                     Assured Forwarding PHB Group

Status of this Memo

Copyright Notice

Abstract

   This document defines a general use Differentiated Services (DS)
   [Blake] Per-Hop-Behavior (PHB) Group called Assured Forwarding (AF).
   The AF PHB group provides delivery of IP packets in four
   independently forwarded AF classes. Within each AF class, an IP
   packet can be assigned one of three different levels of drop
   precedence. A DS node does not reorder IP packets of the same
   microflow if they belong to the same AF class.

1. Purpose and Overview

   There is a demand to provide assured forwarding of IP packets over
   the Internet. In a typical application, a company uses the Internet
   to interconnect its geographically distributed sites and wants an
   assurance that IP packets within this intranet are forwarded with
   high probability as long as the aggregate traffic from each site does
   not exceed the subscribed information rate (profile). It is
   desirable that a site may exceed the subscribed profile with the
   understanding that the excess traffic is not delivered with as high
   probability as the traffic that is within the profile. It is also

important that the network does not reorder packets that belong to
the same microflow, as defined in [Nichols], no matter if they are in
or out of the profile.

Assured Forwarding (AF) PHB group is a means for a provider DS domain
to offer different levels of forwarding assurances for IP packets
received from a customer DS domain. Four AF classes are defined,
where each AF class is in each DS node allocated a certain amount of
forwarding resources (buffer space and bandwidth). IP packets that
wish to use the services provided by the AF PHB group are assigned by
the customer or the provider DS domain into one or more of these AF
classes according to the services that the customer has subscribed
to. Further background about this capability and some ways to use it
may be found in [Clark].

Within each AF class IP packets are marked (again by the customer or
the provider DS domain) with one of three possible drop precedence
values. In case of congestion, the drop precedence of a packet
determines the relative importance of the packet within the AF class.
A congested DS node tries to protect packets with a lower drop
precedence value from being lost by preferably discarding packets
with a higher drop precedence value.

In a DS node, the level of forwarding assurance of an IP packet thus
depends on (1) how much forwarding resources has been allocated to
the AF class that the packet belongs to, (2) what is the current load
of the AF class, and, in case of congestion within the class, (3)
what is the drop precedence of the packet.

For example, if traffic conditioning actions at the ingress of the
provider DS domain make sure that an AF class in the DS nodes is only
moderately loaded by packets with the lowest drop precedence value
and is not overloaded by packets with the two lowest drop precedence
values, then the AF class can offer a high level of forwarding
assurance for packets that are within the subscribed profile (i.e.,
marked with the lowest drop precedence value) and offer up to two
lower levels of forwarding assurance for the excess traffic.

This document describes the AF PHB group. An otherwise DS-compliant
node is not required to implement this PHB group in order to be
considered DS-compliant, but when a DS-compliant node is said to
implement an AF PHB group, it must conform to the specification in
this document.

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

2. The AF PHB Group

   Assured Forwarding (AF) PHB group provides forwarding of IP packets
   in N independent AF classes. Within each AF class, an IP packet is
   assigned one of M different levels of drop precedence. An IP packet
   that belongs to an AF class i and has drop precedence j is marked
   with the AF codepoint AFij, where 1 <= i <= N and 1 <= j <= M.
   Currently, four classes (N=4) with three levels of drop precedence in
   each class (M=3) are defined for general use. More AF classes or
   levels of drop precedence MAY be defined for local use.

   A DS node SHOULD implement all four general use AF classes. Packets
   in one AF class MUST be forwarded independently from packets in
   another AF class, i.e., a DS node MUST NOT aggregate two or more AF
   classes together.

   A DS node MUST allocate a configurable, minimum amount of forwarding
   resources (buffer space and bandwidth) to each implemented AF class.
   Each class SHOULD be serviced in a manner to achieve the configured
   service rate (bandwidth) over both small and large time scales.

   An AF class MAY also be configurable to receive more forwarding
   resources than the minimum when excess resources are available either
   from other AF classes or from other PHB groups. This memo does not
   specify how the excess resources should be allocated, but
   implementations MUST specify what algorithms are actually supported
   and how they can be parameterized.

   Within an AF class, a DS node MUST NOT forward an IP packet with
   smaller probability if it contains a drop precedence value p than if
   it contains a drop precedence value q when p < q. Note that this
   requirement can be fulfilled without needing to dequeue and discard
   already-queued packets.

   Within each AF class, a DS node MUST accept all three drop precedence
   codepoints and they MUST yield at least two different levels of loss
   probability. In some networks, particularly in enterprise networks,
   where transient congestion is a rare and brief occurrence, it may be
   reasonable for a DS node to implement only two different levels of
   loss probability per AF class. While this may suffice for some
   networks, three different levels of loss probability SHOULD be
   supported in DS domains where congestion is a common occurrence.

   If a DS node only implements two different levels of loss probability
   for an AF class x, the codepoint AFx1 MUST yield the lower loss
   probability and the codepoints AFx2 and AFx3 MUST yield the higher
   loss probability.

A DS node MUST NOT reorder AF packets of the same microflow when they
belong to the same AF class regardless of their drop precedence.
There are no quantifiable timing requirements (delay or delay
variation) associated with the forwarding of AF packets.

The relationship between AF classes and other PHBs is described in
Section 7 of this memo.

The AF PHB group MAY be used to implement both end-to-end and domain
edge-to-domain edge services.

3. Traffic Conditioning Actions

A DS domain MAY at the edge of a domain control the amount of AF
traffic that enters or exits the domain at various levels of drop
precedence. Such traffic conditioning actions MAY include traffic
shaping, discarding of packets, increasing or decreasing the drop
precedence of packets, and reassigning of packets to other AF
classes. However, the traffic conditioning actions MUST NOT cause
reordering of packets of the same microflow.

4. Queueing and Discard Behavior

This section defines the queueing and discard behavior of the AF PHB
group. Other aspects of the PHB group's behavior are defined in
Section 2.

An AF implementation MUST attempt to minimize long-term congestion
within each class, while allowing short-term congestion resulting
from bursts. This requires an active queue management algorithm. An
example of such an algorithm is Random Early Drop (RED) [Floyd].
This memo does not specify the use of a particular algorithm, but
does require that several properties hold.

An AF implementation MUST detect and respond to long-term congestion
within each class by dropping packets, while handling short-term
congestion (packet bursts) by queueing packets. This implies the
presence of a smoothing or filtering function that monitors the
instantaneous congestion level and computes a smoothed congestion
level. The dropping algorithm uses this smoothed congestion level to
determine when packets should be discarded.

The dropping algorithm MUST be insensitive to the short-term traffic
characteristics of the microflows using an AF class. That is, flows
with different short-term burst shapes but identical longer-term
packet rates should have packets discarded with essentially equal
probability. One way to achieve this is to use randomness within the
dropping function.

The dropping algorithm MUST treat all packets within a single class
and precedence level identically. This implies that for any given
smoothed congestion level, the discard rate of a particular
microflow's packets within a single precedence level will be
proportional to that flow's percentage of the total amount of traffic
passing through that precedence level.

The congestion indication feedback to the end nodes, and thus the
level of packet discard at each drop precedence in relation to
congestion, MUST be gradual rather than abrupt, to allow the overall
system to reach a stable operating point. One way to do this (RED)
uses two (configurable) smoothed congestion level thresholds. When
the smoothed congestion level is below the first threshold, no
packets of the relevant precedence are discarded. When the smoothed
congestion level is between the first and the second threshold,
packets are discarded with linearly increasing probability, ranging
from zero to a configurable value reached just prior to the second
threshold. When the smoothed congestion level is above the second
threshold, packets of the relevant precedence are discarded with 100%
probability.

To allow the AF PHB to be used in many different operating
environments, the dropping algorithm control parameters MUST be
independently configurable for each packet drop precedence and for
each AF class.

Within the limits above, this specification allows for a range of
packet discard behaviors. Inconsistent discard behaviors lead to
inconsistent end-to-end service semantics and limit the range of
possible uses of the AF PHB in a multi-vendor environment. As
experience is gained, future versions of this document may more
tightly define specific aspects of the desirable behavior.

5. Tunneling

When AF packets are tunneled, the PHB of the tunneling packet MUST
NOT reduce the forwarding assurance of the tunneled AF packet nor
cause reordering of AF packets belonging to the same microflow.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.

fastcase
Smarter legal research.