White Paper

Software Configurable Backhaul Network Sharing

DragonWave

www.dragonwaveinc.com

NOTICE

This document contains DragonWave proprietary information. Use, disclosure, copying or distribution of any part of the information contained herein, beyond that for which it was originally furnished, requires the written permission of DragonWave Inc.

The information in this document is subject to change without notice and relates only to the product defined in the introduction of this document. DragonWave intends that information contained herein is, to the best of its knowledge, correct and accurate. However, any/all liabilities associated with the use or accuracy of the information contained herein must be defined in a separate agreement between DragonWave and the customer/user.

DragonWave®, Horizon® and Avenue® are registered trademarks of DragonWave Inc. ©2015 DragonWave Inc. All rights reserved

Copyright 2015 DragonWave Inc. All rights reserved.

TABLE OF CONTENTS

Copyright 2015 DragonWave Inc. All rights reserved.

1. Overview

To improve profitability, network operators and service providers (SPs) have been seeking ways to increase the utilization of their infrastructures, among which network sharing by multiple service providers has emerged as a key strategy of profitability in recent years.

A good example of network sharing is the mobile backhaul network that is shared from the cell sites all the way to the controller site by multiple mobile service providers. The benefits of network sharing are simple and significant. Network operators get more revenues from multiple service providers, while service providers avoid the investment in the network and thus significantly reduce the entry cost; improve time-to-market, and lower operating costs. Each of the players in network sharing can focus on its core business and competency.

In a sense, mobile network sharing can be perceived as another type of cloud service. Compared with cloud service over server farms, which is the most popular infrastructure-based cloud service and has been widely adopted, mobile network sharing is much less studied and faces many technical unknowns.

This paper will address some of these unknowns and provide possible solutions.

2. The Goals of Mobile Network Sharing

From a technical perspective, the goals of mobile network sharing include:

1. The services of various operators over the shared network are clearly distinguished, individually managed, and can be established from end to end.
2. Security of the services is ensured. Traffic of one mobile operator will not leak to or be seen by any other mobile operators.
3. The quality of service (QoS) is guaranteed. QoS has many parameters. Bandwidth, packet delay and jitter, and packet loss rate are the major QoS parameters. If a service level agreement (SLA) is signed between network operator and service providers, there needs to be a mechanism to ensure and monitor the QoS of each service and produce SLA reports. Particular to network sharing, implementation of the policy of allocating the excessive network resources such as bandwidth among service providers is also required.

The mobile network sharing solution must also take into account the total cost of ownership (TCO) in achieving the technical goals above and the operation cost is absolutely the top concern of network operators.

3. Service Models and Definitions of Network Sharing

Multipath fading refers to a radio signal having multiple paths that can reach the receiver at the same time the Without losing generality, the following example of a service model depicted by Figure 1 reflects the needs of many network operators providing network sharing services.

A network port is shared by several (2, in this example) mobile operators or service providers who have their own service (implemented by a certain type of virtual circuit or VC such as Provider VLAN based on IEEE 802.1ad) with multiple traffic classes each of which corresponds to one of the applications such as voice, video, and data.

Every mobile operator wants to get the guaranteed bandwidth for its services. Understandably, each mobile operator does not want its unused bandwidth to be shared by others. And for this purpose, both the service and each class of traffic within the service are subject to their bandwidth profile.

A bandwidth profile consists of the following four parameters, as defined by Metro Ethernet Forum (MEF):

CIR – Committed Information Rate, which must always be guaranteed from the end to end of a service.

Copyright 2015 DragonWave Inc. All rights reserved.

EIR – Excessive Information Rate, which is the maximum burst rate allowed and whose rate is not guaranteed. EIR is exchangeable with PIR (Peak Information Rate).

CBS – CIR burst size, either the period or the number of bytes that a service is allowed to send packets into the network continuously at CIR.

EBS – EIR burst size, either the period or the number of bytes that a service is allowed to send packets into the network continuously at EIR.

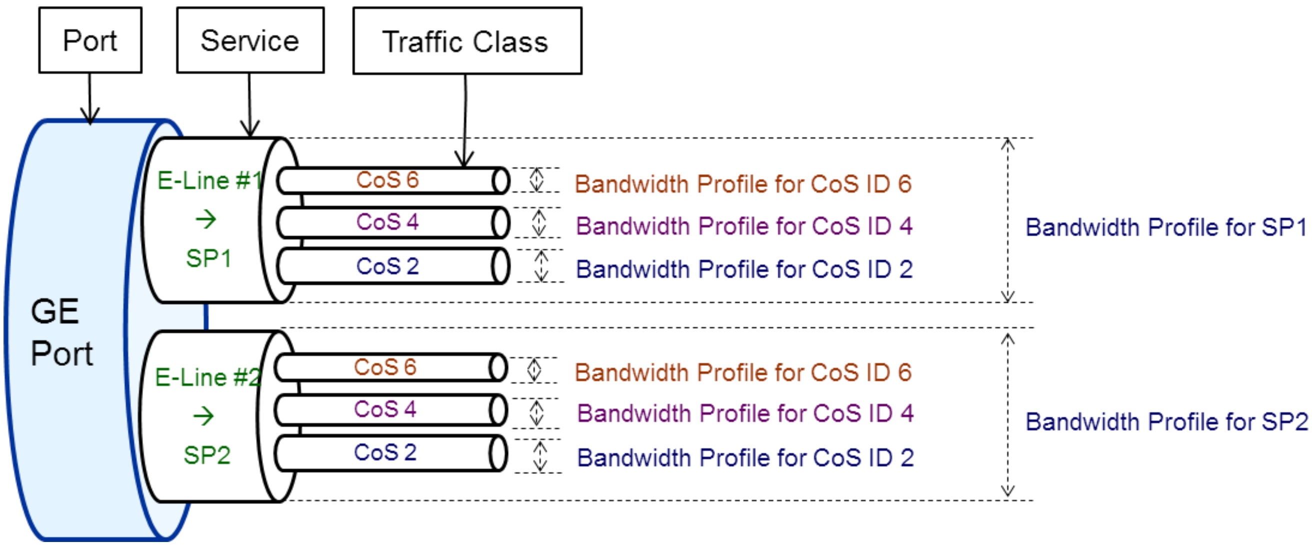

Figure 1: Figure Definition of Network Sharing

Copyright 2015 DragonWave Inc. All rights reserved.

4. Hierarchical QoS (H-QoS) for Network Sharing

The service definition above is naturally implemented by a model that supports QoS on both the service/VC level and the traffic class level. This QoS model is called Hierarchical QoS or H-QoS as illustrated in Figure 2. It is named as such because traffic goes through multiple levels of QoS processing.

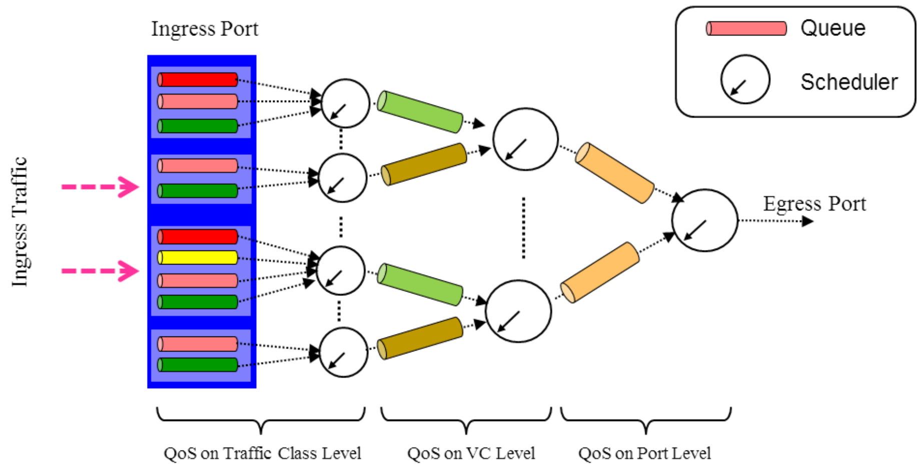

Figure 2: H-QoS Model

For almost two decades, H-QoS has been widely used by BRAS (Broadband Remote Access Server) which is basically a big router with the support for H-QoS, usually used to maintain the PPP sessions and control the traffic for Internet access subscribers (especially when these subscribers are connected via DSLAMs). BRAS is deployed at a central location in front of the core routers and can identify the types of traffic of each subscriber as well as control the traffic rate for each subscriber by H-QoS so that in the downstream direction traffic of low priority gets dropped by BRAS prior to the real-time and other high-priority traffic when the overall received traffic from Internet will jam the DSL line whose rate is known to the BRAS. H-QoS architecture in BRAS reflects how the traffic is aggregated, i.e., queues in BRAS are organized in multiple levels such as the traffic class level, subscriber level and DSLAM level, as shown in Figure 3.

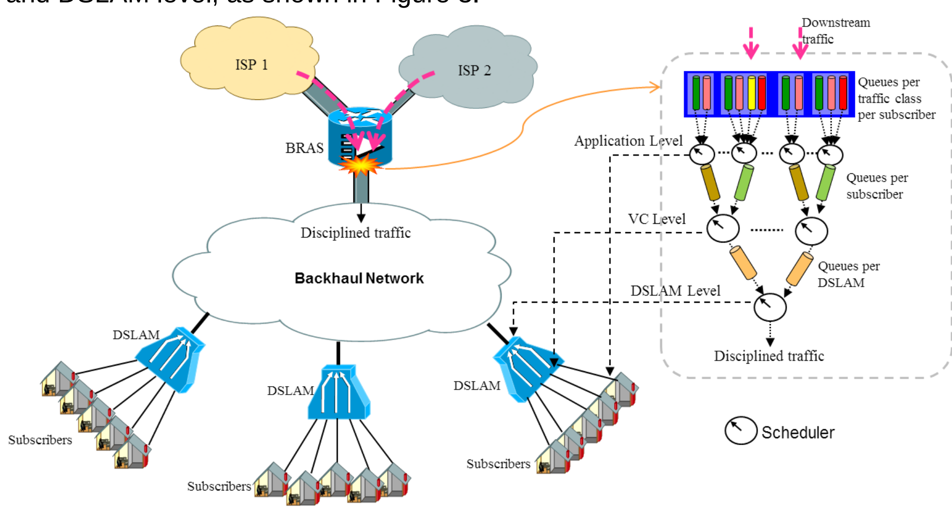

Figure 3: H-QoS in BRAS

Copyright 2015 DragonWave Inc. All rights reserved.

Thanks to H-QoS in BRAS, the DSLAMs and the backhaul network can be kept very simple without the support for sophisticated QoS functions. Thus, TCO is reduced by centralizing the complexity.

This is not the case in network sharing. Simply put, network sharing requires H-QoS control on the traffic by every network device on the paths of the services to fulfill the service definition described by Figure 1. If a single network device on the path of a service does not support H-QoS as other devices do, traffic passing through it will not be handled in the expected way and the objectives of the service definition will not be completely reached.

5. The Reality and Practical Solutions

The likely reality is that most of the installed devices do not support H-QoS. How can network sharing be implemented if the network devices only support basic or plain QoS (i.e., single hierarchy of QoS control)?

Network sharing can be done if we simplify the service definition by applying QoS to only one level. There are two options as shown in Figure 4-1 and Figure 4-2.

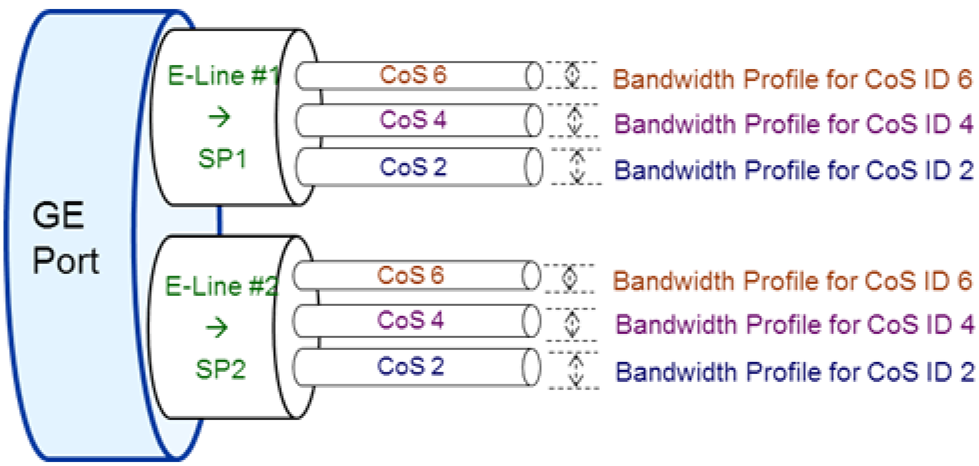

Figure 4-1: Multi-CoS on each service and per-CoS bandwidth profile

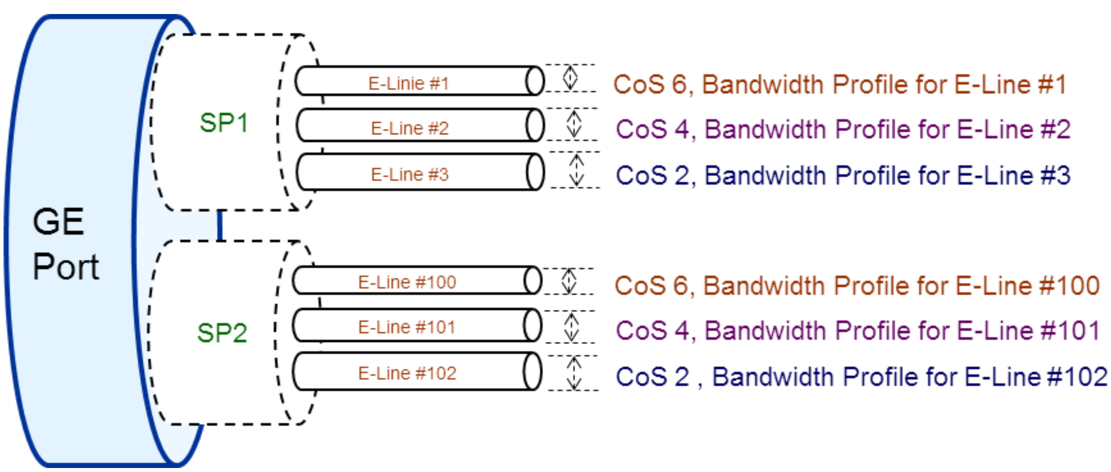

Figure 4-2: Single-CoS on each service and per-service bandwidth profile

Copyright 2015 DragonWave Inc. All rights reserved.

In option 1 (Figure 4-1), an E-Line service is established for each mobile operator or service provider and it is applied with Multi-CoS. Within the service, there are multiple classes of traffic and per-CoS bandwidth profile is applied so that each traffic class has both guaranteed bandwidth (i.e. CIR) and maximum burst rate (i.e. EIR).

The CIR/EIR of the E-Line service can be obtained by the sum of the CIR/EIR of all the traffic classes within the service.

In option 2 (Figure 4-2), a certain number of E-Line services are assigned to each mobile operator or service provider. Single-CoS (traffic of a service belongs to the same class) and per-service bandwidth profile is applied to each service.

The CIR/EIR that a service provider actually gets can be obtained by the sum of the CIR/EIR of all the E-Line services assigned to the mobile operator.

The two mobile operators may have the same traffic classes and thus their traffic can share the same queues. But if there are enough queues, their traffic classes can use different queues. For example, most network devices support 8 queues per port. Two of these queues can be configured as Strict Priority (SP) queues for the traffic of CoS 6, and the other six queues as Weighted Round Robin (WRR) queues for the traffic of CoS 4 and CoS 2 (two WRR queues for each mobile operator in our example).

In these two simplified service models, the CIR of each traffic class and the collective CIR of each mobile operator can be guaranteed if frame coloring and color-aware packet dropping are supported by the network devices, achieving the same goal of the original service definition depicted in Figure 1. Frame coloring means that the ingress packets are marked green if the traffic rate of the service by the time a packet is received on UNI of the service is within CIR, yellow if the traffic rate is between CIR and EIR, and red if the rate is above EIR but red packets are dropped right away at the UNI port so there are no red packets in the network. All the devices on the path of a service will never drop green packets at any time and they will drop yellow packets when only congestion occurs. As all the green packets can get to their destination, CIR is guaranteed, given that the network capacity is not over-subscribed by the total CIR traffic (which must be true even for the original service definition).

The only difference between the original service definition implemented by H-QoS and the simplified ones implemented by plain QoS is in how the excessive bandwidth left by CIR is shared between mobile operators, which affects the actual EIR of each service and traffic class.

In implementing the original service model, H-QoS gives each mobile operator its dedicated bandwidth and this bandwidth can be fully enjoyed by all the traffic classes belongs to the mobile operator for their EIR and won't be given to other mobile operators until there is no packet in all the queues when served that belong to that mobile operator.

In implementing the simplified service model, plain QoS can't control the excessive bandwidth left by CIR in such a small granularity as H-QoS and it treats all the mobile operators equally in allocating the excessive bandwidth. Still, each mobile operator may enjoy expected EIR because the number of yellow packets coming into the queues is proportional to the configured EIR of the services or traffic classes of each mobile operator. Moreover, the network operator can set the weight in WRR (or Deficit-based WRR) of the queues in the ratio of the configured EIR of different traffic classes or services. In the long run the actual excessive bandwidth that each mobile operator gets is close to its configured EIR. In statistical sense, plain QoS achieves roughly the same goal of H-QoS.

In the case where mobile operators want to run service-layer OAM (Operations, Administration and Maintenance) over their services, the existing OAM tools such as Fault Management defined by 802.1ag and Performance Measurement/Monitoring defined by Y.1731 can continue to be used in the network sharing environment. These tools are able to monitor the desired segment of the network designated by MEP (Maintenance Entity Point) and MEG (Maintenance Entity Group) Level configured by the service provider.

Copyright 2015 DragonWave Inc. All rights reserved.

6. Network Sharing Migration to SDN

Thanks to the fast progress in electronics technology, H-QoS is becoming more popular and affordable in networking devices. For example, DragonWave's Harmony Enhanced and Harmony E-Band ODU along with the next-generation Harmony IDU have H-QoS capabilities now. Moving forward, network operators will add these new devices to the network and they need to manage the network mixed with both installed devices that only support basic QoS and the new devices that support H-QoS. Establishing the services over this kind of multi-generation or even multi-vendor network is very complicated if configuration is done manually on each device. In case service models are evolved with more parameters whose implementation requires H-QoS, the service provisioning system must take into account the capability of the devices on the selected path for the service, adding another level of complexity.

This impediment to the automation of service provisioning can be solved by software-defined network (SDN) in the future.

In SDN architecture (Figure 5), the data plane of the network is decoupled from the control plane. The control plane is run by an external server that controls the forwarding path of the packets while packet forwarding is on data plane implemented by the infrastructure layer. Between the two planes is the Control Data Plane Interface. In the north bound, there is the application layer that allows network operators or even users to re-program the network functions for the applications via the standard APIs. This programmability is crucial for service automation.

In the context of network sharing, network operators or mobile operators can program the service models on the application layer, which calls the functions of the control plane in order to push the service path and the configured QoS policy to the data plane. The network devices on the data plane have the virtualization layer that abstracts the networking functions according to the SDN specifications such as OpenFlow so that the differences between the new and the installed network devices can be represented as different network resources (the number of queues, the number of QoS hierarchies, etc.) in the network devices. This way, the SDN control server is able to select the best path that meets the demands of the service defined on the application layer. Network operators do not need to care about the differences between the network devices.

Copyright 2015 DragonWave Inc. All rights reserved.

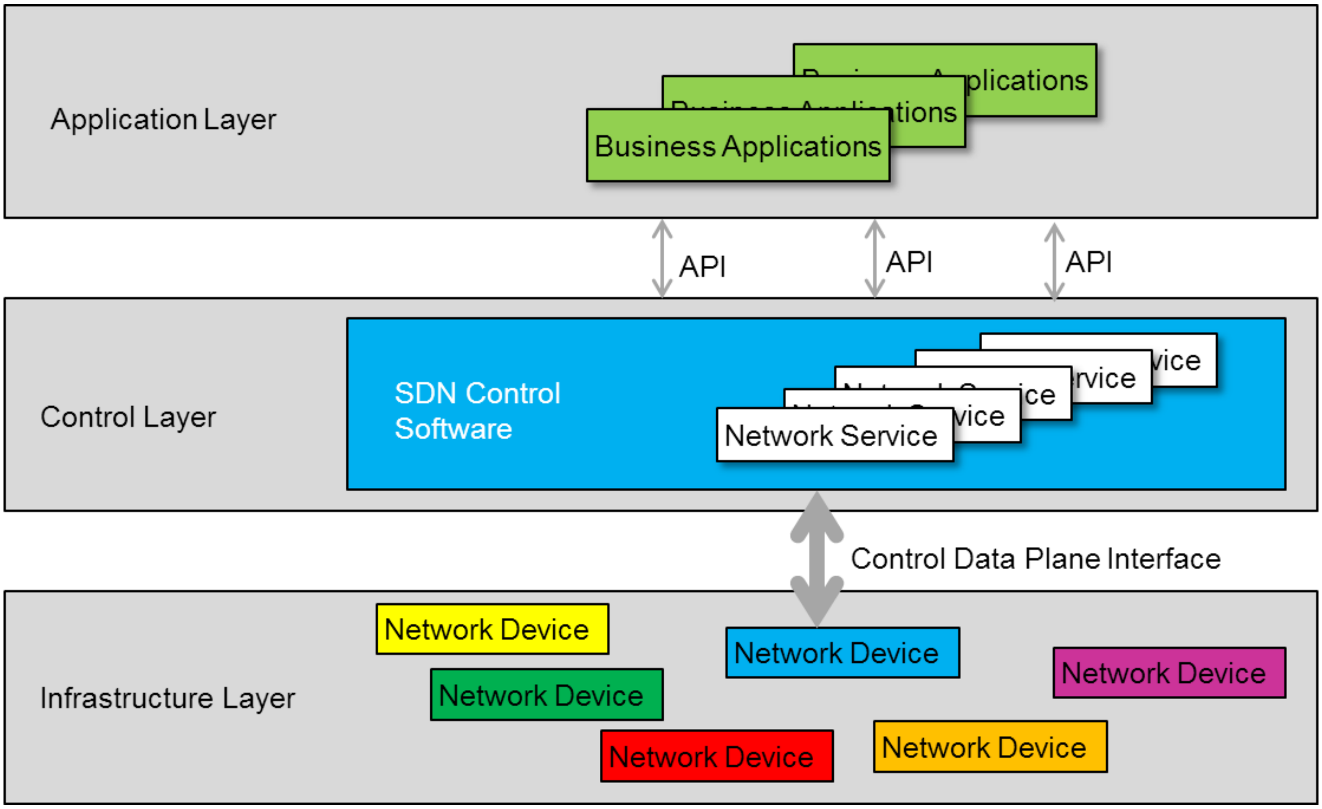

Figure 5: SDN Architecture

Since the SDN interface in the network devices is an add-on function and the traditional control plane of the device (such as xSTP, or MPLS signaling protocols, or IP routing protocols) remains in place to perform traditional switching or routing, SDN-enabled devices can support both OpenFlow traffic and non-OpenFlow traffic, with the mechanisms for determining the path that each traffic flow should follow. Therefore, the migration to SDN is smooth.

DragonWave's Harmony ODUs and IDUs along with its network management and service provisioning system NetViewer adopt the SDN-based software architecture, enabling network sharing and service automation over a heterogeneous network.

7. Summary

By simplifying service models, network sharing can be achieved over the existing network that supports basic QoS functions. More sophisticated service models require advanced H-QoS that is becoming more affordable than before. Still, it takes time to replace the installed base with the new products and network operators will be facing a complicated multi-generation and multi-vendor environment. A practical solution is to add the support for SDN to the network so that the application layer does not need to care about the differences on the infrastructure layer and service automation can be implemented on the application layer in an easy way.

Copyright 2015 DragonWave Inc. All rights reserved.