Experiences on Setting up On-Premise Enterprise Cloud Using OpenStack

R. Ananthalakshmi Ammal(✉), K.B. Aneesh Kumar, Beniwal Alka, and B. Renjith

Broadcast and Communications Group,
Centre for Development of Advanced Computing (CDAC), Thiruvananthapuram, India
{lakshmi,aneesh_kb,alkab,renjithb}@cdac.in

Abstract. Cloud Computing allows users to access a shared pool of computing resources which include networks, servers, storage, applications and services. One of the major advantages of cloud is that the resources can be rapidly provisioned. However, many enterprises are still unwilling to move their IT infrastructure to cloud. Major concerns of enterprises are the initial expenditure on capital and the cost of maintenance of the cloud infrastructure which requires a new set of expertise. The huge licensing cost of proprietary cloud solutions is also a major concern. Hence open source cloud solutions are gaining popularity. In this paper we discuss major challenges faced while implementing an on-premise cloud using OpenStack and the solutions developed to overcome the same. The source code availability and strong support from a wide community made us choose OpenStack. The cloud usage model, introduced for better ROI (Return on Investment), which is capable of providing both Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) is also discussed. Our experience proved that irrespective of all the challenges enterprise can save both time and money with the adoption of an on-premise open source cloud with a suitable service model.

Keywords: OpenStack · ROI · IaaS · Cloud usage model · Experience with OpenStack

1 Introduction

Efficient and fair provisioning of resources for achieving optimal resource utilization is a challenging task for organizations. Ensuring on-time availability of resources for development teams is yet another issue faced by management. Traditional resource allocation and management is time consuming and fails in timely allocation. In contrast to the traditional approach, cloud based solutions provide many benefits in terms of virtualization, scalability, flexibility and provisioning. The advantages of using cloud in resource provisioning and how it differs from the traditional approach is discussed by Kepes [1]. Cloud computing provides primarily three service models, namely Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) [2]. There are primarily three deployment models which are public cloud, private cloud and hybrid cloud [3]. Public cloud offers scalability, elasticity and a pay-per-use model on shared infrastructure. Private Cloud is owned by an organization, hosted and

Y. Zhang et al. (Eds.): CloudComp 2015, LNICST 167, pp. 114–125, 2016.
DOI: 10.1007/978-3-319-38904-2_12

operated internally. Hybrid cloud is a mixture of private cloud and public cloud services with orchestration between the two. Private cloud and public cloud offer similar benefits, but in public cloud there is a concern of data security and compliance. Amid these choices, the decision for a suitable service and deployment model, that is best suited for an organization, is very difficult and depends on the analysis of multitude of factors and requirements.

Cloud Computing frameworks are available from commercial vendors as well as open source community. Some of the popular commercial offerings include 'vCloud' from VMWare and Azure from Microsoft. Open source cloud platforms include Eucalyptus, Nimbus, OpenStack, OpenNebula and CloudStack. Many tradeoff parameters such as price, documentation, features, expertise, etc. have to be considered before finalizing an open source cloud framework compared to commercial offerings. Several studies compare the cloud- frameworks in terms of these factors [3–7]. Studies on open source platforms conclude that OpenStack provides powerful features, especially the support for a wide array of hypervisors. The architecture and features of OpenStack show that it is ideal for large-scale private cloud deployments [8, 9]. Moreover, the support community behind OpenStack is very wide and strong. As a result of the analysis and in consideration of the comparative studies [2–7], we have chosen IaaS deployment model using OpenStack. It also allows leveraging our existing infrastructure to provide PaaS services.

Even though the use cases initially considered for the deployment of cloud system were generic in nature, the cost savings on the infrastructure was one of the prime objectives. In fact, dedicated hardware resources were chosen for the implementation on account of performance [10]. During the deployment we faced multiple challenges to fulfill some of the requirements. However we could overcome those challenges and came up with a successful deployment which is characterized by improved utilization, legitimate provisioning and greater availability of resources. A service portal for users to request specific software and services was introduced and software platform with custom applications were built as part of the cloud image itself. In cloud-networking, we extended existing VLANs (Virtual Local Area Network) of departments as multiple external networks, a feature which was not available by default in OpenStack. As a result of these, a better usage model for the cloud which is acceptable for all users was introduced. This was based on efficient sharing of services and software, which could generate significant ROI for the organization.

The paper is structured as follows: Sect. 2 gives related work and requirements are explained in Sect. 3. The Challenges faced are described in Sect. 4 and solutions developed are explained in Sect. 5. The implementation details are explained in Sect. 6 and results are presented in Sect. 7. Conclusion is given in Sect. 8.

2 Related Work

The main studies on OpenStack are focused on its architecture, components, characteristics and comparison with other cloud solutions, but only a few are focused on the challenges that enterprises have to face during deployment of an open cloud solution.

Keshavarzi et al. [11] discussed six challenges about cloud computing, which include security, Autonomic Resource Management (ARM), adoption, development, benchmarking and big data, in detail but no solutions are developed so far. Huang et al. [12] pointed out several key security problems. The study aims to identify the most vulnerable security threats in cloud computing with emphasis on security of cloud data storage. It also discuss security issues regarding data integrity, data confidentiality, access control and data manipulation to encourage the researchers to come up with some techniques and solutions so that the whole cloud storage system is reliable and trustworthy.

Cloud computing always tries to maximize the utilization of its resources. Private cloud setup requires a lot of investment in hardware implementation. Kashyap et al. [13] proposed a solution of Virtual Machine migration that minimizes the total energy consumption of cloud infrastructure. In cloud, a number of VMs (Virtual machine) are spawn and are destroyed every single minute, resulting in underutilization of physical servers. To solve this problem, cloud-providers use VM migration, in which active VMs are fused to a single physical machine to save energy. Thenceforth, underutilized servers are switched off and the consolidated servers ensure a power efficient green cloud. Raiyani Kashyap has also proposed a VM placement algorithm to migrate VMs based on system load.

Ristov et al. [14] estimated security vulnerabilities of OpenStack cloud framework. They used Nessus 5, the vulnerability and configuration assessment scanner, to exploit OpenStack server node and tenant and Acunetix Web Vulnerability Scanner for OpenStack-dashboard. The study warns about some security issues on deploying an open source cloud. Since intruders can access the cloud source code they can exploit its vulnerabilities. However these vulnerabilities can be mitigated, with software patches. As far as our knowledge, a study that addresses the challenges during and after the deployment of a private OpenStack cloud is not available.

3 Requirements

Embracing the cloud in an organization is a significant step even for mature IT organizations. The first step is identifying the requirements for the cloud. There are some generic advantages of private cloud namely monitoring of resources, flexibility of customization, ability to recover from failure and the ability to scale up or down based on demand.

Each department in our Organisation has their own subnets, IP schemes, network gateway and VLAN. The network configuration has to remain unchanged while migrating to cloud. To achieve the same, extension of existing VLANs into OpenStack network is required. In other words, the objective is to seamlessly integrate the cloud with the existing network without causing any disruption to the existing workflows.

Each department is managing its own hardware and software resources. Many of the servers and storage in each department are underutilized. Moreover these resources are not shared among departments because of physical location, security and authorization constraints. In order to achieve optimal utilization of resources, sharing of resources among departments, without compromising on security, is a requirement that is to be fulfilled.

Security is the biggest concern when it comes to cloud computing. Data security is one of the major concerns today due to which organizations are not fully adopting this technology. Even in private cloud which is built and managed in-house, there are some security issues. Data stored by a department in the cloud should be visible to that department only, denying access even to the cloud administrator.

The demand for resources is always dynamic based on application development needs and most of the time short-lived. So allocation of resources is a complex task affecting the productivity of the Organisation. In many cases provisioning of resources with different operating systems or their variants is a time consuming task which cripples the productivity. There is a strong need to manage the dynamic allocation of resources to various projects.

Organization wants to remove unnecessary overhead, minimize operational cost, maximize resource utilization, and an infrastructure with better ROI. Some of these requirements can be satisfied by the default features of OpenStack while others demand customization. Since the OpenStack code is openly distributed, it can be modified and adapted according to the requirements.

4 Challenges

As discussed, we selected OpenStack, in the light of thorough comparison of multiple cloud platforms. Moving an organization infrastructure from traditional system to cloud is not an easy task. It requires changes in the mindset of people, process and technology. There should be convincing use cases that can rationalize the upfront investment in the cloud and shall be ready to rebut questions like "Is cloud suitable for us? Does the cloud coexist with our existing IT infrastructure and able to meet our performance requirements? Is my data secure in the cloud?" In addition to the above, there are technical challenges to meet specific organization requirements. In this section, we discuss the challenges faced during the implementation, operation and management of a multi-node enterprise OpenStack cloud.

4.1 Selection of Hypervisor

OpenStack is compatible with many hypervisor such as KVM, Xen, LXC, QEMU, UML, Hyper-V etc., which makes difficult to choose from [15]. Selection of hypervisor is always a critical task as it affects the cloud performance and features. Studies were conducted to evaluate performance of different Hypervisors i.e. VMWare ESXi Server, XenServer and KVM in the private cloud using CloudStack and proved that XenServer and ESXi hypervisors exhibit impressive performance in comparison with KVM [16]. It was also proved that there is no perfect hypervisor among Hyper-V, KVM, vSphere and Xen since overheads incurred by each hypervisor can vary significantly depending on the type of application and the resources assigned to it [17]. Different workloads may be best suited for different hypervisors. Our challenge is to find a Hypervisor that is best suited for OpenStack which meets our organization requirements.

4.2 Develop a Cloud Service Usage Model

Although cloud service models and cloud deployment models are well defined, we could not find any best practices or usage models that can be suited for a private cloud deployment. In a highly abstract way, we can say that a usage model based on effective sharing of resource will work for any organization, but to concretize the concept, it requires serious thoughts. It has specific dependency on the service environment of the cloud as well as on the unique requirements of the organization. Hence to build a cloud usage model that adheres to the requirements of an organization is not an easy task and the success of the deployment is much dependent on this model. The main focus of ours was maximizing resource sharing and reducing CAPEX (Capital expenditure), there by maximizing the ROI while maintaining the resource allocation time in its least minimum. Major challenge was to device a usage model for the cloud to address these specific requirements.

4.3 Extension of External VLANs to GRE Based Tenant Network

Departments require individual tenants in cloud and complete isolation in terms of networks and resource. Additionally, the same VLAN network configuration has to be retained in cloud as well. To achieve the same, there should be a mechanism to extend the departmental VLANs to OpenStack Cloud VMs. Although OpenStack provides GRE (Generic Routing Encapsulation)/VLAN based tenant-network isolation, there is no default mechanism to extend existing external VLAN networks to these tenant networks.

4.4 Image Creation for OpenStack

To create instances in OpenStack, Operating System (OS) image has to be built and updated to the image store. There are two ways to obtain a cloud virtual machine image. The simplest way is to download pre-built images. But pre-built images are available only for open-source operating systems. Second way is to create the image manually outside of OpenStack and then upload those images to cloud. There are some drawbacks associated with the former method. Firstly, the pre-built images are not customized according to our needs. Second they have their default administrator username and password. For some OS images, there is no provision to change the administrator username. The most important is the security concern as these pre-built images are uploaded by random users and we cannot trust them. The latter method requires so many steps to be followed to make an image and has to take care of the drivers required for OpenStack environment. This is one of the major challenges faced during the operation of cloud.

4.5 Data Security

The security of data has consistently been cited as the primary barrier to cloud adoption [18]. Storage functionality is provided by three components in OpenStack i.e. Cinder, Swift and Glance. One of the primary concerns is about the security of the data stored in Cinder volumes. There is no provision to encrypt the Cinder volumes for a virtual machine. Encryption of volume can be done in two ways. The users can encrypt it themselves or rely on the cloud provider to do so. On one hand, there is significant security, but high-complexity. On the other, we have ease of use, but limited protection. As per requirement, we require user configurable data encryption.

5 Solution

5.1 Selection of Hypervisor

XenServer and KVM support almost equal features that can be controlled through compute [19]. XenServer is a bare metal Hypervisor whereas KVM is hosted hypervisor. KVM has strong guest isolation with an extra layer of protection against guest breakouts. KVM is rigorously implemented and tested. It is open source software so developers are continuously inspect KVM for flaws. It has the advantage over other x86 hypervisors in terms of lower total cost of ownership and greater flexibility than competing hypervisors [20]. It is part of Linux and uses the regular Linux scheduler and memory management. It makes KVM much smaller and simpler to use; it is also more feature rich. From user's perspective, there is almost no difference in running a Linux OS with KVM disabled and running a Linux OS with KVM enabled, except the speed difference. A user having fair knowledge of Linux can manage KVM Hypervisor. KVM is most widely used hypervisor in OpenStack. Most OpenStack development is done with the KVM hypervisor for which more community support is available Moreover, bugs associated with KVM hypervisor is negligible as compared to other hypervisors [21–23].

5.2 Develop a Cloud Service Usage Model

For generating better ROI, a cloud usage model has been developed. Licensing cost of software components is a big issue for organizations. For efficient utilization of these licenses, a service sharing mechanism is introduced without violating terms and conditions of license providers. A self-service portal that enables users to request infrastructure and platforms as a service is introduced. It contains a service catalog that lists the categories and the services available. The service portal enables reserving as well as requesting the services on demand. Cloud administrator will service the request if enough resources are available and users will be provided with credentials and IP address of the allocated instance. In addition to that, separate storage volumes will be maintained for each of the users and this storage will be attached to the instance based on the request. Users have to use the credentials provided by the administrator to access the system. One VM instance will be shared across multiple users but the volume associated with the instance varies with the user.

5.3 Extension of External VLANs to the GRE Based Tenant Network

Tenant-network provides internal connectivity for its instances (VMs) as well as isolation from other tenant networks. OpenStack supports both GRE and VLAN based tenant networks. Out of which, VLAN based isolation requires the configuration of physical switches to trunk the required VLANs. On the other hand, GRE is an overlay network which uses encapsulation for network traffic. VLAN based tenant network limit the maximum networks to 4096 but GRE has no hard limit on this number. In view of these facts we used GRE for tenant networking with soft-routers provided by OpenStack for external network connectivity. These soft-routers use NAT (Network Address Translation) for linking the IPs of both external and tenant-network. OpenStack generally uses a single flat external network but we require multiple VLANs as external networks.

Available configuration options of OpenStack were insufficient to achieve the above requirement. As a solution, we adopted custom scripting in addition to OpenStack configuration scripts. Figure 1 illustrates the architecture for exposing multiple VLANs as external networks in the cloud. As in the diagram, 'eth1', the physical interface, connected to the trunk-port of the physical switch which is trunked to allow all departmental VLANs. For each of the external VLAN, which has to be extended to the tenant network, a bridge interface 'br-ex*' is created (* represents the network number). These bridge ports (br-ex* and br-eth1) are interlinked through a Linux 'veth-pair' (virtual link) and 'br-eth1' is configured in promiscuous mode so as to allow all network traffic through it. Custom scripting is used to achieve these configurations at the Neutron-node of the OpenStack implementation. OpenStack creates a set of 'veth-pairs' (int-br-ex*, phy-br-ex*) for each of the external networks to the 'br-int' port. For each of these 'veth-pairs' the phy-br-ex* is tagged with the ID of the corresponding external VLAN. Soft-router of the OpenStack uses these virtual links to forward the external network traffic based on the external network to which it is associated. In effect this configuration enables to extend any number of VLANs in the external network to the cloud.

5.4 Image Creation

There is no specific tool to create images for OpenStack cloud. We created images manually since we require customized images that can be directly deployed to OpenStack. Virtual machine images come in different formats. Out of which we selected qcow2 (QEMU copy-on-write version 2) because it supports snapshots and use sparse representation which results in smaller size of the image. Smaller images mean faster uploads. The qcow2 format is commonly used with the KVM hypervisor [24]. We installed KVM hypervisor to create customized images. We have created images for various operating systems along with required software. For Windows images creation, hypervisor specific drivers and tools are required; for example: VirtIO for KVM and XenServer tools for XenServer/XCP. The image created consists of software and applications required by users along with the Operating System. In other words, the image created is all encompassing and the instances which are created using the images are ready to be used with all the required software, tools and frameworks for end users.

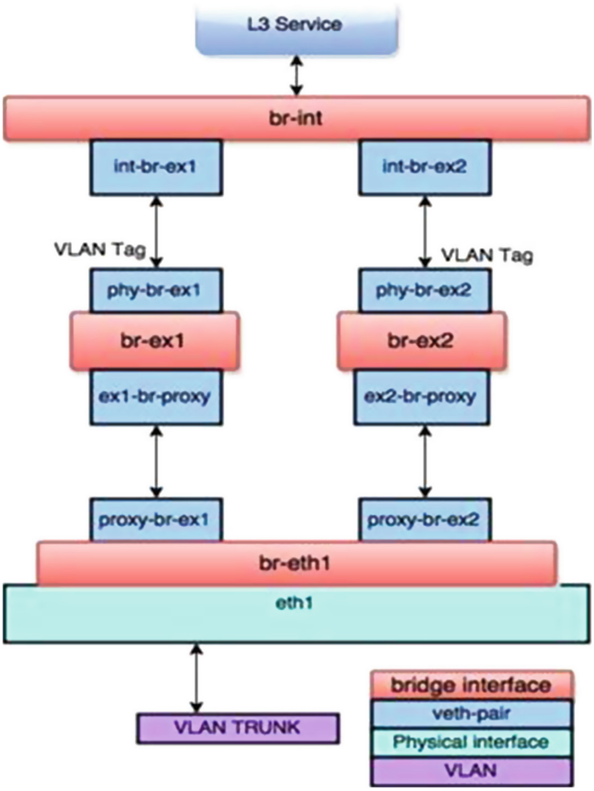

Fig. 1. External networks with multiple VLANs

5.5 Data Security

Encryption of the VM's data prior to writing to disk is provided by the latest release of OpenStack which is Kilo, as of this writing. Still, it is vulnerable for attack since this implementation uses a single, fixed key. It does not provide protection if that key is compromised [25]. Moreover, the fixed key is configured by cloud administrator and users always may not trust cloud administrators. Hence our challenge to allow the cloud user to encrypt their data before sending to the cloud storage remains as an open problem.

6 Implementation

Based on the requirements, we sized cloud infrastructure and used five PMs (Physical Machine). Four of them (PM1, PM2, PM3, PM4) are Dell PowerEdge -C6220 Sled servers, while the fifth one (PM5) is a PowerEdge -R720 server. The corresponding hardware configuration is shown in Fig. 2. On all PMs, we have installed Ubuntu 12.04.4 LTS server. On PM1 we installed an OpenStack controller with Keystone, Glance, Swift, Nova (but without Nova-Compute), MySQL and Horizon services. PM2, PM3 and PM4 are compute nodes (CN). PM5 is configured as a storage node as well as network node. The Compute nodes run KVM, the native Linux VM Hypervisor. All PMs are equipped with 2 Intel Xeon processors with 8 cores each, which are enabled for hyper-threading, resulting in 32 VCPUs (Virtual CPU). On compute nodes, the Nova-Compute service was installed in order to be able to instantiate VMs on top of it.

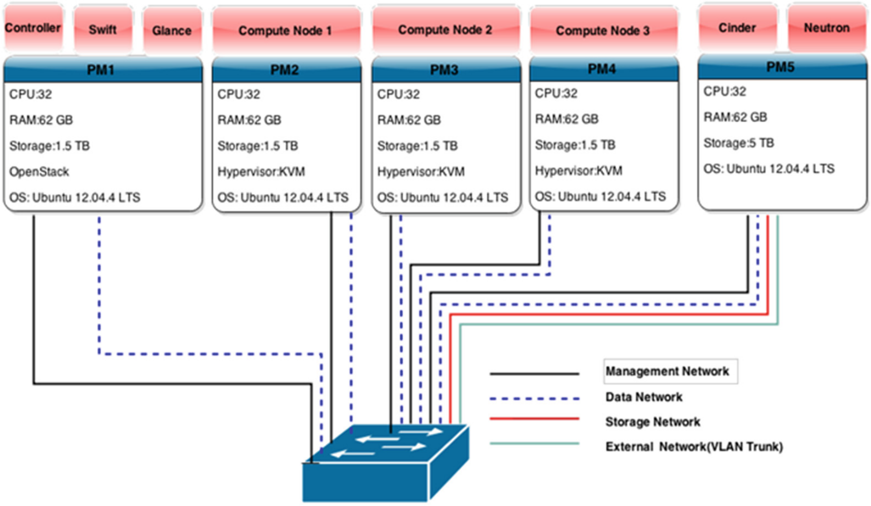

Fig. 2. Cloud infrastructure setup with PMs and hardware/software configuration

We used the management network as 10.176.46.0/24 and data network as 192.168.0.0/16. External networks were multiple VLANs in the range 10.176.1.0/23-10.176.44.0/23. All Compute nodes had two physical NICs (Network Interface Card). One attached to the management network and the other one to data network.

Storage functionality is provided by three components. Swift is the sub-project that delivers object storage. Cinder is the block-storage component that uses standard protocols such as iSCSI. Glance provides a repository for VM images and can use storage from basic file systems or Swift. Using a dedicated storage node or storage subsystem to host, Cinder volumes can be provided. We have used a dedicated storage node of 5 TB.

6.1 Service Request Portal

A user service portal request is also developed which is running outside the cloud framework, and is used by end users to request for any specific platform for a period of time. The service portal is integrated with OpenStack controller using API.

7 Result

Every year we find an increase on demand in the required resources from end users and the computing nodes are increased based on this demand. With the introduction of cloud, currently about 70 % of the physical infrastructure requirement is being met with cloud. Moreover, the time required to fulfill these ever increasing demands either for the initiation of new projects or commissioning of new departments, has been reduced substantially by avoiding delays in procurement and provisioning. In effect, the cloud platform helped the organization to improve utilization, productivity and early adoption of open source technologies with minimal financial overhead. It is also proven that by proper strategy and planning for shared use of software, systems and other services can generate greater cost-saving for organizations. If such a strategy and a cloud usage model are in place, it will have huge impact on the ROI of organizations.

8 Conclusion and Future Work

Successful implementation of cloud computing in an enterprise requires proper planning and understanding of emerging risks, threats, and vulnerabilities and possible countermeasures. We were able to deploy fully operational OpenStack Cloud environment in our enterprise infrastructure successfully with best practices for operating a cloud. There were many challenges to implement a community driven cloud solution but we could successfully overcome those challenges either through customization of OpenStack or introduction of agile models for the cloud usage. Our experience proved that even though a private cloud implementation entails substantial cost from hardware acquisition, deployment, on-going maintenance, management and monitoring, there are corresponding benefits as well. It has reasonably reduced the cost of investment in physical resources and created a noticeable ROI. OpenStack is found to be promising for private cloud implementation and can be adopted by even small organizations.

Overall we feel that, OpenStack's security solution on volume-storage is still in its infancy. Although a few solutions have been developed recently like volume-encryption, all are configurable by cloud service provider only. Therefore a solution, which allows the tenants to configure their encryption mechanism, is highly desirable. This will pose additional challenge on volume-encryption. Hence this research paper will hopefully motivate future researchers to come up with secure user configurable volume encryption algorithms and framework to strengthen the cloud computing security.

Acknowledgements. This work is sponsored by Department of Electronics and Information Technology (DeitY) of the Government of India. We greatly acknowledge Mr. B.M Bhaveja, Group Coordinator, R&D in CC&BT for providing funding support for this project.

References

1. Kepes, B.: Revolution Not Evolution: How Cloud Computing Differs from Traditional IT and Why it Matters. http://broadcast.rackspace.com/hosting_knowledge/whitepapers/Revolution_Not_Evolution-Whitepaper.pdf. Accessed June 2015

2. Gibson, J., Eveleigh, D., Rondeau, R., Tan, Q.: Benefits and challenges of three cloud computing service models. In: Proceedings of 2012 Fourth International Conference on Computational Aspects of Social Networks (CASoN), Sao Carlos, pp. 198–205 (2012)
3. Yadav, S.: Comparative study on open source software for cloud computing platform: Eucalyptus, OpenStack and OpenNebula. Int. J. Eng. Sci. **3** (2013)
4. Thilagavathi, M.: Cloud platforms – a comparison. Int. J. Adv. Res. Comput. Sci. Softw. Eng. **3**, 275–279 (2013)
5. Cloud platform comparison-Comparison of CloudStack, Eucalyptus, vCloud Director and OpenStack. http://www.networkworld.com/article/2189981/tech-primers/cloud-platform-comparison–cloudstack–eucalyptus–vcloud-director-and-openstack.html. Accessed May 2015
6. von Laszewski, G., Diaz, J., Wang, F., Fox, G.C.: Comparison of multiple cloud frameworks. In: Proceedings of 2012 IEEE 5th International Conference on Cloud Computing (CLOUD), Honolulu, HI, pp. 734–741 (2012)
7. Buyyaa, R., Yeoa, C.S., Venugopala, S., Broberga, J., Brandic, I.: Cloud computing and emerging IT platforms: vision, hype, and reality for delivering computing as the 5th utility. Future Gener. Comput. Syst. **5**, 599–616 (2009)
8. Kamboj, R., Arya, A.: OpenStack: open source cloud computing IaaS platform, **4**, 1200–1202 (2014)
9. Sefraoui, O., Aissaoui, M., Eleuldj, M.: OpenStack: toward an open-source solution for cloud computing. Int. J. Comput. Appl. **55**(03), 38–43 (2012). (0975 – 8887)
10. Nasim, R., Kassler, A.J.: Deploying OpenStack: virtual infrastructure or dedicated hardware. In: Proceedings of 2014 IEEE 38th International Conference on Computer Software and Applications Conference Workshops (COMPSACW), Vasteras, pp. 84–89 (2014)
11. Keshavarzi, A., Haghighat, A.T., Bohlouli, M.: Research challenges and prospective business impacts of cloud computing: a survey. In: Proceedings of 2013 IEEE 7th International Conference on Intelligent Data Acquisition and Advanced Computing Systems (IDAACS), Berlin, pp. 731–736 (2013)
12. Huang, C., Qin, Z., Kuo, C.J.: Multimedia storage security in cloud computing: an overview. In: Proceedings of 2011 IEEE 13th International Workshop on Multimedia Signal Processing (MMSP), Hangzhou, pp. 1–6 (2011)
13. Kashyap, R., Chaudhary, S., Jat, P.M.: Virtual machine migration for back-end mashup application deployed on OpenStack environment. In: Proceedings of 2014 International Conference on Parallel, Distributed and Grid Computing (PDGC), Solan, pp. 214–218 (2014)
14. Ristov, S., Gusev, M., Donevski, A.: Security vulnerability assessment of OpenStack cloud. In: Proceedings of 2014 Sixth International Conference on Computational Intelligence, Communication Systems and Networks (CICSyN), Tetova, pp. 95–100 (2014)
15. Hypervisors-OpenStack Hypervisors. http://docs.openstack.org/icehouse/config-reference/content/section_compute-hypervisors.html. Accessed May 2015
16. Reddy, P.V.V., Rajamani, L.: Evaluation of different hypervisors performance in the private cloud with SIGAR framework. Int. J. Adv. Comput. Sci. Appl. (IJACSA) **5**, 60–66 (2014)
17. Hwang, J., Zeng, S., Wu, F., Wood, T.: A component-based performance comparison of four hypervisors. In: Proceedings of 2013 IFIP/IEEE International Symposium on Integrated Network Management (IM 2013), Ghent, pp. 269–276 (2013)
18. Bajwa, M.S., Himani: A concern towards data security in cloud computing. Int. J. Comput. Appl. **114**, 17–19 (2015)
19. Hypervisor Support Matrix. http://docs.openstack.org/developer/nova/support-matrix.html#guest_setup_inject_networking. Accessed May 2015
20. Wilson, G., Day, M., Taylor, B.: KVM hypervisor security you can

21. OpenStack Compute (Nova)-Compute bugs associated with XenServer. https://bugs.launchpad.net/nova/+bugs?field.tag=xenserver,xen. Accessed May 2015
22. OpenStack Compute (Nova)-Compute bugs associated with KVM. https://bugs.launchpad.net/nova/+bugs?field.tag=kvm. Accessed May 2015
23. OpenStack Compute (Nova)-Compute bugs associated with VMware. https://bugs.launchpad.net/nova/+bugs?field.tag=vmware. Accessed May 2015
24. Image guide-OpenStack Image guide. http://docs.openstack.org/image-guide/content/ch_introduction.html. Accessed May 2015
25. Volume Encryption-Volume Encryption with Static Key. http://docs.openstack.org/kilo/config-reference/content/section_volume-encryption.html. Accessed May 2015