

Technical, Commercial and
Regulatory Challenges of

QoS

An Internet Service Model Perspective



XiPeng Xiao

Morgan Kaufmann Publishers is an imprint of Elsevier.
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA

This book is printed on acid-free paper.

© 2008 by Elsevier Inc. All rights reserved.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which Morgan Kaufmann Publishers is aware of a claim, the product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

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, scanning, or otherwise—without prior written permission of the publisher.

Permissions may be sought directly from Elsevier's Science & Technology Rights Department in Oxford, UK: phone: (+44) 1865 843830, fax: (+44) 1865 853333, E-mail: permissions@elsevier.com. You may also complete your request online via the Elsevier homepage (<http://elsevier.com>), by selecting "Support & Contact" then "Copyright and Permission" and then "Obtaining Permissions."

British Library Cataloguing in Publication Data

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

Library of Congress Cataloging-in-Publication Data

Xiao, XiPeng.

Technical, commercial, and regulatory challenges of QoS / XiPeng Xiao.

p. cm. — (The Morgan Kaufmann series in networking)

Includes bibliographical references and index.

ISBN 978-0-12-373693-2 (alk. paper)

1. Computer networks—Quality control. 2. Computer networks—Management. 3. Internet service providers. 4. Internet—Government policy. I. Title. TK5105.5956.X53 2008 004.67'8—dc22

2008026434

ISBN: 978-0-12-373693-2

ISSN: 1875-9351

For information on all Morgan Kaufmann publications,
visit our Web site at www.mkp.com or www.books.elsevier.com

Typeset by Charon Tec Ltd., A Macmillan Company. (www.macmillansolutions.com)

Printed and bound in The United States of America

08 09 10 10 9 8 7 6 5 4 3 2 1

Working together to grow
libraries in developing countries
www.elsevier.com | www.josaid.org | www.sabre.org
ELSEVIER BOOKS INTERNATIONAL Sabre Foundation

Preface

Today, the increasing popularity of mobile phones and VoIP generates a large impact on the revenue of traditional telecom service providers. To maintain their subscriber base and average revenue per user, telecom service providers are eager to offer premium services such as IP TV, online gaming, etc. It is assumed that these services will create a large demand for IP QoS. At the same time, there is a trend to use the Internet as the common carrier for all kinds of services, instead of having many special-purpose networks. It is assumed that this will bring the QoS requirement to the Internet. However, after so many years of research, development, and claimed deployment, QoS is still something of the future in the Internet. Among other issues, the service quality of the Internet can still be unpredictable. What makes QoS so elusive? What is missing? What needs to be done to bring QoS to reality?

The current Net Neutrality debate further complicates the matters on QoS. Since the idea of QoS was formed, it has always been taken for granted that if carriers can provide QoS, they can charge users for it. The Net Neutrality debate cast doubt on this belief for the first time. A Net Neutrality legislation can dramatically change the QoS landscape overnight. What is Net Neutrality? How will the debate shape the evolution of QoS? Why should a common person even care?

In this book, we will discuss the technical as well as commercial and regulatory challenges of QoS, and propose a model to overcome these challenges. We will first define what QoS is, and then discuss:

- What are the QoS requirements of common applications?
- What can the Internet offer in terms of QoS today?

This lets us see the performance gap and thus provides a base for subsequent discussions.

We then review the contemporary QoS wisdom, and discuss its commercial, regulatory, and technical challenges. Some of the important topics include:

- The commercial challenges of the traditional QoS wisdom, regarding:
 - Who should pay for QoS, business or consumers, senders or receivers?
 - Why does this matter?
 - What kind of assurance should carriers provide with QoS?
 - Will the attempt to sell QoS increase customer churn, because it is considered as evidence of poor quality for the basic service?
- The regulatory challenges of the traditional QoS wisdom, regarding:
 - What is Net Neutrality and how does it relate to QoS?
 - Will carriers be allowed to charge for QoS?
 - What is the impact of the uncertain government regulation on QoS?

- The technical challenges of the traditional QoS wisdom, regarding:
 - What cooperation is needed among carriers in order to provide QoS? Can they be done at an acceptable cost?
 - Will various QoS mechanisms introduce too much complexity into the network to hurt network reliability and reduce QoS rather than improve QoS?
 - What are the technical challenges to differentiating a higher CoS to a point that the end users can perceive the difference from Best Effort (so that the end users will be willing to buy the higher CoS)?

We will then propose how to improve the current QoS business model to overcome the commercial, regulatory, and technical challenges. On the commercial side, this involves a change to the QoS pricing scheme. On the technical side, this involves comprehensive consideration of all the options available, increased emphasis on certain mechanisms and deemphasis of some other mechanisms. We will go to great lengths to explain why the proposed pricing scheme is better for the industry and for the users, and back it up with a large amount of historic evidences. These evidences include revenue and usage statistics in the postal industry's 200-year history and in the telephony industry's 100-year history. These statistics establish the evolution trend of pricing schemes for communication services. We believe that our rationale for the proposed model becomes clear in light of the historic trend.

Next, we will present two case studies on real-world QoS deployment. One is about Internet2 (<http://www.internet2.edu/>), the next generation Internet test bed in the United States. This case study is written by Ben Teitelbaum and Stanislav Shalunov, formerly of Internet2 and now of Bit Torrent (<http://www.bittorrent.com/>). The other is about Internap, one of the few network service providers in the United States that successfully commercialized QoS (<http://www.internap.com/>). This case study is written by Ricky Duman of Internap. Because these case studies are written by the network operators who did the actual deployment, the readers can hear directly from the horse's mouth about QoS deployment issues in real-world networks, and the lessons they learned. We will also discuss QoS issues in wireless networks. That chapter is written by Dr. Vishal Sharma of Metanoia Inc. (<http://www.metanoia-inc.com/>), a well-known technology consulting firm in Silicon Valley. The contributions of Ben Teitelbaum, Stanislav Shalunov, Ricky Duman, Vishal Sharma, and Abhay Karandikar are gratefully acknowledged. We draw our conclusions at the end.

Throughout this book, there are a number of advanced technical issues that are discussed but are not fully resolved. These are good topics for further research.

Because this is the first book that covers all three important aspects of QoS—technical, commercial, and regulatory—and each aspect has a broad range of topics, we recognize that it is possible that over time, some of our opinions may turn out to be revisable. But we believe that this won't hurt the main purpose of this book, which is to help people see the big picture of QoS, think critically about QoS, and form their own opinion on QoS. With this recognition, we are eager to

hear back from the readers. A web site has been set up at <http://groups.google.com/group/qos-challenges> for discussion. You can present your view points for other people to see.

AUDIENCE

In this book, we will discuss all three major aspects of QoS—technical, commercial, and regulatory—and how they interact with each other. We will first examine the status quo of QoS to show that the contemporary QoS wisdom has not been able to make QoS a reality for the Internet. We will then provide our explanation for this outcome by discussing the technical, commercial, and regulatory challenges. We will then propose a revision to the QoS model; discuss how it can overcome the commercial, regulatory, and technical challenges; and explain why we believe it is better for the industry.

We believe this book has value for the following audiences:

1. For people who are interested in understanding QoS technology, this book is a one-stop place for various flavors of technical QoS solutions, their pros and cons, the major overlooked factors in the current Diffserv/traffic management-centric solution, and the key trade-offs that must be made for a technical solution to be practical. The description is relatively high level so that most people can understand it without difficulty. For people who are interested in knowing the details, the book provides pointers to other references. The case studies about Internet2 and Internap's QoS deployments enable the readers to see QoS deployment issues in real-world networks. People about to deploy QoS can benefit from the lessons they provided. Academic people may also be interested in a number of advanced technical issues that are discussed but are not fully resolved—these can be good topics for further research.
2. For people who are interested in understanding the commercial issues, this book provides a comprehensive discussion about the commercial challenges in selling QoS. The key issues include “What is the effect of Internet users' Free mentality?,” “What QoS assurance should be provided to attract users to buy QoS, soft or hard?,” “Whom should NSPs charge QoS to, business or consumers? Senders or receivers?,” “Will charging for QoS be considered as a sign of poor service quality for the basic service? Will it trigger customer defection?,” “What should the contractual settlement among NSPs be to facilitate interprovider QoS?,” “How much QoS revenue can realistically be generated?” These discussions are particularly useful for people/companies who plan to invest in QoS, for example, either developing QoS features or deploying QoS mechanisms.
3. For people who are interested in understanding the regulatory issues and the Net Neutrality debate, this book provides a succinct summary of the

key issues, and the opinions of both the proponents and opponents on these issues. This saves the readers from having to spend the time to locate the information and follow the discussions. This would help the readers quickly form their own opinions on Net Neutrality.

ORGANIZATION

This book contains three parts. Part 1 discusses the current situation of Internet QoS, and points out that the contemporary QoS wisdom has not been able to make QoS a reality. Part 2 explains this outcome by discussing the commercial, regulatory, and technical challenges. Part 3 proposes a revised QoS pricing scheme and a technical solution, and discusses how they overcome or relieve the commercial, regulatory, and technical challenges.

Part 1 contains four chapters.

- Chapter 2 discusses what QoS means in this book, common applications' requirements on QoS, and the degree that the current Internet meets those requirements. The purpose of discussing the application requirements is to make the objectives of QoS clear. The purpose of discussing the degree that the current Internet meets those requirements is to clarify the gap between what is needed and what is available, so that we know what else may be needed to deliver QoS.
- Chapter 3 discusses the historic evolution of QoS solutions. The purpose is to provide the readers with some technical background and a historic view on various flavors of technical QoS solutions.
- Chapter 4 discusses the “contemporary QoS wisdom,” including its business model and its technical solution. This is to provide a base for commercial, regulatory, and technical examination.
- Chapter 5 discusses the reality related to QoS, especially from a commercial perspective. This is to give us a sense of how well the traditional QoS wisdom works.

Part 2 contains four chapters.

- Chapter 6 discusses the commercial challenges of the conventional QoS business model.
- Chapter 7 discusses the regulatory challenges.
- Chapter 8 discusses the technical challenges.
- Chapter 9 summarizes the key points discussed in this part, and discusses the lessons that are learned.

The purpose of discussing the commercial, regulatory, and technical challenges is to expose the issues of the conventional QoS model. The purpose of discussing the lessons learned is to point out the direction for possible improvement of the QoS model.

Part 3 contains five chapters.

- Chapter 10 proposes a revised pricing scheme for QoS, and discusses how it overcomes or relieves the most difficult commercial and regulatory challenges. To help the readers see the rationale of this revision, we present a large amount of revenue and usage statistics in the postal industry's 200-year history and in the telephony industry's 100-year history. These statistics establish the evolutionary trend of pricing schemes for communication services. Our rationale for the proposed pricing scheme revision becomes clear in light of the historic trend.
- Chapter 11 discusses the revised technical solution and its benefits.
- Chapter 12 presents two real-world QoS deployments at Internet2 and Internap, and the lessons they learned.
- Chapter 13 discusses QoS in the wireless world. Because network resource is much more limited in the wireless world, QoS approaches are very different too. This is another effort to help the readers see the big picture.
- Chapter 14 concludes the book.

It is recommended that this book be read in its entirety, and in the order the chapters are presented. This allows the big picture to manifest in a way that is easier to understand.

Acknowledgements

Writing the acknowledgements is the enjoyable part of this long and laborious journey. This book would not be possible without the help of a number of friends and experts.

First, I would like to thank Dr. Andrew Odlyzko. His Internet economics theory enlightened me. His research work formed the foundation of Chapter 10. He also reviewed the manuscript, provided valuable feedback, and pointed me to other information and subject matter experts.

I would like to thank Zheng Wang and John Yu for the helpful discussions during the conception of this book. Zheng Wang and Angela Chiu reviewed the proposal. Les Cottrell, K. Claffy, and Al Morton provided valuable information on Internet performance statistics and application QoS requirements for Chapter 2. Louise Wasilewski and Mike Benjamin provided early feedback on Net Neutrality. Jon Aufderheide and Brook Bailey provided information on service provider peering and service provider perspective. Dave Wang and Yakov Rekhter provided feedback on the key messages of the book. Zhisheng Chen and Yong Xue took the time to read the whole manuscript and provided valuable feedback. Bill St. Arnaud reviewed the manuscript and provided detailed comments and valuable suggestions for improvement.

Ben Teitelbaum and Stanislav Shalunov wrote the Internet2 Case Study. Ricky Duman wrote the Internap Case Study. These formed Chapter 12. Vishal Sharma and Abhay Karandikar wrote Chapter 13, "QoS in Wireless Networks." Their generous contributions are greatly appreciated.

Rick Adams, Rachel Roumeliotis, and Greg Chalson at Elsevier managed to drag me to the finish line. Without them, this book would still be in writing. Rick was very involved in the development of the manuscript, and provided useful guidance and candid feedback.

Last but not the least, I would like to thank Amy Chen for taking care of our son Alan Xiao, thus allowing me to concentrate completely on the book during the crunching stage.

XiPeng Xiao

About the Author

XiPeng Xiao has a unique background in QoS. He did his Ph.D. thesis on QoS at Michigan State University. This gave him a strong theoretical background on QoS. The author has product management experience with multiple network equipment vendors in Silicon Valley and network operations experience with a global network service provider. This vendor experience let him know the implementation cost of various QoS mechanisms and the hidden caveats behind those mechanisms. The NSP experience let him understand the practical trade-off that network operators must make between the need for network control and the need for network simplicity. The author also participates in international standard organizations such as the Internet Engineering Task Force (IETF), the standard organization that drives the technology development of IP networking. This let him know what's going on in the industry in terms of QoS, and what other people are thinking and doing. But maybe most importantly, he is a technologist-turned-marketing person. At the network equipment vendors, he is responsible for managing product lines for market success. This forced him to look beyond technology and develop the sensitivity to business and regulatory issues. This unique background with research experience, vendor experience, provider experience, standard experience, and business experience enables him to see the big picture of QoS, which comprises of a technical aspect, a commercial aspect, and a regulatory aspect.

With this unique background, the author has made contribution to QoS before writing this book. In 1999, while deploying QoS and MPLS Traffic Engineering in Global Crossing's network, the author and his colleagues discovered a conflict between Diffserv and TCP. They proposed a solution which was eventually standardized as [RFC2873]. This book is his continuous effort to describe the QoS big picture. In 1999, he published "Internet QoS: A Big Picture" in *IEEE Networks Magazine*. According to Google Scholar (scholar.google.com), it is among the most quoted QoS articles. Over the years, the author published multiple RFCs and journal articles in the fields related to QoS. The author also made many presentations on QoS in many industrial forums and conferences.

Introduction

This chapter will discuss why we need a big picture of Internet Quality of Service (QoS) and will provide a high-level overview of how the Internet works for people who are not familiar with it.

THE BIG PICTURE

In November 2005, in an interview with *Business Week*, Edward E. Whitacre, Jr., Chairman and CEO of SBC (now AT&T), stated:

“Why should they be allowed to use my pipes? The Internet can’t be free in that sense, because we and the cable companies have made an investment and for a Google or Yahoo! or Vonage or anybody to expect to use these pipes free is nuts.” [Bizweek].

This statement triggered a Net Neutrality debate between the Internet Content Providers (ICPs) (e.g., Google, Yahoo!, etc.) and Network Service Providers (NSPs) (e.g., AT&T, Comcast, etc.). The debate soon spread into the general public, the U.S. Senate, the House of Representatives, and the Federal Communications Commission (FCC). To date, controversial Net Neutrality legislation is still being discussed.

One of the important aspects of Net Neutrality is whether NSPs should be allowed to charge a QoS fee, and prioritize users’ traffic accordingly. Many people in the IT/network/telecom industries actively participate in this debate. A large amount of online discussions have been generated.

From following the debates on Net Neutrality/QoS, we frequently see insights mixed with misconceptions, and facts mixed with personal beliefs. It is probably safe to say that the people who take the time to express their opinion on Net Neutrality have above-average knowledge on QoS. After all, they at least formed their own opinion. The fact that their opinions are so divergent indicates that after many years of research, standardization, implementation, and claimed deployment, QoS is still elusive and therefore intriguing.

The following is a list of topics that are vigorously debated in the Net Neutrality controversy. As a reader, if you agree to many of the following points, then this book can be useful for you, because it will provide the opposite views on these topics. This is not to say that the following points are all wrong, but to say that there is the other side of the story. After reading this book, you will get some new perspectives, and you may find that some of these points are not so agreeable anymore.

- The basic goal of QoS is to use Differentiated Service (Diffserv) and traffic management to create one or multiple Classes of Services (CoS's) that are better than Best Effort.
- NSPs need to invest a large amount of resources to provide QoS; they should be allowed to sell QoS to recoup some of their investment.
- If NSPs are not allowed to sell QoS, they will just provide raw bandwidth and will be reduced to “dumb pipers” like utility companies.
- QoS is a commercial and technical matter; government should stay out of it.
- Selling QoS will bring in additional revenue for the NSPs. The money will trickle down the value chain. Therefore, it's good for the whole industry.
- When different traffics are prioritized differently, users will be able to tell the difference between the different CoS's to pick the right CoS for their applications.
- Consumers and enterprises who want better service than today's Best-Effort service will buy a higher CoS if one is available.
- It doesn't matter whether business or consumers, senders or receivers, pay for QoS, as long as somebody pays for it.
- Traffic should be prioritized so that if there is congestion, high-priority traffic can be preferred. This can't be wrong.
- Relying on sufficient capacity to deliver QoS (i.e., over-provisioning) is a more expensive QoS approach compared to others.
- TCP is “greedy” (i.e., it will hunt for the largest possible sending rate). Therefore you can never have too much bandwidth because TCP will eat it up.
- Because data traffic is self-similar (i.e., still very bursty even after a large amount of aggregation), transient congestion is unavoidable.
- Network reliability/availability is important. But it is orthogonal to QoS. QoS and reliability can be tackled separately.
- QoS is not a reality today because the QoS mechanisms provided by network equipment vendors are not matured enough for the NSPs to deploy. When such mechanisms mature, NSPs will deploy them, and QoS will become reality.

From our perspective, the divergent QoS opinions indicate that people have a silo-ed view on QoS. That, in turn, indicates that existing QoS literature didn't present a QoS big picture.

First, there is a lack of coverage on the commercial and regulatory issues. QoS literature is too focused on the technical side. Since 1998, more than ten QoS books have been published [Ferguson][Armitage][Huston][Vegesna][Wang][Jha]

[Park][Szigeti][Alvarez][Soldani][Evans]. But none of them has coverage on the regulatory side and only a few [Ferguson][Huston] have scarce coverage on the commercial side.

The lack of coverage on the commercial side may be a little puzzling, given that there are a fair number of discussions on it. For example, the difficulty in selling QoS as an add-on to network connectivity has long been discussed among network operators. This is considered part of the “free mentality” problem of the Internet (for NSPs): After buying network connectivity, Internet users think that everything else on the Internet should be free. Also, since the inception of QoS, there is always this debate of “QoS vs. no QoS.” (This name can be a little deceiving. The debate is really about “use mechanisms to provide QoS” vs. “use capacity to provide QoS”.) This debate involves technical as well as commercial issues, that is, which approach is more feasible and which approach is more economic. Given these, the lack of commercial coverage in the QoS literature can be a little puzzling. We believe that this is partly because network operators are not very keen on contributing to the QoS literature. Furthermore, certain debates are done in a philosophical or even religious way, not in a logical way. For example, in many cases, each debating party simply takes its own assumptions for granted and does not bother to validate them. The basic argument became “I am right and you are wrong. I know it.” There is little value to add that to the QoS literature.

The lack of coverage on the regulatory side is less of a surprise. After all, the Net Neutrality debate only started at the end of 2005 and the earliest Net Neutrality legislation was proposed in 2006. So it is no wonder that some people are shocked to hear that the government would even consider forbidding NSPs from selling QoS. “Why does the government have any business in QoS?” these people would think. The truth is, as the Internet became an integral part of ordinary people’s daily life, traffic prioritization and the related pricing (i.e., QoS) can also affect ordinary people’s daily lives. Because of its relevance to the mass people, QoS inevitably becomes a topic of interest to government regulators. If we further consider the phenomenon that some people think that introducing traffic prioritization would move the Internet from being neutral today to being unneutral tomorrow, and would deprive Internet users of the freedom to pick their favorite applications and application providers, then the government’s involvement should not be a surprise. This is not to mention that the charging of QoS, or the lack of, involves an interest conflict between the ICPs, such as Google and MSN, and the NSPs, such as AT&T and Verizon, both with deep pockets and strong lobbying power. Therefore, the fact that some people are shocked does not say that government involvement is absurd. It just reflects a void of regulatory coverage in the QoS literature.

Second, even on the technical side, the technical big picture has not been clear. Technical QoS literature has been too focused on using Diffserv and traffic management to deal with traffic congestion. But many important questions are left unanswered. For example, how often will congestion happen? If congestion is too rare or too often, then Diffserv may not be useful. (If it is too often, something else

is fundamentally wrong, and that needs to be fixed first.) Also, when congestion happens, is it better to handle it with Diffserv and traffic management, or is it better to find some other ways to relieve congestion? If the latter can be done most of the time, Diffserv/traffic management won't be useful either. Also, will the network have too much high-priority traffic to render Diffserv ineffective? If video becomes priority traffic, this would become true. The failure to ask and answer these questions indicates a lack of a technical big picture.

To a certain extent, we believe that this lack of a big picture is somewhat inevitable. Modern people are highly specialized. Few people have the opportunity to develop a comprehensive view on QoS. Today, the majority of QoS literature is authored by academic people. But academic people generally do not have the opportunity to develop a comprehensive industrial view. Economic and legal experts may have an industrial view. But they may not have the technical depth to tackle QoS head on. Their works are relatively focused on the economic and legal aspects of Internet services, and may be considered by normal QoS people as irrelevant to QoS. Industrial experts may be the most suitable people to present the big picture. But it is still not easy to find people with a strong background in all three aspects: technical, commercial, and regulatory. The people who have such expertise may not have the time to write either.

The effect of this lack of a big picture is inefficiency. We list a few examples below.

First, without a QoS big picture, it is difficult for people to communicate effectively. It is rare that people clearly list their assumptions when they express their opinions. They either don't realize that their arguments are based on certain assumptions, or assume that the other party accepts the same assumptions. Therefore, if both parties have a silo-ed view, and their silos are different (i.e., their assumptions are significantly different), then they can only wonder why the other party "just doesn't get it." This happens in many online Net Neutrality debates. This is also why some people are surprised to find that there are other people who would disagree with the bulleted points listed at the beginning of this chapter—those points seem like truisms to them.

Second, without a QoS big picture, people may come to a suboptimal decision on QoS. For example, on the technical side, if some network operators think that Diffserv/traffic management is the primary way to enable QoS, they may introduce a lot of Diffserv/traffic management mechanisms into their network. By the time they realize that Diffserv/traffic management alone cannot provide a satisfactory QoS solution, too much complexity may have crept into the network. On the commercial side, if a NSP fails to recognize that there are significant commercial challenges in selling QoS, it will likely fail to prepare a strategy to deal with those challenges. Consequently, its technical investment to enable QoS may not get good return. On the regulatory side, if some NSP people do not fully understand the argument of the Net Neutrality proponents, they may think that Net Neutrality has little validity. In an effort to get it over with, those people may attack Net Neutrality. That will have the exact opposite effect. The harder anybody tries to

discredit Net Neutrality, the stronger its momentum will be. NSP's best strategy to avoid Net Neutrality legislation may be "do nothing"—at least in public.

Third, without a big picture, current QoS models overlooked several key factors. This reduces the commercial viability of QoS. That may be the reason why there are few commercial QoS success stories today. One of these factors is "user perceivability." When traffic prioritization/class of service is proposed, it is automatically assumed that their effect can be clearly perceived by the end users. "User perceivability" is critical because without it, users won't buy QoS (i.e., the higher CoS). The lack of effort to make sure that the effect of a QoS solution can be clearly felt by the end users reflects a lack of commercial consideration, and thus a lack of a big picture. Reliability is often another major overlooked factor. A significant portion of today's QoS problems are caused by human errors and equipment failures. QoS mechanisms can introduce additional complexity and reduce reliability. But existing QoS literature rarely considers the reliability impact of the proposed QoS mechanisms. This reflects a lack of a technical big picture.

Therefore, it is desirable to have a QoS big picture. This book is an effort towards that.

HIGH-LEVEL OVERVIEW OF HOW THE INTERNET WORKS

In this section, we give an overview of how the Internet works for people outside the network industry. This includes what the Internet looks like and how the major control components fit together.

At the highest level, the Internet consists of a number of Autonomous Systems (ASs) interconnected together, as shown in [Figure 1-1](#).

Each AS is basically a NSP, e.g., AT&T or Deutsche Telekom. Each AS has its own customers. For example, many residential users are customers of AT&T or Verizon, while Google and Yahoo! are customers of Level3. Note that on the Internet, each customer is represented by its IP address. Via a protocol called external Border Gateway Protocol (eBGP), each AS tells other ASs who their customers are. Each AS also helps other ASs propagate their words. For example, AS2 in the


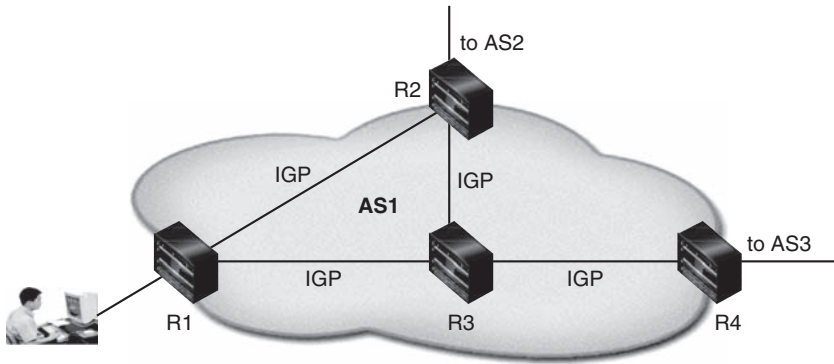


FIGURE 1-1

High-level view of the Internet

**FIGURE 1-2**

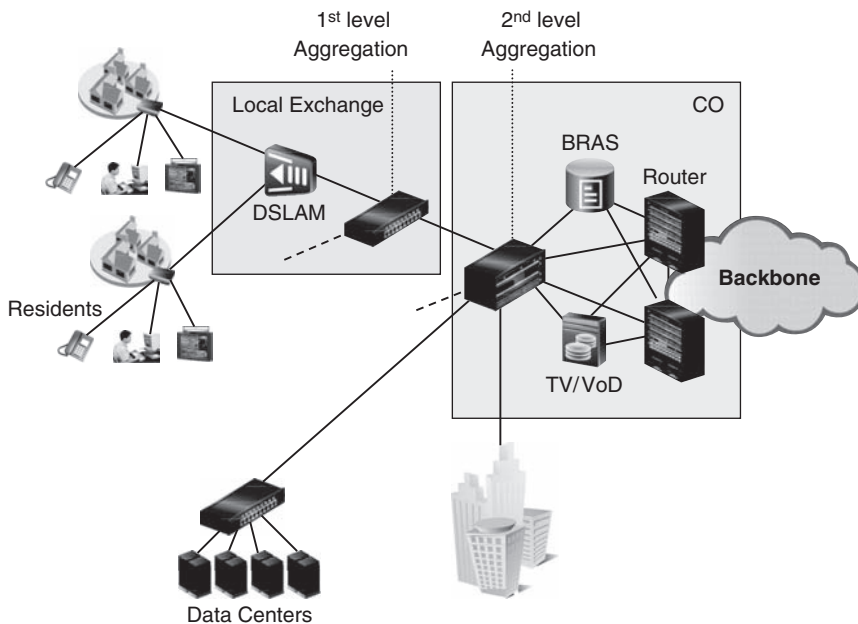
Network and control inside an AS

picture will tell AS1 that Google is a customer of AS5. This way, each AS will eventually learn which customer is in what AS. Because AS1 will hear from both AS2 and AS3 that Google is at AS5, AS1 also knows that it can get to Google through both AS2 and AS3. BGP also tells AS1 that to get to Google through AS2 involves two ASs, AS2 and AS5, while to get to AS5 through AS3 involves at least three ASs. All things being equal, AS1 will pick AS2 because the AS_Path is shorter. However, network operators at AS1 can override this decision. For example, if AS3 charge AS1 less money per unit of bandwidth, AS1 may decide to send all traffic to Google through AS3 instead. The use of a policy to override default routing protocol decision is called Policy-Based Routing.

Now let's zoom into each AS, using AS1 as an example, as shown in Figure 1-2.

For discussion purposes, let's assume that AS1 has four sites, each represented by a router, partially meshed together as showed in Figure 1-2. R1, R2, and R4 have connections to the outside world and have learned some routes to the external destinations. Inside AS1, these routers use the internal BGP (iBGP) to exchange reachability information. With iBGP, R1 will learn that to get to destinations in AS3, its next hop router should be R4 although it is not directly connected. How R1 can get to R4 inside AS1 is resolved by an Interior Gateway Protocol (IGP). The commonly used IGPs are Intermediate System-to-Intermediate System (IS-IS) [RFC1195] and Open Shortest Path First (OSPF) [RFC2328]. The reason both BGP and IGP are used inside an AS is that BGP is for managing external routes (i.e., external destinations) and IGP is for managing internal routes. Inside an AS, both types of routes are involved, therefore both protocols are needed.

BGP and IGP jointly determine the traffic distribution inside an AS. Sometimes the network operators of an AS may want to change the traffic distribution. For example, there may be too much traffic over the R1-R2 path for some reason. Therefore, the network operators want to offload some of the traffic from the R1-R2 path to the R1-R3-R2 path.

**FIGURE 1-3**

Metro network architecture

Usually, the network operators would start with IGP by making the direct R1-R2 path and the R1-R3-R2 path have equal IGP metric (i.e., length). This would cause traffic from R1 to R2 to be distributed on both paths. This practice is called IGP traffic engineering. But sometimes making multiple paths have equal cost can be difficult to do or can have undesirable side effects. This is why Multi-Protocol Label Switching (MPLS) traffic engineering is introduced. MPLS traffic engineering allows the use of multiple paths with different IGP metrics for load sharing. This kind of intra-domain traffic engineering is transparent to the outside world.

Sometimes, the network operators can also use BGP to change traffic distribution. For example, the reason there is too much traffic from R1 to R2 may be AS1 is using AS2 (connected via R2) to get to too many destinations. AS1 may decide to get to some of those destinations via AS3 instead by changing its policy. Naturally, the amount of traffic from R1 to R2 will be reduced. This kind of inter-domain traffic engineering will affect traffic distribution in other NSPs.

The above are concerned with the Internet backbone, i.e., the Wide Area Network (WAN), which covers multiple metropolitan areas. Now let's zoom into the Metro Area Network (MAN). [Figure 1-3](#) provides a detailed metro network diagram. It is provided because not many people know what a metro network looks like.

In the metro, there are three groups of users, residential, enterprises, and data centers. The first two groups are content users while the third are the content

providers. Residential users are connected via Digital Subscriber Line (DSL) or cable modem or Passive Optical Network (PON). The access devices, e.g., the DSL Access Multiplexer (DSLAMs), are connected by an aggregation network to the Central Office. The aggregation network usually has two levels of devices, with the first-level devices in the Local Exchange Office and the second-level devices in the Central Office. In each metro area, there are usually dozens of Local Exchange Offices and one or two Central Offices. Residential aggregation networks used to be based on ATM. They are migrating to Ethernet. Enterprises' access routers are connected to the Central Office via TDM links (e.g., T1/E1's) or Ethernet links. Similarly, servers in Data Centers are first aggregated by an Ethernet switch fabric, and connected via an Ethernet link to the Central Office. Note that although all the links connecting to the Central Office are showed as point-to-point links at Layer 2, they are usually transported over a ring-like optical network at Layer 1, with Ethernet mapped over SONET/SDH if necessary. Sometimes, direct fiber links may be used for Ethernet. This Layer 1 transport network is usually common for residential, enterprise, and data center customers. Also, to increase network availability, each device in the metro aggregation network may be connected to two higher-level aggregation devices, although this is not showed in the diagram. For example, the DSLAM may be connected to two first-level aggregation devices, and each first-level aggregation device may be connected to two second-level aggregation devices.

Inside the Central Offices, in addition to the second-level aggregation devices, there are also routers, Broadband Remote Access Servers (BRAS's), Video on Demand (VoD) servers, and possibly other servers. The BRAS's are the control devices for residential users' Internet access service. Other servers are for their own respective services. These servers are usually dual-homed to the routers. The routers are those shown in [Figure 1-2](#). They are connected to routers in other metro areas.

Metro aggregation networks are usually Layer 2 networks. Therefore, there is no routing protocol involved. From a Layer 3 perspective, customer's access routers (i.e., the residential DSL modems or the enterprise CPE routers) are directly connected to the BRAS and the routers in the Central Office, respectively.

Up to this point, we have presented what the Internet looks like and how the major control protocols fit together to work. To recap, each domain (i.e., AS) has customers connected to it in some way, as shown in [Figure 1-3](#). Each domain will announce its customers' IP address via BGP to other domains so that every domain knows what domains to go through to get to each destination. Inside each domain, IGP tells each router how to send packets to another router. Together, BGP and IGP determine the packet route between any two points. This way, communication between any two end points can ensue.

In relation to QoS, we would like to point out that the likely communication bottlenecks in the Internet are the links between the DSLAMs and the first-level aggregation devices, and the links between ASs. The former can become a bottleneck because NSPs assume that their residential users will rarely send/receive at

maximum speed, or at least not at the same time. Therefore, they allow the sum of all customers' maximum access speed to be much higher than the DSLAM uplink's speed. Sometimes the ratio can be as high as 20. However, peer-to-peer file sharing is causing this assumption to be violated. Consequently, the DSLAM uplinks may become congested. The links between ASs can also become bottlenecks because upgrading them involves agreement between two different organizations. Sometimes that can take a long time. Consequently, traffic outgrows link speed and congestion happens. But, in general, network operators have a mental problem with chronic congestion in their network. So they will try to upgrade link speed to relieve that, no matter how tight the budget is.

Sometimes, the last-mile links, e.g., the DSL local loops or the enterprise access links, can also become bottlenecked. But today, most customers have the option to upgrade to a higher-speed link. Therefore, when congestion happens at the last mile, it is usually by the customer's own choice, mostly to save money.

If you have questions, please ask them at <http://groups.google.com/group/qos-challenges>.

PART

The Status Quo

1

This part contains four chapters:

- Chapter 2 discusses what QoS means in this book, common applications' requirements on QoS, and the degree to which the current Internet meets those requirements.
- Chapter 3 discusses the historic evolution of QoS solutions.
- Chapter 4 discusses the current “mainstream QoS wisdom,” including its business model for QoS and technical solution for QoS.
- Chapter 5 discusses the network reality related to QoS, especially from a commercial perspective.

The purpose of discussing the QoS requirements of the applications is to make the objectives of QoS clear. The purpose of discussing the degree to which the current Internet meets those requirements is to let us know the gap between what is needed and what is available, so that we know what else may be needed to deliver QoS. The purpose of discussing the historic evolution of QoS solutions is to provide some technical background. The purpose of discussing the current QoS business model and its technical solution is to provide a base for commercial, regulatory, and technical examination. The purpose of discussing the commercial reality related to QoS is to give us a sense of how well the traditional QoS wisdom works.