A QoS Enabled Public Ethernet Access Network

(Invited)

V. Souza, B. Tremblay, P. Karlsson, K. Persson, W. Zhao, Z. Ghebretensaé

Broadband & Transport, Ericsson Research
Tel: +46 87190000, Fax: +46 8184085, E-mail: victor.souza@ericsson.com

It is a known fact that today's access networks represent a bottleneck to provide full quality assured services to end users. Metropolitan networks and core backbones are over provisioned and rarely represent a problem from the end user perspective. The advent of triple-play has brought new aspects to be considered, since no quality drops are desired during the execution of videos and phone calls. This paper proposes an architecture to offer guaranteed Quality of Service (QoS) for end users in a public Ethernet access network. The key element of the architecture is an access edge node with strong QoS and traffic-engineering capabilities. A resource manager is responsible for ensuring efficient utilization of available resources, performing policy-based admission control, resource reservation, network elements configuration and monitoring. By allocating resources in the access network, network service providers can assure service delivery featuring high bandwidth and prioritization. Finally, we present the architecture implementation and the installation of a Video on Demand (VoD) and TV over Internet Protocol (TVoIP) services over the developed platform.

1. Introduction

It is a known fact that today's access networks represent a bottleneck to provide full quality assured services to final customers. While metropolitan networks and core backbones are usually over provisioned, access networks hold strict bandwidth limitations mainly associated to the technology deployed. The last mile is an example of a critical point. Copper lines are largely used and quality variations are commonly observed [1]. It is a hard task to control or reduce these variations dynamically but it can be fairly simple to take actions at application level when they occur.

The need to transmit media stream in real time poses high restrictions when it comes to bandwidth and latency. During voice calls latency should be kept as low as possible to make the experience equivalent to traditional phone calls. Bandwidth consumption can be observed when a user wants to access a video stream in high definition.

Regardless, quality of service in core and metropolitan networks has been the center of studies during the last decade. A myriad of solutions and protocols have been proposed and compared. However, without any guarantees in the access network these solutions make little practical difference to user experience. Efficient usage of resources in access networks is therefore of major importance for providing full end-to-end quality of service.

This paper is organized as follows. The Chapter 2 describes an architecture to provide QoS and traffic engineering in access networks. The Chapter 3 presents a prototype developed based in the proposed architecture, highlighting implementation decisions. The installation of a TVoIP and VoD services is also described. Chapter 4 presents the conclusions.

2. Proposed Architecture

The architecture proposed is based on a new network model designed for a multi-access scenario. Simplistically, three different domains are defined with different characteristics and constraints. The first one is the user domain where all customer equipments are located. The second one is the access domain, comprising the access and regional networks. The last one is the service provider domain, where all the servers and necessary infrastructure to provide services are deployed.

The customer premises equipment includes all the hardware and appliances connected to the network residing at the user's premises. The management of this area is performed by the end user in case a local network is present. A residential gateway can be used between this domain and the access network. Optionally, the end user can be directly connected to the access node, without the need for a residential gateway. Several access technologies are allowed in the last mile, including FTTx, xDSL and copper Ethernet.

The second domain is a product of the access and regional networks. In the access network traffic from customers is aggregated, initially in the access node, through the layer-2 aggregation network and finally in the access edge node. The regional network is essentially a layer-3 metropolitan area network.

The last domain is managed and implemented by the service provider. The service logic is implemented in the application servers. Other servers may also be present, implementing functionalities of accounting, media streaming, and authorization. Figure 1 shows the proposed network model.

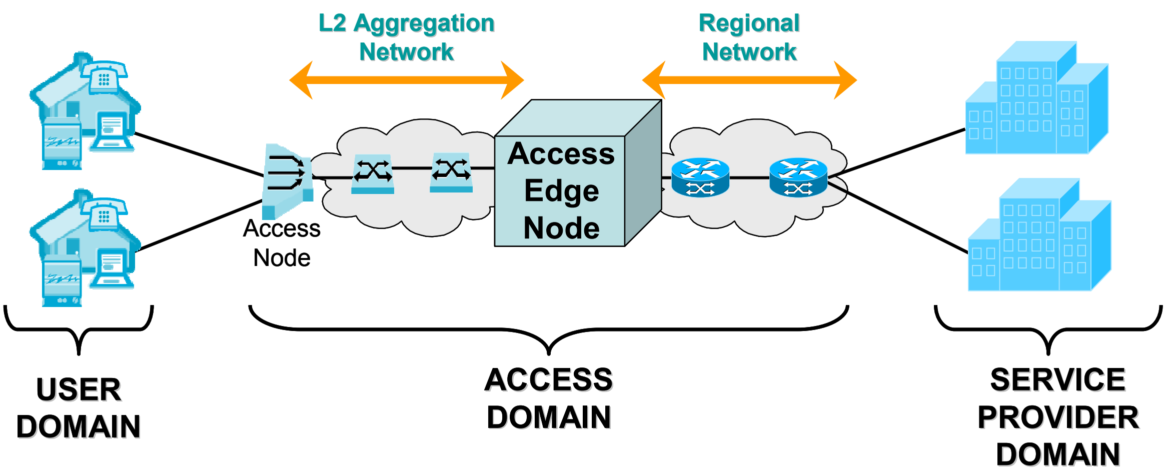

Figure 1: Proposed network model.

In the following sections all the components of the architecture and their functional requirements are discussed.

2.1 Access Edge Node

The Access Edge Node (AEN) is the core component of the architecture, performing the mapping between the layer-2 access network and the layer-3 regional network. Policies that are to be enforced by the resource manager can be pushed down to the access edge node which acts as policy enforcement point. Functionalities such as traffic shaping and monitoring are performed in this point of the network and propagated to the access node when needed.

The access edge node is also responsible for enabling equal access. This is possible due to the creation of service bindings, which can be created by any service provider requesting resources in the network. The role of the access edge node is to manage and control the service bindings by forwarding the L2 frames from the customer premises equipments to a specific service provider. In this sense, the access edge node can be seen as bridge between the L2 and L3 environments.

Several other functionalities are embedded in the AEN since it is located in a strategic location, i.e., the termination of the layer-2 connection from the users. For example, a DHCP server (and proxy) is present there to handle requests for IP addresses. Other functionalities include HTTP-proxy, multicast group management and authentication-proxy.

2.2 Access Node

The Access Node (AN) is the element where individual physical connections to the user domain are attached and maintained. It provides a bridge between the Ethernet switching domain of the access network and the end-user's network domain. Therefore, the access node acts as a gateway between the physical technology used by the user devices to connect to the access node on the last mile (Wi-Fi, xDSL, fiber) and the Ethernet technology used by the aggregation network. It is the first point in the access network where traffic coming from multiple user ports is aggregated towards the access edge node.

The access node has the role to forward and regulate traffic to/from the user ports and to/from the access edge node through the layer-2 Ethernet aggregation network. It also has the responsibility to perform traffic marking (802.1P field) and policing on the upstream of service bindings based on defined policy rules. By using the 802.1P marking scheme traffic with different characteristics and requirements such as latency and jitter can be prioritized in the aggregation network.

Besides, access control lists are installed in the access node by the access edge node upon reception of a new service subscription. In this way, specific protocols from one user device are allowed into the network.

2.3 Resource Manager

The Resource Manager (RM) is the element of the architecture in charge of controlling resources in the first mile and in the access network. It comprises facilities for admission control, resource reservation, monitoring, configuration and policy-based control. The resource manager plays a role equivalent to A-RACS in the ETSI TISPAN architecture [3].

Essentially, the resource manager receives requests for resources from the service manager residing in the service provider's domain. By evaluating the current state of resource usage and application of domain specific policies the RM is capable of deciding upon each request. In order to take decision the RM possesses an internal database that corresponds to the full topology of the access network, including the access nodes, L2 aggregation network and access edge nodes. Moreover, it keeps track of used resources in this topology map in a static approach of resource management.

Alternatively, the RM can poll the network elements to retrieve the current available resources, leading to a dynamic view of currently used resources. The later model is used in the monitoring of first mile copper lines by contacting the Loop Qualification and Monitoring (LQ&M) software responsible for monitoring DSL lines. After a decision has been taken, the configuration of network elements takes place. This happens by pushing down policies, services bindings and access control lists into the network elements (e.g., access edge node and access node). For a more detailed view about the resource manager please refer to [4].

2.4 Service Provider

The Service Manager (SM) is the component of the platform that resides in the service provider domain and is responsible for interfacing with the network resource provisioning system represented by the resource manager. It plays a role equivalent to SPDF in the ETSI TISPAN architecture [3].

When subscribing to a service, the client contacts the application server in the service provider's domain. Fundamentally, the service manager has the function of receiving requests for service access from the application server and translating them into network parameters characterized in a Service Level Specification (SLS). The SM then sends the service definition (e.g., multicast group, resources required) and the authorization for that user to the resource manager. The resource manager checks if enough resources are available to deliver the service in the specified period. If there are enough resources and the user is known to the access domain, the resource manager will allow the service by sending the request to the access edge node; otherwise the resource manager rejects the request and sends a negative response to the service manager.

The mapping between the end-user's request and the SLS sent to the resource manager needs to be configured in the service manager. This file must specify precisely the network resources that are required for optimal service delivering. For example, traffic priority and shaping parameters can be specified.

2.5 Service Bindings

A service binding is a network association that enforces the necessary transport relationship between a user domain and a service provider domain to enable a user device to reach and access the services of a particular provider.

Virtual LANs (VLANs) are used in the architecture as a mean to identify traffic belonging to different service bindings and service providers. A broadcast handler in the access node and in the access edge node ensures that two user domains are never in the same broadcast domain. Therefore, the service providers are able to use private IP addresses for its customers since separate broadcast domains are created for each user domain.

The first VLAN trunk present in the system is connecting the end users to the access node. The client willing to access services from one service provider should have configured in his machine one VLAN per service for the given SP. The second VLAN is created between the access node and access edge node, aggregating traffic towards a given Service Provider in one unique VLAN. Therefore, there is only one VLAN per Service Provider starting in one specific AN towards the AEN. The third and last VLAN starts in the AEN and goes until the first router in the regional network - also Service Provider specific.

Furthermore, the service binding guarantees the service delivery with the right integrity and QoS between a specific user and a specific provider. Besides, it ensures that traffic between different users is kept independent, and that no users are able to view or use resources in another user's network.

3. Platform Description

3.1 Platform Implementation

For assessment of the proposed architecture a fully functional prototype was implemented. All the components were implemented according to the specification and network model presented in the Chapter 2. The developed platform was deployed and tested at TNO (The Netherlands), Ericsson and Acreo (Sweden) in the context of the 6th Framework Project MUSE (MultiService Access Everywhere) [7]. The hardware and services present in each one of the systems differ slightly, conversely they are functionally equivalent.

Two access nodes are included in the prototype. One is an ADSL2+ access node [6] with enhanced software functionalities. The layer-2 aggregation network was implemented by standard Ethernet switches with IGMP (Internet Group Management Protocol [5]) snooping and VLAN capabilities. The central element of the system, the access edge node, was built on commercial-grade network processors and blades that are capable of handling large amount of network traffic. Due to practical reasons the system contains a limited number of client machines. Figure 2 shows the implemented components.

The resource and service managers were implemented in standard blade servers running Linux operating system. Three service providers were deployed in the network in order to prove the multi-service capabilities of the platform. They are the BlueNet, the GreenNet and the IMS Interactive Personalized TV Service Provider. As an example, the BlueNet server's main functionalities are:

- A Darwin Media Server that can stream unicast and multicast in different resolutions;
- A service manager that can negotiate resources with the resource manager;
- An AAA server that handles accounting and authorization of subscribing customers.

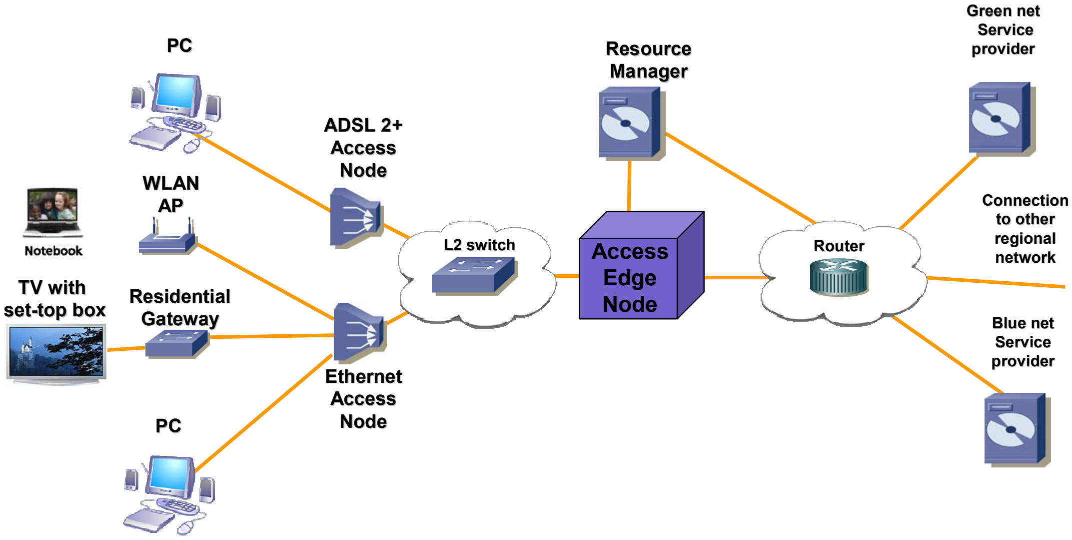

Figure 2: Implemented components.

The regional network in the platform is mainly a router connecting the service providers to the access edge node. A Juniper M5 was used for this purpose. This router is also connected to external networks in the lab environment which extends the actual regional network also into these networks.

3.2 Service Description

As previously exposed the platform implemented features multi-service access. However, in the following sections we focus in only one service from BlueNet. This service provider has two models for the delivery of media to its customers under the movies service, Video on Demand (VoD) and broadcasted (TVoIP style). After subscribed to the movies service the client can decide which one of the models to access. The web page for this service is exposed in the Figure 3.

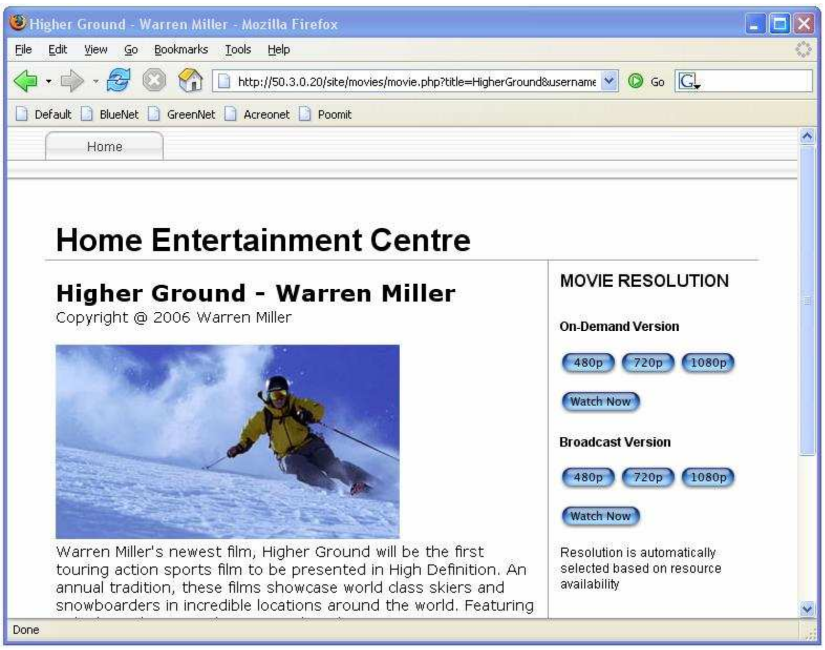

Figure 3: BlueNet movies screen.

In the VoD model the client has a dedicated stream coming to its premises. It uses standard protocols for unicast streaming, namely Real-time Transport Protocol (RTP) and Real Time Streaming Protocol (RTSP). The TVoIP movie service is implemented through a multicast channel controlled by the IGMP protocol. Both models are supported by the proposed architecture.

Both services can be viewed in different resolutions. If a client wants to get maximum possible resolution given the available bandwidth it can use the "Watch now" button. The service manager will negotiate with the resource manager in the access network the maximum possible bandwidth that can be allocated for the given stream.

3.3 System Interaction

The BlueNet VoD service follows the resource allocation mechanism defined in the Chapter 2. In the TVoIP scheme, a special multicast allocation process is used instead. In this scenario, the service provider authorizes the user device to request later for resources through an IGMP Join message handled by the AEN. The service provider will supply to the resource manager the specification of the required resources for delivering the service.

The admission control and the authorization involve the service manager, the resource manager and the access edge node in different moments. When a user subscribes to the service, the service manager will inform the resource manager that the user is authorized to use the multicast service and will also inform about the network resources required to deliver the service.

When a user device launches the multicast service (IGMP Join), the AEN will forward the request to the resource manager that will check that the user device has been authorized to receive the service and that there are enough resources in the network to deliver the service. In the positive case, the AEN will establish the multicast path from the multicast source down to the user device and will forward the multicast traffic. In the negative case, the AEN will reject the request from the user device and no multicast path will be established for this user.

The multicast service delivery for a given user device ends when the first of the following events occurs:

- The user device informs explicitly the AEN that it is not interested anymore in receiving the multicast service by sending an IGMP Leave message;
- The AEN does not receive any positive response (IGMP Report) to its periodic requests (IGMP Query) from that user device;
- The resource manager informs the AEN that the user is not authorized anymore to receive the multicast service due to end of subscription.

4. Conclusion

Quality of service has been shown to be of major importance when it comes to video services offered to triple play clients. While some applications such as web browsing and e-mail access do not pose high requirements to the network, no drop outs are desired in video or phone calls quality.

The available bandwidth in the last mile has increased dramatically in the last few years. However, without a mechanism for differentiating traffic in the network, certain services still present unacceptable quality to end users. Applications that are non delay or jitter sensitive can generate traffic bursts causing degradation of sensitive ones when no control over link usage is applied.

This paper has presented an architecture and implementation of a public Ethernet based access network featuring Quality of Service and strong resource control mechanisms. The key element of the architecture is an access edge node capable of performing traffic marking, shaping and creating service bindings. This node must work in cooperation with the resource manager and access nodes to provide true end-to-end quality of service.

The implementation of a Video on Demand and TVoIP services over the platform was valuable to prove the platform features. Dynamical monitoring of the DSL line supplied useful information to the resource manager which in turn was able to provide optimized video resolution depending on the current available bandwidth. No drops in the quality are perceived even when simultaneous services are accessed.

Acknowledgements

The authors would like to thank the European Commission IST 6th Framework Project MUSE and the design and development team in Ericsson Research Montreal Canada for their support.

References

[1] K. Vanbleu, G. Ysebaert, G. Cuypers, M. Moonen, “Adaptive bit rate maximizing time-domain equalizer design for DMT-based systems”, *IEEE Transactions on Signal Processing*, vol. 54, no. 2, pp. 483-498, February 2006.

[2] J. Rius i Riu, et al., “The IST-MUSE Approach to DSL Loop Qualification and Monitoring”, *Proc. Broadband Europe 2006*, Genève (Switzerland), December 2006.

[3] *Resource and Admission Control Sub-system (RACS), Functional Architecture, version 1.1.1*, ETSI Standard ES 282 003, June 2006.

[4] E. Trojer, M. Berg, Z. Ghebretensaé, J. Olsson, I. Pinilla, J. Rius I Riu, “End-to-end Resource Allocation in a Multi Access Network”, *Proc 12th European Conference on Networks & Optical Communications*, Stockholm (Sweden), June 2006.

[5] B. Cain et al., *RFC3376, Internet Group Management Protocol, Version 3*, IETF Proposed standard, October 2002.

[6] Ericsson Ethernet DSL Access (EDA 1200), www.ericsson.com.

[7] P. Vetter, et al., "Multi Service Access Everywhere", *Proceedings 31st European Conference on Optical Communication*, Glasgow (Scotland), September 2005.