INTERNATIONAL JOURNAL OF PURE AND APPLIED RESEARCH IN ENGINEERING AND TECHNOLOGY

A PATH FOR HORIZING YOUR INNOVATIVE WORK

OVERVIEW OF CLOUD COMPUTING, SECURITY AND RESEARCH CHALLENGES

ANKUR. O. BANG , PRABHAKAR. L. RAMTEKE

Accepted Date:

27/02/2013

Publish Date:

01/04/2013

Keywords

Cloud

Security

Corresponding Author

Mr. Ankur. O. Bang

Student of first year M.E. (Computer Science and Engineering) HVPM's College of Engineering and Technology, Amravati - SantGadge Baba Amravati University.

Abstract

Cloud computing has recently emerged as a new paradigm for hosting and delivering services over the internet .In this paper we focus on the overview about the cloud computing and the different issues about the security of the cloud and the research challenges in the forth coming future .In this paper we provide the introductory information of related technologies about cloud computing and the classification and description about the security domains in cloud computing and also discussed different research challenges in cloud computing .The aim of this paper is to provide better understanding about the security concept in cloud and to boost up the research in different area about the cloud computing .

1. INTRODUCTION

With the rapid development of processing and storage technologies and the success of the Internet, computing resources have become cheaper, more powerful and more ubiquitously available than ever before. This technological trend has enabled the realization of a new computing model called cloud computing, in which resources (e.g., CPU and storage) are provided as general utilities that can be leased and released by users through the Internet in an on-demand fashion. In this paper we provide a introductory information about the cloud computing and the related technologies such as grid computing, utility computing, Autonomic computing .Classification about the security issues and the importance of the security to the cloud are also highlighted- which include the following section under discussion Traditional security ,Availability ,Third-party data control . Along with this we have focused on the research challenges and provided information about them , which we hope will give the directions for the further research work in the cloud computing

The remainder of this paper is organized as follows. In Sect. 2 we provide an overview of cloud computing and other related technologies. In Sect. 3, we describe the security of cloud computing. The key features about the research challenges are detailed in Sect. 4. Section 5. Contains the conclusion and summery of the paper.

2. OVERVIEW OF CLOUD COMPUTING

This section presents a general overview of cloud computing, including its definition and a comparison with related concepts.

Definitions

The main idea behind cloud computing is not a new one .John McCarthy in the 1960s already envisioned that computing facilities will be provided to the general public like a utility [1]. The term “cloud” has also been used in various contexts such as describing large ATM networks in the1990s. However, it was after Google’s CEO Eric Schmidt used the word to describe the business model of providing services across the Internet in 2006, that the term really started to gain popularity. Since then, the term cloud computing has been used mainly as a marketing term in a variety of contexts to represent many different ideas. Certainly, the lack of a standard definition of cloud computing has generated not only market hypes, but also a fair amount of skepticism and confusion. For this reason, recently there has been work on standardizing the definition of cloud computing. As an example, the work in [2] compared over 20different definitions from a variety of sources to confirm a standard definition. In this paper, we adopt the definition of cloud computing provided by The National Institute of Standards and Technology (NIST) [3], as it covers, in our opinion , all the essential aspects of cloud computing:

NIST definition of cloud computing *Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources* (*e.g., networks, servers, storage, applications, and services*) *that can be rapidly provisioned and released with minimal management effort or service provider interaction*.

The main reason for the existence of different perceptions of cloud computing is that cloud computing, unlike other technical terms, is not a new technology, but rather a new operations model that brings together a set of existing technologies to run business in a different way. Indeed, most of the technologies used by cloud computing, such as virtualization and utility-based pricing, are not new. Instead, cloud computing leverages these existing technologies to meet the technological and economic requirements of today’s demand for information technology.

Related technologies

Cloud computing is often compared to the following technologies, each of which shares certain aspects with cloud computing:

Grid Computing: Grid computing is a distributed computing paradigm that coordinates networked resources to achieve a common computational objective. The development of Grid computing was originally driven by scientific applications which are usually computation-intensive. Cloud computing is similar to Grid computing in that it also employs distributed resources to achieve application-level objectives. However, cloud computing takes one step further by leveraging virtualization technologies at multiple levels (hardware and application platform) to realize resource sharing and dynamic resource provisioning.

Utility Computing: Utility computing represents the model of providing resources on-demand and charging customers based on usage rather than a flat rate. Cloud computing can be perceived as a realization of utility computing. It adopts a utility-based pricing scheme entirely for economic reasons. With on-demand resource provisioning and utility base depricing, service providers can truly maximize resource utilization and minimize their operating costs.

Virtualization: Virtualization is a technology that abstracts away the details of physical hardware and provides virtualized resources for high-level applications. A virtualized server is commonly called a virtual machine (VM). Virtualization forms the foundation of cloud computing, as it provides the capability of pooling computing resources from clusters of servers and dynamically assigning or reassigning virtual resources to applications on-demand.

Autonomic Computing: Originally coined by IBM in2001, autonomic computing aims at building computing systems capable of self-management, i.e. reacting to internal and external observations without human intervention. The goal of autonomic computing is to overcome the management complexity of today’s computer systems. Although cloud computing exhibits certain autonomic features such as automatic resource provisioning, its objective is to lower the resource cost rather than to reduce system complexity .In summary, cloud computing leverages virtualization technology to achieve the goal of providing computing resources as a utility. It shares certain aspects with grid computing and autonomic computing but differs from them in other aspects. Therefore, it offers unique benefits and imposes distinctive challenges to meet its requirements.

3. SECURITY OF THE CLOUD

What are the “security” concerns that are preventing companies from taking advantage of the cloud? Numerous studies, for example IDC’s 2008 Cloud Services User Survey [4] of IT executives, cite security as the number one challenge for cloud users. In this section we present a taxonomy of the “security” concerns. The Cloud Security Alliance’s initial report [5] contains a different sort of taxonomy based on 15 different security domains and the processes that need to be followed in an overall cloud deployment. We categorize the security concerns as:

Traditional security

Availability

Third-party data control

Traditional Security

These concerns involve computer and network intrusions or attacks that will be made possible or at least easier by moving to the cloud. Cloud providers respond to these concerns by arguing that their security measures and processes are more mature and tested than those of the average company. Another argument, made by the Jericho Forum [6], is: "It could be easier to lock down information if it's administered by a third party rather than in-house, if companies are worried about insider threats… In addition, it may be easier to enforce security via contracts with online services providers than via internal controls." Concerns in this category include:

TS1. VM-level attacks. Potential vulnerabilities in the hypervisor or VM technology used by cloud vendors are a potential problem in multi-tenant architectures. Vulnerabilities have appeared in VMW are[7], Xen[8], and Microsoft’s Virtual PC and Virtual Server [9]. Vendors such as Third Brigade [10] mitigate potential VM-level vulnerabilities through monitoring and firewalls.

TS2. Cloud provider vulnerabilities. These could be platform-level, such as an SQL-injection or cross-site scripting vulnerability in salesforce.com. For instance, there have been a couple of recent Google Docs vulnerabilities [11] and [12]. The Google response to one of them is here: [13]. There is nothing new in the nature of these vulnerabilities; only their setting is novel. In fact, IBM has repositioned its Rational App Scan tool, which scans for vulnerabilities in web services as a cloud security service (see Blue Cloud Initiative [14]).

TS3. Phishing cloud provider. Phishers and other social engineers have a new attack vector, as the Sales force phishing incident [15] shows.

TS4. Expanded network attack surface. The cloud user must protect the infrastructure used to connect and interact with the cloud, a task complicated by the cloud being outside the firewall in many cases. For instance, [16] shows an example of how the cloud might attack the machine connecting to it.

TS5. Authentication and Authorization. The enterprise authentication and authorization framework does not

naturally extend into the cloud. How does a company meld its existing framework to include cloud resources? Furthermore, how does an enterprise merge cloud security data (if even available) with its own security metrics and policies?

TS6. Forensics in the cloud. This blog posting on the CLOIDIFIN [17] project summarizes the difficulty of cloud forensic investigations: *"Traditional digital forensic methodologies permit investigators to seize equipment and perform detailed analysis on the media and data recovered. The likelihood therefore, of the data being removed, overwritten, deleted or destroyed by the perpetrator in this case is low. More closely linked to a CC environment would be businesses that own and maintain their own multi-server type infrastructure, though this would be on a far smaller scale in comparison. However, the scale of the cloud and the rate at which data is overwritten is of concern."*

Availability

These concerns center on critical applications and data being available. Well-publicized incidents of cloud outages include Gmail (one-day outage in mid-October 2008 [18]), Amazon S3 (over seven-hour downtime on July 20, 2008 [19]), and FlexiScale (18-hour outage on October 31, 2008 [20]).

A1. Uptime. As with the Traditional Security concerns, cloud providers argue that their server uptime compares well with the availability of the cloud user's own data centers.

Besides just services and applications being down, this includes the concern that a third-party cloud would not scale well enough to handle certain applications. SAP's CEO, Leo Apotheker said: *"There are certain things that you cannot run in the cloud because the cloud would collapse…Don't believe that any utility company is going to run its billing for 50 million consumers in the cloud."* (11/24/08, searchSAP.com)

A2. Single point of failure. Cloud services are thought of as providing more availability, but perhaps not – there are more single points of failure and attack.

A3. Assurance of computational integrity. Can an enterprise be assured that a cloud provider is faithfully running a hosted application and giving valid results? For example, Stanford's Folding@Home project gives the same task to multiple clients to reach a consensus on the correct result.

Third-party data control

The legal implications of data and applications being held by a third party are complex and not well understood. There is also a potential lack of control and transparency when a third party holds the data. Part of the hype of cloud computing is that the cloud can be implementation independent, but in reality regulatory compliance requires transparency into the cloud.

All this is prompting some companies to build private clouds to avoid these issues and yet retain some of the advantages of cloud computing. For example, Benjamin Linder, Scalent System's CEO, says [21]: *"What I find as CEO of a software company in this space, Scalent Systems, is that most enterprises have a hard time trusting external clouds for their proprietary and high-availability systems. They are instead building internal "clouds", or "utilities" to serve their internal customers in a more controlled way."*

BL1. Due diligence. If served a subpoena or other legal action, can a cloud user compel the cloud provider to respond in the required time-frame? A related question is the provability of deletion, relevant to an enterprise's retention policy: How can a cloud user be guaranteed that data has been deleted by the cloud provider?

BL2. Auditability. Audit difficulty is another side effect of the lack of control in the cloud. Is there sufficient transparency in the operations of the cloud provider for auditing purposes? Currently, this transparency is provided by documentation and manual audits. Information Security Magazine asks [22]: *"How do you perform an on-site audit when you have a distributed and dynamic multi-tenant computing environment spread all over the globe? It may be very difficult to satisfy auditors that your data is properly isolated and cannot be viewed by other customers."*

A related concern is proper governance of cloud-related activity. It's easy, perhaps too easy, to start using a cloud service [23].

One popular auditing guideline is the SAS 51, which defines guidelines for auditors to assess internal controls, for instance controls over the processing of sensitive information. SOX and HIPAA are other well-known regulations. US government agencies generally need to follow guidelines from FISMA, NIST, and FIPS.

Certain regulations require data and operations to remain in certain geographic locations. Cloud providers are beginning to respond with geo-targeted offerings [24].

BL3. Contractual obligations. One problem with using another company's infrastructure besides the uncertain alignment of interests is that there might be surprising legal implications. For instance, here is a passage from Amazon's terms of use [25]:

10.4. Non-Assertion. During and after the term of the Agreement, with respect to any of the Services that you elect to use, you will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, end users, vendors, business partners (including third party sellers on websites operated by or on behalf of us), licensors, sub licensees or transferees, any patent infringement or other intellectual property infringement claim with respect to such Services.

This could be interpreted as implying that after you use EC2, you cannot file infringement claims against Amazon or its customers suggesting that EC2 itself violates any of your patents. It's not clear whether this non-assert would be upheld by the courts, but any uncertainty is bad for business.

BL4. Cloud Provider Espionage. This is the worry of theft of company proprietary information by the cloud provider. For example, Google Gmail and Google Apps are examples of services supported by a private cloud infrastructure. Corporate users of these services are concerned about confidentiality and availability of their data. According to a CNN article [26]:

For Shoukry Tiab, the vice president of IT at Jenny Craig, which uses Postini and Google Maps, the primary concern is security and confidentiality. "Am I nervous to host corporate information on someone else's server? Yes, even if it's Google."

Note that for consumers, there were initially widespread confidentiality concerns about Gmail (see [27]), but now those concerns seem to have faded. We believe this is an example of the Privacy Hump [28]:

Early on in the life cycle of a technology, there are many concerns about how these technologies will be used. These concerns are lumped together forming a "privacy hump" that represents a barrier to the acceptance of a potentially intrusive technology…. Over time, however, the concerns fade, especially if the value proposition is strong enough.

Consumers at least seem to have decided that, in this case, the dangers of placing their data in the cloud were outweighed by the value they received.

BL5. Data Lock-in. How does a cloud user avoid lock-in to a particular cloud-computing vendor? The data might itself be locked in a proprietary format, and there are also issues with training and processes. There is also the problem of the cloud user having no control over frequent changes in cloud-based services (see [29]). Coghead [30] is one example of a cloud platform whose shutdown left customers scrambling to re-write their applications to run on a different platform. Of course, one answer to lock-in is standardization, for instance GoGrid API [31].

BL6. Transitive nature. Another possible concern is that the contracted cloud provider might itself use subcontractors, over whom the cloud user has even less control, and who also must be trusted. One example is the online storage service called The Linkup, which in turn used an online storage company called Nirvanix. The Linkup shutdown after losing sizeable amounts of customer data, which some say was the fault of Nirvanix [32]. Another example is Carbonite [33], who is suing its hardware providers for faulty equipment causing loss of customer data.

4. RESEARCH CHALLENGES

Although cloud computing has been widely adopted by the industry, the research on cloud computing is still at an early stage. Many existing issues have not been fully addressed, while new challenges keep emerging from industry applications. In this section, we summarize some of the challenging research issues in cloud computing.

6.1 Automated service provisioning One of the key features of cloud computing is the capability of acquiring and releasing resources on-demand. The objective of a service provider in this case is to allocate and de-allocate resources from the cloud to satisfy its service level objectives (SLOs), while minimizing its operational cost. However, it is not obvious how a service provider can achieve this objective. In particular, it is not easy to determine how to map SLOs such as QoS requirements to low-level resource requirement such as CPU and memory requirements. Furthermore, to achieve high agility and respond to rapid demand fluctuations such as in flash crowd effect, the resource provisioning decisions must be made online. Automated service provisioning is not a new problem. Dynamic resource provisioning for Internet applications has been studied extensively in the past [34, 35]. These approaches typically involve: (1) Constructing an application performance model that predicts the number of application instances required to handle demand at each particular level in order to satisfy QoS requirements; (2) Periodically predicting future demand and determining resource requirements using the performance model; and (3) Automatically allocating resources using the predicted resource requirements. Application performance model can be constructed using various techniques, including Queuing theory [34],

Control theory [36] and Statistical Machine Learning [37]. Additionally, there is a distinction between proactive and reactive resource control. The proactive approach uses predicted demand to periodically allocate resources before they are needed. The reactive approach reacts to immediate demand fluctuations before periodic demand prediction is available. Both approaches are important and necessary for effective resource control in dynamic operating environments.

6.2 Virtual machine migration

Virtualization can provide significant benefits in cloud computing by enabling virtual machine migration to balance load across the data center. In addition, virtual machine migration enables robust and highly responsive provisioning in data centers. Virtual machine migration has evolved from process migration techniques [38]. More recently, Xen [39] and VMWare [40] have implemented “live” migration of VMs that involves extremely short downtimes ranging from tens of milliseconds to a second. Clark et al. [41] pointed out that migrating an entire OS and all of its applications as one unit allow avoiding many of the difficulties faced by process level migration approaches, and analyzed the benefits of live migration of VMs. The major benefits of VM migration is to avoid hotspots; however, this is not straightforward. Currently, detecting workload hotspots and initiating a migration lacks the agility to respond to sudden workload changes. Moreover, the in memory state should be transferred consistently and efficiently, with integrated consideration of resources for applications and physical servers.

6.3 Server consolidation

Server consolidation is an effective approach to maximize resource utilization while minimizing energy consumption in a cloud computing environment. Live VM migration technology is often used to consolidate VMs residing on multiple under-utilized servers onto a single server, so that the remaining servers can be set to an energy-saving state. The problem of optimally consolidating servers in a data center is often formulated as a variant of the vector bin-packing problem [42], which is an NP-hard optimization problem. Various heuristics have been proposed for this problem [43, 44]. Additionally, dependencies among VMs, such as communication requirements, have also been considered recently [45]. However, server consolidation activities should not hurt application performance. It is known that the resource usage (also known as the footprint [46]) of individual VMs may vary over time [47]. For server resources that are shared among VMs, such as bandwidth, memory cache and disk I/O, maximally consolidating a server may result in resource congestion when a VM changes its footprint on the server [48]. Hence, it is sometimes important to observe the fluctuations of VM footprints and use this information for effective server consolidation. Finally, the system must quickly react to resource congestions when they occur [47].

6.4 Energy management

Improving energy efficiency is another major issue in cloud computing. It has been estimated that the cost of powering and cooling accounts for 53% of the total operational expenditure of data centers [50]. In 2006, data centers in the US consumed more than 1.5% of the total energy generated in that year, and the percentage is projected to grow 18% annually [51]. Hence infrastructure providers are under enormous pressure to reduce energy consumption. The goal is not only to cut down energy cost in data centers, but also to meet government regulations and environmental standards. Designing energy-efficient data centers has recently received considerable attention. This problem can be approached from several directions. For example, energyefficient hardware architecture that enables slowing down CPU speeds and turning off partial hardware components [52] has become commonplace. Energy-aware job scheduling [54] and server consolidation [53] are two other ways to reduce power consumption by turning off unused machines. Recent research has also begun to study energy-efficient network protocols and infrastructures [55]. A key challenge in all the above methods is to achieve a good trade-off between energy savings and application performance. In this respect, few researchers have recently started to investigate coordinated solutions for performance and power management in a dynamic cloud environment [56].

6.5 Traffic management and analysis

Analysis of data traffic is important for today's data centers. For example, many web applications rely on analysis of traffic data to optimize customer experiences. Network operators also need to know how traffic flows through the network in order to make many of the management and planning decisions. However, there are several challenges for existing traffic measurement and analysis methods in Internet Service Providers (ISPs) networks and enterprise to extend to data centers. Firstly, the density of links is much higher than that in ISPs or enterprise networks, which makes the worst case scenario for existing methods. Secondly, most existing methods can compute traffic matrices between a few hundred end hosts, but even a modular data center can have several thousand servers. Finally, existing methods usually assume some flow patterns that are reasonable in Internet and enterprises networks, but the applications deployed on data centers, such as MapReduce jobs, significantly change the traffic pattern. Further, there is tighter coupling in application's use of network, computing, and storage resources, than what is seen in other settings. Currently, there is not much work on measurement and analysis of data center traffic. Greenberg et al. [57] report data center traffic characteristics on flow sizes and concurrent flows, and use these to guide network infrastructure design. Benson et al. [16] perform a complementary study of traffic at the edges of a data center by examining SNMP traces from routers.

6.6 Data security

Data security is another important research topic in cloud computing. Since service providers typically do not have access to the physical security system of data centers, they must rely on the infrastructure provider to achieve full data security. Even for a virtual private cloud, the service provider can only specify the security setting remotely, without knowing whether it is fully implemented. The infrastructure provider, in this context, must achieve the following objectives:

(1) *confidentiality*, for secure data access and transfer, and (2) *auditability*, for attesting whether security setting of applications has been tampered or not. Confidentiality is usually achieved using cryptographic protocols, whereas auditability can be achieved using remote attestation techniques. Remote attestation typically requires a trusted platform module (TPM) to generate non-forgeable system summary (i.e. system state encrypted using TPM's private key) as the proof of system security. However, in a virtualized environment like the clouds, VMs can dynamically migrate from one location to another, hence directly using remote attestation is not sufficient. In this case, it is critical to build trust mechanisms at every architectural layer of the cloud. Firstly, the hardware layer must be trusted using hardware TPM. Secondly, the virtualization platform must be trusted using secure virtual machine monitors [58]. VM migration should only be allowed if both source and destination servers are trusted. Recent work has been devoted to designing efficient protocols for trust establishment and management [58, 59].

6.7 Software frameworks

Cloud computing provides a compelling platform for hosting large-scale data-intensive applications. Typically, these applications leverage MapReduce frameworks such as Hadoop for scalable and fault-tolerant data processing. Recent work has shown that the performance and resource consumption of a MapReduce job is highly dependent on the type of the application [60, 61, 62]. For instance, Hadoop tasks such as sort is I/O intensive, whereas grep requires significant CPU resources. Furthermore, the VM allocated to each Hadoop node may have heterogeneous characteristics. For example, the bandwidth available to a VM is dependent on other VMs collocated on the same server. Hence, it is possible to optimize the performance and cost of a MapReduce application by carefully selecting its configuration parameter values [63] and designing more efficient scheduling algorithms [64, 65]. By mitigating the bottleneck resources, execution time of applications can be significantly improved. The key challenges include performance modeling of Hadoop jobs (either online or offline), and adaptive scheduling in dynamic conditions. Another related approach argues for making MapReduce frameworks energy-aware [66]. The essential idea of this approach is to turn Hadoop node into sleep mode when it has finished its job while waiting for new assignments. To do so, both Hadoop and HDFS must be made energy-aware. Furthermore, there is often a trade-off between performance and energy-awareness. Depending on the objective, finding a desirable trade-off point is still an unexplored research topic.

6.8 Storage technologies and data management

Software frameworks such as MapReduce and its various implementations such as Hadoop and Dryad are designed for distributed processing of data-intensive tasks. As mentioned previously, these frameworks typically operate on Internet-scale file systems such as GFS and HDFS. These file systems are different from traditional distributed file systems in their storage structure, access pattern and application programming interface. In particular, they do not implement the standard POSIX interface, and therefore introduce compatibility issues with legacy file systems and applications. Several research efforts have studied this problem [4, 40]. For instance, the work in [49] proposed a method for supporting the MapReduce framework using cluster file systems such as IBM's GPFS. Patil et al. [66] proposed new API primitives for scalable and concurrent data access.

6.9 Novel cloud architectures

Currently, most of the commercial clouds are implemented in large data centers and operated in a centralized fashion. Although this design achieves economy-of-scale and high manageability, it also comes with its limitations such high energy expense and high initial investment for constructing data centers. Recent work [53, 51] suggests that small size data centers can be more advantageous than big data centers in many cases: a small data center does not consume so much power, hence it does not require a powerful and yet expensive cooling system; small data centers are cheaper to build and better geographically distributed than large data centers. Geo-diversity is often desirable for response time-critical services such as content delivery and interactive gaming. For example, Valancius et al. [51] studied the feasibility of hosting video-streaming services using application gateways (a.k.a. nano-data centers). Another related research trend is on using voluntary resources (i.e. resources donated by end-users) for hosting cloud applications [63]. Clouds built using voluntary resources, or a mixture of voluntary and dedicated resources are much cheaper to operate and more suitable for non-profit applications such as scientific computing. However, this architecture also imposes challenges such managing heterogeneous resources and frequent churn events. Also, devising incentive schemes for such architectures is an open research problems

5. CONCLUSION

Cloud computing is the most popular notion in IT today; even an academic report [6] from UC Berkeley says “Cloud Computing is likely to have the same impact on software that foundries have had on the hardware industry.” They go on to recommend that “developers would be wise to design their next generation of systems to be deployed into Cloud Computing”. While many of the predictions may be cloud hype, we believe the new IT procurement model offered by cloud computing is here to stay.

However, despite the significant benefits offered by cloud computing, the current technologies are not matured enough to realize its full potential. One of key challenge in this domain is security management, Therefore, we believe there is still tremendous opportunity for researchers to make groundbreaking contributions in this field, and bring significant impact to their development in the industry.

In this paper, we have surveyed cloud computing, security issues as well as research directions. As the development of cloud computing technology is still at an early stage, we hope our work will provide a better understanding of cloud computing, and pave the way for further research in this area.

6. REFERENCES

1. Parkhill D (1966) The challenge of the computer utility. Addison- Wesley, Reading

2. Vaquero L, Rodero-Merino L, Caceres J, Lindner M (2009) A break in the clouds: towards a cloud definition. ACM SIGCOMM computer communications review

3. NIST Definition of Cloud Computing v15, csrc.nist.gov/groups/ SNS/cloud-computing/cloud-def-v15.doc

4. IT Cloud Services User Survey, pt.2: Top Benefits & Challenges. http://blogs.idc.com/ie/?p=210.

5. Security Guidance for Critical Areas of Focus in Cloud Computing. http://www.cloudsecurityalliance.org/guidance/csaguide.pdf.

6. Don't cloud your vision. http://www.ft.com/cms/s/0/303680a6-bf51-11dd-ae63-0000779fd18c.html?nclick_check=1.

7. VMWare vulnerability. http://securitytracker.com/alerts/2008/Feb/1019493.html.

8. Xen vulnerability. http://secunia.com/advisories/26986

9. VirtualPC vulnerability. http://www.microsoft.com/technet/security/bulletin/ms07-049.mspx.

10. Third Brigade. http://www.thirdbrigade.com.

11. Google Docs Glitch Exposes Private Files. http://www.pcworld.com/article/160927/google_docs_glitch_exposes_private_files.html.

12. Security issues with Google Docs. http://peekay.org/2009/03/26/security-issues-with-google-docs/.

13. Google's response to Google Docs concerns. http://googledocs.blogspot.com/2009/03/just-to-clarify.html.

14. Blue Cloud. http://www-03.ibm.com/press/us/en/pressrelease/26642.wss.

15. Salesforce.com Warns Customers of Phishing Scam. http://www.pcworld.com/businesscenter/article/139353/salesforcecom_warns_customers_of_phishing_scam.html.

16. Security Evaluation of Grid Environments. https://hpcrd.lbl.gov/HEPCybersecurity/HEP-Sec-Miller-Mar2005.ppt.

17. CLOIDIFIN. http://community.zdnet.co.uk/blog/0,1000000567,2000625196b,00.htm?new_comment.

18. Extended Gmail outage hits Apps admins. http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9117322.

19. Amazon S3 Availability Event: July 20, 2008. http://status.aws.amazon.com/s3-20080720.html.

20. FlexiScale Suffers 18-Hour Outage. http://www.thewhir.com/web-hosting-news/103108_FlexiScale_Suffers_18_Hour_Outage.

21. Disaster-Proofing The Cloud. http://www.forbes.com/2008/11/24/cio-cloud-disaster-tech-cio-cx_dw_1125cloud.html.

22. How to Secure Cloud Computing. http://searchsecurity.techtarget.com/magOnline/0,sid14_gci1349550,00.html.

23. Storm clouds ahead. http://www.networkworld.com/news/2009/030209-soa-cloud.html?page=1.

24. Amazon EC2 Crosses the Atlantic. http://aws.amazon.com/about-aws/whats-new/2008/12/10/amazon-ec2-crosses-the-atlantic/.

25. Amazon's terms of use. http://aws.amazon.com/agreement.

26. Why Google Apps is not being adopted. http://money.cnn.com/2008/08/19/technology/google_apps.fortune/index.htm.

27. End-User Privacy in Human–Computer Interaction. http://www.cs.cmu.edu/~jasonh/publications/fnt-end-user-privacy-in-human-computer-interaction-final.pdf.

28. Organizations urge Google to suspend GMail. http://www.privacyrights.org/ar/GmailLetter.htm.

29. Cloud computing: Don't get caught without an exit strategy. http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9128665&source=NLT_AM.

30. Cloud Bursts as Coghead Calls It Quits. http://blogs.zdnet.com/collaboration/?p=349.

31. GoGrid API. http://www.gogrid.com/company/press-releases/gogrid-moves-api-specification-to-creativecommons.php.

32. Loss of customer data spurs closure of online storage service 'The Linkup'. http://www.networkworld.com/news/2008/081108-linkup-failure.html?page=1.

33. Latest cloud storage hiccups prompts data security questions.

34. Urgaonkar B et al (2005) Dynamic provisioning of multi-tier Internet applications. In: Proc of ICAC

35. Zhang Q et al (2007) A regression-based analytic model for dynamic resource provisioning of multi-tier applications. In: Proc ICAC

36. Kalyvianaki E et al (2009) Self-adaptive and self-configured CPU resource provisioning for virtualized servers using Kalman filters. In: Proc of international conference on autonomic computing

37. Bodik P et al (2009) Statistical machine learning makes automatic control practical for Internet datacenters. In: ProcHotCloud

38. Greenberg A, Jain N et al (2009) VL2: a scalable and flexible data center network. In: Proc SIGCOMM

39. Osman S, Subhraveti D et al (2002) The design and implementation of zap: a system for migrating computing environments. In: Proc of OSDI

40. XenSourceInc, Xen, www.xensource.com

41. VMWare ESX Server, www.vmware.com/products/esx

42. Clark C, Fraser K, Hand S, Hansen JG, Jul E, Limpach C, Pratt I, Warfield A (2005) Live migration of virtual machines. In: Proc of NSDI

43. Chekuri C, Khanna S (2004) On multi-dimensional packing problems. SIAM J Comput 33(4):837–851

44. Li B et al (2009) EnaCloud: an energy-saving application live placement approach for cloud computing environments. In: Proc of international conf on cloud computing

45. Srikantaiah S et al (2008) Energy aware consolidation for cloud computing. In: Proc of HotPower

46. Meng X et al (2010) Improving the scalability of data center networks with traffic-aware virtual machine placement. In: Proc INFOCOM

47. Sonnek J et al (2009) Virtual putty: reshaping the physical footprint of virtual machines. In: Proc of HotCloud

48. Wood T et al (2007) Black-box and gray-box strategies for virtual machine migration. In: Proc of NSDI

49. Padala P, Hou K-Y et al (2009) Automated control of multiple virtualized resources. In: Proc of EuroSys

50. Ananthanarayanan R, Gupta K et al (2009) Cloud analytics: do we really need to reinvent the storage stack? In: Proc of HotCloud

51. Hamilton J (2009) Cooperative expendable micro-slice servers (CEMS): low cost, low power servers for Internet-scale services In: Proc of CIDR

52. Valancius V, Laoutaris N et al (2009) Greening the Internet with nano data centers. In: Proc of CoNext

53. Brooks D et al (2000) Power-aware microarchitecture: design and modeling challenges for the next-generation microprocessors, IEEE Micro

54. Church K et al (2008) On delivering embarrassingly distributed cloud services. In: Proc of HotNets

55. Vasic N et al (2009) Making cluster applications energy-aware. In: Proc of automated ctrl for datacenters and clouds

56. IEEE P802.3az Energy Efficient Ethernet Task Force, www. ieee802.org/3/az

57. Kumar S et al (2009) vManage: loosely coupled platform and virtualization management in data centers. In: Proc of international conference on cloud computing

58. Krautheim FJ (2009) Private virtual infrastructure for cloud computing. In: Proc of HotCloud

59. Santos N, Gummadi K, Rodrigues R (2009) Towards trusted cloud computing. In: Proc of HotCloud

60. Sandholm T, Lai K (2009) MapReduce optimization using regulated dynamic prioritization. In: Proc of SIGMETRICS/ Performance

61. ZahariaMet al (2009) ImprovingMapReduce performance in heterogeneous environments. In: Proc of HotCloud

62. Chandra A et al (2009) Nebulas: using distributed voluntary resources to build clouds. In: Proc of HotCloud

63. Sandholm T, Lai K (2009) MapReduce optimization using regulated dynamic prioritization. In: Proc of SIGMETRICS/ Performance

64. ZahariaMet al (2009) ImprovingMapReduce performance in heterogeneous environments. In: Proc of HotCloud

65. Patil S et al (2009) In search of an API for scalable file systems: under the table or above it?HotCloud.

Mr.Ankur O. Bang

Student of first year M.E. (Computer Science And Engineering) HVPM's College Of Engineering And Technalogy , Amravati -SantGadge Baba Amravati University

Mr. Prabhakar L. Ramteke

H.O.D (I.T) HVPM's College Of Engineering And Technalogy , Amravati-SantGadge Baba Amravati University