# ANYTHING, ANYTIME, ANYWHERE: OPEN ADVANCED BANDWIDTH MANAGEMENT OF ON-DEMAND SERVICES

Michael Chen
Concurrent Computer Corporation

*Abstract*

*The convergence of new technologies in an affordable fashion has given rise to new features that not only bolster customer demand, but also provide new revenue-generating opportunities. The ability to deliver the full promise of on-demand services -- "Anything, Anytime, Anywhere" -- is finally within reach, and customers are clamoring for their providers to deliver.*

*New features available now, and some of those envisioned for the future, are identified and investigated, as are the issues which face providers and vendors today. Observations and recommendations on the next generation of systems and their architectures are then offered in closing.*

## INTRODUCTION

Increased demand, competitive market forces, and technology advances have placed Gigabit Ethernet at the heart of new cable architectures offering additional revenue opportunities to the Multiple System Operator (MSO).

The adoption of standard Internet protocols has made the pervasive switching and routing capabilities which power the Internet available to these video delivery systems.

These capabilities provide a framework which, combined with new techniques such as network-based personal video recording (PVR), allow the MSO to deliver their customers the full promise of on-demand services – "anything, anytime, anywhere".

To deliver this, the MSO is faced with a bewildering array of challenges, from the selection and installation of compatible equipment to the configuration, management, and maintenance of this new infrastructure.

These issues facing both MSOs and equipment vendors today, as well as other looming issues, are further discussed below. The new features and capabilities of these systems, both at present and in future, are also identified and investigated. Finally, observations and recommendations are made for the design, procurement, and deployment of next-generation architectures and systems.

## GIGABIT ETHERNET ON-DEMAND SYSTEMS

Current on-demand systems are largely being deployed using Gigabit Ethernet output. A typical video-on-demand (VOD) system employing Gigabit Ethernet looks like this:

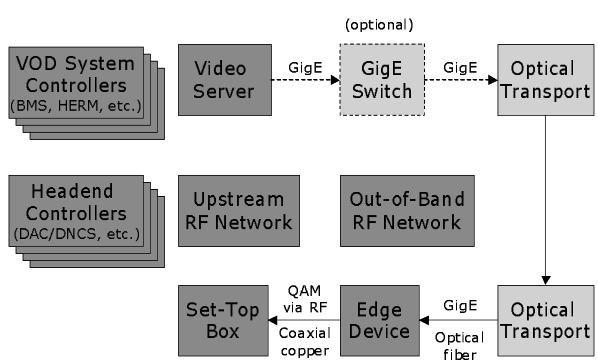


**Figure 1. Typical Gigabit Ethernet VOD system.**

The Gigabit Ethernet output of a streaming server is sent to an edge device, where it is combined and converted to a form suitable for display on digital cable set-top boxes. The server's output may be connected into a switch, and optical transport gear is used when necessary to transmit the signal across large distances.

The streaming server output is encapsulated within User Datagram Protocol (UDP) packets, defined as part of the Internet Protocol (IP) standards to provide low-latency data delivery, while taking advantage of the wide range of products and services the Internet explosion has produced.

Note that the output transmission is often implemented unidirectionally, since this allows the MSO to effectively double the amount of fiber bandwidth available. This one-way connection may require additional effort to configure systems initially, since many standards used with IP protocols assume the existence of a bi-directional network link for proper operation.

Gigabit Ethernet on-demand systems today are usually allocated dedicated network bandwidth for streaming. This often stems from the difficulty of ensuring sufficient quality of service to protect on-demand streams from being damaged by other data traffic. Having dedicated bandwidth for streaming, which the streaming servers then manage among themselves, greatly simplifies the overall system, and has accelerated the availability of Gigabit Ethernet solutions.

## To Switch or Not to Switch?

Gigabit Ethernet switches were used in early deployments to aggregate the outputs of one or more streaming servers, when these servers were unable to generate enough traffic to fill an entire Gigabit Ethernet link.

Since streaming servers can now saturate Gigabit Ethernet links, a switch is no longer technically needed for deployment. However, the use of switches also provides new routing flexibility that was either unavailable or cost-prohibitive with prior output formats, and many of the new features which Gigabit Ethernet enables are built upon this functionality. For this reason, using a switched Gigabit Ethernet transmission framework is still quite advantageous for these on-demand services.

## Asymmetric Deployment and Expansion

The division of labor between the streaming server and the edge device in the Gigabit Ethernet framework offers the MSO a new method for system deployment and expansion. Gigabit Ethernet's switching and routing functionality allows streaming servers and edge devices to be loosely rather than tightly coupled. The MSO can then deploy and expand edge devices separately from the streaming servers, allowing an asymmetrical buildout of the system.

A typical asymmetric buildout will overprovision the radio frequency (RF) edge with more edge devices than necessary to satisfy initial bandwidth demands. This is because installing new edge devices is often difficult to do without impairing the RF signal to a node, and requires more truck rolls to accomplish. The available granularities of optical transport equipment often will favor having more optical transport capacity than initially required, which may prompt the MSO to overprovision with edge devices at the same time.

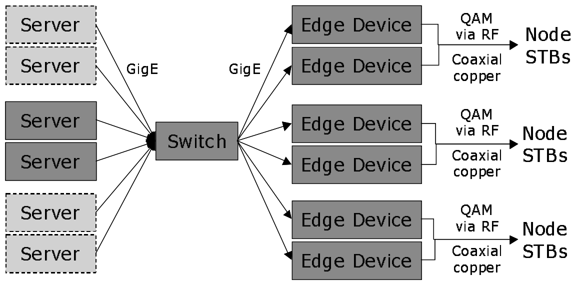


**Figure 2. Asymmetric deployment and expansion.**

Such an asymmetric buildout will generally add only as many streaming servers as required to meet current demand; as demand increases, more servers can be added at the headend and assigned to RF outputs of edge devices.

## REALITY TODAY

What can today's Gigabit Ethernet on-demand solution currently provide?

### "Something, Anytime"

Current solutions have limited on-demand content available. This is often not a storage capacity issue, but rather a rights licensing issue. The limited availability of on-demand content may force the MSO to select content for the on-demand system that is assumed to be more compelling than the broadcast digital cable offerings. This content typically includes movies, special events such as concerts, and popular sporting events.

### Now Playing: "Anything, Anytime"

Personal video recorders (PVRs) such as Tivo can provide a wider selection of on-demand content to the home, but the limited availability of PVRs with integrated digital cable functionality curtails the overall benefit to the customer. PVRs also remove content storage control from the MSO at the home. This raises content protection issues, which tend to ripple back into rights negotiations.

However, successful trials of subscription video-on-demand content indicate that MSOs may not need to supply customers with DVR boxes to satisfy their desire for more varied on-demand content, as long as they can make desirable content available to their subscribers.

### Network-Based PVR

In a network-based PVR approach, broadcast programming is recorded and stored by the MSO at the headend, rather than inside a consumer's set-top box, and is made available to on-demand streaming servers for transmission to customers upon request. Some implementations of network-based PVR allow a customer to pause a program in real time and use standard navigation features such as fast forward and rewind.

The advent of network-based PVR solutions levels the playing field with home PVR boxes, and allows the MSO to provide the full range of broadcast programming on demand, in addition to PVR functionality, without upgrading any customer premises equipment.

However, existing carriage agreements are likely to require renegotiation before broadcast programming will be allowed for on-demand viewing, so MSOs must aggressively pursue content rights to achieve the full potential value of network-based PVR.

### "Many Streams, Each To There"

Current on-demand solutions can be scaled to meet the MSO's streaming capacity needs for their digital subscribers. However, these solutions often suffer from inflexible routing that dates from the previous generation of transmission technology such as

DVB-ASI and integrated quadrature amplitude modulation (QAM) and upconversion. Since this transmission equipment had little or no switching and routing capability, and the capability was often not cost-effective when available, each streaming session had a fixed route to its destination. This meant that only a smaller subset of on-demand servers could stream content to a given customer's set-top box.

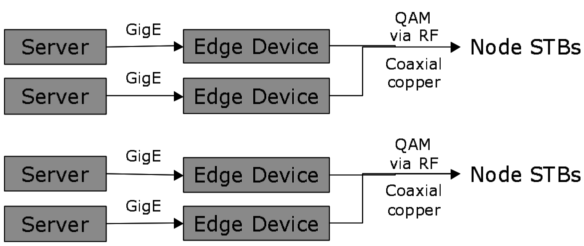


**Figure 3. Fixed server-to-edge device routing.**

For the MSO, these constraints meant that systems had to be designed for and sized to the peak demand expected at each hub, rather than the peak demand expected from the overall system. MSOs responded by defining system pricing in terms of cost per simultaneous stream, independent of service grouping or location. This pushed the cost of additional equipment to satisfy per-hub rather than overall requirements back on the equipment vendors, resulting in lower margins and profit from these sales.

This architecture is workable, but clearly not optimal for either MSOs or equipment vendors. MSOs must deploy larger systems that would otherwise be necessary, which impacts operational and maintenance costs, as well as complicating the issue of failure recovery. Equipment vendors must absorb costs imposed by sizing constraints at each hub, rather than at the overall system level. Performing asymmetric expansion of an on-demand system is further complicated by this routing inflexibility, since the expansions must again be performed at the hub level, not at the overall system level.

## THE PROMISE OF TOMORROW

### "Any Stream Anywhere"

"Any Stream Anywhere" is a phrase used to describe a system where any stream being sent from a streaming server can be directed to any set-top box. Looking from the other direction, this also means that any streaming server can satisfy a stream request from any particular set-top box.

A system with this property has many clear advantages. Since all streaming servers, not only a subset, can satisfy a node of set-top boxes, the total capacity provided by these servers can be sized against the demand of the overall system, instead of sizing each subset individually. This both eliminates unnecessary equipment, and also greatly simplifies the processes for installation and expansion. MSOs can set aside reserve streaming capacity to cover the entire system, rather than separate hubs or nodes.

A switched Gigabit Ethernet transmission framework can easily support the "Any Stream Anywhere" model, using the switching and routing functionality provided to direct traffic from any server to any edge device which transmits to a given set-top box. A conceptual diagram of "Any Stream Anywhere" for a system using a switched Gigabit Ethernet transmission framework is shown below.

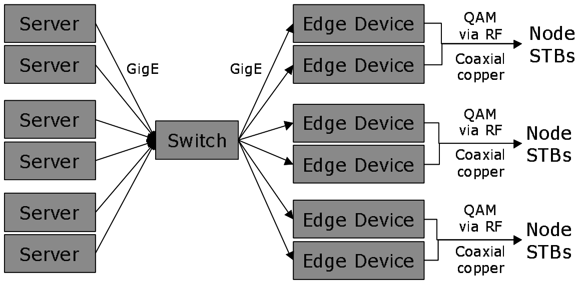


**Figure 4. Gigabit Ethernet Any Stream Anywhere.**

A basic implementation of "Any Stream Anywhere" using switched Gigabit Ethernet can take advantage of the fact that these systems are usually given dedicated network bandwidth. As long as the streaming servers can manage the available bandwidth properly, while taking into account the new switched infrastructure, few if any significant changes should be required to add the ability to support "Any Stream Anywhere" in an existing centralized Gigabit Ethernet system.

"Anything, Anywhere": Sharing Resources Between Multiple Services

A digital cable transmission system using Gigabit Ethernet has at least two distinct networks with resources to manage: the Gigabit Ethernet network used between streaming servers and edge devices, and the RF network between edge devices and set-top boxes. These systems also have an Ethernet network for command and control information, but that network is managed independently and falls outside the scope of this discussion.

RF Resource Sharing

Each RF frequency available has two separate but related resources to manage: the program numbers which can be individually tuned by set-top boxes, and the bandwidth which all programs using the same frequency must share.

A rudimentary level of RF resource sharing is easily achieved with a static partitioning of the available RF frequencies between the services sharing the RF network. This avoids most possibilities of conflict between services, but is clearly not optimal since resources unused by the assigned service are not available for reuse by other services.

An incremental improvement can be gained by changing the partitioning so that program numbers and their associated RF bandwidth can be assigned to services, instead of entire RF frequencies. However, the lack of mechanisms to guarantee quality of service (QoS) at this level makes it possible for an ill-behaved service to disrupt other services which share the same RF frequency.

Dynamic partitioning of these resources is clearly more efficient, but requires a resource management system to arbitrate requests. If the site in question uses the Scientific-Atlanta headend infrastructure, the Digital Network Control System (DNCS) is responsible for performing this function, using the DSM-CC protocol specified in the MPEG-2 standard. However, if the site uses the Motorola headend infrastructure, no such entity manages the RF resources. In this case, VOD system vendors have typically implemented their own internal management to handle resource sharing. Requests from other services for resource sharing can be accommodated by sending these requests to the VOD system for fulfillment.

At present, few services attempt to share RF resources with VOD systems, and the small number of involved parties makes solutions by private arrangement feasible. But as more potential services emerge, and providers begin to call for unified multiple vendor support, open standards should be adopted to define the interactions required for these services to share common resources.

## Gigabit Ethernet Resource Sharing

Resource sharing for the Gigabit Ethernet network is simpler, thanks to both its inherent switching and routing functionality, and the suite of Internet protocols available for use. Like the RF network, Gigabit Ethernet networks have at least two separate resources to be managed: the addresses used to identify each device on the network, and the bandwidth available for data traffic.

Ethernet devices generally have unique Media Access Control (MAC) addresses, so only IP addresses generally need to be directly managed. The Address Resolution Protocol (ARP), part of the standard suite of Internet protocols, handles the matching of IP addresses with appropriate MAC addresses, and the Dynamic Host Configuration Protocol (DHCP) is often used to assign IP addresses to devices, whether on a static or dynamic basis.

Gigabit Ethernet network bandwidth can be statically allocated in a fashion similar to the RF bandwidth allocation described previously to provide a rudimentary level of sharing between services. Without any quality of service guarantees, an ill-behaved or misconfigured service can once again disrupt other services sharing the same network.

The effects of this disruption can be significantly worse for Gigabit Ethernet, since the vastly increased bandwidth available encourages a correspondingly higher number of sessions per link to share the network. But in this case, the Internet comes to the rescue, since mechanisms have been developed to ensure quality of service for IP and Ethernet traffic.

## Gigabit Ethernet Quality of Service

There are several different methods, such as IP precedence, IP Type of Services (ToS), and Differentiated Services Code Point (DSCP), which can be used to specify which quality of service policy should be applied, if any, to IP traffic. Some of these methods overlap, and may conflict with one another if not configured and used carefully.

Fortunately, streaming servers are relatively immune to this problem, since the switch that receives their output can be configured to tag all incoming traffic on an input port with particular QoS settings. The streaming server is therefore not required to know how QoS will be implemented.

Edge devices are not so lucky, and so should be capable of receiving input with QoS tagging. QoS indications are not currently used to signal the relative priority of individual streams; therefore, vendors may note that it is safe for the edge device, as the last device in the chain, to ignore the QoS indications it receives.

Note that although lost data can sometimes be tolerated by other applications, within streaming video server output such losses are almost always clearly visible and objectionable to the customer. In light of this fact, best-effort queuing policies to enforce QoS are much more suitable for digital cable transmission than policies which result in lost traffic.

## Gigabit Ethernet Bandwidth Reservation

The standard Internet protocol used to perform network bandwidth management is the Resource Reservation Protocol (RSVP). This protocol allows a receiver to establish a bandwidth reservation between itself and a specified source. Dynamic partitioning of network bandwidth between services can be readily accomplished with this protocol.

A crucial RSVP feature is its ability to accommodate portions of the network that are not RSVP-aware. This feature enables the gradual introduction of RSVP at sites with existing equipment that predates or otherwise does not support it. Although many existing Gigabit Ethernet switches support RSVP, and optical transport equipment is generally not required to do so, existing streaming servers and edge devices largely do not support RSVP. Even if direct support for RSVP is added, the encapsulation of these messages within UDP multicast packets may be required, as specified in Annex C of RFC 2205. This is due to various operating system and security issues regarding the use of raw sockets.

The fact that many Gigabit Ethernet switches provide RSVP support is again advantageous to streaming servers, since these switches may act as a sender proxy and hide the details of RSVP operation from the connected servers. In this situation, the switch maintains RSVP states, and generates required downstream "Path" messages in response to received streaming input.

Edge devices are, once again, not as lucky and may be required to directly support RSVP. The primary reason for this stems from the fact that unidirectional transport from streaming server to edge device is often employed to better utilize the available optical fiber. For RSVP, the receiver must initiate an upstream request for bandwidth reservation, but it is unclear what upstream path will be available to the edge device.

## Bi-directional Edge Connectivity

Downstream video traffic requires much more bandwidth than upstream control traffic, which is why unidirectional transport from the streaming server to the RF edge is often implemented. However, having bi-directional connectivity at the edge would enable much simpler autodiscovery and autoconfiguration methods, and allow standard protocols used by the Internet such as ARP and RSVP to accomplish their tasks.

The establishment of bi-directional edge connectivity, with only unidirectional transport to the edge, requires a switch to exist between every edge device and the optical transport feeding it. This can become expensive, but an emerging new breed of equipment, combining switching and optical transport capability in the same device, may prove well-suited to this task.

As an alternative, devices may support methods such as the Unidirectional Link Routing (UDLR) protocols specified in RFC 3077 to logically create an upstream network path over a different connection, such as the command and control network.

## Autodiscovery and Autoconfiguration

Automatic discovery and configuration methods are not strictly required for these systems to be deployed. However, for MSOs unfamiliar with the intricacies of these new systems, any automation that can help reduce the probability of misconfiguration, and also simplify system expansion, will clearly be of great value.

However, to implement autodiscovery and autoconfiguration, bi-directional connectivity and support for each device is required. Set-top autodiscovery schemes can use the upstream communications link provided by the RF network to perform these functions, but this makes open standardization difficult. Existing network equipment with full support for autodiscovery and autoconfiguration methods may require modifications to work with unidirectional links, as described above.

Complex Network Topologies

Up to this point, the discussion of "Anything, Anywhere" has been based on a simple centralized model, where streaming sources and switches are located at the master headend, and their output is distributed to hubs and nodes using optical transport. Although the simplicity of this model eases the discussion of issues which are not dependent on topology, real-world systems are much more complicated.

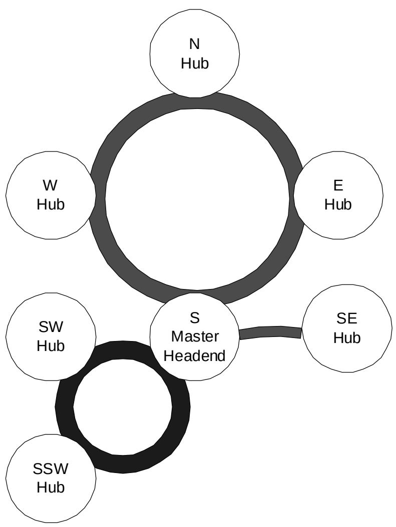


**Figure 5. A slightly more complicated topology.**

Since the cost of putting optical fiber in the ground is prohibitive, the topology of available fiber often dictates that of the services it carries. In other cases, the available space at headend locations may constrain the amount of equipment that can be installed. In addition, headends often also act as hubs to serve local customers. Lastly, redundant equipment is often used to provide failover capabilities. The net result of all this is that most real-world architectures diverge significantly from the ideal centralized model.

The multiple possible paths introduced by complex network topologies make routing and other management tasks much more difficult. However, complex topologies are generally chosen because they can be more flexible, and also more resilient if problems arise. Other possibilities which may drive MSOs to adopt more complex network topologies include the regionalization of functions such as broadcast feed generation, network-based PVR content ingestion, and reserve streaming capacity.

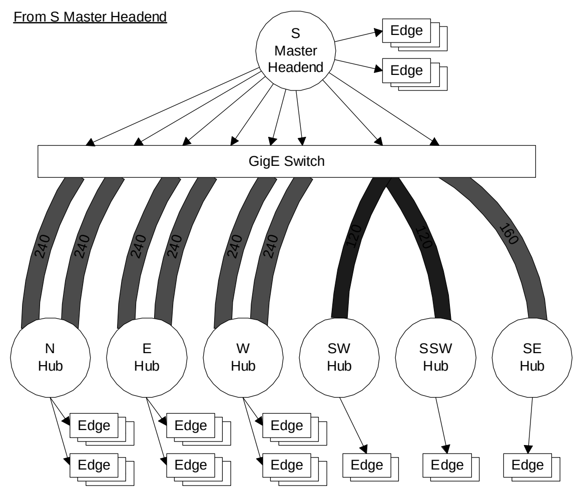


**Figure 6. This example is centralized... almost!**

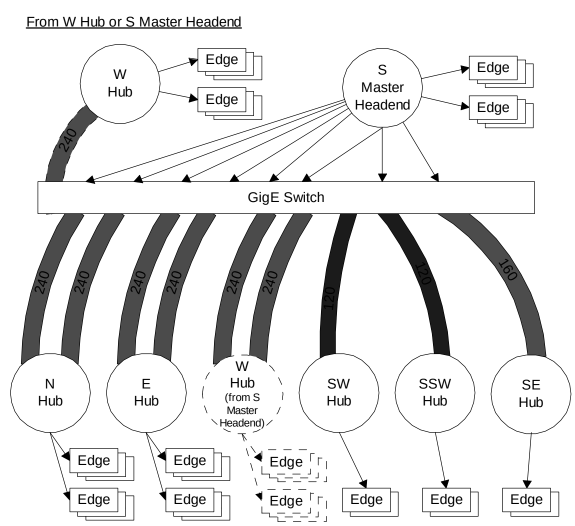


**Figure 7. A server elsewhere on the main ring makes things more complicated.**

Existing network devices and architectures have sturdy mechanisms available to handle failure detection and recovery, as well as other issues such as bandwidth reservation and quality of service. However, in some cases, the Internet solution does not quite fit the digital cable problem. For example, reserving bandwidth for a stream to a set-top box differs from the typical Internet case, due to the separation between the IP network and the RF network. A simpler problem thus becomes complicated in the digital cable space.

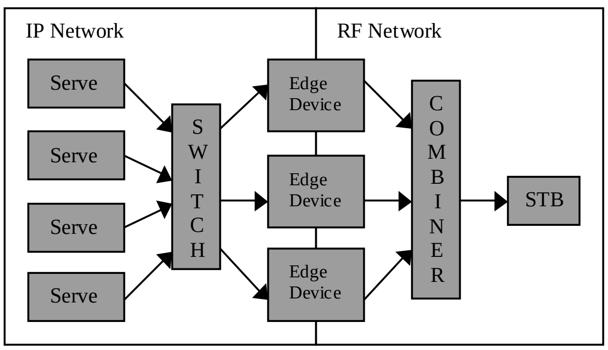


**Figure 8. Gigabit Ethernet IP and RF networks.**

The challenge here is to integrate network management functionality with the resource management of on-demand streaming systems. If this integration is performed at a high enough level, each piece can manage its own responsibilities and ask the others when external resources are required. But if the integration is performed at too low a level, the labor required to properly configure a system with complex topology may require that working autodiscovery and autoconfiguration methods be devised and implemented first.

## High-Definition Video-on-Demand (HDVOD)

The advent of high-definition (HD) content for VOD systems is quickly approaching. In fact, HDVOD may have already arrived! HDVOD content differs from standard VOD content only in video resolution and bit rate, but support for these higher resolutions and bit rates can have ripple effects throughout an on-demand system. Care must be taken in both the underlying infrastructure and devices themselves to ensure that HDVOD content does not cause design limits to be exceeded.

## RECOMMENDATIONS

Considerations for selecting and purchasing equipment for deployment:

Streaming servers should be able to fill a Gigabit Ethernet link so that switch ports and transport bandwidth are fully utilized.

Switches and/or routers should implement port queuing and QoS policy enforcement in a fashion compliant with streaming content requirements. Also, switches, routers, and transport equipment should only minimally modify the nature and timing of streaming media content. This is simplified by selecting equipment verified by vendors to interoperate correctly with other components, including both streaming server and edge device.

Edge devices should be upgradeable to support both variable bit rate (VBR) and high-definition (HD) streaming input. Devices with better buffering and dejittering capabilities are generally preferred over their competitors.

It must be decided up front whether asymmetrical deployment and expansion now merits the increased capital expenditure that it requires at initial rollout. Future cost projections for needed equipment will clearly play a significant role in this decision, as will the bargaining power brought by higher-volume and/or integrated purchases.

Considerations for designing or deploying a system:

The rollout and expansion of proven revenue sources such as on-demand services should not be delayed to wait for the promise of resource sharing with other services. The revenue to be gained now facilitates the expansion for these services later, and is a valuable hedge against the chance that other services may not end up as viable opportunities for additional revenue.

The network topology should not be complicated more than absolutely necessary, unless the benefits of doing so are tangible and compelling.

Component interactions should be kept at a high level when possible to accommodate differing implementation at lower layers. This avoids unnecessary problems that can arise from conflicting decisions made in the design and implementation of individual components.

The use of open standards should be encouraged for interoperability whenever feasible, but may not be required for existing or near-term deployments. This prevents unnecessary and unavoidable delays for acceptance and integration from impacting the timetables for these deployment.

## CONCLUSION

Equipment vendors in this space hold an enviable position; they are poised in a market ready to explode with new business, and are positioned well to capitalize on that fact. The new features needed by MSOs are already being developed and deployed now, while open standards are being refined and proposed to allow smoother integration and interoperability for the future. The acceptance and adoption of these standards will allow vendors to focus on the development of next-generation features to drive the next wave of business.

For the MSO, this is an exciting time to be in the business, due to the convergence of several new technologies in an affordable fashion. This recent development has given rise to new features that not only bolster customer demand, but also provide new revenue-generating opportunities. The wide range of services available to customers has neven been more compelling. The ability to deliver “Anything, Anytime, Anywhere” is finally within reach, and customers are clamoring for the MSO to deliver this promise. The last remaining hurdle is to standardize rollout procedures to make them suitable for mass deployment, and then the MSO can let the good times roll.

## REFERENCES


Chen, Michael, “Evolving the Network Architecture of the Future: The Gigabit Ethernet Story”, SCTE Cable-Tec Expo Proceedings, May 2003.

Braden, R., Ed., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification", RFC 2205, September 1997.

Duros, E., Dabbous, W., Izumiyama, H., Fujii, N. and Y. Zhang, “A Link-Layer Tunneling Mechanism for Unidirectional Links”, RFC 3077, March 2001.

Postel, J., “User Datagram Protocol”, RFC 768, August 1980.

Droms, R., “Dynamic Host Configuration Protocol”, RFC 2131, March 1997.

Plummer, David C., “An Ethernet Address Resolution Protocol“, RFC 826, November 1982.

Nichols, K., Blake, S., Baker, F. and D. Black, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers“, RFC 2474, December 1998.

ISO/IEC International Standard 13818; "Generic coding of moving pictures and associated audio information", November 1994.

Gigabit Ethernet Alliance, “Gigabit Ethernet: Accelerating the Standard for Speed”, 1999. http://www.10gea.org/GEA-Accel1999_rev-wp.pdf