

USING EVENT SIMULATION TO EVALUATE INTERNET PROTOCOL ENHANCEMENTS FOR SPECIAL SERVICES

David A. Garbin

Noblis
3150 Fairview Park Dr South
Falls Church, VA 22042, U.S.A.

Patrick McGregor

Nyquetek, Inc
8397 Piping Rock Ct
Millersville, MD 21108, U.S.A.

Denise M. Bevilacqua Masi

Noblis
3150 Fairview Park Dr South
Falls Church, VA 22042, U.S.A.

ABSTRACT

Disasters can cause extraordinary service demand by the public, while concurrently causing outages that reduce network capacity to serve the surging demand. It is imperative that services supporting disaster response management perform with minimal degradation during such events. Mechanisms exist within Internet Protocol (IP)-based networks to provide preferential treatment for services such as voice and video using Differentiated Services Code Points (DSCP) in the packet headers and Per Hop Behaviors in the routers. However, there is currently no way to identify voice and video packets supporting response management and to ensure their timely delivery during network overload periods. We have applied simulation to evaluate the benefit of additional DSCP markings to be applied to such voice and video packets, and several router configurations. The results demonstrate significant value of the additions in preserving disaster response management performance even when aberration in demand causes ordinary voice and video performance to degrade.

1 INTRODUCTION

Networks based on Internet Protocol (IP) technology are now the primary means of providing data communications in the world. The commercial Internet has become an indispensable part of daily life. Traditionally, IP networks stressed the “best effort” delivery of packets of information, with small delays usually achieved by fast links and relatively rare packet losses corrected by higher layer protocols. Communication was accomplished by indiscriminately passing packets of information through a series of routers interconnected by transmission links; each router would immediately serve incoming packets if the outgoing link was not already busy transmitting a packet and queue packets for service if it was already busy. As IP networks became more ubiquitous and cost-efficient, more types of traffic were targeted for these networks. In particular, real-time traffic such as voice and video could be converted into packets and carried over IP networks. However, the

performance requirements of these services were significantly different from the data services. Timeliness of delivery (on the order of milliseconds) became a primary concern, with rare packet losses more tolerable. It became evident that the standard First Come First Served (FCFS) queuing discipline employed by the routers was not sufficient to guarantee quality for these services. Modifications were implemented in two areas to address this issue: packets were marked according to the type of traffic they were carrying and new queuing disciplines were implemented which gave the appropriate treatment to marked packets. The current state-of-the-art uses Differentiated Services Code Points (DSCP) in the IP header to mark packets according to their class of traffic. Queuing disciplines utilized in routers include Priority Queuing (PQ) which serves all higher priority packets before lower priority ones, Class-Based Weighted Fair Queuing (CBWFQ) which guarantees a minimum amount of bandwidth to traffic classes and shares the remaining bandwidth fairly based on user supplied weights, and Low Latency Queuing (LLQ) which combines PQ and CBWFQ in a single discipline. In an archetypical implementation, packets marked with an Expedited Forwarding (EF) DSCP are assigned to a single priority queue; other packets marked with Assured Forwarding (AF) and default (“best effort”) DSCP values share the remaining bandwidth using CBWFQ with appropriate weights. The EF marking is typically reserved for voice packets which are short in length and have stringent delay delivery requirements.

The public network infrastructure for telecommunications services is transforming into an IP based packet switched integrated services (voice, video, data) infrastructure. The transformation must ensure service continuity of services historically provided by the circuit switched Public Switched Telephone Network (PSTN). Of particular interest in this paper are the services supporting disaster response management.

In the PSTN experience, disasters have been documented to produce demand of up to 10X overload (i.e., the traffic offered to the network is 10 times the traffic that the network is engineered to handle) (Raushe et al. 1993). To

ensure effective disaster response management communications during such conditions, the PSTN has been enhanced with features to give such traffic priority treatment (McGregor et al. 1998, McGregor et al. 1999, Fonash and McGregor 2003). Disaster response management communications must also be ensured in the new IP based packet switched network. Previous work has demonstrated the feasibility of ensuring such service, but left open the identification of appropriate specific techniques (McGregor et al. 2006a). In this paper, specific alternative techniques at the packet level are investigated.

Under disaster conditions producing severe overloads, the EF queues in the network routers will be overwhelmed with packets unless strict Call Admission Control (CAC) is used to curtail the excess call attempts. In today's voice over IP (VoIP) implementations, CAC is either loosely applied or not applied at all. Disaster response management services can apply strict CAC as part of the authorization process; experience has shown that such traffic during an event is no more than 10 percent of the original engineered load (McGregor et al. 2006b). However, without special handling of the marked packets, during severe overload events they would suffer the same high packet loss rates and / or long delays as normal EF packets at the routers.

A request has been made to the Internet Assigned Numbers Authority (IANA) for two additional code points for assignment to voice and video traffic subject to strict CAC (Baker et al. 2007). These are referred to as VOICE-ADMIT and VIDEO-ADMIT. This paper describes the performance of differently marked voice and video traffic during severe overload conditions for several candidate router configurations that give the different markings different treatments. The evaluation metrics will be packet loss, delay, and delay variation (i.e., jitter, defined as the 99.9th percentile of delay). These metrics will be compared against established standards of performance for real-time services.

2 ANALYSIS FRAMEWORK

2.1 Traffic Stream Assumptions

A typical circuit in a converged services IP network will have at least four classes of traffic offered to it. In this study, we analyzed a typical access circuit with a bandwidth limit of 1460 kbps (essentially a T-1 link less its lower-level protocol overhead). The traffic on the circuit consisted of voice, video, mission-critical data traffic and best effort data traffic. The mission-critical data traffic included network and user signaling traffic.

Voice calls are modeled using G.711 codec parameters without silence suppression. A single voice call is modeled as a stream of constant size packets (200 bytes including IP, Real-Time Protocol [RTP], and User Datagram Proto-

col [UDP] overhead) arriving at a constant rate (one packet every 20 ms, or 80 kbps). When multiple voice calls are modeled, the packet arrival distribution is assumed to be Poisson.

All data traffic is modeled as stream of packets with a Poisson arrival distribution and a discrete packet size distribution of 50 percent at 40 bytes, 10 percent at 750 bytes, and 40 percent at 1500 bytes (Masi and Fischer 2005).

Video traffic requires a more detailed structure to model effectively, as discussed below.

2.2 Video Source Model

Video traffic does not lend itself to a simple formulation of interarrival and packet size distribution. Our representative stream is based on the H.264 codec; captured packet traces were analyzed at different bit rates and a scalable generic model was formulated. The parameters of this model follow.

- A constant video frame rate of 15 fps.
- A Gaussian frame size distribution with a variance to mean ratio of 310. The mean frame size depends on the bit rate being modeled. A mean frame size of 3300 bytes was used in this study, representing a 437 kbps bit rate.
- An embedded 29.1 kbps compressed voice stream interspersed among the video frames.

Figure 1 shows an actual frame size distribution of a 1610 kbps stream and the associated Gaussian fit.




Figure 1: Example of Gaussian fit to actual frame size data

In addition to the general parameters discussed above, it is important to model accurately the time distribution of the packets within a frame. Each composite frame consists of a video frame and three embedded voice packets (see Figure 2).

Each video frame is segmented into equal size packets with a maximum size of 1000 bytes. The resulting video packets are transmitted to the router as a "packet train"




Figure 2: Composite frame structure of a video signal

assuming a Fast Ethernet interface. Voice packets are transmitted at 22 ms intervals (i.e., three voice packets per frame) centered within the frame. If multiple video calls are being modeled, composite frames arrive with a Poisson distribution; each composite frame then generates packets according to the generic model.

2.3 Router Model

The generic router model shown in Figure 3 was used in these analyses. Input traffic is classified into traffic streams based on DSCP markings. The router uses a low latency queuing paradigm with three priority queues available. Each priority queue has a bandwidth limit enforced by a policing function, which ideally drops all packets exceeding the specified bit rate. The remaining three queues are rate-based queues implemented using a class based weighted fair queuing paradigm. The particular scheduling technique utilized is a virtual finish time algorithm described in the literature (Golestani 1994). Its characteristics are described in Fischer, Masi and Garbin (2007). The weights assigned to the traffic classes are designed to assure adequate performance for the video and mission-critical data traffic under normal loads. The various im-

plementation scenarios considered in this study consist of adding VOICE-ADMIT and VIDEO-ADMIT traffic streams to the router while increasing the existing traffic streams to simulate public network overload. The assignment of traffic streams to queues can be done in several different ways; the performance of each of these ways is analyzed in the following sections.

2.4 Simulation Tool Structure

A customized discrete event simulator was created for this study using the object-oriented RealBasic programming environment. Six traffic streams can be independently specified. The arrival distribution can be chosen to be Poisson, deterministic, or from a trace read from a file. The packet size distribution can be chosen as exponential, constant, Internet Mix (as described above), or video. A video distribution is a Gaussian frame distribution (the arrival distribution chosen then refers to frames, not packets). A mean packet size is also specified (mean frame size for video).

Six queues are specified by their queue depth (the maximum number of packets that can be held in the queue before packets are discarded). For the rate-based queues, weights are specified. Any of the traffic streams can be assigned to any of the queues.

Figure 4 illustrates the simulator input screen. The output of the simulator provides statistics both at the queue level and at the stream level. Offered load, throughput, packet loss, delay, and jitter values are provided for all entities.

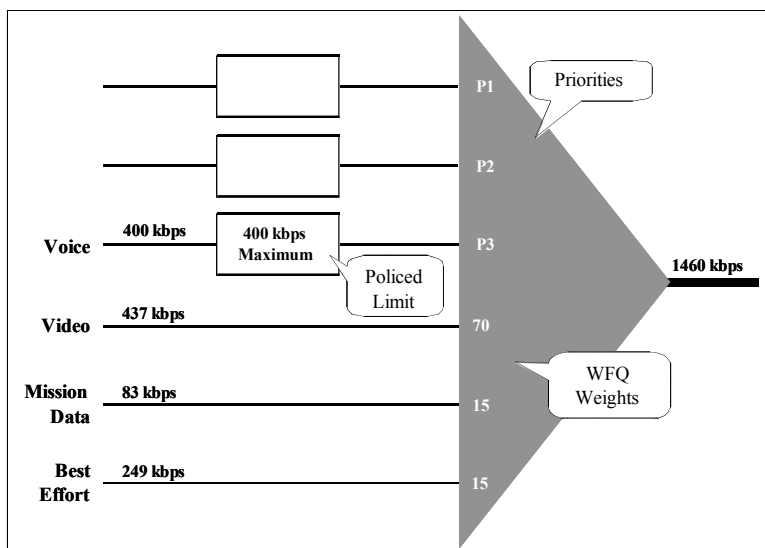


Figure 3: Generic router model for converged network services

QSim
File Edit

Packet Event Simulator

Title: T1_Baseline

Max Pkts: 1000000
Linespeed (kbps): 1460

Debug: n
Trace: n

| Stream 1 | Stream 2 | Stream 3 | Stream 4 | Stream 5 | Stream 6 |
|---------------------------|----------|----------|----------|----------|----------|
| Seed: 694567 | 345678 | 245789 | 903824 | 523489 | 892341 |
| Arrival Distr: m | m | m | d | m | m |
| Pkts per Sec: 0 | 0 | 250 | 15 | 15 | 45 |
| Pkt Size Distr: d | d | d | v | i | i |
| Ave Pkt Size (bytes): 200 | 200 | 200 | 3300 | 0 | 0 |
| Queue Assignment: 1 | 2 | 3 | 4 | 5 | 6 |

| Queue 1 | Queue 2 | Queue 3 | Queue 4 | Queue 5 | Queue 6 |
|--------------------|------------|------------|---------|---------|---------|
| Queue Size: 20 | 20 | 20 | 20 | 512 | 512 |
| Weight: Priority 1 | Priority 2 | Priority 3 | 70 | 15 | 15 |

Run

Figure 4: Custom simulator input screen

3 SCENARIO DESCRIPTIONS

3.1 Baseline Traffic Assumptions

The baseline traffic assumptions follow general engineering guidelines for an access link of this speed carrying multimedia traffic. Real-time traffic (voice and video) are engineered to consume no more than half of the link bandwidth. The mission-critical traffic is small, but must be protected from packet loss. The signaling traffic carried in this class is essential for operation of the network. Five voice calls and a single video call are supported along with 83 kbps of mission-critical data. The remaining data is classified as best effort; the baseline rate of 249 kbps for this traffic combined with the other traffic results in an 80 percent link utilization. This is a typical engineered utilization for access links in the busy hour.

3.2 Overload Traffic Assumptions

The environment modeled in the alternative configurations reflects the load offered to the access link under disaster response conditions. The public voice, video, and best effort data traffic is increased by a factor of ten. This assumes that an effective call admission control policy for the public voice and video traffic is not available (this is the worst case scenario postulated for this study). The mission-critical data and signaling traffic is under an organiza-

tion's control and is assumed to remain at the baseline level. One VOICE-ADMIT call and one VIDEO-ADMIT call are added to the traffic. The performance of these two calls under different router configurations is the main focus of the study.

3.3 Alternative Router Configurations

Seven alternative router configurations were considered for handling the disaster response management traffic. They differ in the queues assigned for VOICE-ADMIT, VIDEO-ADMIT, and public Voice traffic. One, two, or three priority queues are utilized in the configurations. In one configuration, a single rate-based queue is used for all video with preferential treatment of VIDEO-ADMIT packets given through an early discard mechanism. The alternatives are described in the following sections. Each alternative is given a descriptor depending on how many EF code points are defined and how many priority queues are utilized (e.g., 3EF-2Q). Appendix A contains diagrams of the alternatives.

3.3.1 Single Priority Queue Case

This case defines three EF code points, one each for VOICE-ADMIT, VIDEO-ADMIT, and public Voice. Each code point is policed separately before being assigned to the same priority queue. Since the ADMIT traffic classes

are subject to admission control, no packets are lost due to the policing function. However, the policing function for the public Voice traffic restricts the offered packet stream to the queue to 400 kbps (five voice calls). The excess packets resulting from the 10X overload are dropped by the policer, rendering all conversations unintelligible. The Voice packets that are offered to the queue compete with VOICE-ADMIT and VIDEO-ADMIT packets for resources. Even though this queue has strict priority over the rate queues, packets in service from the rate queues can still affect performance (there is no preemption of packets in service in this model). This case has the descriptor 3EF-1Q.

3.3.2 Two Priority Queue Cases

These cases define the same three EF code points, but they utilize two priority queues. Public Voice is again policed at the 400 kbps level. The cases differ in how they combine code points for queue assignment and which code points are given priority over others. The cases considered follow.

- 3EF-2Qa: All ADMIT traffic is put into a single priority queue with preference over the public Voice priority queue. The rationale for this alternative is that the ADMIT traffic has a mission precedence over the public traffic.
- 3EF-2Qb: VOICE-ADMIT and public VOICE are put into a single priority queue with preference over the VIDEO-ADMIT priority queue. This alternative recognizes the necessity of protecting the short voice packets from the jitter affects of the long video packets.
- 3EF-2Qc: VOICE-ADMIT and public VOICE are put into a single priority queue, but VIDEO-ADMIT is given preference over the combined voice queue. This alternative protects the critical video traffic from the adverse affects of the public Voice traffic.

3.3.3 Three Priority Queue Cases

These alternatives define three EF code points and assign each code point to its own priority queue. The alternatives differ in the priority order given to the traffic streams. In both cases, VOICE-ADMIT traffic is given the highest priority. The cases considered follow.

- 3EF-3Qa: VIDEO-ADMIT is given precedence over public Voice. This recognizes the mission-importance of ADMIT traffic.
- 3EF-3Qb: Public Voice is given precedence over VIDEO-ADMIT. This protects the short voice packets from the long video packets.

3.3.4 One Priority Queue Case with Rate-Based VIDEO-ADMIT Control

This alternative only defines two EF code points for the two types of voice traffic. The VIDEO-ADMIT traffic is assigned a special AF code point and is assigned to the same rate-based queue as the public video traffic. There is an existing mechanism in routers to selectively discard packets before they are offered to a queue. The discard algorithm is based on the number of packets in the queue and on the drop priority of the packets. The drop priority can be based on the DSCP of the packets. While this mechanism, known as Weighted Random Early Detection (WRED), was originally intended to gradually cause Transmission Control Protocol (TCP) data sessions to adapt their input rate to network conditions, it is possible to adjust the parameters of this mechanism to ruthlessly drop lower precedence video packets in favor of higher precedence packets. We refer to this technique as Differential Early Discard (DED). For the purpose of this study, we assume ideal operation of this mechanism (i.e., under overload of the video queue, all public video packets are dropped and all VIDEO-ADMIT packets are passed). The dynamic nature of the WRED mechanism and its affect on this assumption will be the subject of a future study. In this configuration, called 2EF-1Q+DED, the VOICE-ADMIT traffic is assigned to the same priority queue as the policed public Voice traffic.

4 RESULTS

The results are analyzed against three measures of performance: packet loss, average queue delay and jitter (variation of delay). Results are displayed for the two types of ADMIT traffic and also for the public Voice traffic.

For real-time traffic, packet loss is one of the primary measures of performance. Unlike data traffic, real-time voice and video traffic is not protected against packet loss by an end-to-end transport protocol. Voice and video use UDP as their transport protocol; UDP does not detect or retransmit lost packets. ITU-T Standard Y.1541 specifies end-to-end targets for impairments. We have allocated these targets to network segments and derived targets specific to the access link in our example. The specification for packet loss is 0.05 percent (99.95 percent of the packets are transmitted).

The second metric of interest is average delay. Absolute delay in itself does not cause any degradation in the voice or video signal, but it does cause a loss of interactivity when conversations are taking place. Beyond a critical value, a person will perceive that the distant party is not responding and will resume talking; the result is a jumble of mixed messages. The major components of delay are the propagation delay of the signal (about 25 ms coast-to-coast in the U.S.) and the codec delay (about 60 ms in voice over

IP systems). The additional delay reported in this study is the average time spent in queue at the router under study. Allocation from Y.1541 gives us targets of 10 ms for voice and 20 ms for video for average queue delay.

Perhaps the most important metric for real-time service performance is jitter. The packets of a voice or video stream must be offered to the receiving decoder at a constant rate. If there is variation in packet delay along a path (due to varying queue depth at the routers along the path), a buffer must be provided in the receiving terminal to compensate for this variation. The larger the buffer, the more variation can be tolerated. The penalty for a large buffer is an increase in absolute delay. Most devices have a maximum value of jitter that can be tolerated; packets arriving outside this window are discarded. Managing network queues to ensure that jitter is within tolerance is the key design issue in providing real-time services. In this study, the jitter performance of the alternative configurations is the discriminator among them. The allocated targets for jitter on this access link are 10 ms for voice and 60 ms for video.

4.1 Packet Loss Results

Figure 5 presents the packets transmitted results for all seven configurations. All of the packet loss in these examples is due to the policing of public Voice traffic from its 10X overload level to its policed limit of 400 kbps. This results in only 10% of the public Voice packets being transmitted. The result is that all ADMIT traffic packets suffer no loss in their respective queues for any of the configurations.

4.2 Average Queue Delay Results

Figure 6 presents the results for average queue delay for all seven configurations. The average queue delay for VOICE-ADMIT traffic suffers if it is in the same queue as VIDEO-ADMIT (configurations 3EF-1Q and 3EF-2Qa) or if video is given priority (configuration 3EF-2Qc). For VIDEO-ADMIT traffic, the average queue delay suffers when public Voice is given priority (configurations 3EF-2Qb, 3EF-3Qb, 2EF-1Q+DED). Nevertheless, the average queue delay for all traffic types is within target objectives for all configurations.

4.3 Jitter Results

Figure 7 presents the results for jitter for all configurations. These results provide a clear delineation among the alternatives. Jitter for VOICE-ADMIT is unacceptable (greater than 10 ms) if it is in the same queue as VIDEO-ADMIT (3EF-1Q, 3EF-2Qa) or if video is given priority (3EF-2Qc). Of the remaining alternatives, configuration 3EF-3Qa gives the best performance for both VOICE-ADMIT and VIDEO-ADMIT traffic (at the expense of public Voice, which is likely to be unviable anyway due to excessive packet loss). The other three configurations (3EF-2Qb, 3EF-3Qb, and 2EF-1Q+DED) provide good public Voice performance and marginal, but acceptable performance for VIDEO-ADMIT. The 2EF-1Q+DED approach has the advantage of being implementable in most of today’s routers. However, further study is needed into the dynamics of WRED in practice (i.e., how the different parameter settings used to achieve DED operation affect the transient behavior of both video streams).

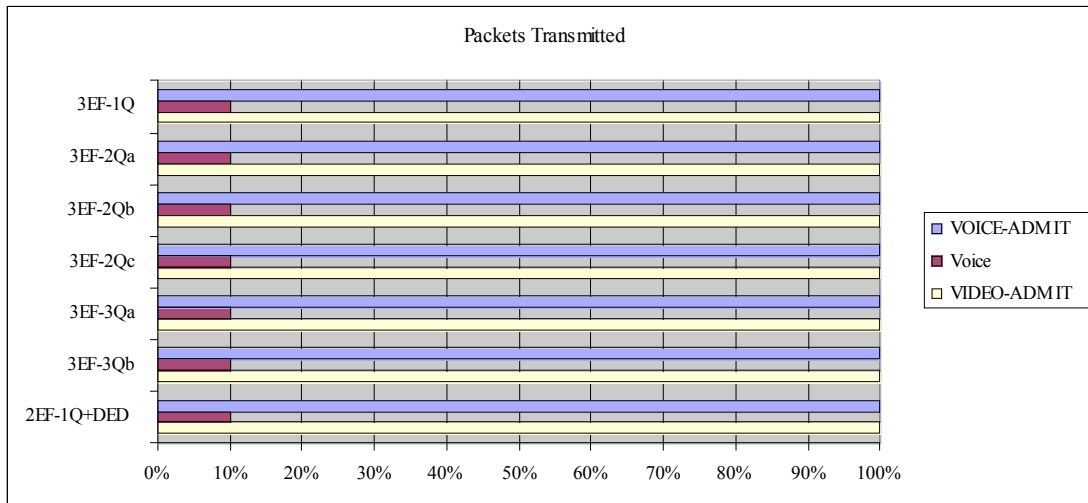


Figure 5: Packet loss results

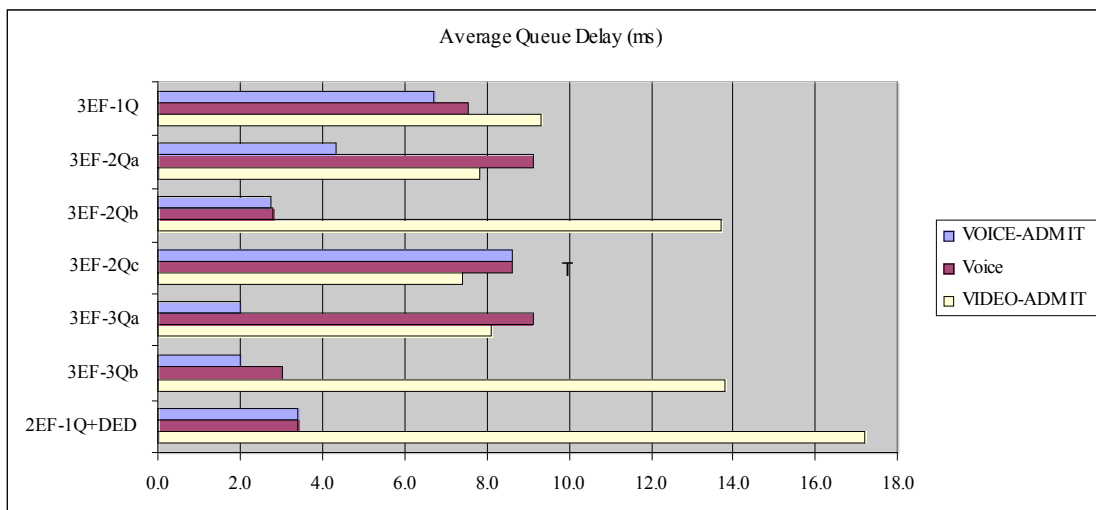


Figure 6: Average queue delay results




Figure 7: Jitter results

5 SUMMARY AND NEXT STEPS

Two additional DiffServ Code Points have been proposed to mark special service voice and video packets for use in disaster response situations. This study has shown that alternative queue configurations exist whereby IP routers can ensure the performance of these packets over typical access link bandwidths in the face of extreme public traffic overloads. When both voice and video disaster response management packets are present, the best performance is achieved when the traffic types are assigned to separate queues. The controlling factor in determining performance is the variation in queuing delay (jitter). In performing this study, a generic model of video traffic was developed which can serve as a flexible tool for future studies.

A technique was proposed to utilize the well-known Weighted Random Early Detection mechanism with extreme discard parameters and apply it to real-time, UDP video traffic. Preliminary results indicate this would be a useful mechanism for differentiating video streams assigned to the same queue. More work is needed to study the detailed transient behavior of the WRED algorithm as a function of its discard parameters before this technique can be recommended.

ACKNOWLEDGMENTS

This work was partially funded by the National Communications System Contract Number NBCH-D-02-0039 (Task Order Number D0200390095). Research

funding was also provided by the Center for Network-Based Systems, a research collaboration of Noblis and George Mason University. We would like to thank Dr. Martin J. Fischer of Noblis who provided valuable analytic support to this project.

APPENDIX A—ROUTER CONFIGURATION DIAGRAMS




Diagram A-1: Configuration 3EF-1Q

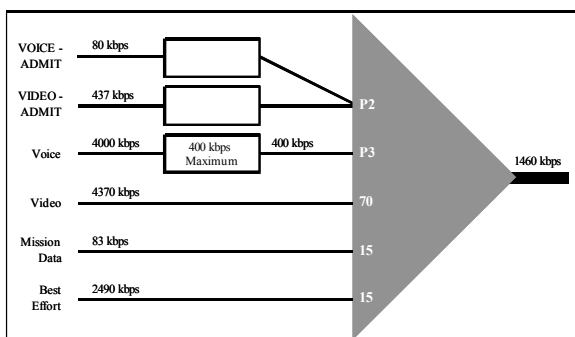


Diagram A-2: Configuration 3EF-2Qa

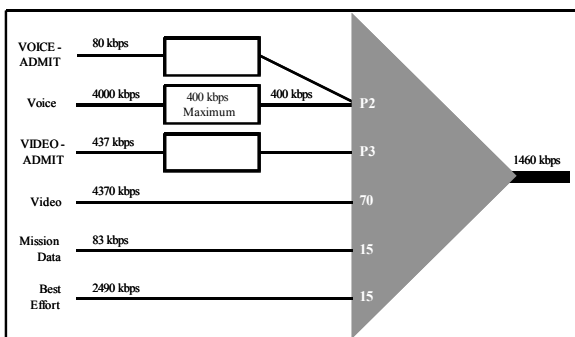


Diagram A-3: Configuration 3EF-2Qb




Diagram A-4: Configuration 3EF-2Qc

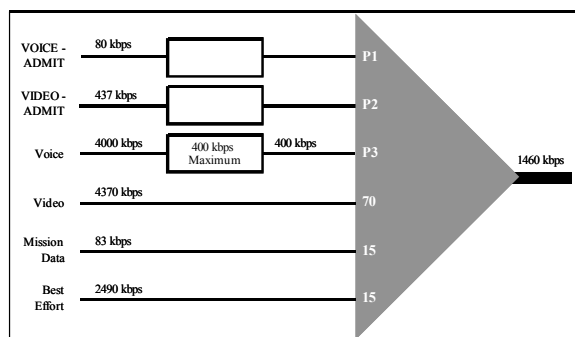


Diagram A-5: Configuration 3EF-3Qa

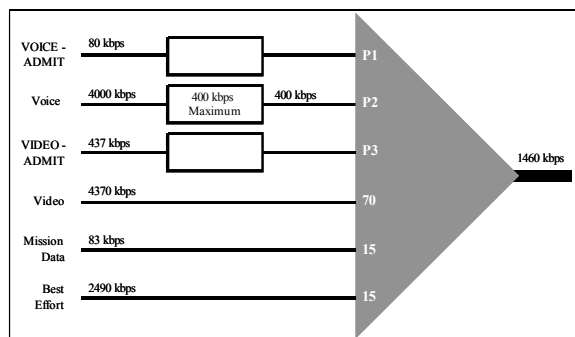


Diagram A-6: Configuration 3EF-3Qb

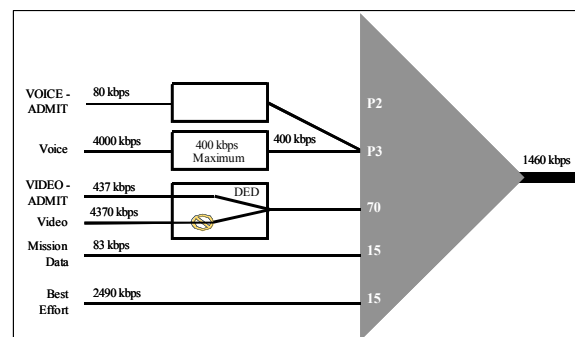


Diagram A-7: Configuration 2EF-1Q+DED

REFERENCES

- Baker, F., J. Polk, and M. Dolly, "DSCPs for Capacity-Admitted Traffic." March 2007. Draft-ietf-tsvwg-admitted-realtime-dscp-01 (work in progress), <<http://www.ietf.org/internet-drafts/draft-ietf-tsvwg-admitted-realtime-dscp-01.txt>>.
- Fonash, P. and P. McGregor, "National Security/Emergency Preparedness Wireless Priority Service," *Proceedings of the 8th International Conference on Intelligence in Next Generation Networks (ICIN 2003)*, Bordeaux, France, March 31–April 3, 2003.
- Golestani, S. J. 1994. "A Self-Clocked Fair Queuing Scheme for Broadband Applications," *Proceedings of the IEEE INFOCOM*.
- Masi, D. M. B. and M. J. Fischer, "Voice over Internet Protocol (VoIP) Performance Models—A Comprehensive Approach," *International Conference on Telecommunication Systems—Modeling and Analysis (ICTSMA)*, Dallas, TX, November 17–20, 2005.
- Masi, D. M. B., M. J. Fischer, and D. A. Garbin, "Modeling the Performance of Low Latency Queuing for Emergency Telecommunications," *Winter Simulation Conference*, Washington, DC, December 9–12, 2007.
- McGregor, P., J. Szeto, and F. Suraci, "Performance Modeling of High Probability Call-Completion Features," *Proceedings of the Seventh International Conference on Computer and Communications Networks*, October 12-15, 1998.
- McGregor, P., J. Graves, and T. Otto. 1999. "Government Emergency Telecommunications Service and the Evolving Public Switched Network," *The 1999 Annual Review of Communications of the International Engineering Consortium*.
- McGregor, P., R. Craighill, and V. Mosely, "Government Emergency Telecommunications Service (GETS) and Wireless Priority Service (WPS) Performance During Katrina," *Proceeding (533) Communications, Internet, and Information Technology-2006, IASTED*, November 29-December 1, 2006.
- McGregor, P., R. Kaczmarek, V. Mosely, D. Dease, and P. Adams, "National Security / Emergency Preparedness and the Next Generation Network", *IEEE Communications Magazine*, 44(5), May 2006, 133-143.
- Raushe, A., J. Coe, J. Gerencser, L. Patterson, and D. Waters, "Traffic Congestion Analysis of the Public Switched Network," *IEEE MILCOM Conference Record*, pp. 917-921, October 3, 1993.

AUTHOR BIOGRAPHIES

DAVID A GARBIN is a senior fellow at Noblis. His 30+ years of experience is in the telecommunications and networking field, focusing on the design and economic analysis of large networks, both for carriers and their customers. His current duties include research into providing quality of service in convergent IP networks and advising government agencies in the acquisition and implementation of VoIP. He holds advanced degrees from MIT and is the co-author of the best-selling New McGraw-Hill *Telecom Factbook*, currently published in its second edition. His email address is <david.garbin@noblis.org>

PATRICK MCGREGOR received a BSEE degree from Purdue University and MSEE and PhD degrees in electrical engineering from Polytechnic Institute, Brooklyn. He is currently president of Nyquetek, Inc., a network modeling and analysis company under sub contract to Computer Sciences Corporation, which in turn is under contract to the National Communications System. Previously, he has held various officer positions for small companies and subsidiaries of large companies. Email: <pat_mcgregor@msn.com>

DENISE M. BEVILACQUA MASI is a senior principal engineer at Noblis. Her experience and research interests include queueing theory and simulation applied to telecommunications networks. She received her doctorate degree in information technology and engineering at George Mason University. Her email address is <dmasi@noblis.org>