Communicate Simply

White Papers by Polycom

Network-based Quality of Service

For IP Video Conferencing

Timothy M. O'Neil
Director of Technical Marketing
Polycom Video Communications

POLYCOM®

Table of Contents

Introduction

This white paper provides a basic, although technical, review of network-based Quality of Service (NQoS) architectures and processes. The white paper specifically discusses the implementations of Polycom applications in a NQoS environment. The discussions are a small subset of the myriad of NQoS services and architectures that exist.

Because there are many possible architectures that are not covered here, see the following sections in this white paper to obtain more information about QoS architectures in general:

- "References" on page 16
- "Additional Information" on page 17.

For a high-level overview of application-based and network-based QoS, see the white paper *Quality of Experience and Quality of Service for IP Video Conferencing*, which is available on the Polycom web site (http://esupport.polycom.com/whitepapers.html).

Information Sources

This white paper has derived some of its information from existing publications, such as other white papers and RFCs. In addition, some terminology and descriptions directly pertain to a specific publication. In these cases, a number, in brackets, follows the term or description. The number identifies the specific publication that contains the information about that term or description. The publications are numbered and listed in "References" on page 16. For example, this white paper contains the term *dynamic bandwidth allocation (DBA) [1]*. The [1] identifies the number of the publication that contains the information about dynamic bandwidth allocation (DBA). In this case, the DBA information is derived from *Dynamic Bandwidth Allocation for Video Conferencing in Lossy Packet Switched Networks*.

Network-based Quality of Service (NQoS) Service Levels[4]

The following NQoS service levels are implemented in layer two (Data-Link) and layer three (Network) of the ISO application model. See Table 1 on page 3 for information about the service levels and their associated services.

- **Best effort service:** This is considered a *basic* service and is also considered soft NQoS. Without a significant over-subscription of bandwidth, best effort service should not be considered a true NQoS service. Best effort service is typically not suited for real-time applications, such as video and VoIP, and does not have the ability to separate flows into a category or class of service. However, most networks use best effort service for their initial video and VoIP services.

  With best effort service, you could successfully use separate network paths for IP video and VoIP, with properly over-subscribed bandwidth.

  Using H.323 to H.320 gateways, instead of corporate IP WAN links, is also considered a NQoS implementation. In this case, the WAN link inherits the guaranteed NQoS of the ISDN network.
- **Differentiated service:** This is also considered soft NQoS. Traffic flows (applications/protocols) are grouped into classes based on the level of service required to operate most efficiently/effectively. Network devices can then differentiate flows and apply services according to the programmed NQoS mechanism for the flow classes. Differentiated service:
  - Allows for the preferential treatment of flows, without an absolute guarantee
  - Works well for bandwidth-intensive data applications like video and VoIP.

- **Guaranteed service:** This is considered hard NQoS. Guaranteed service requires a reservation of network resources in order to meet specified service classification requirements.

Table 1: NQoS Service Levels with Associated Service and Layer Description

Service Level	Associated Service	Layer Description
Best effort	Basic connectivity	• Asynchronous transfer mode (ATM) • Unspecified bit rate (UBR) • Frame-relay • Committed information rate (CIR)=0
Differentiated	• CoS committed access rate (CAS) • Weighted fair queuing (WFQ) • Weighted early random detection (WRED)	IEEE 802.1p
Guaranteed	Resource Reservation Protocol (RSVP)	• Asynchronous transfer mode (ATM) • Constant bit rate (CBR) • Frame-relay • Committed information rate (CIR)

Network-based QoS Functional Areas

Polycom terminals work with network-based QoS (NQoS) in three different areas of the network:

- Admission control
- Resource allocation
- Gatekeepers.

Admission control and resource allocation manage network resources to provide for NQoS. Call admission control (CAC) is an H.323 service provided with gatekeepers. Gatekeepers specifically work with H.323 call admission control (CAC).

The following sections describe the three NQoS functional areas.

Admission Control

Admission control combines bandwidth control and policy control to provide NQoS services. In a typical scenario, applications, such as Polycom video communications terminals, request a particular NQoS for their traffic. The devices in the network, through which this traffic passes, can either grant or deny the request, depending on various factors, such as capacity, load, policies, and so on. If the request is granted, the application has a contract for that service. The contract is honored as long as there are no disruptive events, such as network outages. If the network is disrupted, Polycom's IPriority™ facilities and Polycom Video Error Concealment (PVEC) try to compensate for packet loss. Also, dynamic bandwidth allocation (DBA) [1] reduces bandwidth.

Resource Allocation

Resource allocation is also referred to as *queuing and scheduling*. Traffic flows and traffic classes can be assigned to different queues based on their traffic classification and marking. See “NQoS Architectures” on page 6 for more information about traffic classification and marking.

The placement or switching of packets into different queues handles the diverse service needs of marked packets. Special requirements, such as low delay, can be provided to packets by servicing their queue more often. Queue selection is based on the classification of the marked packets. NQoS queuing determines packet order and priority within the queues.

Priority queuing is an example of a simple queueing model. In this model, packets are forwarded based strictly on the priority of the queue to which the packets are assigned. Packets in the high priority queue are forwarded first. If there are no packets in the high priority queue, the packets in the medium priority queue are forwarded. If there are no packets in the medium priority queue, the packets in the next priority queue are forwarded, and so on.

There are several other types of queuing, for example, weighted fair queuing (WFQ) and round robin (RR).

Call Admission Control - Gatekeepers

One of the most important components in an H.323-based video communications network is the gatekeeper. Polycom's PathNavigator™, which is the Premier Call Processing Server solution, contains a gatekeeper. The gatekeeper is a network administration component that administers terminal and infrastructure components registered both with itself or with other gatekeepers. Although the gatekeeper is an optional component, it has become an integral part of video communications deployments.

Regarding NQoS services, a gatekeeper's main function is admission control. Gatekeepers also can provide alternate NQoS-enabled routes over H.323 to H.320 gateways.

Gatekeepers perform admission control services by managing a pool of available bandwidth. The available bandwidth capacity is the number of potential calls and their bandwidth requirement. The gatekeeper automatically matches the available bandwidth capacity, which is provided during the network architecture phase of deployment, to the bandwidth pool.

The following explains how traditional gatekeepers handle call initiation. During call setup, the gatekeeper subtracts any bandwidth that the new call requests from the bandwidth pool. When the call ends, the gatekeeper adds the bandwidth back to the pool. If there is not enough bandwidth in the pool for the call, the call is rejected.

The gatekeeper provided with Polycom's PathNavigator is capable of extending this negotiation at call setup and renegotiates the bandwidth request. PathNavigator lets Polycom terminals renegotiate to a lower bandwidth, which lets the call connect, instead of being rejected. If you use PathNavigator, you must still do a detailed analysis before gatekeeper deployment to ensure effective bandwidth administration and full access to network resources for authorized terminals (users).

Gatekeeper bandwidth control is static. Gatekeepers currently base their bandwidth decisions on statically configured bandwidth pools. The gatekeeper is not aware of how a particular call flows through the network. The gatekeeper does not know which network devices will service the call, or whether the devices have sufficient capacity. This is another potential failing of network-based QoS that can easily be compensated for by Polycom's IPriority facilities. Polycom Video Error Concealment (PVEC) tries to conceal

packet loss, and dynamic bandwidth allocation (DBA) [1] works to reduce bandwidth.

NQoS Architectures

The use of NQoS elements described in the preceding sections provide different network-based QoS architectures.

Three of the most popular NQoS architectures are:

- IP precedence
- Differentiated services (Diffserv)
- Integrated services (Intserv/RSVP).

IP precedence and Diffserv are discussed in the following sections. For information about integrated services (IntServ/RSVP), see “References” on page 16 and “Additional Information” on page 17 for a list of documents to consult.

IP Precedence

IP precedence is one of several NQoS architectures. The IP precedence field in an IP packet's header indicates the priority with which a packet should be handled. The IP precedence field is composed of three bits in the type of service (ToS) byte of the packet's IP header. Figure 1 shows the construction of precedence bits in the ToS byte. Polycom ViewStation™ video communications systems support IP precedence.

Figure 1:Internet Protocol (IP) Precedence Bits in the Type of Service Byte

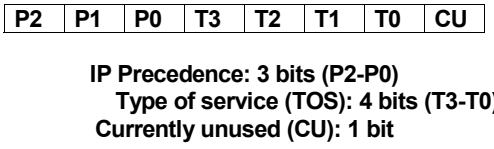

Table 2 shows the IP precedence bits and precedence names associated with IP precedence values. By default, all IP routing protocols use IP precedence bit 6. IP precedence bit 7 is reserved for network traffic control. It is recommended that you not use bits 6 and 7 for user traffic.

Table 2: IP Precedence Values, Bits, and Names

IP Precedence Value	IP Precedence Bits	IP Precedence Names
0	000	Routine
1	001	Priority
2	010	Immediate
3	011	Flash
4	100	Flash Override
5	101	Critical
6	110	Internetwork control
7	111	Network control

Differentiated Services[4] [6] [7]

Differentiated services (DiffServ) defines a set of NQoS mechanisms to implement scalable service differentiation in IP environments. The Diffserv architecture specifies the basic network policy mechanisms for packet handling. Because Diffserv provides multiple services for packet handling, it is referred to as a *multiple service model*.

Diffserv defines a field in the IP header called the Differentiated Services Code point (DSCP). The DSCP is a six-bit field, as shown in Figure 2 on page 8, and spans the fields formerly used as type-of-service and IP precedence fields in the Type of Service (ToS) byte. The inclusion of the packet handling information in the packets themselves is referred to as *in-band signaling*.

Figure 2:DSCP Bits in the Differentiated Services (DS) Byte

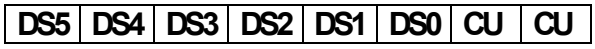

DSCP: 6 Bits (DS5-DS0)
Currently Unused: 2 Bits

Hosts, which send traffic that requires NQoS into a network that supports DiffServ, mark each packet with a DSCP value. Devices within the Diff-Serv network use these values to classify the traffic into distinct service classes. The packets in the stream are queued and scheduled, based on the configuration of the service class. One of the most popular DSCP values is *expedited forwarding* (EF), which marks the packet for low latency, low loss service.

Signaling, or marking of packets, is carried by the packets themselves. This is referred to as in-band signaling. For example, if a packet arrives at a router with an EF marking, and passes the policing and shaping functions, the router schedules the packet into the highest priority queue that the router supports. This allows the packet to be forwarded with minimal queuing delay, which results in very little delay and jitter. EF marking is commonly used by latency-sensitive and loss-sensitive applications, such as Voice over IP (VoIP).

EF marking has good packet handling for VoIP applications. However, the following limitations[1] preclude using EF marking for video applications:

- "Video packets are much larger than voice packets, usually as large as the maximum link MTU size. If video packets are marked as EF, then they will be allowed into the same priority queue as voice. If a small VoIP packet enters the queue behind a large video packet (or worse, a whole bunch of such packets), then the delay in the VoIP packet will increase, possibly substantially. This will adversely affect the performance of VOIP applications."
- "Given that most EF queues are very small, using them for video traffic may lead to dropped packets."
- "Video coders tend to have much longer coding delays than voice coders, and hence giving the audio

1. The limitations, in quotations, are provided by Subha Dhesikan, Software Engineer Enterprise, Cisco Inc.

stream(s) of a videoconference absolute priority only causes them to arrive early and be held in order to achieve lip sync. Therefore, it does not help to put voice packets associated with a videoconference in a queue with better service than that given to the video packets."[1]

IP precedence and DSCP have established mappings between their respective levels of priority, as shown in Table 3.

Table 3: IP Precedence to DSCP Mapping

IP Precedence Value	DSCP
0	0-7
1	8-15
2	16-23
3	24-31
4	32-39
5	40-47
6	48-55
7	56-63

1. The limitations, in quotations, are provided by Subha Dhesikan, Software Engineer Enterprise, Cisco Inc.

The following highlights the advantages of using Diffserv:

- **It's simple:** Diffserv is comparatively simple to adopt. For any application, it usually takes less effort to mark the packets using DSCP than it does to introduce an additional signaling/marking scheme. Although not all Polycom systems currently support DSCP, Polycom is evaluating DSCP marking for the future. For a list of Polycom products that support DSCP, see the white paper *Quality of Experience and Quality of Service for IP Video Conferencing*, which is available on the Polycom web site (http://esupport.polycom.com/whitepapers.html). The information is in Table 2 in the white paper.
- **It's scalable:** With this architecture, the network device does not process the traffic flow-by-flow. Therefore, the network device is not required to maintain any elaborate state information to identify the flows. Instead, the arriving traffic is aggregated into a few traffic classes, based on the DSCP marking. Because this process does not require substantial processing overhead, the network device can handle a large number of flows without experiencing deterioration in performance. This allows deployment of Diffserv without the scalability concerns associated with flow-based systems.

The disadvantage of using Diffserv is the lack of admission control. Diffserv does not have a facility for admission control. For example, it is possible to ask the network device for resources with the DSCP marking in the packet. However, there is no mechanism for the network device to determine whether this traffic would cause overloads, or to report the possibility of overloads back to the application. Therefore, the application is not aware of the service that the network device provided. The application must observe its own end-to-end performance to adapt to the service it is actually receiving from the network device. By trying to use unavailable network resources, the application might disrupt services for other traffic that shares its class. This potential shortcoming of network-based QoS can easily be compensated for by Polycom's IPriority facilities. Polycom Video Error Concealment (PVEC) tries to conceal packet loss, and dynamic bandwidth allocation (DBA) works to reduce bandwidth.

Traffic Types, Traffic Classes, Priority Settings and VLANs

Up until now, the white paper has covered IETF recommendations. This section provides IEEE information about traffic types, traffic classes, and priority settings and virtual LANs (VLANs) [3].

Traffic Types and Traffic Classes

There are seven *traffic types*:

- **Network control traffic:** Required to keep the network operational. Network control traffic must get through the network with the highest priority.
- **Voice traffic:** Requires less than 10 ms delay and jitter.
- **Video traffic:** Requires less than 100 ms delay.
- **Controlled load traffic:** Consists of important business applications that are subject to some form of admission control. The admission control can range from bandwidth pre-provisioning to reservation per flow, which is requested at the flow's start.
- **Excellent effort traffic:** Treated with the "best" *best effort service* the network has.
- **Best effort traffic:** Traffic that is flowing on most networks today, for example, traffic on LANs.

- **Background traffic:** Traffic that can be on the network without impacting other network users, for example, bulk transfers.

The seven traffic types map to traffic classes, which are also known as *queues*. You can have from 1-7 traffic classes. If there are seven traffic classes, the traffic types would map to the traffic classes as shown in Table 4.

Table 4: Mapping Traffic Types

This Traffic Type	Maps to This Traffic Class
Network control	7
Voice	6
Video	5
Controlled load	4
Excellent effort	3
Background	1
Best effort	0

Traffic class 2 is reserved.

If there are less than seven traffic classes, each traffic type does not map to a unique traffic class (queue). The traffic classes would handle more than one traffic type. For example, if there only are two traffic classes, some of the traffic types would map to one of the traffic classes, and the rest would map to the second traffic class.

How traffic types map to specific traffic classes depends on the number of traffic classes available. For more information about traffic types and traffic classes, see the white paper *Quality of Service in Frame-Switched Networks*, from Nortel Networks [3].

Priority Settings

For NQoS, you can classify information flowing across a network by priority. The IEEE 802.1p standard defines how to set priority bits [3], which is beyond the scope of this white paper.

IEEE 802.1p also has a standard to map priority settings to traffic classes. There are eight traffic classes (0-7), although traffic class 2 is reserved. There are eight priority settings, ranging from 0-7. A higher priority setting does not automatically result in a packet receiving higher-priority handling. The priority settings, in order of increasing packet-handling priority, are:

- 1 (least)
- 2
- 0 (default)
- 3
- 4
- 5
- 6
- 7 (most)

Mapping of Frames

The mapping of frames, based on their priority setting, to a specific traffic class (queue) depends on the number of traffic classes (queues) that are available. Table 5 shows this mapping based on the IEEE 802.1p standard. The left side of the table shows the priority setting of the frame, which can be 0-7. Across the top row, the total number of traffic classes that are available is shown at the top of each column (1-8). Each cell in the table shows the specific traffic class to which the frame is mapped, based on the frame's priority setting and the total number of traffic classes available.

Table 5: Mapping of Priority Settings and Traffic Classes

Priority Setting of Frame	Total Number of Traffic Classes (Queues) Available							
	1	2	3	4	5	6	7	8
	Specific Traffic Class (TC) to Which the Frame is Mapped							
0 (Default)	TC 0	TC 0	TC 0	TC 1	TC 1	TC 1	TC 1	TC 2
1	TC 0	TC 0	TC 0	TC 0	TC 0	TC 0	TC 0	TC 0
2	TC 0	TC 0	TC 0	TC 0	TC 0	TC 0	TC 0	TC 1
3	TC 0	TC 0	TC 0	TC 1	TC 1	TC 2	TC 2	TC 3
4	TC 0	TC 1	TC 1	TC 2	TC 2	TC 3	TC 3	TC 4
5	TC 0	TC 1	TC 1	TC 2	TC 3	TC 4	TC 4	TC 5
6	TC 0	TC 1	TC 2	TC 3	TC 4	TC 5	TC 5	TC 6
7	TC 0	TC 1	TC 2	TC 3	TC 4	TC 5	TC 6	TC 7

Here are several examples to help explain Table 5 and how mapping works. If there is one traffic class (one queue) only, all frames are mapped to traffic class 0, regardless of priority settings. This is illustrated in column 1 of the table. With only one traffic class, no differentiated services are available and frames are processed on a first-in first-out (FIFO) basis.

If there are two traffic classes (two queues), frames with a priority setting of 3 or less map to traffic class 0. Frames with a priority setting of 4 or greater map to traffic class 1. This is illustrated in column 2 of the table.

If there are eight traffic classes (queues), each priority setting maps to one traffic class, as shown in column 8 of Table 5. Priority setting 0, which is the default, does not map to traffic class 0. Priority setting 0 maps to traffic class 2. The default priority of 0 is used for interoperability with previous mapping methods.

For more information, see the white paper *Quality of Service in Frame-Switched Networks*, from Nortel Networks [3].

The mapping also provides switches to operate with either of two priority modes:

- Best effort
- Priority.

Best effort mode, which is the default, treats all traffic with the same priority. Packet frames are dequeued for output transmission on a first-in first-out (FIFO) basis.

Priority mode lets high priority traffic flow through the switch fabric using a high priority data path. The output buffers are reserved for high priority traffic. You can prioritize frames by:

- Individual ports
- Specific MAC addresses
- Virtual LANs (VLANs)
- IP flow definitions
- IP filter specifications
- IP precedence/ToS bit setting
- IEEE 802.1Q VLAN tagging.

Frames are serviced as follows:

- For ports configured as *HighPriority=true*, all frames the port receives are assigned to the high priority queue
- For MAC addresses configured as *Priority=high*, all frames that originate from the MAC address are assigned to the high priority queue
- For VLANs configured as *HighPriority=true*, all frames that any of the active ports of the VLAN receive are assigned to the high priority queue. High priority can be set for a:
  - Port-based VLAN
  - Source IP subnet-based VLAN
  - Source MAC-based VLAN
  - Protocol-based VLAN
  - VLAN with a user-defined protocol.

Summary

The goal of IP video/VoIP communications is to provide the highest quality audio and video experience, which is the simplest to deploy and manage, and is the most scalable, for both the user and network administrator. One of the challenges network administrators face in achieving this goal is obtaining acceptable service from the network. The network-based tool that is available to meet this challenge is network-based Quality of Service (NQoS).

Currently, two additional facilities exist to help achieve the highest possible quality user experience:

- H.323 gatekeepers
- Polycom's IPriority initiative for terminal and video infrastructure.

Gatekeepers limit bandwidth access to H.323 sessions. IPriority compensates for situations where NQoS is not always available or when soft NQoS implementations, such as best effort service and Diffserv, cannot hold a network's traffic variability at a constant rate.

Terminology

Table 6 explains the various terms used in this white paper.

Table 6: Terminology

Term	Definition
Admission control	The mechanism that decides whether the network device has sufficient resources to supply the requested NQoS.
Differentiated Services Code Point (DSCP)	The DSCP is a six-bit field, which spans the fields formerly known as the type-of-service (ToS) fields.
Flow	A set of packets belonging to one instance of the application, identified by some combination of source address, source port, destination address, destination port, and protocol identifier.
Jitter	Variation in delay.
Network device	This refers to a device in the network that handles traffic. Routers and switches are examples of network devices.
Over-subscription	This refers to applying more bandwidth to the problem than is required.
Quality of Service (QoS)	The type of service that network devices provide.
Resource	This refers to all the factors in the network device that affect the forwarding of packets, such as bandwidth on an interface, queues, processing power, and so on.
Traffic	This refers to one or more flows that traverse through the network.

References

1. *Dynamic Bandwidth Allocation for Video Conferencing in Lossy Packet Switched Networks*
   Richard Flott, Michael Horowitz
   Polycom, Inc.
   September 2000
2. *A Technical Description of Polycom Video Error Concealment*
   Michael Horowitz
   Polycom, Inc.
   December 2001
3. *Quality of Service in Frame-Switched Networks*
   White Paper
   Nortel Networks
   2000
4. *IP Quality of Service*
   Srinivas Vegesna
   Cisco Press
   2001
5. *Quality of Service for IP Videoconferencing*
   White Paper
   Subha Dhesikan, Software Engineer Enterprise
   Cisco, Inc.
   2001
6. *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (RFC 2474)*
   K. Nichols, S. Blake, F. Baker, D. Black
   December 1998
7. *An Architecture for Differentiated Services (RFC 2475)*
   S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss
   December 1998

Additional Information

1. *Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification [RFC 2205]*
   R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin
   September 1997
2. *An Expedited Forwarding PHB [RFC 2598]*
   V. Jacobson, K. Nichols, K. Poduri
   June 1999
3. *An Assured Forwarding PHB [RFC 2597]*
   J. Heinanen, F. Baker, W. Weiss, J. Wroclawski
   June 1999
4. *Internetworking Technologies Handbook,* Third Edition
   Cisco Press
   December 2000

Contact Information

Corporate Headquarters
Polycom, Inc.
1565 Barber Lane
Milpitas, CA 95035
USA

Phone: +1.408.526.9000
Fax: +1.408.526.9100

European Headquarters
Polyspan Ltd.
Whichford House
Parkway Court
Oxford Business Park South
Oxford 0X4 2JY
United Kingdom

Phone: +44 (0) 1865 335500
Fax: +44 (0) 1865 335501

Asia Pacific Headquarters
Polycom Solutions Pte. Ltd.
16 Raffles Quay
#40-02A, Hong Leong Building
Singapore 048581

Phone: +65.323.3882
Fax: +65.323.3022

For additional information, please visit the Polycom web site at: www.polycom.com