



WHITE PAPER

IP CLASS OF SERVICE FOR MOBILE NETWORKS

The 3rd Generation Partnership Program (3GPP) has augmented the capabilities of General Packet Radio Service (GPRS) access to include the support of quality of service (QoS). To meet 3GPP and Global System for Mobile Communications (GSM) Association requirements, the Differentiated Services (DiffServ) architecture should be straightforward to manage and use. A minimum of packet classifications is defined to ensure that fragmentation of buffer resources does not result in packets received out of order. A simplified approach to packet classification is recommended—one that can accommodate other users of a shared 3GPP IP backbone with minimal fragmentation of buffer resources.

This paper discusses issues related to the IETF's definition of a Differentiated Services architecture for mobile IP networks supporting QoS-enabled universal mobile telecommunications service (UMTS) traffic. Classification and conditioning are analyzed and Per-Hop Behavior requirements are discussed.

SUMMARY

Four different classes of QoS have been defined by 3GPP: conversational, streaming, interactive, and background. Radio access-bearer (RAB) service attributes, which can be used to differentiate the management of radio access bearers within the UMTS Terrestrial Radio Access Network (UTRAN), have also been defined.

These attributes include:

- Traffic class
- Maximum bit rate
- Guaranteed bit rate
- Delivery order
- Maximum Service Data Unit (SDU) size
- SDU format information
- SDU error ratio
- Residual bit error ratio
- Delivery of erroneous SDUs
- Transfer delay
- Traffic handling priority (THP)
- Allocation and retention priority
- Source statistics descriptor
- Signaling indication

The 3GPP does not define how to translate or map any of these attributes for the backbone network. However, if the backbone is IP-based, 3GPP specifies that IETF's DiffServ will be used. Service attributes to DiffServ marking translation and mapping will occur in the network edge or transport edge, using the serving GPRS support node (SGSN), gateway GPRS support node (GGSN), or Radio Network Controller (RNC) to establish a radio access bearer between the handset and base station over the IP network. In addition, 3GPP recognizes that the backbone network service is not specific to UMTS (for example, it may additionally be used for other purposes). In such deployments, it is important to understand that the QoS architecture implemented in the backbone network will be shared between 3GPP flows and other users of the shared backbone, including possible IT, data communication network, and customer VPN traffic.

Importantly, DiffServ definitions only lead to packet-handling discrimination when congestion is experienced on a particular interface. Because in many instances this will be the exception rather than the rule, the effort to micro-define the particular differentiated service behavior should be minimized.

The following is a simple DiffServ architecture for an IP backbone network supporting 3GPP traffic.

SCALING DIFFERENTIATED SERVICES

While 3GPP recognizes that the mapping of UMTS QoS to Differentiated Services Code Point (DSCP) is controlled by individual operators, 3GPP 23.207 provides more detail on the DiffServ edge functions supported by the GGSN. The GGSN is compatible with the capabilities defined for a DiffServ boundary node, supporting traffic classification and conditioning.

RFC 2475 describes both traffic classification and traffic conditioning for a DiffServ architecture defining classification of traffic profiles, metering, marking, discarding, and shaping rules that are applied by the boundary node. Although the congestion conditions where DiffServ architecture results in discriminatory packet management may be rare, it is important to define an architecture that will be stable under severe overload conditions. In order to scale DiffServ-enabled networks, RFC 3086 describes how it is important for per-hop behaviors (PHBs) to be defined so that their characteristics do not depend on the traffic volume of the associated Behavior Aggregate on a router's ingress link or on a particular path through the DiffServ domain taken by the packets.

Specifically, different streams of traffic that belong to the same traffic aggregate merge and split as they traverse the network. If the properties of per-domain behavior using a particular PHB hold regardless of how the temporal characteristics of the marked traffic aggregate change as it traverses the domain, then that per-domain behavior can scale.

The definitions for PHBs and classification schemes for managing UMTS traffic should take into account such recommendations. Although single-link DiffServ domains might exist, per-domain behavior that is invariant with the network size is clearly desirable. IETF has defined informational DiffServ traffic conditioners. The inclusion of rate information in both these single-rate and two-rate three-color marker classifiers ensures scalability of solutions built using such techniques.

GSM ASSOCIATION AND DIFFERENTIATED SERVICES

Because IP traffic can traverse different networks (for example, when a user is roaming), the GSM Association has included guidelines on the operation of inter-Public Land Mobile Network (PLMN) backbone networks which impact the ability to transparently support DiffServ in the Visited PLMN (VPLMN) and Home PLMN networks. GSMA Permanent Reference Document (PRD) IR.34 defines guidelines for inter-PLMN backbone networks.

The GSM Association defines how to map the four traffic classifications defined by 3GPP on to DSCPs as well as the QoS requirements to be met by the GPRS roaming exchange service providers for each type of traffic, as shown in Table 1.

Table 1. QoS Mapping for 3G Data Roaming from GSMA PRD IR.34

3GPP QoS Information				QoS Requirement on GPRS Roaming Exchange (GRX)				
Traffic Class	THP	DiffServ PHB	DSCP	Max Delay	Max Jitter	Packet Loss	SDU Error Ratio	Service Example
Conversational	–	EF	101110	20 ms	5 ms	0.5%	10 ⁻⁶	VoIP, video-conferencing
Streaming	–	AF4 ₁	100010	40 ms	5 ms	0.5%	10 ⁻⁶	Audio/video streaming
Interactive	1	AF3 ₁	011010	250 ms	–	0.1%	10 ⁻⁸	Transactional services
	2	AF2 ₁	010010	300 ms	–	0.1%	10 ⁻⁸	Web browsing
	3	AF1 ₁	001010	350 ms	–	0.1%	10 ⁻⁸	Telnet
Background	–	BE	000000	400 ms	–	0.1%	10 ⁻⁸	E-mail download

RFC 2597 specifies that the Assured Forwarding classes are defined so that each DiffServ-enabled node allocates a certain amount of forwarding resources (such as buffer space and bandwidth) to the behavior aggregate. When these resources are exceeded, the traffic conditioning is defined to classify out-of-contract flows accordingly (for example, using a three-color scheme). According to GSMA, six classes of service are strictly required, but the interactive class contains three subclasses that can be aggregated into a single class. Hence, the minimum requirement is for four classes.

Bandwidth allocation is one of the mechanisms to support the different classes of service. For example, GSM Association PACKET Doc 035/03 describes how separate bandwidth portions for each class may be allocated. For example, an E3 connection may be defined to comprise 35 percent conversational, 30 percent streaming, 20 percent interactive, 10 percent background, and the rest is allocated for network-management traffic.

MANAGING OUT-OF-CONTRACT TRAFFIC

Traffic levels that are beyond what has been agreed to are referred to as out-of-contract traffic and should be defined for each class to make sure that all traffic use the same out-of-contract behavior in that class. If the traffic in one class exceeds what the GPRS/UMTS operator has allocated in the CoS traffic contract, there are various options which can be performed. The conversational traffic class is based on the Expedited Forwarding PHB. Accordingly, to meet the low-loss objective of Expedited Forwarding traffic, a node should be configured to drop out-of-contract Expedited Forwarding. Best-effort traffic class cannot be remarked and so will presumably be dropped should traffic exceed allocated resources.

The issue of correct management of assured forwarding PHBs used to transport the streaming and interactive classes requires careful attention. The streaming traffic class will likely carry predominantly User Datagram Protocol/Real-Time Transport Protocol (UDP/RTP) traffic and interactive traffic class will likely carry mainly TCP traffic. In particular, the benefits of re-marking streaming traffic should be analyzed. Because a re-marked flow may lead to different queuing behavior for the re-marked queue, out-of-order packets may result as the in-contract and out-of-contract packets are managed independently. Hence, for streaming traffic, it is strongly recommended that re-marking not be performed on the streaming traffic classes.

Conversely, for interactive traffic classes, out-of-contract packets should be marked as such (for example, using a single rate or two rate three-color marker scheme that modifies the drop priority bits). The remaining issue is which per-hop behavior to allocate to the re-marked Assured Forwarding traffic. It is proposed that out-of-contract AF4x traffic be dropped and out-of-contract AF3/2/1x traffic be marked for queuing with Weighted Random Early Discard defined to control TCP behavior.

DIFFERENTIATED SERVICE TO SUPPORT UMTS TRAFFIC CLASSES

Proposal Classification Scheme

First consider the Behavior Aggregate classification. On the uplink, the GPRS gateway has done a DiffServ MicroFlow classification or a CoS mapping based on the RAB Service Attribute. Other gateways (for example, Cisco Content Services Gateway [CSG] or Cisco Service Selection Gateway [SSG]) may have also performed classification or reclassification according to a per-subscriber policy. The IP network elements will be able to classify based on a behavior aggregate (DSCP).

CONVERSATIONAL TRAFFIC CLASS

The conversational traffic class is marked as Expedited Forwarding. To preserve the low-latency queuing behavior, out-of-contract conversational class is defined to be dropped. To accommodate the burst nature of conversational traffic, metering with token bucket algorithms can be used to determine whether traffic is out-of-contract.

Using the GSM Association Packet Document as an example, a single-rate three-color marking (srTCM) scheme is proposed for assured forwarding marking.

STREAMING TRAFFIC CLASS

Aligned with GSM Association PRD IR.34, streaming traffic class is marked as Assured Forwarding (AF4x). Video traffic, because of its usually limited burst behavior and large packet size, is more problematic to manage than conversational Expedited Forwarding voice. A better option is a single- or dual-rate three-color marker (sr/drTCM) classification for streaming traffic which is able to meter in-contract, exceeded-contract, and violated-contract traffic. sr/drTCM at a DiffServ Assured Forwarding-compliant edge device would drop any out-of-traffic streaming traffic and pass any in-contract traffic.

This scheme assumes that the source of the streaming traffic does not include any additional capability to further classify streaming packets. According to the application, the capability of the streaming source may be augmented to allow differentiated coloring of packets. For example, a server providing a combined audio and video stream may mark the audio portion as AF41 and the video portion as AF42 or may mark MPEG I, F, and B frames with AF41, AF42, and AF43. In this case, traffic marked with different drop precedences will be policed at the ingress and dropped at the egress accordingly. Some streaming applications are expected to resend dropped packets. Policies shall then be adapted accordingly, to better manage the network resources.

INTERACTIVE TRAFFIC CLASS

Interactive traffic classes use the remaining Assured Forwarding classifications. According to GSM Association PRD IR.34, the difference between AF3x, AF2x, and AF1x traffic classes is the priority of the queuing behavior and the maximum buffer size, which is defined to accommodate the high delay possibilities for the lower-priority traffic. However, the delay bounds are very artificial because the interface buffering requirements at each router become excessive.

Reuse of the srTCM scheme to mark traffic from each of these classes has been proposed:

- AFx1—In-contract interactive traffic class
- AFx2—Exceeded-contract interactive traffic class
- AFx3—Violated-contract interactive traffic class

BACKGROUND TRAFFIC

Background traffic management must be defined. Here the recommendations in IR.34 are not the best option. This is because background traffic, at least in the uplink case, has been transported across the UMTS radio interface. Because radio resources have already been used to put the packet into the IP network and these resources are relatively expensive compared to resources in the backbone IP network, such packets should not be marked “best effort” but should instead use the AF1x marking.

To allow for consistent marking of client-server and peer-to-peer best effort services, use of AF1x for both uplink and down-link traffic is suggested. So as a modification of the DiffServ marking proposed by GSMA IR.34, background traffic would not use a default class but the lowest Assured Forwarding class, AF1x. Within the AF1x, you may mix the lowest-importance interactive traffic marked as AF1x as IR.34 suggests with background traffic. Then you can configure adequately the drop preference in the AF1x class to distinguish interactive from background.

- AF11—In-contract, low-importance interactive traffic class
- AF12—Background traffic class and exceeded-contact, low-importance interactive class
- AF13—Violated-contract, low-level interactive and background traffic classes

Default Class can obviously be used in the IP backbone network for other, non-radio traffic.

MULTICAST TRAFFIC

The 3GPP has defined multicast traffic as being classified as either background or streaming class. Obviously, such traffic will receive RAB service attributes that can be mapped and translated to a DSCP in that class of traffic. However, according to RFC2475, IP multicast traffic should be isolated from unicast traffic because of its specific behavior. One means is to use a separate set of code points for multicast packets.

However, draft-baker-diffserv-basic classes classify some multicast traffic as unicast traffic if they are conveying an equivalent service, for example, broadcast video or audio applications.

PROPOSED PER-HOP BEHAVIOR

Now that the classification has been defined for the differentiated service boundary nodes of the IP backbone network, the PHB for each behavior aggregate will be defined.

Conversational Traffic Class

The Expedited Forwarding behavior aggregate makes use of Low Latency Queuing (LLQ). LLQ brings strict-priority queuing to Class-Based Weighted Fair Queuing (CBWFQ). Strict-priority queuing allows delay-sensitive data such as voice to be dequeued and sent first (before packets in other queues are dequeued), giving delay-sensitive data preferential treatment over other traffic.

Typically, configuration of traffic management will include a bandwidth argument that gives maximum bandwidth for this behavior aggregate. This parameter is used to specify the maximum amount of bandwidth allocated for packets belonging to the class. The bandwidth parameter both guarantees bandwidth to the priority class and restrains the flow of packets from the priority class. When congestion occurs, traffic destined for the priority queue is metered to ensure that the bandwidth allocation configured for the class to which the traffic belongs is not exceeded. If congestion occurs, when the allocated bandwidth for that strict-priority queue is exceeded, policing is used to drop packets.

Streaming Traffic Class

When only AF41 is defined, QoS configurations will use CBWFQ with tail drop. The AF41 class is characterized according to the guaranteed bandwidth delivered to the class during congestion, the weight of the class, and also the maximum queue limit for that class. After a queue has reached its configured queue limit, queuing of additional packets to the class will be configured to cause tail drop. The maximum queue limit should be set according to the amount of streaming traffic class delay allocated proportionately to each node in the network.

The Weighted Random Early Detection (WRED) mechanism may also be used at egress for progressive dropping to enable passing traffic if possible. If traffic is marked with multiple drop precedences (AF41, AF42, AF43), all code points in that class will be managed identically (with no fragmentation of buffer resources), but in periods of congestion the PHB uses a congestion-avoidance mechanism (WRED). WRED will selectively discard lower-priority traffic, based on drop preference, when the interface begins to get congested and provide differentiated performance characteristics within and between Assured Forwarding classes of service.

WRED will randomly and selectively (based on DSCP) drop packets prior to periods of high congestion. For streaming service, carrying predominantly RTP/UDP traffic will usually not enable a decrease in the source transmission rate. However, WRED provides separate thresholds and weights for different precedences, allowing you to provide different QoS for packet dropping based on different traffic types. For streaming AF4x class, AF43 traffic may be dropped more frequently than AF42, and AF42 traffic may be dropped more frequently than AF41 traffic during periods of congestion.

Interactive Traffic Class

Interactive traffic class will mainly contain TCP traffic. TCP congestion-avoidance techniques can therefore be used. CBWFQ can be used as with the streaming traffic class. However, for the interactive traffic class in congestion scenarios after the queue has reached its configured queue limit, queuing of additional packets to the class will be configured to cause packet drop. Packet drop corresponds to using the inherent congestion-avoidance techniques of TCP and involves using WRED drop policy.

Background Traffic Class

As described previously, because of the significant expenditure of radio resources involved in transmitting a packet corresponding to a background traffic class from the user equipment to the core network, the use of the AF1x traffic class to classify the background traffic class has been proposed. The same PHB as defined for the interactive traffic classes applies.

Best-Effort Traffic

Typically, when configuring CBWFQ, the sum of all bandwidth allocations on an interface should not exceed a certain percent of the total available interface bandwidth (usually 75 percent). The remaining 25 percent is used for other overhead, including Layer 2 overhead, routing traffic, and best-effort traffic. In this instance, best-effort traffic may correspond to other uses of a shared 3GPP backbone network (for example, one that is used to simultaneously support a mobile operator's internal e-mail service).

SUMMARY

This white paper has analyzed the requirements for defining the operation of a Differentiated Services-enabled 3GPP backbone core IP network which is able to support 3GPP QoS requirements with tight service-level requirements. While 3GPP defines many QoS parameters, these are defined exclusively for use as radio access-bearer service attributes, used to enable differentiated management of RABs within the UMTS Terrestrial Radio Access Network (UTRAN).

Consequently, the latest recommendations from the GSM Association are used to define a core IP network with four classifications and associated per-hop behavior:

- The UMTS conversational traffic class is policed to ensure that if the corresponding allocated bandwidth is exceeded, out-of-contract traffic will be dropped. In-contract traffic is dequeued using LLQ, providing expedited forwarding per-hop behavior.
- Because of the large packet sizes involved and the likely bursty nature of the service, it is recommended that the UMTS streaming traffic class use a single- or dual-rate policer scheme which allows in-contract traffic to be carried while ensuring that traffic that exceeds or violates the allocated bandwidth is dropped. Depending upon the codec and streaming application, more precise classification, conditioning, and behavior can be defined. The streaming traffic class will be treated as Assured Forwarding, with either tail drop or WRED when congestion occurs at the egress interface.
- Recognizing that expensive radio resources will have been used by the time UMTS background traffic has reached the core network, the UMTS background traffic class should be treated as Assured Forwarding traffic. In this way, a single behavior aggregate can be configured for managing the combination of UMTS interactive and background traffic classes. Because of the principal use of TCP for these services, WRED is intelligently used to trigger TCP congestion avoidance rather than simple tail-drop policies. The use of WRED will be perceived by users as similar to simply allocating more buffer resources to the traffic.
- In the case of dropped packets, TCP congestion avoidance will slow down its transmission rate of packets into the network and increase. In the case of more latency corresponding to increased buffer size, the TCP throughput can be shown to decrease inversely related to the roundtrip time.
- Best-effort traffic is limited to non-UMTS user-plane traffic.
- The same may not be true for GPRS-specific core network elements (Radio Network Controller [RNC], serving GPRS support node [SGSN], and gateway GPRS support node [GGSN]), but this discussion only includes the backbone core IP network requirements.

The simplicity of managing four classes of service can help ensure easy interoperability with other networks (for example, with IP/MPLS-enabled networks where only three bits are available for carrying CoS information). In addition, end-to-end consistent marking can be provided across an inter-PLMN backbone, which will allow GPRS Roaming Exchange (GRX) operators to reuse their existing IETF-defined packet-classification schemes. Finally, it is evident that when making use of CBWFQ, the associated per-hop behavior will be stable and can be applied consistently on all nodes within the network.

FOR MORE INFORMATION

GSM Association Inter-PLMN Backbone Guidelines: <http://www.gsmworld.com/documents/ireg/ir34.pdf> Cf. to Chapter 7.4 “IP QoS Definitions for GPRS Release ‘99 and UMTS”

RFC2474: Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers: <http://www.ietf.org/rfc/rfc2474.txt>

RFC2475: An Architecture for Differentiated Services: <http://www.ietf.org/rfc/rfc2475.txt>

RFC2597: Assured Forwarding PHB Group: <http://www.ietf.org/rfc/rfc2597.txt>

RFC3246: An Expedited Forwarding PHB (Per-Hop Behavior): <http://www.ietf.org/rfc/rfc3246.txt>

RFC3247: Supplemental Information for the New Definition of the EF PHB (Expedited Forwarding Per-Hop Behavior): <http://www.ietf.org/rfc/rfc3247.txt>

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR
Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia
Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R)
He/LW8482 05/05

