

# RETHINKING ISP ARCHITECTURES THROUGH VIRTUALIZATION

Alejandro Cordero and David Fernández

*Departamento de Ingeniería de Sistemas Telemáticos, Universidad Politécnica de Madrid, Madrid, Spain*

**Keywords:** Network virtualization, ISP architectures, Future internet.

**Abstract:** Network virtualization is a powerful technology that has been proposed as a mean to evolve the current Internet, allowing the introduction and testing of future network technologies over present infrastructures in coexistence with our current production networks. However, apart from a path to the Future Internet, network virtualization can also be a key technology that modifies and improves today ISP networks. This paper analyses how network virtualization and all the technologies being developed around it can influence and evolve the present ISP network architectures and business models. Starting from the well-known ISP architecture made of access, distribution and core layer; we present and discuss the benefits that can be achieved by introducing virtualization technologies in each layer.

## 1 INTRODUCTION

Virtualization is a mature technology in the context of computing resources, mainly in what is related to operating systems (platform virtualization). Many data centres nowadays are supported by virtualization suites that allow high flexibility in sharing physical machines and provide capabilities like high performance, availability, energy efficiency, load balancing, etc.

However, in the context of networking, virtualization is a recent technology. Network virtualization consists of creating multiple logical networks over a common physical infrastructure made of nodes and links. Physical nodes are shared using platform virtualization techniques adapted to communication nodes, taking into account their strict real-time requirements, as well as the reliability (carrier-class).

Network virtualization has been proposed as a mean to allow the coexistence over the same network infrastructure of the current Internet and the new network architectures and protocols being proposed for the Future Internet. In fact, several research programs and projects are working in the application of network virtualization to Future Internet (e.g. GENI in USA, FIRE in EU or AKARI in Japan). But it has also a clear role to play in the evolution of current Internet architectures, creating

new roles and business models. All the new technologies developed around virtualization, like router virtualization or dynamic reconfiguration of networks can be applied to current network architectures, giving rise to more flexible and efficient networks.

This paper analyses how network virtualization and related technologies can influence and evolve the present ISP network architectures and business models. The paper is organized as follows. Starting from the well-known ISP architecture (Section 2) made of access (2.1), distribution (2.2) and core (2.3) layers; we present and discuss the benefits that can be achieved by introducing these virtualization technologies in each layer. Finally, we provide some concluding remarks in Section 3.

## 2 VIRTUALIZING ISP ARCHITECTURE

In order to organize our discussion about how network virtualization can be applied to today ISP networks, in this section we will briefly describe the three-layer network hierarchical model (Oppenheimer, 2004), a highly adopted “de facto” standard for ISP (and Enterprise) network design topologies. This model permits traffic aggregation and filtering at three successive routing or switching

levels and makes it very scalable up to large international internetworks levels.

A typical hierarchical network topology consists in:

- A core layer of high-end routers and switches that are optimized for availability and performance.
- A distribution layer of routers and switches that implement policies.
- An access layer that connects users via lower-end switches and wireless access points.

The next subsections provide some insights about how network virtualization can be applied to each of the layers described before.

## 2.1 Access Layer

The access layer provides users on local segments the access to the internetwork. The access layer can include a large variety of equipment, such as routers, switches, bridges, shared-media hubs, wireless access points, etc., using different technologies such as xDSL, FTTH, Ethernet, DOCSIS, UMTS/3G, Wi-Fi, WiMAX, etc, as well as different physical media such as copper (twisted pair and coaxial), fiber and air.

From the point of view of traditional ISP, we divide the access layer in three differentiated components for analyze their virtualization possibilities:

- The Customer Equipment.
- The Provider Equipment.

**Customer Equipment:** The virtualization of customer equipment like residential or home gateways is analyzed in (Royon, Frénot, 2007), (Ibanez et al., 2007). OSGi is a common solution to the CE implementation that permits a fast and easy deployment of new services and protocols. Virtualization of OSGi can be seen as a method for sharing CE among providers.

Another result of the virtualization of CE is decoupling service from devices that permits the creation of new solutions. i.e. a Virtual Home Environment (Berl et al., 2009) architecture created to move and consolidate services to ensure energy efficiency, or a service virtualizer (Häber et al., 2009) where remote devices are presented as if they were in local network.

However, we must not forget the hardware needed to achieve an optimal implementation of a CE with virtualization capabilities. Today residential gateways are focused to Internet access as well as its hardware. Future implementation of a CE must be more flexible to allow joining several services in one

device. This flexibility should permit upgrading to new protocols and implementing new services without changing equipment.

The key of success for this approach depends on the equipment cost versus the flexibility that it grants, but the possibility of sharing CE cost between providers improves the chance to win. Therefore, from a user point of view, virtualization can improve customer equipment allowing:

- Faster deployment of new services.
- Easier implementation of new protocols.
- Sharing CE between providers.
- Consolidating equipment to improve energy efficiency (green compliance).

**Provider Equipment:** Today's ISPs have link layer equipment such as DSLAM, CMTS, ONT or RNC, depending on the transmission technology used. This equipment could be virtualized to allow it to be shared among several VNP or to be aggregated into one to facilitate its management. In this case, it could be possible to easily migrate a user inside the same equipment between VNP or upgrade link layer protocols, but we always have the limitation imposed by physical end-user connection.

Currently, ISP's solve the problem of access layer in two ways:

- Deploying their own access network.
- Paying traffic aggregation of their users to its distribution layer.

Virtualization would allow the deployment of an intermediate model, where an ISP with its own network could become the infrastructure provider of other ISP's behaving as virtual network providers. This model has two direct consequences:

- The VNP has greater control over traffic at limited cost.
- The IP could share its resources and have a new business source. This IP focused on the virtualization of the access layer could be called network access operator (NAO).

Another consequence of virtualization is the faster creation of new VNP networks, where deployment times could be significantly reduced. Moreover, depending on the level of virtualization of the devices, a VNP could implement different versions of the protocols in their virtual instance, but this would imply more resources and added complexity.

However, access technologies are not prepared for virtualization. Today vendors only have done some advances in Ethernet switching. FlowVisor (Sherwood et al., 2009), is a special purpose OpenFlow controller that acts as a transparent proxy between OpenFlow switches and multiple

controllers. This solution may have application for today's MetroEthernet access networks, dividing it in a better way than traditional VLANs, but in order to have enough capabilities for VNP, future implementations need to have a separate control and configuration for each virtual instance.

Therefore, from the provider point of view, the use of virtualization has several benefits:

- More network control like local loop unbundling implementations.
- Easier implementation of new solutions.
- Sharing cost between providers.

## 2.2 Distribution Layer

The distribution layer of the network is the demarcation point between the access and core layers. The distribution layer controls the access to resources and network traffic for security and performance reasons, hiding detailed topology information about the access layer from core routers. Also, the distribution layer allows the core layer to connect access layers that run different protocols or technologies while maintaining high performance. In summary, its key needs are traffic optimization, security, and media transitions.

Network virtualization can improve distribution layer in the same way as shown in the access layer reducing costs by sharing equipment and providing a better way to deploy new solutions and protocols. But in the current implementations of the distribution layer in ISP's, mostly based in MPLS technologies, other improvements can be derived from the virtualization of MPLS PE (Figure 1).



Figure 1: Virtual PE at the distribution layer.

The idea of a virtual PE (vPE) is a logic step in the evolution of virtual private networks, where configurations and traffic are better isolated, increasing security levels. Moreover, virtual PE implementation gives rise to the possibility of router migration, a similar functionality to the virtual machines migration found in virtual computing platforms but applied to routers. In VROOM (Wang et al., 2008) router live migration is analyzed as a management primitive for planned maintenance and service deployment. Virtual PE consolidation (i.e. grouping several vPE's over one real node) could be

another benefit, minimizing power consumption through the hibernation of unused real equipment.

However, new developments are needed in order to deploy new capabilities like programmability for new protocol deployment or mobility primitives.

The migration of virtual routers capability opens the possibility to use dynamic reallocation techniques in the distribution layer. This possibility will be analyzed deeper in next section applied to core layer.

Therefore, network virtualization can improve distribution layer in the following way:

- Improving security through isolation of clients and processes.
- Migrating virtual routers to aggregate access layer in less physical nodes according to network traffic.
- Hibernating portions of distribution layer to optimize power consumption, and ensure green energy compliance.

## 2.3 Core Layer

The core layer is the high-speed backbone of the internetwork. The core should have a limited and consistent diameter, designed with redundant components, in order to be highly reliable and to adapt to changes quickly. Moreover, the core layer needs to be optimized for low latency and packet throughput.

The creation of a Composable Router with virtualization could be a good solution to implement scalable megarouters reducing costs. With this approach we could have high speed core routers by reusing common routers. The basic idea consists on separating control and data planes: in the data plane, composed of an array of routers, traffic is managed by and splitter that realises distribution and aggregation features; in the control plane, management is done by a meta-router that analyses control packets and creates and maintains forwarding tables.

Today ISP networks use redundancy in its design as a solution to failure, creating complex architectures to operate and support. Virtualization can change the way that today designs are made, dividing network design between IP and VNP. If redundancy is provided by IPs, deploying a physical redundant substrate with dynamic reallocation capabilities triggered by network failures or resource depletion, network design will be simplified for VNP.

Therefore, a well-designed IP network would allow simple fault tolerant VNPs networks,

improving scalability and optimizing resource use while minimizing cost. But, how can this be done? Different approaches can be implemented to solve the network-embedding problem, like restricting the problem space (sacrificing flexibility) or proposing heuristic algorithms (sacrificing efficiency). Some of these approaches are taken in (Zhu, Ammar, 2006) (Fan, Ammar, 2006) (Yu et al., 2008), but finding the optimal solution turns out to be a NP-hard problem.

However, today core architectures, composed by simple topologies with a limited number of devices, could permit the implementation of dynamic reallocation. This controlled approach can simplify the NP-hard problem, providing a way to create fully meshed substrate networks based on core networks topologies.

Finally, in core layer, as in all layers, the use of shared equipment among providers is an option. In this case, virtual cores can be created to connect isolated networks, having cheaper and self-managed high-speed transport.

Therefore, network virtualization can improve core layer in the following ways:

- Creating virtual routers that aggregate several real ones. This facilitates management having only one interface for all devices. Also offers scalability on demand, adding devices when more power is needed.
- Implementing virtual dynamic reallocation support fault tolerant networks. This simplifies VNP networks design over redundant IP substrate networks.

### 3 CONCLUSIONS

Virtualization is a mature technique in the computing environment that applied to networks can bring new solutions and business opportunities to the actors involved in today Internet. The first steps in this field have been already taken by research programs and industry vendors, creating an environment for future development. However, there is still much work to be done.

On the other hand, although ISP architectures are very consolidated, virtualization can improve their design providing new possibilities not seen before. In this paper we have presented an analysis about the benefits that network virtualization can bring to each of the layers of an ISP architecture. They can be summarized as follows:

- Consolidation of equipment to improve energy efficiency.

- Faster and easier implementation of solutions and protocols.
- Sharing cost between providers.

Finally, we believe that network virtualization techniques can be implemented in a near future, improving today's ISP networks. However, more insights into the possibilities outlined in this article are needed for a virtual networking environment to become a reality.

### REFERENCES

- Berl, A., Weidlich, R., Schrank, M., Hlavacs, H. & de Meer, H. 2009, "Network Virtualization in Future Home Environments", *Integrated Management of Systems, Services, Processes and People in IT*, pp. 177-190.
- Fan, J. & Ammar, M. 2006, "Dynamic topology configuration in service overlay networks: A study of reconfiguration policies", *Proc. IEEE INFOCOM*Citeseer.
- Häber, A., De Mier, J.G.R. & Reichert, F. 2009, "Virtualization of Remote Devices and Services in Residential Networks", 2009 *Third International Conference on Next Generation Mobile Applications, Services and Technologies*IEEE, , pp. 182.
- Ibanez, M., Madrid, N.M. & Seepold, R. 2007, "Virtualization of Residential Gateways", *Intelligent Solutions in Embedded Systems, 2007 Fifth Workshop on IEEE*, pp. 115.
- Oppenheimer, P. 2004, *Top-down network design*, Cisco Press.
- Royon, Y. & Frénot, S. 2007, "Multiservice home gateways: Business model, execution environment, management infrastructure", *IEEE Communications Magazine*, vol. 45, no. 10, pp. 122-128.
- Sherwood, R., Gibb, G., Yap, K.K., Appenzeller, G., Casado, M., McKeown, N. & Parulkar, G. 2009, "Flowvisor: A network virtualization layer".
- Wang, Y., Keller, E., Biskeborn, B., van der Merwe, J. & Rexford, J. 2008, "Virtual routers on the move: live router migration as a network-management primitive", *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 231-242.
- Yu, M., Yi, Y., Rexford, J. & Chiang, M. 2008, "Rethinking virtual network embedding: Substrate support for path splitting and migration", *Computer Communication Review*, vol. 38, no. 2, pp. 17.
- Zhu, Y. & Ammar, M. 2006, "Algorithms for assigning substrate network resources to virtual network components", *Proc. IEEE INFOCOM*Citeseer.