
DIFFERENTIATED SERVICES FOR THE INTERNET

Kalevi Kilkki



Differentiated Services for the Internet

Kalevi Kilkki

Published by:
Macmillan Technical Publishing
201 West 103rd Street
Indianapolis, IN 46290 USA

Copyright© 1999 by Macmillan Technical Publishing

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

International Standard Book Number: 1-57870-132-5

Library of Congress Catalog Card Number: 99-62122

2002 01 00 99 4 3 2 1

Interpretation of the printing code: The rightmost double-digit number is the year of the book's printing; the rightmost single-digit, the number of the book's printing. For example, the printing code 99-1 shows that the first printing of the book occurred in 1999.

*Composed in Galliard and MCPdigital by Macmillan Computer Publishing
Printed in the United States of America*

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Macmillan Technical Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

This book is designed to provide information about **Internet technologies**. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an as-is basis. The authors and Macmillan Technical Publishing shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

Feedback Information

At Macmillan Technical Publishing, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us at network-tech@mcp.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

PUBLISHER
David Dwyer

EXECUTIVE EDITOR
Linda Ratts Engelman

MANAGING EDITOR
Patrick Kanouse

ACQUISITIONS EDITOR
Karen Wachs

DEVELOPMENT EDITOR
Thomas Cirtin

PROJECT EDITOR
Theresa Wehrle

COPY EDITOR
Keith Cline

INDEXER
Larry Sweazy

**AQUISITIONS
COORDINATOR**
Jennifer Garrett

**MANUFACTURING
COORDINATOR**
Brook Farling

BOOK DESIGNER
Ann Jones

COVER DESIGNER
Karen Ruggles

**PRODUCTION TEAM
SUPERVISOR**
Tricia Flodder

PRODUCTION
Lisa England

PROOFREADER
Elise Walter

Dedication

To: Eija, Olli, and Juho

Acknowledgments

First, I would like to thank my management at Nokia Research Center, in particular Raj Bansal and Antti Ylä-Jääski, for allowing me to work on the book with minimal attention to other tasks. I would like to express my appreciation to several colleagues at Nokia, among them Jussi Ruutu, Ove Strandberg, Ravikanth Rayadurgam, Jarno Rajahalme, Pasi Väänänen, and Matti Alkula. I have consumed their time with numerous intricate questions during the past months.

I also am indebted to the technical reviewers, Shivkumar Kalyanaraman and Jon Crowcroft, for their critical comments and beneficial suggestions. I am grateful to all the experts at Macmillan Technical Publishing for editorial assistance and for keeping the project on track although the schedule and the target appeared unattainable from time to time.

Finally, I heartily thank my family for their support and patience during this project.

About the Author

Kalevi Kilkki is a principal scientist at Nokia Research Center in Burlington, Massachusetts. He holds a master of science and doctor of technology in electrical engineering from Helsinki University of Technology, Finland. His comprehensive background includes research projects related to telephone networks at Helsinki University of Technology, to ATM networks at Telecom Finland (now Sonera), and to IP networks at Nokia Research Center. The ATM research work has been carried out in connection with several European research projects, such as COST242 and Eurescom P105, while the framework for IP research has been the Differentiated Services Working Group in IETF.

During the last three years, Kalevi's main effort has been to develop feasible traffic handling mechanisms required for service differentiation in the Internet. He has given numerous presentations at international conferences about diverse topics, from the application of genetic algorithms in connection admission control to pricing models for the Internet. Dr. Kilkki also has continuously delivered lectures about traffic theory and telecommunication systems at the Helsinki University of Technology and has instructed various master's theses. Finally, one of his main hobbies, philosophy, gives a specific flavor to the treatment of these intricate topics.

About the Technical Reviewers

These reviewers contributed their considerable practical, hands-on expertise to the entire development process for *Differentiated Services for the Internet*. As the book was being written, these folks reviewed all the material for technical content, organization, and flow. Their feedback was critical to ensuring that *Differentiated Services for the Internet* fits our reader's need for the highest quality technical information.

Shivkumar Kalyanaraman is an assistant professor for the department of electrical, computer, and systems engineering at Rensselaer Polytechnic Institute in Troy, New York. He received a B.Tech degree from the Indian Institute of Technology, Madras, India, in July 1993, followed by master of science and doctoral degrees in computer and information sciences at Ohio State University in 1994 and 1997, respectively. His research interests include traffic management, multicast, Internet pricing, multimedia networking, and performance analysis of distributed systems.

Shivkumar is a co-inventor in two patents (the ERICA and OSU schemes for ATM traffic management) and has co-authored several papers, IETF drafts, and ATM forum contributions. He is a member of IEEE-CS and ACM. His World Wide Web site is located at <http://www.ecse.rpi.edu/Homepages/shivkuma>.

Jon Crowcroft is a professor of networked systems in the department of computer science, University College London, where he is responsible for a number of European and U.S. funded research projects in multimedia communications. He has been working in these areas for over 18 years. Jon graduated with a degree in physics from Trinity College, Cambridge University, in 1979. He gained his master of science in computing in 1981 and his doctoral degree in 1993.

Jon is a member of the ACM and the British Computer Society. He is a Fellow of the IEE and a senior member of the IEEE. He also is a member of the IAB and general chair for the ACM SIGCOMM. He is also on the editorial team for the ACM/IEEE Transactions on Networks. With Mark Handley, Jon is the co-author of *WWW: Beneath the Surf* (UCL Press). He also wrote *Open Distributed Systems* (UCL Press/Artech House).

Overview

Introduction

PART I BACKGROUND FOR DIFFERENTIATED SERVICES

- 1 The Target of Differentiated Services
- 2 Traffic Management Before Differentiated Services
- 3 Differentiated Services Working Group

PART II BUILDING A NETWORK DOMAIN BASED ON DIFFERENTIATED SERVICES

- 4 General Framework for Differentiated Services
- 5 Differentiation of Customer Service
- 6 Traffic Handling and Network Management
- 7 Per-Hop Behavior Groups

PART III BUILDING GLOBAL NETWORKS BASED ON DIFFERENTIATED SERVICES

- 8 Interworking Issues
- 9 Implementing Differentiated Services

References

Glossary

Index

Contents

PART I	1
1 The Target of Differentiated Services	3
1.1 <i>The Core of Differentiated Services</i>	4
1.1.1 Basic Entities of Differentiated Services	5
1.1.2 The Relationships Among the Basic Entities	8
1.2 <i>The Four Attributes of Differentiated Services</i>	9
1.2.1 Fairness	9
1.2.2 Technical Issues: Versatility and Robustness	13
1.2.3 Cost Efficiency	14
<i>Summary</i>	15
2 Traffic Management Before Differentiated Services	17
2.1 <i>Fundamental Concepts, Models, and Technologies</i>	17
2.1.1 Customer Service	18
2.1.2 Network with Services	19
2.1.3 Network Operation and Management	20
2.1.4 Traffic Handling	22
2.1.5 Traffic Models for the Internet	23
2.1.6 Technological Progress	25
2.2 <i>Traditional Telecommunication Approaches</i>	26
2.2.1 Circuit-Switched Networks	26
2.2.2 ATM Networks	30
2.2.3 Evaluation of Connection-Oriented Approaches	36
2.3 <i>The Best-Effort Approach</i>	44
2.3.1 Service Model	45
2.3.2 Basic Best-Effort Service Based on TCP	47
2.3.3 Improvements to the Basic TCP Behavior	48
2.3.4 Evaluation of the Best-Effort Approach	50
2.4 <i>Integrated Services Model</i>	53
2.4.1 Customer Service	55
2.4.2 Implementation of Integrated Services	56
2.4.3 Evaluation of the Integrated Services Model	59
2.5 <i>Targets for Differentiated Services</i>	60
2.5.1 Customer Service	60
2.5.2 Network Services	61
2.5.3 Operation and Management	61
2.5.4 Traffic Handling	62
<i>Summary</i>	62

3	<i>Differentiated Services Working Group</i>	65
3.1	<i>A Short History of Differentiated Services</i>	65
3.2	<i>The Position of the Differentiated Services Working Group</i>	69
3.3	<i>Basic Working Group Documents</i>	69
3.3.1	Introduction to Differentiated Services Model	70
3.3.2	The Differentiated Services Field in IPv4 and IPv6 Headers	72
3.3.3	Architecture for Differentiated Services	78
3.3.4	A Framework for Differentiated Services	85
	<i>Summary</i>	89
PART II		91
4	<i>General Framework for Differentiated Services</i>	93
4.1	<i>Basis for the Framework</i>	94
4.1.1	Basic Terminology	94
4.1.2	Limitations of the Differentiated Services Working Group	95
4.1.3	Target of the Framework	96
4.2	<i>Tools for Evaluating Service Models</i>	98
4.2.1	Availability of Quality	98
4.2.2	Levels of Aggregation	101
4.3	<i>Customer Service</i>	103
4.3.1	Fulfilling Consumer Expectations	104
4.3.2	Pricing Models and Predictability of Quality	104
4.4	<i>Operation and Management</i>	107
4.4.1	Predictability of Load and Destination	108
4.4.2	Service Models	109
4.4.3	A Model for Evaluating Statistical Multiplexing	111
4.4.4	Improving Statistical Multiplexing	117
4.4.5	Conclusions About Statistical Multiplexing	124
4.5	<i>Traffic Handling</i>	125
4.5.1	Urgency	125
4.5.2	Importance	127
4.5.3	Bandwidth Division (Virtual Private Networks)	130
4.5.4	Routing	132
4.5.5	Support for Adaptive Applications	133
4.6	<i>Framework for Per-Hop Behaviors</i>	134
4.6.1	Essence of Quality	134
4.6.2	Relative Scales for Importance and Urgency	135

4.6.3	Relative Scales for Bandwidth	137
4.6.4	Predictability of Quality	138
	<i>Summary</i>	<i>139</i>
5	<i>Differentiation of Customer Service</i>	<i>141</i>
5.1	<i>Service Level Agreement</i>	<i>142</i>
5.1.1	Guaranteed Connections	144
5.1.2	Leased-Line Service	145
5.1.3	Resource Sharing	147
5.1.4	Dynamic Importance	148
5.1.5	Comparison Based on Availability of Quality	149
5.2	<i>Requesting Specific Service</i>	<i>151</i>
5.2.1	Dynamic Quality or Bandwidth for Guaranteed Connections	151
5.2.2	Permanent Bandwidth Reservation Versus Permanent Share	151
5.2.3	Dynamic Share	152
5.2.4	Summary of Service Requests	152
5.3	<i>Pricing as a Tool for Controlling Traffic</i>	<i>155</i>
5.3.1	Price of Bandwidth	157
5.3.2	Price of Quality	158
5.3.3	Price of Availability	159
5.3.4	Relationship Between Bandwidth and Availability	160
5.3.5	Relationship Between Quality and Availability	161
5.3.6	The Relationship Between Quality and Bandwidth	161
5.3.7	Effect of Variable Destinations	163
5.3.8	Levels of Pricing	163
5.3.9	Variable Bit Rate	164
	<i>Summary</i>	<i>165</i>
6	<i>Traffic Handling and Network Management</i>	<i>167</i>
6.1	<i>Traffic Handling in Boundary Nodes</i>	<i>168</i>
6.1.1	Classifiers	170
6.1.2	Meters	171
6.1.3	Packet Marking	176
6.1.4	Traffic Shaping	178
6.1.5	Packet Dropping at Boundary Nodes	181
6.2	<i>Traffic-Handling Functions in Interior Nodes</i>	<i>181</i>
6.2.1	Buffering	182
6.2.2	Discarding	193
6.2.3	Feedback Information	198

6.3	<i>Functions Related to a Network Domain</i>	200
6.3.1	Routing	200
6.3.2	Resource Reservation	203
6.3.3	Network Dimensioning	204
6.3.4	Boundary Nodes Between Network Domains	207
	<i>Summary</i>	209
7	<i>Per-Hop Behavior Groups</i>	211
7.1	<i>Systematic Basis for PHB Evaluation</i>	212
7.2	<i>Class Selector PHB Group</i>	215
7.2.1	Description	215
7.2.2	Position in the Framework	217
7.2.3	Required Technical Tools	222
7.2.4	Evaluation of Attributes	223
7.3	<i>Expedited Forwarding PHB</i>	224
7.3.1	Description	224
7.3.2	Position in the Framework	226
7.3.3	Required Technical Tools	230
7.3.4	Evaluation of Attributes	232
7.4	<i>Assured Forwarding PHB Group</i>	232
7.4.1	Description	232
7.4.2	Position of AF-PHB Group in the Framework	233
7.4.3	Required Technical Tools	243
7.4.4	Evaluation of Attributes	245
7.5	<i>Dynamic RT/NRT PHB Group</i>	245
7.5.1	Description	245
7.5.2	Position of DRT-PHB Group in the Framework	246
7.5.3	Required Technical Tools	253
7.5.4	Evaluation of Attributes	255
	<i>Summary</i>	256
PART III		257
8	<i>Interworking Issues</i>	259
8.1	<i>Interworking Among Differentiated Services Models</i>	259
8.1.1	Class Selector PHB Group with Other PHB Groups	260
8.1.2	Expedited Forwarding Versus Assured Forwarding	263
8.1.3	DRT-PHB Group Compared with AF- and EF-PHB Groups	268
8.1.4	Summary of PHB groups	271
8.2	Interworking with Other Internet Schemes	272

8.2.1	Best-Effort Service	272
8.2.2	Integrated Services and RSVP	274
8.2.3	Multiprotocol Label Switching	278
8.3	<i>Interworking with Non-IP Networks</i>	280
8.3.1	Asynchronous Transfer Mode	280
8.3.2	IEEE 802.1p	283
8.3.3	Wireless Networks	286
8.4	<i>Multicast Services</i>	290
	<i>Summary</i>	292
9	<i>Implementing Differentiated Services</i>	295
9.1	<i>Network and Traffic Models</i>	296
9.2	<i>Improving Fairness Using an AF-PHB Group</i>	298
9.2.1	Traffic Model	299
9.2.2	Implementation	300
9.2.3	Performance Evaluation	301
9.2.4	Possible Improvements	304
9.3	<i>Virtual Private Networks by Using an EF-PHB</i>	305
9.3.1	Traffic Model	305
9.3.2	Implementation	307
9.3.3	Performance Evaluation	308
9.3.4	Possible Improvements	310
9.4	<i>Service Differentiation with Three AF PHB Groups</i>	311
9.4.1	Implementation	311
9.4.2	Traffic Model	313
9.4.3	Network Dimensioning	314
9.4.4	Possible Improvements	317
9.5	<i>Total Service on the Basis of a DRT-PHB Group</i>	317
9.5.1	Implementation	318
9.5.2	Traffic Model	319
9.5.3	Performance Evaluation	320
9.5.4	Possible Improvements	322
	<i>Summary</i>	323
	<i>References</i>	327
	<i>Glossary</i>	333
	<i>Index</i>	345

Figures

Figure 1.1 Main entities of the Internet and the relationships among them p. xxx

Figure 2.1 The main building blocks are boundary nodes (A), interior nodes (B), and customer equipment (C) p. xxx

Figure 2.2 Virtual circuits and virtual paths in an ATM network p. xxx

Figure 2.3 Part of a fictitious ATM service bill p. xxx

Figure 2.4 The difference in real service quality of three service classes p. xxx

Figure 2.5 Sawtooth pattern of a TCP connection p. xxx

Figure 3.1 Discussion activity on the integrated services mailing list p. xxx

Figure 3.2 Construction of Differentiated Services p. xxx

Figure 3.3 Per-hop Behavior (PHB) p. xxx

Figure 3.4 The main building blocks of Differentiated Services p. xxx

Figure 3.5 Per-Hop Behavior Group with two PHB classes p. xxx

Figure 3.6 The structure of the Differentiated Services field p. xxx

Figure 3.7 Basic elements of a Differentiated Services network p. xxx

Figure 3.8 Main entities of the Internet and the relationships between them p. xxx

Figure 3.9 Packet classifier and traffic conditioning according to the architecture document p. xxx

Figure 3.10 Three approaches to sharing link resources p. xxx

Figure 4.1 Availability of quality for a packet loss ratio p. xxx

Figure 4.2 Availability of quality with a theoretical curve and two practical examples p. xxx

Figure 4.3 Fairness aspects on different levels of aggregates p. xxx

Figure 4.4 Predictability of traffic and destination p. xxx

Figure 4.5 A network model p. xxx

Figure 4.6 Dimensioning rule based on mean (M) and variance (V) of a prediction distribution p. xxx

Figure 4.7 Traffic from node A and to three nodes: B, C, and D p. xxx

Figure 4.8 The effect of traffic shaping p. xxx

Figure 4.9 Availability of quality with and without reservations p. xxx

Figure 4.10 Availability of quality with delay criterion p. xxx

Figure 4.11	Selection of packets to be discarded	p. xxx
Figure 4.12	VPN categories	p. xxx
Figure 4.13	A VPN network with a possible conflict	p. xxx
Figure 4.14	Relative scales of importance and urgency	p. xxx
Figure 4.15	Changing of the order of two PHBs in relation to packet loss ratio	p. xxx
Figure 4.16	A PHB system with importance and bandwidth scales	p. xxx
Figure 5.1	Four basic approaches for SLA	p. xxx
Figure 5.2	Service model for guaranteed connections	p. xxx
Figure 5.3	Service model for leased line service	p. xxx
Figure 5.4	Service model for resource sharing	p. xxx
Figure 5.5	Service model for resource sharing with dynamic importance	p. xxx
Figure 5.6	Quality function for guaranteed service and resource sharing service	p. xxx
Figure 5.7	An approximate relation between charged unit and costs	p. xxx
Figure 5.8	A tentative relationship between bandwidth and price	p. xxx
Figure 5.9	A tentative relationship between delay variation and price	p. xxx
Figure 5.10	A tentative relationship between availability and price	p. xxx
Figure 5.11	Bandwidth as a function of availability for fixed price and quality	p. xxx
Figure 5.12	Quality as a function of availability for fixed price and bandwidth	p. xxx
Figure 5.13	Bandwidth as a function of quality for fixed price and availability	p. xxx
Figure 5.14	A tentative relationship between traffic variability and available bandwidth	p. xxx
Figure 6.1	Differentiated Services from a customer's perspective	p. xxx
Figure 6.2	The main building blocks of traffic handling	p. xxx
Figure 6.3	Desired measuring principle for guaranteed service and resource-sharing service	p. xxx
Figure 6.4	Conforming and non-conforming flows measured by token bucket and by exponential moving average	p. xxx
Figure 6.5	Two marking principles (M1 and M2) with two values for lost traffic (0 and -1000)	p. xxx
Figure 6.6	Traffic shaping for an MPEG coded video stream	p. xxx
Figure 6.7	Traffic shaping for traffic process with variation on all time scales	p. xxx
Figure 6.8	Effect of shaping on the metering result	p. xxx

-
- Figure 6.9** Bit rates before (R_{11} and R_{21}) and after (R_{12} and R_{22}) packet discarding p. xxx
-
- Figure 6.10** A target for quality differentiation of CBQ system p. xxx
-
- Figure 6.11** Target “higher price more bandwidth” for a service model using CBQ p. xxx
-
- Figure 6.12** The principle of RED p. xxx
-
- Figure 6.13** Random Early Detection with two importance levels p. xxx
-
- Figure 6.14** Primary and secondary routes p. xxx
-
- Figure 6.15** Load balancing using secondary routes p. xxx
-
- Figure 7.1** Criteria for higher probability of timely forwarding p. xxx
-
- Figure 7.2** Fairness criteria for class selector PHBs p. xxx
-
- Figure 7.3** Most relevant service category for class selector PHB group p. xxx
-
- Figure 7.4** Predictability of load and destination for class selector PHB group p. xxx
-
- Figure 7.5** Delay and importance relations of class selector PHBs p. xxx
-
- Figure 7.6** An example of class selector PHB implementation p. xxx
-
- Figure 7.7** Shaping and discarding in EF PHB p. xxx
-
- Figure 7.8** Relevant fairness issues for expedited forwarding p. xxx
-
- Figure 7.9** Primary and secondary service models for EF PHB p. xxx
-
- Figure 7.10** Predictability of load and destination for expedited forwarding PHB p. xxx
-
- Figure 7.11** Expedited forwarding vs. default PHB in delay-importance scale p. xxx
-
- Figure 7.12** One possible implementation of expedited forwarding with default PHB p. xxx
-
- Figure 7.13** Structure of AF-PHB group p. xxx
-
- Figure 7.14** Relevant fairness issues for assured forwarding PHB group p. xxx
-
- Figure 7.15** Possible service model for an AF-PHB class p. xxx
-
- Figure 7.16** Possible service model for three AF-PHB classes p. xxx
-
- Figure 7.17** Predictability of load and destination for AF-PHB group p. xxx
-
- Figure 7.18** Two delay classes made by two AF PHBs p. xxx
-
- Figure 7.19** Six importance levels made by two AF PHBs p. xxx
-
- Figure 7.20** AF implementation based on four queues and three importance levels p. xxx
-
- Figure 7.21** Structure of DRT-PHB group p. xxx
-
- Figure 7.22** Relevant fairness issues for DRT-PHB group p. xxx
-
- Figure 7.23** Primary and secondary service models for DRT-PHB group p. xxx
-

Figure 7.24	Predictability of load and destination for DRT-PHB group	p. xxx
Figure 7.25	Quality model of DRT-PHB group	p. xxx
Figure 7.26	An implementation of DRT-PHB group	p. xxx
Figure 8.1	Applicability areas for PHB groups	p. xxx
Figure 8.2	Predictability characteristics of CS-PHB group compared to other PHB groups	p. xxx
Figure 8.3	CS PHB group compared to other PHB groups regarding importance and urgency scales	p. xxx
Figure 8.4	Service model comparison of EF- PHB and AF-PHB group	p. xxx
Figure 8.5	Comparison of EF-PHB and AF-PHB group regarding importance and urgency scales	p. xxx
Figure 8.6	The location of an AF class and EF PHB in the scale of urgency and importance	p. xxx
Figure 8.7	Possible location of overload situation with EF- and AF-PHB groups	p. xxx
Figure 8.8	Service model of DRT-PHB group compared with EF-PHB and AF-PHB group	p. xxx
Figure 8.9	DRT-PHB group compared with EF and AF regarding predictability of traffic and destination	p. xxx
Figure 8.10	Three AF classes (each with three PHBs) and DRT-PHB group (2*6 PHBs)	p. xxx
Figure 8.11	Three AF classes (each with three PHBs) and DRT-PHB group (2*6 PHBs)	p. xxx
Figure 8.12	Five alternatives for interoperability between Differentiated Services and Integrated Services	p. xxx
Figure 8.13	The primary service model for Integrated Services	p. xxx
Figure 8.14	Integrated Services compared with EF-PHB and AF-PHB group regarding predictability	p. xxx
Figure 8.15	Tentative urgency and importance relationships between rt-VBR, nrt-VBR, and UBR services	p. xxx
Figure 8.16	Urgency and importance relationships of IEEE 802.1p	p. xxx
Figure 8.17	IEEE 802.1p adjusted into the DRT-PHB structure	p. xxx
Figure 8.18	GPRS quality attributes	p. xxx
Figure 8.19	Predictability of load and destination for multicast streams	p. xxx
Figure 9.1	Network structure for implementation examples	p. xxx

Figure 9.2 Auxiliary VPs between nodes 3, 4, and 11 p. xxx

Figure 9.3 The units of four large organizations p. xxx

Figure 9.4 Average bit rates of 26 VPs on link 1-4 p. xxx

Figure 9.5 EF- PHB with best effort (BE) PHB and an intermediate PHB p. xxx

Figure 9.6 AF- PHB structure with three AF classes p. xxx

Figure 9.7 DRT- PHB structure with 14 PHBs p. xxx
