

# **How To Accelerate Your Internet**

**A practical guide to Bandwidth Management and  
Optimisation using Open Source Software**

# How To Accelerate Your Internet

For more information about this project, visit us online at <http://bwmo.net/>

Editor: Flickenger R.

Associate Editors: Belcher M., Canessa E., Zennaro M.

Publishers: INASP/ICTP

© 2006, BMO Book Sprint Team

First edition: October 2006

ISBN: 0-9778093-1-5

Many designations used by manufacturers and vendors to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the authors were aware of a trademark claim, the designations have been printed in all caps or initial caps. All other trademarks are property of their respective owners.

The authors and publisher have taken due care in preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information contained herein.



This work is released under the Creative Commons **Attribution-ShareAlike 2.5** license. For more details regarding your rights to use and redistribute this work, see <http://creativecommons.org/licenses/by-sa/2.5/>

# Contents

<b>Preface</b>	<b>ix</b>
<b>About This Book</b>	<b>xi</b>
<b>Introduction</b>	<b>1</b>
Bandwidth, throughput, latency, and speed.....	2
Not enough to go around.....	3
Where to begin.....	5
<b>Policy</b>	<b>9</b>
The importance of policy.....	10
Explosive network growth at Havensburg.....	10
Bandwidth as a public good.....	11
Desperate measures.....	12
Policy, strategy, rules and regulations.....	13
Real policy development at Havensburg.....	14
Characteristics of good policy.....	15
The new Havensburg network policy.....	16
The policy development process.....	17
Policy is needed in all environments.....	19
Policy pitfalls.....	20
Example policies.....	20
Policy checklist.....	21
References.....	22

## Monitoring & Analysis 25

Networking 101.....	26
Introduction.....	26
Cooperative communications.....	28
The OSI model.....	28
The TCP/IP model.....	31
The Internet protocols.....	32
Networking hardware.....	44
Physical connectivity.....	49
Virtual connectivity.....	58
What is network monitoring?.....	62
An effective network monitoring example.....	63
Monitoring your network.....	66
The dedicated monitoring server.....	67
What to monitor.....	70
How to select tools to monitor the network.....	71
Types of monitoring tools.....	72
Walking around the lab.....	73
Spot check tools.....	74
Log analysers.....	80
Trending tools.....	83
Realtime tools.....	87
Benchmarking.....	89
What is normal?.....	91
How do I interpret the traffic graph?.....	95
Monitoring RAM and CPU usage.....	97
Resources.....	99

## Implementation 101

The importance of user education.....	102
The 5/50 rule.....	102
Providing feedback to users about network load.....	103
General good practices.....	105
Essential services.....	112
Firewall.....	114
Caching.....	134
Mirroring.....	144
Email.....	148
Resources.....	156

## Troubleshooting 159

Proper troubleshooting technique.....	159
---------------------------------------	-----

Preparing for problems.....	160
Responding to a problem.....	160
A basic approach to a broken network.....	161
Common symptoms.....	164
Automatic updates.....	164
Spyware.....	165
P2P.....	165
Email.....	165
Open email relay hosts.....	166
Email forwarding loops.....	167
Open proxies.....	167
Programs that install themselves.....	167
Programs that assume a high bandwidth link.....	167
Windows traffic on the Internet link.....	168
Streaming media / Voice over IP.....	169
Denial of Service.....	170
Rogue DHCP servers.....	170
Port analysis.....	171
Browser prefetch.....	172
Benchmark your ISP.....	172
Large downloads.....	172
Large uploads.....	173
Users sending each other files.....	173
Viruses and worms.....	174

## Performance Tuning 177

Squid cache optimisation.....	178
Cache server hardware.....	179
Tuning the disk cache.....	180
Memory utilisation.....	181
Tuning the hot memory cache.....	182
Cacheable content limits.....	182
Access Control List (ACL) optimisation.....	183
Redirectors.....	184
DansGuardian.....	185
Authentication helpers.....	186
Hierarchical caches.....	187
Configuring delay pools.....	189
More information.....	191
Monitoring your Squid performance.....	192
Graphing Squid metrics.....	195
Traffic shaping.....	196
Linux traffic control and QoS tools.....	196
Traffic shaping with BSD.....	203
Farside colocation.....	205

Choosing a colo or ISP.....	208
Billing considerations.....	208
Protocol tuning.....	209
TCP window sizes.....	209
Link aggregation.....	210
Bonding.....	211
Aggregate routing.....	211
DNS optimisation.....	212
Web access via email.....	214
www4mail.....	215
web2mail.....	215
PageGetter.com.....	216
GetWeb.....	216
Time Equals Knowledge (TEK).....	216
Other useful web-to-email applications.....	217
loband.org.....	217
High Frequency (HF) networks.....	218
Modem optimisation.....	219
Hardware compression.....	219
Software compression.....	220
Bandwidth accounting.....	221
Squid bandwidth accounting.....	221
Bandwidth accounting with BWM tools.....	222
Linux interface bandwidth accounting with RRDtool.....	223
VSAT optimisation.....	223
Use of inclined orbit satellite.....	224
C band, Ku band, and Ka band.....	224
Shared vs. dedicated bandwidth.....	226
Resources.....	232

## Case Studies

235

KENET, Kenya.....	235
Problems.....	236
Analysis.....	236
Solutions.....	236
Site One: firewall & proxy server.....	237
Site Two: proxy & mail server.....	237
Site Three: FOSS traffic shaper.....	238
Aidworld in Accra, Ghana.....	239
BMO in the UK.....	241
JANET, UK.....	241
Blackburn College, UK.....	243
Malawi.....	245
One Bellevue Center.....	247
Carnegie Mellon University.....	248

Workaround #1: Best effort rate limiting.....	248
Getting more than you paid for.....	248
Workaround #2: Fun with rate limiting.....	249
More problems with packet drops.....	249
Requirements and considerations.....	250
Researching hardware rate limiters.....	250
Final solution or new workaround?.....	250
Application layer analysis to the rescue.....	251
Social engineering.....	251
The campus bandwidth usage guidelines.....	252
Human effort.....	253
Positive results.....	253
Conclusion.....	253

## **The Future 255**

Bandwidth consuming technologies.....	255
Trends in developing countries.....	256
New software.....	257
In closing.....	258

## **Resources 259**

Links.....	259
Wikipedia entries.....	267
Relevant RFCs.....	267

## **Squid ACL Primer 269**

ACL elements.....	269
ACL rules.....	271
Examples.....	272
Allow only local clients.....	272
Deny a list of sites.....	273
Block a few clients by IP address.....	273
Allow access to the bad sites only after hours.....	273
Block certain users regardless of their IP address.....	273
Direct certain users to a delay pool.....	273

## **Glossary 275**



# Preface

One measure of the growing disparity between the developed and developing worlds is the speed of the Internet. For example, the speeds of connections from North America to Africa are slower than those to Europe by a factor of 50 or so. Such assessments have been made by measuring the round trip time that it takes for a digital pulse sent over the Internet to return to the sender.

The reasons for this disparity include the availability of Internet access only via slow satellite connections, and the lack of communications infrastructure in the remote parts of the world. Bandwidth and computing equipment are expensive as a result of weak currencies, high transport costs, small budgets and unreasonable tariffs. Bandwidth in some developing countries can be so costly that even their prime universities cannot afford speeds equivalent to the average western household with an ADSL connection. Thus universities and other institutions cannot afford a decent link, or are simply unaware of existing alternatives.

This book attempts to provide practical information on how to gain the largest benefit from existing connections to the Internet, by exposing readers to the latest techniques to optimise the use of low-bandwidth network connections. By applying optimisation techniques based on open source technologies discussed here, the effectiveness of available connections can be significantly improved. Access to more bandwidth will facilitate better exchange of scientific information, data and literature among researchers all over the world. One hopes that the process will enable every scientist to become part of the scientific enterprise no matter where geographically she is located with respect to the main centers of modern science.

While the Internet has helped global communication, and its use is rising everywhere, the fraction of people with access to it is far higher in rich countries than in poor countries. The average per capita income in industrialised nations is about \$27,000 per year, compared with barely \$2,000 or so in the developing

world. Literacy rates approach 100% of the adult population in developed countries, but the figure falls to below 50% in developing nations. Even as the world is becoming more interconnected, it is becoming increasingly divided in these regards.

This book is a collaborative effort enabled by the support of INASP (UK) and ICTP. The effort that has gone into its preparation will be rewarded if the book can reach large audiences of interested readers and assist them in improving the quality of service of the bandwidth available to them. The authors of the book realise that it is a small drop in the huge ocean of bits and bytes, but the value of their service is not in any doubt. I congratulate them on their work and their decision to make the book freely available both in print and on the Internet.

*K.R. Sreenivasan*  
*Abdus Salam Professor*  
*Director, ICTP*

*Trieste*  
*October 2006*

# About This Book

This work is published under a Creative Commons **Attribution-ShareAlike 2.5** license. This allows anyone to make copies or derivative works, and even sell them for a profit, as long as proper attribution is given to the authors and any derivative works are made available under the same terms. Any copies or derivative works **must** include a prominent link to our website, <http://bwmo.net/>. See <http://creativecommons.org/licenses/by-sa/2.5/> for more information about these terms. Consult our website (<http://bwmo.net/>) for details on how to order a printed copy.

## Credits

This book was started as a BookSprint project at the ICTP in Trieste, Italy, in May of 2006. A core team of ten experts in the field of bandwidth management built the initial outline, and developed the book over the course the following months. Throughout the project, the core group has actively solicited contributions and feedback from the Internet community, particularly those who work in the area of bandwidth optimisation in the developing world. The final manuscript was produced by Hacker Friendly LLC in Seattle, WA (USA).

## Contributors

- **Aidworld** (<http://www.aidworld.org/>) is a not-for-profit organisation focussed on information technology for international development. Aidworld's mission is to effectively support the Millennium Development Goals with appropriate ICTs. Aidworld builds bandwidth management solutions and helps NGOs and others make their online services accessible in the developing world. Aidworld has also created an on-line tool (<http://www.loband.org/>) that shrinks web pages so they are accessible over poor internet connections. Aidworld contributors include **Nishant Bhaskar**, **Hamish Downer**, **Alan Jackson**, **Simon Liu**, **Tom Lord**, **Jon Stafford**, **Nick Street**, **Tom Taylor**, and **Chris Wilson**.

- **Martin Belcher** is the Senior Programme Manager for the International Network for the Availability of Scientific Publications (INASP), Lund, Sweden. He can be reached at *mbelcher@inasp.info*.
- **Enrique Canessa** is a PhD Physicist working at the ICTP in Trieste, Italy. His areas of interest are scientific software applications, ICT training, and dissemination of science to/from and within developing countries using open source technologies.
- **Kevin Chege** is the Senior Network Administrator at the Kenya Education Network (KENET). He is an avid user of FreeBSD and an open source enthusiast focusing on improving ICT reach in education using FOSS tools. He can be contacted at *kevin@kenet.or.ke*.
- **Rob Flickenger** was the lead editor of this project, and is the founder of Hacker Friendly LLC. Rob is a long-time supporter of the use of wireless networking to extend the reach of the Internet. He can be reached at *rob@hackerfriendly.com*.
- **Carlo Fonda** is a member of the Radio Communications Unit at the Abdus Salam International Centre for Theoretical Physics in Trieste, Italy.
- **Duncan Greaves** is an Executive Officer at the Tertiary Education Network (TENET), a not-for-profit company supporting higher education in South Africa. Duncan oversees TENET's capacity development programs. He can be contacted at *dbg@tenet.ac.za*.
- **Casey Halverson** is a Network Engineer at Infospace Inc. in Seattle, Washington, USA. He has a broad experience in WAN optimisation, traffic shaping, and other bandwidth management techniques. He is also a member of the SeattleWireless network project. *<http://seattlewireless.net/>*
- **Peter Hill** is a self-titled "Holistic Network Engineer" for the University of Washington. He previously worked in the trenches keeping Amazon's network afloat, and still has fond memories of Carnegie Mellon's network and awesome Network Development team.
- **Nigel Kukard** has a PhD in Computer Science, and has been a passionate supporter of open source (GPL) software for over ten years. He is the founder of LinuxRulz (*[www.linuxrulz.org](http://www.linuxrulz.org)*) and the Linux Based Systems Design group of companies. Can be reached at *nkukard@lbsd.net*.
- **Richard Stubbs** is a technical evangelist who works for the University of KwaZulu-Natal in South Africa. He has been involved with the Internet and associated activities at the Institution for the past 15 years. He can be contacted at *stubbs@ukzn.ac.za*
- **Marco Zennaro** is an electronic engineer working at the ICTP in Trieste, Italy. He has been using BBSes and ham radios since he was a teenager, and is happy to have merged the two together working in the field of wireless networking.

# Additional material

Portions of this work were adapted from:

- Network traffic monitoring and analysis workshop (INASP) by Dick Elleray, AfriConnect, 2006 <http://www.inasp.info/training/bandwidth/bmo-ntmw/>
- Optimising Internet Bandwidth (INASP) by Gerhard Venter, AfriConnect, 2003 <http://www.inasp.info/pubs/bandwidth/index.html>
- The VSAT Buyer's Guide, IDRC, 2005 <http://ictinafrica.com/vsat/>
- Wireless Networking in the Developing World, <http://wndw.net/>

## Funding

The development and publication of this book have been made possible with funding support from:

- Department for International Development, UK (DfID); through funding to INASP for the "Programme for the Enhancement of Research Information (PERI)."
- Flemish Interuniversity Council (VLIR); through funding to INASP for the programme "Optimization of the use and management of bandwidth at university level."
- International Network for the Availability of Scientific Publications (INASP).
- International Development Research Centre (IDRC): through funding to INASP for the programme "Supporting training for the optimization of university bandwidth in Africa" undertaken with financial support from the Canada Fund for Africa.

## Special thanks

The production team would like to thank the ICTP (<http://www.ictp.it/>), INASP (<http://www.inasp.info/>), and everyone else who has made this project possible.





# 1

## Introduction

The Internet has irrevocably invaded many aspects of daily life. What was once an obscure scientific research tool has blossomed into a communications platform used by hundreds of millions of people. Telecom providers use the Internet to carry critical voice communications. Banking institutions use it to provide access to account services and market trading. Airline tickets, hotel reservations, and car rentals can all be booked with a click of the mouse. Whole industries have sprung into existence with business models that depend entirely on Internet infrastructure to reach their customers. More users than ever depend on the Internet to connect with family and colleagues using email, instant messaging, Voice over IP, photo and video sharing services, and online journals. Children born in the last ten years have grown up in a time when the Internet has always been available.

This point of view is popular among Internet users, but it does not necessarily reflect the experience of all, or even most, of the rest of the world. According to the ITU\*, more than half of the users on the Internet are concentrated in the G8 countries (Canada, France, Germany, Italy, Japan, Russia, the UK, and the US). In 2004, less than 3% of Africans used the Internet, compared with an average of one 50% of the inhabitants of the G8 countries. The entire African continent accounts for about 13% of the total world population, yet in 2004 it had fewer Internet users than France alone.

Fortunately, in places where the Internet has not yet penetrated, it is all but certain to happen soon. There is a global push to bridge the so-called digital divide by bringing modern telecommunications to the developing world. State and private investment in public infrastructure, in the form of fibre optic backbones, wireless networks, and satellite connectivity are bringing the Internet to the most remote locations at a pace that is accelerating over time. People all

---

\* Source: <http://www.itu.int/ITU-D/ict/statistics/ict/>

over the globe are beginning to realise that in order to effectively participate in the global marketplace, they need access to the global communications network.

But superhighways aren't built overnight. As with any major undertaking to build infrastructure, extending fast network connections to all of the ends of the earth takes time. Technologies such as VSAT make it possible to install an Internet connection just about anywhere, particularly in the absence of existing wired infrastructure. While this does extend the footprint of the Internet to otherwise unreachable places, the capacity of the connection provided is far from infinite. The cost of these connections is also quite high for many organisations. This often leads to the practice of stretching an insufficient network connection to serve many users simultaneously.

## Bandwidth, throughput, latency, and speed

There are a few technical words used to describe how fast an Internet connection may go. Users often find these terms confusing, so it's best to be clear about their definitions from the beginning.

- **Bandwidth** refers to a measure of frequency ranges, typically used for digital communications. The "band" part of broadband is short for bandwidth, meaning that the device uses a relatively wide range of frequencies. In recent years, the term bandwidth has been popularly used to refer to the **capacity** of a digital communications line, typically measured in some number of bits per second. In its popular usage, you might read that a T1 provides a theoretical maximum "bandwidth" of 1.544 Mbps.

While some purists insist that we should speak of capacity when talking about data transfer speeds and bandwidth when talking about frequency ranges, the popular usage of the term "bandwidth" has been reinforced by years of product marketing and misleading documentation. There simply is no going back now. Therefore, we will use the terms bandwidth and capacity interchangeably in this book.

- **Throughput** describes the actual amount of information flowing through a connection, disregarding protocol overhead. Like bandwidth, it is expressed in some number of bits per second. While a T1 may provide 1.544 Mbps between the endpoints, the protocol spoken on the physical line reduces the effective throughput to about 1.3 Mbps. When you factor in the additional overhead of Internet protocols, the available throughput is even less. When you measure the actual usage of a connection or perform a "speed test" on a line, you are measuring throughput.
- **Latency** refers to the amount of time it takes for a packet to travel from one point on a network to another. A closely related concept is **Round Trip Time (RTT)**, which is the amount of time it takes for a packet to be acknowledged

from the remote end of a connection. Latency is measured as some amount of time, usually in milliseconds. The latency of Ethernet is about 0.3 ms. A T1 connection has a latency of 2 to 5 ms, while a VSAT connection requires at least 500 ms before an acknowledgment can be received, due to the speed of light and the large distances involved. Some factors that contribute to latency are network congestion, overutilised servers, and the distance between the two points.

- **Speed** is an ambiguous term that refers to some combination of these other terms. An Internet connection may "feel slow" when using an interactive service (such as Voice over IP or gaming) on a line with high latency, even if there is sufficient bandwidth. Users will also complain when transferring large files on a connection with insufficient capacity, even if the latency is very low.

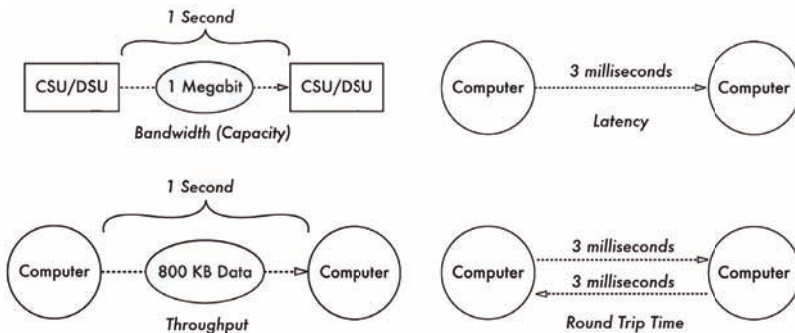


Figure 1.1: Bandwidth, Capacity, Throughput, Latency, and Round Trip Time.

The goal of this book is to show you how to optimise your Internet connection so that it provides the greatest possible throughput and lowest possible latency. By eliminating wasted bandwidth, the cost of operating your network connection will be reduced, and the usability of the network will be improved.

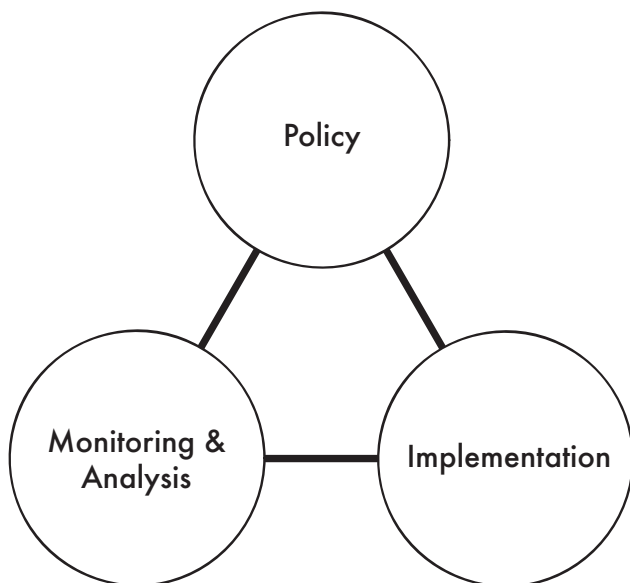
## Not enough to go around

What actually causes a slow Internet connection? Obviously, the capacity of a given connection is finite, so if too many people request information at once, then someone will have to wait. In an ideal world, organisations would simply order more bandwidth to accommodate the increased traffic. But as we all know, Internet access costs money, and most organisations do not have infinite budgets.

It is an interesting fact of online life that users tend to consume more bandwidth over time. It is very rare to find a user who, once they have had access to a broadband connection, is satisfied with going back to a low speed dialup line. As users are exposed to Internet services earlier in life and in a variety of venues (for example at home, at work, at University, or at a cyber-cafe), they be-

come accustomed to using it in a certain way. They are increasingly unlikely to know or care about the bandwidth required to listen to Internet radio, or to download the latest video game, or to watch funny movies on a video sharing service. They "just want it to work," and may complain when the Internet "is slow." Users often have no idea that they can single-handedly bring an organisation's Internet connection to a halt by running a simple file sharing program on their computer.

User education is obviously critical to every stage of implementing a plan to manage your bandwidth. While users can be forced to adhere to certain behaviour patterns, it is always far easier to implement a plan with their voluntary compliance. But how does such a plan come into being? If you simply order people to change their behaviour, little is likely to change. If you install technical hurdles to try to force them to change, they will simply find a way around the obstacles.



*Figure 1.2: Policy, Monitoring & Analysis, and Implementation are three critical (and interdependent) components of bandwidth management.*

In order to effectively manage a network connection of any size, you will need to take a multifaceted approach that includes effective **network monitoring**, a sensible **policy** that defines acceptable behaviour, and a solid **implementation** that enforces these rules. Each component is important for effective bandwidth management in any network that consists of more than a few users. This book includes chapters devoted to each of these three major areas.

A **policy** is a statement of opinions, intentions, actions and procedures that guide the overall use of the network. An **acceptable use policy** is a subset of

this, setting out in technical detail what uses of the network are believed by the network operators to be acceptable, and what they intend to do to anyone who uses it in a manner that they consider unacceptable. It should be a written document that defines acceptable forms of network access, as well as guidelines for how network problems are dealt with, definitions of abuse, and other operational details. The policy also typically includes definitions of legal constraints for network users (such as the exchange of copyrighted material, requesting inappropriate materials, etc.). Having a policy makes it much easier to enforce certain types of network behaviour, as you will be able to hold people to a set of agreed rules.

**Network monitoring** is the ongoing process of collecting information about various aspects of your network operations. By carefully analysing this data, you can identify faults, find cases of waste and unauthorised access, and spot trends that may indicate future problems.

**Implementation** is the step of implementing traffic shaping, filtering, caching, and other technologies within your network to help bring actual usage in line with policy. The actions you need to take are indicated by the data collected through monitoring and analysis, and are constrained by the network policy. Many people expect to begin the task of bandwidth management by starting with this step. But without good monitoring techniques, you are effectively blind to the problem. Without a policy, your users will not understand what you are doing or why, and will complain or subvert your actions instead of helping you to achieve your goal.

Don't underestimate the value of personally interacting with your network users, even at a very large institution. At Carnegie Mellon University (page **248**), social interactions made a far greater impact on bandwidth consumption than did technical constraints. But at an organisation as large as CMU, personal attention could only have had this effect by operating within a well-defined policy, with the support of a good network implementation and watched by careful network monitoring.

## *Where to begin*

Effective bandwidth management can only happen by applying a combination of technical computer skills, effective network monitoring, and a sensible policy that is understood by all users. If your organisation has a small network, one person may need to work on all of these areas. Larger organisations will likely require a team of people to effectively manage busy networks, with each person specialising in a particular area.

This book is designed to be used as both a guide and a reference to anyone who needs to tackle this difficult problem. While you may read it cover-to-cover,

each chapter is designed to stand on its own and address a particular aspect of bandwidth management. If you don't know where to begin, these guidelines should help you find a good starting place.

### Do you need to fix your network immediately?

- Is something wrong with your computers or Internet access?
- Do the problems get in the way of people getting legitimate work done?
- Is your job at risk if you don't do something now?

If you answered **yes** to any of these questions, go to the **Troubleshooting** chapter (page 159). When you've solved the immediate problem, continue with the steps below.

### Do you know what's happening on your network?

- Do you monitor your network?
- Do you know what your bandwidth usage is, on average?
- Do you know who is using your bandwidth?
- Do you know how your bandwidth is being used? How much bandwidth is used for email, as compared to web traffic and peer-to-peer applications?
- Do you know about network outages before your users complain?
- Are you certain that your network only being used for appropriate services, and has not been compromised by a malicious user?

If you answered **no** to any of these questions, take a look at the **Monitoring & Analysis** chapter on page 25. When you have a clear idea of what's happening on your network, continue with the steps below.

### Do you want to change how users behave on your network?

- Is inappropriate user behaviour (e.g. peer-to-peer file sharing or excessive downloads) causing problems on your network?
- Do you need to create a written policy on network usage?
- Do you need to update an existing policy?
- Are your users largely unaware of what the network policy is, and why it is important?
- Do you need to guarantee the availability of certain services on your network?