



VOICE OVER IP ON CONTENTAWARE™ STRATASWITCH™ BCM5600

10/15/99
Revision 5600A-WP201

REVISION HISTORY

<i>REVISION #</i>	<i>DATE</i>	<i>CHANGE DESCRIPTION</i>
5600A-WP201-R	10/15/99	Initial release

Broadcom Corporation
16215 Alton Parkway
P.O. Box 57013
Irvine, California 92619-7013
© 1999 by Broadcom Corporation
All rights reserved
Printed in the U.S.A.

Broadcom[®] and the pulse logo are registered trademarks of Broadcom Corporation and/or its subsidiaries.
All other trademarks are the property of their respective owners.

TABLE OF CONTENTS

Introduction 1

Application Overview — Voice over IP 1

 PC-to-PC Internet Phone/Video Conferencing 1

 IP PBX 2

 Voice Traffic Characteristics in a Packet Network 2

 Special Treatment of VOIP Traffic 3

ContentAware™ Switching 3

Network Requirements 5

Summary 6

LIST OF FIGURES

Figure:	1: Telephony with Video over Existing IP Network	1
Figure:	2: Telephony over IP PBX Network	2
Figure:	3: Voice Transmission over a Lightly-Loaded IP Network	3
Figure:	4: Voice Transmission over a Congested IP Network	3
Figure:	5: Traffic Monitoring Flow Between the BCM5600 and the CPU.....	4
Figure:	6: Network Cluster Bounded by StrataSwitch-Based Switches.....	5
Figure:	7: Network Clusters of 802.1p Domains	6



INTRODUCTION

Broadcom's StrataSwitch™ chip enables development of the next generation of switches that prioritize voice, video and other latency-sensitive traffic, ensuring appropriate Quality of Service (QoS) throughout the entire network. Voice and video applications over IP, which require a data network as the transport, are the type of applications that benefit from traffic prioritization in the network. The BCM5600 is marketed towards enterprise LAN switches and enables QoS support at protocol layers L2 (IEEE 802.1p) and L3 (TOS or DiffServ).

This white paper first presents a brief overview of Voice Over IP (VOIP) and two of its flavors in the enterprise world, the characteristics of voice traffic over an IP network, and a discussion of the ContentAware™ feature in the BCM5600, which prioritizes latency-sensitive traffic. The main issue addressed in this paper is how a switch built on the BCM5600 prioritizes and regulates VOIP traffic in a data network.

APPLICATION OVERVIEW — VOICE OVER IP

Traditionally, a telephone call in a Public Switch Telephone Network (PSTN) requires an allocation of an exclusive, full-duplex hardware connection between the two end-points. Since the circuit is reserved exclusively for this call, any unused bandwidth resources cannot be shared by other calls. In a typical telephone conversation, over 50% of resources are unused, because telephony is simplex in nature (it is rare for both parties to talk at the same time), with silence between sentences. For this reason, the use of exclusive circuits in a traditional phone network takes up more resources than necessary.

However, if packet-switching technology is used for a telephone call, there is no *a priori* allocation of a circuit, thus no reservation of resources that cannot be shared. Packet-switching methodology thus takes full advantage of all available bandwidth in the network. That is how data networks, especially IP-based networks, emerged as an alternate transport for telephone service.

Voice Over IP, and more generally, multimedia over IP, schemes have different characteristics in different business models. The following examples present two common scenarios of VOIP in the enterprise environment.

PC-TO-PC INTERNET PHONE/VIDEO CONFERENCING

The earliest form of VOIP was PC to PC conversation. Equipped with multimedia hardware and software, a PC-based user can converse with another PC-based user if both PCs are connected to an IP network. The call signaling, voice compression, tone generation, etc. are all carried out in the PC. If the calling party knows the called party's IP address, the call can be made directly by entering it. If the called party's IP address is unknown, an intermediate server, known as an Internet Locator Server (ILS), locates the called party if he or she has registered with the server upon login. The PC to PC "call" scenario, with video added, is illustrated in Figure 1. In this mode, the users run standard video conferencing client applications, such as the Intel ProShare Video Conferencing System, Microsoft NetMeeting, and so on.



Figure 1: Telephony with Video over Existing IP Network

IP PBX

Another common scenario for VOIP is IP PBX. In a large corporation with offices nationwide, the corporate data network can carry both voice and data traffic between offices. To support this form of VOIP, the functions of the IP phone gateway, the PBX, and possibly the call center are packaged into one turnkey solution, the IP PBX. Offices installed with IP PBX may also do without wiring for conventional desktop PBX phones by using special IP phones that plug into a standard LAN connection. In this scheme, individual IP phones and the IP PBX are devices on the LAN addressable by standard IP addresses. The IP PBX also interfaces with the standard PSTN to connect to phones in the PSTN.

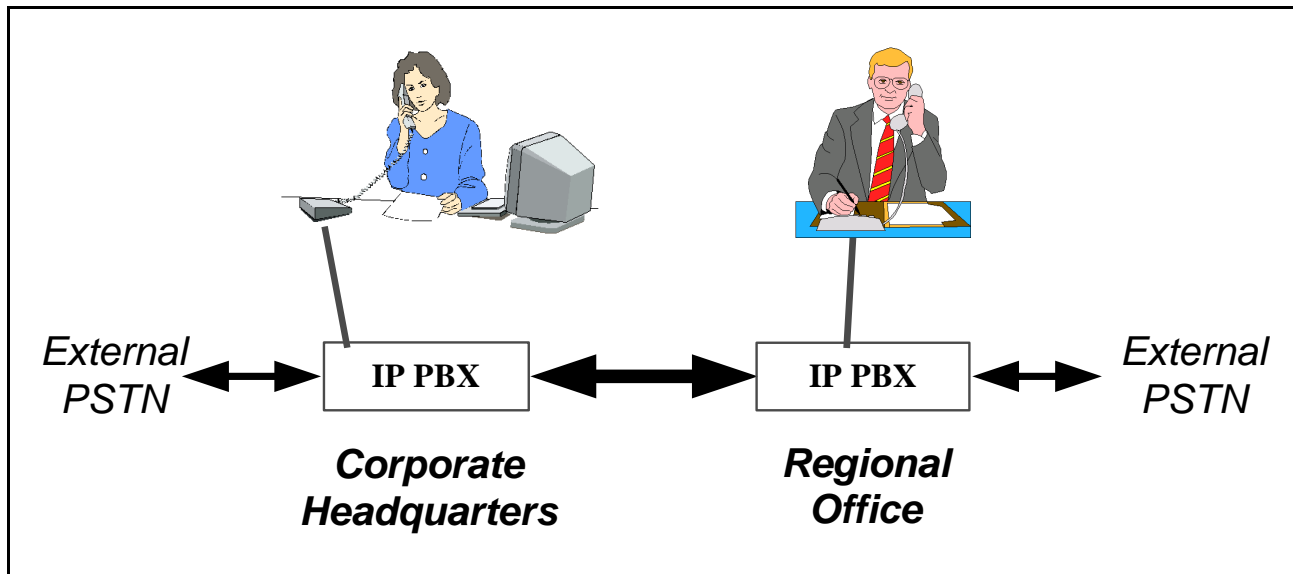


Figure 2: Telephony over IP PBX Network

The call scenario between a corporate headquarter and a regional office via IP PBX is illustrated in Figure 2.

VOICE TRAFFIC CHARACTERISTICS IN A PACKET NETWORK

Voice traffic has a different characteristic from ordinary data traffic. First and foremost, voice traffic is real-time traffic that is sensitive to latency and packet loss. Drops or delay of voice packets in the network affect the quality of the audio, and a long latency adversely affects the quality of conversation, or interactivity. The maximum tolerance of mouth-to-ear latency is about 250 to 300 ms, making this the worst-case for latency in most VOIP products. The larger the time window to play back the voice packet at the destination, the greater chance for it to arrive on time, but the feeling of interactivity is progressively degraded as the time window lengthens. The smaller the time window, the better the subjective feel of interactivity, but the less the chance for the packets to arrive within the time window. In order to make a VOIP call indistinguishable from a regular phone call, the right balance between latency and packet loss tolerance must be maintained.

The parameter that determines the quality of a VOIP call is the speed of voice travel through the network. When the conversation is transmitted, it is divided into small units before being sent through the packet network. These units, called voice frames, consist of a very short duration (from 10 - 30 ms) section of audio. Regular digitized voice without PCM (Pulse Code Modulation) code compression requires 64 KBPS of bandwidth, or 640 bits per 10ms-voice-frame. To reduce bandwidth consumption, compression is used to pack digitized audio into voice frames of smaller size. If the voice is compressed to 8 KBPS using a standard algorithm, each 10ms voice-frame requires only 80 bits. In normal practice, several continuous voice frames are assembled together to form a voice packet before being transmitted over the network. This reduces the band-

width overhead of the packet header. Figure 3 illustrates a voice transmission scenario under typical traffic loading conditions.

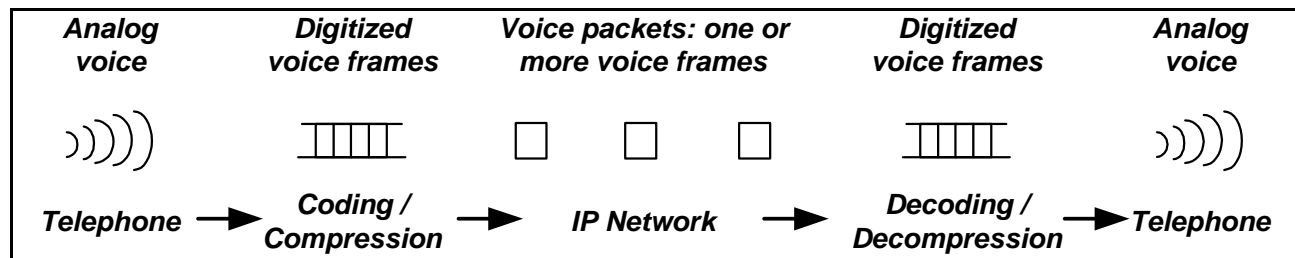


Figure 3: Voice Transmission over a Lightly-Loaded IP Network

Compression requires more processing power in the VOIP equipment and introduces algorithmic delay in the process itself. It is worth the overhead of compression only if the traffic is time sensitive and timely delivery to the destination is critical.

SPECIAL TREATMENT OF VOIP TRAFFIC

Since it is critical for voice traffic to reach its destination in a timely manner, special treatment should be given to it. Because of the time window to play back a certain voice frame, any packet arriving late is treated as lost.

In a heavily congested network voice packets have more difficulty arriving at the destination within the correct time-window. Compared to the process in Figure 3, Figure 4 shows how network congestion can adversely affect the quality of the conversation. If the voice packets fail to arrive at the destination within the time window, the destination decoding / decompression system can run out of voice frames to play back, resulting in silence or unacceptable noise. To prevent degradation, it is essential to minimize voice traffic transmission time in the network.




Figure 4: Voice Transmission over a Congested IP Network

CONTENTAWARE™ SWITCHING

The BCM5600 is the foundation of the next generation intelligent gigabit switch at the leading edge of network performance. With ContentAware automated address learning, the BCM5600 can also distinguish types of traffic and can apply control based on data type to deliver voice and multimedia data to their destinations faster and more accurately than other switching solutions.

As we have seen, in VOIP sessions with audio-only or audio/video calls, late delivery adversely affects call quality. By sorting out voice, video and data traffic, the BCM5600 can assign higher priorities to real-time voice and video traffic at the expense of regular data traffic. The switching of data within the switch is then based on packet priorities. The higher the priority of a packet, the faster the packet can be switched out.



To distinguish voice and video traffic from data traffic, StrataSwitch relies on a CPU “traffic cop”, which monitors control packets of a communication session (e.g. a NetMeeting or an IP phone call) involving real-time traffic. It in turn informs the BCM5600 to treat subsequent real-time traffic differently. A flow diagram of traffic monitoring is illustrated in Figure 5.

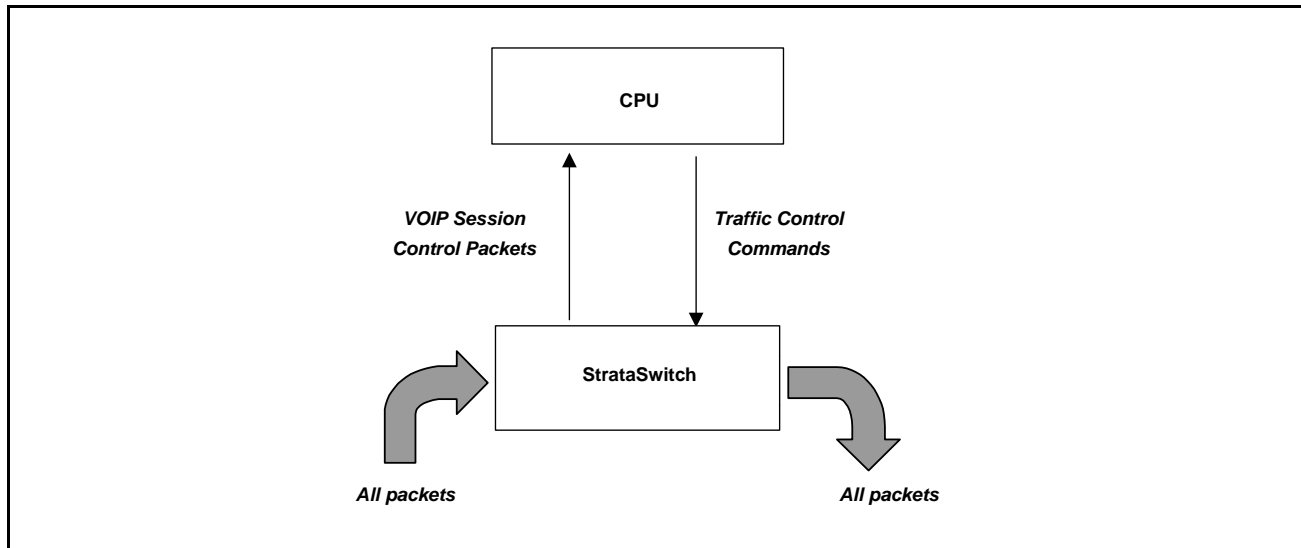


Figure 5: Traffic Monitoring Flow Between the BCM5600 and the CPU

The BCM5600 can distinguish audio, video and data traffic. Besides Voice Over IP, other form of multimedia communications over IP networks can also benefit from the traffic shaping and conditioning the BCM5600 provides:

- IP telephony - IP-PBX, long distance telephony over IP, local phone service over IP
- IP-based video conferencing - desktop, group-based system, standalone set-top system
- streaming audio - live Internet radio, Internet music
- streaming video - Internet telecast, video on demand, video clips

All these applications involve time-sensitive data - untimely delivery makes the audio and video experience undesirable.

To enable VOIP applications, the networking domain must be QoS aware. Switches and routers within the domain must support QoS for this scheme to work.

NETWORK REQUIREMENTS

To take advantage of the QoS features in BCM5600-based switches, the underlying network has to be QoS compliant. This section illustrates a situation that can benefit from the BCM5600's QoS features.

Figure 6 shows a LAN using StrataSwitch-based switches. In this example, all the BCM5600-based edge switches support the IEEE 802.1p priority scheme. By definition, if all the switches or routers within a cluster support 802.1p, this network cluster is classified as an 802.1p domain.

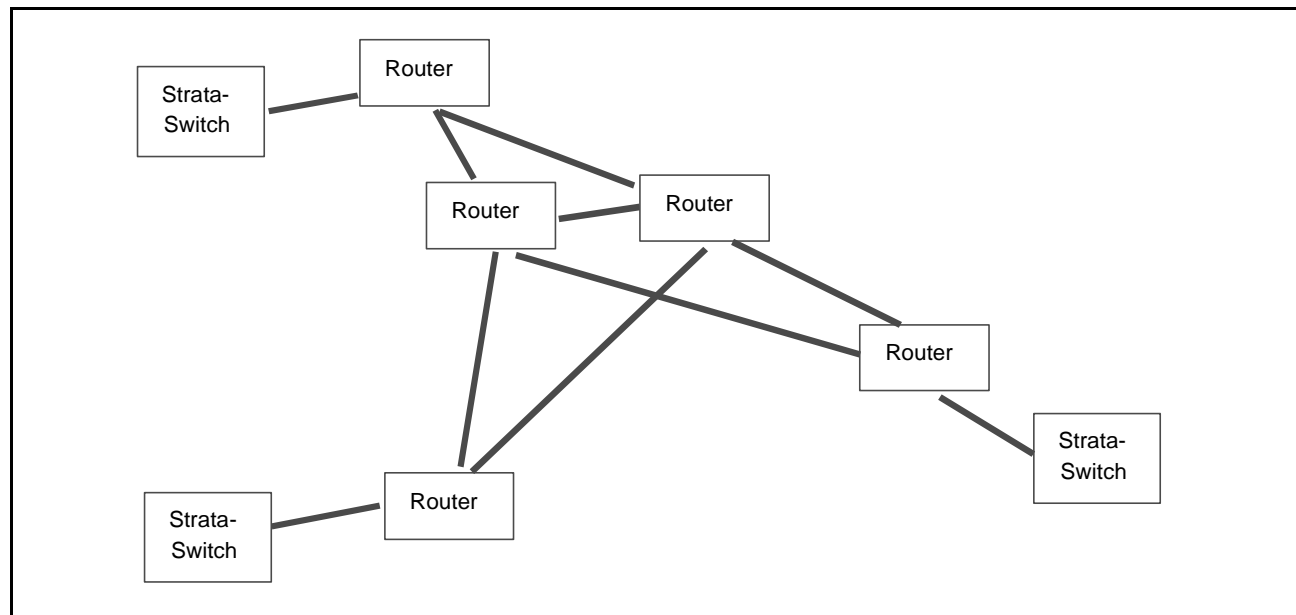


Figure 6: Network Cluster Bounded by StrataSwitch-Based Switches

An 802.1p domain is essential to take advantage of the QoS support provided by the BCM5600. When StrataSwitch processes a stream of Ethernet frames, it assigns a higher priority to the time-sensitive frames. The data with the newly assigned priority stay with the frame while it travels through the 802.1p domain, thus eliminating the need for frame reclassification at the next hop. As a result, the larger the 802.1p domain, the longer the frame can enjoy higher priority switching.

Suppose a cluster of 802.1p domains, as illustrated in Figure 7, form a network. In this case, the network is 802.1p enabled from one end to the other. Any VOIP call within this network can obtain the benefit of 802.1p QoS. Figure 7 illustrates one such connection scenario between host A and host B in an 802.1p - compliant network.




Figure 7: Network Clusters of 802.1p Domains

The BCM5600 also supports IP TOS and DiffServ for service classifications, so QoS support for VOIP traffic is also available in TOS domains or DiffServ domains.

SUMMARY

With ContentAware capability, the BCM5600 chipset brings time-sensitive multimedia applications to a new era. The road-block for quality multimedia applications has been a bandwidth limit that increases latency of multimedia data transmission, degrading the quality of multimedia session. For example, the quality of an IP phone call is currently often not comparable to a POTS call, due to latencies, packet delays, and drops. However, with the BCM5600 ContentAware capability, the biggest obstacle is removed. The BCM5600 can classify and reshape traffic to switch voice and multimedia traffic faster, providing timely delivery of such traffic from origin to destination. As a result, this technology makes the quality of an Internet phone rival a POTS connection.



INDEX

	A		N
audio latency tolerance		network clusters	
audio quality and latency		network congestion and voice quality	
	C	nternet Locator Server	
ContentAware			P
ContentAware automated address learning		packet loss	
	D	Public Switch Telephone Network	
DiffServ		Pulse Code Modulation	
	I		Q
IEEE 802.1p		QoS	
IP PBX		Quality of Service	
	L		T
latency		TOS	
			V
		voice frames	

Broadcom Corporation

16215 Alton Parkway
P.O. Box 57013
Irvine, California 92619-7013
Phone: 949-450-8700
Fax: 949-450-8710

Broadcom Corporation reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.
Information furnished by Broadcom Corporation is believed to be accurate and reliable. However, Broadcom Corporation does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

