

MPLS-Enabled Applications

Emerging Developments
and New Technologies

SECOND EDITION



Ina Minei
Julian Lucek

 WILEY

TN915.04
M664
E.2

MPLS-Enabled Applications

Emerging Developments
and New Technologies

Second Edition

Ina Minei

Juniper Networks

Julian Lucek

Juniper Networks



E2008001428

John Wiley & Sons, Ltd

Copyright © 2008

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester,
West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): cs-books@wiley.co.uk

Visit our Home Page on www.wiley.com

All Rights Reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770620.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The Publisher is not associated with any product or vendor mentioned in this book. All trademarks referred to in the text of this publication are the property of their respective owners.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Other Wiley Editorial Offices

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 42 McDougall Street, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 6045 Freemont Blvd, Mississauga, ONT, L5R 4J3, Canada

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Cataloging in Publication Data

Minei, Ina.

MPLS-enabled applications : emerging developments and new technologies /

Ina Minei, Julian Lucek. — 2nd ed.

p. cm.

Includes index.

ISBN 978-0-470-98644-8 (pbk.)

1. MPLS standard. 2. Extranets (Computer networks) I. Lucek, Julian. II. Title.

TK5105.573.M56 2008

621.382'16—dc22

2008008385

British Library Cataloguing in Publication Data

A catalogue record for this book is available from the British Library

ISBN 978-0-470-98644-8 (PB)

Typeset in 10/12pt Palatino by Integra Software Services Pvt. Ltd, Pondicherry, India.

Printed and bound in Great Britain by Antony Rowe Ltd, Chippenham, England.

This book is printed on acid-free paper.

'While MPLS is in itself simple, its apparent complexity lies in the proliferation of applications, which shows no signs of ceasing. To make things worse, catching up involves reading a large number of documents written by various authors at various times in various styles. Here at last is a single, all-encompassing resource where the myriad applications sharpen into a comprehensible text that first explains the whys and whats of each application before going on to the technical detail of the hows.'

Kireeti Kompella, Juniper Fellow, Juniper Networks

'MPLS-Enabled Applications thoroughly covers the MPLS base technology and applications on MPLS-enabled IP networks. It guides you to a comprehensive understanding of standards, problems, and solutions in networking with MPLS. Before it had been necessary to go through material from many different sources, here we have everything in one place. All the MPLS protocols are covered, as are the applications of these protocols. This should be the textbook for MPLS courses, both for training of experienced networking professionals and for universities.'

Loa Andersson, Acreo AB, IAB-member and IETF MPLS working group co-chair

'This is the MPLS text that the industry has been waiting for. On one hand, the text presents MPLS technology clearly enough that the reader can absorb its content in a few easy sittings. On the other hand, the text provides a sufficiently in-depth treatment that even an MPLS expert can learn from it. The authors offer a clear and complete description of MPLS, its inner workings and its applications, in a manner that could only be achieved by persons who have been significant contributors to the MPLS development effort. Every network operator who has deployed or is considering the deployment of MPLS technology should read this book. It is appropriate reading for everyone from the CTO to the tier 1 NOC engineer.'

Ron Bonica, Juniper Networks, IESG member and IETF L3VPN working group co-chair

'This book provides an excellent overview of MPLS mechanisms and applications. It allows understanding what we can do today and what we will be able to do tomorrow with MPLS, as it covers not only stable and deployed technologies such as MPLS Fast Reroute, L3VPNs, L2VPN, but also new technologies under development within the IETF, such as IP/LDP Fast Reroute, inter-domain Traffic

Engineering and Point-To-Multipoint MPLS. Hence this book will be highly useful for network designers, operators, students as well as anyone interested with the MPLS architecture.'

Jean-Louis Le Roux, Senior MPLS architect, Orange

'This is a highly recommended book for network design engineers who want to update themselves with the latest MPLS development, or those who want to learn this technology thoroughly. In addition to the impressive technology coverage and depth, the book is also a delightful reading!'

Lei Wang, IP Network Architect, Telenor

'*MPLS-Enabled Applications* provides an excellent review of the fundamentals and key applications of the MPLS suite of protocols. Its balanced mix of both the technical and business motivations for using these applications make it a must-read for anyone involved in enterprise or service-provider networks.'

Dave Cooper, VP, Network Architecture, Global Crossing Ltd.

'*MPLS-Enabled Applications* is an excellent read for network engineers involved in the design of MPLS networks and services. It can serve as an introduction to MPLS networking or as a reference book for the advanced engineer. It discusses practical issues that must be considered in the design of MPLS networks and services, including MPLS-TE, MPLS-IPVPNs and MPLS L2VPNs. It also discusses current topics that are still evolving in the industry such as inter-AS/area MPLS-TE, point-to-multipoint LSPs and IPVPN multicast, providing a good overview of the issues being addressed and the current industry direction.'

Nabil N. Bitar, Principal member of Technical Staff and lead network architect, Verizon

'*MPLS Enabled Applications: Emerging Developments and New Technologies* second edition, by Ina Minei and Julian Lucek, presents the current state-of-the-art in the specification, development, and application of MPLS and its related technologies. I believe, the readers will find the book to be a very valuable resource.'

Bijan Jabbari, PhD, Founder of Isocore, and Professor of Electrical Engineering, George Mason University

MPLS-Enabled Applications

WILEY SERIES IN COMMUNICATIONS NETWORKING & DISTRIBUTED SYSTEMS

Series Editor: David Hutchison, *Lancaster University, Lancaster, UK*
Series Advisers: Serge Fdida, *Université Pierre et Marie Curie, Paris, France*
Joe Sventek, *University of Glasgow, Glasgow, UK*

The 'Wiley Series in Communications Networking & Distributed Systems' is a series of expert-level, technically detailed books covering cutting-edge research, and brand new developments as well as tutorial-style treatments in networking, middleware and software technologies for communications and distributed systems. The books will provide timely and reliable information about the state-of-the-art to researchers, advanced students and development engineers in the Telecommunications and the Computing sectors.

Other titles in the series:

- Wright: *Voice over Packet Networks* 0-471-49516-6 (February 2001)
Jepsen: *Java for Telecommunications* 0-471-49826-2 (July 2001)
Sutton: *Secure Communications* 0-471-49904-8 (December 2001)
Stajano: *Security for Ubiquitous Computing* 0-470-84493-0 (February 2002)
Martin-Flatin: *Web-Based Management of IP Networks and Systems*, 0-471-48702-3 (September 2002)
Berman, Fox, Hey: *Grid Computing. Making the Global Infrastructure a Reality*, 0-470-85319-0 (March 2003)
Turner, Magill, Marples: *Service Provision. Technologies for Next Generation Communications* 0-470-85066-3 (April 2004)
Welzl: *Network Congestion Control: Managing Internet Traffic* 0-470-02528-X (July 2005)
Raz, Juhola, Serrat-Fernandez, Galis: *Fast and Efficient Context-Aware Services* 0-470-01668-X (April 2006)
Heckmann: *The Competitive Internet Service Provider* 0-470-01293-5 (April 2006)
Dressler: *Self-Organization in Sensor and Actor Networks* 978-0-470-02820-9 (November 2007)
Berndt: *Towards 4G Technologies: Services with Initiative* 978-0-470-01031-0 (February 2008)

About the Authors

Ina Minei joined Juniper Networks in 2000 and has worked in the JUNOS routing group ever since, first as a software engineer and then as a technical manager. Her focus has been on next-generation network technologies, in particular MPLS protocols and applications. She previously worked at Cisco for two years in various software development projects for routers and switches. Ms Minei is an active participant in industry forums and conferences and has co-filed several patents in the area of IP and MPLS. She holds a Master's degree in computer science from the Technion, Israel.

Julian Lucek joined Juniper Networks in 1999 and is currently the Technical Leader for JUNOS for the Europe, Middle East and Africa region, where he has been working with many service providers on the design and evolution of their networks. He previously worked at BT for several years, at first in the Photonics Research Department and later in the data transport and routing area. During this time, he gained a PhD in ultrahigh-speed data transmission and processing from Cambridge University. He is the holder of several patents in the area of communications technology. He has a Master's degree in Physics from Cambridge University and holds Juniper Networks Certified Internet Expert (JNCIE) certification number 21.

Foreword

Yakov Rekhter, Juniper Fellow, Juniper Networks

Multi-Protocol Label Switching (MPLS) began in the mid-1990s with just two modest design objectives. The first was a better integration of ATM with IP by providing a single IP-based control plane that could span both ATM switches and IP routers. The second objective was to augment the IP control plane with some additional functionality, namely traffic engineering using constraint-based routing that was already present in the ATM control plane.

Not long after it started, MPLS usage was extended to applications such as Circuit Cross Connect (CCC), ATM and Frame Relay service over an IP/MPLS infrastructure (draft-martini), BGP/MPLS VPNs (2547 VPNs) and more recently Virtual Private LAN Services (VPLS). The original constraint-based routing functionality evolved beyond traffic engineering to applications such as fast reroute and Differentiated Services Traffic Engineering (DiffServ-TE).

The idea of a single control plane for both ATM switches and IP routers evolved into Generalized Multi-Protocol Label Switching (GMPLS), which provides a single control plane that could span not only routers and ATM switches but SONET/SDH and optical cross connects as well.

One of the recent MPLS developments deserving of mention here is the use of MPLS in the access network, which is covered in this book. Since the first edition of this book, significant progress has been made in the area of MPLS multicast, IP multicast with BGP/MPLS VPNs and IP multicast with VPLS. The second edition covers the advances in these areas.

It is important to keep in mind that in all of the applications mentioned above MPLS is just one of the components of such applications, albeit a critical one. If we look back at the time when MPLS was created, and

compare its design objectives with what MPLS is used for today, we notice several things. First of all, most of the applications of MPLS that we have today were not conceived of during the original design of MPLS. Furthermore, while the original design goal of a better integration of ATM and IP routers is still with us, the way MPLS today supports this is completely different from the way it was originally conceived. Instead of having a single control plane that could span both ATM switches and routers, 'better integration' today has a dual meaning of being able to offer the ATM service over an IP/MPLS infrastructure that has no ATM switches at all, as well as the ability to interconnect ATM switches over such an infrastructure. While originally MPLS was conceived as a technology solely for the Service Providers, we see today how MPLS is gradually penetrating the enterprise environment. Additionally, over time the whole MPLS concept evolved from *Multi-Protocol Label Switching* to *Multi-Purpose Label Switching*.

A new technology quite often generates opposition, and MPLS was by no means an exception. You may all remember how MPLS was branded by its opponents in negative terms as 'bad', 'evil', 'a social disease' or 'a nightmare of unprecedented proportions'. To put this in a proper perspective, we need to keep in mind that technologies exist not for their own sake but for the purpose of solving business problems. Therefore, talking about 'good' technologies versus 'bad/evil' technologies has little practical relevance; yet what is of great relevance is how well a particular technology meets business needs.

One might wonder how they could judge how well a particular technology, like MPLS, meets business needs. To answer this question I would like to remind you of the expression that 'the proof of the pudding is in the eating' (and not in the debate about the pudding). That being said, the ultimate judge of how well a particular technology meets business needs is the marketplace. It is the judgment of the marketplace that determines whether a particular technology deserves to live or to die; and with respect to MPLS the market made its judgment loud and clear – MPLS is here to stay.

Preface

In the three years since we began the first edition of this book, so many new MPLS developments have taken place that our publisher and many readers suggested that a second edition would be useful. The motivation for the book remains the same: MPLS is moving so fast that some of its new applications have already been deployed in production networks, yet are not described anywhere in book form. In many cases, the only available resources are the IETF drafts which list the extensions needed to produce interoperable implementations. These documents often assume familiarity with the problem at hand and do not explain the details of why a particular solution has been chosen or what are the pros and cons of using it. The second edition of *MPLS-Enabled Applications* attempts to fill this gap and provide the reader with an understanding of both the problem and why the solution looks the way it does.

Therefore, when we describe the mechanisms underpinning an MPLS application, the emphasis is on giving an overview of the protocol machinery without delving into the bits and bytes of each protocol message. This allows us to convey the concepts without making it difficult to see the wood for the trees. Also, some of the mechanisms that we write about are currently being defined, so details of the protocol messages may change, but the concepts are less likely to. References at the end of each chapter point to the documents describing the message formats and processing rules.

Although we both happen to work for the same router vendor, the book is not vendor-specific. Occasionally, we point out some vendor-specific quirks if they are relevant to the discussion, or aid in understanding a particular topic. Many of the topics discussed are still under debate in the IETF, and naturally our personal views on one topic or another may be stated more strongly than the opposing view.

WHO SHOULD READ THIS BOOK?

The intended audience of this book includes employees of service providers and network equipment vendors, customers of service providers who are interested in the mechanisms underpinning the services that they buy, network professionals who want to keep up to date with the latest advances in MPLS and students of network technology. To make the second edition of this book more accessible to both the student and to the practitioner of MPLS, we have added study questions at the end of each chapter.

We assume that the reader has some degree of familiarity with network technology and routing protocols, in particular BGP and the link-state IGPs, but these are not a requirement to benefit from the book. Although our main aim is to cover the cutting-edge developments of MPLS, the Foundation chapter allows the reader unfamiliar with MPLS to get up to speed in order to benefit from the remainder of the book. Even when discussing basic topics such as TE or fast reroute, we also explore the more interesting and advanced aspects of the technology.

WHAT IS NEW IN THE SECOND EDITION?

In this second edition, we aim to capture the latest developments in the field. For this reason, we added two new chapters: Chapter 10 covers multicast in L3VPNs, including discussion of both the PIM/GRE and the new 'next-generation' BGP/MPLS schemes, and Chapter 14 discusses the use of MPLS in access networks. Additional material was added and updated throughout the book, such as IPv6 VPNs, ECMP fast-reroute for LDP, 6PE, interworking of LDP-based and BGP-based VPLS, support of multicast in VPLS, OAM tools for P2MP LSPs, topics in self-healing networks and MPLS in mobile networks. The study questions at the end of each chapter are intended to help readers test their understanding of the topics discussed and can serve to trigger debate on the pros and cons of the applicability of a particular technology to a particular deployment.

HOW THIS BOOK IS ORGANIZED

The book is divided into two parts, each containing several chapters. Part One describes the MPLS infrastructure tools used as the foundation to build services, and Part Two covers the MPLS-based services themselves.

The structure of Part One

Chapter 1, the Foundations chapter, reviews the control plane and forwarding plane mechanisms associated with MPLS. In that chapter, we give an overview of the LDP and RSVP signaling protocols and compare the two.

Chapter 2 discusses MPLS Traffic Engineering, which gives service providers control over the path taken by traffic through their network and the ability to give bandwidth guarantees. In this context, we look at the impact of TE on network scalability, as well as at solutions for TE in LDP networks.

Chapter 3 explores the topic of Protection and Restoration in MPLS networks, essential to allowing MPLS networks to carry mission-critical traffic. We cover link and node protection, their respective scaling properties and the cost of bandwidth protection. We also explore more advanced topics such as fast reroute and the new developments for providing fast restoration in IP and LDP networks.

Chapter 4 presents Differentiated Services (DiffServ) Aware Traffic Engineering, which allows traffic engineering to be applied with per-class granularity, bringing QoS to the network.

Chapter 5 introduces Interdomain Traffic Engineering. This is likely to be of increasing importance in the future as MPLS-based services extend across multiple IGP areas and AS boundaries. Both the signaling and computation aspects are discussed, and path-computation elements are also reviewed.

Chapter 6 is devoted to MPLS multicast functionality. This chapter covers not just P2MP LSP setup with RSVP and LDP but also advanced topics such as upstream label allocation and hierarchies of P2MP LSPs. MPLS multicast is currently of great interest as it allows MPLS to be used in broadcast TV and IPTV applications and because it is an essential part of the next-generation L3VPN multicast solutions that have recently been devised in the IETF.

The structure of Part Two

Chapters 7, 8, 9 and 10 are devoted to Layer 3 VPNs – the most widespread application of MPLS to date. Chapter 7 provides a tutorial on L3VPN and explains the basic concepts, Chapter 8 discusses more advanced topics such as route target filtering and scalability analysis, Chapter 9 covers hierarchical VPNs, while Chapter 10 is entirely dedicated to the different solutions for providing multicast support in a VPN, an area which has seen tremendous change in the last couple of years.

Chapter 11 describes the rapidly growing area of Layer 2 transport over MPLS, including pseudowires and Layer 2 VPNs. These allow service

providers to migrate ATM and Frame Relay services to an IP/MPLS network and to offer Ethernet-based alternatives to those services.

Chapter 12 describes the Virtual Private LAN Service (VPLS). This allows a service provider to offer a very simple-to-use service to enterprise customers, in which the customer's sites appear to be attached to the same LAN. Multicast support over VPLS, an area which has seen a lot of change in recent years, is also discussed.

Chapter 13 covers some aspects of the management and troubleshooting of MPLS networks. The subject of management of MPLS networks could fill an entire book by itself and a single chapter does not do it justice. However, we attempt to show some of the challenges (such as ICMP tunneling) and some of the available tools, such as LSPing.

Chapter 14 provides an overview of the emerging trend of using MPLS in the access network, explains why this technology is taking off and describes the various deployment models.

The final chapter takes a look at the achievements of MPLS to date and how MPLS may in future extend beyond the service provider core to DCNs, mobile Radio Access Networks and enterprise networks.

REFERENCES

At the end of each chapter, there is a list of references. In the body of the text, these references appear in brackets, like this [REF1]. Many of the references are IETF documents. As these documents progress in the IETF process, their revision number and document name may change. Therefore, when looking up a reference online, search by the author and title rather than by the document name.

In some chapters, we have included a section with further reading. These are documents that we thought would be useful for those wanting to broaden their knowledge on a particular topic.

Ina Minei, Sunnyvale, CA
Julian Lucek, Ipswich, UK

Acknowledgements

This book would not have existed if it were not for the following three people: Yakov Rekhter, Aviva Garrett and Patrick Ames, and to them we extend our most heartfelt thanks.

Yakov Rekhter encouraged us to pursue this project and provided valuable insight throughout the writing process, from the book proposal, in-depth discussions of many topics, in particular multicast support in L3VPNs and detailed technical reviews of numerous chapters. Most importantly, his faith in our ability to do this work was one of the main factors that determined us to go ahead with this project.

Aviva Garrett was the first person to hear about this idea, encouraged it and arranged all the required support within Juniper Networks, as well as providing detailed editorial review.

Patrick Ames guided us through the intricate process of bringing a book from proposal to the printing press and provided moral support and appropriate scolding as necessary. We would not have been able to pull this off without him. Patrick also did all the hard work of preparing the manuscript (and in particular the art manuscript) for editing.

Both the first and second editions of this book benefited from the contribution of many people. We would like to thank our following colleagues:

Pedro Marques, for his thorough review of almost all chapters, for many technical discussions and for contributing the analysis of VPN scaling and RR scaling.

Arthi Ayyangar, for her insight on all topics RSVP and TE-related, for many technical discussions throughout our years of working together and for the careful review of numerous chapters.

Steven Lin, for reading and commenting on the entire manuscript, on a very tight schedule.

Der-Hwa Gan, for his mentoring role on TE, RSVP and MPLS, and for his very thorough technical review and comments.

Chaitanya Kodeboyina (CK), for very detailed reviews and discussions on several chapters.

Josef Buchsteiner, for always bringing up tough customer problems and for the timely reviews of selected chapters.

Serpil Bayraktar, for never leaving open ends and for very careful reading of the VPN and interdomain TE chapters.

Amir Tabdili, for always asking the hard questions and for reviewing selected chapters.

Quaizar Vohra, for his insight into Layer 2 circuits and the IGP, and for his technical review of these topics.

Margarida Correia, for always questioning proposed solutions to customer problems, and for technical review of selected chapters.

Hector Avalos, for valuable technical discussions and technical review of selected chapters.

Nischal Sheth, for being a mentor on LDP and for numerous discussions on all topics MPLS-related.

Kireeti Kompella, for many technical discussions, for his insight into all aspects of the MPLS technology and for his review of the MPLS in Access Networks chapter. The discussion of 'Option 1' and 'Option 2' deployment models for MPLS within that chapter is largely based on Kireeti's analysis of this topic.

Hannes Gredler, for his review of the Protection and Restoration chapter.

Nitin Bahadur, for his review of the Management, Protection and Restoration, and Multicast over MPLS chapters.

Amit Shukla, for a detailed review of the VPLS chapter and valuable insights in the LDP/BGP VPLS interworking scheme described there.

Derek Harkness, for his expertise on broadband DSL architectures and his review of the MPLS in Access Networks chapter.

Nils Swart, for his expertise on Metro Ethernet and his review of the MPLS in Access Networks chapter.

João Gomes, for his expertise on mobile networks and his review of the Conclusions chapter.

Pierre Bichon, for his expertise on mobile networks and his review of the Conclusions chapter.

Senad Palislaamovic, for his thorough review of all the study questions and their answers.

As always, any errors and omissions are the responsibility of the authors.

We would also like to acknowledge the support and assistance of Juniper Networks and in particular **Jim Murphy**, **Rob Enns**, **Andy Heffernan**, **Nischal Sheth** and **Hector Avalos** for encouraging us to work on this project and providing the necessary resources.

Last but not least we would like to thank the wonderful team at Wiley: **Sarah Hinton, Anna Smart, Birgit Gruber, Richard Davies, Pat Bateson, Vidya Vijayan** and **Emily Dungey** for their support and guidance.

Finally, the authors would like to express their personal thanks to family and friends:

Ina Minei – First of all, I would like to thank my husband Pedro Marques for being my strongest supporter and harshest critic, for having infinite patience throughout the entire project, not once, but two times around. I would not have been able to do this without his support. Second, I would like to thank my father for the long hours he spent with me in my high school years, teaching me English and writing.

Julian Lucek – I would like to thank my partner Rachel and our daughters Emma and Hannah for their great patience and support during the writing of both editions of this book. Also I would like to thank my parents for their encouragement and for looking after my daughters during my writing sessions.