



IPTV Appliance

Optimizing Provider Networks
October 2006



INTRODUCTION

With Service Providers pursuing new revenue streams from non-traditional sources, the end result is that providers become more and more alike, and yesterday's specialty services becoming today's commodities. Established voice providers are pursuing high-speed Internet access and multimedia (particularly video) distribution. "Gathering more momentum in the ongoing broadband wars, U.S. phone companies netted slightly more high-speed data subscribers in the second quarter than did their cable rivals," ([Cable Digital News](#), Light Reading, 9/14/2006)

Cable and satellite multimedia providers are offering voice services to preserve their standing. "As a result, the Bells continued to nibble away at the cable industry's broadband dominance, further slicing cable's once commanding lead in the market. But, thanks to its strong headstart earlier in the decade, cable retains a sizable edge over the telcos as broadband penetration approaches 45 percent of all U.S. homes." (op. cit.)

Cellular providers seek to layer both traditional data services such as electronic mail as well as multimedia onto their networks in support of customer acquisition.

The common denominator is the Internet Protocol [IP] as the basis for services delivery. IP-based telephony is offered by cable and satellite data providers as a means to avoid tariffs associated with circuit-switched telephony. Circuit-switched networks transition to IP to provide audio and video programming to existing subscribers in hopes of retaining them. IP has been adapted for and applied to the transport of multimedia data since the late eighties, and became especially suited with the development of multicast transport protocols such as the Distance-Vector Multicast Routing Protocol [DVMRP, RFC1075], multicast extensions to the standard routing protocols such as OSPF [MOSPF, RFC1584], and group management protocols [IGMP, RFC1812].

As demand for the new multimedia services grows, mechanisms must be found to maximize the bandwidth allocated to delivering these services. Physical media must provide enough capacity to sustain multi-megabit data streams without congestion from other concurrent traffic. Layer 2 networking must deliver data atoms by the shortest path and without disruption. Layer 3 transports must guarantee in-order arrival and guarantee delivery of all atoms comprising a data stream. Layer 4 applications should accurately track setup, usage, and termination statistics on a per-user basis.

It is axiomatic that network processing should occur as closely as possible to either source or destination of traffic, to minimize (or even eliminate) delay and overhead within the core. In the delivery of IP-based multimedia traffic the processing that occurs at the source is related to the efficient distribution of resource-intensive data streams. It is at the network edge closest to the destination of the multimedia traffic where most of the processing should occur, and also where the most important processing occurs.

It is at the edge where requests for multimedia are inspected to determine if they are from an authorized user. The requests are then validated in terms of permitted services. The requests are then sent to the appropriate source, joining (or forming) groups of like requests to participate in multicast efficiencies in the distribution of streams. Likewise, as group memberships define participation in multimedia distribution, the network edge is where statistics relating to that usage are gathered for billing purposes.

The optimal edge device, then, will be the one that accurately and efficiently allows for the participation in the new multimedia services, like IPTV, and tracks that usage for revenue generation. Further, it will participate in IPTV while imposing the lowest impact on the core of the network, maximizing network resources for the distribution of additional revenue generating multicast data streams.

NETWORKING TODAY

Current networks were not designed to the current set of requirements: multi-transport, multi-purpose, multi-protocol converged distribution of circuit services, point-to-point, point-to-multipoint and multipoint-to-multipoint services. Security threats arise from internal and external sources, comprising theft of identity, theft of services, and theft of data, malicious intrusion and compromising of integrity.

While networks are being re-architected, the process is time-consuming and not complete. Network optimization comes from the equipment placed in the network. Further, that equipment allows for controlled migration of infrastructure to more efficient technologies, maximizing revenue generation during the migration. Migration occurs while preserving the integrity and security of the network, and doing it without disruption to the quality of the experience for the subscriber.

To make new services revenue generating as rapidly as possible, it is then necessary to have equipment in place that allows the introduction of services without placing any undue strain on the network itself. Therefore, new equipment must support multicast protocols and efficient distribution, while moving that support further out to the network edge.

As those services deploy, providers must ensure that contracted services are delivered according to service-level agreements, and that subscribers received only the services for which they have contracted, and nothing extra. Further, services must go only to valid subscribers and not to unauthorized intruders on the network.

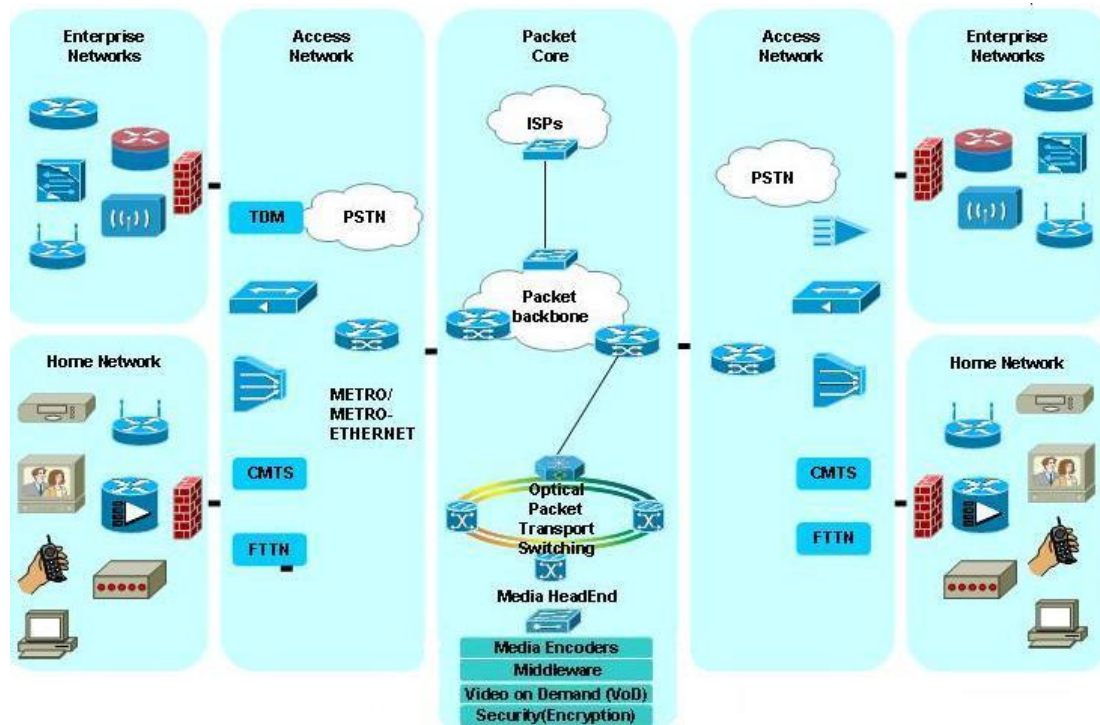


Figure 1 Video Everywhere

Distribution mechanisms

Multicast data, including IPTV is distributed in two main fashions today. One set is managed, and includes both broadcast and on-demand (Pay Per View) video programming. The second set is un-managed, and includes video conferencing (as in teleconferencing) and the more interactive forms of online gaming and TV.

MANAGED

Managed distribution of multicast data comes from content providers, or more commonly a partnership between content-provider and service provider. In both cases, subscribers are known, services have been authorized, and delivery is predictable, reducing the complexities of provisioning the infrastructure.

Broadcast

Broadcast audio and video are those streams that are associated with “regular programming” from content providers. While the overall subscription level may not be known, or even predictable, the streams will always be provided at the same time, on the same days, and thus are suited to multicast distribution involving the designation of a talker, optimally located branch nodes, and even to an extent participating subscribers.

A video content provider such as a cable operator carries a finite offering of selections. Those selections may be global (from a network perspective), regional or local in nature, defining sets of distribution points and physical locations. The process of acquiring subscribers provides a de facto maximum number of consumers for any given offering at any given point in time.

On-Demand (Pay-per-View)

On-demand programming is managed from the perspective that content is finite, defined and distribution paths are predictable. This allows for a degree of certainty when the network is provisioned, and minimizes the probability of over-subscription of any single network element (given that provisioning is to the maximum predicted load, and the resulting idle bandwidth is acceptable).

However, due to the on-demand nature of the service, a multicast distribution mechanism is not necessarily the optimal method to deliver the data. While multicast can be used to “pre-position” on-demand data with the anticipation of some degree of demand, there will remain a subset of content that cannot be anticipated to generate enough demand to justify a multicast network. Instead, a unicast architecture for on-demand services is preferable, eliminating idle branch nodes and underutilized network paths.

As with broadcast audio or video, a content provider such as a cable operator carries a finite offering of selections. Those selections may be global (from a network perspective), regional or local in nature, defining sets of distribution points and physical locations. The process of acquiring subscribers provides a de facto maximum number of consumers for any given offering at any given point in time.

UNMANAGED

Some unmanaged distribution of multicast data (like managed) comes from content providers, or more commonly partnerships between content- and service-provider. In these cases,

subscribers are known, services have been authorized, but the end-points of multicast streams are indeterminate, and delivery is unpredictable owing to the proliferation of intelligent mobile devices (allowing for connections to be made from any location). This has the effect of complicating the provisioning of the infrastructure, as [more or less] static tree structures cannot be defined owing to the dispersion of the leaf nodes. However, unmanaged multimedia services also stem from the low barrier to entry that allows entrepreneurs to come into being as “content providers” (e.g.: the growth of multimedia content in blogging, or mblogs). In this case, demand is sent to, and content is provided from extra-net sources that can have serious negative impact on network provisioning.

Gaming / P2P / Interactive TV

Gaming, peer to peer services and interactive TV make up the bulk of unmanaged multimedia distribution right now. Massively Multi-player Online Game(s) [MMORGs] have subscription numbers that are in the millions, and growing rapidly (in some areas more than others). While the sources of these streams are known and predictable, for the reasons cited above, destination leaf nodes are undefined.

The same complexities arise in the P2P and Interactive TV subsets. P2P and ITV services are aimed at the market segment that is typified by mobile handset connectivity. Audio data (e.g.: music, or ring-tones) are downloaded at any time, anywhere the requestor is able to make a connection. Further, it is these same devices that are (predominantly) targeted by ITV providers, as the end-user has the ability to use the mobile device controls in interacting with the service.

Video Conferencing / Teleconferencing

Business has long taken advantage of communications in providing remote access to centralized services like teleconferencing. Early implementations suffered, however, from costs associated with providing the connectivity and lack of quality in the transmission, yielding a less-than desired user experience.

With the (conversion of and) expansion of the network edge into low-cost media like Ethernet, and growing sophistication of the transport technologies like MPLS, those barriers to adoption are being minimized, or even eliminated. Digital signals can now transit the entire length of the connections from audio and video equipment to end device. As the quality of the experience increases, even greater adoption is to be expected.

Add to this the commoditized nature of the audio and video equipment, and these same services are now within reach of large segments of the “wired” public. Intra-family video conferencing is the modern equivalent of the “party line.” On-line service providers host full-feature connections between subscribers having a mutual interest. Social networking rises to a new height with audio and video added to the sites.

Distribution architecture

The current distribution architecture for multimedia is dependent upon a ‘physical’ transport that provides sufficient bandwidth while maintaining security at the lowest cost, and transmission protocols that ensure on-time in-order delivery of packetized data.

LAN SERVICES

LAN services, as extensions of the traditional Local Area Network, offer the most cost-effective mechanism for providing multimedia services. The physical layer is inexpensive, the technology is well-known and wide-spread, administrative and management costs are incremental instead of additional, and existing host configurations can remain unchanged. Two dominant access strategies exist in today's market, both gaining traction, although the intersection between the two sets is limited. This allows for growth in both connectivity offerings, with little or no risk of cannibalization for the near term.

Ethernet

Ethernet, as the set of physical and logical protocols that add up to a CSMA/CD [Carrier Sense Multiple Access/Collision Detection] architecture do not tend to extend beyond the premises in either business or residential installations. However, Ethernet in the broader sense of those standards that loosely equate to some combination of IEEE 802.1 and 802.3 implementations extend the premises to the edge of the carrier network. Protocols running in the premises now carry information to and from the Carrier network, and do so at 100/1000Mbps, outstripping legacy dial-up or cable network capacities.

WiFi

WiFi [802.11a/b/g] was intended to be used for mobile devices, such as laptops, in LANs comprising a residential network, or corporate intranet. A person with a WiFi device, such as a computer, telephone, or mobile computing device can connect to the Internet when in proximity of an 802.11a/b/g access point. The region covered by one or several access points is called a hotspot. Hotspots can range from a single room to many square miles of overlapping hotspots, including Internet and VoIP phone access in public "hotspots." These hotspots largely are extensions of the carrier network to free or fee-based subscribers within a commercial premise. WiFi can also be used to create a Wireless mesh network. Both architectures are used in a wireless community network like those proposed for San Francisco and metro-scale networks like M-Taipei, or San Jose.

PROTOCOLS

Quality of Service

With contention for bandwidth increasing, providing for fair allocation of bandwidth resources is itself difficult. In placing a greater emphasis on the transport of multimedia data providers require greater control over the bandwidth on a per-stream basis. There are two main mechanisms for exercising this control, using protocol extensions to apply traffic engineering to the network to establish routes having differing characteristics, and filtering packets according to policy criteria to establish "flows" that can be given preferential treatment and allocation to specific routes.

Traffic Engineering

Traffic Engineering refers to the discovery and configuration of selected routes through the network that are guaranteed to provide a specific level of service (e.g.: Bulk, 'best-effort,' expedited, or congestion-controlled versus non-congestion-controlled). The route or path chosen meets the requirement that is requested by the entry point to the network for a particular traffic type. Traffic that has equal requirements takes the same logical path through the network, while dissimilar traffic takes separate logical paths.

All these paths, however, traverse [roughly] the same physical paths, so

DEEP PACKET INSPECTION

MPLS

MultiProtocol Label Switching (MPLS) is a transport mechanism which emulates some properties of a circuit-switched network over a packet-switched network. It was designed to provide a unified data-carrying service for both circuit-based clients and packet-switching clients which provide a datagram service model. It can be used to carry many different kinds of traffic, including IP packets, and native Ethernet frames.

Operating at a lower layer of the transport stack, and using a label-swapping mechanism for switching packetized data through the network, MPLS is inherently more efficient than routed protocols. This makes it ideal as the transport for jitter- and timing-sensitive flows that comprise IP multimedia.

Unicast

Unicast refers to the sending of information packets to a single destination. "Unicast" is in direct contrast to multicast, where packets are delivered to multiple destinations simultaneously. In computer networking, multicasting is used to regain some of the efficiencies of broadcasting.

These terms are also synonymous with streaming content providers' services. Unicast servers provide a stream to a single user at a time, while multicast servers can support a larger audience by serving content simultaneously to multiple users.

Multicast

Multicast is the delivery of information to a group of destinations simultaneously using a dynamic tree structure to represent paths and destinations for data, delivering the messages over each link of the network only once and performing packet copies only when links to the destinations diverge.

Multicast tree structures are created when a content source or "speaker" advertises its presence to a set of specific addresses on the network. Thereafter, hosts indicate their willingness to subscribe, or "listen" to the content. As leaf nodes [hosts] join the network, intervening routers participating in multicast distribution collect those requests, and create "branches" that are grafted onto the trunk of the multicast tree. [It is worth noting that leaf nodes drop off the multicast branches, as well, requiring pruning of branches by designated routers.]

To minimize the overhead associated with maintaining the tree structure, these participating routers "designate" themselves as the focal points for group maintenance, and take the responsibility for sending the group protocol messages back down the tree toward the next lowest graft on the trunk.

Group Management Protocol

IGMP is a protocol used between hosts and multicast routers on a single physical network to establish hosts' membership in particular multicast groups. Multicast routers use this information, in conjunction with a multicast routing protocol, to support IP multicast forwarding across the Internet.

A router receives explicit Join/Prune messages from those neighbors that have downstream group members. In order to join a multicast group, a host conveys its membership information through the Internet Group Management Protocol (IGMP). The router then forwards data packets addressed to a specific multicast group only onto those interfaces on which explicit

joins have been received. A Designated Router [DR] sends periodic Join/Prune messages toward a group-specific Rendezvous Point [RP] for each group for which it has active members. (Note that one router will be automatically or statically designated as the rendezvous point (RP), and all routers must explicitly join through the RP.) A route entry is built where the incoming interface points toward the RP; the outgoing interface(s) point to the neighboring downstream router(s) that have sent Join/Prune messages toward the RP. This state creates a shared, RP-centered, distribution tree that reaches all group members. When a data source first sends to a group, its DR unicasts Register messages to the RP with the source's data packets encapsulated within. If the data rate is high, the RP can send source-specific Join/Prune messages back towards the source and the source's data packets will follow the resulting forwarding state and travel un-encapsulated to the RP. Whether they arrive encapsulated or natively, the RP forwards the source's de-encapsulated data packets down the RP-centered distribution tree toward group members. If the data rate warrants it, routers with local receivers can join a source-specific, shortest path, distribution tree, and prune this source's packets off of the shared RP-centered tree. For low data rate sources, neither the RP, nor last-hop routers need join a source-specific shortest path tree and data packets can be delivered via the shared, RP-tree. Once the other routers which need to receive those group packets have subscribed, the RP will unsubscribe to that multicast group, unless it also needs to forward packets to another router or node. Additionally, the routers will use reverse-path forwarding to ensure that there are no loops for packet forwarding among routers that wish to receive multicast packets.

IGMPv1/2/3

Over time, multiple variations of IGMP were architected with the aim of reducing overhead imposed on the networks from management traffic associated with group participation. The first version, IGMPv1, was a minimal implementation of "join" messages, relying on timeout for leaves. IGMPv2 added "leave" messages to the management protocol, also implementing a "fast leave" concept to free up now unused bandwidth. IGMPv3 was a major rewrite to provide for source and unwanted packet filtering.

MLDv1/2

Multicast Listener Discovery protocol is the IPv6 version of IGMP for IPv4. It allows the router to discover not only directly attached leaf nodes, but also the speakers in which those leaves are interested. Version 2 of the protocol added the ability to allow for source filtering at both include and exclude levels. Over time, multiple variations of IGMP were architected with the aim of reducing overhead imposed on the networks from management traffic associated with group participation. The first version, IGMPv1, was a minimal implementation of "join" messages, relying on timeout for leaves. IGMPv2 added "leave" messages to the management protocol, also implementing a "fast leave" concept to free up bandwidth. IGMPv3 was a major rewrite to provide for source and unwanted packet filtering.

PIM

Protocol-Independent Multicast (PIM) is a family of multicast routing protocols that can provide one-to-many and many-to-many distribution of data over the Internet. The "protocol-independent" part refers to the fact that PIM does not include its own topology discovery mechanism, but instead uses routing information supplied by other traditional routing protocols such as Border Gateway Protocol (BGP). There are four variants of PIM:

- PIM Dense Mode (PIM-DM) implicitly builds shortest-path trees by flooding multicast traffic domain wide, and then pruning back branches of the tree where no receivers are present. PIM-DM generally has poor scaling properties and as such is more ideally suited for small or Enterprise networks where the domain is limited.
- PIM Sparse Mode (PIM-SM) explicitly builds unidirectional shared trees rooted at a Rendezvous Point (RP) per group, and optionally creates shortest-path trees per source. PIM-SM generally scales fairly well for wide-area usage.
- PIM Source Specific Multicast (PIM-SSM) builds trees that are rooted in just one source, offering a more secure and scalable model for a limited amount of applications (mostly broadcasting of content). In SSM, an IP datagram is transmitted by a source [S] to an SSM destination address [G, or group] and receivers can receive this datagram by subscribing to the channel delineated by (S,G).
- Bidirectional PIM explicitly builds shared bi-directional trees. It never builds a shortest path tree, so may have longer end-to-end delays than PIM-SM, but scales well because it needs no source-specific state.

Of the four, PIM-SM has the widest deployment. It is ideal for IP NGN multimedia data as it builds the specific tree at the closest graft (i.e.: the closest Rendezvous Point [RP]) and requires only unidirectional traffic flow, optimizing the network resources for the transport of data, with little control overhead.

IPTV APPLIANCE

With most of the provider infrastructure in place, delivery of multicast data, especially that relating to broadcast and on-demand video (IPTV) is assured *right up to the edge of the network*. What is needed now is development and introduction of a device (or family of devices) that will optimize the last mile of the connection into the premises.

One of the most recent concepts in the distribution of multimedia is the home gateway [GW]. Central to the contemporary multimedia household, the home GW serves as the repository for all multimedia data in the household, to which supplicants connect for content, and to which source devices send the data which they have captured.

The home GW serves all authorized devices within range, and also remote devices connecting via wide-area connections across the Web. In effect, the household then becomes a private or semi-private content provider, filling content requests to all authorized users regardless of connection. It is this application that makes the optimization of that last mile of utmost importance.

The next critical piece of the IPTV story is the optimized, application-specific appliance that guarantees the highest-quality viewing experience for subscribers to the Next Generation Network. Consumers are driving the explosion of multimedia data through the purchase of IP-enabled devices. They will be the primary market for the home GW, securing always-on access to their personal media. It stands to reason the consumer will be the driving force behind the grooming of the last access mile into the home.

NETWORK LOCATION

Network location will be critical to the placement of the IPTV appliance. It must be located where the benefits from the enhanced technologies can provide the maximum benefit to the last mile connection to the provider network edge. Quality of Service must be extended from provider edge to consumer edge. Overhead associated with multicast distribution must be minimized to provide the greatest amount of bandwidth for actual data delivery. Authorization,

authentication and accounting [AAA] must be performed to support provider delivery and prevent theft of services.

The IPTV appliance should lie either within the residence or just without a group of premises (e.g.: multi-tenant unit or small residential association). This will place the device just behind the demarcation between public and private network, taking advantage of a firewall and other security measures in place.

[NOTE: depending on the existing security implementations, there may be changes necessary to active policies to permit multicast traffic to transit the firewall. However, those types of changes would be minimal and offset by security within the appliance itself.]

Multi-Tenant Unit [MTU]

For the multi-tenant implementation the IPTV appliance will be the last location of control for the provider, authenticating requests for participation in multicast streams through local (or remote) policies. It will aggregate all requests for downstream leaf nodes, and convey only those requests for services that are licensed by the consumer back up the tree for registration and participation. It will likewise convey drop requests for more efficient bandwidth reclamation as leaves leave the tree.

As the focal point for group arbitration, the appliance will significantly reduce the amount of overhead associated with joins and leaves by using the most recent enhancements to the group management and multicast distribution protocols. This will also improve delivery of streams as fewer network and processing resources are expended on administrative overhead, and more on delivery of packetized data.

Home

For the home implementation, the appliance serves an offload function for the home GW, providing group administration for all devices within the home, and optimizing delivery of multimedia streams to the requesting device(s) for direct consumption, or storage for future use. Confining the overhead associated with membership administration to the home network, the direct effects are:

- increase in the available bandwidth on the local loop for multimedia data,
- reduction in the group management protocol overhead arriving at the provider edge from subscriber households
- a significant savings in processing overhead associated with the expanded number of devices in the household, as a single authentication is done for the IPTV appliance, which then performs all further downstream authentication.

PROTOCOL SELECTION

Practical experience and seems to indicate that the current standards solution with IGMPv1 and IGMPv2 does not work in an optimal way for professional, commercial distribution of TV with multicast. IGMPv3 addressed a number of the shortcomings in the earlier versions, but has one main concern that has not been addressed: knowing - even through firewalls - the number of subscribers (leaves) to a particular stream.

With today's technology, any number of platforms behind a firewall can watch the same stream because the additional boxes do not have to report to the RP that they want the group, it is sufficient to just listen for the same source port as it has been translated by the firewall.

This has the effect of reducing potential revenues for the provider, if they are trying to capture revenues by subscriber device as opposed to subscriber.

There are additional concerns that are addressed by v3, or that should be addressed:

- Prune (multicast leave request) times are important, ideally sub-second, to maximize available bandwidth. With IGMP today you must do queries to determine remaining listeners, imposing additional overhead possibly causing bandwidth congestion
- Stream quality and statistics reporting are lacking in current implementations. Any reporting mechanism accumulate and return quality related statistics to the provider RP or DR to allow gathering of audit information to ensure SLAs are met, and that bottlenecks do not exist in the distribution tree. It is quite difficult to obtain this kind of information today from multicast networks. A simple audit mechanism is a requirement for the protocol stack in an IPTV appliance.

IGMP & MLD

IGMP and its companion protocol for IPv6, Multicast Listener Discovery protocol are used to manage group membership, and thus the distribution of multicast streams in an IP network. Both are efficient at that management, but incorporate levels of overhead that can no longer be supported on provider networks in support of the current and expected large numbers of participants.

IGMP [RFC3376]

IGMP is the protocol used by IPv4 systems to report their IP multicast group memberships to neighboring multicast routers. Version 3 of IGMP adds support for "source filtering", that is, the ability for a system to report interest in receiving packets **only** from specific source addresses, or from **all but** specific source addresses, sent to a particular multicast address. That information may be used by multicast routing protocols to avoid delivering multicast packets from specific sources to networks where there are no interested receivers.

MLD [RFC2710] is the protocol used by an IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. This protocol is referred to as Multicast Listener Discovery or MLD. MLD is derived from version 2 of IPv4's Internet Group Management Protocol, IGMPv2. One important difference to note is that MLD uses ICMPv6 (IP Protocol 58) message types, rather than IGMP (IP Protocol 2) message types.

Snooping/Proxy

IGMP snooping, as implied by the name, is a feature that allows the switch to "listen in" on the IGMP conversation between hosts and routers. When a switch hears an IGMP report from a host for a given multicast group, the switch adds the host's port number to the list for that group. And, when the switch hears an IGMP Leave, it removes the host's port from the group table entry.

An interface running in IGMP proxy mode acts as a proxy multicast host that sends IGMP Membership Report messages on one interface for IGMP Membership Report messages received on all other interfaces running in IGMP router mode. The upstream router attached to the network of the IGMP proxy mode interface receives the IGMP proxy mode Membership Report packets and adds them to its own multicast tables. In this way, the upstream router knows to forward multicast packets to the network segment of the IGMP proxy mode interface for multicast groups that were registered by hosts attached to network segments of the IGMP proxy mode router.

When the upstream router forwards multicast traffic to the IGMP proxy mode interface network, it is forwarded to the appropriate hosts on the IGMP router mode interface networks by the TCP/IP protocol.

All non-local multicast traffic received on all interfaces running IGMP router mode is forwarded using the IGMP proxy mode interface. The upstream router that receives the

forwarded multicast traffic can elect to forward the multicast traffic or discard it. Using IGMP proxy mode, multicast sources on networks attached to the server running Routing and Remote Access can send multicast traffic to multicast hosts attached to upstream multicast routers.

IGMP proxy mode is designed to pass IGMP Membership Report messages from a single router intranet to a multicast-capable portion of the Internet. The multicast-capable portion of the Internet is known as the Internet multicast backbone, or MBone. With IGMP proxy mode enabled on the Internet interface, hosts on the single router intranet can receive multicast traffic from multicast sources on the Internet and send multicast traffic to hosts on the Internet.

SUMMARY

With Internet access becoming ubiquitous, and multimedia data (particularly IP-based TV) assuming primary importance as a premium revenue-generating service for Providers, networks are in danger of becoming overwhelmed with control traffic to the detriment of actual content.

New thinking has to be applied to the network, or to its participating elements, in order to maximize the existing bandwidth for revenue generation by delivering more streams, more profitably, to more devices. Unproductive overhead must be reduced or eliminated. The ideal candidate for eliminating the bulk of this administrative traffic (which originates primarily from within the home, or concentrations of homes) is with an intelligent device that can initiate, monitor and terminate group memberships in multimedia streams on behalf of many requestors. This device is the IPTV appliance.

Incorporating both IPv4 and IPv6 stacks, the device also incorporates IGMP and MLD in support of group memberships for either v4- or v6-based networks. Performing snooping, and acting as a proxy on behalf of the downstream leaf nodes, the device eliminates multiple control sessions between designated router, or rendezvous point and leaves, freeing large amounts of bandwidth in the public network.

IP Infusion is able to provide embedded software for these Next Generation Network appliances incorporating its market-leading dual IPv4/IPv6 stack, IGMPv3 and MLDv2. Supporting these advanced functions, IP Infusion provides standards-based, robust QoS implementations that are tightly integrated with silicon from leading chip manufacturers so that maximum use is made of available bandwidth in the public (and private) networks.

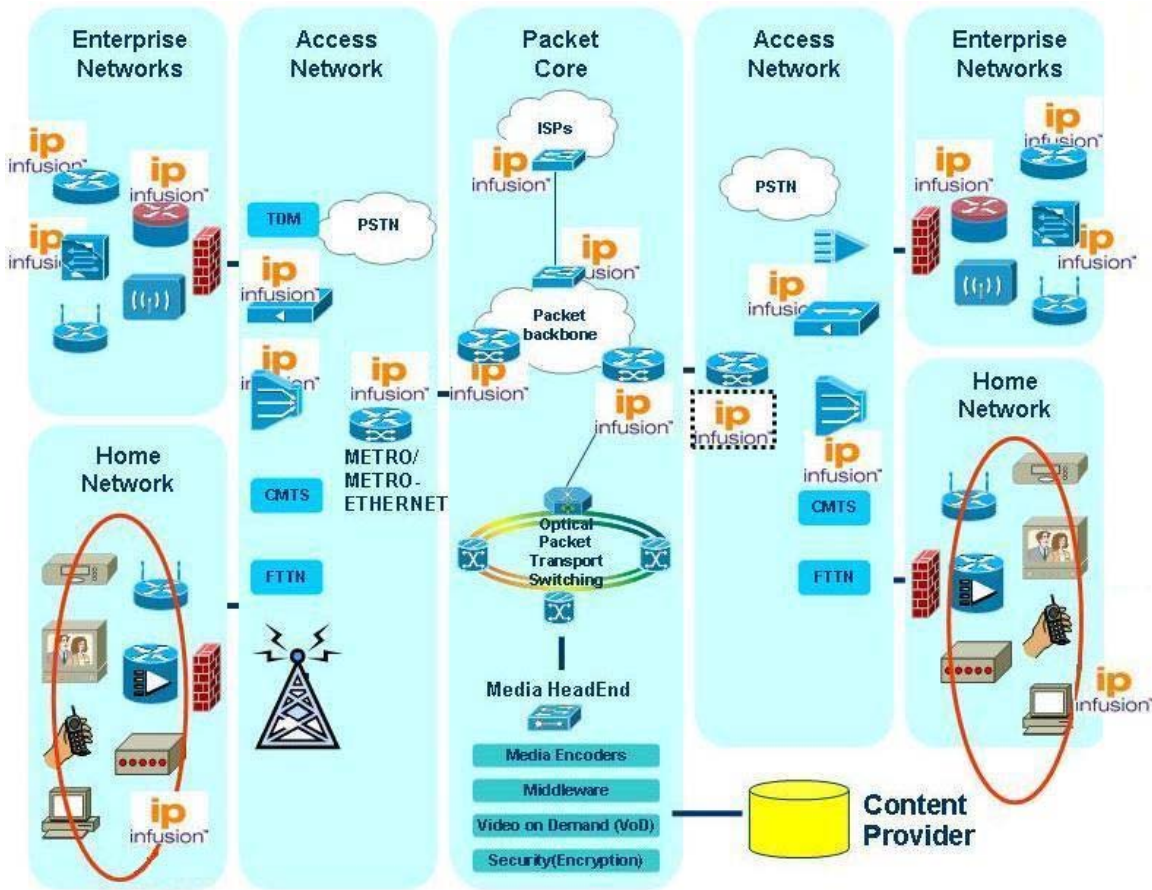


Figure 2 | IP End 2 End

Each protocol module in ZebOS ARS is built on the ZebOS Network Services Module (NSM). The NSM is the base module that simultaneously and independently communicates with every ZebOS ARS routing and switching process. The NSM manages both the route table and each of the enabled protocols; performs route conversion and redistribution; and manages the interface state, routing policies, and filtering. The ZebOS Integrated Management Interface (IMI) provides a command line interface and SNMP capability that can be used by vendors as is—or integrated into that vendor’s existing management infrastructure. APIs defined in both NSM and IMI are accessible and extendable. AIS is an architecture of software platforms called the Hardware Integration Platform (HIP), each of which is created for, and pre-integrated with, an industry-leading merchant silicon and an operating system. These HIPs provide a comprehensive forwarding plane implementation supporting L2, L3 (IPv4 & v6), multicast and MPLS/Traffic Engineering. A ZebOS AIS HIP, when combined with the ZebOS ARS protocol software for the control plane, provides a full system solution for enterprise switching, metro Ethernet, access, edge, mobile wireless and advanced IP services applications. Two key components in the ZebOS architecture enable AIS support—Protocol Abstraction Layer (PAL) and Hardware Abstraction Layer (HAL). PAL is a set of well-defined API that

abstracts operating systems calls, enabling ARS to use a unified set of API to access system resources and services such as memory and timer, on any operating systems.

Hardware Abstraction Layer (HAL) is the well-defined, extensible API between the ARS (control plane) and the AIS (forwarding plane). HAL isolates all hardware platform-specific interactions into a small set of well-defined function calls for the ARS (control plane). The HAL provides a unified interface for the control plane to interact with the forwarding plane for all L2, L3, multicast and MPLS forwarding needs. The function calls above the HAL run unmodified for any switching and routing hardware platform. The result is that customers have full flexibility to select only required protocol modules in the most cost- and code space-effective way.

Please visit the IP Infusion Web site for more information on ZebOS architecture and product offerings www.ipinfusion.com.



© Copyright 2006 IP Infusion Inc. All Rights Reserved.

ZebOS and IP Infusion are registered trademarks and the ipinfusion logo is a trademark of IP Infusion Inc. All other brands or product names are trademarks or registered trademarks of their respective holders. All specifications within this document are subject to change without notice. Contact Sales for current feature availability.