ELIMINATING OPEN ACCESS WOES WITH INTELLIGENT CARRIER-CLASS EDGE ROUTING

Gerry White, Chief Technology Officer
RiverDelta Networks, Inc.

Abstract

Cable operators can embrace Open Access as a wholesale revenue opportunity and as an opportunity to create closer bonds with both residential and corporate subscribers. They can eliminate Open Access woes by deploying intelligent, carrier-class routing at the edge of the cable network to isolate and police individual traffic flows.

Operators can break down traditional barriers to Open Access by implementing carrier-class routing with sophisticated per-flow queuing to support multiple providers of content, applications, and services over shared cable networks. They can use advanced technologies such as MPLS and policy based Routing to deliver end-to-end QoS across the access network and the core networks of multiple revenue-sharing partners.

BREAKING DOWN THE BARRIERS TO OPEN ACCESS

Best-effort data services provide limited revenue growth potential for cable operators. However, by implementing end-to-end Quality of Service (QoS) controls, operators can expand the customer base by offering a wide variety of business and residential services, build increased customer loyalty offering bundled services supporting voice, data, audio, and video traffic, and create multiple revenue streams for the Hybrid Fiber Coax (HFC) network.

However, to fully realize the benefits of Open Access, operators must gain the ability to isolate each traffic flow and police the network infrastructure to ensure that traffic flows are in compliance with established Service Level Agreements (SLAs).

While Open Access was once viewed as a problem by cable operators, many operators today are realizing the opportunity to accelerate subscriber growth, provide a rich and more complete set of value-added services, and establish profitable revenue agreements with third-party providers.

It is important to carefully define Open Access terminology to understand the technical demands that Open Access imposes on the network. The term "Open Access" means the ability of a cable operator to allow multiple providers to deliver services across the shared cable access network. The term "services" should be interpreted broadly to include content, applications, and other profit-making flows of information.

Operators therefore face the challenge of supporting providers that in the past may have more resembled competitors. But the key to the successful delivery of Open Access is to recognize that the more services are made available to the subscriber, the broader the penetration of cable access networks. Subscribers will select cable as the preferred medium for network services, which in turn increases the total market opportunity for cable operators.

MSOs can continue to deliver their own value-added services, but they will be able to create incremental revenue streams by opening up infrastructure to third-party providers—and gain a percentage of revenue from each new service delivered over the shared network. Open Access does not

trivialize the role of the operator as a mere provider of transport; it creates opportunities for complex and creative business models that enable multiple revenue streams and new opportunities to increase both market share and profits.

Operators need to recognize the diverse business models operators can build to support Open Access. They can deliver IP network services—such as transport, naming, routing, etc.—to enable a basic Internet Service Provider (ISP) service offering. They can also continue to deliver services such as Internet access and Web hosting directly to subscribers.

Operators can create tiered data services to enable Gold, Silver, and Bronze offerings of a given service. This approach allows MSOs to charge premium prices for premium services. Similarly, once they've deployed the technology to support tiered services, they can also allow third-party providers to offer tiered services. This creates opportunities for operators to gain increased wholesale revenues from each service provider partner.

Operators—or their partners—can also deliver enhanced services such as Voice over IP (VoIP), and they can allow Application Service Providers (ASPs) to lease business applications over broadband access networks.

The common denominator of all of these service opportunities is the ability to deploy QoS enabled carrier-class routing at the edge of the broadband access network. Without the ability to isolate and police individual traffic flows, operators lack the control over network resources needed to support multiple providers.

NETWORKING REQUIREMENTS FOR OPEN ACCESS

It is instructive to identify the requirements of the MSO's network within the context of Open Access applications.

In the current Internet access model, a service provider manages IP addresses for subscribers and statically or dynamically allocates unique IP addresses that fall within the address space of that provider. Traffic to the subscriber is then routed to the provider's network based on the IP destination address (which is within the providers address space). Traffic from the subscriber is routed to the desired destination via the provider's network.

Open access will involve assigning IP addresses from the address spaces of multiple providers. These may be delivered from multiple servers or multiple address ranges supplied from a single server. . Traffic to the subscriber can still be routed to the provider's network based on the IP destination address. However routing traffic from the subscriber is more complex, as the path will be dependent on the subscriber's service provider as well as the destination address; e.g. should a packet addressed to a given web site, be routed via ISP1's network or ISP2's network?

The service provider can be determined by the source address of the IP packet so that all the required information is present in the packet. In order to operate in this environment systems must be able to make decisions based on multiple fields in the packet header in real time.

In the Open Access model, services are provided to subscribers from multiple sources. Each provider, therefore, must be able to ensure that their services are working correctly for all subscribers. This is a non trivial problem since each service is based on QoS-enabled IP transport over a shared HFC infrastructure rather than over dedicated

PSTN lines. Effective service management requires MSOs to develop sophisticated QoS and availability parameters and offer third-party providers the abilities to test, quantify, and troubleshoot service delivery of multiple services to all of their subscribers — end-to-end, from the cable modem to the backbone network of each provider.

Quality of Service

Operators require the ability to create and enforce a hierarchy of nested QoS domains within the HFC infrastructure (provider, subscriber, service) which requires sophisticated, high-performance packet filtering and forwarding. Open Access also requires the ability to support end-to-end QoS guarantees across both HFC and third-party networks using industry standards such as Multi-Protocol Label Switching (MPLS) and Diff-Serv.

To provide subscribers and third-party providers with predictable levels of service, it is essential that traffic flows be contained at each level of the QoS hierarchies. Overload or misbehavior within the HFC network by any given provider must be contained within the network resources committed to that service provider — and not be allowed to impact other providers sharing the network. The profitable delivery of Open Access requires advanced isolation functionality to prevent unscrupulous or naive providers from massively overselling their service to the detriment of all other providers on the HFC network

Similarly, each service provider must be able to isolate each of its subscribers so that none of them can impact other subscribers sharing a common domain. In addition, any overload or misbehavior within a subscriber service should be isolated to that particular service. For example, a CLEC offering Internet access and voice services must be able to prevent a subscriber's web traffic from impacting that same subscriber's voice calls.

Policing of traffic flows is required to provide the necessary isolation and enable SLA enforcement. Operators need to police traffic flows to make sure that each service provider is compliant with documented SLA parameters. They need the flexibility to ensure that knowledgeable users do not take advantage of the network QoS mechanisms to obtain services for which they have not paid. Traffic that exceeds SLAs should be handled according to SLA policies that determine whether excess flows should be dropped, assigned lower priority levels, or routed at incremental costs.

Carrier-Class Routing

The transition from providing basic Internet access to offering a variety of services from multiple providers moves the MSO from an entertainment provider into a communications carrier. This requires next-generation system that are architected for "carrier-class" reliability, which is usually defined as systems that deliver "99.999%" reliability, which is less than six minutes of unscheduled downtime in a year. Meeting the carrier-class requirements of critical services requires high-levels of redundancy to ensure non-stop operations in the event of a failure of any system component.

Operators must be able to efficiently scale HFC infrastructure to accommodate increased demands for new services and content. This requires next-generation equipment with faster forwarding engines, increased port density, and greater abilities to add network ports so that operators can increase network capacity to support revenue streams from multiple service providers.

As providers aggressively develop partnerships with ISPs and content providers, demand for cable services will escalate. The ability to maximize use of scarce real estate at the distribution hub and regional headend

requires next-generation platforms that provide higher-density RF termination and eliminate the need for external equipment, such as up converters and LAN switches.

Services such as VoIP or streaming multimedia require consistently high-levels of performance, and wire-speed forwarding is required to support a vast array of enhanced services offered by third-party providers. Next-generation, carrier-class edge routing platforms are needed to provide the scalability, density, reliability, and performance needed to support Open Access. Operators need to be able to ensure that carrier-class platforms deliver the guaranteed SLA requirements that they have committed to both provider partners and to subscribers.

Service Provider Selection

A subscriber should be able to select from multiple providers based on the competitive nature of their offerings, such as Internet access from a selection of ISPs, video service from the MSO, and voice service from Competitive Local Exchange Carriers (CLECs) or InterExchange Carriers (IXCs). Both residential and corporate customers should be able to select services either on a subscription or pay-per-use basis. This requires flexible, open systems provisioning and management combined with sophisticated, high-performance routing.

Operators need to support advanced SLA parameters such as maximum bandwidth allocation, minimum bandwidth guarantees, bounded delays, and bounded jitter. They will need the ability to define QoS parameters both statically (e.g., Gold/Silver/Bronze services) and dynamically (e.g., for services such as voice call set-up). At a minimum, operators need the QoS capabilities of DOCSIS 1.1-based equipment, but they also need features beyond these standards to enable enhanced services over both HFC and service provider networks.

Metering/Billing/Reconciliation

Allowing multiple service providers to operate over a shared access network requires robust features for reconciliation and billing. Detailed accounting information needs to be maintained on a per-flow basis to ensure that SLAs are enforced, and the sophistication and complexity of accounting can vary dramatically.

In the simplest case, a provider could define an SLA and the MSO could implement a policing mechanism to ensure that it is not exceeded. However, in most applications both the provider and operator will want to meter the SLA to ensure conformance. If subscribers have access to pay-per- use services such as long-distance phone calls or videoconferences, then the MSO needs to offer metering services that can support dynamic billing. Billing models based on both time-of-use and traffic volume is required with an event-driven mechanism used to initiate and terminate metering at wire speed.

OPTIONS TO PROVIDE OPEN ACCESS

In theory, an MSO could create and maintain multiple RF channels to carry traffic for each provider. Lack of sufficient RF frequencies and the requirement to duplicate CMTS systems per provider render such a solution impractical.

Fortunately more viable alternatives are available. These can be classified into two general categories. Tunnel based solutions in which subscribers are tunneled back to a centralized subscriber management platform responsible for implementing traffic policies and routing subscribers to the appropriate provider networks.

Policy-based routing solutions in which the edge router/CMTS system is responsible for implementing traffic policies and for routing subscriber traffic to the appropriate provider network.

TUNNELING :A CIRCUIT BASED APPROACH

Generally, tunneling is used for dial-up Internet and DSL access. Subscribers connect to a network access server using a modem connected to the public switched telephone network (PSTN) or a DSL circuit. In these networks, a subscriber management system located inside the network manages the traffic flows.

Traffic flows reach the subscriber management system via a tunnel mechanism such as a Point-to-Point Protocol over Ethernet (PPPoE) or Layer Two Tunneling Protocol (L2TP) tunnel built on top of the generic network infrastructure. Once the flow reaches the subscriber management system, the system terminates the tunnel, examines the data received, implements QoS and policing and directs the traffic flow to the required application server.

This mechanism requires client software on the host system to initiate the subscriber end of the tunnel, which can present an ongoing support problem. The most serious drawback to tunneling is that it hides the content of the flow. Because the CMTS cannot recognize what the tunnel carries, the HFC access network cannot use the QoS built into DOCSIS 1.1. Application-based QoS is not available to traffic within the tunnel. Without the ability to give voice or video traffic higher priority, operators will have difficulty meeting the performance guarantees promised for these services.

The “bandwidth tax” associated with tunneling is also significant. Tunneling requires additional headers on top of the DOCSIS protocol. This approach wastes bandwidth in the access segment, where network capacity is most strained.

Because tunneling requires that the subscriber management system be located inside the network, operators must place applications servers even deeper in the network. Thus they negate the benefits gained from moving content for high-bandwidth services closer to the user.

Finally, tunneling deprives the cable operator of one of its most powerful weapons against its DSL competitors--the “always-on” connection. Before a user can have access to even the most basic e-mail services, the tunnel must be established. It is therefore difficult to deliver push services like newscasts. Likewise, multicast services are problematic, because MSOs must convert a multicast into multiple unicast messages at the subscriber management system, which further hogs scarce bandwidth.

POLICY-BASED ROUTING

Policy-based routing differs in that the router manages traffic flows at the edge of the network. The router looks at multiple fields within packets to determine the appropriate routing and QoS.

Each user is provided with an IP address in an address scope associated with his/her selected service provider. Packet routing is partially determined by looking at the source IP address, understanding to which service provider partner the IP address belongs, and then routing the traffic to that partner for their handling. Such examination allows the router to implement more sophisticated QoS policies than are possible by simply looking at the data’s destination address.

Two variants of policy based routing can be considered for cable networks, a centralized model with the policy router located at the regional head end (or other convenient central location) and a distributed model in which the policy routing function is moved to the edges of the network (e.g to a distribution hub).

Both distributed and centralized architectures require a DHCP address management system and a policy based router. The distributed solution places the

policy router with the CMTS and uses an MPLS based metro or wide area network to connect the policy routers to the service provider networks.

DHCP Server

A DHCP server provides each user with an IP address in an address scope associated with his/her selected service provider. This address is used to identify the service provider to which the customer has subscribed. The DHCP server is well understood technology and need not be described further.

MPLS Virtual Networks

MultiProtocol Label Switching (MPLS) is a standard under development by the IETF for efficiently switching IP traffic over IP or ATM core networks. MPLS adds a label to IP packets which instructs network routers and switches where and how to forward the packets.

Today's conventional routers analyze IP packets at each hop in the network, which is a time-consuming process. With MPLS, an intelligent edge router (Label Edge Router or LER in MPLS terminology) looks at the header of the first packet in the traffic flow. Based on the header's contents, the router applies a label to that packet and all subsequent packets in the flow. This label determines where to send the entire data stream and the QoS policies to apply. For example, the label may indicate that the flow contains a Voice over IP (VoIP) call destined for a particular voice service provider. The label would dictate that every packet be placed in a low-delay path with guaranteed delivery to the provider POP. This path is known as a Label Switched Path or LSP. The routers or switches in the core of the MPLS network (Label Switch Routers or LSRs) do not examine the IP headers of the packets but instead switch the packets based on the appended MPLS labels.

Labeling packets at the network's edge eliminates the processing traditionally performed in the core. The core network can therefore focus on switching traffic to its destination as quickly and efficiently as possible. Bottlenecks in the core are reduced or even eliminated.

MPLS enables cable operators to offer a variety of services over shared network infrastructure with differing QoS requirements without overloading the core network with unnecessary processing.

For a distributed Open Access architecture a virtual network is created for each service provider between the service provider points of presence (POPs) and the edge router / CMTS at the distribution hubs. This virtual network is based on MPLS technology and consists of a mesh of MPLS label switched paths (LSP's) which are set up between the CMTS and the POP. Each LSP provides a traffic engineered path over the shared metro or wide area network transport.

CMTS / ER

The CMTS/ER located at the distribution hub looks at multiple fields within packets to determine the appropriate routing and QoS. Packet routing is partially determined by looking at the source IP address, understanding to which service provider partner the IP address belongs, and then routing the traffic to a specific LSP in the virtual network of the partner for their handling. Such examination allows the router to implement more sophisticated QoS policies than are possible by simply looking at the data's destination. The CMTS/ER is functioning as an MPLS label edge router in this case.

The policy information required by the edge router to enable these complex forwarding decisions can be disemminated

using policy extensions to existing routing protocols such as OSPF and BGP4. BGP4 is the preferred routing protocol used for connection between autonomous systems and as such is a favoured approach for the policy router.

For policy-based routing to work in a cable environment, operators need to deploy a high-powered QoS capable routing engine used in combination with a DOCSIS 1.1-capable CMTS in the distribution hub. The CMTS/ER should enable QoS to be maintained from the cable modem to the service provider POP. Thus it must implement QoS on the HFC network by mapping IP flows to DOCSIS 1.1 service flows; it must maintain the QoS through the CMTS / router and then map these flows to MPLS label switching paths with the desired QoS characteristics.

Upstream packets from the HFC network are scheduled by the CMTS/edge router according to DOCSIS 1.1. First, the packets are classified by the cable modem, which requests transmission on the appropriate DOCSIS flow. At the CMTS, the packets are re-classified based on filters and then QoS policy can be applied. Each flow can be assigned its own queue to ensure QoS is maintained through the CMTS /ER, and then packets are forwarded to the LSP required to reach the network of service provider partners based on fields in the packet header such as the source IP address.

Downstream packets are received from the (traffic engineered) LSP, mapped into downstream flows based on the IP header fields and scheduled for transmission onto the DOCSIS downstream channel. By providing sophisticated queuing and scheduling mechanisms QoS may be maintained in this direction also.

Advantages of Policy-Based Routing

Unlike tunneling, policy-based routing is designed for a broadband IP infrastructure. Because the router can look at the individual application flows, it can extend the capabilities of DOCSIS 1.1. It can assign QoS and routing policies based on parameters such as service provider, subscriber, and application. Per-flow queuing enables the router to isolate the traffic of different services and different providers at the edge of the DOCSIS network.

Because policy-based routing conforms to IP standards, current features like transparent IP multicast can be supported. Policy-based routing can also take advantage of future developments in IP standards such as the rollout of MPLS.

DISTRIBUTED vs CENTRALIZED POLICY ROUTING

As operators evaluate the networking requirements for Open Access, they should consider whether the policy routing intelligence should reside at the edge of the network or at a central location.

A centralized solution leads to inefficient bandwidth use. All traffic must be routed to the policy router for treatment—no matter its destination. As subscriber penetration of new services increases, this approach can cause bottlenecks. With intelligent edge routers, operators can contain local traffic on the network and establish the optimum routing path for internetwork traffic.

Pushing intelligence to the edge also allows operators to move application and content servers closer to subscribers. This is a plus for customers because they'll see increased performance for which both operators and partners can realize premium pricing.

The distributed model has much better scaling properties as the work intensive policy decision making is distributed. It also simplifies provisioning multiple connections to provider networks from the MSO regional network to support redundancy and load sharing..

A distributed solution is shown in the figure below.

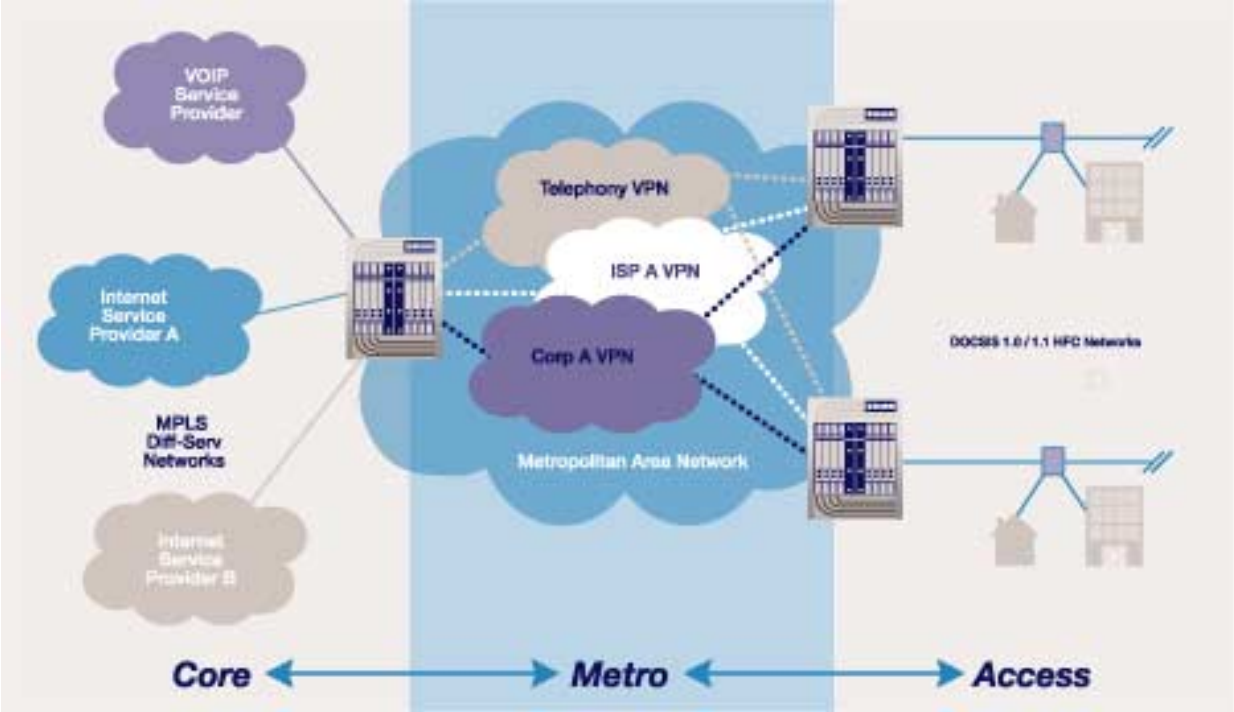

Distributed Policy Based Routers with MPLS based Virtual Networks

REQUIREMENTS FOR A POLICY ROUTER AT THE NETWORK EDGE

Operators need to be able to deploy intelligent policy based routers at the edge of the network to isolate and police traffic flows. These next-generation edge routers need advanced intelligence so they can classify, manage, and police the traffic from the cable modem to the core networks of service provider partners. Both third-party providers and subscribers will demand SLAs based on guaranteed QoS levels, and edge routers must be capable of implementing these QoS guarantees end-to-end.

Carrier-class routing protocol implementations are essential because the edge router should be capable of peering with other ISPs to make Open Access a reality. A full suite of unicast and multicast routing protocols is required to allow interoperation with peer routers, (e.g. RIP v1; RIP v2, OSPF v2, BGP4, IS-IS, DVMRP, PIM-SM). Providing a carrier grade implementation of the routing protocols requires more than a basic implemetation of the minimal functions required for comformance. As well as offering a full feature set it must provide a robust, highly available solution which is resilient to network errors and attacks The

routing software implementation must scale in terms of numbers of routes, interfaces, and peering relationships to support expansion as new services, subscribers, and providers are added. It must also provide support for operation staff to detect and resolve network routing issues.

The edge router must apply policy functions on a per-flow basis and must be able to provide a guaranteed minimum rate per flow to enforce SLA commitments.

The router must isolate the traffic of individual providers, subscribers, and applications. When IP traffic from a provider exceeds its SLA, the edge router implements predetermined policing policies to ensure that each flow receives at least its minimum guaranteed bandwidth. When congestion occurs, the router should drop packets from misbehaving flows instead of dropping traffic that is operating within its SLA.

Operators will need sophisticated accounting and metering systems in a multiprovider environment because wholesale providers will of course require proof that they are receiving committed performance levels. Again, an intelligent edge router can provide these statistics. Because operators will be peering with other ISPs and service providers, it is critical that edge routers conform to the highest standards of interoperability. Partners will require the ability to control their portion of the network; so flexible standards based network management is essential.

Implementations require a high-powered QoS routing engine and a DOCSIS 1.1-compliant CMTS. QoS on the HFC network is provided by mapping IP flows to DOCSIS 1.1 service flows based on contracted service levels, and the QoS on the metropolitan network is provided by mapping IP flows to traffic engineered label switched paths in the MPLS networks. Thus cable operators can guarantee performance to their wholesale partners and generate additional revenue streams.

Edge routers are the transition point from the HFC access network to the regional backbone. They identify and classify traffic flows, apply QoS, implement admission control and efficiently forward traffic to its destination—which can be the core network of one of multiple providers.

By applying this intelligence at the edge of the cable network, MSOs can provide end-to-end QoS across the HFC network and across the backbones of multiple revenue-sharing partners.

AVOIDING OPEN ACCESS WOES

The use of hierarchical per-flow queuing and carrier-class edge routers allows cable operators to benefit from Open Access and ensure maximum control over network resources. By selecting next-generation edge router/CMTS platforms, operators can welcome Open Access as an opportunity for new revenue streams.

They will be able to bind subscribers to the cable network, improve bandwidth utilization, and increase profits. Operators will be able to classify and treat individual traffic flows and deliver QoS guarantees across access, metropolitan, and core networks, and they will be able to allow multiple revenue-sharing partners to offer diverse portfolios of services that will create tighter loyalty to cable networks for both residential and corporate subscribers.

Policy-based routing has not been widely deployed up to this time since legacy routers lack the performance, scalability, and per-flow processing necessary to implement it effectively. They can't perform source based, content-aware routing because the per-flow packet classification and QoS required is beyond their processing capacity. But with today's high-powered silicon and advances in QoS theory, next-generation edge routers can

examine individual traffic flows and provide the forwarding and QoS functions required at wire speed.

ABOUT THE AUTHOR

Gerry White is Vice President and Chief Technical Officer for RiverDelta Networks, which designs, develops, and markets intelligent edge routers for broadband access to optical networks. He is the co-author of several patents and articles on data communications technology and he can be reached at gwhite@riverdelta.com. RiverDelta Networks can be found on the Web at www.riverdelta.com.