NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

QUALITY OF SERVICE FOR IP-BASED NETWORKS

by

Konstantinos Sambanis

March 2001

Thesis Advisors: Gilbert M. Lundy
Rex A. Buddenberg

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 2001	Master's Thesis

4. TITLE AND SUBTITLE: Quality of Service for IP-based Networks	5. FUNDING NUMBERS
6. AUTHOR (S) Sambanis, Konstantinos	
7. PERFORMING ORGANIZATION NAME (S) AND ADDRESS (ES) Naval Postgraduate School Monterey, CA 93943-5000	8.PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING / MONITORING AGENCY NAME (S) AND ADDRESS (ES)	10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the Hellenic Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; distribution is unlimited.	

13. ABSTRACT *(maximum 200 words)*

In recent decades, the networking community has been looking for strategies to converge over a single common network infrastructure carrying voice, video and data. The pervasive and ubiquitous packet-based IP network provides the most convenient platform for the desirable convergence, where resources can be managed in an efficient and dynamic manner.

The gradual convergence into the IP infrastructure introduces multimedia-rich and interactive applications that are bandwidth-intensive and delay-bound, while more sophisticated data applications are deployed that place new demands onto IP networks. The IP-based network is evolving to satisfy the requirements of traffic differentiation and reliable service. Quality of Service (QoS) mechanisms are introduced to meet the traffic expectations and enhance the basic service model of the network in many subtle ways.

This thesis provides a comprehensive examination of QoS mechanisms and protocols that have surfaced to optimize the utilization of network resources, to provide differentiated treatment of traffic and enforce the appropriate policies. The study proposes a balanced approach of bandwidth increase and integration of robust QoS techniques into existing IP network infrastructure to arrive at a convergent, multiservice and scalable telecommunications network. Findings from this thesis can be incorporated into the design and implementation of an integrated network within a large organization that will deliver accurate services and defined level of performances.

14. SUBJECT TERMS Networking, Convergence, Quality of Service, IP Multiservice Network, Policy-based network, Traffic Management	15. NUMBER OF PAGES 94
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE IS INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

QUALITY OF SERVICE IN IP-BASED NETWORKS

Konstantinos Sambanis
Lieutenant, Hellenic Navy
B.S., Hellenic Naval Academy, 1989

Submitted in partial fulfillment of the
requirements for the degrees of

Master of Science in Computer Science
and
Master of Science in Information Technology Management

from the

NAVAL POSTGRADUATE SCHOOL
March 2001

Author: _____
Konstantinos Sambanis

Approved by: _____
Gilbert M. Lundy, Thesis Advisor

_____
Rex A. Buddenberg, Thesis Advisor

_____
Dan Boger, Chairman
Information Systems Academic Group

THIS PAGE IS INTENTIONALLY LEFT BLANK

ABSTRACT

In recent decades, the networking community has been looking for strategies to converge over a single common network infrastructure carrying voice, video and data. The pervasive and ubiquitous packet-based IP network provides the most convenient platform for the desirable convergence, where resources can be managed in an efficient and dynamic manner.

The gradual convergence into the IP infrastructure introduces multimedia-rich and interactive applications that are bandwidth-intensive and delay-bound, while more sophisticated data applications are deployed that place new demands onto IP networks. The IP-based network is evolving to satisfy the requirements of traffic differentiation and reliable service. Quality of Service (QoS) mechanisms are introduced to meet the traffic expectations and enhance the basic service model of the network in many subtle ways.

This thesis provides a comprehensive examination of QoS mechanisms and protocols that have surfaced to optimize the utilization of network resources, to provide differentiated treatment of traffic and enforce the appropriate policies. The study proposes a balanced approach of bandwidth increase and integration of robust QoS techniques into existing IP network infrastructure to arrive at a convergent, multiservice and scalable telecommunications network. Findings from this thesis can be incorporated into the design and implementation of an integrated network within a large organization that will deliver accurate services and defined level of performances.

THIS PAGE IS INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE IS INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS

ACELP	Adaptive Code Excited Linear Prediction
ADPCM	Adaptive Differential Pulse Code Modulation
AF	Assured Forwarding
ATM	Asynchronous Transfer Mode
AVVID	Architecture for Voice, Video and Integrated Data
CLI	Command Line Interface
CODEC	Coder-Decoder
COPS	Common Open Policy Service
DHCP	Dynamic Host Configuration Protocol
DiffServ	Differentiated Services
DNS	Domain Name Service
DSCP	DiffServ Code Point
DSL	Digital Subscriber Loop
EF	Expedited Forwarding
FDDI	Fiber Distributed Data Interface
FIFO	First in First Out
FTP	File Transfer Protocol
Gbps	Gigabits Per Second
HDTV	High Definition Television
HTML	Hypertext Transfer Language
HTTP	Hypertext Transfer Protocol
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IntServ	Integrated Services
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ITU	International Telecommunication Union
Kbps	Kilobits Per Second
LAN	Local Access Network
LDAP	Lightweight Directory Access Protocol
MAC	Media Access Control
Mbps	Megabits Per Second
MPEG	Moving Picture Experts Group
NGI	Next Generation Internet
PBX	Private Branch Exchange
PCM	Pulse Code Modulation
PDP	Policy Decision Point
PEP	Policy Enforcement Point
PHB	Per Hop Behavior
QoS	Quality of Service
PSTN	Public Switched Telephone Network

RAP	Resource Allocation Protocol
RED	Random Early Detection
RF	Radio Frequency
RFC	Request For Comment
RSVP	Resource Reservation Protocol
RTCP	Real Time Control Protocol
RTP	Real Time Transport Protocol
SMTP	Simple Mail Transfer Protocol
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TOS	Type Of Service
UDP	User Datagram Protocol
VLAN	Virtual LAN
WAN	Wide Area Network
WDM	Wave Division Multiplexing
WFQ	Weighted Fair Queuing
WRED	Weighted Random Early Detection

ACKNOWLEDGMENTS

I would like to express my gratitude to professor Gilbert Lundy for his priceless guidance, supervision and support. I would also like to thank professor Rex Buddenberg for his considerate and practical contributions. Without their aid, this thesis would not have been possible.

THIS PAGE IS INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. NETWORKS OF THE 1990S

Traditional voice networks are circuit-switching networks that are built to provide an optimal service for time-sensitive voice conversation, requiring low delay, low delay variation (jitter), and constant but low bandwidth of 3 KHz in analog form and 64 Kbps in digital form. A key characteristic of such networks is that resources are dedicated to a particular connection for the entire duration. This approach is nice for the users, but fails to utilize efficiently the limited network resources. Additionally, these networks do not provide the high bandwidth requirements for video and high-speed data traffic.

The data networks and the Internet, based on a connectionless packet-switching scheme, accomplish a more robust and dynamic handling of resources. The TCP/IP protocol suite, which lies at the heart of public and private data networks, was designed to provide a best-effort service model for the data applications. It is this model that allowed the IP-based data networks to grow exponentially; this also led to a network infrastructure that falls short of delivering tight performance guarantees needed for delay-sensitive communication.

A third type of networks is the cable TV network deployed over coaxial wiring scheme that can deliver much higher bandwidth up to 36 Mbps [Ref. 1], through cable modems. This network was designed and optimized to carry TV signals to users' premises in the downstream direction only. To carry interactive voice and data, proper switching equipment is installed that permits bi-directional communication across the cable.

The recent decades, the networking community has been looking for strategies to converge over a single common network infrastructure carrying voice, video and data.

Since the mid-1980s, the need for a universal packet-based network has been identified, where resources can be managed in a more efficient and dynamic manner. After several convergence efforts, it appears that the pervasive and ubiquitous IP-based data network will provide the common denominator and the most convenient platform for the desirable convergence.

The gradual convergence into the IP infrastructure introduces multimedia-rich and interactive applications that are bandwidth-intensive and delay-bound. Additionally, more sophisticated data applications are deployed that place new demands onto IP networks. It is not only the introduction of a greater volume of traffic, but also a greater diversity of traffic characteristics. Not all of the streams are equally important. Each presents unique performance requirements, but all expect an excellent network service. While these requirements can be quantified using several criteria that incorporate performance, availability, reliability, and security, in the context of this thesis, they are defined in terms of data rate, delay, jitter, and packet loss. The IP-based network is evolving to satisfy the requirements of traffic differentiation. Quality of Service (QoS) mechanisms are introduced to meet the traffic expectations and enhance the basic service model of the network in many subtle ways, enabling the reliable and predictable service to differentiated network traffic.

Delivering QoS in a multiservice and convergent IP-based network presents several challenges. These challenges require a clear understanding of the current network environment, analysis of existing and emerging technologies, potential upgrade of the infrastructure and efficient implementation of flexible mechanisms that can deliver the desired QoS to overcome current performance deficiencies.

B. RECENT ADVANCES IN QUALITY OF SERVICE

The need for QoS deployment in IP network has brought a lot of attention in recent years. Research and education communities have perceived the importance of the

current and future sustainability of the IP-based networks. The U.S. government under the Next Generation Internet (NGI) is playing a key role in supporting research and development of high-performance network technologies and services.

The Internet II project, led by the private sector and universities, runs in parallel with the federal NGI initiative. One of the most important studies in the Internet II project is on the QoS. The two projects stress the necessity of network resource allocation and management for different kinds of applications.

From the standards perspective, Internet Engineering Task Force (IETF), the principal body engaged in the development of new Internet standard specifications, is actively embracing the issues around QoS, especially the "Operations and Management Area" working group. Recent work in the IETF has led to the development of several standards for a QoS-enabled network.

The networking industry, reacting to the growing demand of enterprises that struggle with the issue of how best to ensure that applications receive the service quality that they require, has already begun to put QoS technologies into practice.

C. OBJECTIVE

The objective of this thesis is to investigate the methods of deploying QoS in IP-based networks, identify key issues of QoS implementation and propose a robust QoS deployment that will allow a unified convergent IP-based network to provide reliable, consistent and guaranteed service in a plethora of applications within an organization.

D. ASSUMPTIONS AND LIMITATIONS

This thesis attempts to describe the ongoing effort for improvement of the IP network infrastructure, to explore proposed mechanisms and evaluate enhanced service

models. The findings reported will provide the reader with an understanding of key technologies and techniques that enable QoS and traffic management. Because these are relatively new developments, they need to be brought to the attention of the reader to help him better evaluate existing IP network design and operation and determine how to improve upon traditional best-effort service.

The author assumes the reader has a basic understanding of fundamental networking principles, models and devices. A thorough presentation of all issues related with QoS in IP networks entails employment of a great deal of technical detail, but a systematic effort has been made to present the topic in a simple and clear manner.

Given the breadth of the topic, the scope of this study is limited to QoS in a large private network. This large network can consist of the following three components:

- Campus networks, which connect a building or group of buildings with one or more Local Access Network (LAN) and backbone connections.
- Wide Area networks (WAN), which connect campuses together and may also include radio-WANs that reach to mobile platforms such as ships, which increasingly have campus networks within.
- Remote connections that link remote offices, units and users to a central location or campus.

This enables a comprehensive and holistic analysis of QoS that meets the needs of a large organization, whose QoS-enabled multiservice network can support strategic/tactical operations, rapid access to information, effective decision making, efficient data distribution and knowledge management.

E. ORGANIZATION OF THESIS

In Chapter II, a first task is to analyze the nature of IP-based and voice networks. Then, the trend for convergence of voice and data networks into an IP infrastructure is examined and the needs of differentiated traffic are analyzed.

In Chapter III, QoS is defined and effective QoS mechanisms and techniques are examined that provide prioritization, signaling, congestion control and policing in the IP-based network.

In Chapter IV, the study provides a solid explanation and evaluation of the QoS service models that have been developed to integrate all the QoS mechanisms in a network-wide implementation.

In Chapter V, an integrated end-to-end QoS deployment is proposed that is essentially a balanced approach of bandwidth increase and QoS implementation with minimal overhead and change to the stateless and connectionless nature of the IP network, providing the necessary efficiency and dynamic traffic management. Several current commercial solutions are also examined.

In Chapter VI, conclusions, recommendations and suggested topics for further study are discussed.

THIS PAGE IS INTENTIONALLY LEFT BLANK

II. CONVERGENCE OF IP-BASED AND VOICE NETWORKS

The purpose of this chapter is to analyze the nature of IP-based and voice networks, examine the trend of converging these two networks and understand the requirements imposed in the integrated network by different voice, video and data applications. First, it is necessary to define the network metrics used to measure the performance of the network.

A. NETWORK PERFORMANCE METRICS

The way in which any network handles traffic can be characterized through a set of metrics. These parameters provide a quantitative picture of the traffic performance in the network and in the context of this thesis, the following metrics have been specified:

- Data rate (bandwidth) is the raw data carrying capacity of a network. It is the rate at which traffic is carried by the network from one host to another and is usually measured in bits per second.
- Delay (latency) is the amount of time it takes the network to deliver a packet to its destination and is often measured in milliseconds. While it is very small in voice networks, the majority of the end-to-end delay in data networks is attributed to the switching delay that is introduced by the networking devices.
- Jitter (delay variation) is a term used to describe the variation in arrival times to the destination for different packets and is measured in milliseconds. Jitter is particularly disruptive to voice and video communications that require transmissions at a constant rate. Receiving devices compensate for jitter by setting up local buffers to playback voice and smooth out variability in packet arrival times. There is a tradeoff though because the use of buffering increases latency. Jitter is mainly introduced in the packet switching networks and is due to the different

paths that packets may follow and different queuing times packets experience in the internetworking devices.

- Packet loss is measured in terms of the percentage of the total packets sent. Data applications, such as file transfers, are very sensitive to packet loss, while voice and video applications allow a certain percentage of packet loss. Packet loss may occur because of network link failures, introduction of noise in wireless and RF networks and in packet networks, in case of congestion, the network devices fill up and start discarding packets.

Based on these metrics, the requirements of voice, video and data traffic are summarized below in Figure 1.

Application	Bandwidth	Delay	Jitter	Packet loss
Interactive Voice	Low 5-64 Kbps	200 ms	30 ms	1%
Interactive Video	Medium - High	400 ms	30 ms	5%
Data	Adaptive	Seconds Minutes	Allowed	0%

Figure 1. Requirements of Voice, Video and Data Traffic.

B. THE NATURE OF IP-BASED DATÁ NETWORK

IP-based networks were designed to support data applications that are characterized by bursty traffic with occasional high bandwidth demand and longer delays. These data applications are consolidated over a packet-switching network, built around network devices, such as routers, switches, bridges and hubs as shown in Figure 2.

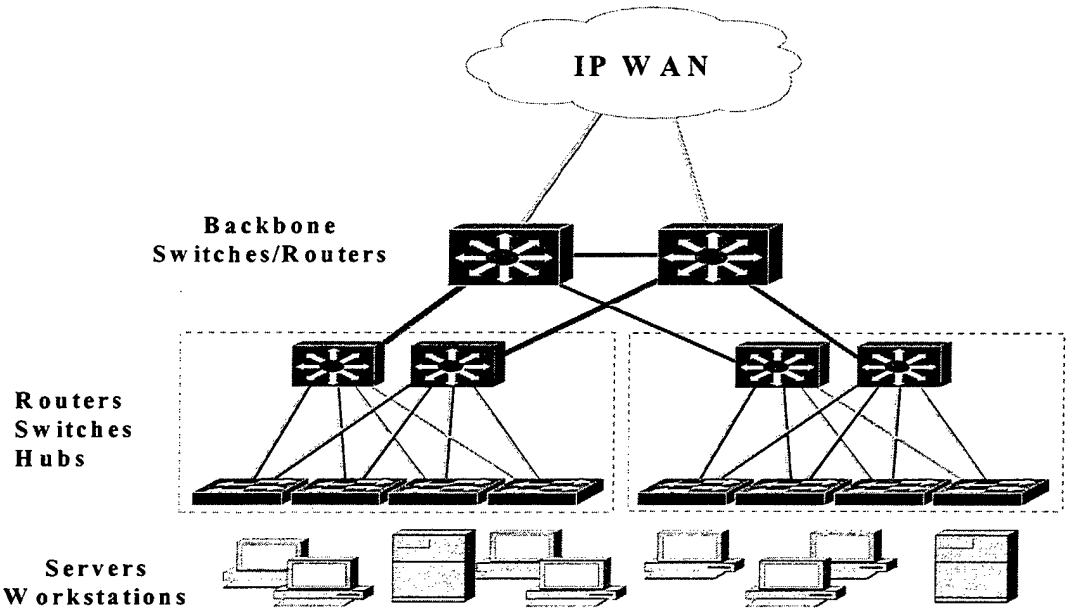

Figure 2. Data Network Infrastructure

1. Internet Protocol

The rapid adaptation of the TCP/IP protocol suite in the enterprise intranets and the global Internet has resulted in the indisputable dominance of Internet Protocol (IP) as the most widely used internetworking protocol. IP represents the third conceptual layer of the TCP/IP protocol suite, shown in Figure 1, and has provided a consistent service interface that has remitted the relatively independent development of applications and underlying networking technology.

IP provides a connectionless service between end systems, routing packets from network to network. It specifies the format of packets sent across the network as well as the mechanisms to forward packets. Currently, the most widely deployed version of this protocol is version 4 (IPv4). However, an enhanced version (IPv6) has been standardized and is ready for deployment.

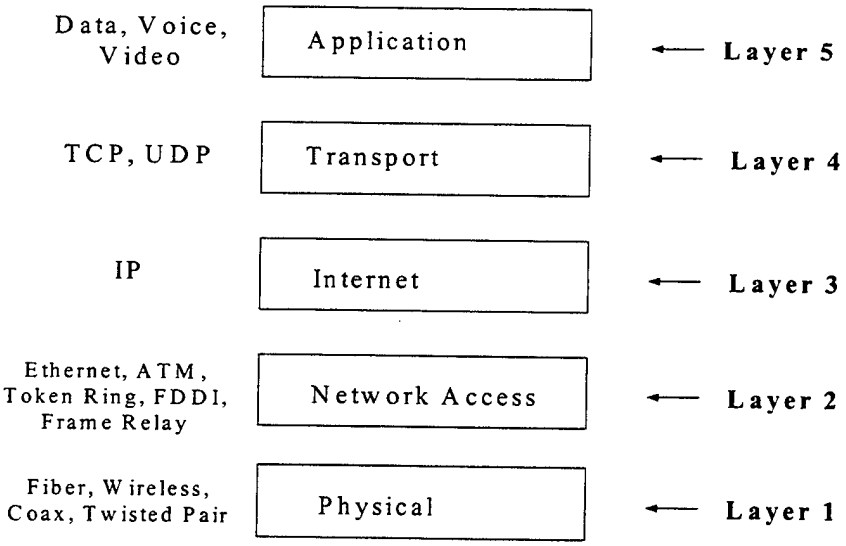

Figure 3. Layers of TCP/IP Protocol Suite.

2. Transport Protocols

The transport protocols provide the basic end-to-end service of transferring data between end hosts. They are the interfaces between the internet and the application layer. The TCP/IP protocol suite includes two transport protocols, the Transmission Control Protocol (TCP), which is connection oriented, and the User Datagram Protocol (UDP), which is connectionless.

TCP is a connection-oriented protocol and is responsible for the reliable, in-order delivery of a stream of data. Since IP is a connectionless protocol that makes no effort to correct transmission errors, TCP is deployed to guarantee that the stream of data leaving the sender will be reassembled intact at the receiver end. It accomplishes that through a system of delivery acknowledgements. It is also responsive to packet loss allowing efficient retransmission of lost information.

TCP is designed to be adaptive to data rates allowed by the network. A TCP sender constantly adjusts itself based on the current level of network performance and introduces traffic into the network accordingly. When the TCP sender determines that a

packet did not make it through the network, it slows the transmission rate. These features make TCP a very reliable protocol, but also a source of overhead for the network.

In contrast to the adaptive, error correcting nature of TCP, UDP is a connectionless transport protocol that makes no effort to guarantee delivery and does not adapt to network congestion. It sacrifices these in return for no setup overhead, no acknowledgement and sequencing of traffic. Because it is connectionless, UDP has very little to do and essentially it only adds a port addressing capability to IP.

3. Connectionless and Stateless Nature

The IP-based network is connectionless and stateless because the IP protocol by nature is a connectionless protocol. The minimal function required from the network is mere connectivity; that is delivery of datagrams from a source to a destination. There are no predefined circuits or fixed paths among the network devices and the end nodes. TCP is normally used for setting up a confirmed connection between two communicating end hosts, but networking devices along the path never look this deep into the passing packets that they route. This principle allows complexity to stay in the end-hosts, so the network can remain relatively simple.

The term stateless means the nodes along the path of the traffic flow do not maintain specific information about the state of each flow. Each datagram is treated independently and equally, with no reference to datagrams that have been processed before. The routers maintain routing tables and forward datagrams according to these tables without keeping track of whether a particular datagram is part of several in a flow from one node to another. Successive datagrams of the same flow may follow entirely different routes to the same destination.

Given knowledge of the ultimate destination of a datagram, the network finds, if possible, a path through any available links to the destination. If network failures or congestions occur (created by lack of resources), there is unpredictable response from the network (a datagram may be delayed, duplicated or discarded). There are no guarantees that any given packet will reach its destination at all and the time it takes the network to achieve packet delivery is a secondary consideration and heavily depends on the switching delay introduced in the networking devices. Essentially, the packet-switching IP-based network is a network of queues that uses a "store-and-forward" approach. An arriving packet is placed in the queue of a networking device until the processor of the device can process and forward the packet to its destination. The processing of a packet may include several steps, such as looking up the routing table, deciding for the correct forwarding interface and manipulating the packet (changing the encapsulation type, changing the hop count). Furthermore, as more traffic is introduced to the network, service demands eventually exceed resources. Nevertheless, the network does not deny service, but instead it degrades its performance gracefully. This service model is called best effort because, although the network makes every effort to deliver datagrams, it makes no guarantees.

This scheme has led in part to the success of the IP-based network and the Internet, providing increased flexibility, resource sharing, robustness, responsiveness and scalability. Nevertheless, it does not guarantee a bounded service with respect to timeliness and preservation of temporal ordering. The IP network focuses more on "where" to send datagrams and little on the "when"[Ref. 2].

C. THE NATURE OF VOICE NETWORKS

1. General

Traditional voice and telephone networks are built to provide an optimal service for time-sensitive voice applications requiring low delay, low jitter and constant but low bandwidth of 3 KHz in analog form and 64 Kbps in digital form.. These networks are

built in a connection-oriented and circuit-switched approach. Communication is achieved by dedicating a communication channel, which is set up prior to information transfer, for the duration of the connection between two nodes.

Initially, analog transmission was used for the transfer of voice. But gradually, the voice network is migrating to digital transmission using pulse code modulation (PCM) or adaptive differential pulse code modulation (ADPCM). In both cases, analog voice signal is converted into digital form by sampling the analog signal 8000 times per second and converting each sample into a numeric code. That results in a data rate of 64 Kbps for a digital voice channel that has been standardized throughout the world.

2. Structure

a. Public Switched Telephone Network

Public Switched Telephone Network (PSTN) is the global telephone voice network and represents the collection of the switching and networking equipment that belongs to carriers who provide telephone services. It refers primarily to the wireline telephone network and its access points to wireless networks, such as cellular and satellite communications. The basic hardware elements of the PSTN are telephone sets, premises wiring, local exchange switches, interexchange carrier switches and trunks between switches [Ref. 3]. The PSTN distinguishes between

- Access lines (also called local loop) that connect telephone sets to local switches.
- Trunking networks that multiplex voice channels between voice switches for long distance access.

The access line is typically a single pair of twisted-pair copper wires, while the backbone trunking network can be carried on coaxial cable, fiber optic cable or microwave towers.

In organizations, the telephone devices are anchored around a voice switch, named Private Branch Exchange (PBX), as illustrated in Figure 4, that is used to interconnect telephones within a building and enable connection into PSTN infrastructure or a dedicated leased line for private use. The links from the PBX to individual telephones are access lines, where voice exists either in analog or digital form, and from the PBX to the central office are direct trunks that carry voice in digital form. Once voice reaches the central office, it exists in digital form on the network, on time division multiplexed channels of 64 Kbps each.

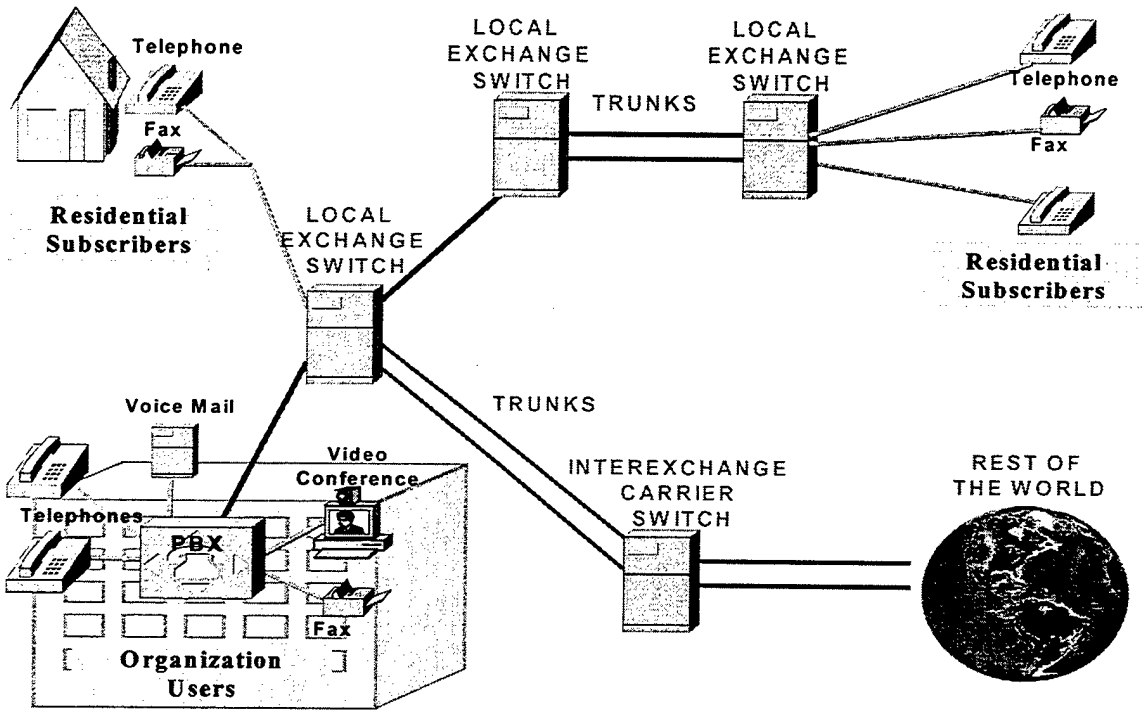

Figure 4. PSTN Infrastructure.

b. Wireless Networks

Electromagnetic radiation can also be used to transmit information. Cellular networks are deployed extensively to support voice communications. These networks do not require a direct physical connection between users. Each participating unit attaches to an antenna, which can both transmit and receive RF. Base stations communicate through radio signals with end users. RF technology can be combined with

satellites to provide communication across longer distances. The wireless networks present difficulties inherent in wireless media, such as interference and limited capacity.

3. Observations

In voice networks, the quality of a call has never been a negotiable parameter. However, delivering this quality comes at a cost because the circuit-switched nature of the network ties up network resources for the duration of an entire call, regardless of the actual bandwidth utilized [Ref. 4].

Furthermore, the whole network has been originally designed and implemented to support small, fixed-bandwidth needs in increments of 64Kbps. Data can be transmitted by the use of proper dialup modems. Although technologies like Integrated Services Digital Network (ISDN) and Digital Subscriber Loop (DSL) allow the transmission of data at higher data rates, the inherent limitations of the twisted pair wiring places upper bounds to the rates at which data can be sent.

Voice networks achieved their dominant position because they were well suited to the analog and digital transmission of voice signals. However, their connection-oriented nature and their support for relatively low bandwidth does not make them well suited today with the increased need for high-speed digital access and the expansion of intensive data and video applications.

D. THE TREND FOR CONVERGENCE IN IP-BASED NETWORK

The recent decades, the networking community has been looking for strategies to merge the two different networks over a single common network infrastructure carrying voice, video and data. Since the mid-1980s, it has been viewed as necessary to replace the TDM public network with a universal packet-based network, where resources can be managed in a more efficient and dynamic manner. After several convergence efforts

(ISDN, ATM), it appears that the industry has landed on a common platform of the IP-based network. The irresistible logic is that digitized voice and video is just another kind of packet data that can be handled by the IP network in a robust and bandwidth efficient manner. IP's pervasiveness and ubiquity in personal computers, servers, workstations, routers and switches makes it the common denominator and the most convenient platform for the support of integrated traffic.

Using a converged network, an organization with geographically dispersed offices and units can reap the following benefits:

- Reduce operating costs by eliminating redundant hardware and wiring required to support separate voice and data infrastructures.
- Enable rapid access to information and effective decision-making.
- Increase network manageability and interoperability.
- Enhance productivity, mobility and efficiency.

Converged IP networks could also enable new compelling applications, such as multimedia call centers that integrate customer messages coming in any format - telephone, fax, paging, voice mail, or e-mail - into a single, centralized and unified messaging system. Other applications that can take advantage of the convergence are distance learning, powerful collaboration tools, integrated directory services and device portability with user ID services.

The merger of voice, video, and data networking is in the early stages of adoption today. Applications, such as Voice over IP (VoIP) and desktop video conferencing that take advantage of these converged communications capabilities have appeared in large organizations. Packetized voice and video is on the rise in the WAN access, successfully bypassing expensive long-haul PSTN connections, as illustrated in Figure 5.

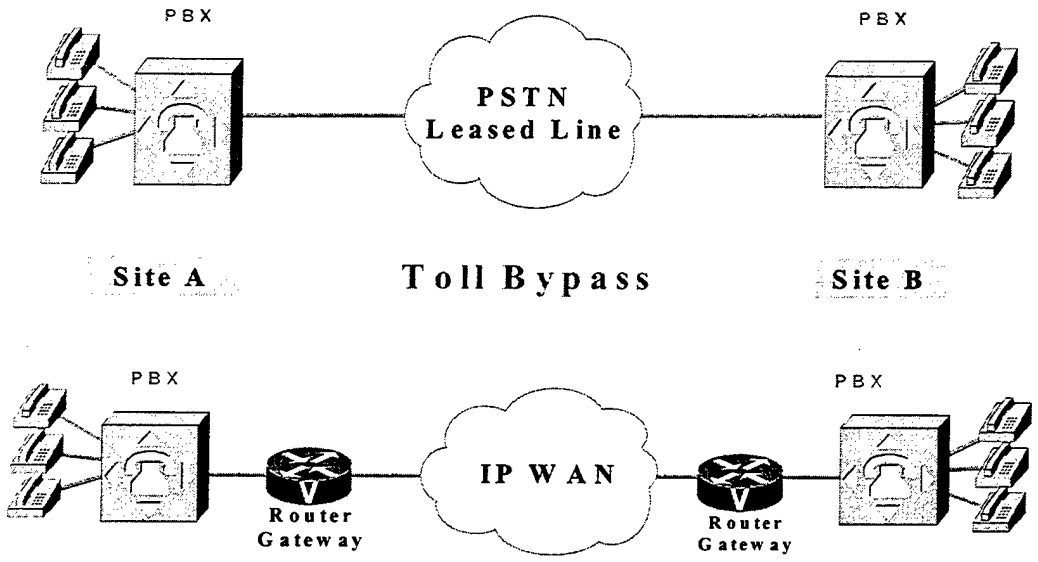

Figure 5. Bypassing Long-Haul PSTN Connections Using IP WAN Access.

Organizations take incremental steps towards an integrated, multiservice network:

- Upgrade to an IP-enabled PBX. This upgrade can be done to the existing traditional PBX inside the organization and essentially involves only the addition of IP line cards with LAN interfaces to a PBX.
- Introduction of IP Gateway that bridges the gap between the IP data network and PSTN, allowing voice calls to traverse voice and data networks seamlessly.
- Introduction of IP Gatekeeper on the IP network that provides call functionality, address translation, admission control and bandwidth allocation.
- Introduction of call servers that perform voice-related applications, such as voicemail.

The IP-based convergent network, as shown in figure 6, is increasingly being adopted in new facilities and offices, where there are no legacy systems to amortize.

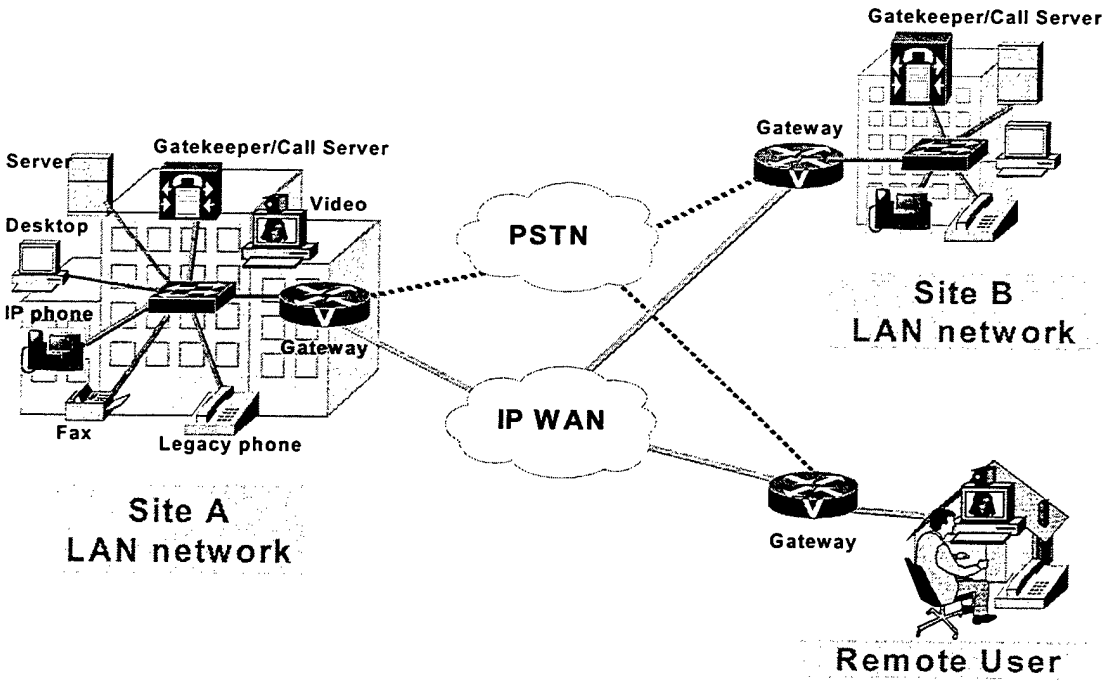

Figure 6. End-to-End Converged IP Network.

E. APPLICATIONS REQUIREMENTS

The gradual integration of voice and video in IP networks, the emergence of new applications and the increased sophistication of existing ones have forced the IP network to carry traffic with diverse requirements. Traditional applications, such as e-mail, file transfer and web browsing and emerging ones, such as voice over IP, desktop video-conferencing, live or on demand streaming media, all compete for network resources. They generate traffic at varying rates and expect the network to handle it in an appropriate and sufficient way.

This evolution places new demands onto the IP infrastructure. The network is expected to accommodate all these applications and provide reliable and deterministic service to the various traffic types. Therefore, the basic service model of the network must be enhanced in many subtle ways to achieve the required performance guarantees.

It is important to understand the requirements of the various applications in order to function correctly. These requirements are expressed using the quantifiable network parameters that have been defined earlier in this chapter and include bandwidth, delay, jitter and packet loss. For some applications, there is always a need for fixed bandwidth, others adapt to consume the maximum available bandwidth, and for others it may vary with time. Certain applications are more or less tolerant of traffic delays in the network and of variation in delay. Other applications can tolerate some degree of traffic loss while others cannot.

1. Elastic / Inelastic Applications

There are a number of categories under which application traffic may fall, depending on how tolerant or intolerant the application is to network congestion and other inconsistent network behavior. Broadly speaking, the two fundamentally different traffic types on datagram networks are elastic and inelastic traffic [Ref. 5].

Elastic traffic originates from applications that run on top of TCP, such as HTTP, FTP and SMTP and can be described as data applications. They always wait for all data to arrive and TCP provides the required reliability and adaptability. Services that run in the background without any user interaction are also typically elastic, such as network services (routing tables updates, network management).

On the other hand, inelastic applications need to have their data communicated in a deterministic and consistent manner. Such time-sensitive data must arrive at its destination on schedule or within a bounded delay. There are further subdivided into tolerant and intolerant, based on the way they react to delay. Inelastic tolerant applications, such as streaming media and video-on-demand do expect their data to arrive in a timely fashion, but they do not impose serious timing constraints. Occasional delayed packets do not cause unacceptable operation. The intolerant applications, such as

interactive voice, are the most demanding and have tight timing constraints. If their traffic is not handled consistently and precisely at all times, they degrade unacceptably.

2. Voice Considerations

a. Digitization and Compression

The irresistible logic is that digitized voice and video is just another kind of packet data. Several coding schemes have been developed that compress speech by sending only simplified information about voice transmission; reducing the required bandwidth [Ref. 6]. Essentially, compression is a balancing act between voice quality, local computation power, delay, and network bandwidth required. The greater the bandwidth reduction, the higher the computational cost and delay for a given level of perceived clarity.

Coding techniques are standardized by the ITU-T in its G-series recommendations. The most popular coding standards for telephony and voice packet are:

- G.711, which describes the 64-kbps PCM voice coding technique used in PSTN
- G.723.1, which describes a compression technique that can be used for compressing speech or audio signal components at a very low bit rate. It has two bit rates associated with it: 6.3 and 5.3 kbps.
- G.726---Describes ADPCM coding at 32 kbps. ADPCM-encoded voice can be interchanged between packet voice, PSTN, and PBX networks if the PBX networks are configured to support ADPCM.
- G.728, which describes 16-kbps low-delay voice compression.
- G.729, which describes voice compression into 8-kbps streams.

b. Compression Delay

One of the most important design considerations in implementing interactive voice is minimizing one-way, end-to-end delay. For human requirements, an acceptable one-way delay for a voice conversation (mouth-to-ear delay) is approximately 150 milliseconds, while toll quality is achieved when the delay is less than 100 ms [Ref. 7] and [Ref. 8]. Apart from the delay imposed by the network, specifically for voice, compression delay must also be considered. It is the delay induced by the devices that handle voice information. Figure 7 shows the bit rate and the compression delay introduced by different coding schemes.

CODEC	Bit Rate (kbps)	Payload Size (bytes)	Compression Delay (ms)
G.711 PCM	64	160	0.75
G.726 ADPCM	32	80	1
G.728	16	40	3 to 5
G.729	8	20	10
G.723.1 MP-MLQ	6.3	30	30
G.723.1 ACELP	5.3	30	30

Figure 7. Bit Rates, Payload Sizes and Delays Induced By Coding Schemes.

c. Packet size

Packet size is also a key issue that must be taken into account in a network that will accommodate voice traffic. Data achieves maximum throughput when packet sizes are large, minimizing the overhead of the headers. However, voice cannot use such large packets because it will introduce extreme handling delays and additional echo problem. Furthermore, if a packet is lost in the network, the small packet is less likely to contain significant parts of a speech signal. The packet size was a big issue in standardization of ATM cell size that was eventually agreed in 53 bytes. The payload sizes of voice packets vary from 20-160 bytes based on the coding scheme that is used to generate the packet, as shown in previous figure 6.

d. *Real-Time Transport Protocol (RTP) / Real-Time Control Transport Protocol (RTCP)*

The transport of voice and video on IP-based networks requires the existence of a protocol that can provide a notion of time network-wide and allow control over jitter, desequencing and delay. Real-time Transport Protocol (RTP) was developed from IETF to fulfill this gap and is now being used as the core protocol for transport of inelastic applications. Its primary role is to act as a simple, functional and scaleable interface between inelastic applications and existing transport layer protocols. It is described in RFC 1889 specification as being a thin protocol providing support for applications with time sensitive properties, including timing reconstruction, loss detection, security and content identification. While its specification does not dictate which of the underlying transport and network layer to use, typically it is utilized on top of UDP/IP as shown in Figure 8.

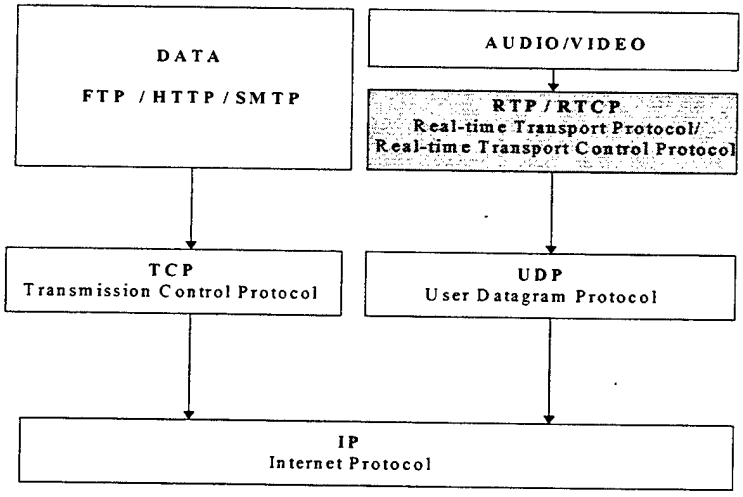

Figure 8. RTP / RTCP Protocols.

RTP is used to send data in one direction with no acknowledgement, but with an inherent notion of time. This allows receivers to compensate for the jitter and desequencing introduced by IP-based networks. RTP adds a new 12-byte header on top of UDP/IP headers to describe each datagram. This header contains:

- Time stamp, so the recipient can reconstruct the timing of the original data. This timestamp contains relative timing information that represents

timing relations between packets, not absolute points in time. Therefore, sender and receiver do not need to be synchronized.

- Sequence number, which lets the recipient reassemble the data and deal with missing, duplicate or out-of-order datagrams.
- Payload type, which describes the type of data, such as voice, audio or video and how it is encoded and compressed.
- Source ID, which helps a recipient distinguish multiple, simultaneous streams, using a unique sender-generated value.

The functionality of RTP is enhanced with another protocol called the Real Time Control Protocol (RTCP), which provides a mechanism for session control and monitoring of the RTP data. RTCP is based on the periodic transmission of control packets (limited to a small and known fraction of the session bandwidth, at most 5%) to all participants in the session, using the same distribution mechanism as the data packets. As a compliment to RTP, it performs four main functions and these are:

- Feedback Information. RTCP packets contain information such as the number of RTP packets sent, the number of packets lost, etc., which the receiving application or any other third party program can use to monitor network problems. The application might then change the transmission rate of the RTP packets to help reduce any problems.
- Participant identification, used to keep track of each of the participants in a session.
- Transmission Interval Control, which ensures that the control traffic will not overwhelm network resources.
- Minimal Session Control Information, an optional function which can be used to convey a minimal amount of information to all session participants, e.g. to exchange and display personal names of users joining or leaving an informal session.

It must be clarified that RTP and RTCP do not guarantee real-time delivery or prevent out-of-order delivery. They do not control quality of service in any way. This always requires the support of lower layers that actually have control over resources in internetworking devices. The network can drop, delay or desequence an RTP packet like any other IP packet. RTP and RTCP simply allow receivers to recover from network jitter by appropriate buffering and sequencing and to provide more information on the network so that appropriate corrective measures can be adopted.

3. Video Considerations

Packetized video traffic can be generated by interactive videoconference applications, which require low delay and jitter, and streaming video or video-on-demand that do not present tight timing constraints. A typical video application relies on TCP control channel and two UDP data channels, one for voice and one for video images.

The transmission of images, moving or still, is one of the largest consuming applications. Simple digitization of a video signal can yield from 10Mbps for traditional full-motion television video to up to 1 Gbps for High Definition Television (HDTV). The size of each image, which is called frame, depends on the resolution of the picture. For example, an image with resolution of 352 by 240 pixels, with each pixel represented by 24 bits of information, as would be the case for 24-bit color, results in a frame size of 247.5 Kbytes. To provide TV video quality, the images must be provided in a frame rate of 25-30 frames per second that results in a bandwidth requirement of 59.4 Mbps [Ref. 9].

Thankfully, differentially compression algorithms and supporting hardware reduce the bandwidth requirements by about 100-fold or better. The compression consists of key frames that describe the entire image and intermediate frames that describe changes from the original frame. The high degree of video compression is achieved by

losing data (lossy compression) and by requiring a big amount of time to do the compression. Better image quality requires more information.

H.261 is a video codec used in H.320 video conferencing to encode the image over several 64 Kbps ISDN connections. It is intended for compressed rates between 40 Kbit/s and 2 Mbps. H263, another ITU compression standard, was designed for low bit rate, as low as 20 Kbps. Streaming video applications use MPEG-1 and MPEG-2 compression standards that produce variable data rates in the range of 1.5 Mbps and 6Mbps respectively. The recent years, development of low-bit-rate streaming media proprietary codecs (Real Player, Windows Media Player and Apple's QuickTime) has led to data rates that vary from 28Kbps to 768 Kbps, depending on the desired video quality. It is expected that MPEG-4 will manage to standardize and integrate them in an interoperable format.

4. Synopsis Of Application Requirements

While it's widely understood that inelastic applications are more critical because of human perception and sensitivity to network delay, elastic applications also require a certain level of service from the network to operate effectively. With the introduction of very demanding voice and video applications, it is no surprise that data-oriented applications are also beginning to cry out for proper service. Timely delivery is also necessary for even the classical bulk-data applications. Delays may be fatal to critical file or image transfers. A user downloading a web page would likely find excessively slow loads unacceptable, significantly limiting the usability of applications.

Figure 9 provides a synopsis of the specific requirements of existing and emerging applications inside an organization.

APPLICATION	BANDWIDTH	LATENCY	JITTER	PACKET LOSS
File Transfer-Email-Web browsing	Low-Medium	> 1s	Allowed -	0%

Shared-Interactive Data Applications	Low-Medium	400 ms	Allowed -	0%
Interactive Voice Conversation	Low 5.3-64Kbps	150-200 ms	30 ms	<5%
Interactive Videoconferencing Streaming Media	Medium 28-768Kbps	400 ms	30 ms	<10%
Video on demand/TV	High 1.5-6Mbps	> 1s	Allowed	<10%
Imaging	High 8-100Mbps	> 1s	Allowed	0%
Virtual Reality- Tele-immersion	Very High > 100 Mbps	> 1s	Allowed	0%

Figure 9. Application Requirements In Terms of Performance Metrics.

III. QUALITY OF SERVICE PRINCIPLES AND FUNCTIONS

A. QUALITY OF SERVICE DEFINITION

Quality of Service (QoS) refers to the ability of a network to provide reliable and predictable service to selected network traffic. It refers to the ability of a network user or application to have some level of assurance that its traffic requirements can be satisfied.

While QoS can be quantified using several criteria that incorporate performance, availability, reliability, and security, in the context of this thesis, QoS is defined as a measure of the service provided by the network, which is a composite of the previously analyzed four measurable components:

- Data rate.
- Delay.
- Jitter.
- Packet loss.

As it was shown in the previous chapter, different applications have varying needs for QoS. Thus, to determine if a network offers proper QoS with respect of a specific connection, it is sufficient to determine if the performance traits are satisfied. Alternatively, we may also define intermediate levels of QoS that categorize it as excellent, good, fair, poor, or non-existent. For example, for a voice conversation, an end-to-end delay of less than 200 ms, with data rate of 64 Kbps, zero jitter, and packet loss of less than 5% could be considered excellent QoS, while a delay of 300 ms (others parameters the same) might be considered fair QoS. A delay of 500 ms or more would fall into the poor or nonexistent category. Similarly, for data applications, a file transfer with zero packet loss, 1 Mbps data rate and a reasonable timing delay would be considered excellent QoS.

QoS embraces a number of functions (prioritization, signaling, congestion control, policing) that intelligently match the needs of users and applications to the network resources available. These functions allow the treatment of certain packets in a preferential way, but also ensure that the less privileged flows do not starve and do get their fair service. Essentially, the goal of QoS is to enhance the current IP "best-effort" service, alleviate the inefficiencies and provide some level of performance guarantees.

B. QUALITY OF SERVICE PRINCIPLES OF IMPLEMENTATION

Applying a robust and scalable network-wide QoS solution in IP-based networks requires the cooperation of all network layers and a degree of coordination between end nodes and internetworking devices that does not exist today. That cooperation can be achieved by proper network engineering, based on the following principles:

- Minimal change to the IP simplistic nature. It is necessary to preserve the connectionless and stateless orientation of IP, avoiding major overhauls and fundamental changes.
- Weakest link principle. End-to-end QoS is determined by the weakest node encountered by a particular flow in the path between sender and receiver.
- Adding the necessary intelligence in each place and applying it as close as possible to the source of the problem. This usually implies the end systems, the applications and the border networking devices. For example, it makes sense to enable QoS in an application that will be afforded some kind of preferential treatment, but not in every application that is satisfied by the best-effort service model.
- QoS should be as non-disruptive and transparent as possible to existing network operations and integrated in a single device that could be remotely manageable.

The overall objective is to meet all of these principles in a practical, affordable and cost-effective fashion, minimizing complexity of the network and involvement of network operators and end users.

The successful implementation of QoS in the network will bring:

- Control and efficient use of resources. Without creating additional network resources, existing ones can be allocated and managed more effectively, especially when the network is under heavy load and there are not enough resources to meet every demand. For instance, the network administrator can limit bandwidth consumed by data transfers over a backbone link and give priority to an important voice conversation.
- Tailored services. QoS offers carefully tailored grades of service differentiation to the members of the organization.
- Harmonic coexistence of elastic and inelastic applications.

C. QUALITY OF SERVICE VS OVERPROVISIONING

Until now, the need for service quality has been addressed by simply augmenting the capacity of the network (overprovisioning). There is still a big argument in favor of overprovisioning. Odlyzko [Ref. 10] argues that overprovisioning data networks is a viable and sustainable response to the demands for service quality. It supports that increasing the available bandwidth is technically and economically superior to implementing complicated QoS techniques that add overhead to the network.

The objections against the implementation of a QoS scheme are based on the fact that eventually the entire system will become more complicated, increasing the computational burden on the networking devices and increasing the numbers and lengths of queues. On the contrary, as bandwidth becomes cheaper and more available with

technologies such as Gigabit Ethernet and high-speed optical WAN circuits, networks will operate so fast that will be able to satisfy the needs of all applications.

The capacity increase is a necessary step, but not sufficient to overcome all hindrances for the following reasons:

- It is a fundamental law of economics that demand will always expand beyond the supply of resources to cause congestion and queues. There will always be new bandwidth-hungry applications that will overrun the network capacity limits, especially during high-demand periods, no matter how much bandwidth the network can provide.
- Despite the fact that bandwidth becomes a commodity, it will take a long time and capital to provide sufficient bandwidth everywhere.
- Increasing bandwidth only is not enough to satisfy low delay and jitter. Voice traffic does not require much bandwidth (a conversation can be compressed to 8 kbps), but it has very strict requirements with respect to delay and jitter.
- Wireless networking, which is expanding rapidly, presents restricted bandwidths and will always be lagging in terms of bandwidth in comparison with terrestrial links and wired campus networks. Especially for a military organization, microwave RF and satellite systems, which mainly serve the communication channels, present important bandwidth mismatches, as depicted in Figure 10.

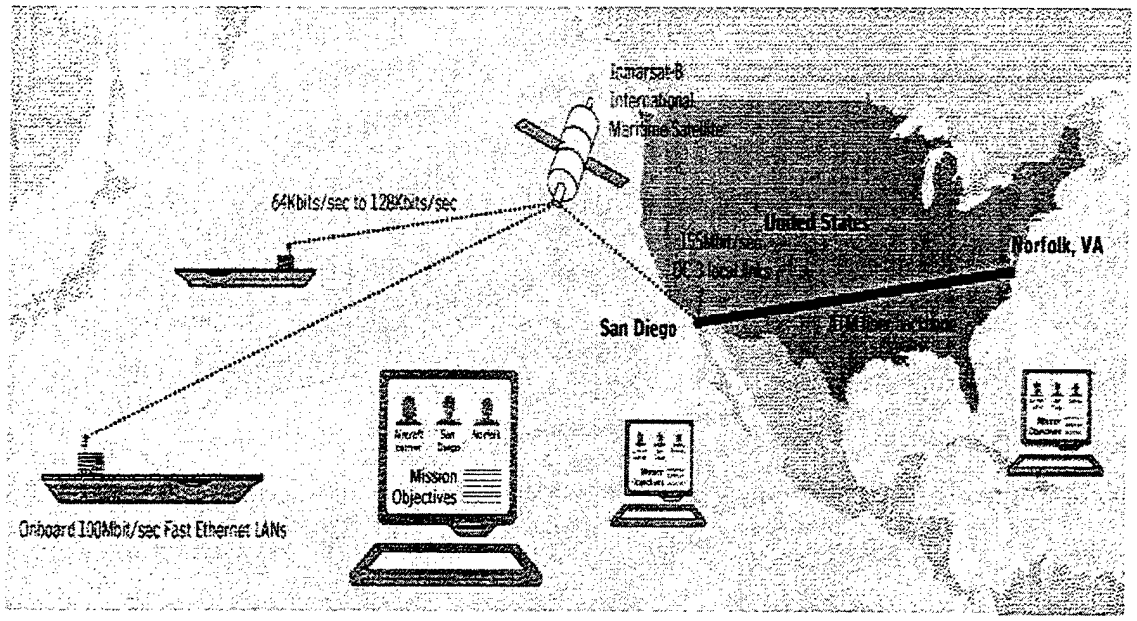

Figure 10. Bandwidth Mismatches In Military Network. "From Ref. 11"

Therefore, it is acceptable to throw bandwidth at the network, when it is cost effective to do so. However, without meaningful control mechanisms, even a network of enormous bandwidth cannot guarantee sufficient performance and is likely to saturate. QoS remains an essential technology, no matter how complicated, to deal effectively with an environment of finite resources, link capacity mismatches and diverse traffic patterns. The real challenge is to keep an optimum balance between the two approaches and have them complement each other, so that a network can give each application the resources it needs.

D. QUALITY OF SERVICE FUNCTIONS

QoS provides better and more predictable network service by accomplishing the following functions:

- Prioritizing traffic
- Signaling traffic requirements

- Congestion control
- Policing (admission control)

1. Prioritizing Traffic

One of the basic building functions of QoS is the ability to partition network traffic into different service classes or priority levels. To accomplish this, the packets must carry an explicit classification field and network devices must be able to identify it throughout the network and treat it accordingly. The classification process may be based on the following options [Ref. 12]:

- User desire: Giving packets preferential treatment whenever the user decides he would like preferred service. Obviously, there could be an abuse problem with this strategy, if it is not accompanied by complementary traffic shaping and admission control mechanisms.
- User privilege: Giving packets preferential treatment if they are associated with an entity that has been designated as eligible for priority treatment. The reason a particular user's end-system might be entitled to such privilege could be related to the user's rank, or affiliation with a special project, or by virtue of having subscribed to a preferred network service level.
- Application need: Giving packets preferential treatment, if they are sent by an application, whose need for certain network characteristics are network-wide known. It should be noted that application need is a relative and not an absolute concept. The actual need for any particular application may differ with circumstances. For example, high quality for a desktop videoconference may be more important in a commanders' brief than in the case of a sailor communicating with his family.

The network must be able to mark and prioritize certain fields of the packets entering the system and the networking devices will then use these fields in deciding how to treat these packets. The following two ways have been widely used so far:

a. Type of Service (TOS) Byte in IP Header

Originally, the IPv4 header included a byte, named Type of Service (TOS), which provided a way for differentiation and classification, but it had remained unimplemented and unused. The need for prioritizing resulted in TOS byte activation. Initially, the first three bits were determined to indicate relative priority of the packet, known as IP precedence, and the values these bits can take are described in RFC 791. As shown in figure 11, these three bits allow partitioning of traffic in up to 8 priority levels.

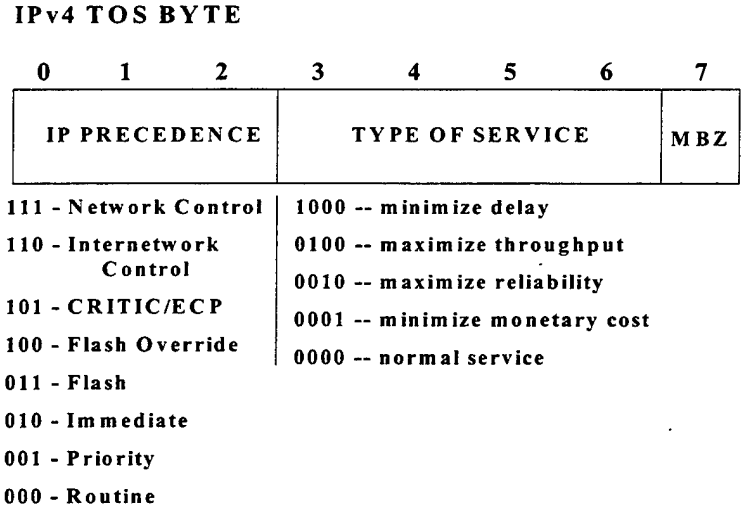

Figure 11. *IPv4* Type Of Service (TOS) Byte.

The next four bits of the TOS byte represent 4 metrics (delay, throughput, reliability and monetary-cost), as defined in RFC 1349, that denote how the network should make tradeoffs between throughput, delay, reliability, and cost. The final bit is unused and must be set to zero.

The Differentiated Services model (examined in the following chapter) renames the TOS byte to DiffServ field and defines new priorities based on the first six bits of the DiffServ field, while the remaining two are unused. The latest version of operating systems let applications and networking devices to manipulate this byte, allowing the flexibility to define up to 64 different priority levels.

b. 802.1p Bits in MAC Header

Most local area networks (LAN) are based on IEEE 802 technology. These include Ethernet, Token-ring, FDDI and other variations of shared media networks. 802.1p defines three bits in the MAC header of 802 packets that can carry one of eight priority values, corresponding to one of eight possible service levels in the LAN network. Typically, hosts or routers sending traffic into a LAN can mark the MAC header bits of each transmitted packet with the appropriate priority value. LAN devices, such as switches, bridges and hubs, are expected to treat the packets accordingly. It is important to mention that the value of this prioritization technique is limited because the scope of the 802.1p priority mark is limited to the LAN environment. Once packets are carried off the LAN, through a layer-3 device, the 802.1p priority is removed or has to be mapped in a higher layer priority scheme.

2. Signaling Traffic Requirements

Another way to request certain service from the network is to signal traffic requirements before data transmission. The Resource Reservation Protocol (RSVP), described in RFC 2205, was developed to be the standard mechanism to precisely signal QoS requirements to the network infrastructure. It should be noted that RSVP does not provide additional resources; it only reserves a portion of the existing ones. RSVP allows an application (or a network device on behalf of an application) to dynamically reserve network resources. The application signals its service requirements to all devices in the network that will handle the associated traffic, ensuring transfer of traffic in a deterministic and consistent manner.

RSVP is not a routing protocol, nor is it part of the routing architecture of the network devices involved in packet forwarding. It does not perform its own routing, but instead it uses underlying routing protocols to determine where it should carry reservation requests. RSVP messages are transmitted directly on top of the IP protocol as opposed to being transmitted over TCP or UDP.

RSVP requests the QoS on behalf of a particular flow. The protocol follows the receiver-based model. Figure 12 illustrates how RSVP works. Each RSVP sender sends a description of the characteristics of the traffic flow it intends to generate and reserve, usually in terms of bandwidth and latency. These resource requests are encoded in parameters within the RSVP "PATH" message and are transmitted from the sender to the receiver along the data path, provided by the routing protocol. At each node, RSVP attempts to make a resource reservation for the flow, determining whether the node has sufficient available resources to supply the requested QoS. In case of finding available local resources to support QoS requests, the "PATH" messages store path state in each node along the way. This path state includes at least the IP address of the previous hop node, which is used to route the "RESV" messages hop by hop in the reverse direction.

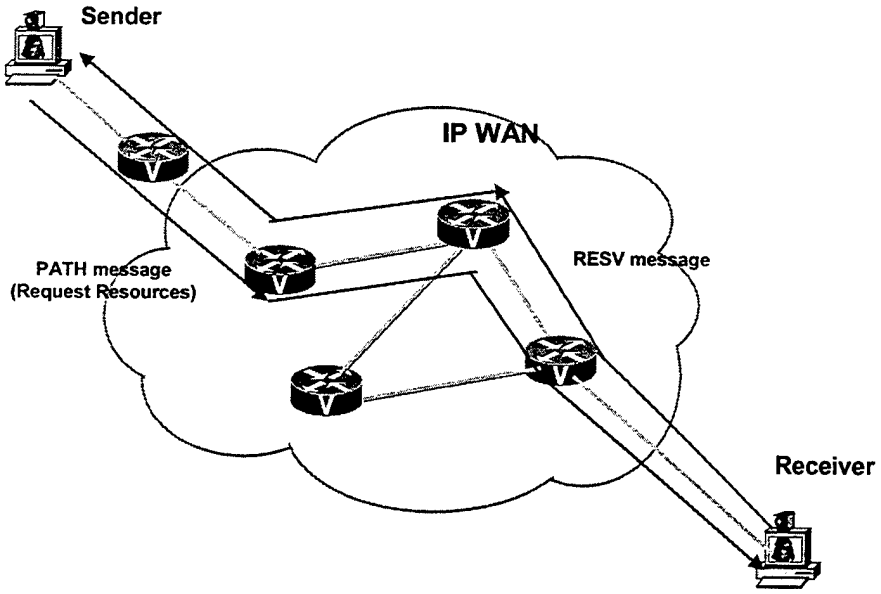

Figure 12. RSVP Signaling Process.

The receiver, as well as every RSVP-capable device along the path, is aware of the sender's traffic. The receiver of the data flow determines what QoS the flow will actually receive, sending the RSVP "RESV" message upstream along the same path toward the sender. This determination may be dependent on the receiver's capabilities, the application requirements or other administrative considerations. The "RESV" messages create and maintain reservation state in each node along the path. "RESV" messages are finally delivered to the sender, so that the host can set up appropriate traffic control parameters for the first hop.

Thus, for RSVP, characterization of the flow is the sender's responsibility, while the receiver specifies its particular service requirements. "PATH" messages are sent with the same source and destination addresses as the data, so that they will be routed correctly through non-RSVP clouds. On the other hand, "RESV" messages are sent hop by hop. Each RSVP-aware node forwards a "RESV" message to the address of a previous RSVP hop.

RSVP messages install state on devices along a data path. But to manage states across a network, the soft state model is used. RSVP sends periodic refresh messages to maintain the state along the reserved paths. In the absence of a timely refresh message, the state automatically times out and is deleted. As routing changes paths to adapt to topology changes, RSVP adapts its reservation to the new paths wherever reservations are in place. This technique is powerful in that it solves the problems associated with deallocation of reservations, lost packets and route changes (RFC 2210).

3. Congestion Control

a. Managing Congestion With Queuing Techniques

Congestion management is used to control congestion after it occurs. Network devices handle an overflow in arriving traffic by using smart queuing techniques that efficiently prioritize and handle traffic in a more effective manner than the traditional first-in-first-out (FIFO) technique [Ref. 13]. These techniques are:

- Priority queuing
- Custom queuing
- Weighted Fair Queuing

Priority queuing is a basic scheme that gives designated higher priority traffic absolute preferential treatment over low-priority traffic. It ensures that important traffic is queued ahead of other traffic and gets the fastest available handling. It provides no means of controlling the allocation of bandwidth, and often results in all but the highest priority applications being starved of bandwidth.

Custom queuing handles traffic by assigning different amounts of queue space to the various classes of traffic and then servicing the queues in a round-robin fashion. While a particular traffic can be assigned more queue space, it can never

monopolize all the bandwidth. All different streams of data are guaranteed a minimum quantity of bandwidth. This feature serves well traffic with specific minimum bandwidth or delay requirements, while still permitting other network applications to run effectively.

Weighted Fair Queuing (WFQ), the most sophisticated queuing technique, differentiates among bandwidth-hogging applications and those that need less bandwidth, and distributes the bandwidth to all applications in equal amounts. It is a flow-based queuing algorithm that does two things simultaneously. It schedules interactive traffic to the front of the queue to reduce response time, and it fairly shares the remaining bandwidth among high-bandwidth flows. WFQ ensures that queues do not starve for bandwidth, and that traffic gets predictable service. Low-volume traffic streams, which comprise the majority of traffic, receive preferential service, transmitting their entire offered loads in a timely fashion. High-volume traffic streams share the remaining capacity proportionally between them. WFQ is efficient in that it uses whatever bandwidth is available to forward traffic from lower-priority flows if no traffic from higher-priority flows is present.

b. Avoiding Congestion by Proactively Dropping Packets

There are certain congestion avoidance techniques that monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks through packet dropping. Random early detection (RED) has emerged as the standard congestion avoidance method. In basic form, RED randomly drops packets as queues fill up, causing end stations to decrease their transmission rates so queues will not overflow. If the forwarding device is not configured with RED, it uses the cruder default packet drop mechanism called "tail drop". Tail drop treats all traffic equally and does not differentiate between classes of service. Queues fill during periods of congestion. When the output queue is full and tail drop is in effect, packets are dropped until the congestion is eliminated and the queue is no longer full.

The RED mechanism was proposed by Sally Floyd and Van Jacobson in the early 1990s to address network congestion in a proactive rather than reactive manner. Underlying the RED mechanism is the premise that most traffic runs on top of TCP, which is sensitive to loss and will temporarily slow down when some of their traffic is dropped. The main goal of RED is to improve the efficiency of TCP congestion control. TCP, which responds appropriately to traffic drop by slowing down its traffic transmission, effectively allows RED traffic-drop behavior to work as a congestion-avoidance signaling mechanism. Given the ubiquitous presence of TCP, RED offers a widespread, effective congestion-avoidance mechanism.

Weighted Random Early Detection (WRED) combines the capabilities of the RED algorithm with IP Precedence to provide for preferential traffic handling of higher priority packets. WRED can selectively discard lower priority traffic when the interface begins to get congested and provide differentiated performance characteristics for different classes of service. WRED can also work with interfaces configured to use RSVP, where WRED chooses packets from other flows to drop rather than the RSVP flows.

c. Traffic Shaping / Rate Limiting

It is often the case that a particular link has sufficient capacity for the offered load if the load is more evenly spaced in time. In reasonably provisioned networks, congestion usually results from the peaks. When demand peaks at certain times, the link is not capable of handling the instantaneous demand. Traffic shaping means modifying the timing of a sequence of packets so as to reduce burstiness. It does not reduce total network demand, but it smoothes out the peak demands, shifting demand from peak times to off-peak times.

The primary reasons traffic shaping is used are to control access to available bandwidth, to ensure that traffic conforms to the policies established for it, and

to regulate the flow of traffic in order to avoid congestion that can occur when the transmitted traffic exceeds the access speed of its remote interface. Hence, traffic shaping can be a very important part for congestion avoidance, eliminating bottlenecks especially in topologies with data-rate mismatches. For example, if one end of the link in a Frame Relay network runs at 256 kbps and the other end of the link runs at 128 kbps, sending packets at 256 kbps could cause link congestion.

Rate limiting provides the means to allocate bandwidth commitments and limitations to traffic sources and destinations, while specifying the actions for handling traffic that exceeds these allocations. Essentially, it dictates a maximum amount of bandwidth that a particular application can consume. For example, the first 100 Kbps of video traffic generated by a video streaming application can be allowed to go through in a preferential manner, but traffic above the first 100 Kbps by the same application can drop to lower priority or be discarded. Similarly, file transfer traffic can be limited to 20% of all available bandwidth so that it does not starve out other applications.

d. Packet Size Optimization and Fragmentation

Time-sensitive traffic is susceptible to increased delay and jitter when the network processes large packets, especially when the packets are queued on slower links. In a slow link, a large packet can make the connection unavailable to other packets for a considerable amount of time. As shown in figure 13, a large frame of 1500 bytes (typical Ethernet packet) takes 215 ms to traverse a 56-kbps line, which exceeds the overall delay requirement for time-sensitive traffic (150-200 ms). Therefore, to limit the delay of time-sensitive packets on relatively slow bandwidth links, a method for fragmenting larger packets and queuing smaller packets between fragments of the large packet is needed.

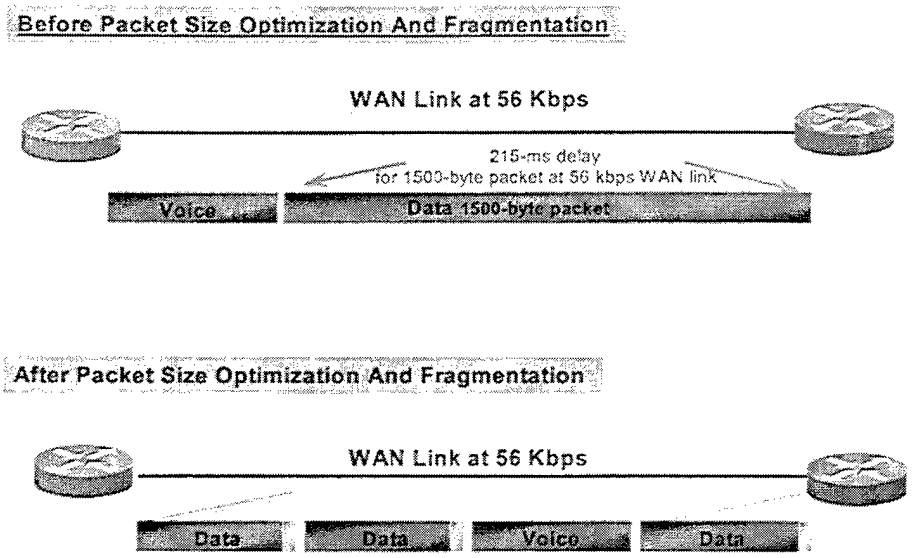

Figure 13. Packet Size Optimization And Fragmentation.

Packet size optimization and fragmentation provides a method of splitting, recombining, and sequencing packets that reduces transmission delay across slow bandwidth WAN links. Large datagrams are multilink encapsulated and fragmented to packets of a size small enough to satisfy the delay requirements of the delay-sensitive traffic, while small delay-sensitive packets are interleaved with the smaller packets resulting from the fragmented datagram.

e. Protocol Header Compression

The header portion of IP/UDP/RTP is considerably large. As shown in figure 15, the minimal 12 bytes of the RTP header, combined with 20 bytes of IP header and 8 bytes of UDP header, create a 40-byte IP/UDP/RTP header. For audio applications, the packet payload typically varies from 20 bytes to 160 bytes, depending on the CODEC used. Given the size of the IP/UDP/RTP header combination, it is inefficient to transmit the IP/UDP/RTP header without compressing it. To avoid the unnecessary consumption of available bandwidth, the protocol header compression feature is used on a link-by-link

basis. It compresses the IP/UDP/RTP packet header from 40 bytes to approximately 2 to 5 bytes, as shown in figure 14.

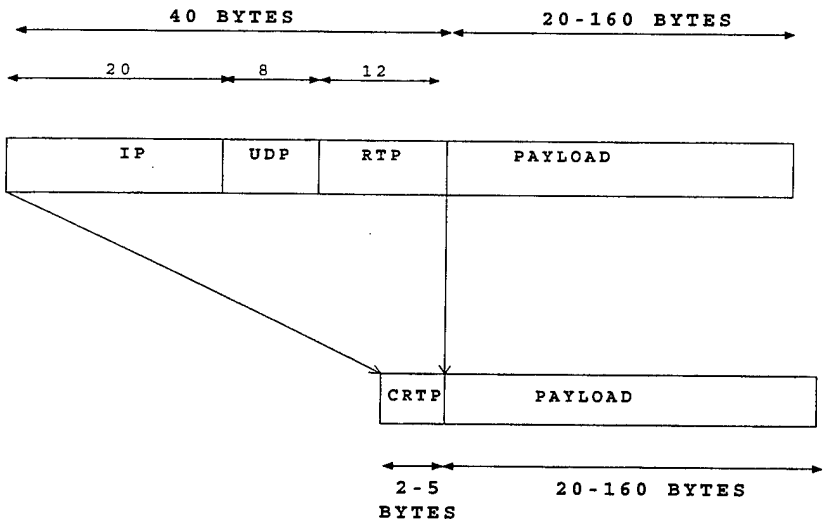

Figure 14. Protocol Header Compression.

This feature accrues major gain in terms of packet compression, because although several fields in the header change in every packet, the difference from packet to packet is often constant, and therefore the second-order difference is zero. The decompressor can reconstruct the original header without any loss of information.

The overhead reduction for multimedia RTP traffic results in a corresponding reduction in delay. The header compression is especially beneficial when the payload size is small, for example, for compressed audio payloads of 20 to 50 bytes. It is recommended on any WAN interface, where bandwidth is a concern and there is a high portion of RTP traffic. It should not be used on any high-speed interfaces, anything over T1 speed, because the trade-offs are not desirable.

4. Policing

Since enabling QoS on an IP network effectively means that some users will get better network service than others, it creates some incentive to steal or abuse the

resources. Therefore, there is a need for a policing and admission control scheme to authenticate those that request the better service levels and to verify the identity of traffic "owners" on a per-packet basis. In the absence of a policing scheme, traffic owners would mark their packets as desirous of preferential treatment, all packets will in fact would become high-priority packets, and the outcome would be another "best-effort" network.

Since there are varying circumstances in which traffic owners (end-users, applications, host machines) are entitled to the services they request, there is a need for a set of rules, a need to decide when these rules apply and a need to enforce them. The rules, the judging and enforcing devices, all comprise a policy system that is an essential component of a QoS-enabled network [Ref. 14].

Policy is one or more rules that describe the actions to occur when specific conditions exist. Policies determine which applications and users are entitled to varying amounts of resources in different parts of the network. Policy rules, conditions and actions, must be unambiguous and verifiable. There should be only one correct rule appropriate for any specific set of conditions so that network personnel can configure QoS mechanisms subject to these rules. Several parameters can be selected to reflect a defined policy, such as IP or MAC address, application port, user, time of day, or location within the network.

Even network management teams, who have well-planned QoS policies, find that implementing and enforcing them throughout a large network is a complicated and sometimes overwhelming task. Even with advanced, intelligent network devices and the latest management tools, the fact remains that in today's networks, the configuration is often very complex. It is often done by hand, through a command-line interface (CLI), one device at a time. Few network administrators have the time or experience to correctly

implement end-to-end QoS policy in every device. This shortfall often results in inconsistent policy implementation and a lack of dynamic application control.

Clearly, there is a need for a simpler, higher-level way to implement policy without requiring a detailed understanding of the mechanisms and the lower-level configuration and maintenance management. The Resource Allocation Protocol (RAP) Working Group in the IETF was originally assigned to "establish a scalable policy control model". The Policy Framework they designed [Ref. 15] was quickly accepted by the networking industry and recognized as being applicable to the whole network. Indeed, it has since been recognized as a generally useful model for other technologies that need policy support, such as network security (for firewalls, IP Security, Virtual Private Networks).

The framework is relatively simple, as shown in figure 15, and comprised by the following policy architectural components:

- Policy decision point (PDP), which translates network-wide higher layer policies into specific configuration information for individual network devices. PDP also inspects resource requests carried in RSVP messages and accepts or rejects them based on a comparison against policy data.
- Policy enforcement point (PEP), which acts on the decisions made by PDP. This is typically a network device that either does or does not grant resources to arriving traffic.
- A policy data repository, typically stored in a directory, which contains the policy data, represented as data structures so they can be stored and retrieved, such as user names, applications, and the network resources to which these are entitled.
- Protocols that enable communication between the data repository, PDP and PEP.

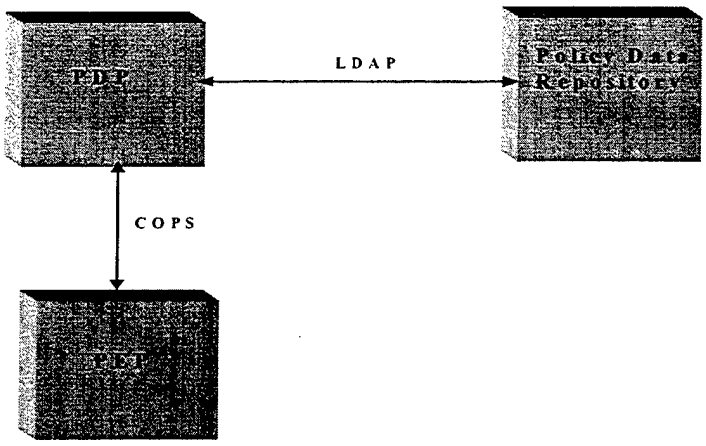

Figure 15. Policy Framework.

In a sense, PDP is the judge that makes decisions based on the policies it retrieves from the policy data repository and PEP is the enforcing device that applies the decisions of PDP.

To make its decisions, the PDP consults the data repository for the rules established by the network manager and decides, based on current network conditions, how traffic and access rights should be enforced. A protocol is required for the communication between the PDP and the policy data repository. Since the data repository tends to take the form of a distributed directory, Lightweight Directory Access Protocol (LDAP) is commonly used for this purpose, so multiple network applications can share and make decisions based on the information.

Once the PDP has made a decision on how to treat network traffic, it communicates the instructions to the PEPs (routers, switches, and gateways) via a protocol named Common Open Policy Service (COPS). It is a simple query/response protocol that has been developed in the context of QoS. It was initially targeted as an RSVP-related policy protocol but has recently been pressed into service as a general configuration protocol. COPS is preferred because it is connection-oriented and reliable,

and includes locking mechanisms to prevent multiple PDPs from simultaneously attempting to update the same PEP.

The separation of PEP and PDP is a logical one, based on functionality and not necessarily a physical separation. In certain cases, the PEP and the PDP can be co-located in the same networking device. In other cases, the PDP may be separated from the PEP in the form of a policy server. A single policy server may reside between the directory and multiple PEPs, as shown in figure 16. Although many policy decisions can be made trivially by co-locating the PDP and the PEP, scalability issues can be better addressed by the use of a separate policy server. Recently, work is proceeding on the concept of a bandwidth broker [Ref. 16]. The bandwidth broker concept is similar to a PDP in the sense that it makes decisions regarding bandwidth provisioning. However, bandwidth brokers tend to operate at a higher level than PDPs, operating at the edges between domains and being less aware of the topologies within domains.

Gradually, DNS, DHCP and authentication servers can be integrated that will dynamically update the directory with IP address-to-device and address-to-user association information. Ultimately, the objective is the network to dynamically learn of changes from the directory and to reconfigure itself to ensure that QoS policies are appropriately applied.

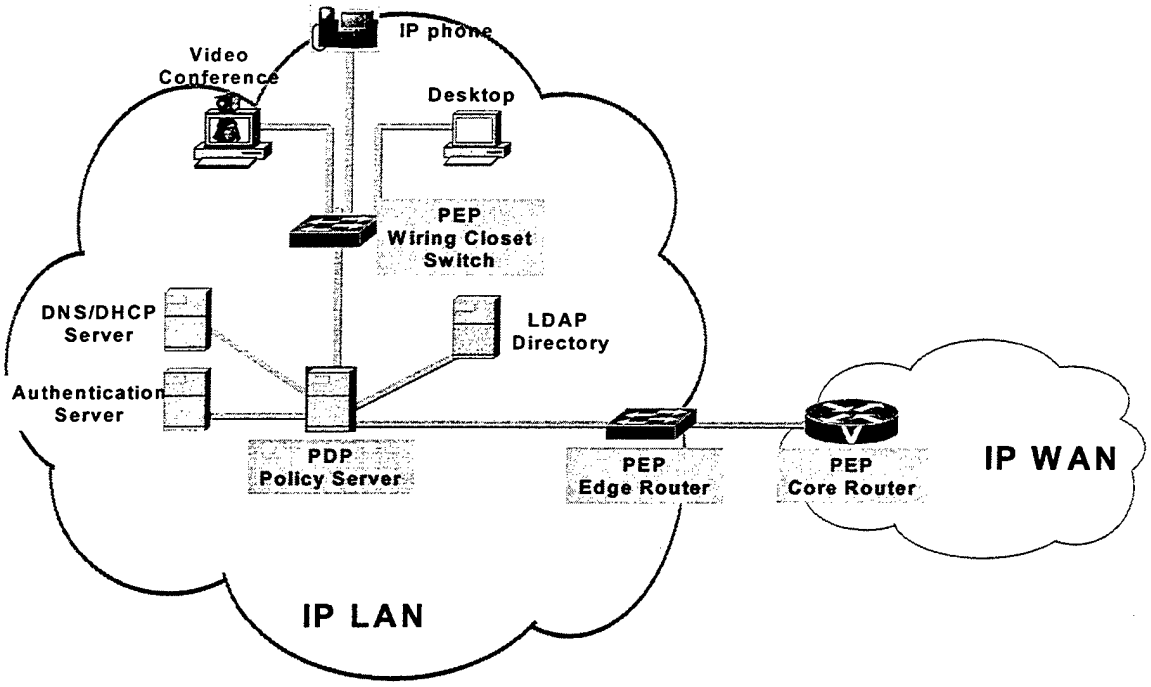

Figure 16. Policy-based Network Environment.

As issues of scalability, interoperability and ease of use are resolved, this framework can enable automatic implementation of defined policies in the network. This will substantially increase network integrity while reducing dramatically the required human resources to implement and maintain QoS. While it will require the integration of multiple servers, directories, protocols and network devices, the above policy framework is expected to be the only reasonable way to police a converged voice/data/video network in a dynamic and efficient way.

THIS PAGE IS INTENTIONALLY LEFT BLANK

IV. QUALITY OF SERVICE MODELS

In the previous chapter, several mechanisms and functions were examined that can contribute to the QoS implementation. To deliver end-to-end service guarantees and create reliable performance outcomes, these mechanisms must be applied in concert. Essentially, there are two basic approaches that provide QoS capabilities network-wide and IETF has developed two standards that reflect these approaches. Integrated Services (IntServ) provide QoS to individual applications or flows and Differentiated Services (DiffServ) provide QoS to aggregated traffic. The purpose of this chapter is to examine the principles of each QoS model, gain an understanding of their respective strengths and weaknesses and examine the recent development of a hybrid model that combines the strengths of both models.

A. INTEGRATED SERVICES MODEL

The Integrated Services (IntServ) model was the initial approach of IETF. This model relies on the creation of a reservation state within the network that corresponds to a service request and maintains this state for the duration of the associated flow.

In this model, it is the application that requests through explicit signaling a specific QoS from the network before sending data. The application informs the network of its traffic profile and requests a particular kind of service that can encompass its bandwidth and delay requirements. The application is expected to send data that lies within its described traffic profile only after it gets a confirmation from the network.

The network performs admission control, based on information from the application and available network resources. It also commits to meeting the QoS requirements of the application as long as the traffic remains within the profile specifications. The network fulfills its commitment by maintaining per-flow state and

then performing packet classification, policing, and intelligent queuing based on that state.

RSVP (examined in chapter III) fulfills the role of explicit signaling. This end-to-end signaling protocol provides a way to communicate the application's requirements to network elements along the path between sender and receiver, and to convey QoS management information network-wide. This imposes flow-specific state in the network elements, which represents an important and fundamental change to the simple IP model.

IntServ defines three levels of service:

- Guaranteed service (RFC 2212), with bandwidth, maximum bounded delay, and no-loss guarantees. Service is guaranteed to be within these limits and that allows applications to meet their requirements. For example, a Voice over IP (VoIP) application can reserve 32 Kbps data rate end-to-end along the path.
- Controlled load service (RFC 2211), which approximates best-effort service in a lightly loaded network and allows applications to have low delay and high throughput even during times of congestion. For example, streaming media applications can use this kind of service.
- Best-effort service, similar to what the IP network currently provides under a variety of load conditions.

1. IntServ Strengths

An important strength of the IntServ model is guaranteed delay bounds for an individual flow. If all nodes are RSVP-aware, then there will be an absolute upper bound on the network delay of the traffic. This can apply not just to one hop, but network-wide.

Furthermore, the IntServ model avoids service degradation. RSVP provides the ability to reject connections that if admitted would receive unacceptable QoS and also degrade the QoS of other reservations in progress. This capability derives from the messages RSVP sends along the path that packets will travel if the connection is established.

RSVP is designed to operate with current and future unicast and multicast routing protocols. Since the membership of a large multicast group and the resulting multicast tree topology are likely to change with time, RSVP sends periodic refresh messages to maintain the state along the reserved path. In the absence of refresh messages, the state automatically times out and is deleted.

Controlled link sharing is another benefit of the IntServ model. It is feasible not to put bounds on delay, but to limit overload shares (oversubscription) on a link, while allowing any mix of traffic to proceed if there is spare capacity.

2. IntServ Weaknesses

The most problematic issue is that IntServ maintains individual flow states to the network links for all accepted reservations that have been made. This represents a fundamental change to the stateless and connectionless IP architecture that was founded on the concept that all flow-related state should remain in the end systems.

Lack of scalability is another IntServ weakness, especially in high-speed backbone networks. Indeed, the amount of resources that a router needs for RSVP processing and storage increases proportionally with the number of IntServ flows. Traffic measurements show [Ref. 17] that most end-to-end IP connections are very short-lived, and that there are several thousand active connections at any time in a backbone router. Consequently, numerous IntServ flows on a high-bandwidth link place an excessive

burden on routers. Furthermore, if a topology change occurs, the reservations would need to be renegotiated simultaneously.

B. DIFFERENTIATED SERVICES MODEL

Differentiated services (DiffServ), described in RFC 2474, is a service model that was designed to satisfy differing QoS requirements, operating on layer 3 (IP) of the TCP/IP protocol stack, preserving the stateless and connectionless nature of the IP network. To overcome the limitations of the IntServ model, DiffServ merges individual flows into fewer and finite aggregates. In contrast to IntServ orientation, there is no need for per-flow state and explicit signaling at every hop inside a DiffServ network.

For differentiated service, the network tries to deliver a particular kind of service based on classification of each packet. This classification takes place in the DiffServ field, which supersedes the existing definitions of the IPv4 Type of Service (ToS) byte (RFC 1349) and the IPv6 Traffic Class byte (RFC 2460). As shown in Figure 17, six bits of the DiffServ field (formerly TOS byte) are used as the DiffServ Code Point (DSCP), while the rest two remain unused. Each DSCP is a six-bit value that identifies a particular Per-Hop Behavior (PHB) that a network element applies to each packet. DiffServ has subsumed IP precedence, but maintains backward compatibility.

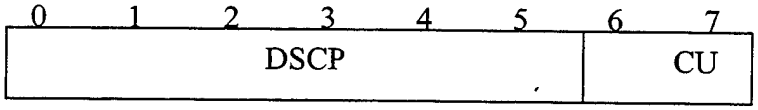

Figure 17. DiffServ Field.

PHBs are at the heart of the DiffServ architecture. A PHB is selected at each node by mapping the DSCP in each received packet. The PHB is the means by which a node allocates resources to behavior aggregates. PHBs are implemented in nodes by means of some buffer management and packet scheduling mechanisms. The IETF has recently specified two PHBs for standardization:

- Expedited Forwarding (EF), specified in RFC 2598. Packets marked with EF (DSCP 101110 is recommended) are forwarded with minimal delay, low jitter, low loss and assured bandwidth at each hop through DiffServ domains.
- Assured Forwarding (AF), specified in RFC 2597. The AF PHB group provides IP packet delivery in four independently forwarded AF classes. Within each AF class, an IP packet can be assigned a certain amount of forwarding resources (buffer space and bandwidth) and one of three levels of drop precedence. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class.

The DiffServ model achieves scalability by aggregating the traffic classification state, which is conveyed by the marking of the DiffServ field of each IP packet. Packets are classified and marked to receive a particular PHB on nodes along their paths. Sophisticated classification, marking, policing, and shaping operations need only be implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service-provisioning policies that govern how traffic is marked and conditioned upon entry to a DiffServ-capable network, as well as how that traffic is forwarded within that network. A wide variety of link characteristics-bandwidth, delay, jitter, and/or loss-can be controlled and adjusted accordingly.

Packet markers set the DiffServ field of a packet to a particular DSCP, adding the marked packet to a particular DiffServ behavior aggregate. Shapers delay some or all of the packets in a traffic stream to bring the stream into compliance with a traffic profile. A shaper usually has a finite-size buffer, and packets may be discarded if there is insufficient buffer space to hold the delayed packets. Droppers discard some or all of the packets in a traffic stream to bring the stream into compliance with a traffic profile.

1. DiffServ Strengths

Unlike RSVP, no QoS requirements are exchanged between the source and the destination, eliminating the inherent setup costs associated with RSVP. Short-lived flows benefit from DiffServ because the absence of QoS setup costs improves responsiveness and reduces the overhead required for a quick discussion with another host.

In case of congestion, flows will adapt their traffic to the available resources and continue operating, albeit at lower levels of service. The benefit is higher overall efficiency-more flows get through with greater simplicity, minimal signaling support, and simple data-path mechanisms

2. DiffServ Weaknesses

DiffServ only maps services with different levels of "sensitivities" to delay and loss, without being associated with explicit values or guarantees. It does not attempt to guarantee a level of service. Instead, it strives for a relative ordering of aggregations, such that one traffic aggregation will receive better or worse treatment relative to other aggregations, based on the behavioral rules of each aggregation.

C. AN INTEGRATED/DIFFERENTIATED HYBRID MODEL

Together, IntServ and DiffServ can facilitate QoS deployment of applications. A hybrid framework, as shown in figure 18 has been proposed [Ref. 18] and seems to have strong support and momentum. It assumes a model in which peripheral networks are IntServ-aware and interconnected by DiffServ networks.

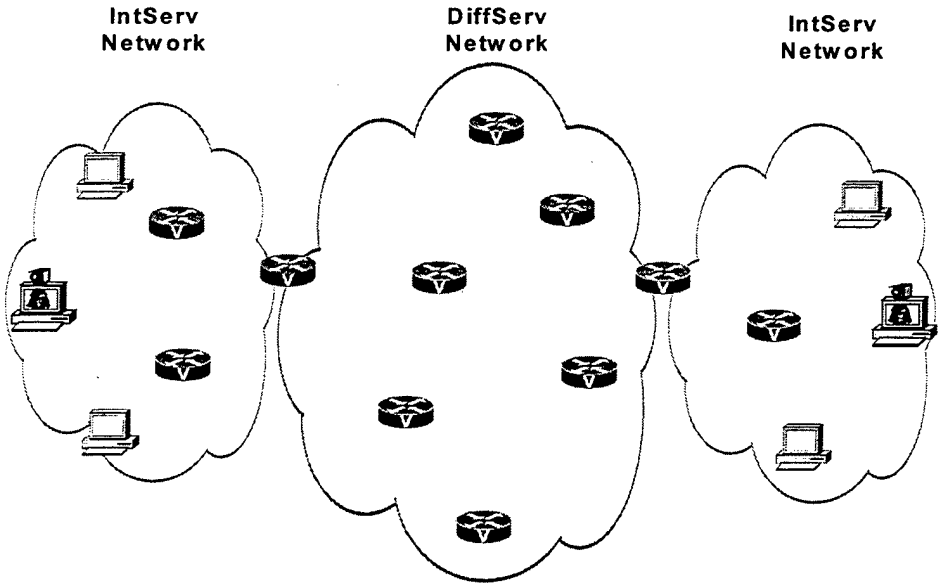

Figure 18. An Integrated/Differentiated Hybrid Model.

In this model, the scalability of DiffServ networks extends the reach of IntServ/RSVP networks. Intervening DiffServ networks appear as a single RSVP hop to the IntServ/RSVP networks. Hosts attached to the peripheral IntServ/RSVP networks signal to each other for per-flow resource requests across the DiffServ networks. Standard IntServ/RSVP processing is applied within the IntServ/RSVP peripheral networks. RSVP signaling messages are carried transparently through the DiffServ networks. Devices at the boundaries between the IntServ/RSVP networks and the DiffServ networks process the RSVP messages and provide admission control based on the resource availability within the DiffServ network.

Adequate IntServ and DiffServ mapping at the boundaries and suitable resource provisioning in the core are essential to ensure that the performance across the transit network does not defeat the end-to-end QoS. IntServ is implemented at the edge of organization LAN networks, where user flows can be managed at the desktop level. DiffServ plays a key role in the core WAN network where it eliminates the scalability concerns of IntServ/RSVP networks. With this model, it is possible maintain the

fundamental principle of IP network that leaves complexity at the 'edges' and keeps the network 'core' simple.

An important driver for IntServ in the vicinity of the end hosts is the implementation of RSVP and QoS capabilities in modern desktop operating systems. The use of RSVP signaling provides admission control to the DiffServ network, based on resource availability and policy decisions. It also greatly simplifies the configuration of DiffServ classifiers, policies, and other traffic conditioning components.

V. FUTURE NETWORK INFRASTRUCTURE PROPOSITION

The network infrastructure must be able to support today's needs while preparing for tomorrow's technological changes. This chapter discusses the alternatives and then points to the preferred direction to move forward. Essentially, there are three options:

1. Maintain voice and data networks separately. The voice network provides the quality of service that is desired only for voice. Data networks can sufficiently serve bandwidth intensive data applications.

2. Move towards an IP-based integrated voice, video and data network by adding bandwidth to the network. Up to now, this was mainly the approach that had been followed by the industry to overcome the network bottlenecks.

3. Move towards an IP-based integrated voice, video and data network with the deployment of extensive QoS and traffic management mechanisms.

The connection-oriented nature of the voice network and the support for relatively low bandwidth does not make it suitable for a network that can accommodate the increased need for high-speed digital access and the expansion of intensive data and video applications. As analyzed in Chapter II, the benefits of a converged network are enormous and therefore it is inevitable to integrate in a robust and dynamic IP-based network, because it has the potential to serve as the basis of a multiservice network. Therefore, following the first option of keeping the two networks separately does not present a scalable, economic and robust solution. Nevertheless, the organizations have invested heavily in enterprise voice equipment that will take several years to depreciate. The result is that organizations view IP migration as a multiple year project requiring incremental steps and major outlays of resources. During this transition, voice networks will be around serving voice communications but eventually will fade away, replaced by voice over IP.

The second option, throwing capacity only in the IP infrastructure, is a necessary step, but not sufficient to overcome all network performance deficiencies of a multiservice network, as analyzed in Chapter III. It is desirable to increase the bandwidth of the network, but it is not wise to consider it as panacea, because it is not always economically and technically feasible.

The third approach, implementing QoS only, compounds management complexity and creates additional overhead and latency in the network. Consequently, the entire network may become more complicated, increasing the computational burden on the networking devices. QoS is not enough if it is not supported by sufficient capacity.

Therefore, the proposition of this thesis is that a combination of the second and third approach is the proper direction to be followed. The proposed approach is to keep an optimum balance between bandwidth increase and QoS deployment and have them complement each other, so that the IP-based network can deal effectively with diverse traffic patterns, user differentiation and link capacity mismatches.

This approach can overcome all the major deficiencies of the IP-based network and raise it in a higher degree of maturity and reliability. Deploying QoS on the IP-based network should be approached systematically, with scalability and robustness in mind and in tandem, proper redesign and upgrade of the network will provide the necessary capacity. The migration and road to convergence will be a bumpy ride and will take time and careful planning. New techniques and applications must be deployed and thoroughly tested on a large scale to mitigate the risks associated with availability, fault tolerance, and redundancy The remainder of this chapter explains how to achieve the proposed approach.

A. UPGRADE OF NETWORK INFRASTRUCTURE THROUGH FIBER OPTICS DEPLOYMENT

As it was mentioned, bandwidth increase was the first reaction of the industry to the IP network performance problems. The increased demand for capacity, quickly consumed all available bandwidth afforded by wire-based media such as twisted pair and coaxial cable. To alleviate this problem, the networking industry has turned to a new medium, fiber optic.

Fiber has also undergone tremendous growth in the past decade. Communication companies have realized the advantages of fiber optics and have invested heavily into replacing wire-based circuits with this medium. Fiber has been deployed extensively in network backbones, especially at WANs, campus backbones and between floors and buildings, but fiber to the desktop or to the home still remains an expensive proposition. Copper is still much more prominent for the end hosts than fiber.

Fiber presents several advantages over the other physical media. Certainly the biggest gain is its enormous bandwidth capability enabling it to handle tremendous digital data rates with near error-free transmission. Other distinct benefits include:

- Reduced size and weight, 20 times lighter and 5 times smaller than equivalent copper cable.
- Transmission over longer distances.
- Permanence of the cable plant (network upgrades are limited to network electronics and software).
- Added security due to its resistance to unauthorized tapping.
- Superior quality attributes (fiber is free from signal cross-talk and electromagnetic interference).

Until recently, the cost of fiber technology made it unsuitable for all but the largest communications systems and networks. Advances in technology however, have changed this. Improvements in fiber optic quality, light sources, light detectors, and the procedures used to splice circuits have reduced the cost of implementing this medium. Fiber optics is now a feasible solution to the bandwidth limitation problems in residential areas and networks that support hundreds of nodes.

Despite being up to 50% more expensive than copper category 5, the cost of installing fiber-to-the desktop is dropping rapidly and that will drive fiber further into the organization. Ironically, the physical fiber cable is the least expensive element. Electronics and density connectors constitute the most expensive elements. These are also responsible for the creation of the capacity limitations. In theory, fiber can be used to transmit between 50 and 75 terabits per second [Ref. 19]. Currently deployed transmission technology supports 10 Gbps fiber channels. However, the deployment of Wave Division Multiplexing (WDM) is expected to lift the transmission rates to between 100 Gbps and 600 Gbps per fiber cable [Ref. 20].

Fiber's declining costs and distinct performance advantages, relative to the proposed available copper upgrade options (Category 6 and 7), make it appealing to large networks and the wiser choice when the time comes for an organization to install or upgrade the cabling infrastructure. Already, companies, such as Western Integrated Networks and Optical Solutions, install fiber to the home and build broadband networks to the residential premises.

B. QUALITY OF SERVICE IMPLEMENTATION

As analyzed earlier, QoS mechanisms must be introduced to handle differentiated traffic and control resources of the network environment. Therefore, a robust QoS

scheme must be implemented to ensure proper service. Before applying an extensive set of QoS policies and QoS-aware devices, it is necessary to identify the particular characteristics of the network (total required capacity, existing bottlenecks, traffic patterns) and the applications running on it. Furthermore, it is also important to determine what applications or users will be given priority, how the classification will be implemented and what mechanisms will be used to satisfy the requirements.

The deployment of QoS mechanisms can be more effective with the implementation of the following:

- Full adaptation of switched networks instead of shared and use of Virtual LAN (VLAN).
- Multicasting, which intersects with QoS, allows the distribution of the same datagram to many destinations without replicating the stream multiple times. The implementation of multicasting results in increased network efficiency. To deliver IP multicast traffic to the desktop efficiently, a solid-switched infrastructure is a necessity.
- Caching and local mirroring, through the deployment of proper servers, constitutes another way of implementing effective traffic engineering and highly complements multicasting. Local access does not require the generation of additional traffic and the network is more responsive and interactive to the end user.
- A non-hardware approach should involve the end-users. User training and discipline is imperative for the efficient use of the network. It is necessary to establish essential operational procedures, minimize wasteful ones and make users understand the existing bandwidth constraints. Bandwidth hungry applications, such as an online PowerPoint presentation, can be easily avoided with optional techniques that perform the same operation in a more efficient way (for example doing the presentation in HTML format).

QoS is too costly and complex to implement everywhere, and even the most robust QoS capabilities cannot overcome poor network design. Furthermore, it should be taken into consideration that many nodes on the network are fully satisfied with the current service they receive and do not require an additional QoS overhead to perform their operations.

It becomes increasingly apparent that the wiring closets in campus networks and the backbone LAN-to-WAN aggregation points constitute the network performance bottlenecks. A typical campus network runs on switched Ethernet at 10/100 Mbps to the end hosts and close to 1 Gbps from the wiring closet to the backbone, either Gigabit Ethernet or ATM. Also, there are big bandwidth mismatches between LAN and WAN networks and it is usually the case that the aggregate LAN traffic destined to the WAN is greater than the WAN bandwidth capacity. A good example is the routers at the border of a military radio WAN that definitely requires additional intelligence to handle efficiently the limited resources.

The model of combined DiffServ/IntServ or pure DiffServ is more appealing than IntServ because it preserves the connectionless nature of the network and keeps per-flow state out of the core network. As it was explained in the previous chapter, the DiffServ model improves the scalability of QoS provisioning by pushing state and complexity to the edges of the network and keeping classification and packet handling functions in the core network as simple as possible. The core devices only perform traffic handling and congestion control. With this approach, intelligence is applied primarily at the edge devices, where the network experiences the most severe bottlenecks, while the core devices remain simple and fast. Therefore, it makes the most sense to implement QoS as follows:

- Classification and prioritization by the end systems (applications/users).

- Mapping, traffic shaping, policing and admission control by the edge network devices.
- Traffic handling and congestion control by the core network devices.
- Supervision and administration control by smart network devices (bandwidth brokers) whereby each of them overviews part of the network.

1. Classification and Prioritization by End Systems

Within a network, classification and prioritization is made by QoS-enabled applications that can dynamically signal their QoS requirements with several approaches. This can be achieved at layer 3 (IP), by setting properly the DSCP of the DiffServ field or sending RSVP messages, and at layer 2, by setting the 802.1p bits of the MAC header.

Despite the fact that prioritization schemes allow the usage of 8 (IP precedence, 802.1p) or 64 different service classes (DiffServ DSCP), the proposed QoS deployment specifies three service levels that categorize applications types, as shown in Figure 19. Improving user and application visibility by the network with the appropriate integration of policing, authentication and network management tools, will further allow to incorporate distinct user needs and privileges at the defined service levels.

LEVELS	TYPES OF APPLICATIONS	EXAMPLES
Level A	Inelastic Intolerant Applications With low delay and jitter	Interactive Voice Over IP Interactive Videoconference Critical Network Services
Level B	- Inelastic Tolerant Applications - Data Intolerant Applications	Streaming Media Video on demand/TV Shared Interactive Data

Level C	Elastic Tolerant Data Applications	File Transfer-Email-Web browsing- Imaging

Figure 19. Proposed Service Levels.

This decision is made for the sake of

- Flexible manageability.
- Network device simplicity (need for configuration of three queues only).
- Achieving end-to-end packet prioritization under heterogeneous internetworking, by smooth mapping among the existing QoS approaches of IP networks in LAN, campus backbones and WAN environments, shown in Figure 20.

SERVICE LEVELS	DIFFSERV	INTSERV	IP Precedence	ATM	802.1P
A	Expedited Forwarding	Guaranteed Service	4,5,7	Constant Bit Rate Variable Bit Rate Real-Time	6
B	Assured Forwarding	Controlled Load Service	3,6	Variable Bit Rate Non Real Time	3,4,5
C	Best Effort	Best Effort	0,1,2,	Unspecified Bit Rate AvailableBit Rate	0,1,2

Figure 20. QoS Mappings For Defined Service Levels.

Adding network-wide QoS that can operate seamlessly across heterogeneous link-layer technologies and a variety of host platforms is complicated. Considering the difficulty of just ensuring the proper operation of the current best effort IP-based network, it makes sense to introduce a small scale of traffic differentiation to the network. Limiting the number of DSCPs to 3, the number of PHBs is limited and consequently this mechanism allows for a large number of individual flows to be aggregated from the point of view of the core device.

2. Functions Performed by Edge Devices

The edge devices, implementing the PDP and PEP functionality in a single box, are configured to perform mapping to DiffServ (the backbone Layer 3 switches map 802.1P prioritization to IP Precedence or DSCP, the edge routers map RSVP messages or IP Precedence to DSCP), traffic shaping, policing and admission control. In addition to the end hosts' priorities, the edge devices can also associate the following fields as the basis for decisions:

- Application TCP/UDP port number (layer 4)
- Source/destination IP address (layer 3)
- Date/ Time of the day

Within an integrated policy-based network, proper network servers, dedicated to perform PDP's control and decision-making operations, can undertake part of the edge devices' tasks. Furthermore, a final association can be made to entitled users with the proper use of authentication servers and directory services.

The edge devices can be QoS-enhanced routers or other QoS-aware devices. Cisco Systems, who is the leading data networking company, has implemented in the latest versions of routers and switches most of the QoS techniques discussed in Chapter 3. With CiscoAssure and Quality Manager, it offers a total policy-based solution that

supports dynamic QoS implementation. Other vendors, such as Nortel, Foundry, Extreme Networks and Juniper, have incorporated advanced QoS features in their products.

Furthermore, there are several software and hardware products made by new companies that specialize in QoS, which are deployed at the edge of the network and provide an integrated solution for traffic conditioning and optimal use of network resources. They also optimize the performance of existing edge routers by offloading them of QoS processing overhead. Companies such as Sitara Networks, Packeteer and Allot Communications deliver QoS solutions through platforms that integrate highly scalable bandwidth management and end-to-end comprehensive QoS. These platforms are positioned into the network between WAN routers and LAN backbone switches, as shown in Figure 21, monitor and analyze real-time traffic and are capable of a broad set of point-solution functions: TCP rate-shaping, classification of IP precedence and DiffServ fields, all types of sophisticated queuing, packet-size optimization, dynamic allocation of bandwidth, HTTP caching and policy management.

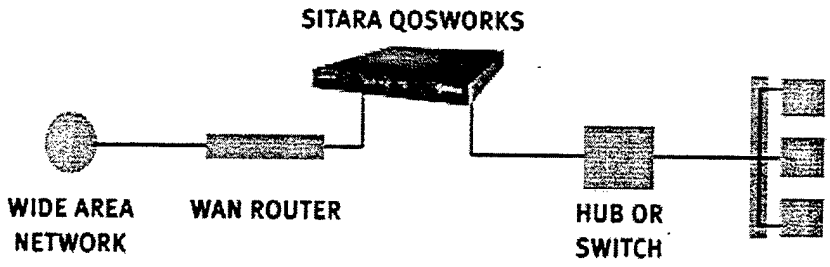

Figure 21. Sitara QoS Integrated Solution.

Policy management of devices can be achieved individually or by using a central policy manager that enables administrators to set global policies, as well as monitor, enforce and refine them dynamically.

3. Incorporation of Bandwidth Brokers

A scalable direction for dynamic provisioning and QoS implementation is to incorporate bilateral QoS-aware devices that will enable communication for exchanging

bandwidth requirements and QoS policy information between adjacent networks and domains. The concept of bandwidth brokers seems appropriate to fill that role. Each bandwidth broker, operating at a higher layer than PDP within a network, takes administrative control of a domain and controls how traffic flows through one's network so as to optimize resource utilization and network performance. Bandwidth brokers peer to ask for and answer with admission control decisions for aggregates and exchange traffic.

Adjacent domain bandwidth brokers negotiate in order to determine the nature and extent of traffic that will traverse across their common boundaries [Ref. 21]. As part of this process, each bandwidth broker describes its requested level of service to its neighbor's bandwidth broker. The latter provides an admission decision based on its resource availability, bilateral arrangements as well as the set of administrative policies in effect. The decision is enforced by monitoring incoming flows into each domain.

In general, a bandwidth broker may receive a resource allocation request either from an element in the domain that the bandwidth broker controls, or a request from a peer (adjacent) bandwidth broker. In any case, the bandwidth broker responds to this request with a confirmation of service or denial of service. The request may have certain effects upon the network, such as altering the router configurations at the access inter-domain borders, and/or internally within the domain, and possibly generating additional messages requesting downstream resources.

To illustrate the operation of bandwidth brokers practically, in Figure 22, suppose that an end user from network A signals to bandwidth broker A that wants to have a videoconference with a user in network D. This signaling can be triggered either by explicit reservation or by proper prioritization of packets. Bandwidth broker A receives the request and checks whether it can handle this request through network B. Bandwidth broker B performs admission control based on whether there are sufficient premium

resources available, and whether the predefined agreements allow such requests. If the request is not admitted, i.e. network B does not have enough resources to accommodate the need for the videoconference, bandwidth broker A tries the path through C-D. If bandwidth brokers C and D agree, the request can be sent through A-C-D, instead of being rejected. The edge routers are being instructed by bandwidth brokers to allow the given traffic to flow through A-C-D.

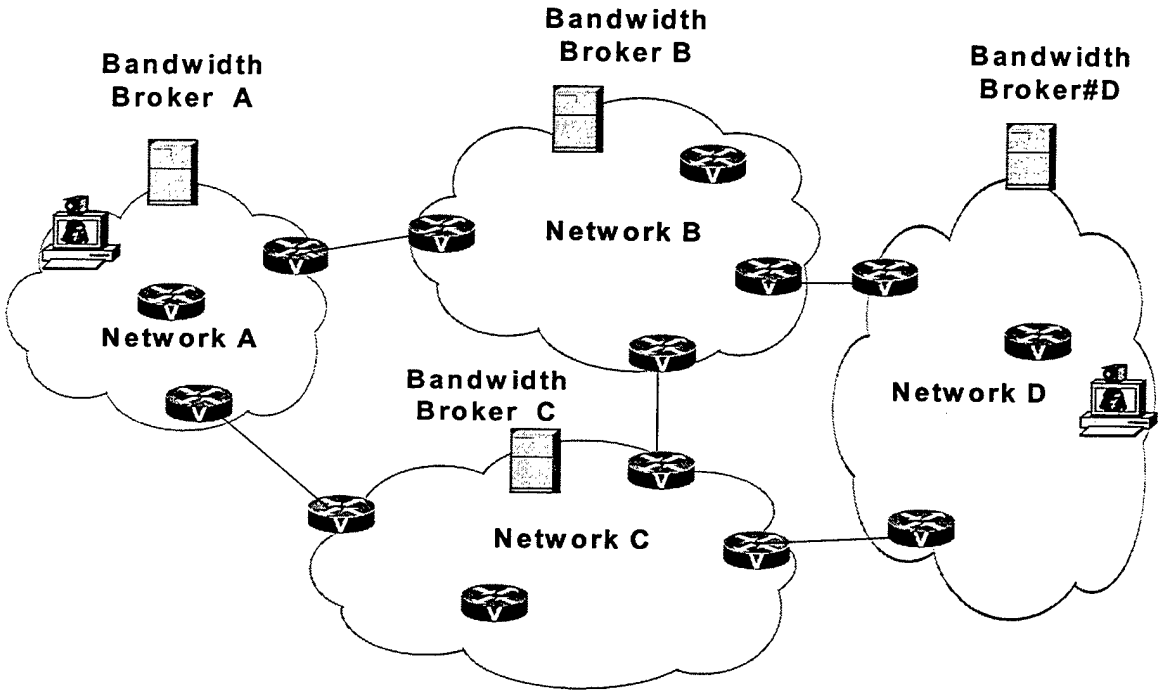

Figure 22. Bandwidth Brokers Implementation.

The concept of bandwidth brokers is very appealing, but has to be stabilized and fully standardized before being implemented. While it provides a sophisticated and scalable approach to resource management, the generated signaling overhead for the communication between bandwidth brokers must be taken into account.

VI. CONCLUSIONS AND RECOMMENDATIONS

A. CURRENT STATE OF NETWORKS

This thesis explored issues relating to the nature of voice and IP data networks, convergence, application demands, QoS, and bandwidth management that will shape networking architecture within an organization in the future. It took a holistic approach to the deployment of IP-based network of an organization that needs to support differentiated traffic of data, voice and video applications.

The connection-oriented nature of the voice network and the support for relatively low bandwidth provided an optimal service for time-sensitive voice conversation, but it is not to provide high-speed digital access and accommodate the expansion of intensive data and video applications. The data network, founded on a connectionless packet-switching scheme, is based on the TCP/IP protocol suite that accomplishes a more robust and dynamic handling of resources. This IP-based network provides a best-effort service model for data applications, but falls short of delivering tight performance guarantees needed for delay-sensitive communication.

The networking community, realizing the advantages of delivering data, voice and video over a single network infrastructure, has started to converge on the IP-based network that is the most convenient platform for the support of integrated traffic and has the potential to serve as the basis of a multiservice network. The gradual integration of voice and video in IP networks, the emergence of new applications and the increased sophistication of existing ones have forced the IP network to carry traffic with diverse requirements that all compete for network resources.

To meet the traffic expectations, the IP-based network infrastructure must evolve and enhance its basic service model in many subtle ways, enabling the reliable and predictable service to differentiated network traffic.

B. RECCOMENDATIONS AND EXPECTATIONS FOR FUTURE NETWORKS

This thesis proposed an organizational IP network infrastructure that can address the current demands and is flexible enough to support future requirements. Essentially, a balanced approach between broadband network capacity delivered at users' doorsteps (through proper deployment of fiber cabling), and robust QoS implementation on the network will manage to accommodate current and future needs.

Fiber cabling was proposed because it enables the greatest number of current and future value-added applications to the users. It is the fiber at the desktop that will revolutionize the capacity and will provide ample bandwidth. Additionally, QoS is essential to ensure sufficient performance and resource management. The proposed QoS implementation attempted to integrate and combine the advantages of the current QoS solutions. It determined the QoS functions to be performed as follows:

- Classification and prioritization by the end systems (applications/users).
- Mapping, traffic shaping, policing and admission control by the edge network devices.
- Traffic handling and congestion control by the core network devices.
- Supervision and administration control by smart network devices (bandwidth brokers) that each of them overviews part of the network.

The intention was to affect minimally the connectionless and stateless nature of the IP infrastructure, and alleviate the concerns regarding the technical complexity, stability and scalability of QoS implementation. This implementation needs to be further tested practically and evaluated empirically.

The introduction of QoS is especially important for organizations, such as the military, where there is an increased need for user differentiation and the extensive use of wireless and RF WANs makes overprovisioning much more difficult than terrestrial

infrastructure. The military organization should push aggressively towards the development of QoS and traffic engineering at the borders of the networks, in order to enhance the IP service model for advanced and guaranteed services. With all these in place, the IP-based network will be able to accommodate current and future needs of various applications and operate as a reliable multiservice convergent platform.

There are still many research issues concerning the development and deployment of technology for a QoS-enabled IP infrastructure. Standards continuously evolve, but QoS has already been moved from an academic topic to an essential element of the network. The deployment of QoS mechanisms in IP networks is essential to allow control of various types of traffic, distinguish between several users and provide predictable and deterministic treatment of packet flows. As with any new technology, it will simply take time to evolve and mature.

C. TOPICS FOR FURTHER RESEARCH

As organizations gradually converge to a pure IP infrastructure, it becomes apparent that organizations will become more reliant upon a broadband, QoS-enabled and policy-aware network. Possible topics for further research include the following:

- Design and implementation of a pure IP-based integrated campus network that could handle data, voice and video traffic, based on commercial solutions, such as CISCO AVVID solution and other networking companies. That could provide the test bed for a broader implementation in other military facilities.
- Examination of the technical and economic issues surrounding the proper deployment of fiber optics close to the end users, at the campus network and at residential areas to subscriber users.
- Implementation of an efficient multicasting file broadcast scheme, in combination with the introduction of caching servers in remote locations

that could sufficiently reduce the traffic demand and increase the information availability.

LIST OF REFERENCES

1. Huston, G., *Internet Performance Survival Guide, QoS Strategies for Multiservice Networks*, Wiley, 2000.

2. Armitage, G., *Quality of Service in IP Networks*, MacMillan, 2000.

3 Goralski, W. & Kolon, M., *IP Telephony*, McGraw-Hill, 1999.

4. Douskalis, B., *IP Telephony, the Integration of Robust VoIP Services*, Prentice Hall, 2000.

5. Durham, D. and Yavatkar, R., *Inside the Internet's Resource Reservation Protocol: foundations for Quality of Service*, Wiley, 1999.

6. Caputo, R., *Cisco Packetized Voice and Data Integration*, McGraw-Hill, 2000.

7. Black, U., *Voice Over IP*, McGraw-Hill, 1998.

8. Hersent, O., Gurle, D., & Petit, J.P., *IP Telephony, Packet-based Multimedia Communications Systems*, Addison-Wesley, 2000.

9. Peterson, L and Davie, B., *Computer Networks, A Systems Approach*, Morgan-Kaufmann, 2000.

10. Odlyzko, A., *The Economics of the Internet: Utility, Utilization, Pricing, and Quality of Service*, 1998.

11. Ma, T. & Shi B., *Bringing Quality Control to IP QoS*, Network Magazine, November 2000.

12. Gray, T., *Enterprise QoS Survival Guide*, University of Washington, 1999.

13. *Cisco IOS 12.0 Quality of Service*, Cisco Systems, 1999.

14. Rajan, R., Verma,D., Kamat,S., *A Policy Framework for Integrated and Differentiated Service in the Internet*, 2000

15. *Introduction to QoS Policies*, www.qosforum.com, 2000.

16. Croll, A. and Packman, E., *Managing Bandwidth, Deploying QoS in Enterprise Networks*, Prentice Hall, 2000.

17. Killki, K., *Differentiated Services for the Internet*, MacMillan, 1999.

18. Bernet, Y., *The Complementary Roles of RSVP and Differentiated Services in the Full-Service QoS Network*, IEEE Communications Magazine, February 2000

19. Cole, M., *Introduction to Telecommunications: Voice, Data, and the Internet*, Prentice Hall, 2000.

20. Faynberg, I., Gabuzda, L. and Lu, H.L., *Converged Networks and Services: Internetworking IP and the PSTN*, Wiley, 2000.

21. Concalves, M., *Voice over IP Networks*, McGraw-Hill, 1998.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
   8725 John J. Kingman Road, Ste 0944
   Fort Belvoir, VA 22060-6218

2. Dudley Knox Library 2
   Naval Postgraduate School
   411 Dyer Road
   Monterey, CA 93943-5101

3. Professor Dan Boger, Code CS 1
   Naval Postgraduate School
   Monterey, CA 93943-5106

4. Professor Gilbert Lundy, Code CS/Ln 1
   Naval Postgraduate School
   Monterey, CA 93943-5100

5. Professor Rex Buddenberg, Code IS/Bu 1
   Naval Postgraduate School
   Monterey, CA 93943-5118

6. Dr. William Richter 1
   SPAWAR Systems Center
   Charleston, SC 29403

7. Captain Dave Glenn 1
   Commandant US Coast Guard (G-OC)
   Washington, DC 20593

8. Konstantinos Sambanis 5
   Ayt. Irakliou 26,
   Marousi, Athens 15122
   GREECE