order to fully understand its potential, we first provide a more precise definition of the term, then introduce a reference model for Cloud computing, and briefly sketch the challenges that lies ahead.

2.1.Cloud Definition

This definition puts Cloud computing into a market oriented perspective and stresses the economic nature of this phenomenon. The key feature, emerging from above two characterizations is the ability to deliver both infrastructure and software as services that are consumed on a pay-per-use-basis. Previous trends were limited to a specific class of users, or specific kinds of IT resources; the approach of Cloud computing is global and encompasses the entire computing stack. It provides services to the mass, ranging from the end-users hosting their
personal documents on the Internet to enterprises outsourcing their entire IT infrastructure to external data centers. Service Level Agreements (SLAs), which include QoS requirements, are set up between customers and Cloud providers. An SLA specifies the details of the service to be provided in terms of metrics agreed upon by all parties, and penalties for violating the expectations. SLAs act as a warranty for users, who can more comfortably move their business to the Cloud. As a result, enterprises can cut down maintenance and administrative costs by renting their IT infrastructure from Cloud vendors. Similarly, end-users leverage the Cloud not only for accessing their personal data from everywhere, but also for carrying out activities without buying expensive software and hardware.
Figure 2 shows the high level components of the service-oriented architectural framework consisting of client's brokering and coordinator services that support utility- driven management of Clouds: application scheduling, resource allocation and migration of workloads. The architecture cohesively couples the administratively and topologically distributed storage and compute capabilities of Clouds as parts of a single resource leasing abstraction . The system will ease the cross-domain integration of capabilities for on-demand, flexible, energy-efficient, and reliable access to the infrastructure based on emerging virtualization technologies .
The Cloud Exchange (CEx) acts as a market maker for bringing together service producers and consumers. It aggregates the infrastructure demands from the application brokers and evaluates them against the available supply currently published by the Cloud Coordinators. It aims to support trading of Cloud services based on competitive economic models such as commodity markets and auctions. CEx allows the participants to locate providers and consumers with fitting offers. Such markets enable services to be commoditized and thus, can pave the way for the creation of dynamic market infrastructure for trading based on SLAs. The availability of a banking system within the market ensures that financial transactions pertaining to SLAs between participants are carried out in a secure and dependable environment. Every client in the Cloud platform will need to instantiate a

Cloud brokering service that can dynamically establish service contracts with Cloud Coordinators via the trading functions exposed by the Cloud Exchange.

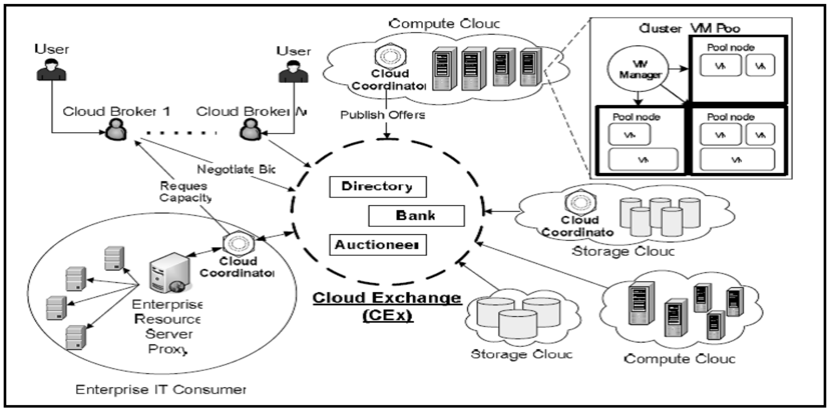

Figure 2: Utility-Oriented Clouds And Their Federated Network Mediated By Cloud Exchange

2.2.Open Challenges

Cloud computing introduces many challenges for system and application developers, engineers, system administrators, and service providers. Fig. 3 identifies some of them. Virtualization enables consolidation of servers for hosting one or more services on independent virtual machines in a multi-tenancy manner. When a large number of VMs are created they need to be effectively managed to ensure that services are able to deliver quality expectations of users. That means, VMs need to be migrated to suitable servers when QoS demand on services is high and later get consolidated dynamically to a fewer number of physical servers. One of the major concerns when moving to Clouds is related to security, privacy, and trust. Security in particular, affects the entire cloud computing stack. The Cloud computing model promotes massive use of third party services and infrastructures to host important data or to perform critical operations. In this scenario, the trust towards providers is fundamental to ensure the desired level of privacy for applications hosted in the Cloud. At present, traditional tools and models used to enforce a secure and reliable environment from a security point of view are the only ones available.

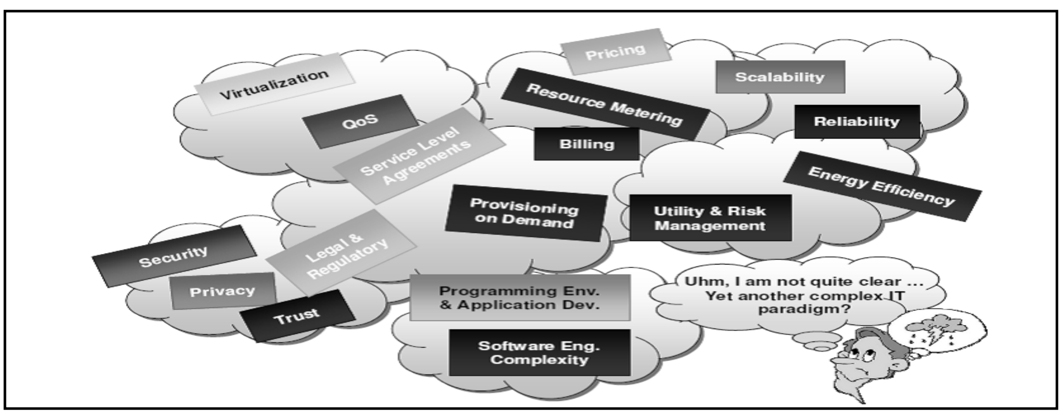

Figure 3: Cloud Computing Challenges

Besides security, there are legal and regulatory issues that need to be taken care of. When moving applications and data to the Cloud, the providers may choose to locate them anywhere on the planet. The physical location of data centers and clusters determines the set of laws that can be applied to the management of data. For example, specific cryptography techniques could not be used because they are not allowed in some countries. Simply, specific classes of users, such as banks, would not be comfortable to put their sensitive data into the Cloud, in order to protect their customers and their business. At present, a conservative approach is taken for what concerns hosting sensitive data. An interesting initiative is the concept of availability zones6 promoted by Amazon EC2. Availability zones identify a set of resources that have a specific geographic location. Currently there are two regions grouping the availability zones: US and Europe. Although this initiative is mostly concerned with providing of better services in terms of isolation from failures, network latency, and service downtime, it could be an interesting example for exploring legal and regulatory issues.

Data centers are expensive to operate as they consume huge amount of electricity. For instance, the combined energy consumption of all data centers worldwide is equivalent to the power consumption of Czech Republic. As a result, their carbon footprint on the environment is rapidly increasing. In order to address these issues, energy efficient resource allocation and algorithms need to be developed. In addition, practical and engineering problems are yet to be solved. Cloud computing infrastructures need to be scalable and reliable. In order to support this, a large number of application service consumers from around the world, Cloud infrastructure providers (i.e., IaaS providers) have been establishing data centers in multiple geographical locations to provide

redundancy and ensure reliability in case of site failures. Cloud environments need to provide seamless/automatic mechanisms for scaling their hosted services across multiple, geographically distributed data centers in order to meet QoS expectations of users from different locations. The scaling of applications across multiple-vendor infrastructures requires protocols and mechanisms needed for the creation of InterCloud environments. From applications' perspective, the development of platform and services that take full advantage of the Cloud Computing model, constitute an interesting software engineering problem.
These are some of the key challenges that need to be addressed for a successful adoption of the Cloud computing paradigm into the mainstream IT industry. R&D initiatives in both academia and industry are playing an important role in addressing these challenges. In particular, the outcome of such research in terms of models, software frameworks, and applications constitute the first tools that can be used to experience Cloud computing. The Cloudbus Toolkit is a step towards this goal.

2.3.Cloud Computing Reference Model

Fig. 4 provides a broad overview of the scenario envisioned by Cloud computing. This scenario identifies a reference model into which all the key components are organized and classified. As previously introduced, the novelty of this approach intercepts the entire computing stack: from the system level, where IT infrastructure is delivered on demand, to the user level, where applications transparently hosted in the Cloud are accessible from anywhere. This is the revolutionary aspect of Cloud computing that makes service providers, enterprises, and users completely rethink their experience with IT.
The lowest level of the stack is characterized by the physical resources, which constitute
the foundations of the Cloud. These resources can be of different nature: clusters, data centers, and desktop computers. On top of these, the IT infrastructure is deployed and managed. Commercial Cloud deployments are more likely to be constituted by data centers hosting hundreds or thousands of machines, while private Clouds can provide a more heterogeneous environment, in which even the idle CPU cycles of desktop computers are used to leverage the compute workload. This level provides the "horse power" of the Cloud.

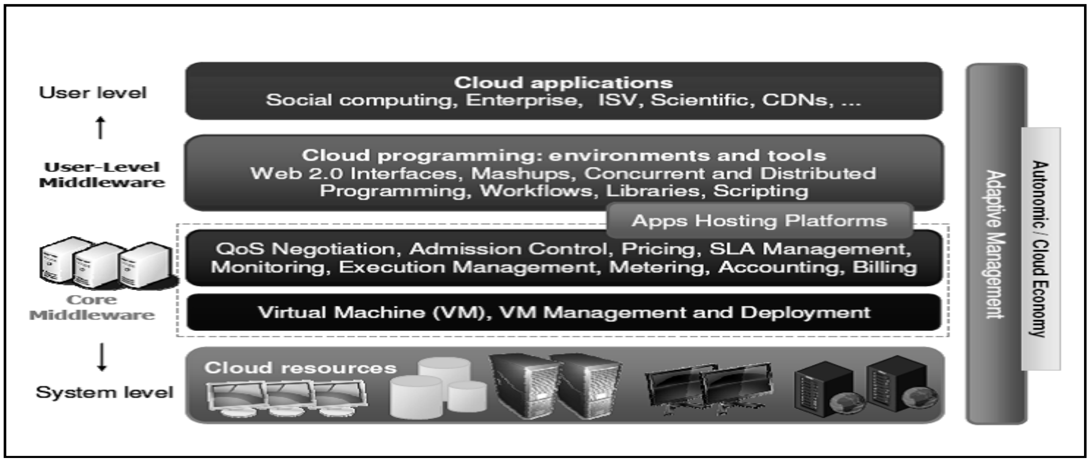

Figure 4: Cloud Computing Reference Model

The physical infrastructure is managed by the core middleware whose objectives are to provide an appropriate runtime environment for applications and to utilize the physical resources at best. Virtualization technologies provide features such as application isolation, quality of service, and sandboxing. Among the different solutions for virtualization, hardware level virtualization and programming language level virtualization are the most popular. Hardware level virtualization guarantees complete isolation of applications and a fine partitioning of the physical resources, such as memory and CPU, by means of virtual machines. Programming level virtualization provides sandboxing and managed executions for applications developed with a specific technology or programming language (i.e. Java, .NET, and Python). Virtualization technologies
help in creating an environment in which professional and commercial services are integrated. These include: negotiation of the quality of service, admission control, execution management and monitoring, accounting, and billing.
Physical infrastructure and core middleware represent the platform where applications are deployed. This platform is made available through a user level middleware, which provides environments and tools simplifying the development and the deployment of applications in the Cloud. They are: web 2.0 interfaces, command line tools, libraries, and programming languages. The user-level middleware constitutes the access point of applications to the Cloud.
It is quite uncommon for a single value offering to encompass all the services described in the reference model. More likely, different vendors specialize their business towards providing a specific subclass of services that address the needs of a market sector. It is possible to characterize the different solutions into three main classes:
Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure /
Hardware as a Service (IaaS/HaaS)

3.Cloudbus Vision And Architecture

Fig.5 provides a glimpse in the future of Cloud computing. A Cloud marketplace, composed of different types of Clouds such as computing, storage, and content delivery Clouds, will be available to end-users and enterprises.

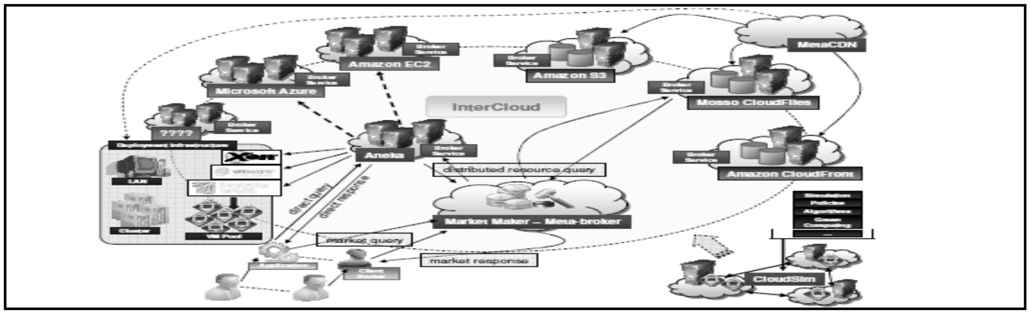

Figure 5: Cloud Computing Marketplace

Users can interact with the Cloud market either transparently, by using applications that leverage the Cloud, or explicitly, by making resource requests according to application needs. At present, it is the responsibility of the users to directly interact with the Cloud provider. In the context of a real Cloud marketplace, users will indirectly interact with Cloud providers but they will rely on a market maker or meta-broker component, which is in charge of providing the best service according to the budget and the constraints of users. A Cloud broker client, directly embedded within applications, or available as a separate tool, will interact with the market maker by specifying the desired Quality of Service parameters through a Service Level Agreement. As a result of the query, the meta-broker will select the best option available among all the Cloud providers belonging to the Cloud marketplace. Such interaction will take place through native interfaces exposed by the provider or via standardized brokering services.

In order to increase their chances of providing a better service to customers, different Cloud providers could establish peering arrangements among themselves in order to offload to (or serve from) other providers' service requests. Such peering arrangements will define a Cloud federation and foster the introduction of standard interface and policies for the interconnection of heterogeneous Clouds. The integration of different technologies and solutions into a single value offering will be the key to the success of the Cloud marketplace. PaaS solutions, such as Aneka [20], could rely on different providers for leveraging the workload and balance the use of private resources by provisioning virtual resources from public Clouds. This approach not only applies for compute intensive services, but also for storage and content delivery. MetaCDN [8], which is a Content Delivery Cloud, aims to provide a unified access to different storage Clouds in order to deliver a better service to end-users and maximize its utility.

The Cloudbus Toolkit is a collection of technologies and components that comprehensively try to address the challenges involved in making this vision a concrete reality. Fig. 6 provides a layered view of the entire toolkit and puts it into the context of a real Cloud marketplace. At the top of the stack, real life applications belonging to different scenarios leverage the Cloud horse power. Resources available in the Cloud are acquired by means of third party brokering services that mediate the access to the real infrastructure.

The Cloudbus toolkit mostly operates at this level by providing a service brokering infrastructure and a core middleware for deploying applications in the Cloud. For what concerns the brokering service, the Market maker is the component that allows users to take full advantage of the Cloud marketplace. The Market maker relies on different middleware implementations to fulfill the requests of users: these can be Cloudbus technologies or third parties implementations. Fig. 6 provides a breakdown of the components that constitute the Cloudbus middleware. Technologies such as Aneka or Workflow Engine provide services for executing applications in the Cloud.

These can be public Clouds, private intranets, or data centers that can all be uniformly managed within an InterCloud realm. In the following sections, we will present more details about the Cloudbus toolkit initiative and describe how they can integrate with each other and existing technologies in order to realize the vision of a global Cloud computing marketplace.

4.Cloudbus / Clouds Lab Technologies

The CLOUDS lab has been designing and developing Cloud middleware to support science, engineering, business, creative media, and consumer applications on Clouds. A summary of various Cloudbus technologies is listed in Table 2. We briefly describe each of these technologies in the following sub-sections.

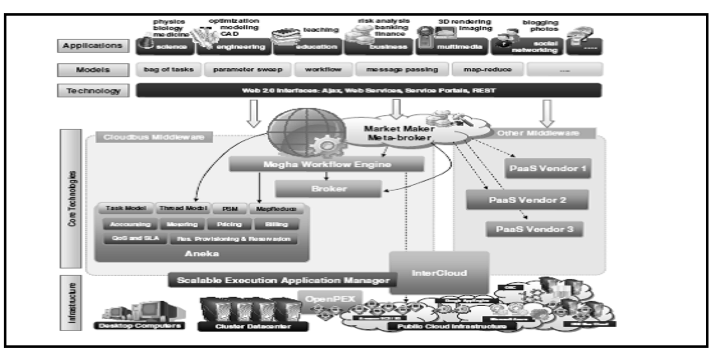

Figure 6: The Cloudbus Toolkit

The picture represents a layered view of the collection of technologies and components for market oriented Cloud computing available within the Cloudbus Toolkit.

4.1.Aneka

Aneka [20] is a "Platform as a Service" solution for Cloud computing and provides a software platform for developing and deploying applications in the Cloud. The core features of Aneka are: a) a configurable software container constituting the building blocks of the Cloud; b) an open ended set of programming models available to developers to express distributed applications; c) a collection of tools for rapidly prototyping and porting applications to the Cloud; d) a set of advanced services that put the horse power of Aneka in a market oriented perspective. One of the elements that make Aneka unique is its flexible design and high level of customization allowing it to target different application scenarios: education, engineering, scientific computing, and financial applications. The Aneka container, which is the core of the component of any Aneka based Cloud, can be deployed into any computing resource connected to the Internet whether it be physical or virtual. This makes the integration with public and private Clouds transparent; and specific services for dynamic provisioning of resources are built into the framework in order to exploit the horse power of the Cloud.

4.2.Broker

The Grid Service Broker [9] mediates access to distributed physical and virtual resources by (a) discovering suitable data sources for a given analysis scenario, (b) selecting suitable computational resources, (c) optimally mapping analysis jobs to compute resources, (d) deploying and monitoring job execution on selected resources, (e) accessing data from local or remote data source during job execution and (f) collating and presenting results. It provides a platform on which enhanced resource brokering strategies can be developed and deployed.

The broker supports various application models such as parameter sweep, workflow, parallel and bag of tasks. It has plug-in support for integration with other middleware technologies such a Globus [21], Aneka [20], Unicore [22] and ssh plug-in for accessing Condor [23], Unix based platforms via fork, PBS [24] and SGE [25]. The broker can provision compute and storage services in Cloud resources via SSH. It also provides QoS parameters in its service description for applications requiring a mix of public and private Cloud resources. For e.g. part of an application workload can be offloaded to Amazon EC2 and rest to local resources dynamically.

4.3.Market Maker/Meta-Broker

Market Maker/Meta-broker [13,14] is a part of Cloud infrastructure that works on behalf of both Cloud users and Cloud service providers. It mediates access to distributed resources by is covering suitable Cloud providers for a given user application and attempts to optimally map users' jobs and requirements to published services. It is a part of a global marketplace where service providers and consumers join to find suitable match for each other. It provides various services to its customers such as resource discovery, meta-scheduler, reservation service, queuing service, accounting and pricing services.

4.4.From Intergrid To Intercloud

In the coming years, users will be able to see a plethora of Cloud several providers around the world desperate to provide resources such as computers, data, and instruments to scale science, engineering, and business applications. In the long run, these Clouds may require sharing its load with other Cloud service providers as users may select various Cloud services to work on their applications, collectively. Therefore, dispersed Cloud initiatives may lead to the creation of disparate Clouds with little or no interaction between them. The InterCloud model will: (a) promote interlinking of islands of Clouds through peering arrangements to enable inter-Cloud resource sharing; (b) provide a scalable structure for Clouds that allow them to interconnect with one another and grow in a sustainable way; (c) create a global Cyberinfrastructure to support e-Science and e-Business applications.

4.5.MetaCDN

MetaCDN [8] is a system that exploits "Storage Cloud" resources offered by multiple IaaS vendors, thus creating an integrated overlay network that provides a low cost, high performance CDN for content creators. It removes the complexity of dealing with multiple storage providers, by intelligently matching and placing users' content onto one or many storage providers based on their quality of service, coverage and budget preferences. By using a single unified namespace, it helps users to harness the performance and coverage of numerous "Storage Clouds".

4.6.Cloudsim

The CloudSim toolkit [17] enables users to model and simulate extensible Clouds as well as execute applications on top of Clouds. As a completely customizable tool, it allows extension and definition of policies in all the components of the software stack. This makes it suitable as a research tool as it can relieve users from handling the complexities arising from provisioning, deploying, configuring real resources in physical environments.

CloudSim offers the following novel features: (i) support for modeling and simulation of large scale Cloud computing infrastructure, including data centers on a single physical computing node; and (ii) a self-contained platform for modeling data centers, service brokers, scheduling, and allocations policies. For enabling the simulation of data centers, CloudSim provides: (i) virtualization engine, which aids in creation and management of multiple, independent, and co-hosted virtualized services on a data center node; and (ii) flexibility to switch between space-shared and time-shared allocation of processing cores to virtualized services. These features of CloudSim would speed up the development of new resource allocation policies and scheduling algorithms for Cloud computing.

CloudSim evolved from GridSim [18], a Grid simulation toolkit for resource modeling and application scheduling for parallel and distributed computing. GridSim provides a comprehensive facility for creating different classes of heterogeneous resources that can be aggregated using resource brokers for solving compute and data intensive applications. It provides a framework for incorporating failures, advance reservations, allocation policies, data models, network model extensions, background traffic and load, and so forth, which are also present in the CloudSim toolkit.

5.Related Technologies, Integration, And Deployment

The Cloudbus toolkit provides a set of technologies completely integrated with each other. More importantly, they also support the integration with third party technologies and solutions. Integration is a fundamental element in the Cloud computing model, where enterprises and end-users offload their computation to third party infrastructures and access their data anytime from anywhere in a ubiquitous manner.

Many vendors provide different solutions for deploying public, private, and hybrid Clouds. At the lowest level of the Cloud computing reference model, virtual server containers provide a management layer for the commodity hardware infrastructure: VMWare10, Xen [7], and KVM11 (Kernel-based Virtual Machine) are some of the most popular hypervisors available today. On top of these, “Infrastructure as a Service” solutions such as Amazon EC2, Eucalyptus [28], and OpenNebula [29] provide a high level service to end-users. Advanced resource managers such as OpenPEX [27] complete the picture by providing an advance reservation based approach for provisioning virtual resources on such platforms. Technologies such as Aneka and the Workflow Engine can readily be integrated with these solutions in order to utilize their capabilities and scale on demand. This applies not only for compute type workloads, but also for storage Clouds and CDNs, as demonstrated by the MetaCDN project. At a higher level, the Market maker and the Grid Service Broker are able to pro-vision compute resources with or without a SLA by relying on different middleware implementations and provide the best suitable service to end-users.

The Cloudbus toolkit is a work in progress, but several Cloudbus technologies have been already put into action in real scenarios. A private Aneka Cloud has been deployed at GoFront12 in order to increase the overall productivity of product design and the return of investment of existing resources. The Workflow Engine has been used to execute complex scientific applications such as functional Magnetic Resonance Imaging (fMRI) workflows on top of hybrid Clouds composed of EC2 virtual resources and several clusters in the world [10, 11]. Various external organizations, such as HP Labs are using CloudSim for industrial Cloud computing research.

6.Future Trends

In the next two decades, service-oriented distributed computing will emerge as a dominant factor in shaping the industry, changing the way business is conducted and how services are delivered and managed. This paradigm is expected to have a major impact on service economy, which contributes significantly towards GDP of many countries, including Australia. The service sector includes health services (e-health), financial services and government services. With the increased demand for delivering services to a larger number of users, providers are looking for novel ways of hosting their application services in Clouds at lower cost while meeting the users’ quality of service expectations. With increased dependencies on ICT technologies in their realization, major advances are required in Cloud Computing to support elastic applications
offering services to millions of users, simultaneously.

Software licensing will be a major hurdle for vendors of Cloud services when proprietary
software technologies (e.g. Microsoft Windows OS) have to be made available to millions of users via public virtual appliances (e.g. customized images of OS and applications). Overwhelming use of such customized software would lead to seamless integration of enterprise Clouds with public Clouds for service scalability and greater outreach to customers. More and more enterprises would be interested in moving to Clouds for cooperative sharing. In such scenarios, security and privacy of corporate data could be of paramount concern to these huge conglomerates. One of the solutions would be to establish a globally accredited Cloud service regulatory body that would act under a common statute for certifying Cloud service providers,
standardizing data formats, enforcing service level agreements, handling trust certificates
and so forth.

As the technology is gradually changing from Cluster and Grid computing to Cloud computing, the Cloudbus toolkit is also evolving towards being more robust and scalable to support the hype. We are continuously consolidating our efforts to enhance the toolkit such that it is able to support more and more users.

7.References

1. Kleinrock, L.: A Vision for the Internet. ST Journal of Research, vol. 2, Issue 1, pp. 4--5, November 2005.
2. Armbrust, M., Fox, A. et. al.: Above the Clouds: A Berkeley View of Cloud Computing. Technical Report No. UCB/EECS-2009-28, University of California at Berkley, USA,
Feb. 10, 2009.
3. London, S.: Inside Track: The high-tech rebels. Financial Times, Sept. 6, 2002.
4. VMware: Migrate Virtual Machines with Zero Downtime. http://www.vmware.com/.
5. Broberg, J., Buyya, R., and Tari, Z.: MetaCDN: Harnessing 'Storage Clouds' for High
Performance Content Delivery, Journal of Network and Computer Applications, 32(5):1012-1022, Elsevier, Amsterdam, The Netherlands, Sept. 2009.
6. Venugopal, S., Nadiminti, K., Gibbins, H., and Buyya, R.: Designing a Resource Brokerfor Heterogeneous Grids, Software: Practice and Experience, 38(8):793-825, ISSN: 0038-
Wiley Press, New York, USA, 2008.

7. Pandey, S., Voorsluys, W., Rahman, M., Buyya, R., Dobson, J., and Chiu, K.: A GridWorkflow Environment for Brain Imaging Analysis on Distributed Systems, Concurrency and Computation: Practice and Experience, Wiley Press, New York, USA, 2009.
8. Pandey, S., Dobson, J.E., Voorsluys, W., Vecchiola, C., Karunamoorthy, D., Chu, X., and Buyya, R.: Workflow Engine: fMRI Brain Image Analysis on Amazon EC2 and S3 Clouds, The Second IEEE International Scalable Computing Challenge (SCALE 2009) in conjunction with CCGrid 2009, Shanghai, China, 2009.