

Web-Based Management of IP Networks and Systems

WILEY SERIES IN COMMUNICATIONS NETWORKING & DISTRIBUTED SYSTEMS

Series Editor: David Hutchison, *Lancaster University, UK*

Series Advisers: Harmen van As, *TU Vienna*
Serge Fdida, *University of Paris*
Joe Sventek, *Agilent Laboratories, Edinburgh*

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 in networking, middleware and software technologies for communications and distributed systems.

The books will provide timely, accurate and reliable information about the state-of-the-art to researchers and development engineers in the Telecommunications and Computing sectors.

Other titles in the series:

Wright: *Voice over Packet Networks* 0-471-49516-6 (February 2001)

Jepsen: *Java in 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 (January 2002)

Web-Based Management of IP Networks and Systems

Jean-Philippe Martin-Flatin



JOHN WILEY & SONS, LTD

Copyright © 2003 John Wiley & Sons, Ltd.,
The Atrium, Southern Gate, Chichester,
West Sussex, PO19 8SQ, England

Telephone (+44) 1243 779777

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

Visit our Home Page on www.wileyeurope.com or 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 4LP9HE, UJK, 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 e.mailed to permreq@wiley.co.uk, or faxed to (+44) 1243 770571.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with 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, 33 Park Road, 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, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

British Library Cataloguing in Publication Data

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

ISBN 0-471-48702-3

Produced from LaTeX files supplied by the author.

Printed and bound in Great Britain by T. J. International Ltd, Padstow, Cornwall.

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

To Béatrice

To the memory of Louise and Henri

Γηράσκω δ'αἰεὶ διδασκόμενος

The older I get, the more I learn

Solon

CONTENTS

List of Figures	xv
List of Tables	xvii
Foreword	xix
Preface	xxi
1 Introduction	1
1.1 Background	1
1.2 Context of This Work	4
1.3 <i>Caveat</i>	6
1.4 Organization	7
2 Terminology	9
2.1 IP World vs. Telecommunications World	10
2.2 Different Types of Enterprise Management	10
2.3 Manager, Agent, Management Application, Management Platform, and Management System	14
2.4 Manager vs. Administrator	15
2.5 Agent, Mobile Agent, and Intelligent Agent	15
2.6 Proxy vs. Gateway	15
2.7 Delegation	16
2.8 Paradigm vs. Technology	17
2.9 Architecture vs. Framework	18
2.10 Firewalls and Trust Domains	18
2.11 Regular Management vs. <i>Ad Hoc</i> Management	20
2.12 Ontologies, Schemas, and Models	22
3 Problem Statement	25
3.1 Characteristics of SNMP-Based Management	25
3.2 Strengths of SNMP-Based Management	33
3.3 Problems with SNMP-Based Management	33

3.3.1 Scalability and efficiency issues	33
3.3.2 Missing features	46
3.3.3 Non-technical problems in SNMP-based management	49
3.4 We Need a New Solution to a New Problem	53
3.5 Summary	56
4 Overview of the Solution Space	57
4.1 Simple Taxonomy of Network and Systems Management Paradigms	58
4.1.1 Centralized paradigm	60
4.1.2 Weakly distributed hierarchical paradigms	61
4.1.3 Strongly distributed hierarchical paradigms	63
4.1.4 Strongly distributed cooperative paradigms	69
4.1.5 Synthetic diagram	71
4.2 Enhanced Taxonomy of Network and Systems Management Paradigms	72
4.2.1 A stroll through organization theory	72
4.2.2 Other criteria for our enhanced taxonomy	78
4.2.3 Synthetic diagram	85
4.3 Summary	86
5 Analysis of the Solution Space	87
5.1 No Win-Win Solution	87
5.2 Do Not Focus on Technology Immediately	89
5.3 Reality Check: Support and Technical Maturity	90
5.4 The My-Middleware-Is-Better-Than-Yours Syndrome	92
5.5 Mobile Code and Security	94
5.6 Distribution	95
5.7 Web-Based Management	95
5.8 Summary	96
6 State of the Art in Web-Based Management	97
6.1 Browser-Based Usage Monitoring	98
6.2 Browser-Based Metamanagement	99
6.2.1 Online problem reporting	100
6.2.2 Online management procedures, online documentation	100
6.2.3 Online troubleshooting assistance	101
6.3 Browser-Based Management	101
6.3.1 Troubleshooting scripts executed via a Web browser	102
6.3.2 Configuration management	102
6.3.3 Java applet with an SNMP stack	103
6.4 Three-Tier Management	103
6.4.1 Deri <i>et al.</i> : SNMP-to-URL mapping	104
6.4.2 Kasteleijn: HTTP and SNMP	104
6.4.3 Integration of a Web browser in the SNMP-based management platform	105
6.4.4 Deri: Java RMI and SNMP	105
6.5 HTTP-Based Management	105
6.5.1 CLI wrappings and mappings	107

- 6.5.2 Embedded HTML pages and CGI programs 115
- 6.5.3 Embedded management applications 116
- 6.5.4 Low-footprint embedded HTTP servers 116
- 6.6 XML-Based Management 117
 - 6.6.1 Web-Based Enterprise Management (WBEM) 117
 - 6.6.2 John *et al.*: XNAMI 122
- 6.7 Distributed Java-Based Management 123
 - 6.7.1 Java Management Application Programming Interface (JMAPI) 123
 - 6.7.2 Java Management eXtensions (JMX) 124
 - 6.7.3 Federated Management Architecture (FMA) 124
 - 6.7.4 Anerousis’s Marvel 124
- 6.8 Commercial Products 125
- 6.9 Summary 126
- 7 A New Management Architecture: WIMA 127
 - 7.1 Main Architectural Decisions (Analysis Phase) 128
 - 7.1.1 One management architecture, four models 128
 - 7.1.2 No need to define yet another information model 131
 - 7.1.3 Dissociation of the communication and information models 131
 - 7.1.4 A new organizational model: WIMA-OM 132
 - 7.1.5 A new communication model: WIMA-CM 137
 - 7.2 Main Design Decisions (Design Phase) 138
 - 7.2.1 Web technologies 138
 - 7.2.2 Three-tier architecture 139
 - 7.2.3 Management server: COTS components and object-oriented frameworks 140
 - 7.2.4 Management-data transfers across firewalls 142
 - 7.2.5 Data repository independent of the management platform 142
 - 7.2.6 Bulk transfers of regular management data 144
 - 7.2.7 Dealing with legacy systems 144
 - 7.2.8 Richer semantics 144
 - 7.2.9 Easy to deploy 145
 - 7.3 Regular Management and Notification Delivery: The Push Model 145
 - 7.3.1 Publication phase 146
 - 7.3.2 Discovery phase 146
 - 7.3.3 Subscription phase 152
 - 7.3.4 Distribution phase 158
 - 7.3.5 Distribution 163
 - 7.3.6 Migration path: four-tier architecture 168
 - 7.4 *Ad Hoc* Management: The Pull Model 168
 - 7.4.1 Two-tier architecture (no firewall) 169
 - 7.4.2 Three-tier architecture (with firewall) 171
 - 7.4.3 Distribution 172
 - 7.4.4 Migration path 173
 - 7.5 Summary 173
- 8 A New Communication Model: WIMA-CM 175
 - 8.1 Two Communication Problems 175

8.1.1	Simplified terminology	176
8.1.2	Communication problem for pushed data	176
8.1.3	Communication problem for pulled data	177
8.2	Main Design Decisions	177
8.2.1	Dissociation of the communication and information models	177
8.2.2	Reliable transport protocol: TCP	178
8.2.3	Persistent TCP connections	181
8.2.4	Firewalls: persistent connections must be created by the manager	191
8.2.5	Reversed client and server roles	193
8.3	The Sockets API as a Manager-Agent Communication API	194
8.4	HTTP-Based Communication for Pushed Data: WIMA-CM-push	195
8.4.1	MIME multipart	196
8.4.2	Some notifications are more equal than others	197
8.4.3	Specifying the information model in the MIME header	200
8.4.4	Optional compression of management data	202
8.4.5	Example of HTTP and MIME-part headers	203
8.4.6	Simplifications in case we do not have a firewall	204
8.5	Timeouts and Reconnections	205
8.5.1	Timeouts by the operating systems	205
8.5.2	Timeouts by the applications	215
8.5.3	Synthesis	219
8.6	HTTP-Based Communication for Pulled Data: WIMA-CM-pull	221
8.7	Summary	222
9	XML in Integrated Management	223
9.1	Why Use XML in NSM?	223
9.1.1	Overview of XML	224
9.1.2	Who uses XML?	224
9.1.3	Advantages of using XML in general	225
9.1.4	Advantages of using XML in NSM	226
9.1.5	Advantages of using XML in integrated management	227
9.1.6	Disadvantages of using XML in general	228
9.2	XML for Representing Management Data: Model- and Metamodel-Level Mappings	228
9.2.1	Model-level mapping	229
9.2.2	Metamodel-level mapping	232
9.2.3	Comparison between model- and metamodel-level mappings	235
9.3	XML for Dealing with Multiple Information Models	235
9.4	XML for High-Level Semantics	238
9.4.1	Transfer of an entire SNMP MIB table	238
9.4.2	Suppression of “holes” in sparse SNMP tables	239
9.4.3	Time series of a MIB variable	239
9.4.4	Distributed object-oriented programming with XML	241
9.5	XML for Integrated Management: a Unified Communication Model	242
9.6	Summary	244
10	A WIMA-Based Research Prototype: JAMAP	245
10.1	Overview of JAMAP	246

- 10.1.1 Key design decisions at a glance 246
- 10.1.2 More on the design of JAMAP 247
- 10.1.3 Advanced Java technologies in JAMAP 248
- 10.1.4 Overview of the communication aspects 249
- 10.1.5 Distribution phase for monitoring and data collection 251
- 10.1.6 Distribution phase for notification delivery 252
- 10.2 Management Station 253
 - 10.2.1 SNMP MIB data subscription applet 253
 - 10.2.2 SNMP notification subscription applet 254
 - 10.2.3 Rule edition applet and rule mapping applet 254
 - 10.2.4 Event log applet 255
- 10.3 Management Server 255
 - 10.3.1 Pushed-data collector servlet 255
 - 10.3.2 Notification collector servlet 257
 - 10.3.3 Event manager servlet 257
- 10.4 Agent 258
 - 10.4.1 Subscription servlet 258
 - 10.4.2 Push dispatcher servlet 258
 - 10.4.3 Notification dispatcher servlet 259
- 10.5 Reusability 259
- 10.6 Summary 260
- 11 How Does Our Solution Compare with Others? 261
 - 11.1 Comparison with SNMP-Based Management 262
 - 11.1.1 Almost all of the problems in SNMP have been solved 262
 - 11.1.2 Complementarities between WIMA and SNMP 264
 - 11.2 Comparison with WBEM 264
 - 11.2.1 Similarities: HTTP and XML 264
 - 11.2.2 Problems with WBEM 265
 - 11.2.3 Complementarities between WIMA and WBEM 267
 - 11.3 Comparison with JMX 267
 - 11.4 Known Problems with WIMA 268
 - 11.5 Summary 270
- 12 Conclusion 271
 - 12.1 Summary 271
 - 12.2 Directions for Future Work 274
- Acronyms 277
- References 283
- Appendix A: The Interfaces Group in SNMP MIB-II 299
- Appendix B: Metamodel-Level XML Mapping of the Interfaces Group in SNMP MIB-II 309
- Appendix C: Metamodel-Level XML Mapping of a Simple CIM Class 319
- Appendix D: Remote Method Invocation of a CIM Object 325
- Index 327

LIST OF FIGURES

Fig. 1. Delegation and cooperation.	16
Fig. 2. Management across firewalls: the real picture	19
Fig. 3. Management across firewalls: the simplified picture	20
Fig. 4. Management tasks and their associated data flows.	21
Fig. 5. The effect of size on the enterprise structure	73
Fig. 6. Enhanced taxonomy of NSM paradigms	85
Fig. 7. HTTP-based management in attended mode: Web browser.	106
Fig. 8. HTTP-based management in attended mode: application.	106
Fig. 9. HTTP-based management in unattended mode: application.	107
Fig. 10. Example of CLI wrapping into XML	108
Fig. 11. Example of CLI wrapping into HTML	109
Fig. 12. Network overhead of an SNMP <code>get</code>	134
Fig. 13. Three-tier architecture	139
Fig. 14. Push model: discovery and subscription phases (without firewall)	153
Fig. 15. Push model: discovery and subscription phases (with firewall)	157
Fig. 16. Push model: distribution phase for notification delivery and event handling.	159
Fig. 17. Push model: distribution phase for data collection and monitoring	160
Fig. 18. Distribution: management server viewed as a distributed system	164
Fig. 19. Mid-level manager = management server + agent + glue.	165
Fig. 20. Distributed hierarchical management: delegation by domain.	166
Fig. 21. Migration path: four-tier architecture	168
Fig. 22. Pull model: <i>ad hoc</i> management without firewall	169
Fig. 23. Pull model: <i>ad hoc</i> management with firewall	171
Fig. 24. Communication problem for pushed data	176
Fig. 25. Communication problem for pulled data	177
Fig. 26. Memory overhead of a socket: neighbors vs. destination hosts	183
Fig. 27. Reversed client and server roles	193
Fig. 28. The sockets API as a manager-agent communication API.	194
Fig. 29. Distribution via HTTP with a firewall	195
Fig. 30. TCP payload of the infinite HTTP response.	196
Fig. 31. Two connections per agent	198
Fig. 32. Multiple connections per agent	199

Fig. 33. Formal definition of the content type202

Fig. 34. Push: HTTP and MIME-part headers of the agent’s reply203

Fig. 35. Distribution via HTTP without firewall204

Fig. 36. Pull: HTTP header of the agent’s reply.222

Fig. 37. Model-level XML mapping of the Interfaces Group in SNMP MIB-II: DTD230

Fig. 38. Model-level XML mapping of the Interfaces Group in SNMP MIB-II: example of
XML document231

Fig. 39. Model-level XML mapping of a simple CIM class232

Fig. 40. Multiple information models: generic integration236

Fig. 41. Multiple information models: specific integration237

Fig. 42. Time series with XML240

Fig. 43. XML: Hierarchically distributed integrated management.243

Fig. 44. JAMAP: Communication between Java applets and servlets for monitoring and data
collection.250

Fig. 45. JAMAP: Communication between Java applets and servlets for notification delivery
and event handling250

Fig. 46. Push model in JAMAP: monitoring and data collection.251

Fig. 47. Push model in JAMAP: notification delivery and event handling252

LIST OF TABLES

Table 1. Simple taxonomy of NSM paradigms	71
Table 2. Semantic richness of the information model	81
Table 3. Memory overhead per socket	185
Table 4. Memory overhead per destination host	186
Table 5. Memory overhead per neighbor	186
Table 6. Memory overhead of TCP receive buffers	186
Table 7. Mapping between Linux and 4.4BSD keepalive-control kernel variables.	209
Table 8. Per-kernel TCP keepalives: default and recommended values	210
Table 9. Keepalive control: mapping between TCP socket options and kernel variables	213

FOREWORD

Management is about monitoring the current state of network and system components, receiving event reports and controlling the components in order to provide the required service expected by users. Network management has concentrated on management of comparatively simple hardware devices such as routers and hubs. Telecommunications management has been concerned with managing more complex switches and intelligent network services. Systems management is concerned with managing resources such as a disk store or memory and setting scheduling or security policies. However there is now a trend towards management of end-to-end services and applications rather than individual network and system components. These services may make use of many different network and system components, so a common standards based approach to management is essential.

Standards for management of networks and systems have been an elusive goal pursued by many in the Management Community over the past 10 years. In the early '90s we saw the war of words between protagonists for the Simple Network Management Protocol (SNMP), being developed for the Internet, and the various International Organization for Standardization (ISO) management standards being pushed by Telecommunications organisations. In spite of the so-called convergence between networking and telecommunications, SNMP was not widely used in telecommunications and the ISO standards were not used by the Internet community. Systems management was all vendor specific. The Distributed Management Environment (DME) from the Open Software Foundation was going to be the saviour of all the world's management problems in the mid '90s but it quickly disappeared without trace. SNMP-based management is currently the dominant approach for network management and is being proposed as being suitable for systems and service management as well.

In this book, Jean-Philippe Martin-Flatin gives some interesting insights into the problems of SNMP-based management and evaluates many of the alternatives such as Java, mobile code and Web-based approaches. He proposes a Web-based Integrated Management Architecture (WIMA) which has the advantage that it makes use of widely used web

standards such as the HTTP protocol and XML for data representation, rather than inventing something specific for management. The prototype implementations described in the book indicate the practicality of the work and that it can make use of distributed processing concepts which will be essential for management of large-scale systems. Although it is very difficult for an individual to influence the many large organisations developing management solutions, I hope that the simplicity of the approach described in this book does make some impact.

Professor Morris Sloman

Imperial College, London, UK

PREFACE

The management of most IP networks is currently based on an open management architecture named after its communication protocol: the Simple Network Management Protocol (SNMP). We analyze a number of problems with this approach. Some are related to the efficiency, scalability, latency, and expressiveness of SNMP; others are due to the way the management-platform market evolved in the 1990s.

Although IP systems can also be managed with SNMP, in practice, their management usually relies on proprietary mechanisms. This defeats the purpose of open management and does not facilitate the integration of network management with systems management.

After reviewing the numerous alternatives that are currently investigated by the management community (e.g., mobile code and intelligent agents), we propose to base the next generation of management applications on a new management architecture: WIMA, the Web-based Integrated Management Architecture. WIMA is based on standard Web technologies. It relies on a push-based organizational model for regular management (i.e., data collection for offline processing, and monitoring over a long period of time) and notification/event delivery, and a pull-based organizational model for *ad hoc* management (data retrieval over a very short time period).

Its communication model is characterized by (i) the use of a standard communication protocol (HTTP) instead of a domain-specific protocol (SNMP); (ii) the use of persistent connections between agents and managers (and between mid- and top-level managers in distributed hierarchical management); (iii) the independence of the information model (SNMP MIB, CIM schema, etc.) and the communication model; and (iv) a reversed client-server architecture that facilitates crossing firewalls.

In WIMA, the preferred method for representing management data in transit is XML. It is well suited to distributed hierarchical management, and unifies the communication model across the range of disciplines covered by enterprise management (i.e., network, systems, application, service, and policy-based management). This facilitates the integration of management and offers high-level semantics to management-application designers.

All the major problems that we identified in SNMP are solved in WIMA. Our architecture is validated by a research prototype, JAMAP, which is described in Chapter 10.

Audience of this book

This book presents rather advanced material. It is not a primer. We assume that the reader already knows the basics of network management and object-oriented analysis & design, and introductory textbooks are referenced. Great care was taken to make this material readable to a vast audience, from practitioners in industry to students, from programmers to designers, and from administrators to managers.

The terminology defined in Chapter 2 clarifies frequent misunderstandings and should be valuable to all. Chapters 3, 4, and 6 can be used by academics for teaching B.S. and M.S. courses in distributed network and systems management, or by consultants for training administrators. They can also be instructive to MIS directors who wish to update their knowledge of enterprise management. Chapter 3 should benefit engineers in charge of designing and developing SNMP-based management applications. In Chapter 5, our analysis of the solution space should be interesting to managers and administrators investigating the current market. It explains why Web-based management is not only trendy, but also makes technical sense. Chapters 7, 8, and 10 should be very useful to start-ups and corporate organizations developing management applications, and to M.S. and Ph.D. students embarking on a Web-based management project. They contain a wealth of technical hints and tips that could save software designers and developers months of work. Finally, by explaining the advantages of using XML for representing management data, Chapter 9 should enlighten those professionals and students who still believe that Web-based management is only about adding a Web browser to an existing management application.

Index Keys

IEEE Web Thesaurus

<<http://www.ieee.org/web/developers/webthes/>>

- Communication System Operations and Management
- Web and Internet Services
- Management Information Systems
- Computer Networks

Computing Research Repository (CoRR)

<<http://www.acm.org/pubs/corr/>>

Area: Computer Science

Subject Classes: Networking and Internet Architecture; Distributed Computing

Copyright Notices

This book is derived from the author's Ph.D. dissertation, which was published by the Swiss Federal Institute of Technology, Lausanne (EPFL) in October 2000 [163]. Some parts are copied *verbatim*, but many sections have been updated or rewritten.

A fraction of the material presented in Section 3.3.1 was published in March 1999 in *The Simple Times*, the online magazine of the SNMP community [243]. This article was co-authored by Ron Sprenkels, who contributed several ideas explained there. This section also benefited from discussions that took place during the first meeting of the IRTF Network Management Research Group.

Early versions of the material presented in Chapter 4 were presented by the author in two articles: one published by Plenum Press (now part of Kluwer Academic) in the *Journal of Network and Systems Management* [160], and another published jointly by the University of Technology, Sydney, Australia and the University of Western Sydney, Nepean, Australia in the proceedings of the DSOM'97 workshop [158].

In Chapter 4, Fig. 5 is adapted from Weinshall and Raveh [285, pp. 56–57], published by Wiley.

Early versions of the figures in Chapter 7 and Chapter 8 appeared in a paper by the author published by the IEEE in the proceedings of the IM'99 symposium [161].

An early version of the material presented in Chapter 10 appeared in an article by the author published by Springer Verlag in the proceedings of the DSOM'99 workshop [162].

In Appendix A, we give the IETF's definition of the Interfaces Group in RFC 1213 [173, pp. 16–23]. This material is reproduced with the permission of the Internet Society.

In Appendix C, we include an example of CIM-to-XML metamodel-level mapping given by the DMTF [73]. This material is reproduced with the permission of the DMTF.

In Appendix D, we include an example of remote method invocation of a CIM object given by the DMTF [77, Section A.10]. This material is reproduced with the permission of the DMTF.

Finally, JAMAP, the software described in Chapter 10, was released by the author under the GNU General Public License of the Free Software Foundation [97].

Acknowledgments

I would like to thank my editors at Wiley UK, Sally Mortimore and Birgit Gruber, for patiently explaining to me all the steps involved in the publication of this book.

I am grateful to the Swiss Federal Institute of Technology, Lausanne (EPFL) and the Swiss National Science Foundation (FNRS) for funding the research work that led to this book. Thanks also to AT&T Labs Research for funding extra developments on the research

prototype described in Chapter 10 and for allowing release 0.4 of this software to be freely available under the GNU General Public License.

I thank my former Ph.D. supervisor, Jean-Pierre Hubaux, for welcoming me on his team at EPFL, for offering me exceptional working conditions, and for giving me a lot of freedom and independence while conducting my research. Thanks also to the members of my Ph.D. committee (Roland Balter, Subrata Mazumdar, André Schiper, and Morris Sloman) for the time they spent reviewing my dissertation and their valuable feedback.

I gratefully acknowledge that most of the implementation of JAMAP, the research prototype described in Chapter 10, was performed by Laurent Bovet [31] and Claire Ledrich [147] in the course of their M.S. theses, under my supervision. Laurent's work was funded by EPFL, Claire's by AT&T Labs Research.

I had the chance to work or discuss with many great people over the years and cannot name them all. I would particularly like to thank Simon Znaty, who taught me how to do research; Andrea Westerinen, who thoroughly reviewed this book and suggested useful changes and corrections; Jürgen Schönwälder, who reviewed several chapters and with whom I love to argue about the future of network management; Werner Almesberger, who helped me quantify the memory footprint of TCP connections in the Linux kernel and reviewed Chapter 8; Luca Deri, George Pavlou, and Gian Pietro Picco, who sent me precious comments on early versions of the material presented herein; Ditlef Martens, Otto Pesonen, and Baudouin Raoult, who taught me so much about the Unix internals; Tony Bakker and Dick Dixon, who taught me the basics of networking; Walter Zwiefelhofer and Claus Hilberg, who taught me rigor when “selling” a technical solution to management; Holly Cogliati, who patiently improved my English, taught me the differences between American and British English, and proofread this book; and Patrick Lachaize, who encouraged me to move into IP networking back in 1989.

I express my deepest gratitude to Gilles Maignaud, Alain Martinez, and François Martin for teaching me, when I was 16–18, that mathematics, physics, and science at large can be great fun. Even in difficult times, I have always strived to take some pleasure in doing my job. Their advice has served me throughout my career.

Finally, I am very grateful to my family and friends for their unwavering support and positive feedback over the years. Most of all, I am indebted to my wife, Béatrice, who put up with the long nights and weekends that I spent writing this book. I knew that I had married a rare pearl...

Jean-Philippe Martin-Flatin, Lausanne, Switzerland, July 7, 2002