Universidad Carlos III de Madrid
e-Archivo

Institutional Repository

The Medianet Project: Integration of Multimedia Services for the next Generations Business oriented Internet

J. Garcia-Reinoso*, G. Ibañez†, J.L. Vazquez-Poletti‡ and I.M. Llorente‡
*Dpto. de Ingenieria Telematica
Universidad Carlos III de Madrid
Leganes, Spain
†line 1 (of Affiliation): dept. name of organization
Universidad de Alcala
Alcala de Henares, Spain
‡Dpto. de Arquitectura de Computadores y Automtica
Universidad Complutense de Madrid
Madrid, Spain

Abstract—This article gathers the foundational premises of the MEDIANET project as well as intermediate results obtained along its firsts two years. MEDIANET is a Spanish project founded by the Comunidad de Madrid government, which strives for a significant scientific advance in the future media Internet where important advances are necessary to allow end-users to perceive a good quality of experience. The network technologies objectives consist of the definition and validation of new proposals for the efficient transport of high bandwidth, real-time data flows in a decentralized way where the network provides mechanisms to seamlessly request and configure devices to increase the quality of experience perceived by end-users. Furthermore, new experiences with layer 2 networks and a cross-layer design will be tested with high bandwidth demanding media services. An important objective is to develop, evaluate in depth and implement on diverse platforms , a new low latency transparent bridge protocol based in on-demand path set up, suitable for campus and data center networks. The global result will be an integrated and independent advancement in future media Internet protocols, algorithms, switching architectures and standards.

Keywords-Multimedia;Business-oriented;Cloud computing

I. INTRODUCTION

In our context, media is the main mean of mass communication. No matter what kind of media we want to transmit (plain text, voice, video or a combination of all) we need a medium or carrier to let end-users access this information. Apart old methods, nowadays the preferred way to interchange this media is using the Internet infrastructure. Internet started like a set of experiments trying to test the viability of data packet switching techniques and, nobody at those days could imagine the high impact that Internet has nowadays. Internet applications rely on the old IPv4 not reliable protocol designed to route packets from one network to others but with a poor quality of service mechanism. It is true that new proposals like DiffServ and IntServ try to improve the quality provided by the network but these solutions are not available in all Internet yet. Real time applications, like audio and video, have big constraints in terms of delay or jitter imposing some restrictions on the underlying network. It is important to notice that this kind of traffic is totally different than web or file transfer traffic where the goal is to get all pieces no matter the jitter introduced by the network. Although there exists mechanisms to distinguish between these two different kinds of traffic, routers do not prioritize multimedia traffic against background traffic.

Recently, virtualization has brought about a new utility computing model called cloud computing, for the on-demand provision of virtualized resources as a service ("Infrastructure-as-a-Service"). This new paradigm has been introduced to better respond to dynamic computation demand by emerging applications and workload, such a web 2.0, and to better utilize computing resources and energy in aggregation to serve a large group of users. Our position is that cloud computing is perfect for future media internet applications, such as video transcoding. Those applications are CPU-intensive, and their tasks usually come in waves. So the Cloud paradigm gives enormous CPU resources that can be scaled up, and down, as needed, to meet the fluctuation demands. For example, video web sites publishing large quantities of content in-house have very difficult to scale their infrastructure. During slow traffic periods their encoding servers are underutilized, resulting in wasted resources. When traffic is heavy, however, the servers can become overwhelmed, decreasing quality-of-service and frustrating end users. Since Amazon EC2 started to popularize IaaS clouds, several solutions have been developed to create clouds. Some of these solutions, like Globus Nimbus and Eucalyptus, fall under the category of cloud toolkits, offering turn-key solutions to transforming existing infrastructure into an IaaS cloud. Eucalyptus is compatible with Amazon EC2 interface and is designed to support additional client-side interfaces. Globus Nimbus exposes EC2 and WSRF interfaces and offers self-configuring virtual cluster support.

These tools focus on the client perspective, being fully functional with respect to cloud compatible interfaces and providing higher-level functionality for security, contextualization and VM disk image management. However, they do not support dynamic allocation and load balancing of computing resources among virtual machines to meet the scalable and dynamic computing requirements of enterprise datacenters. Moreover, they do not support the federation of cloud infrastructures.

If we want to evolve to the future media Internet, some gaps have to be filled from bottom to top layers. Luckily, there already exists some initiatives trying to standardize the future Internet or the Next Generation Networks (NGN). Sadly, there are as many NGN definitions as standardization bodies working on this field. The important point is that all definitions have something in common: control plane and data planes are totally decoupled. Moreover, the transport infrastructure has to provide some quality of service that has to be configured by the control plane. In this respect, 3GPP has standardized the IMS (IP Multimedia Subsystem), which is an architectural framework for delivering multimedia on top of IP. IMS was designed to work on UMTS but nowadays ETSI TISPAN is working to achieve a fix-mobile convergence. An important point here is that IMS relies on the IETF SIP protocol in order to establish multimedia connections between end-users. Using these SIP messages, the IMS core is capable to infer data requirements and then configuring the whole path to reserve necessary resources. Furthermore, a working group in TISPAN proposed some scenarios to include P2P video streaming in the IPTV architecture for IMS in order to increase the scalability of the solution. The goal in TISPAN is to define a general NGN architecture to totally decouple the data plane from the control plane (IP is always there, of course). This architecture allows us to define new access and core network protocols with different capabilities and more aligned with multimedia necessities. In this respect, an important issue is to define the mechanisms to assure quality of service for several layer 2 protocols in both wireless and wired networks.

Ethernet is the dominant technology in enterprise networks and its importance in access and metropolitan is growing, due to its excellent price/performance ratio, configuration convenience and backward compatibility and its high performance. The current paradigm of transparent bridge with spanning tree limits severely the scalability of Ethernets. Transparent learning bridges with spanning tree protocols exhibit loop free operation with zero configuration, without requiring complex routing information, IP address and segment administration. However, spanning tree protocols also limit severely the scalability and performance of Ethernet networks because they block all links exceeding the number of network nodes minus one. The resulting routes along the spanning tree between two arbitrary hosts are far from optimum due to the low connectivity and concentrate traffic on the active links. In addition, the resulting infrastructure is vulnerable to hardware failures or configuration errors in the sense that these events may produce broadcast storms and even network meltdown of the switched domain.

The goal is to explore, validate, simulate and implement if economically effective advanced self configuring Ethernet bridge architectures like the recently proposed (at concept level) of FastpathUD, a protocol that operates in the data plane to set up paths. In this proposal, near-optimum paths are found in the data plane by flooding broadcast frames (for example, issued as a result of ARP). These paths are then selected as a bidirectional, symmetric path upon reception of the unicast reply from destination host or bridge. Loop prevention is achieved by the prohibition for broadcast traffic of down-up turns at nodes, as used by Up/Down [1], much less stringent in terms of available paths than spanning tree. The resulting architecture requires zero-configuration, uses standard Ethernet frame format, is fully transparent to hosts by NAT of MAC or frame encapsulation and compatible with 802.1D bridges.

The rest of the paper is organized as follows: Section II presents the overall objectives of the MEDIANET project while Section III deeply describes those objectives including the intermediate results obtained in the project during the first two years. Section IV concludes the paper and presents our future work.

II. OVERALL OBJECTIVES

The goal of this project is to present a complete and valid proposal to allow high-bandwidth demand and delay/jitter restrictive multimedia services. Several phases are necessary to success on this project: analyze requirements, cross-layer design, in-layer improvements and prototype implementation when necessary.

We are convinced that it is important to learn from our errors. Computer and telematic sciences are flood of bunches of errors and bad decisions that have a common factor: lack of foresight and imagination. It is important to start analyzing all multimedia requirements, not for current services but for the future ones. Another important key phase is to decide how layers have to be designed in order to glue all together. Current models use a layered division of the whole problem to transmit information between distant (or not so) devices. Cross-layer design is a new technique or vision to try to put everything together working in a harmonic way.

Obviously, the hardest task in this project is to introduce in-layer improvements in order to achieve the proposed goal. It is important to define milestones for this phase, because all partners have to receive feedback from the intermediate results from other partners.

Campus and data center networks and switched networks in general need are currently limited in performance by the

spanning tree protocol and the proposals for shortest path bridges are too complex and based on a hybrid of switch and layer two router. An evolution of the bridging mechanisms is needed to achieve efficient, zero configuration shortest path bridges without routing. Prototype implementation phase is also important to refine or even reject design proposals from the previous stages. This will help us to disseminate our research results to the scientific community.

III. DETAILED OBJECTIVES AND INTERMEDIATE RESULTS

Next, we deeply describe the objectives of the MEDIANET project.

A. Generic study of Advanced Data-Plane Ethernet switch architectures (ADPE)

Activities are oriented to study the architectures, mechanisms and protocols of new Data-Plane Ethernet switches, as an alternative to hybrid architectures like routing bridges and to evaluate their performance., Ethernet is the preferred technology for enterprise and campus networks due to its self configuration capability, its excellent price/performance ratio and its upgradeability. But Ethernet does not scale due to the limitation of the spanning tree protocol. Current proposals under standardization like Shortest Path Bridges at IEEE 802.1 Group and Routing Bridges (TRILL) at IETF are based on routing protocol on layer two. They have limited scalability due to complexity and limited self configurability. They also need additional mechanisms for synchronization of port transitions to forwarding state.

B. Specific study, research and development of a new Fast-path based Data-Plane Ethernet Architecture

Activities are oriented to requirement definition and detailed development including design of new mechanisms for a new Data-Plane Ethernet architecture using the recently conceived Fastpath switches, object of a patent application on February at OEPM (reference P20090058). Other switches like evolved versions of TRAIN routing (tree based) will also be proposed and evaluated.

Ethernet is increasingly the technology of choice for campus and metropolitan networks. But Ethernet does not scale and the proposals based on hybrid devices like routing bridges do not satisfy all requirements like scalability and low complexity. Fastpath [1] (also known as ARP Path) [2] switches are a new born proposal with high potential due to its evolved bridge concept, that removes many of the limitations imposed by the spanning tree protocol, that blocks all redundant links, precluding shortest paths. With ARP-Path, on-demand, low latency paths are obtained with this approach.

This architecture may self configure, is compatible with Ethernet frame format, may interoperate with standard bridges (RSTP) in core-island mode and keeps transparency


Figure 1. ARP-Path path confirmation with snooped ARP Reply

Figure 2. ARP Path throughput in a small data center network (b)

to hosts and routers. ARP-Path Ethernet Switches set up an on-demand path between two hosts just reusing and flooding the standard ARP Request frame through all links and confirming the (fastest and unique) path reaching to the destination host with the ARP reply frame as shown in Figure 1. ARP-Path uses the standard Ethernet frame format, is fully transparent to hosts and does not require spanning tree or link state protocol. Simulation results [2] show superior performance to spanning tree and throughput similar to shortest path routing, with lower complexity as shown in Figure 1. Implementations [3] show protocol robustness.

C. On-demand Provision of Virtualized Resources

Cloud computing enables the deployment of an entire IT infrastructure without the associated capital costs, paying only for the used capacity. The new "Infrastructure as a Service" (IaaS) paradigm has been introduced to better respond to changing computing demands, so allowing to add and remove capacity in order to meet peak or fluctuating service demands [4]. Cloud was arguably first popularized in 2006 by Amazons Elastic Compute Cloud (or EC2)[1] which started offering virtual machines for $0.10/hour using both a simple web interface and a programmer-friendly API. Although not the first to propose a utility computing model, where computational resources are provisioned on demand as a service, Amazon EC2 contributed to popularizing the "Infrastructure as a Service" paradigm, which became closely tied to the notion of Cloud Computing. An ecosystem of providers, users, and technologies has coalesced around this IaaS cloud model. More IaaS cloud providers, such as

[1]http://aws.amazon.com/

GoGrid[2], FlexiScale[3], and ElasticHosts[4] have emerged over time, and several companies base their IT strategy on cloud-based resources, spending little or no capital to manage their own IT infrastructure. Some providers, such as Elastra[5] and Rightscale[6], focus on deploying and managing services on top of IaaS clouds, including web and database servers that benefit from the elastic capacity of IaaS clouds, allowing their clients to provision services directly, instead of having to provision and setup the infrastructure themselves. Other providers offer products that facilitate working with IaaS clouds, such as rPaths rBuilder[7], which allows dynamic creation of software environments to run on a cloud. There is now a growing interest in open-source cloud-computing tools that provide commercial cloud compatible interfaces, and let organizations build and customize their own cloud infrastructure. Eucalyptus[8] and Globus Nimbus [5] are examples of flexible tools that allow organizations to transform clusters into "IaaS" clouds. Eucalyptus is compatible with Amazons EC2 interface. Globus Nimbus exposes EC2 and Grid community interfaces and offers self-configuring virtual cluster support. These relevant components provide complete support for remote access, including high level functionality for security, contextualization, and image management. However, the management model offered by these systems cannot be easily extended and adapted to a given physical infrastructure; they do not consider the management of generic services as a whole; and they do not support the dynamic allocation and balance of computing resources. This activity focuses on the research required to build a cloud infrastructure to dynamically provision the resources needed to support the service workloads in the project. The activity will conduct research to meet the main challenges in the dynamic and scalable management of virtual machines in datacenters:

- Reference model for the management of services on distributed infrastructures. The service, and not the VM, will be first-class management entity. A service is a group of VMs with an inter-connection relationship (usually a network), placement constraints and context. The management will be service agnostic, to preserve the generality of the model. The management model must integrate different areas, namely: storage, networking and virtualization.

- Advance reservation of capacity. Existing cloud use immediate provisioning, the VMs must be allocated right away or not at all or best-effort provisioning, the VMs are queued until resources can be allocated. However,

[2]http://www.gogrid.com/
[3]http://www.flexiscale.com/
[4]http://www.elastichosts.com
[5]http://www.elastra.com/
[6]http://www.rightscale.com/
[7]http://www.rpath.org/
[8]http://www.eucalyptus.com/

the media services considered in the project requires more complex policies and the ability to reserve a given capacity to meet service demands [6].

- Dynamic contextualization. The service can be deployed in any cloud, so it is necessary to provide the service components with an infrastructure context awareness. In this way, a service and its components will automatically configure themselves to operate in a given environment. The management layer must provide support to generate and push these context data into the VMs.

- Service Management. The deployment of a service implemented as a set of interconnected virtual machines requires special support form the underlying management entity. In particular, it is necessary to specify the service components (including networks) and their relationship. Moreover, special placement rules could be applied at service level.

- Support for service elasticity. Elasticity support is required to meet variations in service workload. Elasticity could be provided in a number of forms that tune the capacity assigned to the service: adding or removing VMs of a given type to the service; increasing the physical resources (e.g. CPU or memory) allocated to a VM; or varying the total capacity assigned to the interconnection components [7].

- Heuristics for local placement of virtual machines. The placement of Virtual Machines in physical resources may follow different targets: balance workload among to improve efficiency; or consolidate workload to a reduced number of physical systems, so reducing space, power and cooling requirements.

D. Federation of Cloud Infrastructures

Cloud providers are opening new infrastructure centers at different geographical locations and it is clear that no single facility/provider can create a seemingly infinite infrastructure capable of serving massive amounts of users at all times, from all locations. In this way, there are emerging some interesting initiatives that explore the idea of federating different cloud infrastructures. For example, RightScale is leading an interesting commercial federation initiative, by broadening its cloud management platform to support emerging clouds from new vendors, including FlexiScale and GoGrid, while continuing its long-standing support for Amazon EC2. There are also some interesting works that integrate a local virtual infrastructure with external resources from a cloud provider. For example, BioTeam has deployed the hybrid setup, which combines local physical nodes with virtual nodes deployed in the Amazon EC2. However, the federation of different virtual infrastructures and/or clouds providers involves some important challenges and problems that have yet to be solved. First, the degree of interoperability for different infrastructures is very limited

or null because different providers use proprietary interfaces to deploy, monitor, and access the virtual resources. This makes it difficult to have compatibility among the different providers [8]. A standardization effort should be made to unify the different interfaces that access different providers. Networking and security issues are also major concerns in the federation of virtual infrastructures. For example, deploying a complex service that spans multiple sites or migrating a public service from one site to another involve important technological challenges related to network management. Furthermore, there are some important security problems that remain to be deeply analyzed. Some examples include managing users rights to deploy, monitor, or shut down virtual machines in a remote provider, controlling the access to remote virtual resources, or isolating the traffic of different virtual networks. Finally, performance-related issues are also an important factor to consider in federated environments, since deploying and accessing a virtual machine on a remote site always involves higher delays and overheads than deploying on a local one, which can have a negative impact on performance. In this context, it is essential to identify the main performance degradation sources (image transfer delays, extra network latencies, etc.) to optimize the behavior of the federated infrastructure. This activity conducts research in the following topics:

- Architectures for cloud federation. We envision two main federation scenarios: (1) A local infrastructure (datacenter) interface with one or more clouds to supplement its capacity or to better deal with service requests (for example moving the service to a provider closer to the users); and (2) Brokering of different clouds, where a service can be deployed using resources from different providers. This two basic architectures can be easily combined in hybrids setups.
- Service aware management. When different providers are available the management of the underlying architecture can be improved if service specific behavioral metrics are incorporated in the system. For example: it would be possible to develop elasticity heuristics for the service, so effectively allocating more VMs in the cloud beforehand; deploying part of the service closer to the users; or developing effective replication policies [9].
- Heuristics for remote placement of VMs. Several alternatives for remote placement will be examined, always driven by the idea that remote placement versus local placement should take into account the level of coupling between sites. The heuristics could be initially based on the following criteria: (1) VM requirements, in terms of capacity level agreements, such as hardware requirements, network requirements, maximum waiting time for VM deployment, affinity/anti-affinity rules, proximity rules; (2) local resource availability in terms

of physical resources and network; (3) remote resource availability in terms of physical resources and network availability; (4) performance parameters, such as VM deployment time, or connection establishment time; or (5) incomes, expenses and benefits/losses.
- Network management across sites. When a service is deployed in a federated infrastructure, as described in 2.2.1. It is necessary to guarantee the connectivity between all the service VMs, no matter where these components have been deployed. Moreover, the security and functionality of the links must be preserve. This usually will require the deployment of virtual private networks with different topologies among different cloud providers.
- Image management and contextualization across sites. Services are usually build of clonable components. These components can be replicated when additional capacity is needed. The use of different cloud providers makes the management of these images challenging. We will explore different strategies to minimize the cost of strong the same image in different providers, reduce the effort to distribute updates of the image, and deal with different cloud storage interfaces.

E. Towards a scalable high-definition TV service in the future Internet

The Internet is growing day after day with new services, protocols and users. This growing has an enormous impact on its performance so, some improvements or even modifications are necessary to allow the deployment of new services with high requirements which comes from out the IP world. TV and Video on Demand (VoD) are two of such kind of services, which wants to land on the Internet field but, due to its real time requirements, are not prepared to be massively deployed yet. There are some initiatives to enhance the quality of experience (QoE) perceived by the users of Internet and, at the same time, to create an all-IP platform to transport whatever kind of service. In this respect, 3GPP (the third Generation Partnership Project) has standardized the IMS (IP Multimedia Subsystem), which is an architectural framework for delivering IP multimedia services. It was first envisioned for the mobile world but nowadays, initiatives like the TISPAN working group from ETSI are working to extend its functionality to take fixed networks into account, developing a true NGN (Next Generation Network). ETSI-TISPAN releases one document introducing IPTV support to the NGN. This is an important milestone for the NGN specification, because of the new business it provides to third party providers. IPTV providers will be able to construct their services on top of this NGN guaranteeing a good QoE to their users (even better than the traditional broadcasting networks used nowadays).

But there still are some technical issues to be solved before the NGN-based IPTV deployment. For example, it is
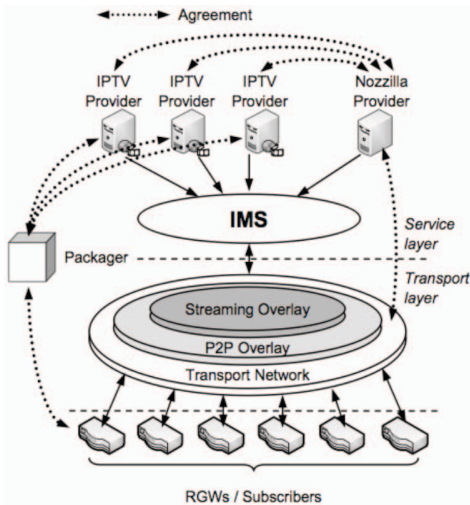
Figure 3.   P2P IPTV over IMS

important to decide the mechanism used to transport the data from the service provider server to the end user devices. It is clear that the best mechanism, in terms of used bandwidth in the whole network, is using multicast IP. But using multicast IP could have some drawbacks in this kind of environment with hundred of channels and, potentially, millions of users. In [10] the authors present a study to analyze the benefits of introducing P2P techniques in IPTV live streaming services. That study concludes that although for popular channels is better to use multicast IP, for unpopular channels (which are around the 80% of the total number of channels following a Pareto principle) the best solution in terms of load is to use P2P. With this idea in mind, it is worth to propose new mechanisms to include P2P techniques in NGN-based IPTV as in [11], where Figure 3 presents the architecture of this proposal. Moreover, in [12] the authors analyze the whole signaling with SIP and 802.21 protocols to achieve a seamless and transparent device mobility, so to provide a better performance for itinerant users.

IV. CONCLUSION

The MEDIANET project has obtained several promising intermediate results, but we still have two more years to improve and to integrate them in a global and unique architecture for the video delivering in next generation networks. We envision the future Internet as an all-IP network with strong relationship with the Ethernet layer 2 protocol to guarantee an end-to-end quality of service. Over this improved network, service providers may construct a video streaming service by using the a federation of clouds and virtualization in order to better use their resources. Single IP subnet layer, zero configuration, two bridged networks are also important for data center and campus networks to achieve the desired end-to-end performance

REFERENCES

[1] G. Ibáñez, "Arp-path ethernet switching: On-demand efficient transparent bridges for data center and campus networks."

[2] G. Ibáñez, J. Carral, J. Arco, D. Rivera, and A. Montalvo, "Arp-path: Arp-based, shortest path bridges," *IEEE Communication Letters*, 2011.

[3] G. Ibáñez, J. Naous, E. Rojas, D. Rivera, J. Carral, and J. Arco, "A simple, zero-configuration, low latency, bridging protocol." The 35th IEEE Conference on Local Computer Networks (LCN), 2010.

[4] J. Vazquez-Poletti, G. Barderas, I. Llorente, and P. Romero, "A model for efficient onboard actualization of an instrumental cyclogram for the mars metnet mission on a public cloud infrastructure," in *Proc. PARA2010: State of the Art in Scientific and Parallel Computing, Reykjavik (Iceland), June 2010*, ser. Lecture Notes in Computer Science, vol. 7133. Springer Verlag, 2012, pp. 33–42.

[5] J. Bresnahan, T. Freeman, D. LaBissoniere, and K. Keahey, "Managing appliance launches in infrastructure clouds." Teragrid 2011, 2011.

[6] R. Montero, R. Moreno-Vozmediano, and I. Llorente, "An elasticity model for high throughput computing clusters," *Journal of Parallel and Distributed Computing*, vol. 71, no. 6, pp. 750–757, June 2011.

[7] R. Mian, P. Martin, and J. Vazquez-Poletti, "Provisioning data analytic workloads in a cloud," *Future Generation Computer Systems*, vol. in press, 2012.

[8] J. Tordsson, R. Montero, R. Moreno-Vozmediano, and I. Llorente, "Cloud brokering mechanisms for optimized placement of virtual machines across multiple providers," *Future Generation Computer Systems*, vol. 28, no. 2, pp. 358–367, 2012.

[9] J. Lucas-Simarro, R. Moreno-Vozmediano, R. Montero, and I. Llorente, "Scheduling strategies for optimal service deployment across multiple clouds," *Future Generation Computer Systems*, vol. in press, 2012.

[10] A. Bikfalvi, J. García-Reinoso, I. Vidal, F. Valera, and A. Azcorra, "P2p vs. ip multicast: Comparing approaches to iptv streaming based on tv channel popularity," *Computer Networks*, 2010.

[11] A. Bikfalvi, J. García-Reinoso, I. Vidal, and F. Valera, "A peer-to-peer iptv service architecture for the ip multimedia subsystem," *International Journal of Communication Systems*, vol. 23, no. 6-7, pp. 780–801, 2010.

[12] I. Vidal, J. Garcia-Reinoso, A. de la Oliva, A. Bikfalvi, and I. Soto, "Supporting mobility in an ims-based p2p iptv service: A proactive context transfer mechanism," *Computer Communications*, vol. 33, no. 14, pp. 1736–1751, 2010.