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

First edition: October 2006

ISBN: 0-9778093-1-5

Contents

Monitoring & Analysis 25

Implementation 101

Troubleshooting 159

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

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*) 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.

The Abdus Salam
International Centre
for Theoretical Physics