A Comparative Study of Network Bandwidth Allocation Protocols for Data Center Networks in Cloud

B Gayathri Kamath
Research Scholar, Dept of CS&E
JNTU, Hyderabad

Ramesh Babu H S
Dept of IS&E
Acharya Institute of Technology, Bangalore

ABSTRACT
Like CPU and memory, network is a most significant and shared resource in the cloud environment. Due to the popularity gained by virtualization technology, shared resources in the cloud are accessed by the tenant son the rent basis. But the network will not be shared as per payment and also will not provide bandwidth guarantee. Sharing the network bandwidth is very difficult because it depends not only on Virtual Machines (VMs) on the same machine but also on VMs that communicate with different machine. Hence the shared nature of the network in the cloud implies that network performance for tenants can vary significantly. Motivated by these factors, in this paper we discuss about different bandwidth allocation protocols and its performance based on key properties like scalability, flexibility, bandwidth guarantee, job completion time and network utilization. A Comparative study on these protocols is done based on the key properties and the outcome of the study reveals the best bandwidth allocation protocol.

Keywords
Bandwidth Guarantee, Datacenter, Network Performance metrics, Bandwidth Allocation Protocols

1. INTRODUCTION
A Data center is a federal warehouse used in companies for storing, managing and broadcasting data and information that support their business. Data center can be physical or virtual. Data centers are used to provide huge amount of compute and storage resources needed by the end users. Cloud Computing is a platform for deploying and running many of today's businesses, in which large-scale data centers will multiplex computing, storage and network resources across multiple tenants. In data centers, the scarce network bandwidth is shared across multiple tenants without any performance guarantees. Hence effective management of network bandwidth will be crucial to handle the network intensive applications. The desirable solution for sharing the cloud networks should meet needs of both cloud users and cloud datacenter providers. Hence various network bandwidth allocation protocols are designed. Bandwidth allocation protocol shares available bandwidth in a cloud in an effective and in a fair manner.

To meet the goal of providing the expected network performance, various bandwidth allocation protocols are measured against the key properties, which the bandwidth allocation protocols should satisfy to provide a desirable solution for sharing the cloud network resources in effective way.

Terry Lam and George Varghese [1] focused on performance guarantees in cloud by fair sharing the unused bandwidth proportionately. Infrastructure As A Service (IAAS) providers propose on-demand computing resources by allocating VMs in data center without performance guarantee to the tenants. Due to this there will be decrease in application productivity, satisfaction of customers and high cost. To overcome the above problem. Microsoft Research team [2] came up with Oktopus, in which virtual network abstractions allow tenants to depict their network requirements. SecondNet [3] objective is to provide bandwidth guarantees among multiple VMs in Virtualized data center (VDC).Advantage of using the VDC is that the resources allocated can be changed as the needs of the tenants. SecondNet uses a forwarding scheme called port-switching source routing (PSSR) that forwards packets using pre-defined port numbers instead of MAC addresses. SecondNet can be build on top of different network topologies.

The first Seawall protocol was implemented only on Windows 7 and Hyper-V.Seawall [4] allows IAAS providers to define the way the bandwidth is shared in a data center network with multiple tenants. Bandwidth is allocated depending on weights of network entities that generate traffic in a proportional way. It prevents malicious tenants from consuming all network resources.

Rodrigues et al. [5] devised a new scheme called Gatekeeper, which meets the requirements of network performance isolation. The authors view with respect to number of VMs, network performance, capability of tenants in malicious environment, and min-max performance guarantees should be scalable and flexible. Gatekeeper objective was to provide guaranteed bandwidth among VMs in a multiple-tenant datacenter, thereby attaininghigh bandwidth consumption. Latency issue in not considered in most of the bandwidth allocation protocol. Experimental evaluation of Gatekeeper was done with two tenants and six machines, which shows that Gatekeeper works fine for simple scenarios.

JianGuo*et al.*[10] et al proposed a bandwidth allocation protocol to fairly share the network resources through the Nash bargaining game approach [11]. In Nash Bargaining game [15], for example in a business, seller and buyer value a product in different way. This creates a surplus. The bargaining solution [9] is that both the seller and the buyer get the win-win solution. That is the surplus created is divided between the two in a fair way. The same analogy holds well for bandwidth allocation at the data center, so that the VMs are allocated with the bandwidth in a fair manner. Nash Bargaining Solution leads to Pareto efficiency or Pareto Optimality. Pareto Optimality allocates resources in a fairly manner [14] without making one individual better off and without making at least one individual worse off. By applying the above approach Falloc achieves high utilization and maintains fairness among VMs in data centers.

The paper emphasizes on the bandwidth allocation protocols for the data center networks in cloud environment. In section II, various network performance metrics are discussed. Different bandwidth allocation techniques are deliberated in Section III. Section IV emphasizes on the comparison of the different bandwidth allocation technique based on the network performance metrics.

2. NETWORK PERFORMANCE METRICS

The key properties are:

2.1 Bandwidth Guarantee

The first property [8] that the bandwidth allocation protocol should meet is to provide tenants guarantee on the minimum network bandwidth they can expect for each VM they buy. Such guarantees can be achieved easily for resources like CPU and memory, but for network to achievethe same guarantee are at lower bounds for the worst-case performance of an application. With bandwidth guarantees for VMs, that runs network sensitive application attains expected performance. For Example, a web service can provide fast data delivery to users if the data transfer between the servers's is guaranteed.

2.2 Scalability

A cloud datacenter supports thousands of physical servers hosting large number of tenants and VMs. The VMs and tenants in the datacenter are tossed dynamically with a high rate. Each datacenter network link is potentially shared by large number of set of VMs. Such a set up must not degrade the network performance and must work fine with large scales and able to manage large amount of state change at high speed.

2.3Flexibility

The basic need of the customers is to achieve deterministic guarantees ensuring expected performance independent of VM placement and migration, traffic and blend of other tenants. Due to the scarce resource in datacenters, bandwidth demand by each VM cannot be guaranteed. Deterministic guarantees can be achieved by allocating network resource to VMs at the cost of underutilization of the physical resource.A practical approach is to guarantee a certain bandwidth considering both the tenant's budget as well as the bandwidth demand of applications. By providing flexibility to tenants to offer service levels to access minimum guarantees by the cloud provider, one can achieve maximum resource efficiency. That is unused bandwidth should be able share with the active sources.

2.4 Network Utilization

Network Utilization **is** also referred to as high utilization. It aims to maximize network utilization without leaving network resources underutilized when there are unsatisfied demands. For example [8], we would like an application to use the entire bandwidth when no other application is active. This can significantly improve the performance for applications with bursty traffic patterns. High utilization is important for throughput sensitive applications.

3. BANDWIDTH ALLOCATION TECHNIQUES

3.1 Netshare

Netshareproposes a hierarchical max-min bandwidth allocation. It divides a network into slices, such that each network slice has the delusion of network all to itself with guaranteed bandwidth. Bandwidth unused by a network slice can also be shared proportionately among active network slices. Netshare uses weighted fair queuing of links to fair sharing of networks.
Netshare slices network bandwidth in 2 steps:

Step1: Ignores customer slices, and only considers the network as being shared between incoming and outgoing router. Step1 reduces the network to set of virtual links between every pair of ingress and egress points.

Step2: Divides the bandwidth of each virtual link between customer's slices.

Step1&2 combine to slice network bandwidth between customers. Customers have guaranteed minimum bandwidth as well as a probable advantage from bandwidth unused by others

Netshare algorithm [1] is being implemented by extending the existing link state routing protocol. It provides a higher-level abstraction that can provide both minimum guarantees and statistical multiplexing. Netshare can also handle adversarial sources and requires no state in the core. Netshare protocol runs well even on large networks and used to provide bandwidth and throughput guarantee. NetSharesupports network sharing through the use of per-tenant users that aresteady throughout the network. The main drawback of NetShare allocation scheme is it provides resilience in the face of uncooperative sources and statistical multiplexing. This model can be used to implement a form of link proportionality.

3.2 Oktopus

Oktopus implements the virtual network abstractions that allow tenants to expose their network requirements. Oktopus [2] uses mapping technique to map tenant virtual networks to the physical network. The provider maintains a datacenter containing physical machines with slots where tenant VMs can be placed. With Oktopus, tenants requesting VMs can opt for a virtual cluster to connect their VMs. It also supports tenants who do not want a virtual network, and are satisfied with the status quo where they simply get some share of network resources. A logically centralized network manager upon receiving a tenant request performs admission control and maps the request to physical machines. NM also accounts for network resources and maintain bandwidth reservations across the physical network. Oktopus scales to large datacenter.

3.3 SecondNet

SecondNetvirtualizes [7] datacenter network for resource allocation for multiple tenants in the cloud. The resources allocated to VDC's can be rapidly adjusted to the needs of the tenants.SecondNet introduces a centralized VDC allocation algorithm [3] that provides deterministic bandwidth guarantees for traffic between each VM pair. SecondNet proposes static reservations throughout the network to implement bandwidth guarantees. The advantage of this reservation system is that they can achieve more complex virtual topologies regardless of physical location of the VMs. SecondNet achieves scalability by distributing all the virtual-to physical mapping, routing, and bandwidth reservation state in server hypervisors. The main drawback of the reservation system is that they do not achieve high utilization of datacenter networks, since the unused bandwidth is not shared between the tenants. From the tenant perspective the above protocol is not ideal one because; tenants do not understand their applications communication patterns to specify their bandwidth requirements between each pair of VMs. As well the communication pattern is dynamic and the amount of data exchanged between any pair of VM will vary significantly over time.

3.4 Seawall

Seawall describes a mechanism that allocates bandwidth on every link of a datacenter network by dividing network capacity based on an administrator specified policy. Seawall's goal is to partition the bandwidth in each congested network link according to weights associated with each VMs sending traffic through that link. Seawall [4] implements the allocation of bandwidth by channeling through congestion controlled, point to multipoint, edge to edge tunnels. The resultant of allocations remainsconstant regardless of number of flows, protocols or destinations in application traffic mix. Seawall scales to large numbers of tenants by reducing the network sharing problem to an instance of distributed congestion control. Efficiency is achieved in seawall by proportionally reallocating unused shares to active sources.The design space that the seawall occupies appears well-suited to emerging hardware trends in data center and virtualization hardware.

3.5 Gatekeeper

A system designed to meet the practical needs of both cloud users and cloud data center providers. Gatekeeper uses the hypervisor-based mechanism. Hypervisor is a software, firmware or hardware that creates and run the VM. Gatekeeper system provides network isolation for multi-tenant datacenters using distributed mechanisms implemented at the virtualization layer and each datacenter server. Gatekeeper [5] controls the usage of each server’s network access link. It provides per-virtual Network interface (vNIC) link bandwidth guarantees in both directions of the network link at each physical server, i.e. for both I/O traffic. Minimum bandwidth guarantees are achieved using an admission control mechanism that limits the sum of guarantees to the available physical link bandwidth. Each vNIC can surpass its guaranteed allocation when extra bandwidth is available at both transmitting and receiving end points. To provide deterministic behavior Gatekeeper limits each vNIC bandwidth to a maximum rate thereby assuring the minimum guarantee. Gatekeeper achieves scalability using a simple point-to-point protocol and minimal datacenter-wide control state. Gatekeeper can take advantage of unused bandwidth both at the transmit and receive side up to maximum rate specified by the system administrator of each vNIC.

3.6 Falloc

To satisfy the needs of the tenants in various network applications, fairness must be considered in sharing the limited bandwidth resource in data centers. The vital role of fairness in bandwidth allocation is to protect each application’s performance in the competition of various network intensive applications. The two basic requirements for data center networks in designing a bandwidth allocation policy are minimum bandwidth guarantee and proportional bandwidth share. Fair Network Bandwidth Allocation (Falloc)[6]is used to allocate the network bandwidth in an effective way. It is an application layerprotocol used to attain VM-based fairness across the datacenter. The two main objectives: i) guarantee bandwidth for VMs based on their base bandwidth requirements ii) share residual bandwidth in proportion to weights of VMs. To achieve this objective the first step is to model the bandwidth allocation process in datacenters as asymmetric weighted Nash bargaining game [10][12], where all VMs are supportive without effecting the other VMs benefit.In falloc, each VM is assigned with base bandwidth and a weight. The protocol can guarantee the bandwidth of a VM when its bandwidth requirements are less than the base bandwidth and share the residual bandwidth among VMs in proportion to their weights. Falloc uses bargaining game approach[13], to allocate the bandwidth in a fair manner and the result achieves a weighted Nash bargaining solution.

4. ANALYSIS BASED ON NETWORK PERFORMANCE METRICS

After investigating the different bandwidth allocation technique, the comparative study was done based on the network performance metrics against the bandwidth allocation protocols. The Table 1 summarizes the comparison offew bandwidth allocation protocols for datacenter networks in cloud against the key properties which measures the network performance analysis.

Netshare is the only mechanism that allocates the relative among tenants based on their weights instead of sender VMs and provides constant proportionality throughout the network. It relies on a centralized bandwidth allocator which is difficult to scale to large datacenters and to deal with workload changes and the high rate of tenant and VM churn of cloud datacenters.

Table 1Comparison of bandwidth allocation protocol against key properties

Protocol	Key Properties			
	Bandwidth Guarantee	**Scalability**	**Flexibility**	**Network Utilization**
NetShare	Based on the weights of source VMs	×	√	high
Oktopus	deterministic	√	×	low
SecondNet	deterministic	√	×	low
Seawall	Based on the weights of source VMs	√	√	high
Gatekeeper	Minimum	√	√	high
Falloc	Based on the base bandwidth	√	√	high

Oktopus and SecondNet focuses on providing static reservations throughout the network to implement bandwidth guarantees for the VMs. They allocate VMs into servers based on VMs bandwidth requirements, and by enforcing reservations in both hosting servers and switches; they can ensure the bandwidth of inter-VM network and achieve predictable network performance for the applications in these VMs. The main drawback of reservation policies is that they may not be able to achieve high utilization of datacenter networks, since the unused bandwidth is not shared between tenants. The advantage of the reservation policy is that they can achieve more complex virtual topologies irrespective of the physical locations of the VMs.

Seawall provides a hypervisor-based mechanism to partition the bandwidth in each congested network link according to weights associated with each VMs sending traffic through that link according to weights associated with each VMs sending traffic through that link. Experimental results as specified in [4] revealed the two advantage of this protocol, (i) Seawall is scalable.(ii) The use of feedback from receivers allows the traffic to be choked at the sources before they use network resources, and prevents a malicious VM to take over bandwidth in the network. The main drawback is that it does not satisfy the predictable service level requirement. While Seawall can provide minimum guarantees if the maximum weight connected with each link is restricted to a maximum value, it cannot enforce maximum rates to support deterministic behavior. Also Seawall’s bandwidth allocation divides the link bandwidth among tenants using the total number of VMs sending traffic through that link.

Gatekeeper useshypervisor-based mechanisms, which provides minimum bandwidth guarantee for VMs by determining the traffic of VMs, but the unused bandwidth is shared in a fair manner. Gatekeeper works well in easy situations.. Still the research work is going on for the behavior of the Gatekeeper for large configuration and dynamic workloads.

Falloc, bandwidth allocation protocol, fairly share network resources at VM-level in datacenters. It assures the bandwidth requirement based on the base bandwidth for each VM, and shares the outstanding available bandwidth in a proportional way according to VM’s weight. Falloc provides flexible fairness for VM’s by balancing the transaction between bandwidth guarantee and proportional bandwidth share. Falloc accomplishes high network utilization and job completion time is good in datacenter network and adapt to dynamic traffic.

5. CONCLUSION

In summary, we discussed various bandwidth allocation protocol and key properties for the network performance. Each of these protocols finds use in different applications. The outcome of the comparative study of these protocols against the key properties reveals that Falloc is the best protocol among the protocols we have considered. Falloc provides bandwidth guarantee and share the residual bandwidth in proportion and achieves high utilization and maintains fairness among VMs in datacenters. Falloc shows no advantage on providing fairness for different jobs with different bandwidth.

6. REFERENCES

[1] T. Lam and G. Varghese, “Netshare: Virtualizing bandwidth within the cloud,” UCSD, Tech. Rep., 2009.

[2] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, “Towards predictable datacenter networks,” in ACM SIGCOMM, 2011.

[3] C. Guo, G. Lu, H. Wang, S. Yang, C. Kong, P. Sun, W. Wu, and Y. Zhang, “Secondnet: A data center network virtualization architecture with bandwidth guarantees,” in ACM CoNEXT, 2010.

[4] A. Shieh, S. Kandula, A. Greenberg, C. Kim, and B. Saha, “Sharing the data center network,” in USENIX NSDI, 2011.

[5] H. Rodrigues, J. Santos, Y. Turner, P. Soares, and D. Guedes, “Gatekeeper: Supporting bandwidth guarantees for multi-tenant datacenter networks,” in USENIX Workshop on I/O Virtualization, 2011.

[6] JianGuo, F Liu, H Tang, Y Lian, Hai Jin and John C.S. Lui, “Falloc:Fair Network Bandwidth Allocation in IaaS Datacenters Via a Bargaining Game Approach” in IEEE, 2013

[7] J. Zhu, D. Li, J. Wu, H. Liu, Y. Zhang, and J. Zhang, “Towards bandwidth guarantee in multi-tenancy cloud computing networks,” in IEEE ICNP, 2012.

[8] L. Popa, G. Kumar, M. Chowdhury, A. Krishnamurthy, S. Ratnasamy, and I. Stoica, “Faircloud: Sharing the network in cloud computing,” in ACM SIGCOMM, 2012.

[9] J. Guo, F. Liu, D. Zeng, J. C. Lui, and H. Jin, “A cooperative game based allocation for sharing data center networks,” in IEEE INFOCOM2013.

[10] J. Guo, F. Liu, H. Tang, Y. Lian, H. Jin, and L. John C.S., “Falloc: Fair network bandwidth allocation in iaas datacenters via a bargaining game approach,” Technical report, HUST, http://grid.hust.edu.cn/fmliu/ icnp2013falloc.pdf, August 2013.

[11] H. Boche, M. Schubert, N. Vucic, and S. Naik, “Non-symmetric nash bargaining solution for resource allocation in wireless networks and connection to interference calculus,” in Proc. 15th European Signal Processing Conference, 2007.

[12] N. Nisan, Algorithmic game theory. Cambridge UnivPr, 2007.

H. Yaiche, R. Mazumdar, and C. Rosenberg, “A game theoretic frame-work for bandwidth allocation and pricing in broadband networks,” IEEE/ACM Transactions on Networking (TON), vol. 8, no. 5, 2000.

[13] Z. Fang and B. Bensaou, “Fair bandwidth sharing algorithms based on game theory frameworks for wireless ad-hoc networks,” in IEEE INFOCOM, 2004.

[14] www.cse.iitd.ac.in/~rahul/cs905/lecture15/index.html