




RHK

IPTV multicast: many different ways to skin the cat

04 April 2006

Mark Seery



www.ovum.com

| | |
|--|----|
| Table of Contents..... | 1 |
| IPTV multicast: many different ways to skin the cat..... | 2 |
| Multicast: what is it and why use it?..... | 3 |
| IP multicast routing to BRAS..... | 5 |
| IP multicast routing to IP access aggregation..... | 7 |
| IP multicast routing to VPLS access aggregation..... | 9 |
| MPLS multicast switching to BRAS..... | 12 |
| Multicast VPNs..... | 13 |
| Impacts on optical transport..... | 14 |
| Traffic engineering..... | 14 |
| Fast reroute..... | 14 |
| OAM..... | 15 |
| Scalability..... | 15 |
| High availability..... | 15 |
| Life after multicast..... | 16 |
| Summary and conclusions..... | 16 |



IPTV multicast: many different ways to skin the cat

While the future of video may be unicast, today multicast is the mainstay of video distribution to servers and users. As if it were not difficult enough for network operators to work through an array of options for supporting their own video services, another set of IP VPN-based options has to be considered when supporting wholesale customers. In general, PIM (protocol independent multicast), an IP-based multicast mechanism, is at the core of most IP-based video multicast networks. How video is then distributed from the core to the end user varies by operator, with VPLS (an MPLS-based approach) gaining recent attention. Also at the core, a purely MPLS-based mechanism may emerge, with one large PTT already considering this option. Today the IP multicast router plays a key role in many IPTV networks, but the emergence of MPLS-based mechanisms is still plausible.

Multicast significantly reduces the capacity needed to support a broadcast / linear programming model of video distribution (i.e., where hundreds of prescheduled and fixed start/stop-time programs/channels are offered). In analyzing various approaches, we make the following observations:

- This subject is too complex from a technical perspective for us to comprehend completely in a single report.
- Financial sector usage of IP multicast where there are reportedly millions of receivers suggests the mechanism is inherently scaleable for even mission-critical applications. Service providers will need to consider this within the context of network elements performing a range of VPN, Internet access, and subscriber management functions as well.
- We believe that the transition from very lightly loaded and highly over-subscribed DSLAM links to gigabit Ethernet links supporting consistently high video traffic loads creates a general incentive for operators to have at least some subscriber management awareness in the aggregation network. In addition, we believe the volume of multicast signaling and the amount of traffic between the multicast edge nodes and the access nodes can be reduced by multicast intelligence in aggregation nodes. The need to keep this additional intelligence cost-effective, and the fact that some operators will also want to aggregate Ethernet traffic supporting business services, supports the argument that VPLS as an aggregation technology is more likely to increase rather than decrease. This scenario could conceivably change if full-featured multicast-enabled subscriber management routers become more cost-effective, or a VPLS alternative like PBT (provider backbone transport) comes to market with a strong value proposition.
- Native MPLS multicast is rare today, but is worth keeping an eye on, as there may be operational synergies with pure MPLS unicast networks.



Protocol independent multicast (PIM) is a procedure created by the IETF to support the creation of multicast trees in IP networks. PIM—often referred to as a routing protocol, though technically not (it is actually a procedure for using topology information created by unicast routing protocols)—is used between multicast routers that form the core of IP/MPLS routed networks. IGMP (Internet group management protocol) is used by end users (multicast receivers) to signal to the closest multicast router what video channel is desired. However, a practice of intercepting IGMP messages has evolved such that any network element, regardless of technology implementation, in between the end user and the IP multicast router can read the information in the IGMP packets and create multicast trees within non-IP routed networks such as VPLS (virtual private LAN service), Ethernet, and MPLS.

In this report, we examine some of the combinations of core and aggregation approaches as well as explore a multicast mechanism based on MPLS that does not require any IP routing to be run in the core of the network. Service providers then are left with numerous options. We can clearly see that just as VPLS has been growing for the transport of Internet traffic and to support business services, its use in IPTV networks as a video multicast transport is also growing.

As service providers increasingly look to base converged networks on MPLS for unicast traffic, the possibility of MPLS multicast increases over time. Today, point-to-multipoint MPLS circuits (label switched paths—LSPs) are most strongly endorsed by Juniper, leaving some operators with the perception that this may not be a good approach because it lacks broad industry support. Ovum-RHK believes BT's endorsement of point-to-multipoint LSPs will lead numerous vendors to support this approach, even if some do not support it vigorously in their public messaging; additionally, two vendors have told us they intend to support point-to-multipoint LSPs. Thus, two to three years from now, this could be a legitimate alternative to an IP-multicast routed core for more than just demanding carriers such as BT. In MPLS, network operators see a networking paradigm that is operationally similar to other models they have used in the past, such as ATM and TDM. Over time, operational familiarity may exert a greater influence on carriers' decision-making than the technical pros and cons of different approaches.

While there are many approaches to building a multicast network, we believe the following four are emblematic of most IPTV deployments:

- IP multicast routing to BRAS
- IP multicast routing to IP access aggregation
- IP multicast routing to VPLS access aggregation
- MPLS multicast switching to BRAS

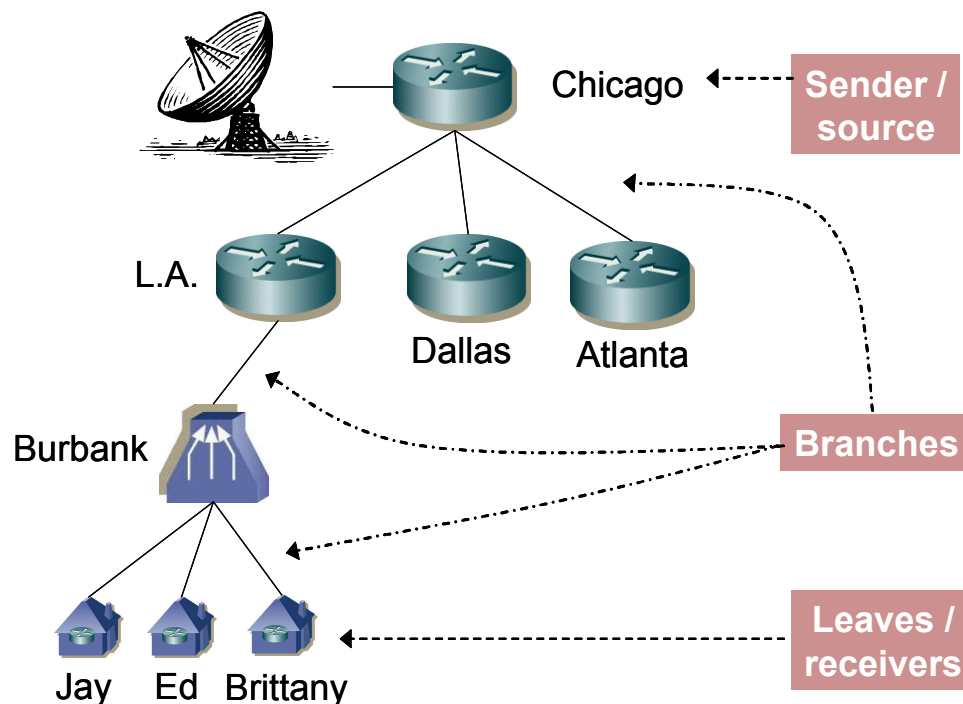
Multicast: what is it and why use it?

There are two predominant styles of video delivery being contemplated. The first is broadcast video or linear programming, which mimics the predominant distribution in cable and satellite networks, where all receivers receive all channels in a package,



and the receiver tunes to the channel that is to be viewed; the second model is video on demand (VoD), where a receiver selects content from a library of 3,000 hours+ of programming (some reports suggest Comcast will be up to 10,000 hours by the end of the year) and only receives the content that was selected. The same equipment supports both models. Today, the first model benefits most from multicast technology—VoD server content can be multicast to the server when initially sent, and unicast from the server as requested by end users, and while this reduces capacity requirements, it does not reduce bandwidth by the same amount as multicast for broadcast programming.

Figure 1 An example of a multicast tree



Source: Ovum

In the first model, transmitting 200–300 channels of video content to all subscribers could easily translate to 500 Gbps or more of traffic in an area of 100,000 homes. Therefore, instead of sending all content to all homes, multicast sends one copy of the content to all parts of the network that need it. This dramatically reduces the traffic requirements of a large package of content over the core/metro transport network. Access equipment can achieve further gains by only sending the content being currently viewed to each house. Both of these gains are achieved through multicast routing protocols in the core/metro packet transport network and by intercepting multicast requests in the access network. Instead of requiring the transport network to support hundreds of gigabits of capacity, a large package of standard-definition and high-definition channels can be transmitted for a few gigabits.



The details of how multicast works varies from protocol to protocol, but the concepts are similar. There is a content source, a video / TV channel, for example. Receivers indicate that they want to receive that content, and a path from the receiver to the sender is created throughout the network by storing state in each network element: specifically, when content is received it is copied to interfaces that have other network elements or content receivers that have expressed an interest in the content. The content is copied to output interfaces only once, instead of for each receiver (for example, each set-top box).

Figure 1 above depicts an example of a multicast tree. If the viewers in Burbank—Jay, Ed, and Brittany—all want to watch “American Idol” at the same time, one copy of the program is sent from Chicago to Los Angeles and one copy is sent from Los Angeles to Burbank. No direct session is established among Jay, Ed, and Brittany and the video source, and neither do they each get their own individual copy of the video stream—they all get the same stream. In this case, the traffic load has been reduced from 6 Mbps to 2 Mbps (assuming that the video stream requires 2 Mbps as is forecast with MPEG4 systems currently being deployed). As more houses are added, the capacity/bandwidth savings increase. One of the most scaleable aspects of multicast is that you pay for the infrastructure once, and as new subscribers are added, the demand on the core of the network does not change significantly (the greater the population of receivers the more chance of variation among what channels are being watched and received, which leads to small increases).

There are clear bandwidth savings and cost reduction associated with using multicast to simultaneously distribute hundreds of video channels to thousands of houses.

Readers should take particular note that many of the diagrams shown in this report do not properly reflect the underlying fiber / optical transport network. It is important to keep this aspect of the network in mind because it does have significant impacts on some decisions, which we will discuss later in the report.

During the development of this report, it became clear that while the concepts of IP and MPLS multicast are quite simple, there are many subtleties of each approach that we were not able to fully explore in preparation for this report. As a result, further research may be required to more fully understand the technical details. Therefore we would direct readers to focus on appreciation of the various approaches that carriers might take, and the drivers, rather than using this report as a technical scorecard of which approach is better in an absolute sense.

IP multicast routing to BRAS


While some incumbent network operators are building new access networks to support IPTV, many are leveraging existing DSL networks, perhaps upgrading the type of DSL to provide additional bandwidth. In this case, there are numerous operational benefits from keeping the same basic architecture in place.

Most BRAS (broadband remote access service) routers support IP multicast routing techniques such as PIM (protocol independent multicast). Over the long evolution of



IP multicast, various protocols have come and gone: DVMRP (distance vector multicast routing protocol), PIM dense mode (PIM-DM), and the any source multicast (ASM) model of PIM sparse mode (PIM-SM), to name a few. All of these have fallen out of favor either due to scaling concerns (DVMRP and PIM dense mode) or in the case of PIM-SM/ASM it may be argued that this approach may have more benefits for many-to-many multicast than for one-to-many multicast applications such as IPTV. As a result, we don't discuss any of these other approaches in this report. PIM-SM/ASM has been deployed to support video applications, and they need to be maintained, but it was our sense from interviews we conducted for this report that the industry is moving toward the source-specific multicast (SSM) model of PIM-SM (which we refer to in this report as PIM-SSM). Both the comments we received from some vendors and also discussion within the IETF indicate that opinions differ about whether the PIM-ASM model may have some significant technical issues. Ovum-RHK is undecided on the technical merit of arguments against PIM-ASM; however, we think the winds of favor have shifted to PIM-SSM and therefore we have focused on this approach here. However, we were referred to implementations of significant size (many millions of receivers) in demanding verticals such as the financial industry that use PIM-ASM; therefore, there are data points supporting the scalability and usability of PIM-ASM.

Figure 2 IP multicast routing to BRAS



Source: Ovum

As shown in *Figure 2*, IP multicast routing is performed between the core router and the BRAS. PIM controls the IP multicast routing, and a protocol called IGMP (Internet Group Management Protocol) indicates to the IP multicast network the channel a set-top box (STB) wishes to view. All the equipment in between the BRAS and the set-top box that do not wish to participate in the IP multicast routing domain (i.e., implement PIM and support IP unicast routing protocols), but still want to build their own multicast trees can examine IGMP packets (“snooping”) and do so. IGMP snooping is done to lower the cost of network elements by not having to implement a full IP multicast router, but without losing the ability to support multicast topologies.

Historically DSL networks have aggregated DSLAMs by using an ATM switching network. Over time, DSLAMs will migrate to implementing Ethernet framing on DSL links, which will position them to implement Ethernet switching or VPLS within the DSLAM as an internal optimization. Even before that though, as DSLAMs migrate to



Ethernet uplinks, service providers are using Ethernet switches or VPLS-enabled edge routers to aggregate DSLAMs. These aggregation devices do not have to support IGMP snooping, but they may if by doing so they can reduce bandwidth requirements by not sending video channels to DSLAMs that do not have viewers currently watching those video channels.

In Ethernet aggregation architectures, typically there will be one VLAN for all the video / multicast traffic, and per-subscriber VLANs for unicast traffic (Internet, gaming, VPNs, etc.). Some operators use individual VLANs for a specific stream of content.

IP multicast routing to a BRAS is a good architecture for network operators focused on residential services, where the combination of multicast and unicast (including the growing amount of video on demand, which will eventually surpass multicast traffic) does not exceed the capacity of the BRAS. This architecture is also good for network operators that want to continue to use PPP for Internet access and also want to use PPP for video. Fastweb, a network operator that uses this type of an architecture, has one of the biggest IPTV networks and uses Juniper E-series routers for the BRAS and Cisco 7600s for DSLAM aggregation.

One concern operators have with this approach is that it forces the BRAS to process a greater number of individual multicast streams as it copies video to each downstream access network, and the more central the architecture, the greater the exposure. This also drives additional IGMP processing (which may be done on a centralized route processor) and might force the implementation of IGMP proxy into the DSLAM. A thorough examination of this concern is beyond the scope of this report, but we can say that from a technical perspective, reduction of IGMP flows is a good thing, as is the reduction of multicast stream destinations through IGMP proxy and IGMP snooping in decentralized aggregation architectures. These concerns could be addressed by either an Ethernet-switched or VPLS aggregation architecture; indeed, BRAS vendors have traditionally been content to let this business go to other vendors, though we think this is changing and will continue to change over the next year, with increased competition between distributed BRAS and Ethernet aggregation architectures.

An additional concern with the BRAS architecture is that with the existing PPPoE-based approach, the concept of subscriber is often not present in the aggregation network, and therefore the aggregation network may not be able to enforce subscriber-aware policy. This may lead to greater capacity requirements between the subscriber management point in the network and the access network. (One network operator told us that this is the case at least in some applications.)

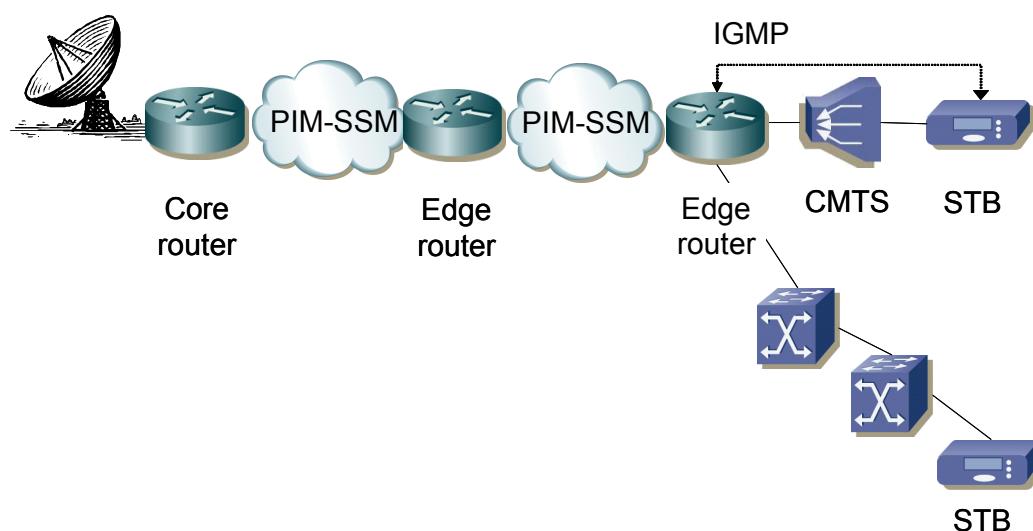
IP multicast routing to IP access aggregation

Technically, IP multicast routing techniques (the combination of unicast IP routing protocols and PIM, for example) do not extend all the way to the service user / STB; they only extend to the "multicast router" closest to the STB. End-to-end IP multicast here refers to the IP multicast routing domain extending to the access aggregation



portion of the network. In this architecture, we refer to IP multicast routing (IP unicast routing and PIM) operating at the access aggregation node.

Figure 3 IP multicast routing to IP access aggregation



Source: Ovum

We expect this architecture to be popular with network operators that do not wish to use MPLS in their networks, and do not need to offer transport services to businesses over the same aggregation network. *Figure 3* shows two types of access networks. The first is based on a CMTS, which is representative of a cable operator / MSO, such as Comcast. The second access network is based on IP/Ethernet switch/routers that are located in the basement and individual floors of high-rise apartments and business locations. A fiber-centric network operator in an Asian country, such as HKBN (Hong Kong Broadband Network Ltd) might pursue this architecture. HKBN users a substantial amount of Cisco equipment, including routers, Ethernet switches, and optical transport equipment.

Comcast is an example of a network operator that has the option of not using optical transport rings in all locations and as a result can explore the option of backhauling over dark-fiber CMTS traffic to a regional or metro edge router. In the future, a product like Cisco's CRS-1—which would have the scalability to aggregate many serving offices, thereby eliminating a layer of aggregation equipment—could play a role in this architecture. Typically, incumbent network operators have preferred to aggregate serving offices (end offices) with SONET/SDH fiber rings, but if that was to change, operators could explore different aggregation architectures; certainly, as BT has stated, one of the best ways to reduce operating costs is to eliminate real estate—specifically central offices.

An important aspect of Comcast's current network designs and the options it is exploring for the future is that the company has so far eliminated the need for MPLS



in its network. Comcast may provide some business services, but it does not have legacy transport services to support. The company is building its own 40 Gbps infrastructure, and residential services will dominate its revenue / service mix. In short, Comcast will have the network capacity and the service mix to pursue a native IP network.

IP multicast routing to VPLS access aggregation

VPLS was created to emulate an Ethernet LAN over an MPLS (or IP/GRE/L2TP) network. The original intent of VPLS was that it would be deployed to support business Ethernet services; however, a number of VPLS proponents have positioned it as an important part of IP video and triple-play networks. Some operators, having already deployed VPLS for business services and/or Internet access aggregation before IPTV was deployed, naturally looked to leverage the architecture they already had in place—an operator like Telefonica, which uses equipment from Juniper, Lucent, and Riverstone, would fit this description. Alcatel and Riverstone are two vendors that have had success with VPLS aggregation approaches to IPTV networks, with Alcatel being involved in high-profile wins such as AT&T (formerly SBC) and Telstra Australia, and Riverstone supplying leading IPTV subscriber networks such as PCCW and Telefonica.

VPLS is not necessarily a total replacement for IP multicast, but simply an aggregation option, such that IP multicast (e.g., PIM-SSM) would be used in the core and VPLS would be used between the core and serving offices. Alcatel informed us that it also supports some networks where VPLS is used even in the core, but we were unable to verify this before publication.

As *Figure 4* shows, using VPLS moves the IP multicast cloud further away from the access network and subscriber; as a result, IGMP packets (multicast joins and leaves) need to travel further into the network before they are terminated on the IP multicast network. To reduce the number of IGMP messages, operators can use a technique known as IGMP proxy, which tracks and discards duplicate requests.

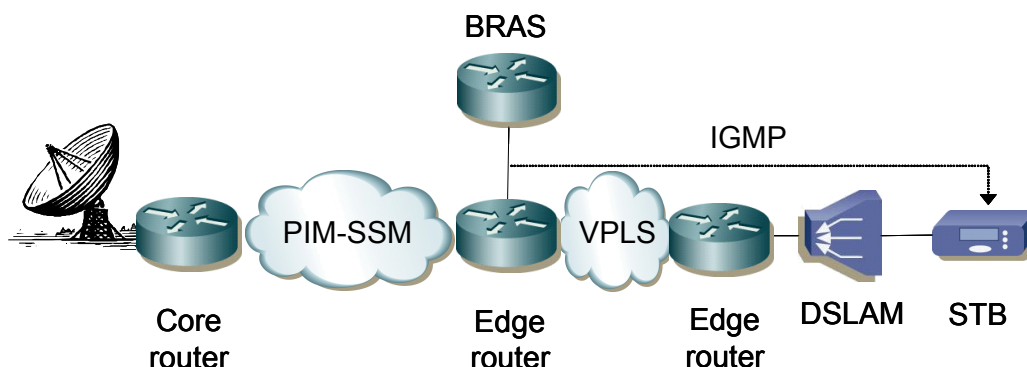
There are a number of reasons for operators to consider using VPLS:

- Common aggregation architecture for business and residential services
- MPLS-based fast reroute
- MPLS-based traffic engineering
- Less expensive equipment

For network operators that want to broadly deploy Ethernet business services, having a VPLS aggregation network provides a common aggregation architecture for business and residential services. A common architecture leads to common OAM methods and procedures, which has operational benefits at that level of the network. Of course, the flip side could be argued: there are now two distinct OAM domains (IP multicast routing at the core and VPLS aggregation which is actually one of the arguments in favor of MPLS-based multicast discussed later).



 Figure 4 IP multicast routing to VPLS access aggregation



Source: Ovum

As VPLS is based on MPLS (VPLS emulates an Ethernet LAN over an MPLS network), it can use MPLS traffic engineering, which is required by some large operators. In addition, operators can use MPLS-based fast reroute to reduce the recovery time for network link failures.

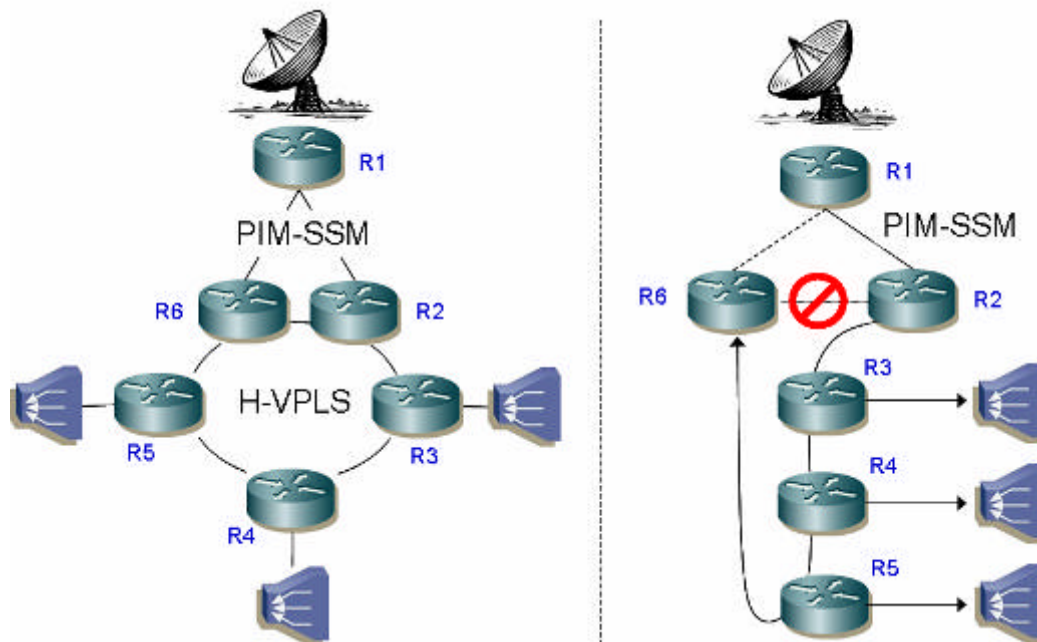
During our research for this report, one large PTT told us that it sees VPLS as a services infrastructure and not a transport infrastructure; also, the PTT considers VPLS a niche service (compared to point-to-point Ethernet) and therefore not a strong enough driver to use as the basis for an infrastructure architecture decision. So, the applicability of VPLS will vary from operator to operator.

Using standard VPLS, especially over a logical ring topology, introduces inefficiencies into the network; for example, multiple copies of the same packet occur because of the internal mesh within a VPLS network, and also because IP multicast packets are mapped to Ethernet multicast frames and by rule have to be broadcast to all participants in the VPLS. In a ROADM architecture, where each serving office may have its own wavelength, the broadcast nature of VPLS would not be liability, but an asset; the liability arises when the entire ring is one logical channel. The additional problem with standard VPLS is that when any new VPLS node is added to the ring, every existing VPLS node has to be touched to update the full mesh (MPLS and PWE3) between nodes, which raises the potential for configuration errors and opex increases as well (which can be diminished with a good management system).

An alternative to VPLS is H-VPLS (hierarchical VPLS), which was first developed as a way of scaling VPLS networks for business services. H-VPLS incorporates the concept of a “spoke,” which is attached to a VPLS rather than a full participant in the internal VPLS mesh (see *Figure 5*). The benefit of this approach in this application is that multiple H-VPLS nodes can be connected to form a “drop and continue” architecture around a ring. In this configuration, one copy of a packet is inserted into the ring and is copied from VPLS node to VPLS node as it passes around the ring.



Figure 5 H-VPLS—physical topology (left) versus logical topology (right)



Source: Ovum

This approach reduces the processing load on the VPLS node doing the initial packet injection (one packet injection instead of one for every VPLS node on the ring), while slightly increasing the processing load on each VPLS node on the ring.

In technologies like RPR, SONET/SDH, and OADMs, traffic not being processed at a node is allowed to bypass it. RPR, for example, was designed to facilitate this and is a true “EAST-WEST” media access control (MAC) layer. Whenever equipment that is not designed to do this is attached externally to a ring structure, and that equipment tries to form a logical ring, or a logical drop-and-continue function, packets have to go in an ingress port, then switched across the backplane and processed out an egress port. This has the potential to increase latency and jitter as the number of nodes on a ring increases. This would be the case whether the architecture was based on H-VPLS or IP multicast routing.

Performance metrics we have received from some vendors indicate that latency would perhaps not be a significant issue with each node adding in the range of 20 microseconds (up to 1,518-byte frames). We do not have a good feel for the potential for jitter in this configuration, but vendors have assured us that this has not been a real-world problem in production deployments. We should also add that no service provider has suggested to us that this is a problem. We therefore mention this issue to add further evidence to our suspicion that once IPTV / VoD networks really start to scale to millions of three-TV households and subscribers watching high-definition content, network operators are more likely to move either to wavelength-based drop




and continue or to simply route individual wavelengths to each serving office. We also suggest to vendors that they should be aware of the potential for this configuration to create performance problems and to make sure they have these issues covered.

MPLS multicast switching to BRAS

While PIM-SSM is a good choice for an IP network, converged multiservice networks based on MPLS have another option referred to as point-to-multipoint (P2MP) LSPs (draft-ietf-mpls-rsvp-te-p2mp). This option is based on MPLS paths that are established using the RSVP-TE signaling protocol. These extensions to RSVP-TE to enable management of P2MP LSPs support not only the initial establishment of a multicast tree, but also modification to the multicast tree by refreshing path information; this is referred to as “grafting,” when adding branches, and “pruning,” when removing them. An LDP-based standard is being developed as well: draft-minei-wijnands-mpls-ldp-p2mp, but LDP in general is not seen as having the same traffic-engineering and fast reroute capabilities as MPLS networks based on RSVP-TE.

As stated in the introduction of this report, the advantage of MPLS-based multicast is clear. If an operator is developing operational procedures around all unicast traffic going over MPLS, then it may want to do the same for multicast traffic and leverage some of that investment. There is an additional driver as well. One large network operator we spoke to said it does not want any IP routing in the core as a general architecture philosophy. However, as specifically related to IP multicast, the operator does not want any control plane activity in the core that is directly driven by the actions of subscribers (IGMP join and leave requests eventually end up as IP multicast routing updates in the cases where state needs to be updated). This operator saw the removal of IP in general, but IP multicast specifically, from the core as a “security” enhancement. Ovum-RHK notes that this mindset is similar to that of some operators when ATM cores were used for IP networks (even though for many the primary driver at the time was price / performance). What is evident in this mindset is the view that MPLS is not a service technology but a transport technology; and therefore the transport layer should support multicast.

Figure 6 **MPLS multicast routing to Ethernet edge**



Source: Ovum



Using P2MP MPLS LSPs is a good way of distributing multicast traffic if the network operator values traffic engineering and fast reroute. By building static pre-engineered P2MP MPLS LSPs, operators can also take the dynamism out of the core multicast trees (in the case of an IPTV application, operators know ahead of time where the content has to go), and then just let the aggregation network prune / graft channels as needed. Operators can construct static multicast trees with standard IP multicast as well, but by deciding to standardize on MPLS, operators also gain the benefits of fast reroute and traffic engineering (which can be argued to be less of an issue in a static environment).

P2MP MPLS LSPs might be attractive for another reason, at least compared to H-VPLS- or even VPLS-based methods, specifically in national business services where a service may transit multiple core nodes, on which the service provider may not wish to turn on edge functions (which are required for VPLS). Some people believe this may even become an issue as aggregation networks scale.

One equipment vendor we interviewed suggested that P2MP MPLS LSPs do not have some of the benefits suggested because leaf joins are unavoidable. We were not able to evaluate this issue fully for this report, but we suspect that the operator we interviewed who is examining this approach believes that activity can be pushed out to the BRAS through static multicast LSPs, and therefore the core can remain stable.

Multicast VPNs

Even though a network operator may find P2MP MPLS LSPs, VPLS, or native IP multicast the right approach for distributing its own content, if it is wholesaling services to other ISPs, building hundreds of separate multicast trees for each wholesale customer creates problems. As a result, there is work under way to multicast-enable IP VPNs so that each wholesale customer can have its own VPN (which could be an MPLS/BGP VPN, for example) and run their own multicast tree within the VPN to the retail customer. Standards work in this area includes:

- draft-ietf-l3vpn-2547bis-mcast-01.txt
- draft-rosen-vpn-mcast-08.txt

The first draft above is applicable to removing IP multicast state from the core of an MPLS network. The second draft above has been implemented in networks since 1999 and has been established according to comments we received from Cisco to be stable and scaleable in production networks. The second approach does lead to IP multicast state in the core.

One operator we interviewed mentioned that its intention was to have the multicast VPNs implemented in the BRAS. This operator believes this approach will lead to a greater need for bandwidth in the access aggregation network, but the operator viewed this as the best option because including multicast VPN capabilities in DSLAMs would lead to, in their opinion, too much cost and complexity for that type of a device.



Impacts on optical transport

Though we briefly alluded to issues related to the underlying fiber / optical networks earlier, it is worth repeating that the nature of the underlying network can have a big impact on the options available to an operator. If there are many serving offices with connections to point-to-point fiber, an operator can explore a more regionalized aggregation architecture based on point-to-point gigabit or 10 gigabit Ethernet. If each serving office is connected to a ROADM and is receiving its own wavelength, the network is logically a tree and not a ring, which means that every multicast packet has to go on every wavelength anyway, so standard VPLS (as opposed to H-VPLS) or a cost-effective multicast router at the serving office might make sense.

Traffic engineering

Traffic engineering is the ability to precisely control the path upon which traffic flows. An operator can combine traffic engineering with QoS mechanisms to do fine-grained capacity planning and management and guarantee a specific level of service to the customer.

IP networks are generally presumed to be incapable of supporting intra-domain traffic engineering, and this applies to PIM-SSM networks as well. MPLS networks that establish paths with the RSVP-TE signaling protocol do support traffic engineering, and as such, both MPLS P2MP multicast and VPLS implementations that support RSVP-TE support traffic engineering.

Fast reroute

IP networks have historically been presumed to not support fast reroute capabilities, though there are some fast reroute standards for unicast IP networks currently working their way through standards bodies (the construction of unicast routing tables is an important foundation of IP multicast "routing"). Not supporting fast reroute means that IP networks have to rely on the convergence times of routing protocols, which can take seconds at best and tens of seconds at worst (though there continues to be R&D investment in reducing this time). There are some approaches for reusing multicast state to quickly rebuild trees after a convergence event, but this is an area where operators (and their wholesale customers) might choose to see a standard IP multicast network as having a deficiency.

MPLS and VPLS (itself based on MPLS) can both make use of MPLS-based fast reroute.

One argument against the importance of fast reroute in video distribution is the fact that even a few milliseconds of recovery time could lead to an I-frame being dropped and the end user losing multiple seconds of video as a result. Therefore the argument goes that a slower convergence time would be acceptable. Ovum-RHK accepts that the loss of a single I-frame could result in a noticeable video loss from the end user's perspective, but still believes that operators will tend to want approaches that provide



the fastest possible restoration; using of course SONET/SDH as one benchmark, simply because that has cemented in operators' minds what is possible.

Some approaches to end-to-end FEC (forward error correction) might make the continually reducing IP routing protocol convergence times less significant in video delivery for viewing experiences that are not degraded by up to a few seconds of latency. Technology from companies like Digital Fountain is worth keeping an eye on.

OAM

OAM includes the ability to perform break-fix operations (identifying and locating a break and suggesting how it can be fixed) as well as ongoing performance and fault management monitoring to capture transient errors and provide data that supports SLA management.

The ability to monitor what data is actually being sent to a subscriber can be enhanced through traffic mirroring.

Scalability

Scalability is expressed in multiple dimensions (traffic; control plane processing while multiple functions are executing, such as VPNs, Internet access, IPTV; and carrier-class operational procedures). Therefore, we must conclude that scalability is an unknown for all multicast approaches. Some very large IP multicast networks have been built for enterprise networks, specifically in the financial sector, and some of these are said to be larger than some small service provider networks, so there are at least a few data points supporting the idea of IP multicast scalability in an enterprise environment.

High availability

High availability is achieved by both network element and end-to-end network capabilities and architecture choices. Most approaches can use redundant elements, in one way or another, and reduce recovery time by statically configuring the delivery of multicast traffic to multiple network elements.

Most network elements support nonstop forwarding for unicast traffic and are extending this capability to multicast traffic. One vendor supports nonstop routing (stateful failover) for unicast, multicast, and multiple VPLS instances.

Availability of video servers for VoD content is also important. While this is not a multicast issue per se, an IP addressing/routing approach known as "anycast" allows multiple video servers to use the same IP address such that if one fails, the network can automatically use another without any impact.



Life after multicast

In the history of service provider data communications, multicast—without a widespread application driving its use—has largely been a solution in search of a problem, with some successful niches. IPTV presents exciting opportunities for multicast technologies to be built at scale and prove their worth. Today, there is much that is unknown about multicast: if and where it will reach its scaling limits, which technology will be the most stable in the presence of network elements supporting multiple functions, and how to operationalize the technology at scale. As a result, it is an important area of design consideration.

However, the future of public video looks every year to be more dominated by unicast business models. In addition, triple-play networks will support voice, business VPNs, and high-speed internet. As a result, total traffic demands for NGN networks may be driven more by unicast than by multicast, and therefore during the writing of this report, the vendors and operators we interviewed cautioned us more than once not to assume that multicast should be the sole criterion for how a network should be architected and implemented.

A single packet transport network can support multiple forwarding modes, including both a multicast and unicast capability. So, while unicast may be the predominant traffic type, multicast remains an important tool for controlling transport costs as long as large linear programming content packages need to be broadly deployed.

Summary and conclusions

During our research, we were unable to clearly determine a winner, either technically or in terms of market momentum. However, *Figure 7* shows a high-level view of some interesting points of comparison.



Figure 7 Multicast comparison

| | IP multicast (PIM-SSM) | MPLS multicast (P2MP LSPs) | IP multicast + VPLS |
|-------------------------|---------------------------|-------------------------------|-----------------------------|
| Automated tree creation | | Not when traffic engineered | Not when traffic engineered |
| Fast reroute | Unknown | | |
| Traffic engineering | | | |
| Nonstop forwarding | | | |
| Nonstop routing | Vendor dependent | Vendor dependent | Vendor dependent |
| Widely supported | | | |

Source: Ovum

Perhaps the clearest dividing line is where MPLS is and is not being used. MPLS is not an end-to-end reality in the Internet today, so it has limited use for purely Internet-based services. However, for network operators delivering services over their own infrastructure, MPLS-based approaches to multicast might make sense if they are already converging unicast traffic over MPLS.

Of the two MPLS-based approaches to multicast, VPLS has broader support today than P2MP MPLS LSPs. However, we expect that P2MP LSPs will grow in the future, and we know of one high-profile operator that is already using them in a business service, and we believe they are pushing hard for them in their residential infrastructure as well.

Those operators that use VPLS as an aggregation technology will likely also use IP multicast routing at the core in the short term, with some of them moving to P2MP MPLS LSPs in the longer term.

**MPLS disrupts multicast as well**

Despite the continued referral to NGN packet transport networks as “IP/MPLS” networks, the fact remains that for some network operators, the networks are not IP networks: they are MPLS transport networks that support multiple services including IP services. For these operators, it is intuitive for them to standardize on MPLS-based multicast if they have already standardized on MPLS-based unicast; with the caveat that there are always exceptions due to the individual passions of different decision makers. Multicast MPLS is rare today, but we highlight it as an important trend to watch in the future.

IP multicast remains the only form of multicast that fits the “Internet” model, but for deploying multicast over a closed portion of a network operator’s network, or for some select negotiated interconnection scenarios, other multicast mechanisms for aggregation as a complement to an IP multicast core has been gaining favor. VPLS is mostly an aggregation option, with P2MP MPLS LSPs having the potential to be an end-to-end option within a network domain/area as that technology matures.

Ovum does not endorse companies or their products. Ovum operates under an Independence Charter. For full details please see www.ovum.com/about/charter.asp.

For full details of Ovum's citation policy, see www.ovum.com/media/citation.asp.

Whilst every care is taken to ensure the accuracy of the information contained in this material, the facts, estimates and opinions stated are based on information and sources which, while we believe them to be reliable, are not guaranteed. In particular, it should not be relied upon as the sole source of reference in relation to the subject matter. No liability can be accepted by Ovum Europe Limited, its directors or employees for any loss occasioned to any person or entity acting or failing to act as a result of anything contained in or omitted from the content of this material, or our conclusions as stated. The findings are Ovum's current opinions; they are subject to change without notice. Ovum has no obligation to update or amend the research or to let anyone know if our opinions change materially.