

Network Working Group
Request for Comments: 4594
Category: Informational

J. Babiarz
K. Chan
Nortel Networks
F. Baker
Cisco Systems
August 2006

Configuration Guidelines for DiffServ Service Classes

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes service classes configured with Diffserv and recommends how they can be used and how to construct them using Differentiated Services Code Points (DSCPs), traffic conditioners, Per-Hop Behaviors (PHBs), and Active Queue Management (AQM) mechanisms. There is no intrinsic requirement that particular DSCPs, traffic conditioners, PHBs, and AQM be used for a certain service class, but as a policy and for interoperability it is useful to apply them consistently.

Table of Contents

1. Introduction	3
1.1. Requirements Notation	4
1.2. Expected Use in the Network	4
1.3. Service Class Definition	5
1.4. Key Differentiated Services Concepts	5
1.4.1. Queuing	6
1.4.1.1. Priority Queuing	6
1.4.1.2. Rate Queuing	6
1.4.2. Active Queue Management	7
1.4.3. Traffic Conditioning	7
1.4.4. Differentiated Services Code Point (DSCP)	8
1.4.5. Per-Hop Behavior (PHB)	8
1.5. Key Service Concepts	8
1.5.1. Default Forwarding (DF)	9
1.5.2. Assured Forwarding (AF)	9
1.5.3. Expedited Forwarding (EF)	10
1.5.4. Class Selector (CS)	10
1.5.5. Admission Control	11
2. Service Differentiation	11
2.1. Service Classes	12
2.2. Categorization of User Service Classes	13
2.3. Service Class Characteristics	16
2.4. Deployment Scenarios	21
2.4.1. Example 1	21
2.4.2. Example 2	23
2.4.3. Example 3	25
3. Network Control Traffic	27
3.1. Current Practice in the Internet	27
3.2. Network Control Service Class	27
3.3. OAM Service Class	29
4. User Traffic	30
4.1. Telephony Service Class	31
4.2. Signaling Service Class	33
4.3. Multimedia Conferencing Service Class	35
4.4. Real-Time Interactive Service Class	37
4.5. Multimedia Streaming Service Class	39
4.6. Broadcast Video Service Class	41
4.7. Low-Latency Data Service Class	43
4.8. High-Throughput Data Service Class	45
4.9. Standard Service Class	47
4.10. Low-Priority Data	48
5. Additional Information on Service Class Usage	49
5.1. Mapping for Signaling	49
5.2. Mapping for NTP	50
5.3. VPN Service Mapping	50
6. Security Considerations	51

7. Acknowledgements	52
8. Appendix A	53
8.1. Explanation of Ring Clipping	53
9. References	54
9.1. Normative References	54
9.2. Informative References	55

1. Introduction

To aid in understanding the role of this document, we use an analogy: the Differentiated Services specifications are fundamentally a toolkit. The specifications provide the equivalent of band saws, planers, drill presses, and other tools. In the hands of an expert, there is no limit to what can be built, but such a toolkit can be intimidating to the point of being inaccessible to a non-expert who just wants to build a bookcase. This document should be viewed as a set of "project plans" for building all the (diffserv) furniture that one might want. The user may choose what to build (e.g., perhaps our non-expert doesn't need a china cabinet right now), and how to go about building it (e.g., plans for a non-expert probably won't employ mortise/tenon construction, but that absence does not imply that mortise/tenon construction is forbidden or unsound). The authors hope that these diffserv "project plans" will provide a useful guide to Network Administrators in the use of diffserv techniques to implement quality-of-service measures appropriate for their network's traffic.

This document describes service classes configured with Diffserv and recommends how they can be used and how to construct them using Differentiated Services Code Points (DSCPs), traffic conditioners, Per-Hop Behaviors (PHBs), and Active Queue Management (AQM) mechanisms. There is no intrinsic requirement that particular DSCPs, traffic conditioners, PHBs, and AQM be used for a certain service class, but as a policy and for interoperability it is useful to apply them consistently.

Service class definitions are based on the different traffic characteristics and required performance of the applications/services. This approach allows us to map current and future applications/services of similar traffic characteristics and performance requirements into the same service class. Since the applications'/services' characteristics and required performance are end to end, the service class notion needs to be preserved end to end. With this approach, a limited set of service classes is required. For completeness, we have defined twelve different service classes, two for network operation/administration and ten for user/subscriber applications/services. However, we expect that network administrators will implement a subset of these classes

relevant to their customers and their service offerings. Network Administrators may also find it of value to add locally defined service classes, although these will not necessarily enjoy end-to-end properties of the same type.

Section 1 provides an introduction and overview of technologies that are used for service differentiation in IP networks. Section 2 is an overview of how service classes are constructed to provide service differentiation, with examples of deployment scenarios. Section 3 provides configuration guidelines of service classes that are used for stable operation and administration of the network. Section 4 provides configuration guidelines of service classes that are used for differentiation of user/subscriber traffic. Section 5 provides additional guidance on mapping different applications/protocols to service classes. Section 6 addresses security considerations.

1.1. Requirements Notation

The key words "SHOULD", "SHOULD NOT", "REQUIRED", "SHALL", "SHALL NOT", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

1.2. Expected Use in the Network

In the Internet today, corporate LANs and ISP WANs are generally not heavily utilized. They are commonly 10% utilized at most. For this reason, congestion, loss, and variation in delay within corporate LANs and ISP backbones is virtually unknown. This clashes with user perceptions, for three very good reasons.

- o The industry moves through cycles of bandwidth boom and bandwidth bust, depending on prevailing market conditions and the periodic deployment of new bandwidth-hungry applications.
- o In access networks, the state is often different. This may be because throughput rates are artificially limited or over-subscribed, or because of access network design trade-offs.
- o Other characteristics, such as database design on web servers (that may create contention points, e.g., in filestore) and configuration of firewalls and routers, often look externally like a bandwidth limitation.

The intent of this document is to provide a consistent marking, conditioning, and packet treatment strategy so that it can be configured and put into service on any link that is itself congested.

1.3. Service Class Definition

A "service class" represents a set of traffic that requires specific delay, loss, and jitter characteristics from the network. Conceptually, a service class pertains to applications with similar characteristics and performance requirements, such as a "High-Throughput Data" service class for applications like the web and electronic mail, or a "Telephony" service class for real-time traffic such as voice and other telephony services. Such a service class may be defined locally in a Differentiated Services (DS) domain, or across multiple DS domains, possibly extending end to end.

A service class as defined here is essentially a statement of the required characteristics of a traffic aggregate. The required characteristics of these traffic aggregates can be realized by the use of defined per-hop behavior (PHB) [RFC2474]. The actual specification of the expected treatment of a traffic aggregate within a domain may also be defined as a per-domain behavior (PDB) [RFC3086].

Each domain may choose to implement different service classes or to use different behaviors to implement the service classes or to aggregate different kinds of traffic into the aggregates and still achieve their required characteristics. For example, low delay, loss, and jitter may be realized using the EF PHB, or with an over-provisioned AF PHB. This must be done with care as it may disrupt the end-to-end performance required by the applications/services. This document provides recommendations on usage of PHBs for specific service classes for their consistent implementation. These recommendations are not to be construed as prohibiting use of other PHBs that realize behaviors sufficient for the relevant class of traffic.

The Default Forwarding "Standard" service class is REQUIRED; all other service classes are OPTIONAL. It is expected that network administrators will base their choice of the level of service differentiation that they will support on their need, starting off with three or four service classes for user traffic and adding others as the need arises.

1.4. Key Differentiated Services Concepts

The reader SHOULD be familiar with the principles of the Differentiated Services Architecture [RFC2474]. We recapitulate key concepts here only to provide convenience for the reader, the referenced RFCs providing the authoritative definitions.

1.4.1. Queuing

A queue is a data structure that holds packets that are awaiting transmission. The packets may be delayed while in the queue, possibly due to lack of bandwidth, or because it is low in priority. There are a number of ways to implement a queue. A simple model of a queuing system, however, is a set of data structures for packet data, which we will call queues, and a mechanism for selecting the next packet from among them, which we call a scheduler.

1.4.1.1. Priority Queuing

A priority queuing system is a combination of a set of queues and a scheduler that empties them in priority sequence. When asked for a packet, the scheduler inspects the highest priority queue and, if there is data present, returns a packet from that queue. Failing that, it inspects the next highest priority queue, and so on. A freeway onramp with a stoplight for one lane that allows vehicles in the high-occupancy-vehicle lane to pass is an example of a priority queuing system; the high-occupancy-vehicle lane represents the "queue" having priority.

In a priority queuing system, a packet in the highest priority queue will experience a readily calculated delay. This is proportional to the amount of data remaining to be serialized when the packet arrived plus the volume of the data already queued ahead of it in the same queue. The technical reason for using a priority queue relates exactly to this fact: it limits delay and variations in delay and should be used for traffic that has that requirement.

A priority queue or queuing system needs to avoid starvation of lower-priority queues. This may be achieved through a variety of means, such as admission control, rate control, or network engineering.

1.4.1.2. Rate Queuing

Similarly, a rate-based queuing system is a combination of a set of queues and a scheduler that empties each at a specified rate. An example of a rate-based queuing system is a road intersection with a stoplight. The stoplight acts as a scheduler, giving each lane a certain opportunity to pass traffic through the intersection.

In a rate-based queuing system, such as Weighted Fair Queuing (WFQ) or Weighted Round Robin (WRR), the delay that a packet in any given queue will experience depends on the parameters and occupancy of its queue and the parameters and occupancy of the queues it is competing with. A queue whose traffic arrival rate is much less than the rate

at which it lets traffic depart will tend to be empty, and packets in it will experience nominal delays. A queue whose traffic arrival rate approximates or exceeds its departure rate will tend not to be empty, and packets in it will experience greater delay. Such a scheduler can impose a minimum rate, a maximum rate, or both, on any queue it touches.

1.4.2. Active Queue Management

Active Queue Management, or AQM, is a generic name for any of a variety of procedures that use packet dropping or marking to manage the depth of a queue. The canonical example of such a procedure is Random Early Detection (RED), in that a queue is assigned a minimum and maximum threshold, and the queuing algorithm maintains a moving average of the queue depth. While the mean queue depth exceeds the maximum threshold, all arriving traffic is dropped. While the mean queue depth exceeds the minimum threshold but not the maximum threshold, a randomly selected subset of arriving traffic is marked or dropped. This marking or dropping of traffic is intended to communicate with the sending system, causing its congestion avoidance algorithms to kick in. As a result of this behavior, it is reasonable to expect that TCP's cyclic behavior is desynchronized and that the mean queue depth (and therefore delay) should normally approximate the minimum threshold.

A variation of the algorithm is applied in Assured Forwarding PHB [RFC2597], in that the behavior aggregate consists of traffic with multiple DSCP marks, which are intermingled in a common queue. Different minima and maxima are configured for the several DSCPs separately, such that traffic that exceeds a stated rate at ingress is more likely to be dropped or marked than traffic that is within its contracted rate.

1.4.3. Traffic Conditioning

In addition, at the first router in a network that a packet crosses, arriving traffic may be measured and dropped or marked according to a policy, or perhaps shaped on network ingress, as in "A Rate Adaptive Shaper for Differentiated Services" [RFC2963]. This may be used to bias feedback loops, as is done in "Assured Forwarding PHB" [RFC2597], or to limit the amount of traffic in a system, as is done in "Expedited Forwarding PHB" [RFC3246]. Such measurement procedures are collectively referred to as "traffic conditioners". Traffic conditioners are normally built using token bucket meters, for example with a committed rate and burst size, as in Section 1.5.3 of the DiffServ Model [RFC3290]. The Assured Forwarding PHB [RFC2597] uses a variation on a meter with multiple rate and burst size measurements to test and identify multiple levels of conformance.

Multiple rates and burst sizes can be realized using multiple levels of token buckets or more complex token buckets; these are implementation details. The following are some traffic conditioners that may be used in deployment of differentiated services:

- o For Class Selector (CS) PHBs, a single token bucket meter to provide a rate plus burst size control.
- o For Expedited Forwarding (EF) PHB, a single token bucket meter to provide a rate plus burst size control.
- o For Assured Forwarding (AF) PHBs, usually two token bucket meters configured to provide behavior as outlined in "Two Rate Three Color Marker (trTCM)" [RFC2698] or "Single Rate Three Color Marker (srTCM)" [RFC2697]. The two-rate, three-color marker is used to enforce two rates, whereas the single-rate, three-color marker is used to enforce a committed rate with two burst lengths.

1.4.4. Differentiated Services Code Point (DSCP)

The DSCP is a number in the range 0..63 that is placed into an IP packet to mark it according to the class of traffic it belongs in. Half of these values are earmarked for standardized services, and the other half of them are available for local definition.

1.4.5. Per-Hop Behavior (PHB)

In the end, the mechanisms described above are combined to form a specified set of characteristics for handling different kinds of traffic, depending on the needs of the application. This document seeks to identify useful traffic aggregates and to specify what PHB should be applied to them.

1.5. Key Service Concepts

While Differentiated Services is a general architecture that may be used to implement a variety of services, three fundamental forwarding behaviors have been defined and characterized for general use. These are basic Default Forwarding (DF) behavior for elastic traffic, the Assured Forwarding (AF) behavior, and the Expedited Forwarding (EF) behavior for real-time (inelastic) traffic. The facts that four code points are recommended for AF and that one code point is recommended for EF are arbitrary choices, and the architecture allows any reasonable number of AF and EF classes simultaneously. The choice of four AF classes and one EF class in the current document is also arbitrary, and operators MAY choose to operate more or fewer of either.

The terms "elastic" and "real-time" are defined in [RFC1633], Section 3.1, as a way of understanding broad-brush application requirements. This document should be reviewed to obtain a broad understanding of the issues in quality of service, just as [RFC2475] should be reviewed to understand the data plane architecture used in today's Internet.

1.5.1. Default Forwarding (DF)

The basic forwarding behaviors applied to any class of traffic are those described in [RFC2474] and [RFC2309]. Best-effort service may be summarized as "I will accept your packets" and is typically configured with some bandwidth guarantee. Packets in transit may be lost, reordered, duplicated, or delayed at random. Generally, networks are engineered to limit this behavior, but changing traffic loads can push any network into such a state.

Application traffic in the internet that uses default forwarding is expected to be "elastic" in nature. By this, we mean that the sender of traffic will adjust its transmission rate in response to changes in available rate, loss, or delay.

For the basic best-effort service, a single DSCP value is provided to identify the traffic, a queue to store it, and active queue management to protect the network from it and to limit delays.

1.5.2. Assured Forwarding (AF)

The Assured Forwarding PHB [RFC2597] behavior is explicitly modeled on Frame Relay's Discard Eligible (DE) flag or ATM's Cell Loss Priority (CLP) capability. It is intended for networks that offer average-rate Service Level Agreements (SLAs) (as FR and ATM networks do). This is an enhanced best-effort service; traffic is expected to be "elastic" in nature. The receiver will detect loss or variation in delay in the network and provide feedback such that the sender adjusts its transmission rate to approximate available capacity.

For such behaviors, multiple DSCP values are provided (two or three, perhaps more using local values) to identify the traffic, a common queue to store the aggregate, and active queue management to protect the network from it and to limit delays. Traffic is metered as it enters the network, and traffic is variously marked depending on the arrival rate of the aggregate. The premise is that it is normal for users occasionally to use more capacity than their contract stipulates, perhaps up to some bound. However, if traffic should be marked or lost to manage the queue, this excess traffic will be marked or lost first.

1.5.3. Expedited Forwarding (EF)

The intent of Expedited Forwarding PHB [RFC3246] is to provide a building block for low-loss, low-delay, and low-jitter services. It can be used to build an enhanced best-effort service: traffic remains subject to loss due to line errors and reordering during routing changes. However, using queuing techniques, the probability of delay or variation in delay is minimized. For this reason, it is generally used to carry voice and for transport of data information that requires "wire like" behavior through the IP network. Voice is an inelastic "real-time" application that sends packets at the rate the codec produces them, regardless of availability of capacity. As such, this service has the potential to disrupt or congest a network if not controlled. It also has the potential for abuse.

To protect the network, at minimum one SHOULD police traffic at various points to ensure that the design of a queue is not overrun, and then the traffic SHOULD be given a low-delay queue (often using priority, although it is asserted that a rate-based queue can do this) to ensure that variation in delay is not an issue, to meet application needs.

1.5.4. Class Selector (CS)

Class Selector provides support for historical codepoint definitions and PHB requirement. The Class Selector DS field provides a limited backward compatibility with legacy (pre DiffServ) practice, as described in [RFC2474], Section 4. Backward compatibility is addressed in two ways. First, there are per-hop behaviors that are already in widespread use (e.g., those satisfying the IPv4 Precedence queuing requirements specified in [RFC1812]), and we wish to permit their continued use in DS-compliant networks. In addition, there are some codepoints that correspond to historical use of the IP Precedence field, and we reserve these codepoints to map to PHBs that meet the general requirements specified in [RFC2474], Section 4.2.2.2.

No attempt is made to maintain backward compatibility with the "DTR" or Type of Service (TOS) bits of the IPv4 TOS octet, as defined in [RFC0791] and [RFC1349].

A DS-compliant network can be deployed with a set of one or more Class Selector-compliant PHB groups. Also, a network administrator may configure the network nodes to map codepoints to PHBs, irrespective of bits 3-5 of the DSCP field, to yield a network that is compatible with historical IP Precedence use. Thus, for example, codepoint '011000' would map to the same PHB as codepoint '011010'.

1.5.5. Admission Control

Admission control (including refusal when policy thresholds are crossed) can ensure high-quality communication by ensuring the availability of bandwidth to carry a load. Inelastic real-time flows such as Voice over Internet Protocol (VoIP) (telephony) or video conferencing services can benefit from use of an admission control mechanism, as generally the telephony service is configured with over-subscription, meaning that some users may not be able to make a call during peak periods.

For VoIP (telephony) service, a common approach is to use signaling protocols such as SIP, H.323, H.248, MEGACO, and Resource Reservation Protocol (RSVP) to negotiate admittance and use of network transport capabilities. When a user has been authorized to send voice traffic, this admission procedure has verified that data rates will be within the capacity of the network that it will use. Many RTP voice payloads are inelastic and cannot react to loss or delay in any substantive way. For these voice payloads, the network SHOULD police at ingress to ensure that the voice traffic stays within its negotiated bounds. Having thus assured a predictable input rate, the network may use a priority queue to ensure nominal delay and variation in delay.

Another approach that may be used in small and bandwidth-constrained networks for limited number of flows is RSVP [RFC2205] [RFC2996]. However, there is concern with the scalability of this solution in large networks where aggregation of reservations [RFC3175] is considered to be required.

2. Service Differentiation

There are practical limits on the level of service differentiation that should be offered in the IP networks. We believe we have defined a practical approach in delivering service differentiation by defining different service classes that networks may choose to support in order to provide the appropriate level of behaviors and performance needed by current and future applications and services. The defined structure for providing services allows several applications having similar traffic characteristics and performance requirements to be grouped into the same service class. This approach provides a lot of flexibility in providing the appropriate level of service differentiation for current and new, yet unknown applications without introducing significant changes to routers or network configurations when a new traffic type is added to the network.

2.1. Service Classes

Traffic flowing in a network can be classified in many different ways. We have chosen to divide it into two groupings, network control and user/subscriber traffic. To provide service differentiation, different service classes are defined in each grouping. The network control traffic group can further be divided into two service classes (see Section 3 for detailed definition of each service class):

- o "Network Control" for routing and network control function.
- o "OAM" (Operations, Administration, and Management) for network configuration and management functions.

The user/subscriber traffic group is broken down into ten service classes to provide service differentiation for all the different types of applications/services (see Section 4 for detailed definition of each service class):

- o Telephony service class is best suited for applications that require very low delay variation and are of constant rate, such as IP telephony (VoIP) and circuit emulation over IP applications.
- o Signaling service class is best suited for peer-to-peer and client-server signaling and control functions using protocols such as SIP, SIP-T, H.323, H.248, and Media Gateway Control Protocol (MGCP).
- o Multimedia Conferencing service class is best suited for applications that require very low delay and have the ability to change encoding rate (rate adaptive), such as H.323/V2 and later video conferencing service.
- o Real-Time Interactive service class is intended for interactive variable rate inelastic applications that require low jitter and loss and very low delay, such as interactive gaming applications that use RTP/UDP streams for game control commands, and video conferencing applications that do not have the ability to change encoding rates or to mark packets with different importance indications.
- o Multimedia Streaming service class is best suited for variable rate elastic streaming media applications where a human is waiting for output and where the application has the capability to react to packet loss by reducing its transmission rate, such as streaming video and audio and webcast.
- o Broadcast Video service class is best suited for inelastic streaming media applications that may be of constant or variable rate, requiring low jitter and very low packet loss, such as broadcast TV and live events, video surveillance, and security.

- o Low-Latency Data service class is best suited for data processing applications where a human is waiting for output, such as web-based ordering or an Enterprise Resource Planning (ERP) application.
- o High-Throughput Data service class is best suited for store and forward applications such as FTP and billing record transfer.
- o Standard service class is for traffic that has not been identified as requiring differentiated treatment and is normally referred to as best effort.
- o Low-Priority Data service class is intended for packet flows where bandwidth assurance is not required.

2.2. Categorization of User Service Classes

The ten defined user/subscriber service classes listed above can be grouped into a small number of application categories. For some application categories, it was felt that more than one service class was needed to provide service differentiation within that category due to the different traffic characteristic of the applications, control function, and the required flow behavior. Figure 1 provides a summary of service class grouping into four application categories.

Application Control Category

- o The Signaling service class is intended to be used to control applications or user endpoints. Examples of protocols that would use this service class are SIP or H.248 for IP telephone service and SIP or Internet Group Management Protocol (IGMP) for control of broadcast TV service to subscribers. Although user signaling flows have similar performance requirements as Low-Latency Data, they need to be distinguished and marked with a different DSCP. The essential distinction is something like "administrative control and management" of the traffic affected as the protocols in this class tend to be tied to the media stream/session they signal and control.

Media-Oriented Category

Due to the vast number of new (in process of being deployed) and already-in-use media-oriented services in IP networks, five service classes have been defined.

- o Telephony service class is intended for IP telephony (VoIP) service. It may also be used for other applications that meet the defined traffic characteristics and performance requirements.
- o Real-Time Interactive service class is intended for inelastic video flows from applications such as SIP-based desktop video conferencing applications and for interactive gaming.

- o Multimedia Conferencing service class is for video conferencing solutions that have the ability to reduce their transmission rate on detection of congestion. These flows can therefore be classified as rate adaptive. As currently two types of video conferencing equipment are used in IP networks (ones that generate inelastic traffic and ones that generate rate-adaptive traffic), two service class are needed. The Real-Time Interactive service class should be used for equipment that generates inelastic video flows and the Multimedia Conferencing service class for equipment that generates rate-adaptive video flows.
- o Broadcast Video service class is to be used for inelastic traffic flows, which are intended for broadcast TV service and for transport of live video and audio events.
- o Multimedia Streaming service class is to be used for elastic multimedia traffic flows. This multimedia content is typically stored before being transmitted. It is also buffered at the receiving end before being played out. The buffering is sufficiently large to accommodate any variation in transmission rate that is encountered in the network. Multimedia entertainment over IP delivery services that are being developed can generate both elastic and inelastic traffic flows; therefore, two service classes are defined to address this space, respectively: Multimedia Streaming and Broadcast Video.

Data Category

The data category is divided into three service classes.

- o Low-Latency Data for applications/services that require low delay or latency for bursty but short-lived flows.
- o High-Throughput Data for applications/services that require good throughput for long-lived bursty flows. High Throughput and Multimedia Streaming are close in their traffic flow characteristics with High Throughput being a bit more bursty and not as long-lived as Multimedia Streaming.
- o Low-Priority Data for applications or services that can tolerate short or long interruptions of packet flows. The Low-Priority Data service class can be viewed as "don't care" to some degree.

Best-Effort Category

- o All traffic that is not differentiated in the network falls into this category and is mapped into the Standard service class. If a packet is marked with a DSCP value that is not supported in the network, it SHOULD be forwarded using the Standard service class.

Figure 1, below, provides a grouping of the defined user/subscriber service classes into four categories, with indications of which ones use an independent flow for signaling or control; type of flow behavior (elastic, rate adaptive, or inelastic); and the last column provides end user Quality of Service (QoS) rating as defined in ITU-T Recommendation G.1010.

Application Categories	Service Class	Signaled	Flow Behavior	G.1010 Rating
Application Control	Signaling	Not applicable	Inelastic	Responsive
Media-Oriented	Telephony	Yes	Inelastic	Interactive
	Real-Time Interactive	Yes	Inelastic	Interactive
	Multimedia Conferencing	Yes	Rate Adaptive	Interactive
	Broadcast Video	Yes	Inelastic	Responsive
	Multimedia Streaming	Yes	Elastic	Timely
Data	Low-Latency Data	No	Elastic	Responsive
	High-Throughput Data	No	Elastic	Timely
	Low-Priority Data	No	Elastic	Non-critical
Best Effort	Standard	Not Specified		Non-critical

Figure 1. User/Subscriber Service Classes Grouping

Here is a short explanation of the end user QoS category as defined in ITU-T Recommendation G.1010. User traffic is divided into four different categories, namely, interactive, responsive, timely, and non-critical. An example of interactive traffic is between two humans and is most sensitive to delay, loss, and jitter. Another example of interactive traffic is between two servers where very low delay and loss are needed. Responsive traffic is typically between a human and a server but can also be between two servers. Responsive traffic is less affected by jitter and can tolerate longer delays than interactive traffic. Timely traffic is either between servers or servers and humans and the delay tolerance is significantly longer than responsive traffic. Non-critical traffic is normally between servers/machines where delivery may be delayed for a period of time.

2.3. Service Class Characteristics

This document provides guidelines for network administrators in configuring their network for the level of service differentiation that is appropriate in their network to meet their QoS needs. It is expected that network operators will configure and provide in their networks a subset of the defined service classes. Our intent is to provide guidelines for configuration of Differentiated Services for a wide variety of applications, services, and network configurations. In addition, network administrators may choose to define and deploy other service classes in their network.

Figure 2 provides a behavior view for traffic serviced by each service class. The traffic characteristics column defines the characteristics and profile of flows serviced, and the tolerance to loss, delay, and jitter columns define the treatment the flows will receive. End-to-end quantitative performance requirements may be obtained from ITU-T Recommendations Y.1541 and Y.1540.

Service Class Name	Traffic Characteristics	Tolerance to		
		Loss	Delay	Jitter
Network Control	Variable size packets, mostly inelastic short messages, but traffic can also burst (BGP)	Low	Low	Yes
Telephony	Fixed-size small packets, constant emission rate, inelastic and low-rate flows	Very Low	Very Low	Very Low
Signaling	Variable size packets, some what bursty short-lived flows	Low	Low	Yes
Multimedia Conferencing	Variable size packets, constant transmit interval, rate adaptive, reacts to loss	Low - Medium	Very Low	Low
Real-Time Interactive	RTP/UDP streams, inelastic, mostly variable rate	Low	Very Low	Low
Multimedia Streaming	Variable size packets, elastic with variable rate	Low - Medium	Medium	Yes
Broadcast Video	Constant and variable rate, inelastic, non-bursty flows	Very Low	Medium	Low
Low-Latency Data	Variable rate, bursty short-lived elastic flows	Low	Low - Medium	Yes
OAM	Variable size packets, elastic & inelastic flows	Low	Medium	Yes
High-Throughput Data	Variable rate, bursty long-lived elastic flows	Low	Medium - High	Yes
Standard	A bit of everything	Not Specified		
Low-Priority Data	Non-real-time and elastic	High	High	Yes

Figure 2. Service Class Characteristics

Notes for Figure 2: A "Yes" in the jitter-tolerant column implies that data is buffered in the endpoint and that a moderate level of network-induced variation in delay will not affect the application. Applications that use TCP as a transport are generally good examples. Routing protocols and peer-to-peer signaling also fall in this class; although loss can create problems in setting up calls, a moderate level of jitter merely makes call placement a little less predictable in duration.

Service classes indicate the required traffic forwarding treatment in order to meet user, application, or network expectations. Section 3 defines the service classes that MAY be used for forwarding network control traffic, and Section 4 defines the service classes that MAY be used for forwarding user traffic with examples of intended application types mapped into each service class. Note that the application types are only examples and are not meant to be all-inclusive or prescriptive. Also, note that the service class naming or ordering does not imply any priority ordering. They are simply reference names that are used in this document with associated QoS behaviors that are optimized for the particular application types they support. Network administrators MAY choose to assign different service class names to the service classes that they will support. Figure 3 defines the RECOMMENDED relationship between service classes and DS codepoint assignment with application examples. It is RECOMMENDED that this relationship be preserved end to end.

Service Class Name	DSCP Name	DSCP Value	Application Examples
Network Control	CS6	110000	Network routing
Telephony	EF	101110	IP Telephony bearer
Signaling	CS5	101000	IP Telephony signaling
Multimedia Conferencing	AF41, AF42 AF43	100010, 100100 100110	H.323/V2 video conferencing (adaptive)
Real-Time Interactive	CS4	100000	Video conferencing and Interactive gaming
Multimedia Streaming	AF31, AF32 AF33	011010, 011100 011110	Streaming video and audio on demand
Broadcast Video	CS3	011000	Broadcast TV & live events
Low-Latency Data	AF21, AF22 AF23	010010, 010100 010110	Client/server transactions Web-based ordering
OAM	CS2	010000	OAM&P
High-Throughput Data	AF11, AF12 AF13	001010, 001100 001110	Store and forward applications
Standard	DF (CS0)	000000	Undifferentiated applications
Low-Priority Data	CS1	001000	Any flow that has no BW assurance

Figure 3. DSCP to Service Class Mapping

Notes for Figure 3: Default Forwarding (DF) and Class Selector 0 (CS0) provide equivalent behavior and use the same DS codepoint, '000000'.

It is expected that network administrators will base their choice of the service classes that they will support on their need, starting off with three or four service classes for user traffic and adding others as the need arises.

Figure 4 provides a summary of DiffServ QoS mechanisms that SHOULD be used for the defined service classes that are further detailed in Sections 3 and 4 of this document. According to what applications/services need to be differentiated, network administrators can choose the service class(es) that need to be supported in their network.

Service Class	DSCP	Conditioning at DS Edge	PHB Used	Queuing	AQM
Network Control	CS6	See Section 3.1	RFC2474	Rate	Yes
Telephony	EF	Police using sr+bs	RFC3246	Priority	No
Signaling	CS5	Police using sr+bs	RFC2474	Rate	No
Multimedia Conferencing	AF41 AF42 AF43	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Real-Time Interactive	CS4	Police using sr+bs	RFC2474	Rate	No
Multimedia Streaming	AF31 AF32 AF33	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Broadcast Video	CS3	Police using sr+bs	RFC2474	Rate	No
Low-Latency Data	AF21 AF22 AF23	Using single-rate, three-color marker (such as RFC 2697)	RFC2597	Rate	Yes per DSCP
OAM	CS2	Police using sr+bs	RFC2474	Rate	Yes
High-Throughput Data	AF11 AF12 AF13	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Standard	DF	Not applicable	RFC2474	Rate	Yes
Low-Priority Data	CS1	Not applicable	RFC3662	Rate	Yes

Figure 4. Summary of QoS Mechanisms Used for Each Service Class

Notes for Figure 4:

- o Conditioning at DS edge means that traffic conditioning is performed at the edge of the DiffServ network where untrusted user devices are connected or between two DiffServ networks.
- o "sr+bs" represents a policing mechanism that provides single rate with burst size control.
- o The single-rate, three-color marker (srTCM) behavior SHOULD be equivalent to RFC 2697, and the two-rate, three-color marker (trTCM) behavior SHOULD be equivalent to RFC 2698.
- o The PHB for Real-Time Interactive service class SHOULD be configured to provide high bandwidth assurance. It MAY be configured as a second EF PHB that uses relaxed performance parameters and a rate scheduler.
- o The PHB for Broadcast Video service class SHOULD be configured to provide high bandwidth assurance. It MAY be configured as a third EF PHB that uses relaxed performance parameters and a rate scheduler.
- o In network segments that use IP precedence marking, only one of the two service classes can be supported, High-Throughput Data or Low-Priority Data. We RECOMMEND that the DSCP value(s) of the unsupported service class be changed to 000xx1 on ingress and changed back to original value(s) on egress of the network segment that uses precedence marking. For example, if Low-Priority Data is mapped to Standard service class, then 000001 DSCP marking MAY be used to distinguish it from Standard marked packets on egress.

2.4. Deployment Scenarios

It is expected that network administrators will base their choice of the service classes that they will support on their need, starting off with three or four service classes for user traffic and adding more service classes as the need arises. In this section, we provide three examples of possible deployment scenarios.

2.4.1. Example 1

A network administrator determines that he needs to provide different performance levels (quality of service) in his network for the services that he will be offering to his customers. He needs to enable his network to provide:

- o Reliable VoIP (telephony) service, equivalent to Public Switched Telephone Network (PSTN).
- o A low-delay assured bandwidth data service.
- o Support for current Internet services.

For this example, the network administrator's needs are addressed with the deployment of the following six service classes:

- o Network Control service class for routing and control traffic that is needed for reliable operation of the provider's network.
- o Standard service class for all traffic that will receive normal (undifferentiated) forwarding treatment through the network for support of current Internet service.
- o Telephony service class for VoIP (telephony) bearer traffic.
- o Signaling service class for Telephony signaling to control the VoIP service.
- o Low-Latency Data service class for the low-delay assured bandwidth differentiated data service.
- o OAM service class for operation and management of the network.

Figure 5 provides a summary of the mechanisms needed for delivery of service differentiation for Example 1.

Service Class	DSCP	Conditioning at DS Edge	PHB Used	Queuing	AQM
Network Control	CS6	See Section 3.1	RFC2474	Rate	Yes
Telephony	EF	Police using sr+bs	RFC3246	Priority	No
Signaling	CS5	Police using sr+bs	RFC2474	Rate	No
Low-Latency Data	AF21 AF22 AF23	Using single-rate, three-color marker (such as RFC 2697)	RFC2597	Rate	Yes per DSCP
OAM	CS2	Police using sr+bs	RFC2474	Rate	Yes
Standard	DF(CS0) +other	Not applicable	RFC2474	Rate	Yes

Figure 5. Service Provider Network Configuration Example 1

Notes for Figure 5:

- o "sr+bs" represents a policing mechanism that provides single rate with burst size control.
- o The single-rate, three-color marker (srTCM) behavior SHOULD be equivalent to RFC 2697.
- o Any packet that is marked with DSCP value that is not represented by the supported service classes SHOULD be forwarded using the Standard service class.

2.4.2. Example 2

With this example, we show how network operators with Example 1 capabilities can evolve their service offering to provide three new additional services to their customers. The new additional service capabilities that are to be added are:

- o SIP-based desktop video conference capability to complement VoIP (telephony) service.
- o TV and on-demand movie viewing service to residential subscribers.
- o Network-based data storage and file backup service to business customers.

The new additional services that the network administrator would like to offer are addressed with the deployment of the following four additional service classes (these are additions to the six service classes already defined in Example 1):

- o Real-Time Interactive service class for transport of MPEG-4 real-time video flows to support desktop video conferencing. The control/signaling for video conferencing is done using the Signaling service class.
- o Broadcast Video service class for transport of IPTV broadcast information. The channel selection and control is via IGMP mapped into the Signaling service class.
- o Multimedia Streaming service class for transport of stored MPEG-2 or MPEG-4 content. The selection and control of streaming information is done using the Signaling service class. The selection of Multimedia Streaming service class for on-demand movie service was chosen as the set-top box used for this service has local buffering capability to compensate for the bandwidth variability of the elastic streaming information. Note that if transport of on-demand movie service is inelastic, then the Broadcast Video service class SHOULD be used.
- o High-Throughput Data service class is for transport of bulk data for network-based storage and file backup service to business customers.

Figure 6 provides a summary of the mechanisms needed for delivery of service differentiation for all the service classes used in Example 2.

Service Class	DSCP	Conditioning at DS Edge	PHB Used	Queuing	AQM
Network Control	CS6	See Section 3.1	RFC2474	Rate	Yes
Telephony	EF	Police using sr+bs	RFC3246	Priority	No
Signaling	CS5	Police using sr+bs	RFC2474	Rate	No
Real-time Interactive	CS4	Police using sr+bs	RFC2474	Rate	No
Broadcast Video	CS3	Police using sr+bs	RFC2474	Rate	No
Multimedia Streaming	AF31 AF32 AF33	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Low-Latency Data	AF21 AF22 AF23	Using single-rate, three-color marker (such as RFC 2697)	RFC2597	Rate	Yes per DSCP
OAM	CS2	Police using sr+bs	RFC2474	Rate	Yes
High-Throughput Data	AF11 AF12 AF13	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Standard	DF(CS0) +other	Not applicable	RFC2474	Rate	Yes

Figure 6. Service Provider Network Configuration Example 2

Notes for Figure 6:

- o "sr+bs" represents a policing mechanism that provides single rate with burst size control.
- o The single-rate, three-color marker (srTCM) behavior SHOULD be equivalent to RFC 2697, and the two-rate, three-color marker (trTCM) behavior SHOULD be equivalent to RFC 2698.

- o Any packet that is marked with DSCP value that is not represented by the supported service classes SHOULD be forwarded using the Standard service class.

2.4.3. Example 3

An enterprise network administrator determines that they need to provide different performance levels (quality of service) in their network for the new services that are being offered to corporate users. The enterprise network needs to:

- o Provide reliable corporate VoIP service.
- o Provide video conferencing service to selected Conference Rooms.
- o Support on-demand distribution of prerecorded audio and video information to large number of users.
- o Provide a priority data transfer capability for engineering teams to share design information.
- o Reduce or deny bandwidth during peak traffic periods for selected applications.
- o Continue to provide normal IP service to all remaining applications and services.

For this example, the enterprise's network needs are addressed with the deployment of the following nine service classes:

- o Network Control service class for routing and control traffic that is needed for reliable operation of the enterprise network.
- o OAM service class for operation and management of the network.
- o Standard service class for all traffic that will receive normal (undifferentiated) forwarding treatment.
- o Telephony service class for VoIP (telephony) bearer traffic.
- o Signaling service class for Telephony signaling to control the VoIP service.
- o Multimedia Conferencing service class for support of inter-Conference Room video conferencing service using H.323/V2 or similar equipment.
- o Multimedia Streaming service class for transfer of prerecorded audio and video information.
- o High-Throughput Data service class to provide bandwidth assurance for timely transfer of large engineering files.
- o Low-Priority Data service class for selected background applications where data transfer can be delayed or suspended for a period of time during peak network load conditions.

Figure 7 provides a summary of the mechanisms needed for delivery of service differentiation for Example 3.

Service Class	DSCP	Conditioning at DS Edge	PHB Used	Queuing	AQM
Network Control	CS6	See Section 3.2	RFC2474	Rate	Yes
Telephony	EF	Police using sr+bs	RFC3246	Priority	No
Signaling	CS5	Police using sr+bs	RFC2474	Rate	No
Multimedia Conferencing	AF41 AF42 AF43	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Multimedia Streaming	AF31 AF32 AF33	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
OAM	CS2	Police using sr+bs	RFC2474	Rate	Yes
High-Throughput Data	AF11 AF12 AF13	Using two-rate, three-color marker (such as RFC 2698)	RFC2597	Rate	Yes per DSCP
Low-Priority Data	CS1	Not applicable	RFC3662	Rate	Yes
Standard	DF(CS0) +other	Not applicable	RFC2474	Rate	Yes

Figure 7. Enterprise Network Configuration Example

Notes for Figure 7:

- o "sr+bs" represents a policing mechanism that provides single rate with burst size control.
- o The single-rate, three-color marker (srTCM) behavior SHOULD be equivalent to RFC 2697, and the two-rate, three-color marker (trTCM) behavior SHOULD be equivalent to RFC 2698.
- o Any packet that is marked with DSCP value that is not represented by the supported service classes SHOULD be forwarded using the Standard service class.

3. Network Control Traffic

Network control traffic is defined as packet flows that are essential for stable operation of the administered network as well as for information that may be exchanged between neighboring networks across a peering point where SLAs are in place. Network control traffic is different from user application control (signaling) that may be generated by some applications or services. Network control traffic is mostly between routers and network nodes that are used for operating, administering, controlling, or managing the network segments. Network Control Traffic may be split into two service classes, i.e., Network Control and OAM.

3.1. Current Practice in the Internet

Based on today's routing protocols and network control procedures that are used in the Internet, we have determined that CS6 DSCP value SHOULD be used for routing and control and that CS7 DSCP value SHOULD be reserved for future use, potentially for future routing or control protocols. Network administrators MAY use a Local/Experimental DSCP; therefore, they may use a locally defined service class within their network to further differentiate their routing and control traffic.

RECOMMENDED Network Edge Conditioning for CS7 DSCP marked packets:

- o Drop or remark CS7 packets at ingress to DiffServ network domain.
- o CS7 marked packets SHOULD NOT be sent across peering points.
Exchange of control information across peering points SHOULD be done using CS6 DSCP and the Network Control service class.

3.2. Network Control Service Class

The Network Control service class is used for transmitting packets between network devices (routers) that require control (routing) information to be exchanged between nodes within the administrative domain as well as across a peering point between different administrative domains. Traffic transmitted in this service class is very important as it keeps the network operational, and it needs to be forwarded in a timely manner.

The Network Control service class SHOULD be configured using the DiffServ Class Selector (CS) PHB, defined in [RFC2474]. This service class SHOULD be configured so that the traffic receives a minimum bandwidth guarantee, to ensure that the packets always receive timely service. The configured forwarding resources for Network Control service class SHOULD be such that the probability of packet drop under peak load is very low in this service class. The Network

Control service class SHOULD be configured to use a Rate Queuing system such as defined in Section 1.4.1.2 of this document.

The following are examples of protocols and applications that SHOULD use the Network Control service class:

- o Routing packet flows: OSPF, BGP, ISIS, RIP.
- o Control information exchange within and between different administrative domains across a peering point where SLAs are in place.
- o LSP setup using CR-LDP and RSVP-TE.

The following protocols and applications SHOULD NOT use the Network Control service class:

- o User traffic.

The following are traffic characteristics of packet flows in the Network Control service class:

- o Mostly messages sent between routers and network servers.
- o Variable size packets, normally one packet at a time, but traffic can also burst (BGP).
- o User traffic is not allowed to use this service class. By user traffic, we mean packet flows that originate from user-controlled end points that are connected to the network.

The RECOMMENDED DSCP marking is CS6 (Class Selector 6).

RECOMMENDED Network Edge Conditioning:

- o At peering points (between two DiffServ networks) where SLAs are in place, CS6 marked packets SHOULD be policed, e.g., using a single rate with burst size (sr+bs) token bucket policer to keep the CS6 marked packet flows to within the traffic rate specified in the SLA.
- o CS6 marked packet flows from untrusted sources (for example, end user devices) SHOULD be dropped or remarked at ingress to the DiffServ network.
- o Packets from users/subscribers are not permitted access to the Network Control service classes.

The fundamental service offered to the Network Control service class is enhanced best-effort service with high bandwidth assurance. Since this service class is used to forward both elastic and inelastic flows, the service SHOULD be engineered so that the Active Queue Management (AQM) [RFC2309] is applied to CS6 marked packets.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth, and the max-threshold specifies the queue depth above which all traffic is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold CS6 < max-threshold CS6
- o max-threshold CS6 <= memory assigned to the queue

Note: Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

3.3. OAM Service Class

The OAM (Operations, Administration, and Management) service class is RECOMMENDED for OAM&P (Operations, Administration, and Management and Provisioning) using protocols such as Simple Network Management Protocol (SNMP), Trivial File Transfer Protocol (TFTP), FTP, Telnet, and Common Open Policy Service (COPS). Applications using this service class require a low packet loss but are relatively not sensitive to delay. This service class is configured to provide good packet delivery for intermittent flows.

The OAM service class SHOULD use the Class Selector (CS) PHB defined in [RFC2474]. This service class SHOULD be configured to provide a minimum bandwidth assurance for CS2 marked packets to ensure that they get forwarded. The OAM service class SHOULD be configured to use a Rate Queuing system such as defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the OAM service class:

- o Provisioning and configuration of network elements.
- o Performance monitoring of network elements.
- o Any network operational alarms.

The following are traffic characteristics:

- o Variable size packets.
- o Intermittent traffic flows.
- o Traffic may burst at times.
- o Both elastic and inelastic flows.
- o Traffic not sensitive to delays.

RECOMMENDED DSCP marking:

- o All flows in this service class are marked with CS2 (Class Selector 2).

Applications or IP end points SHOULD pre-mark their packets with CS2 DSCP value. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475].

RECOMMENDED conditioning performed at DiffServ network edge:

- o Packet flow marking (DSCP setting) from untrusted sources (end user devices) SHOULD be verified at ingress to DiffServ network using Multifield (MF) Classification methods, defined in [RFC2475].
- o Packet flows from untrusted sources (end user devices) SHOULD be policed at ingress to DiffServ network, e.g., using single rate with burst size token bucket policer to ensure that the traffic stays within its negotiated or engineered bounds.
- o Packet flows from trusted sources (routers inside administered network) MAY not require policing.
- o Normally OAM&P CS2 marked packet flows are not allowed to flow across peering points. If that is the case, then CS2 marked packets SHOULD be policed (dropped) at both egress and ingress peering interfaces.

The fundamental service offered to "OAM" traffic is enhanced best-effort service with controlled rate. The service SHOULD be engineered so that CS2 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery. Since this service class is used to forward both elastic and inelastic flows, the service SHOULD be engineered so that Active Queue Management [RFC2309] is applied to CS2 marked packets.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth for each DSCP, and the max-threshold specifies the queue depth above which all traffic with such a DSCP is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold CS2 < max-threshold CS2
- o max-threshold CS2 <= memory assigned to the queue

Note: Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4. User Traffic

User traffic is defined as packet flows between different users or subscribers. It is the traffic that is sent to or from end-terminals and that supports a very wide variety of applications and services. User traffic can be differentiated in many different ways; therefore,

we investigated several different approaches to classifying user traffic. We looked at differentiating user traffic as real-time versus non-real-time, elastic or rate-adaptive versus inelastic, sensitive versus insensitive to loss as well as traffic categorization as interactive, responsive, timely, and non-critical, as defined in ITU-T Recommendation G.1010. In the final analysis, we used all of the above for service differentiation, mapping application types that seemed to have different sets of performance sensitivities, and requirements to different service classes.

Network administrators can categorize their applications according to the type of behavior that they require and MAY choose to support all or a subset of the defined service classes. Figure 3 provides some common applications and the forwarding service classes that best support them, based on their performance requirements.

4.1. Telephony Service Class

The Telephony service class is RECOMMENDED for applications that require real-time, very low delay, very low jitter, and very low packet loss for relatively constant-rate traffic sources (inelastic traffic sources). This service class SHOULD be used for IP telephony service.

The fundamental service offered to traffic in the Telephony service class is minimum jitter, delay, and packet loss service up to a specified upper bound. Operation is in some respect similar to an ATM CBR service, which has guaranteed bandwidth and which, if it stays within the negotiated rate, experiences nominal delay and no loss. The EF PHB has a similar guarantee.

Typical configurations negotiate the setup of telephone calls over IP, using protocols such as H.248, MEGACO, H.323, or SIP. When a user has been authorized to send telephony traffic, the call admission procedure should have verified that the newly admitted flow will be within the capacity of the Telephony service class forwarding capability in the network. For VoIP (telephony) service, call admission control is usually performed by a telephony call server/gatekeeper using signaling (SIP, H.323, H.248, MEGACO, etc.) on access points to the network. The bandwidth in the core network and the number of simultaneous VoIP sessions that can be supported needs to be engineered and controlled so that there is no congestion for this service. Since the inelastic types of RTP payloads in this class do not react to loss or significant delay in any substantive way, the Telephony service class SHOULD forward packets as soon as possible. Some RTP payloads that may be used in telephony applications are adaptive and will not be in this class.

The Telephony service class SHOULD use Expedited Forwarding (EF) PHB, as defined in [RFC3246], and SHOULD be configured to receive guaranteed forwarding resources so that all packets are forwarded quickly. The Telephony service class SHOULD be configured to use a Priority Queuing system such as that defined in Section 1.4.1.1 of this document.

The following applications SHOULD use the Telephony service class:

- o VoIP (G.711, G.729 and other codecs).
- o Voice-band data over IP (modem, fax).
- o T.38 fax over IP.
- o Circuit emulation over IP, virtual wire, etc.
- o IP Virtual Private Network (VPN) service that specifies single-rate, mean network delay that is slightly longer than network propagation delay, very low jitter, and a very low packet loss.

The following are traffic characteristics:

- o Mostly fixed-size packets for VoIP (60, 70, 120 or 200 bytes in size).
- o Packets emitted at constant time intervals.
- o Admission control of new flows is provided by telephony call server, media gateway, gatekeeper, edge router, end terminal, or access node that provides flow admission control function.

Applications or IP end points SHOULD pre-mark their packets with EF DSCP value. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475].

The RECOMMENDED DSCP marking is EF for the following applications:

- o VoIP (G.711, G.729 and other codecs).
- o Voice-band data over IP (modem and fax).
- o T.38 fax over IP.
- o Circuit emulation over IP, virtual wire, etc.

RECOMMENDED Network Edge Conditioning:

- o Packet flow marking (DSCP setting) from untrusted sources (end user devices) SHOULD be verified at ingress to DiffServ network using Multifield (MF) Classification methods, defined in [RFC2475].
- o Packet flows from untrusted sources (end user devices) SHOULD be policed at ingress to DiffServ network, e.g., using single rate with burst size token bucket policer to ensure that the telephony traffic stays within its negotiated bounds.

- o Policing is OPTIONAL for packet flows from trusted sources whose behavior is ensured via other means (e.g., administrative controls on those systems).
- o Policing of Telephony packet flows across peering points where SLA is in place is OPTIONAL as telephony traffic will be controlled by admission control mechanism between peering points.

The fundamental service offered to "Telephony" traffic is enhanced best-effort service with controlled rate, very low delay, and very low loss. The service MUST be engineered so that EF marked packet flows have sufficient bandwidth in the network to provide guaranteed delivery. Normally traffic in this service class does not respond dynamically to packet loss. As such, Active Queue Management [RFC2309] SHOULD NOT be applied to EF marked packet flows.

4.2. Signaling Service Class

The Signaling service class is RECOMMENDED for delay-sensitive client-server (traditional telephony) and peer-to-peer application signaling. Telephony signaling includes signaling between IP phone and soft-switch, soft-client and soft-switch, and media gateway and soft-switch as well as peer-to-peer using various protocols. This service class is intended to be used for control of sessions and applications. Applications using this service class require a relatively fast response, as there are typically several messages of different sizes sent for control of the session. This service class is configured to provide good response for short-lived, intermittent flows that require real-time packet forwarding. To minimize the possibility of ring clipping at start of call for VoIP service that interfaces to a circuit switch Exchange in the Public Switched Telephone Network (PSTN), the Signaling service class SHOULD be configured so that the probability of packet drop or significant queuing delay under peak load is very low in IP network segments that provide this interface. The term "ring clipping" refers to those instances where the front end of a ringing signal is altered because the bearer path is not made available in time to carry all of the audible ringing signal. This condition may occur due to a race condition between when the tone generator in the circuit switch Exchange is turned on and when the bearer path through the IP network is enabled. See Section 8.1 for additional explanation of "ring clipping" and Section 5.1 for explanation of mapping different signaling methods to service classes.

The Signaling service class SHOULD use the Class Selector (CS) PHB, defined in [RFC2474]. This service class SHOULD be configured to provide a minimum bandwidth assurance for CS5 marked packets to ensure that they get forwarded. The Signaling service class SHOULD

be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the Signaling service class:

- o Peer-to-peer IP telephony signaling (e.g., using SIP, H.323).
- o Peer-to-peer signaling for multimedia applications (e.g., using SIP, H.323).
- o Peer-to-peer real-time control function.
- o Client-server IP telephony signaling using H.248, MEGACO, MGCP, IP encapsulated ISDN, or other proprietary protocols.
- o Signaling to control IPTV applications using protocols such as IGMP.
- o Signaling flows between high-capacity telephony call servers or soft switches using protocol such as SIP-T. Such high-capacity devices may control thousands of telephony (VoIP) calls.

The following are traffic characteristics:

- o Variable size packets, normally one packet at a time.
- o Intermittent traffic flows.
- o Traffic may burst at times.
- o Delay-sensitive control messages sent between two end points.

RECOMMENDED DSCP marking:

- o All flows in this service class are marked with CS5 (Class Selector 5).

Applications or IP end points SHOULD pre-mark their packets with CS5 DSCP value. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475].

RECOMMENDED conditioning performed at DiffServ network edge:

- o Packet flow marking (DSCP setting) from untrusted sources (end user devices) SHOULD be verified at ingress to DiffServ network using Multifield (MF) Classification methods defined in [RFC2475].
- o Packet flows from untrusted sources (end user devices) SHOULD be policed at ingress to DiffServ network, e.g., using single rate with burst size token bucket policer to ensure that the traffic stays within its negotiated or engineered bounds.
- o Packet flows from trusted sources (application servers inside administered network) MAY not require policing.
- o Policing of packet flows across peering points SHOULD be performed to the Service Level Agreement (SLA).

The fundamental service offered to "Signaling" traffic is enhanced best-effort service with controlled rate and delay. The service SHOULD be engineered so that CS5 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery and low delay. Normally, traffic in this service class does not respond dynamically to packet loss. As such, Active Queue Management [RFC2309] SHOULD NOT be applied to CS5 marked packet flows.

4.3. Multimedia Conferencing Service Class

The Multimedia Conferencing service class is RECOMMENDED for applications that require real-time service for rate-adaptive traffic. H.323/V2 and later versions of video conferencing equipment with dynamic bandwidth adjustment are such applications. The traffic sources in this service class have the ability to dynamically change their transmission rate based on feedback from the receiver. One approach used in H.323/V2 equipment is, when the receiver detects a pre-configured level of packet loss, it signals to the transmitter the indication of possible on-path congestion. When available, the transmitter then selects a lower rate encoding codec. Note that today, many H.323/V2 video conferencing solutions implement fixed-step bandwidth change (usually reducing the rate), traffic resembling step-wise CBR.

Typical video conferencing configurations negotiate the setup of multimedia session using protocols such as H.323. When a user/endpoint has been authorized to start a multimedia session, the admission procedure should have verified that the newly admitted data rate will be within the engineered capacity of the Multimedia Conferencing service class. The bandwidth in the core network and the number of simultaneous video conferencing sessions that can be supported SHOULD be engineered to control traffic load for this service.

The Multimedia Conferencing service class SHOULD use the Assured Forwarding (AF) PHB, defined in [RFC2597]. This service class SHOULD be configured to provide a bandwidth assurance for AF41, AF42, and AF43 marked packets to ensure that they get forwarded. The Multimedia Conferencing service class SHOULD be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the Multimedia Conferencing service class:

- o H.323/V2 and later versions of video conferencing applications (interactive video).

- o Video conferencing applications with rate control or traffic content importance marking.
- o Application server-to-application server non-bursty data transfer requiring very low delay.
- o IP VPN service that specifies two rates and mean network delay that is slightly longer than network propagation delay.
- o Interactive, time-critical, and mission-critical applications.

The following are traffic characteristics:

- o Variable size packets.
- o The higher the rate, the higher the density of large packets.
- o Constant packet emission time interval.
- o Variable rate.
- o Source is capable of reducing its transmission rate based on detection of packet loss at the receiver.

Applications or IP end points SHOULD pre-mark their packets with DSCP values as shown below. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475] and mark all packets as AF4x. Note: In this case, the two-rate, three-color marker will be configured to operate in Color-Blind mode.

RECOMMENDED DSCP marking when performed by router closest to source:

- o AF41 = up to specified rate "A".
- o AF42 = in excess of specified rate "A" but below specified rate "B".
- o AF43 = in excess of specified rate "B".
- o Where "A" < "B".

Note: One might expect "A" to approximate the sum of the mean rates and "B" to approximate the sum of the peak rates.

RECOMMENDED DSCP marking when performed by H.323/V2 video conferencing equipment:

- o AF41 = H.323 video conferencing audio stream RTP/UDP.
- o AF41 = H.323 video conferencing video control RTCP/TCP.
- o AF41 = H.323 video conferencing video stream up to specified rate "A".
- o AF42 = H.323 video conferencing video stream in excess of specified rate "A" but below specified rate "B".
- o AF43 = H.323 video conferencing video stream in excess of specified rate "B".
- o Where "A" < "B".

RECOMMENDED conditioning performed at DiffServ network edge:

- o The two-rate, three-color marker SHOULD be configured to provide the behavior as defined in trTCM [RFC2698].
- o If packets are marked by trusted sources or a previously trusted DiffServ domain and the color marking is to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Aware mode.
- o If the packet marking is not trusted or the color marking is not to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Blind mode.

The fundamental service offered to "Multimedia Conferencing" traffic is enhanced best-effort service with controlled rate and delay. For video conferencing service, typically a 1% packet loss detected at the receiver triggers an encoding rate change, dropping to the next lower provisioned video encoding rate. As such, Active Queue Management [RFC2309] SHOULD be used primarily to switch the video encoding rate under congestion, changing from high rate to lower rate, i.e., 1472 kbps to 768 kbps. The probability of loss of AF41 traffic MUST NOT exceed the probability of loss of AF42 traffic, which in turn MUST NOT exceed the probability of loss of AF43 traffic.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth for each DSCP, and the max-threshold specifies the queue depth above which all traffic with such a DSCP is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold AF43 < max-threshold AF43
- o max-threshold AF43 <= min-threshold AF42
- o min-threshold AF42 < max-threshold AF42
- o max-threshold AF42 <= min-threshold AF41
- o min-threshold AF41 < max-threshold AF41
- o max-threshold AF41 <= memory assigned to the queue

Note: This configuration tends to drop AF43 traffic before AF42 and AF42 before AF41. Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4.4. Real-Time Interactive Service Class

The Real-Time Interactive service class is RECOMMENDED for applications that require low loss and jitter and very low delay for variable rate inelastic traffic sources. Interactive gaming and video conferencing applications that do not have the ability to change encoding rates or to mark packets with different importance

indications are such applications. The traffic sources in this traffic class do not have the ability to reduce their transmission rate according to feedback received from the receiving end.

Typically, applications in this service class are configured to negotiate the setup of RTP/UDP control session. When a user/end-point has been authorized to start a new session, the admission procedure should have verified that the newly admitted data rates will be within the engineered capacity of the Real-Time Interactive service class. The bandwidth in the core network and the number of simultaneous Real-time Interactive sessions that can be supported SHOULD be engineered to control traffic load for this service.

The Real-Time Interactive service class SHOULD use the Class Selector (CS) PHB, defined in [RFC2474]. This service class SHOULD be configured to provide a high assurance for bandwidth for CS4 marked packets to ensure that they get forwarded. The Real-Time Interactive service class SHOULD be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document. Note that this service class MAY be configured as a second EF PHB that uses relaxed performance parameter, a rate scheduler, and CS4 DSCP value.

The following applications SHOULD use the Real-Time Interactive service class:

- o Interactive gaming and control.
- o Video conferencing applications without rate control or traffic content importance marking.
- o IP VPN service that specifies single rate and mean network delay that is slightly longer than network propagation delay.
- o Inelastic, interactive, time-critical, and mission-critical applications requiring very low delay.

The following are traffic characteristics:

- o Variable size packets.
- o Variable rate, non-bursty.
- o Application is sensitive to delay variation between flows and sessions.
- o Lost packets, if any, are usually ignored by application.

RECOMMENDED DSCP marking:

- o All flows in this service class are marked with CS4 (Class Selector 4).

Applications or IP end points SHOULD pre-mark their packets with CS4 DSCP value. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475].

RECOMMENDED conditioning performed at DiffServ network edge:

- o Packet flow marking (DSCP setting) from untrusted sources (end user devices) SHOULD be verified at ingress to DiffServ network using Multifield (MF) Classification methods defined in [RFC2475].
- o Packet flows from untrusted sources (end user devices) SHOULD be policed at ingress to DiffServ network, e.g., using single rate with burst size token bucket policer to ensure that the traffic stays within its negotiated or engineered bounds.
- o Packet flows from trusted sources (application servers inside administered network) MAY not require policing.
- o Policing of packet flows across peering points SHOULD be performed to the Service Level Agreement (SLA).

The fundamental service offered to "Real-Time Interactive" traffic is enhanced best-effort service with controlled rate and delay. The service SHOULD be engineered so that CS4 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery. Normally, traffic in this service class does not respond dynamically to packet loss. As such, Active Queue Management [RFC2309] SHOULD NOT be applied to CS4 marked packet flows.

4.5. Multimedia Streaming Service Class

The Multimedia Streaming service class is RECOMMENDED for applications that require near-real-time packet forwarding of variable rate elastic traffic sources that are not as delay sensitive as applications using the Multimedia Conferencing service class. Such applications include streaming audio and video, some video (movies) on-demand applications, and webcasts. In general, the Multimedia Streaming service class assumes that the traffic is buffered at the source/destination; therefore, it is less sensitive to delay and jitter.

The Multimedia Streaming service class SHOULD use the Assured Forwarding (AF) PHB, defined in [RFC2597]. This service class SHOULD be configured to provide a minimum bandwidth assurance for AF31, AF32, and AF33 marked packets to ensure that they get forwarded. The Multimedia Streaming service class SHOULD be configured to use Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the Multimedia Streaming service class:

- o Buffered streaming audio (unicast).
- o Buffered streaming video (unicast).
- o Webcasts.
- o IP VPN service that specifies two rates and is less sensitive to delay and jitter.

The following are traffic characteristics:

- o Variable size packets.
- o The higher the rate, the higher the density of large packets.
- o Variable rate.
- o Elastic flows.
- o Some bursting at start of flow from some applications.

Applications or IP end points SHOULD pre-mark their packets with DSCP values as shown below. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475], and mark all packets as AF3x. Note: In this case, the two-rate, three-color marker will be configured to operate in Color-Blind mode.

RECOMMENDED DSCP marking:

- o AF31 = up to specified rate "A".
- o AF32 = in excess of specified rate "A" but below specified rate "B".
- o AF33 = in excess of specified rate "B".
- o Where "A" < "B".

Note: One might expect "A" to approximate the sum of the mean rates and "B" to approximate the sum of the peak rates.

RECOMMENDED conditioning performed at DiffServ network edge:

- o The two-rate, three-color marker SHOULD be configured to provide the behavior as defined in trTCM [RFC2698].
- o If packets are marked by trusted sources or a previously trusted DiffServ domain and the color marking is to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Aware mode.
- o If the packet marking is not trusted or the color marking is not to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Blind mode.

The fundamental service offered to "Multimedia Streaming" traffic is enhanced best-effort service with controlled rate and delay. The service SHOULD be engineered so that AF31 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery. Since the AF3x traffic is elastic and responds dynamically to packet loss, Active Queue Management [RFC2309] SHOULD be used primarily to reduce forwarding rate to the minimum assured rate at congestion points. The probability of loss of AF31 traffic MUST NOT exceed the probability of loss of AF32 traffic, which in turn MUST NOT exceed the probability of loss of AF33.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth for each DSCP, and the max-threshold specifies the queue depth above which all traffic with such a DSCP is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold AF33 < max-threshold AF33
- o max-threshold AF33 <= min-threshold AF32
- o min-threshold AF32 < max-threshold AF32
- o max-threshold AF32 <= min-threshold AF31
- o min-threshold AF31 < max-threshold AF31
- o max-threshold AF31 <= memory assigned to the queue

Note: This configuration tends to drop AF33 traffic before AF32 and AF32 before AF31. Note: Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4.6. Broadcast Video Service Class

The Broadcast Video service class is RECOMMENDED for applications that require near-real-time packet forwarding with very low packet loss of constant rate and variable rate inelastic traffic sources that are not as delay sensitive as applications using the Real-Time Interactive service class. Such applications include broadcast TV, streaming of live audio and video events, some video-on-demand applications, and video surveillance. In general, the Broadcast Video service class assumes that the destination end point has a de jitter buffer, for video application usually a 2 - 8 video-frame buffer (66 to several hundred of milliseconds), and therefore that it is less sensitive to delay and jitter.

The Broadcast Video service class SHOULD use the Class Selector (CS) PHB, defined in [RFC2474]. This service class SHOULD be configured to provide high assurance for bandwidth for CS3 marked packets to ensure that they get forwarded. The Broadcast Video service class SHOULD be configured to use Rate Queuing system such as that defined in Section 1.4.1.2 of this document. Note that this service class

MAY be configured as a third EF PHB that uses relaxed performance parameter, a rate scheduler, and CS3 DSCP value.

The following applications SHOULD use the Broadcast Video service class:

- o Video surveillance and security (unicast).
- o TV broadcast including HDTV (multicast).
- o Video on demand (unicast) with control (virtual DVD).
- o Streaming of live audio events (both unicast and multicast).
- o Streaming of live video events (both unicast and multicast).

The following are traffic characteristics:

- o Variable size packets.
- o The higher the rate, the higher the density of large packets.
- o Mixture of variable rate and constant rate flows.
- o Fixed packet emission time intervals.
- o Inelastic flows.

RECOMMENDED DSCP marking:

- o All flows in this service class are marked with CS3 (Class Selector 3).
- o In some cases, such as those for security and video surveillance applications, it may be desirable to use a different DSCP marking. If so, then locally user definable (EXP/LU) codepoints in the range '011xx1' MAY be used to provide unique traffic identification. The locally user definable (EXP/LU) codepoint(s) MAY be associated with the PHB that is used for CS3 traffic. Furthermore, depending on the network scenario, additional network edge conditioning policy MAY be needed for the EXP/LU codepoint(s) used.

Applications or IP end points SHOULD pre-mark their packets with CS3 DSCP value. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475].

RECOMMENDED conditioning performed at DiffServ network edge:

- o Packet flow marking (DSCP setting) from untrusted sources (end user devices) SHOULD be verified at ingress to DiffServ network using Multifield (MF) Classification methods defined in [RFC2475].
- o Packet flows from untrusted sources (end user devices) SHOULD be policed at ingress to DiffServ network, e.g., using single rate with burst size token bucket policer to ensure that the traffic stays within its negotiated or engineered bounds.

- o Packet flows from trusted sources (application servers inside administered network) MAY not require policing.
- o Policing of packet flows across peering points SHOULD be performed to the Service Level Agreement (SLA).

The fundamental service offered to "Broadcast Video" traffic is enhanced best-effort service with controlled rate and delay. The service SHOULD be engineered so that CS3 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery. Normally, traffic in this service class does not respond dynamically to packet loss. As such, Active Queue Management [RFC2309] SHOULD NOT be applied to CS3 marked packet flows.

4.7. Low-Latency Data Service Class

The Low-Latency Data service class is RECOMMENDED for elastic and responsive typically client-/server-based applications. Applications forwarded by this service class are those that require a relatively fast response and typically have asymmetrical bandwidth need, i.e., the client typically sends a short message to the server and the server responds with a much larger data flow back to the client. The most common example of this is when a user clicks a hyperlink (~ few dozen bytes) on a web page, resulting in a new web page to be loaded (Kbytes of data). This service class is configured to provide good response for TCP [RFC1633] short-lived flows that require real-time packet forwarding of variable rate traffic sources.

The Low-Latency Data service class SHOULD use the Assured Forwarding (AF) PHB, defined in [RFC2597]. This service class SHOULD be configured to provide a minimum bandwidth assurance for AF21, AF22, and AF23 marked packets to ensure that they get forwarded. The Low-Latency Data service class SHOULD be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the Low-Latency Data service class:

- o Client/server applications.
- o Systems Network Architecture (SNA) terminal to host transactions (SNA over IP using Data Link Switching (DLSw)).
- o Web-based transactions (E-commerce).
- o Credit card transactions.
- o Financial wire transfers.
- o Enterprise Resource Planning (ERP) applications (e.g., SAP/BaaN).
- o VPN service that supports Committed Information Rate (CIR) with up to two burst sizes.

The following are traffic characteristics:

- o Variable size packets.
- o Variable packet emission rate.
- o With packet bursts of TCP window size.
- o Short traffic bursts.
- o Source capable of reducing its transmission rate based on detection of packet loss at the receiver or through explicit congestion notification.

Applications or IP end points SHOULD pre-mark their packets with DSCP values as shown below. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475] and mark all packets as AF2x. Note: In this case, the single-rate, three-color marker will be configured to operate in Color-Blind mode.

RECOMMENDED DSCP marking:

- o AF21 = flow stream with packet burst size up to "A" bytes.
- o AF22 = flow stream with packet burst size in excess of "A" but below "B" bytes.
- o AF23 = flow stream with packet burst size in excess of "B" bytes.
- o Where "A" < "B".

RECOMMENDED conditioning performed at DiffServ network edge:

- o The single-rate, three-color marker SHOULD be configured to provide the behavior as defined in srTCM [RFC2697].
- o If packets are marked by trusted sources or a previously trusted DiffServ domain and the color marking is to be preserved, then the single-rate, three-color marker SHOULD be configured to operate in Color-Aware mode.
- o If the packet marking is not trusted or the color marking is not to be preserved, then the single-rate, three-color marker SHOULD be configured to operate in Color-Blind mode.

The fundamental service offered to "Low-Latency Data" traffic is enhanced best-effort service with controlled rate and delay. The service SHOULD be engineered so that AF21 marked packet flows have sufficient bandwidth in the network to provide high assurance of delivery. Since the AF2x traffic is elastic and responds dynamically to packet loss, Active Queue Management [RFC2309] SHOULD be used primarily to control TCP flow rates at congestion points by dropping packets from TCP flows that have large burst size. The probability of loss of AF21 traffic MUST NOT exceed the probability of loss of AF22 traffic, which in turn MUST NOT exceed the probability of loss

of AF23. Explicit Congestion Notification (ECN) [RFC3168] MAY also be used with Active Queue Management.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth for each DSCP, and the max-threshold specifies the queue depth above which all traffic with such a DSCP is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold AF23 < max-threshold AF23
- o max-threshold AF23 <= min-threshold AF22
- o min-threshold AF22 < max-threshold AF22
- o max-threshold AF22 <= min-threshold AF21
- o min-threshold AF21 < max-threshold AF21
- o max-threshold AF21 <= memory assigned to the queue

Note: This configuration tends to drop AF23 traffic before AF22 and AF22 before AF21. Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4.8. High-Throughput Data Service Class

The High-Throughput Data service class is RECOMMENDED for elastic applications that require timely packet forwarding of variable rate traffic sources and, more specifically, is configured to provide good throughput for TCP longer-lived flows. TCP [RFC1633] or a transport with a consistent Congestion Avoidance Procedure [RFC2581] [RFC3782] normally will drive as high a data rate as it can obtain over a long period of time. The FTP protocol is a common example, although one cannot definitively say that all FTP transfers are moving data in bulk.

The High-Throughput Data service class SHOULD use the Assured Forwarding (AF) PHB, defined in [RFC2597]. This service class SHOULD be configured to provide a minimum bandwidth assurance for AF11, AF12, and AF13 marked packets to ensure that they are forwarded in a timely manner. The High-Throughput Data service class SHOULD be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the High-Throughput Data service class:

- o Store and forward applications.
- o File transfer applications.
- o Email.
- o VPN service that supports two rates (committed information rate and excess or peak information rate).

The following are traffic characteristics:

- o Variable size packets.
- o Variable packet emission rate.
- o Variable rate.
- o With packet bursts of TCP window size.
- o Source capable of reducing its transmission rate based on detection of packet loss at the receiver or through explicit congestion notification.

Applications or IP end points SHOULD pre-mark their packets with DSCP values as shown below. If the end point is not capable of setting the DSCP value, then the router topologically closest to the end point SHOULD perform Multifield (MF) Classification, as defined in [RFC2475], and mark all packets as AF1x. Note: In this case, the two-rate, three-color marker will be configured to operate in Color-Blind mode.

RECOMMENDED DSCP marking:

- o AF11 = up to specified rate "A".
- o AF12 = in excess of specified rate "A" but below specified rate "B".
- o AF13 = in excess of specified rate "B".
- o Where "A" < "B".

RECOMMENDED conditioning performed at DiffServ network edge:

- o The two-rate, three-color marker SHOULD be configured to provide the behavior as defined in trTCM [RFC2698].
- o If packets are marked by trusted sources or a previously trusted DiffServ domain and the color marking is to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Aware mode.
- o If the packet marking is not trusted or the color marking is not to be preserved, then the two-rate, three-color marker SHOULD be configured to operate in Color-Blind mode.

The fundamental service offered to "High-Throughput Data" traffic is enhanced best-effort service with a specified minimum rate. The service SHOULD be engineered so that AF11 marked packet flows have sufficient bandwidth in the network to provide assured delivery. It can be assumed that this class will consume any available bandwidth and that packets traversing congested links may experience higher queuing delays or packet loss. Since the AF1x traffic is elastic and responds dynamically to packet loss, Active Queue Management [RFC2309] SHOULD be used primarily to control TCP flow rates at congestion points by dropping packets from TCP flows that have higher

rates first. The probability of loss of AF11 traffic MUST NOT exceed the probability of loss of AF12 traffic, which in turn MUST NOT exceed the probability of loss of AF13. In such a case, if one network customer is driving significant excess and another seeks to use the link, any losses will be experienced by the high-rate user, causing him to reduce his rate. Explicit Congestion Notification (ECN) [RFC3168] MAY also be used with Active Queue Management.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth for each DSCP, and the max-threshold specifies the queue depth above which all traffic with such a DSCP is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold AF13 < max-threshold AF13
- o max-threshold AF13 <= min-threshold AF12
- o min-threshold AF12 < max-threshold AF12
- o max-threshold AF12 <= min-threshold AF11
- o min-threshold AF11 < max-threshold AF11
- o max-threshold AF11 <= memory assigned to the queue

Note: This configuration tends to drop AF13 traffic before AF12 and AF12 before AF11. Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4.9. Standard Service Class

The Standard service class is RECOMMENDED for traffic that has not been classified into one of the other supported forwarding service classes in the DiffServ network domain. This service class provides the Internet's "best-effort" forwarding behavior. This service class typically has minimum bandwidth guarantee.

The Standard service class MUST use the Default Forwarding (DF) PHB, defined in [RFC2474], and SHOULD be configured to receive at least a small percentage of forwarding resources as a guaranteed minimum. This service class SHOULD be configured to use a Rate Queuing system such as that defined in Section 1.4.1.2 of this document.

The following applications SHOULD use the Standard service class:

- o Network services, DNS, DHCP, BootP.
- o Any undifferentiated application/packet flow transported through the DiffServ enabled network.

The following is a traffic characteristic:

- o Non-deterministic, mixture of everything.

The RECOMMENDED DSCP marking is DF (Default Forwarding) '000000'.

Network Edge Conditioning:

There is no requirement that conditioning of packet flows be performed for this service class.

The fundamental service offered to the Standard service class is best-effort service with active queue management to limit overall delay. Typical configurations SHOULD use random packet dropping to implement Active Queue Management [RFC2309] or Explicit Congestion Notification [RFC3168], and MAY impose a minimum or maximum rate on the queue.

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth, and the max-threshold specifies the queue depth above which all traffic is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold DF < max-threshold DF
- o max-threshold DF <= memory assigned to the queue

Note: Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

4.10. Low-Priority Data

The Low-Priority Data service class serves applications that run over TCP [RFC0793] or a transport with consistent congestion avoidance procedures [RFC2581] [RFC3782] and that the user is willing to accept service without guarantees. This service class is specified in [RFC3662] and [QBSS].

The following applications MAY use the Low-Priority Data service class:

- o Any TCP based-application/packet flow transported through the DiffServ enabled network that does not require any bandwidth assurances.

The following is a traffic characteristic:

- o Non-real-time and elastic.

Network Edge Conditioning:

There is no requirement that conditioning of packet flows be performed for this service class.

The RECOMMENDED DSCP marking is CS1 (Class Selector 1).

The fundamental service offered to the Low-Priority Data service class is best-effort service with zero bandwidth assurance. By placing it into a separate queue or class, it may be treated in a manner consistent with a specific Service Level Agreement.

Typical configurations SHOULD use Explicit Congestion Notification [RFC3168] or random loss to implement Active Queue Management [RFC2309].

If RED [RFC2309] is used as an AQM algorithm, the min-threshold specifies a target queue depth, and the max-threshold specifies the queue depth above which all traffic is dropped or ECN marked. Thus, in this service class, the following inequality should hold in queue configurations:

- o min-threshold CS1 < max-threshold CS1
- o max-threshold CS1 <= memory assigned to the queue

Note: Many other AQM algorithms exist and are used; they should be configured to achieve a similar result.

5. Additional Information on Service Class Usage

In this section, we provide additional information on how some specific applications should be configured to use the defined service classes.

5.1. Mapping for Signaling

There are many different signaling protocols, ways that signaling is used and performance requirements from applications that are controlled by these protocols. We believe that different signaling protocols should use the service class that best meets the objectives of application or service they control. The following mapping is recommended:

- o Peer-to-peer signaling using SIP/H.323 is marked with CS5 DSCP (use Signaling service class).

- o Client-server signaling as used in many implementation for IP telephony using H.248, MEGACO, MGCP, IP encapsulated ISDN, or proprietary protocols is marked with CS5 DSCP (use Signaling service class).
- o Signaling between call servers or soft-switches in carrier's network using SIP, SIP-T, or IP encapsulated ISUP is marked with CS5 DSCP (use Signaling service class).
- o RSVP signaling depends on the application. If RSVP signaling is "on-path" as used in IntServ, then it needs to be forwarded from the same queue (service class) and marked with the same DSCP value as application data that it is controlling. This may also apply to the "on-path" Next Steps in Signaling (NSIS) protocol.
- o If IGMP is used for multicast session control such as channel changing in IPTV systems, then IGMP packets should be marked with CS5 DSCP (use Signaling service class). When IGMP is used only for the normal multicast routing purpose, it should be marked with CS6 DSCP (use Network Control service class).

5.2. Mapping for NTP

From tests that were performed, indications are that precise time distribution requires a very low packet delay variation (jitter) transport. Therefore, we suggest that the following guidelines for Network Time Protocol (NTP) be used:

- o When NTP is used for providing high-accuracy timing within an administrator's (carrier's) network or to end users/clients, the Telephony service class should be used, and NTP packets should be marked with EF DSCP value.
- o For applications that require "wall clock" timing accuracy, the Standard service class should be used, and packets should be marked with DF DSCP.

5.3. VPN Service Mapping

"Differentiated Services and Tunnels" [RFC2983] considers the interaction of DiffServ architecture with IP tunnels of various forms. Further to guidelines provided in RFC 2983, below are additional guidelines for mapping service classes that are supported in one part of the network into a VPN connection. This discussion is limited to VPNs that use DiffServ technology for traffic differentiation.

- o The DSCP value(s) that is/are used to represent a PHB or a PHB group should be the same for the networks at both ends of the VPN tunnel, unless remarking of DSCP is done as ingress/egress processing function of the tunnel. DSCP marking needs to be preserved end to end.

- o The VPN may be configured to support one or more service classes. It is left up to the administrators of the two networks to agree on the level of traffic differentiation that will be provided in the network that supports VPN service. Service classes are then mapped into the supported VPN traffic forwarding behaviors that meet the traffic characteristics and performance requirements of the encapsulated service classes.
- o The traffic treatment in the network that is providing the VPN service needs to be such that the encapsulated service class or classes receive comparable behavior and performance in terms of delay, jitter, and packet loss and that they are within the limits of the service specified.
- o The DSCP value in the external header of the packet forwarded through the network providing the VPN service may be different from the DSCP value that is used end to end for service differentiation in the end network.
- o The guidelines for aggregation of two or more service classes into a single traffic forwarding treatment in the network that is providing the VPN service is for further study.

6. Security Considerations

This document discusses policy and describes a common policy configuration, for the use of a Differentiated Services Code Point by transports and applications. If implemented as described, it should require that the network do nothing that the network has not already allowed. If that is the case, no new security issues should arise from the use of such a policy.

It is possible for the policy to be applied incorrectly, or for a wrong policy to be applied in the network for the defined service class. In that case, a policy issue exists that the network SHOULD detect, assess, and deal with. This is a known security issue in any network dependent on policy-directed behavior.

A well-known flaw appears when bandwidth is reserved or enabled for a service (for example, voice transport) and another service or an attacking traffic stream uses it. This possibility is inherent in DiffServ technology, which depends on appropriate packet markings. When bandwidth reservation or a priority queuing system is used in a vulnerable network, the use of authentication and flow admission is recommended. To the author's knowledge, there is no known technical way to respond to an unauthenticated data stream using service that it is not intended to use, and such is the nature of the Internet.

The use of a service class by a user is not an issue when the SLA between the user and the network permits him to use it, or to use it up to a stated rate. In such cases, simple policing is used in the

Differentiated Services Architecture. Some service classes, such as Network Control, are not permitted to be used by users at all; such traffic should be dropped or remarked by ingress filters. Where service classes are available under the SLA only to an authenticated user rather than to the entire population of users, authentication and authorization services are required, such as those surveyed in [AUTHMECH].

7. Acknowledgements

The authors thank the TSVWG reviewers, David Black, Brian E. Carpenter, and Alan O'Neill for their review and input to this document.

The authors acknowledge a great many inputs, most notably from Bruce Davie, Dave Oran, Ralph Santitoro, Gary Kenward, Francois Audet, Morgan Littlewood, Robert Milne, John Shuler, Nalin Mistry, Al Morton, Mike Pierce, Ed Koehler Jr., Tim Rahrer, Fil Dickinson, Mike Fidler, and Shane Amante. Kimberly King, Joe Zebarth, and Alistair Munroe each did a thorough proofreading, and the document is better for their contributions.

8. Appendix A

8.1. Explanation of Ring Clipping

The term "ring clipping" refers to those instances where the front end of a ringing signal is altered because the bearer channel is not made available in time to carry all the audible ringing signal. This condition may occur due to a race condition between when the tone generator located in the circuit switch Exchange is turned on and when the bearer path through the IP network is enabled. To reduce ring clipping from occurring, delay of signaling path needs to be minimized. Below is a more detailed explanation.

The bearer path setup delay target is defined as the ISUP Initial Address Message (IAM) / Address Complete Message (ACM) round-trip delay. ISUP refers to ISDN User Part of Signaling System No. 7 (SS7), as defined by ITU-T. This consists of the amount of time it takes for the ISUP Initial Address Message (IAM) to leave the Transit Exchange, travel through the SS7 network (including any applicable STPs, or Signaling Transfer Points), and be processed by the End Exchange thus generating the Address Complete Message (ACM) and for the ACM to travel back through the SS7 network and return to the Transit Exchange. If the bearer path has not been set up within the soft-switch media gateway and the IP network that is performing the Transit Exchange function by the time the ACM is forwarded to the originating End Exchange, the phenomenon known as ring clipping may occur. If ACM processing within the soft-switch media gateway and delay through the IP network is excessive, it will delay the setup of the bearer path, and therefore may cause clipping of the ring tone to be heard.

The intra-exchange ISUP IAM signaling delay value should not exceed 240ms. This may include soft-switch, media gateway, router, and propagation delay on the inter-exchange data path. This value represents the threshold where ring clipping theoretically commences. It is important to note that the 240ms delay objective as presented is a maximum value. Service administrators are free to choose specific IAM delay values according to their own preferences (i.e., they may wish to set a very low mean delay objective for strategic reasons to differentiate themselves from other providers). In summary, out of the 240-ms delay budget, 200ms is allocated as cross-Exchange delay (soft-switch and media gateway) and 40ms for network delay (queuing and distance).

9. References

9.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC1349] Almquist, P., "Type of Service in the Internet Protocol Suite", RFC 1349, July 1992.
- [RFC1812] Baker, F., "Requirements for IP Version 4 Routers", RFC 1812, June 1995.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J., and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, April 1998.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, December 1998.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [RFC3662] Bless, R., Nichols, K., and K. Wehrle, "A Lower Effort Per-Domain Behavior (PDB) for Differentiated Services", RFC 3662, December 2003.

9.2. Informative References

- [AUTHMECH] Rescorla, E., "A Survey of Authentication Mechanisms", Work in Progress, September 2005.
- [QBSS] "QBone Scavenger Service (QBSS) Definition", Internet2 Technical Report Proposed Service Definition, March 2001.
- [RFC1633] Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC2581] Allman, M., Paxson, V., and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.
- [RFC2697] Heinanen, J. and R. Guerin, "A Single Rate Three Color Marker", RFC 2697, September 1999.
- [RFC2698] Heinanen, J. and R. Guerin, "A Two Rate Three Color Marker", RFC 2698, September 1999.
- [RFC2963] Bonaventure, O. and S. De Cnodder, "A Rate Adaptive Shaper for Differentiated Services", RFC 2963, October 2000.
- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, October 2000.
- [RFC2996] Bernet, Y., "Format of the RSVP DCLASS Object", RFC 2996, November 2000.
- [RFC3086] Nichols, K. and B. Carpenter, "Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification", RFC 3086, April 2001.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.
- [RFC3175] Baker, F., Iturralde, C., Le Faucheur, F., and B. Davie, "Aggregation of RSVP for IPv4 and IPv6 Reservations", RFC 3175, September 2001.

- [RFC3290] Bernet, Y., Blake, S., Grossman, D., and A. Smith, "An Informal Management Model for Diffserv Routers", RFC 3290, May 2002.
- [RFC3782] Floyd, S., Henderson, T., and A. Gurtov, "The NewReno Modification to TCP's Fast Recovery Algorithm", RFC 3782, April 2004.

Authors' Addresses

Jozef Babiarz
Nortel Networks
3500 Carling Avenue
Ottawa, Ont. K2H 8E9
Canada

Phone: +1-613-763-6098
Fax: +1-613-765-7462
EMail: babiarz@nortel.com

Kwok Ho Chan
Nortel Networks
600 Technology Park Drive
Billerica, MA 01821
US

Phone: +1-978-288-8175
Fax: +1-978-288-8700
EMail: khchan@nortel.com

Fred Baker
Cisco Systems
1121 Via Del Rey
Santa Barbara, CA 93117
US

Phone: +1-408-526-4257
Fax: +1-413-473-2403
EMail: fred@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

