IPv6 DEPLOYMENT AND MANAGEMENT

Michael Dooley Timothy Rooney







Includes an introduction by

IPv6 DEPLOYMENT AND MANAGEMENT

IEEE Press 445 Hoes Lane Piscataway, NJ 08854

IEEE Press Editorial Board 2013

John Anderson, Editor in Chief

Linda Shafer George W. Arnold Ekram Hossain Om P. Malik Saeid Nahavandi David Jacobson Mary Lanzerotti George Zobrist Tariq Samad Dmitry Goldgof

Kenneth Moore, Director of IEEE Book and Information Services (BIS)

IPv6 DEPLOYMENT AND MANAGEMENT

By

MICHAEL DOOLEY TIMOTHY ROONEY









Copyright © 2013 by The Institute of Electrical and Electronics Engineers, Inc.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey. All rights reserved. Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permission.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data:

Rooney, Tim.
IPv6 deployment and management / Timothy Rooney, Michael Dooley. pages cm
ISBN 978-1-118-38720-7 (pbk.)
1. TCP/IP (Computer network protocol) 2. Internet addresses. I. Title. TK5105.585.R66 2013 004.6'2068—dc23

2012041248

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Michael would like to dedicate this book to his parents

Timothy would like to dedicate this book in memory of his mother, Kathryn "Kitty" Rooney

CONTENTS

ACKNOWLEDGMENTS			XI
IN	INTRODUCTION		
1	IPve	5 DEPLOYMENT DRIVERS	1
	1.1	The Internet: A Success Story	
		1.1.1 Supply-Side Issues	3
		1.1.2 Internet at a Crossroads	6
		1.1.3 Which Internet Are You On?	7
	1.2	Emerging Applications	7
	1.3	IPv6 Business Case	10
2	IPve	5 OVERVIEW	13
	2.1	2.1 IPv6 Key Features	
	2.2	The IPv6 Header	14
		2.2.1 IPv6 Extension Headers	15
	2.3	IPv6 Addressing	17
		2.3.1 Address Notation	18
		2.3.2 Address Structure	19
		2.3.3 IPv6 Address Allocations	20
		2.3.4 Internet Control Message Protocol for IPv6 (ICMPv6)	27
		2.3.5 IPv6 Ping	28
		2.3.6 Multicast Listener Discovery	28
		2.3.7 Multicast Router Discovery	31
		2.3.8 Neighbor Discovery Protocol	31
		2.3.9 Secure Neighbor Discovery (SEND)	33
		2.3.10 Inverse Neighbor Discovery	33
		2.3.11 Router Renumbering	34
		2.3.12 Node Information Query	34
	2.4	IPv6 Address Autoconfiguration	35
		2.4.1 Modified EUI-64 Interface Identifiers	36
		2.4.2 Duplicate Address Detection (DAD)	37
	2.5	Mobile IPv6	
	2.6	Reserved Subnet Anycast Addresses	
	2.7	Required Host IPv6 Addresses	
	2.8	8 IPv6 Routing	

3 IPv4/IPv6 CO-EXISTENCE TECHNOLOGIES

	3.1	Dual Stack	44
		3.1.1 Implementing Dual Stack	44
		3.1.2 Which Address Is Used?	45
		3.1.3 DNS Considerations	47
		3.1.4 DHCP Considerations	48
	3.2	Tunneling Approaches	49
		3.2.1 Tunneling Scenarios for IPv6 Packets Over IPv4 Networks	49
		3.2.2 Tunnel Types	51
		3.2.3 Tunneling Scenario for IPv4 Packets Over IPv6 Networks	62
		3.2.4 Tunneling Summary	63
	3.3	Translation Approaches	63
		3.3.1 IP/ICMP Translation	65
		3.3.2 Bump in the Host (BIH)	72
		3.3.3 Network Address Translation for IPv6/IPv4 (NAT64)	74
		3.3.4 Other Translation Techniques	75
	3.4	Application Support of IPv6	78
	3.5	Service Provider IPv4/IPv6 Co-Existence	78
		3.5.1 Reference Architecture	79
		3.5.2 Deployment Approaches Overview	80
		3.5.3 Routing Infrastructure Deployment Approaches	80
		3.5.4 Comparison of Deployment Approaches	87
	3.6	Addressing and DNS Considerations	87
4	IPve	5 READINESS ASSESSMENT	91
	4.1	Putting a Plan in Place	92
	4.2	IP Network Inventory	93
		4.2.1 IPv6 Readiness	93
		4.2.2 Discovery	93
		4.2.3 IPv6 Assessment	94
	4.3	IPv6 to do List	106
	4.4	IPv6 Readiness Assessment Summary	106
5	IPve	5 ADDRESS PLANNING	109
	51 14 40 44		
	5.1	Internet Registries	109
		5.1.1 KIK Address Allocation Efficiency	111
	5 2	J.1.2 Address Allocation Efficiency	112
	53	IPv6 Address Allocation Methods	112
	5.5	5.3.1 Best. Fit Method	113
		5.3.1 Dest-1 ti Method	114
		5.3.2 Sparse Allocation	110
		5.3.4 DHCPv6 Prefix Delegation	117
		5.3.5 Unique Local Address Space	110
	54	Defining Your IPv6 Address Plan	110
	5.4	Multihoming and IP Address Space	
	5.5	IP Address Planning Summary	122
			120

43

6	IPv6	5 SECURITY PLANNING	127
	6.1	The Good News: IP Is IP	127
	6.2	The Bad News: IPv6 Is Not IPv4	128
	6.3	Update Your Security Policy	129
	6.4	Network Perimeter Monitoring and Intrusion Prevention	129
		6.4.1 IPv6 Address Filtering	130
		6.4.2 ICMPv6 Messages	131
	6.5	Extension Headers	132
	6.6	Internal Network Protection	133
		6.6.1 Network Reconnaissance	133
		6.6.2 Network Access	134
		6.6.3 DHCPv6	135
		6.6.4 DNS	135
		6.6.5 Anycast Addressing	136
		6.6.6 Internal Network Filtering	136
	6.7	Network Device Security Considerations	137
	6.8	Mobile IPv6 Security	138
		6.8.1 Mobility Extension Header	139
		6.8.2 Mobile IPv6 Vulnerabilities	143
	6.9	IPv4/IPv6 Coexistence Measures	144
		6.9.1 Securing Tunneling Implementations	145
	<i></i>	6.9.2 Securing Translation Implementations	146
	6.10	Summary	148
7	IPv6	5 NETWORK MANAGEMENT PLANNING	149
	7.1	149	
	7.2	Network Management Scope	150
		7.2.1 Network Inventory	150
		7.2.2 IP Address Inventory	151
		7.2.3 The Management Network	151
	7.3	The Simple Network Management Protocol (SNMP)	152
		7.3.1 Configuration Management	153
		7.3.2 Fault Management	153
		7.3.3 Accounting Management	154
		7.3.4 Performance Management	154
	7.4	Methods and Procedures	154
	7.5	Summary	155
8	MAN	NAGING THE DEPLOYMENT	157
	8.1	157	
	8.2	Project Management	159
	8.3	Testing Deployment	160
	8.4	Production Deployment	161
9	MAN	NAGING THE IPv4/IPv6 NETWORK	163
	9.1	Common Network Management Tasks	163
	9.2	Configuration Management	163

		9.2.1	Network Allocation-Related Tasks	164
		9.2.2	Adding a New Device	166
		9.2.3	Deletion Tasks	167
		9.2.4	Address Renumbering or Movement Tasks	168
		9.2.5	Block/Subnet Splits	171
		9.2.6	Block/Subnet Joins	172
		9.2.7	DHCPv6 Server Configuration	173
		9.2.8	DNS Server Configuration	174
		9.2.9	Prefix Renumbering	175
9	9.3	Fault l	Management	176
		9.3.1	Fault Detection	176
		9.3.2	Troubleshooting and Fault Resolution	177
9	9.4 Accounting Management		177	
		9.4.1	Inventory Assurance	177
		9.4.2	Address Reclamation	180
9	9.5	Perfor	mance Management	181
		9.5.1	Services Monitoring	181
		9.5.2	Application Performance Management	182
		9.5.3	Auditing and Reporting	182
9	9.6	Securi	ity Management	183
9	9.7	Disast	er Recovery/Business Continuity	183
40	ID.		THE ELITIDE INTEDNET	105
10	IP V0) AND		165
	10.1	Tech	nology Enablers	185
	10.2	The l	Internet's Dark Side	187
	10.3	The l	Internet's Bright Future	187
		10.3.	1 Living Smarter	187
		10.3.	2 Keeping Track	188
		10.3.	3 Extensible Healthcare	188
		10.3.4	4 Public Safety	188
		10.3.	5 Credit Cards of the Future	188
		10.3.	6 Consumer Applications	188
	10.4	Conc	clusion	189
APP		IX		101
			1)1	
BIBLIOGRAPHY			193	
INDEX			199	

ACKNOWLEDGMENTS

We would both like to thank Vint Cerf for the introduction to this book; we are humbled and honored. We would also like to thank Thomas Plevyak, our series editor at IEEE Press, as well as Michael Vincent and Jeff Schmidt for their time spent reviewing drafts of this book and providing extremely useful feedback and comments.

From Michael: I would also like to thank my family, my wife Suzanne, my son Michael, and my daughter Kelly, for all their love and support and allowing me not to be distracted at home while I was working on this book. And I can't forget my puppy Bailey as well, who nudged me at every opportunity to pet her instead of letting me write. I would also like to thank the following individuals who are my friends and coworkers. I have had the pleasure to work with some of the best and brightest people in the world, and I am truly blessed. In no particular order, I thank Karen Pell, Steve Thompson, Greg Rabil, John Ramkawsky, Alex Drescher, Brian Hart (aka Billy Bond), Bob Lieber, David Cross, and Al Hilton. I would also like to acknowledge the original Quadritek leadership team that I had the privilege to work with as we helped to define and create the IP Address Management market back in the early years, specifically including Arun Kapur, Keith Larson, and Leah Kelly. And a special thanks to Joe D'Andrea whose leadership has had a profound impact on my life and my career.

From Timothy: I would also like to thank my family, my wife LeeAnn, and my daughters Maeve and Tess, for their love and support during the development of this book! I would also like to thank the following individuals with whom I have had the pleasure to work and from whom I have learned tremendously about communications technologies and IPv6: Greg Rabil, John Ramkawsky, Andy D'Ambrosio, Alex Drescher, David Cross, Marco Mecarelli, Brian Hart, Frank Jennings, and those I have worked with at BT Diamond IP, INS, and Lucent. From my formative time in the field of networking at Bell Laboratories, I thank John Marciszewski, Anthony Longhitano, Sampath Ramaswami, Maryclaire Brescia, Krishna Murti, Gaston Arredondo, Robert Schoenweisner, Tom Walker, Ray Pennotti, and especially Thomas Chu.

INTRODUCTION

Nearly 14 years have passed since RFC2460 was published, specifying the IPv6 packet format. Authored by Steve Deering and Bob Hinden, this document represented nearly 8 years of debate beginning in the early 1990s over how the Internet's 32-bit IPv4 address space could be expanded. There were four proposals for what was called "IPng" for IP next generation. I won't catalog them here except to say they varied dramatically in their functionality. There was even a fifth proposal to adopt the OSI connectionless networking protocol format (CLNP) that provoked howls of outrage from many passionate engineers in the Internet Engineering Task Force (IETF) where this problem was near the top of the agenda.

After all the debate, the cochairs of the IPng Working Group, Deering and Hinden, recorded the results in December 1998 and submitted them as RFC 2460 to the Internet Engineering Steering Group (IESG) for release to the RFC editor. Many of us hoped there would be an immediate effort to implement this protocol. There was great concern that the rate of consumption of the Internet address space was accelerating during the period now known as the "dot-boom." New Internet companies were popping up like mushrooms after a spring rain. But at the same time that the IPng debates were taking place, another effort to restrain IPv4 address consumption, through reinterpretation of the bits of the address structure, was in full swing. The so-called classless interdomain routing system made much more efficient use of address space by allowing any bit boundary in the address structure to mark the dividing line between "network" and "host." In addition, the concept of autonomous system (AS) was introduced through which to associate indicators (masks) illustrating where this boundary lay. The Border Gateway Protocol was revised to take into account the masks marking network and host extents in the address format. Together with rules to guide very conservative IPv4 address allocations by the Regional Internet Registries, the rate of consumption of IPv4 address space was substantially curtailed. So much so that the pressure to implement IPv6 generally dissipated.

Network address translation (NAT) functionality was also introduced to allow multiple devices using private IP address numbering to share a single public address space. Port numbers were used in the NAT boxes to map to/from public addresses and the private addresses associated with individual devices in a local network. This practice attracted cable and telecommunications providers who were offering Internet service because they could now maximize the number of devices that could share one "public" IP address. This improved the absolute number of customers they could sign up to be subscribers to their Internet service.

These various practices actually stretched the use of IPv4 addresses until February 2011 when the Internet Assigned Numbers Authority (IANA), operating under the auspices of the Internet Corporation for Assigned Names and Numbers, announced that it had exhausted the supply of IPv4 addresses at the source of its allocation. The Regional Internet Registries (ARIN, LACNIC, RIPE-NCC, AFRINIC, APNIC) still had allocations but APNIC soon exhausted its supply in April 2011 and RIPE-NCC has announced that it has exhausted its supply in September 2012. A market for IPv4 address space has formed but it cannot possibly solve the real need.

The "Internet of Things" is upon us. Mobiles using LTE for data transfer will need end-to-end communication capability. The same may be said for set-top boxes, sensor devices, Internet-enabled automobiles, countless household and office appliances, and, eventually, personal devices that may even be embedded or attached in some way to our bodies. The only sensible solution is to implement IPv6 addressing capability *in parallel* with IPv4. We cannot simply "throw a switch" to convert every device on the Internet from IPv6 to IPv6 addressing. The transition will take years.

This long transition leads to the need for very thoughtful design and implementation of control and management systems that can deal with both IPv4 and IPv6 operating concurrently in the network and in many devices. We cannot even try to form enclaves that are IPv4 only or IPv6 only for "simplicity." Devices that are mobile or portable will regularly encounter both IPv4 and IPv6 and mixed environments. There is also a very good chance that areas of the Internet will be IPv6 only for lack of IPv4 address space. Complex environments involving NATted IPv4 and endto-end IPv6 will also be encountered. It is no wonder that a book of this kind, written by Michael Dooley and Timothy Rooney, will be needed on every Internet engineer's bookshelf (or in his laptop or pad or mobile, cloud client, and digital reader).

Configuration and network management are hard. Dealing with them in a mixed IP packet format environment is even harder. Error messages will be generated for both protocols even if a common fault, for example, a fiber cut, is the proximate cause. Network management systems will need to become much smarter about filtering, correlating, and sorting various error and status or warning messages emerging from a mixed IP addressing environment. The mere fact that the packet headers are potentially larger in IPv6 will create the potential for fragmentation or at least complicate the discovery of the minimum packet size needed to avoid fragmentation of blockage. These are just a few of the questions that need answering. Any system architect preparing to cope with a dual-stack environment will find this book a useful companion and source of advice.

It is not too late to start implementation but it is surely timely. The rest of this decade will see major changes and extensions to the Internet in many dimensions, not the least of which is a massive increase in the number of devices that can be attached and referenced in the system.

Some ISPs have been heard to say "customers are not asking for IPv6" as an excuse to delay implementation. From where I sit, customers should have to know nothing about IPv6. They should have a reasonable expectation that their ISPs will implement dual stack without their asking. It is irresponsible not to move rapidly to deployment of dual stack before there is no more IPv4 address available, even through NAT mechanisms. We must complete the transition to a fully connected IPv6 network as soon as possible. This does not mean we have to abandon the use of IPv4, only that we need connectivity as complete with IPv6 as we have had with IPv4—and we need it now.

CHAPTER

IPv6 DEPLOYMENT DRIVERS

1.1 THE INTERNET: A SUCCESS STORY

The Internet has come a long way. Invented in the late 1960s as a resilient interconnected network of networks for the U.S. Department of Defense, it has evolved into a global communications phenomenon. With the invention of the World Wide Web by Tim Berners-Lee, defining the hypertext linking of information over a network such as the Internet through the use of a web browser, this innovation of simple point-and-click user interface brought the Internet out of government and science laboratories and into ordinary people's lives. Email was the second key Internet application that contributed to the widespread adoption of Internet services during the mid-1990s. Today's Internet users generally find this ubiquitous availability of wide variety of information and applications indispensable in their day-to-day lives. If popular Internet applications like Facebook, YouTube, Twitter, Google, Blogger, shopping, and news sites, and even good old email were suddenly rendered unavailable, most people would not know what to do with themselves!

But the abundance of information and applications on the Internet is not universally available worldwide today. Figure 1-1 illustrates the statistics reported by Internet World Stats indicating the penetration of Internet users as a percentage of overall population in various regions of the world as of mid-2012. Just over one-third of the world's population has access to and use the Internet from work, home, mobile, or wireline. Penetration in North America is highest among the measured regions at more than 78% with Europe second at 63%. Among the Asian population, penetration is only about 28%.

Looking at the same data from a raw numbers perspective, Figure 1-2 illustrates the distribution of Internet users throughout the world. Comparing Figures 1-1 and 1-2, note that while Internet penetration in Asia is less than 28%, the number of users in Asia tops 1 billion, representing 45% of global Internet users, estimated at 2.4 billion by Internet World Stats.

With worldwide penetration of Internet users at just 34%, there seems to be plenty of room for an expanding Internet population. And with the likelihood of multiple devices required per user, this expansion will create accelerated demand for Internet Protocol (IP) addresses. But what circumstances would facilitate such expansion? A recent study by the World Bank concluded that in low- and middle-income countries, for every 10% increase in Internet penetration, the country's average economic growth increases by 1.12% as measured by gross

IPv6 Deployment and Management, First Edition. By Michael Dooley and Timothy Rooney.

^{© 2013} by The Institute of Electrical and Electronics Engineers, Inc. Published 2013 by John Wiley & Sons, Inc.



Figure 1-1. Internet penetration by region [1].

domestic product (GDP) [2]. A 10% increase in broadband penetration yields an average GDP increase of 1.38%. The report also extols other socio-economic benefits for broadband deployment including higher employment, expanding entrepreneurial opportunities, providing social contacts, and delivering public information-based services. While some governments may desire to restrict free access to certain content or applications, the economic correlation to Internet growth is difficult to ignore.

And trends leading up to the present time indicate strong growth over the last decade. Figure 1-3 illustrates growth in Internet users and penetration. These appear to be directly proportional. The compound annual growth rate (CAGR) over this time period, by region, is represented in Table 1-1 and averaged 18% worldwide.



Figure 1-2. Worldwide Internet users by region [1].



Figure 1-3. Internet users history [1].

The number of content providers or "producers" on the Internet as measured by the number of new websites is also growing at an increasingly rapid rate according to Netcraft [3], an Internet research and security services firm. As of June 2012, the total number of discovered unique website hostnames reached nearly 700 million, while active (non-template, based solely on domain registration) sites approached 200 million as shown in Figure 1-4. Both metrics have increased at an accelerated rate over the past two years. New organizations beyond some point in time desiring to operate their own websites or hosting providers will eventually only have access to IPv6 address space for publishing web content.

1.1.1 Supply-Side Issues

Given this history of sustained growth of Internet users and content suppliers, the relatively modest penetration rates, and the economic benefits of broadband deployment and Internet access, it's reasonable to predict that Internet user and producer

Internet users (M)	2000	2011	CAGR (%)
Africa	4.5	139.9	37
Asia	114.3	1016.8	22
Europe	105.1	500.7	15
Middle East	3.3	77.0	33
North America	108.1	273.1	9
Latin America/Caribbean	18.1	235.8	26
Oceania/Australia	7.6	23.9	11
Worldwide total	361.0	2267.2	18

TABLE 1-1. Worldwide Internet User Growth per Region



Figure 1-4. Measured quantity of Internet websites [3].

demand will continue to grow. Unfortunately, the currently available capacity of IPv4 addresses to support this growth is insufficient. With IPv4 address space depleted, for all intents and purposes, the only Internet Protocol available to support this demand is IPv6.

When the Internet Assigned Names and Numbers Authority (IANA) announced, on February 3, 2011, that it had allocated its last remaining IPv4 address space to the Regional Internet Registries (RIRs), that day was the beginning of the end of the Internet as we know it. Figure 1-5 illustrates the IP address space "food chain," with IANA allocating base address blocks to RIRs. IANA is a department within the Internet Corporation for Assigned Names and Numbers (ICANN), which itself is a not-forprofit public-benefit corporation with participants from around the world. IANA is the centralized coordinating body for Internet domain names managing the Domain Name System (DNS) root and several other top-level domains, Internet number resources (IP addresses), and protocol assignments (protocol-specific parameters, e.g., Dynamic Host Configuration Protocol (DHCP) option number assignments).

The RIRs are organizations responsible for allocation of address space within their respective global regions from their corresponding space allotments from IANA:

- AfriNIC (African Network Information Centre)—Africa region
- APNIC (Asia-Pacific Network Information Centre)—Asia-Pacific region
- ARIN (American Registry for Internet Numbers)—North America region including Puerto Rico and some Caribbean Islands



Figure 1-5. The IP address space hierarchy [4].

- LACNIC (Regional Latin-American and Caribbean IP Address Registry)— Latin America and some Caribbean Islands
- RIPE NCC (Réseaux IP Européens Network Coordination Centre)—Europe, the Middle East, and Central Asia.

The RIR system was established with the following goals for IP address allocation to National Internet Registries, Local Internet Registries (LIRs), and Internet Service Providers (ISPs):

- *Uniqueness*. Each IP address must be unique worldwide for global Internet routing.
- *Aggregation.* Hierarchical allocation of address space assures proper routing of IP traffic on the Internet. Without aggregation, routing tables become fragmented which could ultimately create tremendous bottlenecks within the Internet.
- *Conservation.* With IPv4 in particular but also for IPv6 space, address space needs to be distributed according to actual usage requirements.
- *Registration*. A publicly accessible registry of IP address assignments eliminates ambiguity and can help when troubleshooting. This registry is called the *whois* database. Today, there are many whois databases, operated not only by RIRs but also by LIR/ISPs for their respective address spaces.
- *Fairness*. Unbiased address allocation based on true address needs and not on long-term "plans."

Despite efforts to extend the lifetime of IPv4 through technologies such as Network Address Translation (NAT) and Classless Inter-Domain Routing (CIDR) as

well as RIR policies enabling sales and transfers of IPv4 address space, eventually the RIRs will each allocate their last vestige of IPv4 address space to their constituents. The ISPs then in turn will eventually exhaust their IPv4 resources for distribution to their customers, generally enterprise businesses. The APNIC and RIPE NCC RIRs have already exhausted their respective IPv4 space.

1.1.2 Internet at a Crossroads

So what does all this mean? When IPv4 address space runs out among the ISPs within a given RIR's region, any new organization requiring IP address space or existing organizations requiring supplemental IP address space in that region will, by necessity, receive an IPv6 address space allocation. As new organizations initiate web presences, they will be accessible only by IPv6. As new "IPv6-only" organizations join the web, the composition of the Internet itself will slowly change from the homogeneous IPv4 Internet available today to a mixed IPv4/IPv6 Internet in the future. How quickly and to what levels this IPv6 density within the Internet will rise is uncertain.

The growth in number of IPv6-only users is expected to surface initially in Asia. The economies of major Asian countries, particularly China and India, have grown more rapidly in recent years than the rest of the world according to the International Monetary Fund [5] as illustrated in Figure 1-6. From 2000 through 2011, the average annual GDP growth in China was 10.2% and India was 7.1%, while the world average was 2.7%. This in turn has led to rising disposable incomes and government infrastructure investment in communications technologies such as broadband and wireless. Point Topic Ltd estimated that nearly half of the world's broadband net additions for the first half of 2012 were in Asia [6]. Point Topic's report also mirrored that of Internet World Stats, which reported that while Asia boasts the most Internet users at over one billion, its Internet penetration rate is below the world average (33%) at 26%.



Figure 1-6. GDP growth 2006–2010 [7].