International Journal of Information, Control and Computer Sciences
ISSN: 2517-9942
Vol:9, No:6, 2015

Networking the Biggest Challenge in Hybrid Cloud Deployment

Aishwarya Shekhar, Devesh Kumar Srivastava

***Abstract*—**Cloud computing has emerged as a promising direction for cost efficient and reliable service delivery across data communication networks. The dynamic location of service facilities and the virtualization of hardware and software elements are stressing the communication networks and protocols, especially when data centres are interconnected through the internet. Although the computing aspects of cloud technologies have been largely investigated, lower attention has been devoted to the networking services without involving IT operating overhead. Cloud computing has enabled elastic and transparent access to infrastructure services without involving IT operating overhead. Virtualization has been a key enabler for cloud computing. While resource virtualization and service abstraction have been widely investigated, networking in cloud remains a difficult puzzle. Even though network has significant role in facilitating hybrid cloud scenarios, it hasn't received much attention in research community until recently. We propose Network as a Service (NaaS), which forms the basis of unifying public and private clouds. In this paper, we identify various challenges in adoption of hybrid cloud. We discuss the design and implementation of a cloud platform.

***Keywords*—**Cloud computing, networking, infrastructure, hybrid cloud, open stack, Naas.

I. Introduction

HYBRID cloud is defined as a cloud infrastructure composed of two or more cloud infrastructures (private, public, and community clouds) that remain unique entities, but are bound together via technologies and approaches for the purposes of application and data portability. Hybrid cloud is becoming a standard operating model for many organizations specifically; public and private cloud entities will be discussed for a hybrid cloud approach. The approach is based on extension of virtual Open Systems Interconnection (OSI) Layer 2 switching functions from a private cloud and to public clouds, tunneled on an OSI Layer 3 connection. As a result of this hybrid cloud approach, virtual workloads can be migrated from the private cloud to the public cloud and continue to be part of the same Layer 2 domain as in the private cloud, thereby maintaining consistent operational paradigms in both the public and private cloud .Cloud computing offers flexible IT solutions that suites the growingly dynamic organizational demands [1]. Different types of cloud deployment models offer distinctive solutions based on the security requirements of the organization. Hybrid clouds are considered the prominent deployment model that offers the extended flexibility of public cloud resources and the security of private clouds. The three modeling and evaluating firewall deployment models for hybrid clouds are no firewall, regular firewalls and special firewalls that block web traffic. Experimental results illustrated that having a regular firewall would be sufficient for small hybrid clouds (≤150 nodes) to give high point-to-point utilization without compromising the response time. However, for larger networks (up to 300 nodes) a firewall that blocks web traffic is more efficient. Cloud Computing is often described as "resources accessed via a browser over the Internet." However, this definition has become increasingly insufficient to characterize the breadth of applications and use cases for the cloud, and the networks that must support them.

Aishwarya Shekhar and Devesh Kumar are with the School of Computing and Information Technology, Manipal University Jaipur, India (e-mail: aishwaryashekhar26@gmail.com, devesh988@yahoo.com).

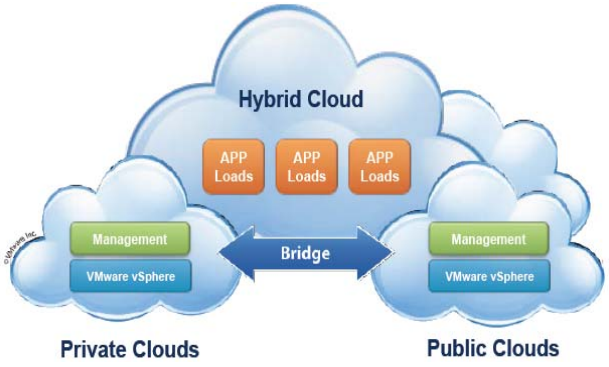

Fig. 1 Hybrid Cloud

A broadening range of end points are accessing the cloud: browser-free device apps, multimedia endpoints such as video and game consoles, sensor networks, servers, and storage. The wire line and wireless network requirements-e.g., jitter, latency, packet loss, protocol support-for these uses vary, and imply that a variety of network capabilities are sometimes necessary: e.g., MPLS for quality of service via class of service to support interactive high definition video in the cloud; optical transport for native protocols such as Fiber Channel for data integration in hybrid cloud scenarios; route control for country compliance issues [2]. Also, distributed topologies and optimized routing are required due to application latency constraints. Moreover, wireless sensor networks and hybrid cloud scenarios such as cloud bursting that require a variety of complex distributed data approaches are driving new transport requirements: guaranteed bandwidth, dynamic bandwidth on demand, and usage-sensitive pricing for fine-grained quantities and duration of bandwidth. Cloud Computing, either as an integrated service or in support of pure-play customers must drive service providers'

international telecommunications infrastructure evolution as well as BSS/OSS.

II. NaaS

Network as a service (NaaS) describes services for network transport connectivity. NaaS involves the optimization of resource allocations by considering network and computing resources as a unified whole. The term "Network as a service" (NaaS) is often used along with other marketing terms like cloud computing, along with acronyms such as infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS) [3]. NaaS sometimes includes the provision of a virtual network service by the owners of the network infrastructure to a third party. Often this includes network virtualization using a protocol such as Open Flow. Network as a Service (NaaS) is sometimes listed as a separate Cloud provider along with Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). This factors out networking, firewalls, related security, etc. from IaaS as is shown in Fig. 2. NaaS can include flexible and extended Virtual Private Network (VPN), bandwidth on demand, custom routing, multicast protocols, security firewall, intrustions detection and prevention, Wide Area Network (WAN), content monitoring and filtering, and antivirus. Some of the models are.

A. Virtual Private Network (VPN)

It extends a private network and the resources contained in the network across networks like the public Internet. It enables a host computer to send and receive data across shared or public networks as if it were a private network with the functionality and policies of the private network.

B. Bandwidth on Demand (BoD)

It is a technique by which network capacity is assigned based on requirements between different nodes or users. Under this model link rates can be dynamically adapted to the traffic demands of the nodes connected to the link.

C. Mobile Network Virtualization

It is a model in which a telecommunications manufacturer or independent network operator builds and operates a network (wireless, or transport connectivity) and sells its communication access capabilities to third parties (commonly mobile phone operators) charging by capacity utilization. A mobile virtual network operator (MVNO), is a mobile communications services provider that does not own the radio spectrum or wireless network infrastructure over which it provides services. Commonly a MVNO offers its communication services using the network infrastructure of an established mobile network operator.

III. Challenges in the Adoption of Hybrid Cloud

Any vendor or consultant who promises a “seamless and straightforward” hybrid cloud is not being realistic. There are technical, cultural, and logistical challenges [4].

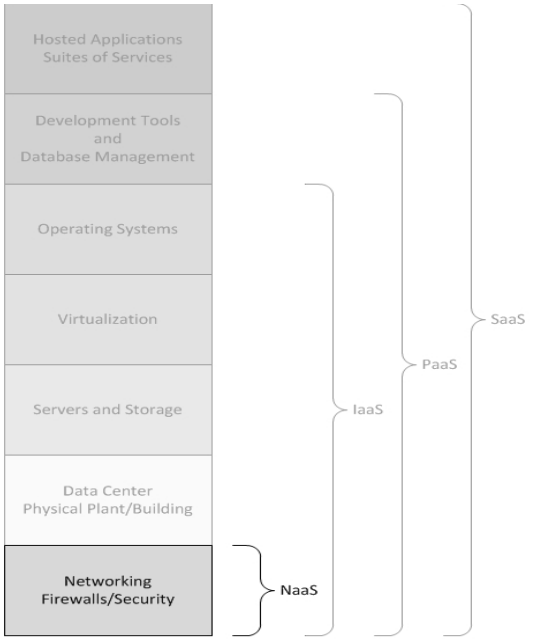

Fig. 2 NaaS

A. Security:

This is usually the first item in any list regarding cloud computing. “Security” is an umbrella term for a wide range of considerations that may impact your vendor choice and implementation strategy. The security concerns themselves vary based on industry, services delivered, compliance and auditing requirements, and other factors, just as they do with a traditional non-cloud approach. However, some security concerns are fairly typical for any organization considering cloud adoption: the security of data when migrating to the cloud; the protection of data that resides in the cloud; the potential impact to application availability if data protection and disaster recovery practices fail; the ability to meet the security requirements of application regulatory compliance standards; and whether the IT organization will be able to maintain enough visibility into and power over their security stance.

B. Networking

Meaningful hybrid integration requires thoughtful network design. Cloud providers are increasingly offering sophisticated networking options, but you will likely find it challenging to natively extend your existing topology to the cloud. Cloud networking is an essential component of cloud computing and forms the foundation for the hybrid cloud. Every vCloud Air service includes a connection to the Internet, one or more public IP addresses, and critical networking capabilities such as load balancing, a firewall, network address translation (NAT), and VPN connectivity via the Edge Gateway. Direct Connect is also available if you are looking for higher speed, secure private line connectivity or if you are cross-connecting to co-located infrastructure.

C. Data and Application Integration

This is one of the first areas of integration between public and private environments that organizations focus on, but you still will face challenges when doing hybrid integration.

D. System Management

Lifecycle management of hybrid cloud systems can be gruesome if done incorrectly. Loud management means the software and technologies designed for operating and monitoring applications, data and services residing in the cloud. Cloud management tools help ensure cloud computing-based resources are working optimally and properly interacting with users and other services

E. Compatibility

There's a good chance that your public and private clouds are running different infrastructure and software stacks. If you have an existing dependency on a particular hypervisor, you may face challenges when dealing with a public cloud that uses a different hypervisor – or doesn't expose one to you at all The world of cloud computing APIs has been constantly evolving since this highly-scalable architecture first gained attention less than five years ago. It is an area with great expectations but little commanding consensus of architecture - so far. Meanwhile, use of services is a common trend, as is use of REST interfaces.

F. Portability

Many start down the hybrid cloud path with visions of moving workloads easily between hosts as the business need dictates. Moving virtual machines and applications between clouds has gotten easier, but you will struggle to move metadata and configurations seamlessly between environments. If the hybrid cloud is based on identical platforms on both ends, this won't be as big of a challenge, but if there's any compatibility mismatch, this will turn into an area of frustration.

G. Tooling and Skills

Hybrid cloud skills and cloud skills in general are in high demand. Some are finding it very difficult to find people with the architectural skills needed to deploy a successful hybrid cloud. A hybrid cloud plan requires expertise in infrastructure configuration, network architecture, application design, and business process automation.

IV. Implementation of Hybrid Cloud

Successful hybrid cloud implementation assumes a well architected private cloud as opposed to simply a well-built traditional IT infrastructure. Adoption of hybrid cloud starts with the transition from a traditional on premises environment to one that includes concepts and supporting technologies to enable functionality normally associated with public cloud: self-provisioning for application owners, dynamic resource scaling, a chargeback model for lines of business, orchestration for automating repeatable tasks and a high visibility management platform to monitor how and where services get deployed. It's the familiarity with the very nature of the public cloud model that has fueled the business and technical requirements in the enterprise for what is essentially an IT as a Service (ITaaS) framework that allows for agile self-service, provisioning and consumption monitoring while simplifying the load on application owners [5]. Because many legacy data center environments were not built with these principles in mind, transitioning can be a challenge. Hybrid cloud also opens the possibility for workload overflow processing or cloud bursting so that applications can bring up new instances as needed in the public part of the hybrid cloud once data center capacity is reached. Application load balancing among these dynamic instances often serves as a core supporting mechanism. It can also be difficult to deterministically know where data center capacity exhausts, however, as well as how many external resources will be needed. Additionally, applications built with the capabilities to traverse public and private cloud boundaries bring about the additional challenges of ensuring that the underlying data is in the right place at the right time as well as dealing with enforcement of governance and security policies regardless of where active instances are operating.

A. Capacity Planning for Hybrid Cloud

To address data center capacity planning, load testing against a proposed infrastructure configuration, trending based on previous growth and building analytic models can help enterprises create accurate estimates of when on premises capacity will start to struggle and require public cloud overflow processing. Of these, taking the time to throw simulated application traffic at a preproduction application environment can be one of the best ways to create awareness of where your environment is in relation to current and projected application processing needs. When properly used, this data plays a key role in helping to determine the amount of on premises resources needed to start with and the sweet spot for the public cloud architecture. While these methods can be time consuming and, in some cases, costly, for organizations committed to a true hybrid cloud strategy, the benefits and long term cost savings of proactive planning as opposed to reactive rearchitecture far outweigh the investment.

B. Hybrid Cloud Data Location and Networking

Ensuring that the correct data is in the right place at the right time can present yet another set of challenges. Enterprises may have requirements for applications to operate in various parts of their private and public cloud for resilience, scalability and flexibility. Oftentimes, applications assume close proximity to the associated data. In hybrid cloud deployments, this is not always practical, so detailed network planning to keep latency at acceptable levels between application front ends servicing client requests and data stores hosting the underlying data must be considered. Traditionally, this has been far easier to control with a fully in house infrastructure, since public cloud offerings tend to provide less flexibility when it comes to networking (e.g. inability to assign more than one or a static IP address to a virtual machine

instance, constraints on how virtual networks can be spanned within the infrastructure, requirements for 'internal' traffic to route 'externally' in certain scenarios, etc.) [6]. Fortunately, this has been steadily improving. Pioneers such as Microsoft with their Azure offering and VMware with vCloud Hybrid Service (vCHS) are providing new dedicated, private networking solutions allowing customers to securely extend their existing data centers while keeping throughput and response levels high. These solutions can keep hybrid cloud connections from going over the public internet, resulting in higher reliability, faster speed, lower latency and improved security. When coupled with advancements in application delivery technology that make it possible to use complex traffic steering algorithms across a fabric of private and public clouds based on business rules, organizations are now equipped with far more control and flexibility than previously possible. The key thing that organizations must keep in mind is to weigh cost versus performance and put workloads and data in the places that make the most sense for the applications, based on how and from where users typically interact with them.

V. Open Stack

OpenStack is a set of software tools for building and managing cloud computing platforms for public and private clouds. Backed by some of the biggest companies in software development and hosting, as well as thousands of individual community members, many think that OpenStack is the future of cloud computing. OpenStack is managed by the OpenStack Foundation, a non-profit which oversees both development and community-building around the project. OpenStack is a free and open-source cloud computing software platform. Users primarily deploy it as an infrastructure as a service (IaaS) solution. The technology consists of a series of interrelated projects that control pools of processing, storage, and networking resources throughout a data center—which users manage through a web-based dashboard, command-line tools. OpenStack lets users deploy virtual machines and other instances which handle different tasks for managing a cloud environment on the fly. It makes horizontal scaling easy, which means that tasks which benefit from running concurrently can easily serve more or less users on the fly by just spinning up more instances. For example, a mobile application which needs to communicate with a remote server might be able to divide the work of communicating with each user across many different instances, all communicating with one another but scaling quickly and easily as the application gains more users. And most importantly, OpenStack is open source software, which means that anyone who chooses to can access the source code, make any changes or modifications they need, and freely share these changes back out to the community at large. It also means that OpenStack has the benefit of thousands of developers all over the world working in tandem to develop the strongest, most robust, and most secure product that they can [7].

A. Components of Open Stack

OpenStack is made up of many different moving parts. Because of its open nature, anyone can add additional components to OpenStack to help it to meet their needs. But the OpenStack community has collaboratively identified nine key components that are a part of the "core" of OpenStack, which are distributed as a part of any OpenStack system and officially maintained by the OpenStack community.

1. Nova

Nova is the primary computing engine behind OpenStack. It is a "fabric controller," which is used for deploying and managing large numbers of virtual machines and other instances to handle computing tasks.

2. Swift

Swift is a storage system for objects and files. Rather than the traditional idea of a referring to files by their location on a disk drive, developers can instead refer to a unique identifier referring to the file or piece of information and let OpenStack decide where to store this information.

3. Cinder

Cinder is a block storage component, which is more analogous to the traditional notion of a computer being able to access specific locations on a disk drive.

4. Neutron

Neutron provides the networking capability for OpenStack. It helps to ensure that each of the components of an OpenStack deployment can communicate with one another quickly and efficiently.

5. Horizon

Horizon is the dashboard behind OpenStack. It is the only graphical interface to OpenStack, so for users wanting to give OpenStack a try, this may be the first component they actually “see.” Developers can access all of the components of OpenStack individually through an application programming interface (API), but the dashboard provides system administrators a look at what is going on in the cloud, and to manage it as needed.

6. Keystone

Keystone provides identity services for OpenStack. It is essentially a central list of all of the users of the OpenStack cloud, mapped against all of the services provided by the cloud which they have permission to use. It provides multiple means of access, meaning developers can easily map their existing user access methods against Keystone.

7. Glance

Glance provides image services to OpenStack. In this case, "images" refers to images (or virtual copies) of hard disks. Glance allows these images to be used as templates when deploying new virtual machine instances.

8. Ceilometer

Ceilometer provides telemetry services, which allow the cloud to provide billing services to individual users of the cloud. It also keeps a verifiable count of each user's system usage of each of the various components of an OpenStack cloud.

9. Heat

Heat is the orchestration component of OpenStack, which allows developers to store the requirements of a cloud application in a file that defines what resources are necessary for that application. In this way, it helps to manage the infrastructure needed for a cloud service to run.

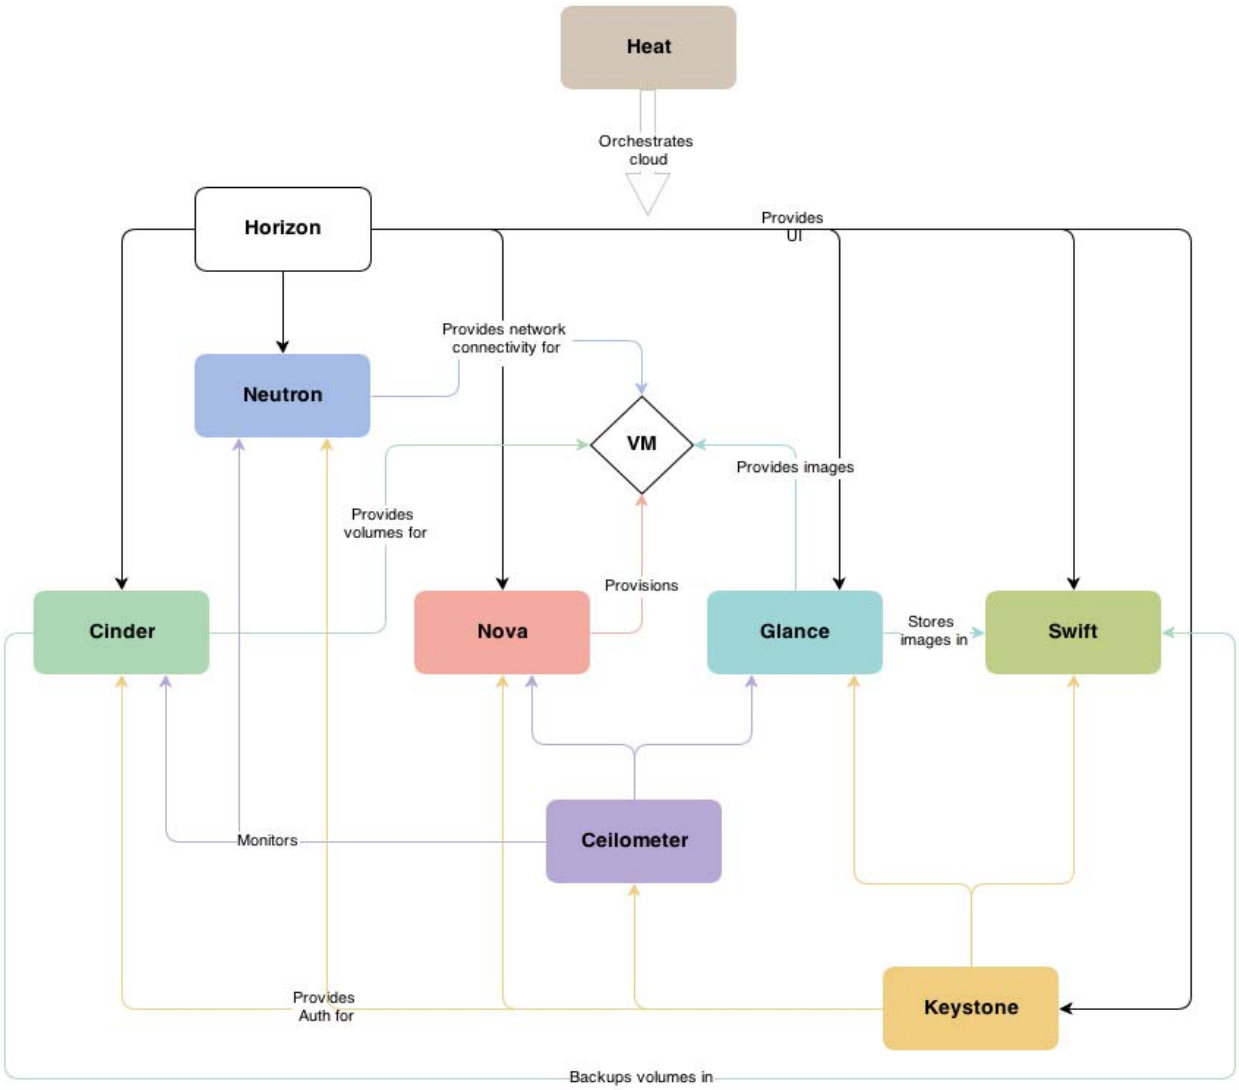

Fig. 3 Open Stack

VI. Conclusion

Compliance and privacy rules may mean keeping some data out of the public cloud, especially for organizations in verticals such as finance and medical. Just as classification, data leak protection, archiving and governance are important in on premises data centers, they are equally if not more critical when public cloud enters the equation. To meet this challenge, cloud management platforms have come to market capable of propagating a unified set of policies across cloud borders. Vendors such as VMware, Microsoft and IBM have launched many new offerings to help companies build better private clouds with the visibility that CIOs dream about and then extend these same principles into public infrastructure. By exposing metadata from the underlying infrastructure, the applications being used and the content being accessed, and then processing these into a usable format, the challenge of compliance and governance can also be met in hybrid cloud. These enablers all have driven the adoption of hybrid cloud strategy in the enterprise, and the outlook is positive. Modern IT has more demands than ever to provide ITaaS solutions that give enterprise lines of business increased agility, rapid provisioning and quicker time to market of application

services. This, combined with the current gap left by an all in public cloud model, all mean one thing: Hybrid cloud is here to stay.

REFERENCES

[1] Rongbo Zhu, Zhili Sun and Jiankun Hu, "Special section: Green computing," *Future Generation Computer Systems* 28, pp. 368-370, 2012.
[2] Jerry Gao, Xiaoying Bai, and Wei-Tek Tsai "Testing as a Service (TaaS) on Clouds" IEEE Seventh International Symposium on Service-Oriented System Engineering 2013.
[3] Pankaj Goyal, 2009. Policy-based Event-driven Services-oriented Architecture for Cloud Services Operation & Management. Prceedings of IEEE International Conference on Cloud Computing (Cloud'09), pp135-138.
[4] L. Riungu-Kalliosaari, O. Taipale, K. Smolander, "Testing in the Cloud: Exploring the Practice,"Special issue on Software Engineering for Cloud Computing, IEEE Software, March/April 2012.
[5] Massoud Pedram, "Energy-Efficient Datacenters," *IEEE Transactions on Computer –Aided Design of Integrated Circuits and Systems,* Vol. 31, N0. 10, pp. 1465-1484, October 2012.
[6] Marios D. Dikaiakos, George Pallis, Dimitrios Katsaros, Pankaj Mehra, and Athena Vakali, "Cloud Computing: Distributed Internet Computing for IT and Scientific," *1089-7801/09 IEEE Computer Society*, 2009.
[7] J Gaurav Kakariya, Prof. Sonali Rangdale, "A Hybrid Cloud Approach For Secure Authorized De-duplication", International Journal of Computer Engineering and Applications, Volume VIII, Issue I, October 14.