Network Working Group
Request for Comments: 2475
Category: Informational

S. Blake
Torrent Networking Technologies
D. Black
EMC Corporation
M. Carlson
Sun Microsystems
E. Davies
Nortel UK
Z. Wang
Bell Labs Lucent Technologies
W. Weiss
Lucent Technologies
December 1998

An Architecture for Differentiated Services

Status of this Memo

Copyright Notice

Abstract

This document defines an architecture for implementing scalable service differentiation in the Internet. This architecture achieves scalability by aggregating traffic classification state which is conveyed by means of IP-layer packet marking using the DS field [DSFIELD]. Packets are classified and marked to receive a particular per-hop forwarding behavior on nodes along their path. Sophisticated classification, marking, policing, and shaping operations need only be implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service provisioning policies which govern how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network. A wide variety of services can be implemented on top of these building blocks.

Table of Contents

1. Introduction

1.1 Overview

This document defines an architecture for implementing scalable service differentiation in the Internet. A "Service" defines some significant characteristics of packet transmission in one direction across a set of one or more paths within a network. These

characteristics may be specified in quantitative or statistical terms of throughput, delay, jitter, and/or loss, or may otherwise be specified in terms of some relative priority of access to network resources. Service differentiation is desired to accommodate heterogeneous application requirements and user expectations, and to permit differentiated pricing of Internet service.

This architecture is composed of a number of functional elements implemented in network nodes, including a small set of per-hop forwarding behaviors, packet classification functions, and traffic conditioning functions including metering, marking, shaping, and policing. This architecture achieves scalability by implementing complex classification and conditioning functions only at network boundary nodes, and by applying per-hop behaviors to aggregates of traffic which have been appropriately marked using the DS field in the IPv4 or IPv6 headers [DSFIELD]. Per-hop behaviors are defined to permit a reasonably granular means of allocating buffer and bandwidth resources at each node among competing traffic streams. Per-application flow or per-customer forwarding state need not be maintained within the core of the network. A distinction is maintained between:

- the service provided to a traffic aggregate,
- the conditioning functions and per-hop behaviors used to realize services,
- the DS field value (DS codepoint) used to mark packets to select a per-hop behavior, and
- the particular node implementation mechanisms which realize a per-hop behavior.

Service provisioning and traffic conditioning policies are sufficiently decoupled from the forwarding behaviors within the network interior to permit implementation of a wide variety of service behaviors, with room for future expansion.

This architecture only provides service differentiation in one direction of traffic flow and is therefore asymmetric. Development of a complementary symmetric architecture is a topic of current research but is outside the scope of this document; see for example [EXPLICIT].

Sect. 1.2 is a glossary of terms used within this document. Sec. 1.3 lists requirements addressed by this architecture, and Sec. 1.4 provides a brief comparison to other approaches for service differentiation. Sec. 2 discusses the components of the architecture

in detail. Sec. 3 proposes guidelines for per-hop behavior specifications. Sec. 4 discusses interoperability issues with nodes and networks which do not implement differentiated services as defined in this document and in [DSFIELD]. Sec. 5 discusses issues with multicast service delivery. Sec. 6 addresses security and tunnel considerations.

1.2 Terminology

This section gives a general conceptual overview of the terms used in this document. Some of these terms are more precisely defined in later sections of this document.

Behavior Aggregate (BA)	a DS behavior aggregate.
BA classifier	a classifier that selects packets based only on the contents of the DS field.
Boundary link	a link connecting the edge nodes of two domains.
Classifier	an entity which selects packets based on the content of packet headers according to defined rules.
DS behavior aggregate	a collection of packets with the same DS codepoint crossing a link in a particular direction.
DS boundary node	a DS node that connects one DS domain to a node either in another DS domain or in a domain that is not DS-capable.
DS-capable	capable of implementing differentiated services as described in this architecture; usually used in reference to a domain consisting of DS-compliant nodes.
DS codepoint	a specific value of the DSCP portion of the DS field, used to select a PHB.
DS-compliant	enabled to support differentiated services functions and behaviors as defined in [DSFIELD], this document, and other differentiated services documents; usually used in reference to a node or device.

DS domain	a DS-capable domain; a contiguous set of nodes which operate with a common set of service provisioning policies and PHB definitions.
DS egress node	a DS boundary node in its role in handling traffic as it leaves a DS domain.
DS ingress node	a DS boundary node in its role in handling traffic as it enters a DS domain.
DS interior node	a DS node that is not a DS boundary node.
DS field	the IPv4 header TOS octet or the IPv6 Traffic Class octet when interpreted in conformance with the definition given in [DSFIELD]. The bits of the DSCP field encode the DS codepoint, while the remaining bits are currently unused.
DS node	a DS-compliant node.
DS region	a set of contiguous DS domains which can offer differentiated services over paths across those DS domains.
Downstream DS domain	the DS domain downstream of traffic flow on a boundary link.
Dropper	a device that performs dropping.
Dropping	the process of discarding packets based on specified rules; policing.
Legacy node	a node which implements IPv4 Precedence as defined in [RFC791,RFC1812] but which is otherwise not DS-compliant.
Marker	a device that performs marking.
Marking	the process of setting the DS codepoint in a packet based on defined rules; pre-marking, re-marking.
Mechanism	a specific algorithm or operation (e.g., queueing discipline) that is implemented in a node to realize a set of one or more per-hop behaviors.

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.