# Yield Management: Turning Bandwidth into Bucks

Robert F. Cruickshank III
Daniel J. Rice
Stargus, Inc.

*Abstract*

*This technical paper examines how analyzing customer bandwidth consumption and optimizing MSO RF network capacity can provide detailed insight and actionable information to increase the yield (i.e. utilized traffic carrying capacity) and profitability of DOCSIS™ networks.*

*This paper begins with a discussion of current bandwidth measurement approaches and their challenges. An alternative is then discussed and results are presented—from a software-based approach that uses the deployed cable modems and CMTSs as network 'sensors' to collect and analyze the upstream and downstream usage (and many other variables) on a per-subscriber level in hourly increments. This information is critical for enabling cable operators to assess the impact of Peer-to-Peer (P2P) and other applications and overall network performance for Voice over Internet Protocol (VoIP).*

*This paper then provides a detailed overview of the process for optimizing DOCSIS networks to improve network capacity, performance, availability, and reliability by recommending and setting the appropriate configuration parameters for each CMTS RF Interface. Given the dynamic nature of HFC networks, deployment experience demonstrates that a single channel width, modulation and error correction configuration (i.e. a single modulation profile) is not effective for all RF interfaces. Field results show that network optimization has at least a doubling effect on the capacity and performance of DOCSIS networks.*

## INTRODUCTION

Cable Operators (MSOs) around the world have built a far-reaching DOCSIS-based infrastructure that increasingly supports voice, video and data services. In the first quarter of 2003 about 1.5 cable modems were installed every second during business hours worldwide[i]. By the end of 2003 there were over 30 million modems deployed worldwide.[ii] As an industry, to maximize yield and profitability, we must understand the Customer Experience, what our customers' do to our networks, and how our networks react to our customers' actions.

## CURRENT BANDWIDTH MEASUREMENT APPROACHES AND THEIR CHALLENGES

Consumption monitoring in DOCSIS networks describes the process of measuring the amount of bandwidth resources consumed by each subscriber in the network, then processing those measurements in order to support operational, accounting or other business practices. There are four general steps common to any Internet-based consumption monitoring technology:

Collection: The process of accurately and reliably harvesting bi-directional consumption data from the network element layer on a per-host or per-subscriber basis.

Processing: The application of algorithms to the raw collected data to ensure data-integrity, remove measurement protocol overhead and compress element data into a processed consumption detail record.

Persistence: The data storage required for consumption information in order to support business practices. This persistence timescale can range from the order of minutes to years.

Presentation: The ability to present consumption information to various organizations or facilities within the OSS infrastructure. This includes the automated distribution of consumption records in order to enable billing mediation, traffic engineering, capacity planning, marketing research, abuse detection and other accounting management practices.

Although some technologies currently exist to address these requirements in the Internet today, almost all of them fail to meet the unique technical challenges inherent in the monitoring of large-scale DOCSIS networks. This section compares three methods for the monitoring of per-subscriber, bi-directional data consumption in order to determine the optimal method for DOCSIS networks.

The three methods include:

1. Intrusive packet capture
2. Non-intrusive packet capture
3. Non-intrusive element polling

Intrusive Packet Capture

The intrusive packet capture method collects network usage data by use of passive "probes" inserted directly into the data path. All packets passing through the probes are captured and processed in order to determine per-host or per-user consumption.

*Strengths*:

DOCSIS protocol independent: The DOCSIS 1.0 protocol does not include a mechanism to furnish the capture of application type (by port number). Intrusive packet capture does not rely on the DOCSIS protocol and therefore provides application layer visibility in pre- DOCSIS 1.1 and proprietary cable modem networks.

Vendor independent: Does not rely on any proprietary functionality in the CMTS or backbone network elements (switches, routers). Data can be collected without interfacing to network elements.

*Weaknesses*:

A lot of missed traffic: Although much of DOCSIS network traffic is forwarded to/from the Internet, a majority of traffic may remain local to the CMTS, forwarded either within the same MAC domain, or to other MAC domains contained in the CMTS (e.g. on college campuses). This local traffic will only increase as next generation DOCSIS CMTSs evolve to higher densities (up to 100,000 CMs) and more applications, such as IP telephony, peer-to-peer, and video conferencing, run over the network. Because the intrusive probe typically captures traffic upstream of the CMTS, it cannot see this local intra-CMTS traffic.

Scalability: Because intrusive capture is typically based on inserting a probe in all possible paths to each CMTS, a 1:1 relationship between CMTSs and probes results. This leads to a very high number of additional network elements making this an operationally complex and economically unattractive method.

Point of failure: Because the intrusive probe is inserted directly into the data path, installation requires a scheduled network outage. After installation, network availability becomes dependant on the availability of the probe. Any redundancy or hot-standby solution for high availability would double the number of probes, resulting in two probes per CMTS.

Dynamic addressing: Because the packets are captured by the probe beyond the DOCSIS segment, only a CPE's IP address can be used to determine the packet's origin. Due to the nature of DOCSIS networks, CPE IP addresses are generally dynamic which does not create an authoritative relationship between packets and their CM origin (Hardware Address). In order to resolve this, integration with the provisioning system is required, adding additional cost to the solution and creating an opportunity for poor usage-data origin integrity.

Physical co-location: This method requires that packet capture probes are co-located with each CMTS. If the MSO supports a distributed HFC architecture (digital hubs, micro-head ends) additional rack space, installation and maintenance costs are incurred. Port unknown. Although application layer visibility is provided through this method, a large percentage of packets are not associated with well-known ports. In addition, popular peer-to-peer applications use configurable ports making the application type invisible.

DOCSIS unaware: Because the packet capture is conducted upstream of the DOCSIS segment at layer 3 and above, this method does not provide any layer 2 (DOCSIS MAC) visibility. Both Service Identifiers (SID, DOCSIS 1.0) and Service Flows (SF, DOCSIS 1.1) are invisible to the probe. As a result, each packet cannot be associated with a DOCSIS service profile or packet classifier in order to determine the quality of service it has been assigned.

While it is possible to move the intrusive packet probes further upstream in the network to minimize the number of probes required, this drastically increases the amount of traffic that the probes miss and defeats the purpose of consumption monitoring. It is also not clear whether the probes' I/O cards can handle the faster line speeds upstream in the network without turning into a bottleneck.

Non-Intrusive Packet Capture

The non-intrusive packet capture method assumes that all CMTS traffic is aggregated at an edge switch or router. The probe is attached to a port on the switch or router and promiscuously captures packets passing through all interfaces (i.e. port spanning). Note that probes could be attached to the CMTS devices themselves, but this configuration would result in a 1:1 relationship between probes and CMTSs, offering a solution that is not economically viable.

*Strengths*:

DOCSIS protocol independent: Has the capability to capture application layer traffic independent of DOCSIS protocol version.

No point of failure: Unlike the intrusive approach, this method does not introduce a point of failure into the network.

*Weaknesses*:

A lot of missed traffic: Like the intrusive packet capture method, the non-intrusive method is unable to view the traffic local to the CMTS. As stated, a configuration capable of capturing all packets would require a single probe provisioned for each CMTS. Again, this would prove economically unfeasible.

Third party vendor dependent: Assuming the probe is attached to the edge switch aggregating CMTS traffic for a headend, the probe depends on proprietary functionality in the switch that forwards all CMTS traffic to the capture probe. The extent to which all CMTS, router and switch vendors implement and support a proprietary form of this functionality is not certain and introduces the

possible need for further development and integration work.

Network performance impact: The additional resource burden on the aggregation switch to direct all inbound packets to the capture port is not negligible. This effort may impact the performance of the switch or router's packet forwarding capability in large-scale production networks.

Scalability: Although this method requires fewer boxes than the intrusive approach, there are still a high number of probes required to ensure that all inter-CMTS traffic is captured.

Dynamic addressing resolution: Like the intrusive approach, the mapping of traffic to source CM requires integration with the provisioning system and introduces potential data integrity issues.

Physical co-location: This method requires that the packet capture probe is co-located with the aggregation switch or first upstream router. If the MSO supports a distributed HFC architecture (digital hubs, micro-headends) additional rack space, installation and maintenance costs are incurred.

Port unknown: Although application layer visibility is provided through this method, a large percentage of packets are not associated with well-known ports. In addition, popular peer-to-peer applications use configurable ports making the application type invisible.

DOCSIS unaware: Because the packet capture is conducted upstream of the DOCSIS segment at layer 3 and above, this method does not provide any layer 2 (DOCSIS MAC) visibility. Both Service Identifiers (SIDs, DOCSIS 1.0) and Service Flows (SF, DOCSIS 1.1) are invisible to the probe. As a result, each packet cannot be associated with a DOCSIS service profile or packet classifier in order to determine the quality of service it has been assigned.

While it is possible to move the non-intrusive packet probes further upstream in the network to minimize the number of probes required, this drastically increases the amount of traffic that the probes miss and defeats the purpose of consumption monitoring. It is also not clear whether the probes' I/O cards can handle the faster line speeds upstream in the network without turning into a bottleneck.

Non-Intrusive Element Polling

Because of the limitations of both intrusive and non-intrusive Packet Capture, we propose a new method known a non-intrusive element polling. Non-intrusive element polling leverages the network management instrumentation already embedded in the DOCSIS network and the remote collection of consumption data through the use of the Simple Network Management Protocol (SNMP).

*Strengths*:

No missed traffic: Unlike both intrusive and non-intrusive packet capture methods, the polling method derives consumption data directly from the entire DOCSIS network down to the individual CPE device (CM, MTA, etc.) level. As a result, there are no missed intra-CMTS packets.

No point of failure: Unlike the intrusive approach, this method does not introduce additional points of failure into the network.

Scalability: Relative to the other methods, a small number of probes are required to collect data from the entire network.

No physical co-location: Remote polling is conducted over IP allowing for the probes to exist at any location within the network

infrastructure. This results in reduced configuration, operations and maintenance costs.

Vendor independent: Non-intrusive element polling is independent of vendor infrastructure (CMTS, switches, routers) and does not rely on proprietary vendor functionality to monitor consumption.

DOCSIS aware: Unlike both packet capture methods, the element polling method has been adopted by the cable industry for the monitoring of DOCSIS networks. As a result, polling provides visibility into layer 2 of the DOCSIS network and can monitor service flows and quality of service in the DOCSIS segment. With the adoption of DOCSIS 1.1, application layer visibility will be embedded in the CMTS.

Functionally extensible: Because of the nature of SNMP and the extensibility provided by MIBs, the polling method can be used to gather other operationally significant data from the DOCSIS network related to traffic consumption (interface utilization, etc.).

*Weaknesses*:
DOCSIS protocol dependent: Application-layer visibility is introduced in DOCSIS 1.1 networks. Meanwhile, because DOCSIS 1.0 limits traffic visibility to layer 2 (octets in/out by SIDs),

Performance impact: If not engineered properly, management traffic generated through polling the DOCSIS network can negatively impact network performance.

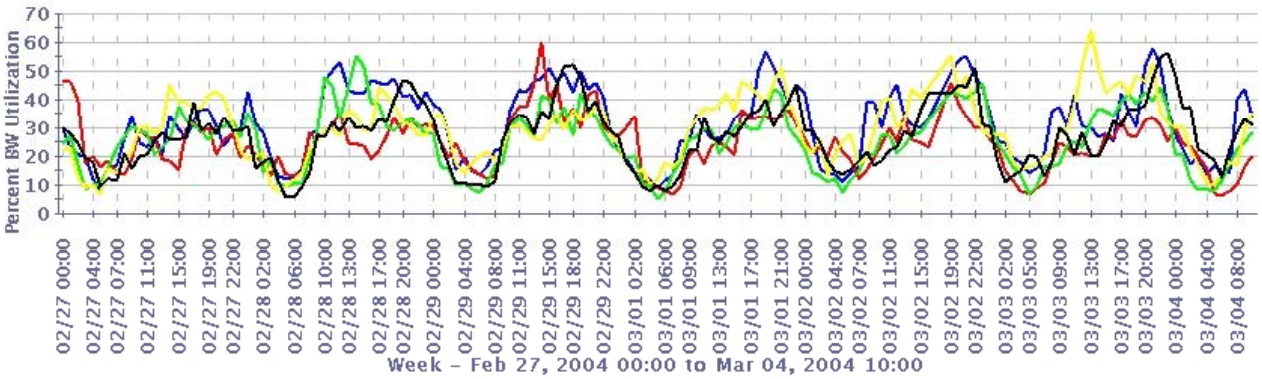


Figure 1: Bandwidth Utilization of Top 5 Most Congested of 88 Downstream Interfaces

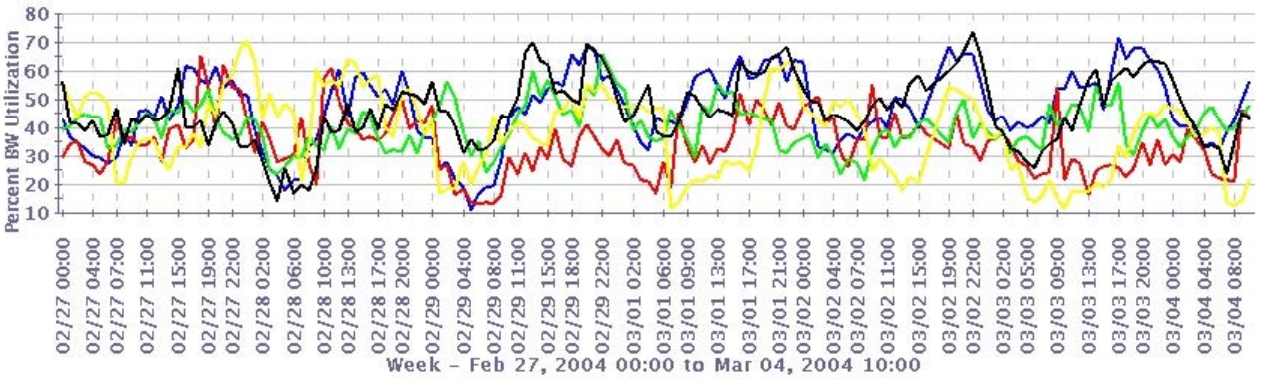


Figure 2: Bandwidth Utilization of Top 5 Most Congested of 450 Upstream Interfaces

## CURRENT BANDWIDTH MEASUREMENTS

Once Consumption data is collected with non-intrusive element polling, analytics may be applied to determine which downstream and upstream interfaces are most congested. Figures 1 & 2 show an MSO's most congested channels; note how much of the time the channels are way underutilized! Notice the daily peaks and valleys in utilization. During the hours after midnight and before sunrise the DOCSIS network is way underutilized. Considering that Figures 1 & 2 show the most congested channels, every other channel is even less utilized than the ones shown here. In summary, there certainly is a lot of "fallow" capacity throughout this network (empty transmission opportunities that are not being used to transport customer data). Surely the pipes can be "filled up" more often and yield management increased. It would make business sense to find ways to "fill the pipes" during off-peak periods, as Telephone Companies did years ago when they used to discount, incent and encourage off-peak usage (in their case after 5 PM).

## OTHER IMPORTANT FINDINGS

Non-Intrusive element polling may be used to collect other DOCSIS parameters (in addition to consumption) that further show how the network reacts to customers' actions. For example, every packet traveling upstream from the customer typically has Forward Error Correction (FEC) redundant coding applied which attempts ensure that bit-wise packet errors (due to noise, etc.) can be "recovered" without retransmission.

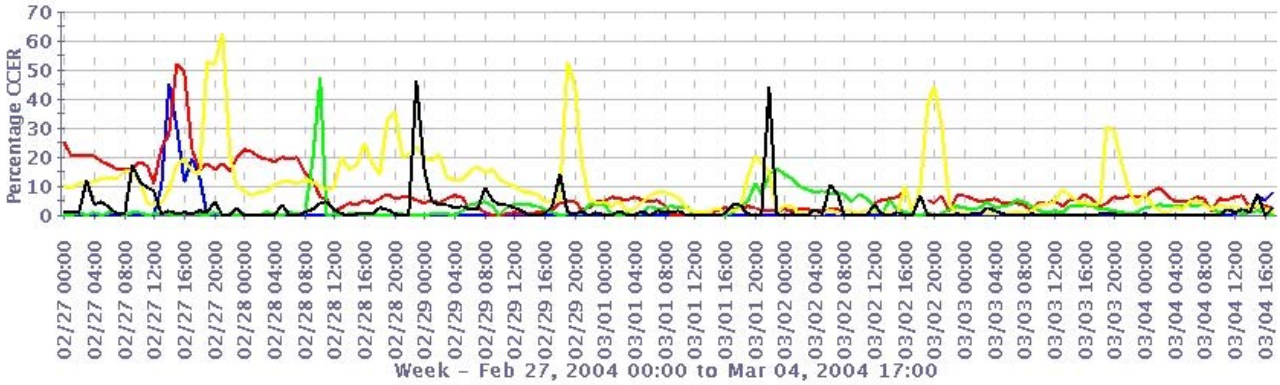


Figure 3: Correctible Error Rate of Top 5 Most Errored of 450 Upstream Interfaces

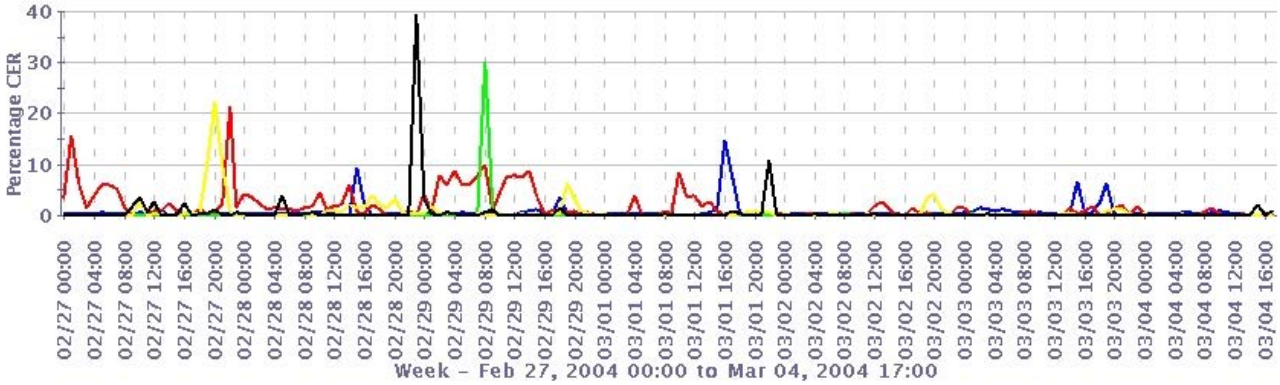


Figure 4: Uncorrectable Error Rate of Top 5 Most Errored of 450 Upstream Interfaces

The derived DOCSIS parameter known as Correctible Codeword Error Rate (CCER) shows the percentage of packets that have been "recovered" from an errored state. Figure 3 shows the worst upstream CCER interfaces. Notice that at times as many as 40%–60% of all packets on certain channels were being recovered automatically by the DOCSIS network.

Another derived DOCSIS parameter known as Uncorrectable Codeword Error Rate (CER) shows the percentage of packets that have been "unrecovered" and are forever lost in the VoIP case (or possibly retransmitted for other TCP/IP applications). Figure 4 shows the worst upstream CER interfaces. Any time that any of the interfaces peaks above 3% CER, VoIP call opportunities are lost[iii]. Notice that the interface indicated by the top (red) line in Figure 4 would not have been able to support any VoIP calls for any subscribers during most of the first half of the week.

## OPTIMIZING DOCISIS NETWORKS

Today many last mile networks are operating with impairments and inefficiencies that are tolerated by email, web surfing, and other less data intensive or less time critical applications. Because of the FEC built into DOCSIS and the resilience of Transmission Control Protocol (TCP), these errors are largely unobserved by and unaffecting to broadband customers that are unaccustomed to the intermittent start-stop nature of web surfing and email. However, advanced, IP-based services such as VoIP, video-conferencing and streaming audio/video are rendered inoperable by such errors and inefficiencies.

Error levels often get worse over time and are typically associated with transient noise and interference due to HFC plant problems. While email and web surfing mask these low error levels, these latent and worsening HFC plant problems are often undetected for months until customers become exasperated with the performance of their applications. A degraded subscriber experience results in operating inefficiencies such as heightened customer care and network maintenance costs, as well as subscriber churn.

By using the capabilities built into DOCSIS, the errors and inefficiencies can be minimized through configuration optimization, thereby allowing MSOs to defer capital expenditures, reduce operational expenses, and ready their infrastructure for advanced IP-based services, particularly VoIP.

### Optimization Goals

When designing or operating digital data communication systems, there are several goals that help drive system optimization:

Goal 1: Transmit as much data in the shortest amount of time possible through the system.

Goal 2: Transmit this high rate of data using as little of the physical resources (spectral bandwidth and power) as possible.

Goal 3: Transmit this data reliably at a much lower rate of errors than will impact the performance or reliability of any of the services.

Goal 4: Develop and operate this system with as little expense and complexity as possible.

The challenge is that these four goals are not completely independent. The error performance and the capacity of the network are interdependent and must be managed together for a quality customer experience.

Method for Optimization

DOCSIS Cable Modems (CM), Multi-media Terminal Adaptors (MTA), Advanced Set Top Boxes (ASTB), and Cable Modem Termination Systems (CMTS) can all be utilized to detect and manage errors while providing bandwidth intelligence data. In so doing, the DOCSIS network can be configured to "four wheel drive" through most service affecting errors – errors that otherwise result in degraded or complete loss of service to subscribers - while maintaining optimal bandwidth capacity. This can be accomplished while notifying operators of degraded network quality before it becomes service impacting, as it occurs, and also providing isolation and identification of the faults.

CMTS vendors ship their equipment with default modulation profiles that are extremely conservative, and significant opportunities exist to reap additional capacity and error protection from DOCSIS networks based on actual network conditions. The DOCSIS 1.0 – 1.1 – 2.0 specifications compromise a progression of features that result in ever-increasing efficiency and capacity, but configuring these networks' elements has become increasingly complex. As a result, capacity and error protection optimization techniques have become increasingly critical to successful deployments of VoIP services.

The technical expertise required to manually adjust DOCSIS parameters is significant. Moreover, RF levels within HFC networks are prone to fluctuations, both periodic and random, and continual use of an automated system to monitor levels can reduce the operational expense of attaining and maintaining optimal configurations. There are many different "knobs" and "levers" that are available in DOCSIS networks that can be tuned to enable capacity optimization and packet error protection, including many parameters for the downstream and upstream signal path. All these parameters are dependent on one another and optimizing them needs to be considered as a collective task. For example, increasing the symbol rate without optimally setting the mini-slot size and codeword structure will result in much less capacity gain than would be expected. Additionally, setting the mini-slot size incorrectly can make large PDUs impossible to transmit.

This process is iterative and continual – certain parameters such as FEC ("k" codeword size and "t" correctable bytes), SLC (Shortened Last Codeword), and max burst size should be adjusted and the results analyzed over a period of days. After the results are known, further adjustments can be planned, including the analysis of modulation order, minislot size, and symbol rate. Over time as RF levels fluctuate, periodic adjustments will ensure optimal network performance. Using this technique of methodical and measured change, an optimal network configuration can be reasonably and cost-effectively attained and maintained without undue risk of network instability, but is necessary to constantly collect and analyze network data.

## PROTECTED TRAFFIC AND NEWLY AVAILABLE TRAFFIC CAPACITY

Adjusting the many DOCSIS configuration "knobs" and "levers" in an automated way results in the right level of packet error protection for all traffic and a huge gains in overall traffic capacity. Field experience shows that small portions of the HFC network require extreme error protection and low transmission speeds – while the vast majority of the network can run at extremely high speeds. Figure 5 shows a relatively noise-free interface whose configuration, over the course

of a month, was changed from QPSK @ 1.6 MHz to 16QAM @ 3.2 MHz. The upper jagged (pink) line represents the channel utilization, indicated on the right hand vertical axis, which initially peaks at 80%-90%. With the added capacity from optimization utilization is reduced to routinely below 25%. Moving downward, the flat (thin light blue) line represents the unique CMs (subscribers) on this upstream, the jagged (thick dark blue) line shows latency due to high utilization (peaking every evening as expected), and lastly the lower oscillating (thick red) line shows the number of CMs (subscribers) active over time.

Experience shows that virtually all upstreams can be optimized to operate error-free for high-quality, high availability VoIP service. In addition, as much as 2/3 of all HFC upstreams can be optimized to operate at much higher transmission speeds—which further incents MSOs to consider ways to increase yield management.

## CONCLUSION

In order to maximize yield and profitability, we must understand what our customers' do to our networks, and how our networks react to our customers' actions. Non-intrusive element polling is the best way to measure the impact of customer traffic—and the corresponding response of our networks. Measurements show that our networks are often underutilized and that we as an industry have opportunities to increase yield management.

Non-Intrusive polling can also collect DOCSIS performance metrics that can be utilized by an automated system that optimizes HFC packet error protection and network capacity so that VoIP works reliably throughout the network. This optimization results in overall double or better capacity gains—resulting in even greater opportunities for yield management.

---

[i] Source: Cable DataComm News Report, 2003

[ii] Source: In-Stat/MDR Report, 2003

[iii] "Network Tolerance of Highly Degraded Conditions for an H.323-based VoIP Service", Peter Holmes, Lars Aarhus, Eirik Maus, Norwegian Computing Center, P.O. Box 114, Blindern, Oslo, Norway

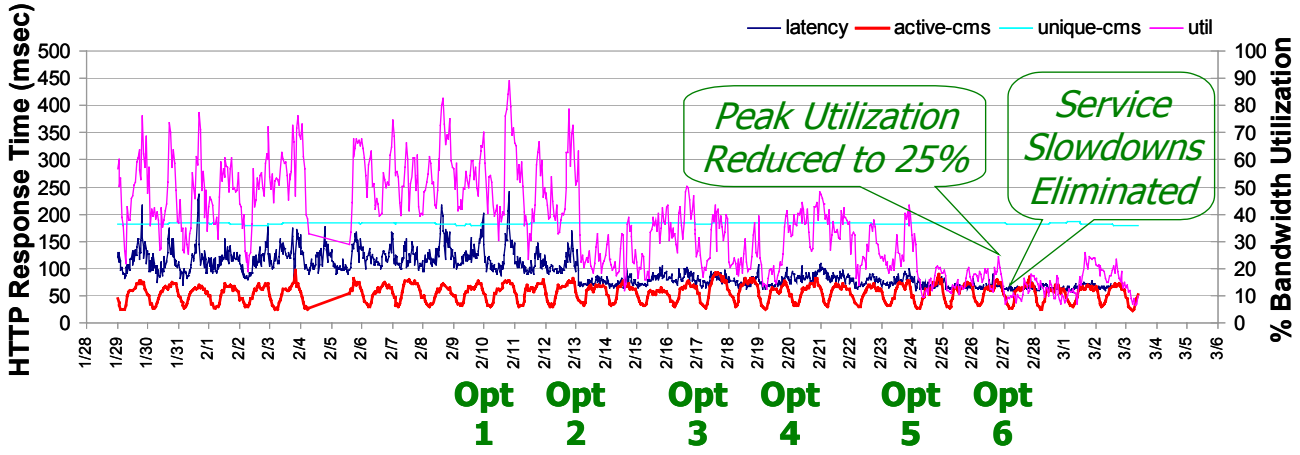


Figure 5: Bandwidth Utilization (right), HTTP Response Time, Unique & Active CMs (left) of a Typical Optimized Upstream Interface Undergoing 6 Iterative Optimizations