

Preface

Computer networks are like any complex engineering project. A small network can be slapped together quite successfully with minimal experience. But a larger network requires careful thought and planning. As with many types of engineering projects, this planning and design phase is best served by an organized and disciplined design philosophy. The trouble with design is that it is difficult to differentiate between personal or near-religious biases and sound ecumenical strategies that can result in better usability, stability, security, and manageability.

Everyone has religious biases when it comes to network design. This is because most networks are so complex that a feeling of black magic falls over anybody trying to understand them. They tend to be too large and too intricate to hold in your mind all at once. So when some particular incantation appears to work miracles, it is adopted as an article of faith. And when a vendor's equipment (or support engineer) saves the day in some important way, it can turn into a blanket belief in that vendor as savior.

So, in the interests of making plain my assumptions and biases, let me explain right from the start that I am a network agnostic. I have used equipment from most of the major vendors, and I believe that every individual piece of gear has its pluses and minuses. I prefer to use the gear that is right for the job, rather than expressing a blind devotion to one or another. So this book is vendor neutral.

I will discuss some proprietary protocols and standards because these are often the best for a particular situation. But in general I will try to lead the reader towards open industry standards: I believe that it is unwise to lock your technology budget to one particular vendor.

In the mainframe-computing era, many firms spent large amounts of money on one company's equipment. Then they found that this required them to continue spending their hardware budget with that company unless they wanted to abandon their initial investment. All incremental upgrades merely reinforced their dependency on

this one vendor. This was fine unless another manufacturer came along with gear that would be better (cheaper, faster, more scalable, etc.) for important business requirements of the company. It is wise to avoid the “fork-lift” upgrade where the entire infrastructure has to be replaced simultaneously to improve performance.

In practice, most LANs are multivendor hybrids. This may be by design or by chance. In many cases a best-of-breed philosophy has been adopted so that a particular type of Ethernet switch is used in the wiring closets, another type at the backbone, with routers from another vendor, while ATM switches and long-haul equipment are provided by still other vendors. In other cases, the multivendor nature of the network is more of an historical accident than intention. And there are also cases where all or nearly all of the network hardware comes from the same manufacturer. If this is the case, then the choice should be made consciously, based on solid technical and business reasons. Having stated my biases here, I leave the reader to make these decisions freely.

Because computer networks are large and complex engineering projects, they should be designed carefully and deliberately. There are many important questions to ask about how a network should function and what purposes it needs to serve. And there are even more questions to ask about how best to meet these objectives. This book will serve as a guide to this process.

Audience

This book is intended for anybody who needs to build or maintain a large-scale network. It is not a theoretical book for classroom use, and it isn’t intended to help programmers with designing applications. Instead it is a hands-on set of rules, guiding principles, and useful tips for people who build networks.

So it should be useful for network-operations people who need to understand the overall logic of their network. It should also be helpful to engineers who need to think about upgrading parts of an existing network in a logical way. And it is particularly relevant for network designers and architects.

In short, this is the book that I always wanted to read when I was starting to work on large networks.

Organization

The book begins in Chapter 1, *Networking Objectives*, by discussing the most important question of all: why build a network in the first place? The answers to this question shape everything that follows. This first chapter also reviews several of the basic networking concepts used throughout the book.

Chapter 2, *Elements of Reliability*, discusses reliability in networks, the factors that lead to a solid network. This includes discussion of how to find the single points of failure in your network. It also includes an important section on how to evaluate the stability of your network and its components statistically.

Chapter 3, *Design Types*, describes many of the most successful design types and their strengths and weaknesses. This chapter should help you to decide on the large-scale shape of the network you build. It includes many ideas for ensuring both reliability and scalability in a large-scale network.

Chapter 4, *Local Area Network Technologies*, delves into the technologies commonly used on LANs. This is intended as a guide to the network designer and implementer, rather than the engineer building the hardware. There are many other books that provide that higher level of detail, which is beyond the scope of this book. The intent here is to provide the information that a network designer needs to make appropriate decisions.

The same is true of Chapter 5, *IP*, which begins a three-chapter discussion of the various Layer 3 and 4 protocols that are commonly used on LANs, focusing in particular on TCP/IP (see Chapter 6, *IP Dynamic Routing*) and IPX (see Chapter 7, *IPX*). At the same time, these chapters look at good ways of implementing networks based on these protocols, as well as appropriate dynamic routing protocols.

Then in Chapter 8, *Elements of Efficiency*, I turn to efficiency. What is meant by efficiency in a network? How is it achieved? In particular, I discuss how to implement Quality of Service (QoS) through a network to ensure that low-priority traffic doesn't interfere with the delivery of important data.

Chapter 9, *Network Management*, is devoted to network management. But, rather than looking at how to manage a network, this section focuses on how a network's design can make it either easier or more difficult to manage. There are several key design decisions that have implications on how the network will later be managed. Since manageability is one of the keys to reliability, it is important to design networks so that they can be managed effectively.

Chapter 10, *Special Topics*, discusses other important considerations that may not be relevant to every network. These include issues of LAN security, designing for multicast capabilities, and inclusion of the new IPv6 protocol.

The Appendix describes in some mathematical detail how to combine statistical probabilities. This is important for estimating failure rates in a complex network. This information is particularly used in Chapter 2 in the section on calculating "Mean time between failures."

The Glossary is a listing of networking terms. Networking has unfortunately become bogged down with highly technical jargon that often makes it difficult to understand things that are actually relatively simple. So this glossary is included to help prevent that from being a problem.

Conventions Used in This Book

I have used the following formatting conventions in this book:

Italic

Used for emphasis and the first use of technical terms, as well as email and URL addresses.

Constant Width

Used for MAC and IP addresses.

Comments and Questions

Please address comments and questions concerning this book to the publisher:

O'Reilly & Associates, Inc.
1005 Gravenstein Highway North
Sebastopol, CA 95472
(800) 998-9938 (in the United States or Canada)
(707) 829-0515 (international/local)
(707) 829-0104 (fax)

There is a web page for this book, which lists errata, examples, or any additional information. You can access this page at:

<http://www.oreilly.com/catalog/lgscalelans/>

To comment or ask technical questions about this book, send email to:

bookquestions@oreilly.com

For more information about books, conferences, Resource Centers, and the O'Reilly Network, see the O'Reilly web site at:

<http://www.oreilly.com>

Acknowledgments

There are several people who were instrumental in helping me get this book completed. Most important was my editor, Jim Sumser. Jim was extremely encouraging, which was important because of the aggressive schedule that I wanted to follow. He also really helped me organize and present these (often scattered) ideas in a rational way.

I'd like to thank my technical reviewers, who each pointed out more than one serious blunder. Andre Paree-Huff made a number of encouraging comments about the book. Ian J. Brown, CCIE# 3372, showed his vast depth of technical knowledge in his comments, which allowed me to make significant improvements to the quality of



the information presented. And Scott M. Ballew offered many important structural, as well as technical, suggestions for which I am grateful.

In addition, I'd like to thank the following people at O'Reilly: Sheryl Avruch, Mary Brady, Claire Cloutier, Emma Colby, Tatiana Apandi Diaz, Jeffrey Holcomb, Mihaela Maier, Jessamyn Read, Ann Schirmer, Ellie Volckhausen, Melanie Wang, Rachel Wheeler, and Sue Willing.

And, I am particularly grateful to Sherry Biscope for her unconditional support and encouragement throughout this project.



