



High Availability Network Fundamentals

By [Chris Oggerino](#)

.....
 Publisher: **Cisco Press**
 Pub Date: **May 08, 2001**
 Print ISBN-10: **1-58713-017-3**
 Print ISBN-13: **978-1-58713-017-5**
 Pages: **250**
 Slots: **2.0**

[Table of Contents](#) | [Index](#)
Overview

High Availability Network Fundamentals discusses the need for and the mathematics of availability, then moves on to cover the issues affecting availability, including hardware, software, design strategies, human error, and environmental considerations. After setting up the range of common problems, it then delves into the details of how to design networks for fault tolerance and provides sample calculations for specific systems. Also included is a complete, end-to-end example showing availability calculations for a sample network.



High Availability Network Fundamentals

By [Chris Oggerino](#)

.....
 Publisher: **Cisco Press**
 Pub Date: **May 08, 2001**
 Print ISBN-10: **1-58713-017-3**
 Print ISBN-13: **978-1-58713-017-5**
 Pages: **250**
 Slots: **2.0**

[Table of Contents](#) | [Index](#)

[Copyright](#)
[About the Author](#)
[About the Technical Reviewers](#)
[Acknowledgments](#)
[Introduction](#)
 [Goals and Methods](#)
 [Who Should Read This Book?](#)
 [Useful References](#)
 [How This Book Is Organized](#)
 [Icons Used in This Book](#)
 [Command Syntax Conventions](#)
[Part I: Introduction to Availability](#)

 Chapter 1.

Introduction to
High Availability
Networking

Why Do We
Need High
Availability?
Presenting and
Describing High
Availability
Measurements
Additional
Considerations
in Measuring
Availability

 Chapter 2.

The Basic
Mathematics of
High Availability
Determining the
Availability of
Network Device
Components
Determining the
Availability of a
Single
Component
Determining the
Availability of
Multiple
Components
Determining
Data Flow in a
Network: Path
Analysis

 Chapter 3.

Network
Topology
Fundamentals
Serial Topology
Parallel
Topology
Serial/Parallel
Topology
Summary

Part II: Predicting
Availability

 Chapter 4.

Factors That
Affect Availability
Predicting
Hardware
Availability
Predicting

Software
 Availability
 Predicting
 Availability
 Considering the
 Environment
 Including
 Human Error
 and Process in
 Availability
 Calculations
 Network Design
 Summary
 References
 Used in This
 Chapter

▣ Chapter 5.

Predicting
 End-to-End
 Network
 Availability: The
 Divide-and-Conquer
 Method
 The
 Divide-and-Conquer
 Steps
 A VoIP
 Network
 Example
 Designing
 Networks for
 Availability
 Goals
 Summary

Part III: Examples
 of Analyzing
 Real-World
 Availability

▣ Chapter 6.

Three Cisco
 Products: An
 Availability
 Analysis
 Cisco uBR 924
 Availability
 Calculations
 Cisco uBR 7246
 Availability
 Calculations
 Cisco 12000
 Availability
 Calculations

▣ Chapter 7. A
 Small ISP

Network: An
 Availability
 Analysis
 The Small
 Internet Service
 Provider
 Network
 Scenario 1 of
 The Small ISP
 Example
 Scenario 2 of
 The Small ISP
 Example
 Summary
 Chapter 8. An
 Enterprise
 Network: An
 Availability
 Analysis
 System Level
 Calculations for
 an Enterprise
 Network
 The Downtime
 from Lost
 Power for an
 Enterprise
 Network
 Network
 Calculations for
 an Enterprise
 Network
 Human Error
 and Process
 Contribution to
 Downtime in an
 Enterprise
 Network
 Summary
 Chapter 9. A
 Large VoIP
 Network: An
 Availability
 Analysis
 A VoIP over
 Cable Network
 The Availability
 Scenarios of
 VoIP over HFC
 The System
 Level
 Calculations for
 VoIP over HFC
 Network Level

[Availability](#)
[Calculations](#)
[Major Network](#)
[Service](#)
[Construct](#)
[Availability](#)
[Calculating](#)
[Scenario](#)
[Availability](#)
[Summary](#)
 [Appendix A.](#)
[The Contents of the](#)
[CD](#)
[Computer](#)
[Requirements](#)
[Basic Instructions](#)
[for Using the CD](#)
[Chapter by](#)
[Chapter Contents](#)
[Using the SHARC](#)
[Spreadsheet](#)
[Index](#)

Copyright

Copyright© 2001 Cisco Press

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:

Cisco Press

201 West 103rd Street

Indianapolis, IN 46290 USA

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

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

Library of Congress Cataloging-in-Publication Number: 00-105374

First Printing April 2001

Warning and Disclaimer

This book is designed to provide information about predicting high availability for Cisco product-based networks. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The author, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

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

Feedback Information

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

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

We greatly appreciate your assistance.

Credits

Publisher

John Wait

Editor-In-Chief

John Kane

Cisco Systems Program Manager

Bob Anstey

Managing Editor

Patrick Kanouse

Acquisitions Editor

Amy Lewis

Development Editor

Christopher Cleveland

Copy Editor

Chuck Gose

Technical Editor(s)

Scott Cherf, Oliver Day, BJ Favaro

Team Coordinator

Tammi Ross

Book Designer

Gina Rexrode

Cover Designer

Louisa Klucznik

Production Team

Steve Gifford

Indexer

Tim Wright

Corporate Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 526-4100

European Headquarters

Cisco Systems Europe

11 Rue Camille Desmoulins

92782 Issy-les-Moulineaux

Cedex 9

France

<http://www-europe.cisco.com>

Tel: 33 1 58 04 60 00

Fax: 33 1 58 04 61 00

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-7660

Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems Australia, Pty., Ltd

Level 17, 99 Walker Street

North Sydney

NSW 2059 Australia

<http://www.cisco.com>

Tel: +61 2 8448 7100

Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2000, Cisco Systems, Inc. All rights reserved. Access Registrar, AccessPath, Are You Ready, ATM Director, Browse with Me, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCSI, CD-PAC, CiscoLink, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco Systems Networking Academy, Fast Step, FireRunner, Follow Me Browsing, FormShare, GigaStack, IGX, Intelligence in the Optical Core, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, iQuick Study, iQ Readiness Scorecard, The iQ

Logo, Kernel Proxy, MGX, Natural Network Viewer, Network Registrar, the Networkers logo, Packet, PIX, Point and Click Internetworking, Policy Builder, RateMUX, ReyMaster, ReyView, ScriptShare, Secure Script, Shop with Me, SlideCast, SMARTnet, SVX, TrafficDirector, TransPath, VlanDirector, Voice LAN, Wavelength Router, Workgroup Director, and Workgroup Stack are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, Cisco, the Cisco Certified Internetwork Expert Logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Collision Free, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Post-Routing, Pre-Routing, Registrar, StrataView Plus, Stratm, SwitchProbe, TeleRouter, are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. and certain other countries.

All other brands, names, or trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0010R)

Dedication

This book is dedicated to Annette, Alex, and our soon-to-be-born daughter.

About the Author

Chris Oggerino has been employed by Cisco Systems, Inc., for over five years and is currently a Serviceability Design Engineer. As a Serviceability Design Engineer at Cisco Systems, Chris spends his time improving the reliability, availability, serviceability, and usability of Cisco products. Prior to his employment with Cisco Systems, Chris spent six years doing technical support of UNIX and internetworking products, three years in microcomputer corporate sales, and four years as a programmer at a variety of companies in the Bay Area. Chris currently resides in Los Gatos, California. You can reach Chris Oggerino at chris@oggerino.org.

About the Technical Reviewers

Scott Cherf is a Technical Leader in Cisco's IOS Technologies Division. Scott pioneered the measurement and estimation of Cisco's IOS software reliability in 1996. Since then he has presented his work both to Cisco's customers and to professional organizations. In 1999 he was invited to present this work at the IEEE's 10th annual International Symposium on Software Reliability Engineering. Scott was recently awarded a U.S. patent in the area of high availability systems design and continues to provide direction to Cisco engineers working on highly available Systems. Prior to joining Cisco in 1994, Scott spent 12 years as an engineer for Tandem Computers, Inc., where he participated in the design and implementation of the first commercial fault tolerant operating system. Scott has authored several internationally published papers on software metrics and has served as a technical editor for IEEE Computer magazine. Scott lives with his wife and two children in the town of Jackson, Wyoming.

Oliver Day is the founder and CTO of Digital Knox, a high security storage service provider. Oliver's research focuses on NAS, SAN, encryption, and network security. He has spent the last four years deploying enterprise solutions for Fortune 1000 companies.

B.J. Favaro Jr., ASQ Certified Reliability Engineer #4642, is a Senior Quality Engineer in Cisco's Manufacturing organization. B.J. is responsible for reliability prediction methods, reliability testing, and field reliability analysis at Cisco. Prior to joining Cisco in 1997, B.J. worked in a variety of management and engineering positions at Seagate Technology and taught mathematics at the high school and college level. B.J. was the principal developer of the SHARC spreadsheet tool provided with this book.

Acknowledgments

I want to thank B.J. Favaro and Scott Cherf first. Without their technical help and mentoring, this book would not have been possible. Scott's work on Cisco IOS Software availability was the catalyst for going from some availability research to an actual book. B.J.'s subject matter mentoring was a crucial part of the preparation and research for this book. B.J. and his team in Cisco Corporate Quality created the SHARC spreadsheet included on the CD.

Thanks to Terry Mar, Michael Shorts, and the other members of the Serviceability Design Engineering Team that supported my efforts.

Additional thanks go to Amy Lewis and Chris Cleveland from Cisco Press for their assistance in the production of this book. Amy, I think I could have kept this on the back burner for another year or two if you hadn't convinced me to get it done in time for the tradeshow. Chris, thanks for the graphics ideas and the solution to our e-mail problems!

I should also thank all the technical and grammatical editors that took the pieces that were delivered and made a book out of it.

Finally, a special thanks to my wife, Annette, for all those times she couldn't get my attention while I worked on this book thanks for your patience.

Introduction

Modern day networks have become crucial to our pursuit of life and happiness. Police departments, hospitals, businesses, and virtually anything we depend on runs on their networked computer systems. The more we depend on these networked computer systems, the more it affects us when they stop working.

For those of us that are involved in the planning, designing, building, and operation of these networks, the more we depend on the network, the more we need to predict problems in advance. Predicting problems in advance allows us to reduce the impact of these problems. With predictions of network availability, we can make sure our networks are going to service people satisfactorily before we build them.

Goals and Methods

The most important and somewhat obvious goal of this book is to show you how to predict availability of proposed network designs. Armed with that knowledge, you can make excellent decisions between alternative designs.

Historically, the methods used to compute availability include the use of advanced mathematics and complex processes. Those methods are the most accurate and provide the best results. However, most people that are responsible for networks do not have that mathematical background or the inclination to follow such complex, advanced processes.

This book presents simple methods for predicting availability that are acceptable in accuracy. By using only arithmetic, algebra, and a small amount of statistics, we can predict availability accurately enough to make good network design decisions

The CD included with this book includes examples and has built-in equations in order to allow you to perform network availability prediction with the least possible amount of math. What you will need to retain from the book will be the method to put your numbers into the spreadsheets.

Who Should Read This Book?

Anyone that is involved with the planning, designing, implementation, or operation of highly available networks is an appropriate reader for this book. The only exception to this list would be those that are already experts on high availability prediction and measurement.

If you are responsible for running a highly available network, you will find that this book includes only introductory material on operations. Operating a network for high availability is a science all by itself and the single section in this book is merely to remind the readers that the subject exists. Other authors are likely writing entire books on that subject.

Measuring the availability of a network is another area that deserves an entire book. Predicting the availability of a network on paper and measuring a real-life network are two completely different things. If your primary goal is to learn how to measure a network's availability, you will get only introductory material in this book. People measuring network availability, however, should definitely understand how to predict network availability. Comparing predictions with measurements enables identification of weak components in the network.

Useful References

Some of the books that I referred to for information about reliability/availability during the writing of this book targeted the advanced reliability engineer that really wants to be an expert (and use lots of calculus). In addition, I also referred to a couple of books that simply provided great information about how to build or run networks for availability (and don't use so much math). The first two books cited in the list that follows are for the advanced reliability engineer, while the second two books are geared more to building or running networks for ability.

- Lyu, Michael R., Editor. *The Handbook of Software Reliability Engineering*. McGraw Hill (ISBN: 0-07-039400-8)
- Kececioglu, Dimitri. *Reliability Engineering Handbook*. Prentice Hall (ISBN: 0-13-772302-4)
- Marcus, Evan and Hal Stern. *Blueprints for High Availability*. Wiley (ISBN: 0-47-135601-8)
- Jones, Vincent C. *High Availability Networking with Cisco*. Addison-Wesley (ISBN: 0-20-170455-2)

How This Book Is Organized

This book is divided into three major sections. [Chapters 1 through 3](#) are introductory in nature. They introduce the reader to the concepts of predicting availability at relatively high level. [Chapters 4 and 5](#) represent the tutorial portion of the book. These two chapters present the basic mathematics and processes for predicting availability and are presented in detail. [Chapters 6 through 9](#) are increasingly complex examples. [Chapter 6](#) takes the reader through the availability analysis of some routers. [Chapter 9](#) takes the reader through the availability analysis of a fairly complex network.

Because many of the chapters in this book call upon concepts presented in previous chapters, I recommend reading the entire book in the order presented. [Chapters 6 through 9](#) should be useful as reference material and reminders of how to perform a variety of common tasks in predicting network availability.

Summarized on a chapter-by-chapter basis, the book covers the following topics:

- **Chapter 1, "Introduction to High Availability Networking"** This chapter introduces the reader to the basic concepts used in the book. "Availability," "Mean Time Between Failure," "Mean Time To Repair," "Parallel Redundancy," and "Serial Redundancy" are the key concepts the reader should understand after reading [Chapter 1](#).
- **Chapter 2, "The Basic Mathematics of High Availability"** This chapter introduces the reader to the equations that will be used in the book. "The Availability Equation," "The Serial Availability Equation," and "The Parallel Availability Equation" are presented in this chapter. This basic presentation about the mathematics used in availability analysis should prepare the reader for the tutorial sections presented in [Chapters 4 and 5](#).
- **Chapter 3, "Network Topology Fundamentals"** This chapter is intended to refresh the reader on network topology as related to high availability. Network design that represents parallel and serial network design is presented to remind the reader of what a parallel construct looks like in a network and what a serial construct looks like in a network. These basic building blocks will be used to create larger and larger network designs as the reader moves through the book.
- **Chapter 4, "Factors that Affect Availability"** This chapter examines each of the five major things that contribute to network downtime. All too often, laypeople consider only hardware contributions to network downtime in their availability analysis. This chapter shows concrete methods for including hardware, software, environmental (power), human error and standard process, and network design issues in availability predictions.
- **Chapter 5, "Predicting End-to-End Network Availability: The Divide-and-Conquer Method"** This chapter presents a method that enables a layperson to take a large network and divide it up into conquerable sections. When faced with analyzing the availability of a large network, dividing it into smaller easier parts will be required. This chapter presents the specific processes to perform this work.
- **Chapter 6, "Three Cisco Products: An Availability Analysis"** This chapter begins our example section of the book. Starting with the simplest of devices, we work through example calculations showing how to perform the math and processes learned in the previous section. This chapter is meant to start the reader with the simplest of analysis and prepare the reader for the subsequent example, which will be more complex.
- **Chapter 7, "A Small ISP Network: An Availability Analysis"** This chapter takes our reader beyond the calculation of individual boxes and into the analysis of a small network. While the network is fairly simple, the chapter takes the reader through the basic process of analyzing the individual boxes and then combining the results in the network calculations. Subsequent chapters will use the techniques in this chapter. This chapter also introduces the use of the SHARC spreadsheet in order to speed up the simpler calculations previously done manually.
- **Chapter 8, "An Enterprise Network: An Availability Analysis"** This chapter presents an example that is similar in difficulty to the example in chapter seven. This example, however, is based on an enterprise type network as opposed to a service provider type network. While similar in process, networks for businesses look different from networks for large service providers and this chapter merely shows that a different looking network can be analyzed using the same tools and techniques.

- **Chapter 9, "A Large VoIP Network: An Availability Analysis"** This chapter presents the culmination of everything in the book. The analysis completed in this chapter uses every equation, tool, and process learned in the book. A full understanding of the availability analysis in this chapter will enable the reader to analyze virtually any network for predicted network availability.
- **Appendix A, "The Contents of the CD"** Appendix A describes the contents of the CD included with this book. The instructions for using the SHARC spreadsheet, included on the CD, are presented so that the reader understands how to use this application.

Icons Used in This Book



Router



Bridge



Hub



DSU/CSU

Catalyst
switch


Multilayer switch

ATM
switch



ISDN switch

Communication
server

Gateway



Access server



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets [] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.
- Boldface indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that

are manually input by the user (such as a show command).

- Italics indicate arguments for which you supply actual values.

Part I: Introduction to Availability

Chapter 1 Introduction to High Availability
Networking

Chapter 2 The Basic Mathematics of High
Availability

Chapter 3 Network Topology Fundamentals

Chapter 1. Introduction to High Availability Networking

Recently, the growth of the Internet and the use of computing systems to run businesses have blossomed in a way that few would have expected. In 1990, contacting individuals at some companies (mostly computer companies) via electronic mail was possible. Just a decade later, thousands of highly reputable companies were offering consumers the ability to actually purchase products from their web sites on the Internet. This fantastic growth of the Internet and networking has and is changing our lives in many ways.

As the Internet and networking become more a part of our lives, we are becoming dependent on them. We rely on them, and thus, we need them to be highly reliable. When you say that you want a highly reliable network, you are saying that you want your network to work all the time you want it to be highly available.

As you proceed through this book, you will find that I have attempted to cover the subject of high availability using only arithmetic and algebra. This is by design and is the key reason I believe many people will be able to understand the material. All the other books I have read about reliability and availability use advanced mathematics such as calculus. While most of us can perform calculus when absolutely required, calculus is not easy to remember, nor is it something most enjoy. I hope that you, the reader, appreciate that what we give up in accuracy and process, we make up for in ease of understanding and accomplishment.

For those of you that want to move on to advanced reliability and availability topics, references in the front of this book will guide you.

Why Do We Need High Availability?

We have come to depend on the use of computers, access to the Internet, and the help of our favorite Internet sites. Many people regularly shop for things on the Internet. They expect to be able to go Internet shopping 24 hours per day, seven days per week.

If you have ever purchased anything online, you have probably felt a bit worried as you typed in your name, address, and credit card number you feel as though you are giving out information that will result in yet more junk mail and the addition of your name on even more lists that cold callers use. The first-time registration process at most Web sites involves entering a wealth of personal information. On subsequent visits to the same Internet merchant, only a small portion of this information is normally requested. If you go back to the

same merchant, you are unlikely to be added to more of those lists. Your privacy is retained. This first-time registration is a small but important barrier to closing the deal for the Internet merchant.

Imagine that you are a merchant on the Internet and someone visits your site and your competitor's site while shopping. Subsequently that person decides to purchase an item that you and your competitor are offering for a similar price. If the person registers on your site and buys the item, he or she very likely will be back to your site the next time a similar item is needed because he or she has already gone through the registration process. If the person registers with your competitor's site, then it is likely that he or she will buy that item from your competitor the next time it is needed. Now imagine that the consumer decided to buy from your competitor because your site was down. You not only lost the sale during the downtime, but you also lost the customer. In sales, this is about as bad as it gets.

Customers don't care why they were unable to access your site to place an order. Whether it was a Web server, network, service provider, or some other problem means nothing to them. All they care about is the fact that your site was unavailable and your competitor's site was available. Keep this story in the back of your head, and you will have no trouble understanding the cost of downtime. When the time comes to make the decision between price and availability, it will be much easier when you understand both the cost of downtime, as well as the cost of making downtime go away. You might say that competition is only a mouse click away!

Today's networks don't just carry sales transactions and business information. In fact in some places, you can pick up a telephone, dial a telephone number, and never even realize that the telephone call went over a data network instead of through the traditional telephone company infrastructure. Voice traffic is becoming yet another part of data networks, and data networks are becoming part of the telephone system.

In some cities, consumers can purchase their telephone service from a cable company instead of a traditional telephone company. These people depend on networking equipment to deliver emergency services. Imagine if you picked up your telephone to dial 911 and didn't get a dial tone. When someone's life depends on a network, network availability is crucial—it can be a matter of life or death. Network availability is now a member of an exclusive club, consisting of hospitals, ambulances, and doctors—things that can save your life.

What Is High Availability?

Have you ever been surfing the Internet and had to wait for a Web page for a minute or two? Have you ever noticed what seemed like a few minutes before anything happening? Now I don't mean waiting for a page to load because of your slow connection and pretty pictures on the page; that is a result of a slow connection. I am talking about those times when everything is going just fine and then, all of a sudden, things just seem like they stopped for a minute.

If you have experienced this, then you understand what it is like to notice some sort of network downtime. Of course, we all know this sort of thing happens all the time. A minute here, a minute there, and nobody actually worries about it too much. But imagine how you would feel if the page stopped loading and didn't budge for an hour. Imagine this happened while you were doing research for a project that needed to be done in 30 minutes. If you are like most folks, a minute here or a minute there won't bother you very much, but an hour might be a different story. An hour might irritate you enough to call someone and ask him or her what is happening with your Internet service.

Now think about your telephone. Do you remember the last time you picked up your telephone and there was no dial tone? Do you remember the last time you were on the telephone and you were disconnected? Some folks remember these things. However, most people have a hard time remembering the last time they had telephone trouble. It's hard to remember because it doesn't happen very often. If you do remember the last time you had no dial tone, how did you feel during that experience? In most cases, the feeling of having no dial tone is considerably more traumatic than having no access to the web. No dial tone equals no help in an

emergency. No dial tone is a big deal to everyone. No web access is usually a minor inconvenience.

If you are about to build a network that provides dial tone, then you are going to have to build a highly available network. People will not stand for loss of dial tone because dial tone is something that they require all the time. Networks that provide dial tone must be highly available, or they will not be used. The customer will find a way to replace the network with something more available if they sense the network is not reliable.

Attracting and Keeping Internet Customers

Imagine that you are the customer and that you have a choice of buying your telephone service from your traditional telephone company, your cable television company, or your Internet service provider (ISP). Most people would say that they want to buy their telephone service from the telephone company. After all, the telephone company has been the most reliable service of all their services. The telephone company is even more reliable than the power company in most cities.

Now let us complicate the issue. Say that you can buy your phone service from your ISP for 60 percent less money than you have been paying for your phone service from the phone company. Most people will consider changing services for that much of a discount. The offer is even more attractive if you have two telephone lines. Maybe you can move the second line to the ISP, but leave the first line with the older, more established service. Perhaps you want to just try out the new discounted service, but only on your extra line.

Everyone I know would be willing to put his or her second line on the alternate source for a while. But everyone I asked said that he or she would switch back to the telephone company if the new service proved unreliable. Even at a 60 percent cost discount, high availability is a requirement. Reliability is not something that most people are willing to give up, no matter how much less money they have to spend for their phone service.

High Availability and Government Regulation

In the United States, telephone service is regulated by the government. Telephone companies are required to report downtime of their systems to the Federal Communications Commission (FCC). If a phone company has been negligent or had too much downtime, the government imposes a fine.

Life or Death and High Availability

A friend of mine used to work for a telephone company as a manager of a team that fixed the telephone network whenever it broke. He was under pressure to fix the network very fast whenever it wasn't working.

He once told me, "You haven't ever felt pressure and pain until you've been standing in front of an FCC review panel explaining why the phone system was down when a citizen needed to use it for a 911 phone call. There is no feeling as bad as knowing that someone died because your network was not working."

If you are planning to build a network that carries 911 traffic, you can expect that the government is going to be regulating the network at some point. For now, data networks are not regulated. But someday soon as more data networks carry voice and more of them make it possible for people to call 911 over them these networks will be regulated. Network repair folks and network managers are going to have to report downtime to the government.

Network downtime is going to be a big deal yet another reason why high availability has to be designed into a network.

Presenting and Describing High Availability Measurements

There are two main ways to state the availability of a network: the percentage method and the defects per million method. Once you understand the basics of these methods, you will understand what you are reading when you get an availability analysis report. Both of these methods use figures like MTBF (Mean Time Between Failure) and MTTR (Mean Time To Repair). The following sections describe these two methods as well as the terms MTBF and MTTR.

The Percentage Method

You have probably heard the term five 9s in relationship to the availability of a network. When someone says this, he or she is really saying that the device or network is 99.999 percent available. In fact, 99.999 percent availability is a sure sign that the person is using the percentage method.

NOTE

When you are actually doing the math of calculating availability, you are likely to see a number that looks like 0.99999, rather than 99.999 percent. Remember that you have to multiply by 100 to come up with a percentage, which is how most people are used to seeing availability described.

The essential use of the availability percentage is to figure out how much downtime you are going to have over a year-long period. You determine downtime by multiplying the number of minutes in a year by the percentage of availability. This gives you the minutes per year that you will be operational. The balance is the downtime you can expect.

Because there are 365 days per year, 24 hours per day, and 60 minutes per hour, we can calculate that there are 525,600 minutes per year. However, this does not account for leap years, which have an extra day. The way that we will account for leap years, since they happen every fourth year, is to add one fourth of a day to every year. This results in 525,960 minutes per year, which is the number that is used in all the calculations in this book. 525,960 minutes per year is important enough, in availability calculations, that you will likely have it memorized before long as you become well versed in performing availability calculations.

In addition to the number of minutes per year, annual reliability should be understood. Annual reliability is the number of times each year that a device fails. When you know the MTBF for a device, you can divide that MTBF by the number of hours in a year (8766) to predict the average number of failures per year. We will be using this knowledge when we predict how many minutes a network is out of service while it switches from a broken device to a working device in a redundant situation.

Because we know the number of minutes in a year and because we now understand that availability is a percentage, we can calculate downtime for a year based on the availability number. [Table 1-1](#) describes how the number of 9s relates to uptime and downtime.

Table 1-1. Number of 9s; Uptime and Downtime

Number of Nines	Availability Percentage	Minutes of Uptime per Year (Percentage * 525,960)	Minutes of Downtime per Year (525,960 - Uptime)	Annual Downtime
1	90.000%	473,364	52,596	36.5 days
2	99.000%	520,700.45	259.63	5.5 days
3	99.900%	525,434.05	25.968	5.5 hours
4	99.990%	525,907.45	2.5961	5.5 hours
5	99.999%	525,954.75	.25965	5.5 minutes
6	99.9999%	525,959.50	.5259632	5.5 seconds

As you can see, for each 9 in the availability percentage, a significant increase in performance is achieved. It is often rumored that after the second 9, each additional 9 costs twice as much. That is to say, if you want to go from three 9s to four 9s, the amount of money you spend building your network is going to double! But remember, double the money buys you 10 times more availability.

The Defects per Million Method

The second way to state availability is by using the defects per million (DPM) method. Using this method, we describe the number of failures that have occurred during a million hours of running time for a device or a network. It is common to see this method used for existing large networks.

With the DPM method, we can report issues of reliability that the percentage method would have difficulty tracking. Because DPM is often used for existing networks, we can use it to measure partial and full network outages. We can also measure the million hours in terms of hours of operation of the network, the hours of operation of the devices (added together) that comprise the network, or perhaps even the hours of use that the users get from the network.

In order to clarify the DPM method, let us work through a couple of short examples. Assume that your network consists of 1000 hubs, switches, and routers. Assume that you count any degradation in performance as an outage. Also assume that we base our figures on 8766 hours per year (accounting for leap years) and that our failure reports are done monthly. To determine hours of operation per month, multiply 1000 devices times 8766 hours per year and divide by 12 to compute 730,500 hours of operation each month. As you can see in [Figure 1-1](#), you compute one million hours divided by the number of operating hours to get the number of defects per million hours for a single defect. Then you multiply this result by the number of defects to get the total defects per million hours. So if we had two failures during a month, we would report that as 2.74 DPM for the month.

Figure 1-1. Determining Availability Given Two Failures in One Month

$$\begin{aligned}
 &\text{Hours per year} = 8766 \text{ (Accounts for leap years)} \\
 &\text{Number of Devices} = 1000 \\
 &\text{Accumulated Hours per Year} = 8,766,000 \text{ hours} \\
 &\text{Accumulated Hours per Month} = \frac{8,766,000}{12} \\
 &\qquad\qquad\qquad = 730,500 \text{ hours} \\
 &\text{Converting 2 Defects per 730,500 hours:} \\
 &\qquad\qquad\qquad \frac{1,000,000}{730,500} = 1.3689 \\
 &\qquad\qquad\qquad 2 * 1.3689 = \boxed{2.74 \text{ Defects per Million}}
 \end{aligned}$$

Another way of reporting failures using the DPM method would be to base the million hours on network use. Let us assume that our network is large and constantly growing. Let us also assume that we bill our customers for each hour they use the network. Let us say that we have several thousand customers and that over the period of a month, they accumulate 1,200,000 hours of operation. In this network, if we had two failures, then the resulting DPM would be 1.67 DPM. [Figure 1-2](#) shows you the calculations.

Figure 1-2. Another DPM Calculation Example

$$\begin{aligned}
 &\text{Total Network Hours} = 1,200,000 \\
 &\text{Total Network Failures} = 2 \\
 &\frac{1 \text{ million}}{\text{Actual time}} = \frac{1,000,000}{1,200,000} \\
 &= .83333 \\
 &\text{DPM} = .83333 * 2 \\
 &= 1.67
 \end{aligned}$$

MTBF, MTTR, and Availability

So far, you have learned about failures per million hours of operation and about downtime as a percentage. In order to have a feeling of the health of a network, relating the number of failures and the length of each failure to each other is helpful.

MTBF is a number that you have probably seen on product documentation or in some other specification for a product. It describes the number of hours between failures for a particular device. A similar term is Mean Time to Failure (MTTF), and it describes the amount of time from putting a device into service until the device fails.

Many people confuse these two terms. Using MTBF instead of MTTF and MTTF instead of MTBF in most situations makes very little difference. Technically, the mathematical equation for calculating availability that we will use in this book should use the term MTTF, according to the historic standards on the subject. You will nearly always be able to get an MTBF number about products you wish to purchase. Finding MTTF numbers about products is difficult. Technically, the companies stating MTBF about their products are very likely to be giving you MTTF and not even know it.

In keeping with this minor industry oversight and to simplify our lives, this book uses MTBF in place of MTTF and discards MTTF completely because it will make very little difference in our calculations. Before getting into an example that shows exactly how little difference it makes to switch MTBF and MTTF numbers, you have learn about MTTR and the availability equation.

MTTR is the amount of time (on average) that elapses between a network failing and the network being restored to proper working order. In most cases, MTTR includes a little bit of time to notice that the network has failed. Then it includes some time to diagnose the problem. Finally, MTTR includes some time to perform the appropriate action to fix the network and a small amount of time for the repairs to bring the network into proper working order. In an ideal world, the timing to detect, diagnose, and repair a network problem will be measured in minutes. However, sometimes things happen in the night and no one notices for hours. Sometimes the first diagnosis is wrong and several hours are wasted fixing something that isn't broken. The key point here is to remember that there are three phases to fixing a network problem: