

HOP A RIDE ON THE METRO GETTING YOUR IP PACKETS FROM HERE TO THERE

Erik Metz and David Grubb
Motorola Broadband Communications Sector
John Brenneman and Daniel Sills
Geyser Networks Inc.

ABSTRACT

Next generation Synchronous Optical Networking equipment is paired with data centric interfaces and packet switching engines to provide reliable Metropolitan networking for Headend interconnects. This integration has lead to new developments in multiplexing structures and bandwidth allocations. By creating more granular packet containers, and dynamic allocation, services are more efficiently packaged onto the transport network. Implementing priority based queuing allows for services to be expedited onto the network. Services bound by Quality of Service constraints in delay and

delay variance can be individually managed across the transport network. Differentiated Services, Open Access and Carrier Class transmission will merge into one consolidated backbone serving entire metropolitan regions.

WHAT SEEMS TO BE THE PROBLEM?

With the increased deployments in cable modem services, and the looming explosion of Voice over Internet Protocol (VOIP) telephony services, getting all of the packets to their destinations on time will play a critical role in the Metropolitan Network. Multiple System Operators (MSO) have many interconnected Headends serving geographical

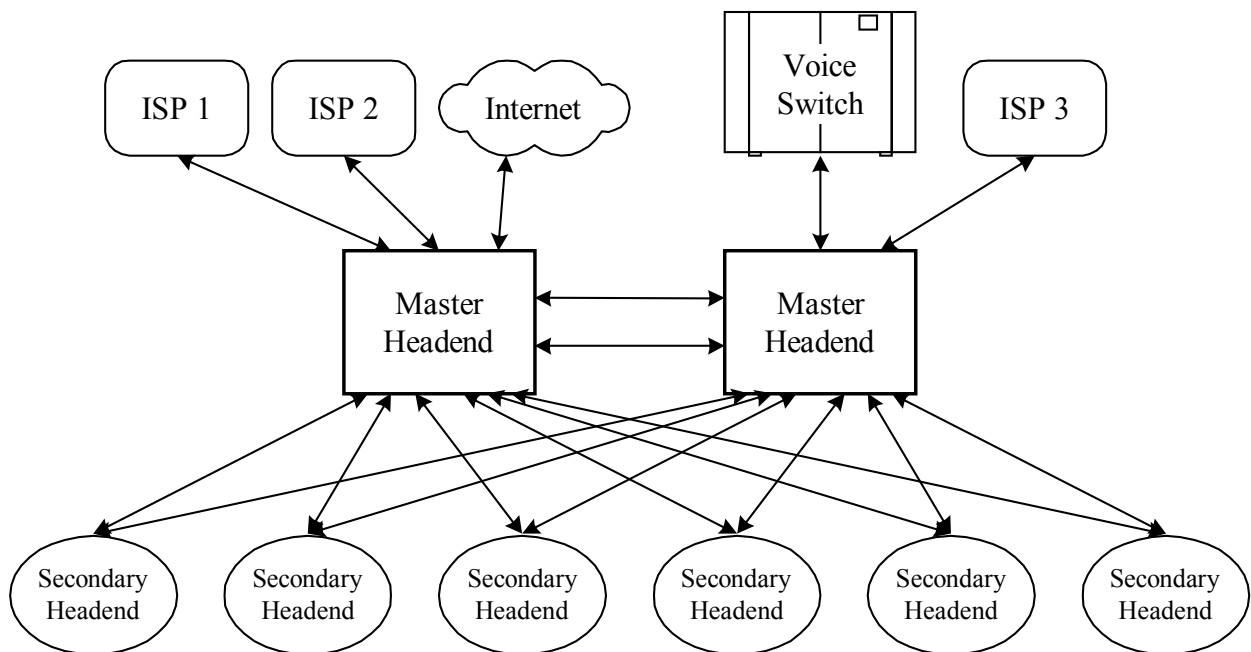


Figure 1

regions that make up the metropolitan network. One or more of these Headends may provide access to the Internet and one or more of these may have access to a voice switch for telephony. These Headends act as gateways to their respective services. These will be called the Master Headends. The remaining Headends will be called Secondary Headends. The Secondary Headends terminate the upstream and downstream traffic to/from the Hybrid Fiber COAX (HFC) plant via the Cable Modem Termination Shelf (CMTS) devices. The input and output from the CMTS to the Metropolitan Network is Internet Protocol (IP) packets. These packets need to be separated based on the type of service they are carrying and by destination Internet Service Provider (ISP).

The traffic that flows in this type of arrangement can be categorized as “Hub and Spoke” traffic (Figure 1). This implies that there is a large amount of traffic flow into and out of the Master Headends, and less in each of the Secondary Headends. This is a logical star. If there are two Master Headends, then it will look like two overlapping stars. Each Secondary Headend will be generating traffic that is different for each type of service it is providing. Internet traffic is bursty when looked at over short periods of time and when aggregated, can have sustained flows that change dramatically over the period of a single day. VOIP traffic is not necessarily bursty, but changes deterministically with each phone call that is placed. These aggregated traffic patterns will also change dramatically throughout the period of a single day.

Clearly there needs to be a solution that can efficiently aggregate the traffic from the Secondary Headends and reliably connect the Master Headends for access to the necessary gateways.

NEXT GENERATION SONET/SDH

Next generation Synchronous Optical NETworking (SONET)/ Synchronous Digital Hierarchy (SDH) platforms incorporate the latest technologies in data networking with the reliable transport from the telecommunications arena. SONET/SDH has been the benchmark for reliability for the Regional and Long Haul telecommunication networks. It is based on self healing rings and highly available equipment shelves with sub 50 ms recovery times from catastrophic failures such as a fiber cut. It is standardized world wide and has been deployed since the 1980's.

Traditional SONET/SDH was developed with the transport of multiplexed voice signals in mind. The basic building block is a Synchronous Transport Signal 1 (STS-1), or 55 Mbps, or Synchronous Transport Multiplex 1 (STM-1) for SDH, or 155 Mbps containers. SONET/SDH Networks are very scalable, and are available with different line rate interfaces and different trunk rate outputs (Table 1). Many of today's SONET networks are OC-48, (Optical Carrier 48 or 48 STS-1's) with a line rate of 2.455 Gbps. These large statically assigned containers make it difficult to implement today's packet services efficiently.

SONET Hierarchy		
Signal	Bit Rate	Capacity
STS-1, OC-1	51.840 Mb/s	1 DS3 or 28 DS1s
STS-3, OC-3	155.520 Mb/s	3 DS3s or 84 DS1s
STS-12, OC-12	622.080 Mb/s	12 DS3s or 336 DS1s
STS-48, OC-48	2488.320 Mb/s	48 DS3s or 1344 DS1s
STS-192, OC-192	9953.280 Mb/s	192 DS3s or 5376 DS1s
STS-768, OC-768	39,813.12 Mb/s	768 DS3s or 21,504 DS1s

Table 1

Next Generation SONET/SDH networks are implementing significantly more granular transport containers. These containers are allocated and de-allocated at the physical layer and the allocation process is managed by hardware based controllers that operate in near

real-time. This near real-time allocation allows for packet based data to be statistically multiplexed onto the transport network. That is to say, that as one service has traffic bursts, that more of the transport network can be allocated to it, over another service that may be temporarily idle. With many packet services connecting to the transport network, the hardware can provide a more statistically packed service, thus minimizing the amount of wasted bandwidth.

This process does not operate unmanaged. The bandwidth is managed on a per service basis, and is allocated using traffic descriptors. This allows differentiated services, such as VOIP, to have a significantly higher priority, thus preserving the Quality of Service (QoS) that is required for voice

traffic. The traffic descriptors can be allocated to provide a three tiered approach to the bandwidth allotment on a service by service basis. Each service will be assigned a minimum rate, that is the amount necessary to sustain the lowest amount of traffic estimated, or just enough to maintain a circuit. The next tier is the guaranteed maximum rate. This is the amount of the transport network that will be made available to the service with no statistical degradation. The last tier is the maximum best effort. This is the maximum amount of bandwidth that any service attain from the transport network, if it is available. This is also the bandwidth that must be relinquished if other services are requesting guaranteed bandwidth. This could return the service to its maximum guaranteed amount.

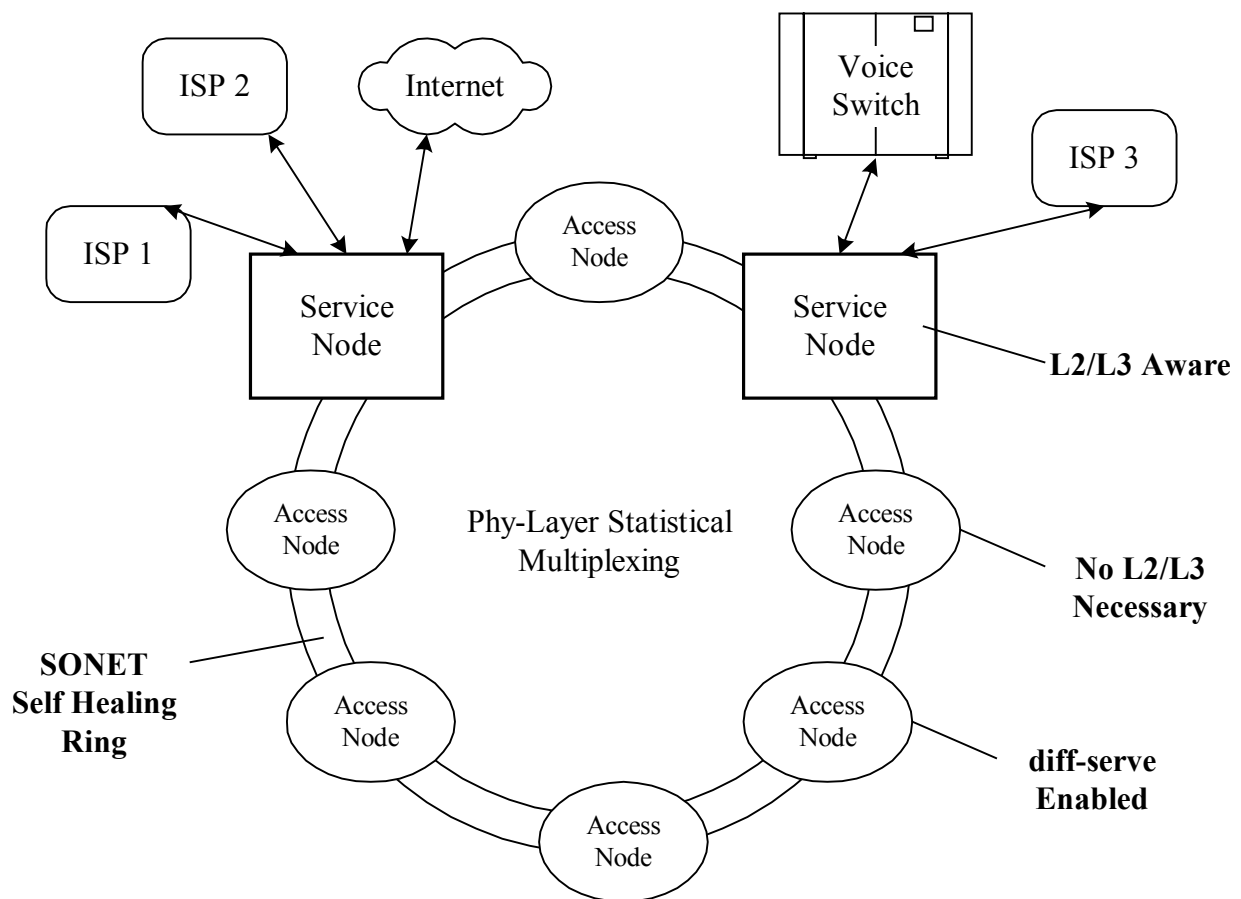


Figure 2

As an example, a 100 BaseT data service may be allocated a minimum of 10 Mbps, a guaranteed rate of 50 Mbps and a best effort rate of 100 Mbps. For a 100 BaseT VOIP service, the minimum rate may be set at 30 Mbps, a guaranteed rate of 100 Mbps, and no best effort. This would be necessary in order to never “drop” a phone call. By building these traffic descriptors on a service by service basis, the operator is able to tailor the statistical nature of the multiplexing on the transport network, and still preserve QoS.

RING SOLUTION

Using SONET/SDH rings to connect each of the Headends, the overlapping logical star becomes a single physical ring connection (Figure 2). With the Hub and Spoke service solution, the Headends become node sites on the ring. The Master Headends act as Service Nodes, and the Secondary Headends act as Access Nodes. The Access Nodes provide multi-protocol, multi-service I/O interfaces for connection with the traffic generating HFC equipment such as the CMTSSs, or Host Digital Terminals (HDTs). The traffic from the Access Nodes is Classified, then policed and shaped prior to being multiplexed onto the transport network. The individual services are sent to the appropriate Service Node. The Service Node provides the necessary Layer 2 switching, or Layer 3 routing of the incoming services to the appropriate destination.

Since the Access Nodes can greatly outnumber the Services Nodes, and there is no initial need to provide L2/L3 capability at the Access Node (because most of traffic is destined for service node), the management of

this network can be significantly reduced. Furthermore, since the multiplexing is managed at the hardware level, and there is a system wide manager of such, the allocation of the bandwidth to services is performed in an unbiased nature. This also greatly simplifies the management which is already a benefactor of the existing SONET/SDH configuration policies.

The Service Nodes essentially become the brains of the Metropolitan network. They are the interface to the critical services, and provide the demarcation for multiple ISPs. If necessary, any Access Node is capable of becoming a Service Node. Also, Service Nodes can either provide unique functionality at each site, or provide replication of services from another Service Node, or simply be a hot standby.

Since the statistical multiplexing takes place at the physical layer, transmission through intermediate nodes around the ring does not compromise QoS. With input classification, prioritization, and traffic descriptors building the transport network rules and regulations, it is possible to provide effective and efficient utilization of the Metropolitan Network bandwidth. In doing so, minimizing the amount of necessary equipment and increasing the ease of management of the network, operators can capitalize on this emerging technology today, and scale it to future needs as they arise.

Erik C. Metz
Systems Engineer
Motorola BCS
101 Tournament Drive
Horsham, Pennsylvania 19044
(215) 323-2284 (phone)
(215) 323-2349 (fax)
emetz@gi.com

John Brenneman
Director of Systems Engineering
Geyser Networks
535 Del Rey Avenue
Sunnyvale, California 94085
408-331-7133 (phone)
408-830-9099 (fax)
jbrenneman@geysernetworks.com