Distributed Reallocation Scheme for Virtual Network Resources

Clarissa Marquezan, Jéferson Nobre, Lisandro Granville
Institute of Informatics - UFRGS
Porto Alegre, RS, Brazil
Email: {clarissa, jcnobre, granville}@inf.ufrgs.br

Giorgio Nunzi, Dominique Dudkowski, Marcus Brunner
NEC Europe Network Laboratories
Heidelberg, Germany
Email: {nunzi,dudkowski,brunner}@nw.neclab.eu

***Abstract*—Network virtualization is an emerging technology for cost-effective sharing of network resources. The key strategy in network virtualization is of slicing physical resources (links, CPU, memory, and storage) to create virtual networks that are assigned to different operators. One important challenge on network virtualization is the efficient use of the physical resources. To accomplish such efficient use the management of the physical resources should be transparent to the applications running within the virtual networks, and should be executed at runtime in order to deal with the variation on the load requests of different virtual networks. Traditional resource allocation schemes use offline, centralized, and global view strategies to manage the use of physical resources. In contrast to these strategies, we propose a runtime, distributed, local view approach to manage physical resources. In this paper we introduce a virtual network architecture and an associated self-organizing algorithm to reallocate virtual network resources along different physical nodes in order to equalize the bandwidth, and storage consumption on the physical nodes. We developed a virtual network model based on Omnet++ to simulate the designed self-organizing algorithm. An IPTV testbed scenario is presented and initial experiments, about the interruption time of the application inside the IPTV virtual network, are described.**

I. Introduction

The increasing demand of multimedia services over the Internet is pushing for new methods to allocate resources in future networks. For example, IPTV services are expected to become more and more popular and integrated offers, like the triple-pay packages, require cost-effective strategies for resource allocation. In fact, a typical IPTV network infrastructure requires significant investments for the distribution network, in terms of guaranteed bandwidth as well as available storage capacity. Normally, these resources need to be planned and well dimensioned in advance, before upper services can be actually deployed [1].

The costs of deploying a physical infrastructure may prevent many service providers to get into the market, like in the case of IPTV services [2]. Nevertheless, recent works in the field of virtual networks offer a viable alternative that promises to cut costs by sharing the infrastructure among different service providers [3]. The key on network virtualization is of dividing the physical network infrastructure into several slices and associating them to different virtual providers. The deployment of virtual networks must observe two different perspectives. The former is the perspective of a virtual provider, who wants the accomplishment of the contracted resources (SLAs must be maintained), while the later regards to the physical infrastructure provider, who wants to save as much as possible its physical resources in order to maximize revenues.

Being this, efficient algorithms to allocate physical resources (links, CPU and storage capacity) must be put in place by physical providers, otherwise punctual high loads on multiplexed physical resources may create resource scarcity that can prevent the deployment of new virtual networks. Traditionally, physical resources are allocated in the initial planning phase: a planning tool [4] provides the estimated dimensioning of network components given a certain SLA and resources are allocated based on this output. This approach can be applied for small virtual environments, but in large scale deployments a static allocation cannot take in account the mass imbalance of users requests between different locations.

In order to efficiently consume resources of the physical infrastructure, this paper proposes a real-time reallocation of virtual network resources. The main contribution of this paper is twofold. First, we propose a new approach to the deployment of services of virtual networks: with this approach, resources can be dynamically moved within the virtual layer to maximize over time the consumption of physical resources. Second, we define a distributed algorithm based on self-organizing techniques to implement a real-time reallocation scheme for virtual networks.

The proposed virtual model was implemented in the Omnet++ simulator and we defined an IPTV scenario with virtual providers in order to test the self-organizing reallocation scheme. The objective of the simulation, in this paper, is identifying the impact of the moving process in terms of interruption of IPTV services.

The reminder of this paper is organized as follows. Section 2 brings the related work on self-organization and virtual networks. Section 3 presents the proposed virtual network model and section 4 presents the designed reallocating scheme. Section 5 presents the implementation of the proposed model using the Omnet ++ simulator, while section 6 describes the testbed scenario. The evaluation and the associated results are discussed in Section 7. Finally, the conclusions and future work are presented in Section 8.

II. Related Work

On a first sight, the proposed self-organization reallocation scheme might be seen as an extension of existing virtual machine live migration. Recently, self-organization techniques have been employed on server virtualization scenarios [5][6]. In these cases, the virtual machines are self-organized according to the workloads of the physical nodes, and generally, this self-organization is accomplished migrating virtual machines to physical ones with lower workloads. However, the metrics traditionally used to determine the workload of virtual machines are CPU and memory, and in a virtual network the bandwidth consumption is one major metric to be considered in the migration process. Beyond these metrics, virtual network live migration is different from virtual machine migration because it has also to deal with virtual topology issues and routing connections reconfigurations.

The research presented by Yuichi Ohsita et al. [7] is able to make the reconfiguration of the virtual network topology in order to cope with the traffic demands on a given moment. The authors use traffic matrix estimation, and partial vision of the virtual nodes to make the reconfiguration decisions. A sub set of the authors from this first paper, Takashi Miyamura et al. [8], enhanced the previous research and defined a centralized server devoted to identify traffic on demand fluctuation and network failures. Based on this, a virtual network reconfiguration is activated. Both cases, the re-configuration process is just restrict to links of a virtual network and does not consider that this process might involve migration of an entire virtual device, like a router.

A very recent research on virtual router migration is presented by Yi Wang et al. [9]. In this paper the authors proposed VROOM, a virtual router migration mechanism, where the virtual interfaces of the routers are not directly mapped to physical ports and in this sense it is possible to migrate a router among different physical devices. The authors presented the migration mechanism itself and argued about the advantages of using this approach to deal with management changes, planning, and new service deployment. However, nothing was mentioned about the analysis to trigger the router migration, and how this approach can help to reduce punctual high loads on the physical infrastructure.

Based on the aforementioned we believe that current researches do not address the problem of reallocate virtual resources at runtime, using local view, and based on a distributed approach. The next sections present the proposed solution.

III. Virtual Network Model

According to recent researches, virtualization is a promising technique to deploy future networks [3][10][11]. Its key idea is the identification and separation of two roles: a physical provider, who owns and maintains the physical network, and a virtual provider, who builds its own infrastructure by renting slices of resources from the physical provider. If we look at a virtual provider as an entity selling services to final customers, the advantage of virtualization relies on the fact that costs in running a physical infrastructure can be outsourced to an external provider.

For the purpose of this paper, it is important to describe the main characteristics of an architecture for virtualization; it should be noted that this paper presents the minimal assumptions, and further details can be found in specific projects like GENI or 4WARD [12]. The physical resources of a node are sliced into different virtual nodes: each virtual node is assigned to a different customer. Physical resources include CPU power, memory space, eventually storage capacity, network interfaces and bandwidth.

An important aspect in the architecture of virtual networks is the transparency: virtual nodes cannot see or exchange any type of information, in order to assure isolation of the networks of different providers. Additionally, the data exchanged in the virtual network is transparent to the physical provider to preserve the privacy of the customers. Nevertheless, some minimal primitives to inspect the activity of the different slices are normally available: as an example, primitives to allow the controller of the physical resources to know the actual usage of computational resources and traffic consumption. Figure 1 shows the architectural view of a node, where resources are sliced and assigned to different virtual providers.

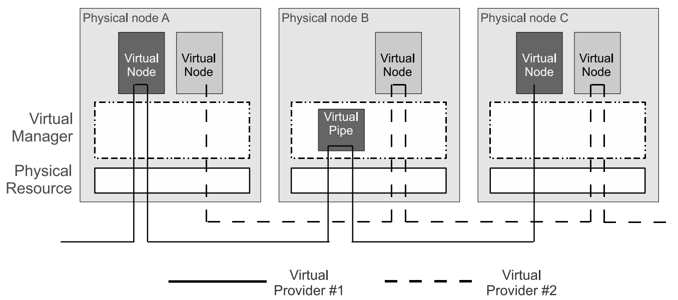

Fig. 1. Virtual node

A physical node is composed of: physical resources, virtual manager, virtual nodes, and virtual pipes. The virtual manager is responsible for receiving the requests for deployment of a virtual node or pipe and managing locally the connections and resources associated to virtual network. The virtual manager can be seen as the “hypervisor” concept used on virtual machines technologies, for instance.

A virtual node is a slice of the physical node comprehending: CPU power, memory space, application(s), storage capacity (if necessary), network interface and bandwidth multiplexing. A virtual pipe regards all virtual node features but application and storage capacity. The introduction of the virtual pipe concept supports the creation of virtual links between non-adjacent virtual nodes. Figure 2 illustrate the differences of using virtual nodes and virtual pipes considering the physical view and the virtual views.

The technology for creating virtual links is already available on current routers [9], but we believe that it is necessary to have some mechanism to determine the amount of traffic passing by physical connections that compose the virtual link,

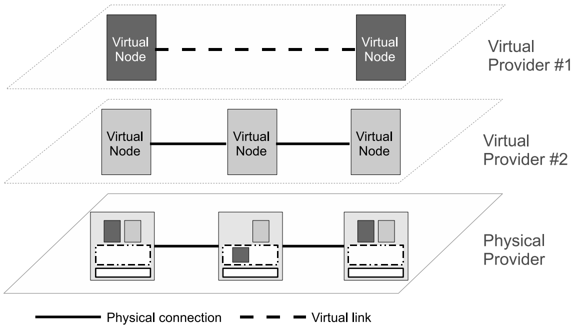

Fig. 2. Virtual link representation

in order to enable a better management of the virtual networks resources. For example, the employment of virtual pipes allows the virtual manager to identify forward traffic inside a physical node without inspecting the packets belonging to a virtual network associated to this traffic. In our model, the information associated to this kind traffic is one of the inputs used to analyze the necessity of reallocating virtual resources. In the next section we present our solution for efficient resources reallocation.

IV. Distributed Reallocation Scheme

As mentioned before the major contributions of our proposal are the employment of distribution, local view, and online features on the reallocation of resources of virtual networks. Some assumptions must be observed in order to provide such features in the new scheme, and they are described below.

- The initial deployment of a virtual network is not addressed by this work, and we assume that a different, external planning tool analyzes the conditions of physical resources and then choose the best initial placement for the slices of the new virtual network.
- We assume that the virtual topology defined by the first placement will not change during the lifetime of the virtual network, even after the reallocation of virtual slices among physical nodes.
- The reallocation of slices must be as transparent as possible for the virtual node. In the current stage of this research, the reallocation of the virtual slices is transparent in the sense of avoiding to exchange any kind of information between the virtual application inside the moving slice and the virtual managers of the physical nodes involved in the reallocation operation. However, we introduce an interruption time on the execution of the application running inside the moving virtual slice.

Based on these assumptions, and inspired on self-organization techniques presented in [13] we defined a reallocation scheme that is executed locally by each virtual manager inside the physical nodes. The main objective of this mechanism is to approximate the virtual node that is generating a great amount of traffic to the destination virtual node. The approximation is done moving the source virtual node from its physical device to another physical device near the destination virtual node. Figure 3 illustrates the reallocation of a virtual router of an IPTV virtual provider (details about the IPTV infrastructure can be found in [2]).

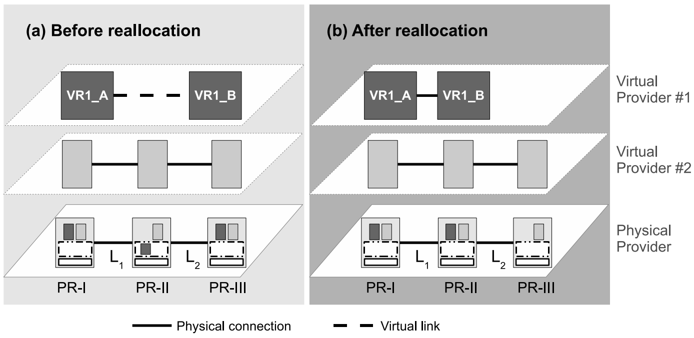

Fig. 3. Reallocation scheme

The algorithm used to accomplish the reallocation scheme is divided in five stages. First, locally each virtual manager analyzes the existence of some traffic associated to a virtual node with characteristics to be moved. Due to space limitations the heuristic developed to identify the virtual nodes to be moved will not be presented. On the second stage, the physical neighbors exchange information about the virtual nodes that must be received or moved. Locally, on the third stage, each neighbor analyzes the exchanged information, and the physical node that must move a virtual resource decides to whom the virtual resources might be moved. The forth stage is the announcement of the decision and the reservation of the resources at the target physical neighbor. Finally, the virtual resources are moved.

During the third and fifth stages the application(s) running inside the virtual node are suspended, and all packets related to this virtual node are queued by the virtual manager. As soon as the virtual node is reestablished on the physical neighbor the packets are unqueued and sent to the virtual node on the new physical location. As aforementioned, the proposed reallocation scheme imposes an interruption time on the application running on the virtual node. This interruption time depends on the nature of virtual node that is being moved. For example, if the virtual node is an IPTV router, the interruption time might be higher because the storage associated to the IPTV router must be also moved. On the other hand, if the virtual node is a common router the interruption time should not be prohibitive because less resources should be moved. Discussions about the routing process during the migration of virtual routers are out of the scope of this paper.

According to the description provided above, it is possible to observe that our proposal does not need a global view of the physical topology to identify the overloaded physical resources, like links or devices. Just using the local information retrieved from the controllers of network interfaces, CPU, memory, and disk, our heuristic is able to identify possible virtual candidates to be moved. Moreover, we also do not need any centralized entity to make the decision of reallocating resources. Our approach is completely distributed and based on information exchanged among the physical neighbors the reallocating scheme is triggered. On the next sections we present the implementation, testbed, and evaluation of the proposed reallocation scheme.

V. Implementation

To validate our reallocation scheme we implemented a new module for Omnet++ Simulator. This new module is presented in Figure 4. The network presented in this figure is composed of 5 physical nodes and 2 virtual IPTV providers ("vnetA" and "vnetB"). Most of the parameters of this virtual module are configurable, like for example the number of physical devices, virtual nodes, and pipes, and also the features of these elements. However, the current version of this module does not support the definition of different network topologies, and only ring network topology can be described in this version.

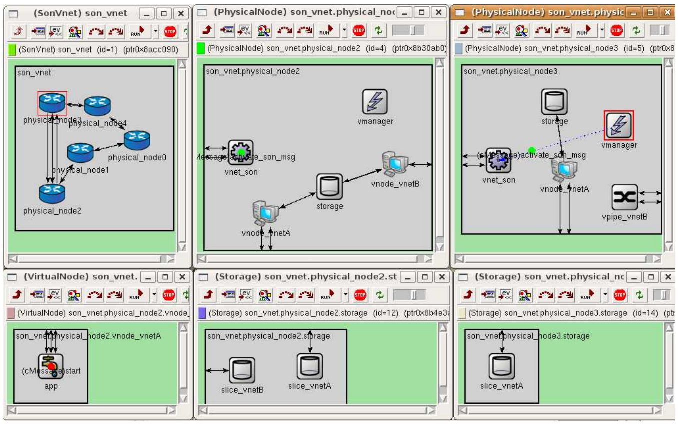

Fig. 4. Virtual module for Omnet++ simulator

So far, in this paper we show the execution of the proposed distributed virtual reallocation scheme in the light of the reallocation of virtual resources from an IPTV virtual provider. To accomplish this we defined a virtual network where an IPTV provider deploys the required infrastructure to attend the requests for movie streams of their costumers, and the associated testbed is described in the sequence.

VI. Testbed

We consider a scenario where the IPTV provider requires routers connecting costumers, and the planning tool, responsible for defining the first placement of the virtual resources, has allocated two virtual routers in different physical routers. Furthermore, storage slices have been attached to each virtual router. These two virtual nodes are connected through a virtual link, and this scenario is illustrated in Figure 3.

The experiments consider users connected to the virtual router "VR1_A" requesting movies located on the storage connected to virtual router "VR1_B". There is a traffic passing by the physical links "L_1" and "L_2", and this traffic transforms the virtual node "VR1_B" (at physical device "PR-III") in a candidate to be moved to the physical device "PR-II".

To run the simulation some main parameters are required and Table I presents these parameters and the associated values used on the simulation.

Parameter	*Associated value*
Datarate of links associated to each virtual network	1 Gbps
Delay of of links associated to each virtual network	1 ms
Datarate of storage	100 Mbps
Delay of storage	1 ms

TABLE I
Parameters of the simulation

Being this, the initial experiments proposed in this paper investigate the interruption time to move the virtual router "VR1_B" and the storage directly connected to it, when the size of the storage varies from 1GB to 10GB. We also discuss the compromise between maintaining a low interruption time and the number of virtual resources necessary to keep the fixed interruption time.

VII. Evaluation

Figure 5 presents the graphic of the interruption time associated to the scenario described above. The interruption time is composed of time to: (a) exchange the control messages between the physical neighbors in order to reserve the required resources for the reallocation process; (b) read and send data from one storage to the storage of the physical neighbor; and (c) write data on the storage of the new physical location of virtual router "VR1_B". The interruption time increases linearly with the increase of the storage size, as expected.

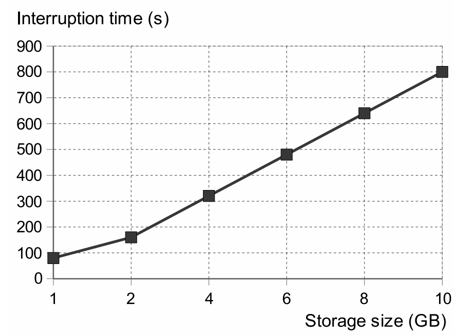

Fig. 5. Interruption time

The analysis of the interruption time is not so interesting when it is done in an isolated fashion. However, the analysis of the interruption time in the light of the amount of virtual resources used by a virtual provider can become a business metric at the moment that a virtual provider is contracting a virtual network. For example, let's consider the scenario where an IPTV provider requires a storage capacity of 20 GB, and the features described in Table I are being used. In this case, the smallest interruption time (i.e, 80s considering that the minimum storage slice is 1GB) is guaranteed when 20 virtual nodes are used.

In this sense, the maintenance of low interruption time during the reallocation scheme imposes the deployment of more virtual resources to a single virtual provider. This information can be used by both sides, physical and virtual providers, to determine the behavior of the reallocation process. For instance, if the virtual provider contracting a virtual network does not desire high interruption time on the applications running inside the virtual network, it can force the physical provider not to employ the reallocation scheme to this virtual network. However, the physical provider can increase the prices for virtual providers that want more fixed constrains on the maintenance of the virtual network operation. We believe that this tradeoff is a metric to be agreed on the SLA between the physical and the virtual providers before the deployment of the virtual network.

VIII. Conclusion

This paper presented the definition of a distributed reallocation scheme for virtual network resources, a high level architecture for virtual networks, and first experiments using the reallocation scheme on IPTV scenario. The main objectives of the experiments were presenting the correct execution of the reallocation mechanism, identifying the interruption time of the applications (inside the virtual nodes) imposed by the resource reallocation process, and analyzing the relationship between interruption time and virtual resources composing the virtual network.

The major outcome of this initial experiment is the utilization of the interruption time and number of virtual resources in order to determine the terms of the SLA between the physical and virtual providers. If application outages are not a constrain for the virtual provider, it is possible to firm an SLA giving more flexibility to the physical provider reallocates the virtual resources, and this flexibility might be translated into a reduction on the price of the virtual network deployment. The major benefit in this case stays with the physical provider, that can reallocate the virtual resources in order to efficiently use the physical resources. On the other hand, virtual provides that require restrict reallocation policies and low interruption time, would not allow the employment of the reallocation scheme, as a consequence the costs for the virtual provider might be increased.

As future work we intend to extend the experiments using the IPTV virtual networks, and employ a user request model to verify the full operation of the reallocation scheme. The next evaluation scenario aims to identify the costs of moving virtual resources, considering the relationship between interruption time and the saved bandwidth on the physical links after the execution of the reallocation scheme. We also intend to test the behavior of other kind of applications on top of the virtual model, for instance, network management applications. Furthermore, we intend to investigate how self-* features, like self-healing, self-configuration, self-awareness, self-monitoring, can improve the management of the virtual networks on top of the physical network.

Acknowledgement

This work was partly funded by the Brazilian Ministry of Education (MEC/CAPES, PDEE Program, process 4436075), and by the European Union through the 4WARD project in the 7th Framework Programme. The views expressed in this paper are solely those of the authors and do not necessarily represent the views of their employers, the 4WARD project, or the Commission of the European Union.

References

[1] N. Degrande, K. Laevens, D. D. Vleeschauwer, and R. Sharpe, "Increasing the User Perceived Quality for IPTV Services," *IEEE Communications Magazine*, vol. 46, no. 2, pp. 94–99, 2008.

[2] S. Han, S. Lisle, and G. Nehib, "IPTV Transport Architecture Alternatives and Economic Considerations," *IEEE Communications Magazine*, vol. 46, no. 2, pp. 70–77, 2008.

[3] N. Niebert, I. E. Khayat, S. Baucke, R. Keller, R. Rembarz, and J. Sachs, "Network Virtualization: A Viable Path Towards the Future Internet," *Journal Wireless Personal Communications*, vol. 45, no. 4, pp. 511–520, June 2008.

[4] D. Agrawal, M. S. Beigi, C. Bisdikian, and N. Lee, "Planning and Managing the IPTV Service Deployment," in *Proceeding of 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, 2007, pp. 353 – 362.

[5] M. Steinder, I. Whalley, D. Carrera, I. Gaweda, and D. Chess, "Server virtualization in autonomic management of heterogeneous workloads," in *Proceeding of 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)*, May 2007, pp. 139–148.

[6] X. Y. Wang, D. Lan, X. Fang, M. Ye, and Y. Chen, "A resource management framework for multi-tier service delivery in autonomic virtualized environments," in *IEEE/IFIP Network Operations and Management Symposium (NOMS 2008)*, 2008, CDROM.

[7] Y. Ohsita, T. Miyamura, S. Arakawa, S. Ata, E. Oki, K. Shiomoto, and M. Murata, "Gradually Reconfiguring Virtual Network Topologies Based on Estimated Traffic Matrices," in *Proceedings of 26th IEEE International Conference on Computer Communications. IEEE INFOCOM 2007*, May 2007, pp. 2511 – 2515.

[8] T. Miyamura, E. Oki, I. Inoue, and K. Shiomoto, "Enhancing bandwidth on demand service based on virtual network topology control," in *IEEE Network Operations and Management Symposium Workshops, 2008. NOMS Workshops 2008.*, April 2008, pp. 201–206.

[9] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, "Virtual routers on the move: Live router migration as a network-management primitive," in *ACM SIGCOMM Computer Communications Review*, August 2008, pp. 231–242.

[10] N. Feamster, L. Gao, and J. Rexford, "How to lease the internet in your spare time," in *ACM SIGCOMM Computer Communications Review*, Jan 2007.

[11] A. Berl, A. Fischer, H. de Meer, A. Galis, and J. Rubio-Loyola, "Management of Virtual Networks," in *Proceedings of Fourth IEEE/IFIP International Workshop on End-to-End Virtualization and Grid Management, EVGM 2008*, September 2008, pp. 197 – 202.

[12] 4WARD, "The fp7 4ward project," 2008, available at http://www.4ward-project.eu/.

[13] K. Hermann, "Self-organizing replica placement - a case study on emergence," in *IEEE First International Conference on Self-Adaptative and Self-Organizing Systems (SASO 2007)*, 2007.