


HOST IDENTITY PROTOCOL (HIP)

TOWARDS THE SECURE MOBILE INTERNET

ANDREI GURTOV

 WILEY



TN915.04
G981

Host Identity Protocol (HIP): Towards the Secure Mobile Internet

Andrei Gurtov

Helsinki Institute for Information Technology (HIIT), Finland



A John Wiley & Sons, Ltd, Publication



E2008001639

This edition first published 2008
© 2008 John Wiley & Sons Ltd

Registered office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ,
United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. 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 or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Gurtov, Andrei.

Host Identity Protocol (HIP) : Towards the Secure Mobile Internet / Andrei Gurtov.
p. cm.

Includes bibliographical reference and index.

ISBN 978-0-470-99790-1 (cloth)

1. Wireless Internet—Security measures. 2. Host Identity Protocol (Computer network protocol) I. Title.

TK5103.4885.G87 2008

005.8—dc22

2008006092

A catalogue record for this book is available from the British Library.

ISBN 978-0-470-99790-1 (H/B)

Set in 10/12pt Times by Sunrise Setting Ltd, Torquay, UK.

Printed in Great Britain by CPI Antony Rowe, Chippenham, England.

Host Identity Protocol (HIP)

WILEY SERIES IN COMMUNICATIONS NETWORKING & DISTRIBUTED SYSTEMS

Series Editors: David Hutchison, *Lancaster University, Lancaster, UK*
Serge Fdida, *Université Pierre et Marie Curie, Paris, France*
Joe Sventek, *University of Glasgow, Glasgow, UK*

The 'Wiley Series in Communications Networking & Distributed Systems' is a series of expert-level, technically detailed books covering cutting-edge research, and brand new developments, as well as tutorial-style treatments in networking, middleware and software technologies for communications and distributed systems. The books will provide timely and reliable information about the state-of-the-art to researchers, advanced students and development engineers in the Telecommunications and the Computing sectors.

Other titles in the series:

Wright: *Voice over Packet Networks* 0-471-49516-6 (February 2001)
Jepsen: *Java for Telecommunications* 0-471-49826-2 (July 2001)
Sutton: *Secure Communications* 0-471-49904-8 (December 2001)
Stajano: *Security for Ubiquitous Computing* 0-470-84493-0 (February 2002)
Martin-Flatin: *Web-Based Management of IP Networks and Systems* 0-471-48702-3 (September 2002)
Berman, Fox, Hey: *Grid Computing. Making the Global Infrastructure a Reality* 0-470-85319-0 (March 2003)
Turner, Magill, Marples: *Service Provision. Technologies for Next Generation Communications* 0-470-85066-3 (April 2004)
Welzl: *Network Congestion Control: Managing Internet Traffic* 0-470-02528-X (July 2005)
Raz, Juhola, Serrat-Fernandez, Galis: *Fast and Efficient Context-Aware Services* 0-470-01668-X (April 2006)
Heckmann: *The Competitive Internet Service Provider* 0-470-01293-5 (April 2006)
Dressler: *Self-Organization in Sensor and Actor Networks* 0-470-02820-3 (November 2007)
Berndt: *Towards 4G Technologies: Services with Initiative* 0-470-01031-2 (February 2008)
Jacquenet, Bourdon, Boucadair: *Service Automation and Dynamic Provisioning Techniques in IP/MPLS Environments* 0-470-01829-1 (February 2008)
Minei, Lucek: *MPLS-Enabled Applications: Emerging Developments and New Technologies. Second Edition* 0-470-98644-1 (April 2008)

About the Author

Andrei Gurtov received his M.Sc and Ph.D. degrees in Computer Science from the University of Helsinki, Finland, in 2000 and 2004. At present, he is a Principal Scientist leading the Networking Research group at the Helsinki Institute for Information Technology focusing on the Host Identity Protocol and next generation Internet architecture. He is a co-chair of the IRTF research group on HIP and teaches as an adjunct professor at the Telecommunications and Multimedia Laboratory of the Helsinki University of Technology. Previously, his research focused on the performance of transport protocols in heterogeneous wireless networks. In 2000–2004, he served as a senior researcher at TeliaSonera Finland contributing to performance optimization of GPRS/UMTS networks, intersystem mobility, and IETF standardization. In 2003, he spent six months as a visiting researcher in the International Computer Science Institute at Berkeley working with Dr. Sally Floyd on simulation models of transport protocols in wireless networks. In 2004, he was a consultant at the Ericsson NomadicLab. Dr. Gurtov is a co-author of over 35 publications including research papers, patents, and IETF RFCs. Admitted to the Finnish Union of University Professors at 28, he remains the youngest member at the time of book publication.

Foreword

Jari Arkko

Bob Moskowitz pioneered the idea of the Host Identity Protocol (HIP) in the late 1990s. While the general idea of separating identifiers and locators had been around for a long time, nothing quite as detailed and comprehensive had actually been designed before. The initial designs were little more than sketches but the architecture was appealing. In particular, many people were drawn to the idea of HIP due to the way that security was naturally integrated as a part of the design, or the fact that the same powerful concepts appeared to be capable of solving many different problems, be they about mobility, multihoming, address translation, or stable anchor points.

Despite a lot of interest, HIP was not an overnight success, largely for two reasons. First, in general, deploying any new functionality in the IP layer has proven slow at best. We live in a world with a continuous stream of new Internet innovations, and it may come as a surprise to some that certain parts of the Internet technology are not easy to replace or upgrade. The HIP designers set out to deal with these issues, and today the result is a protocol that does not require upgrades in routers between two communicating hosts, is capable of communicating with legacy hosts, and can support HIP-unaware middleboxes.

Second, the HIP designers took an approach where they wanted to first see their ideas verified and improved in implementations and detailed protocol specifications written, before attempting to acquire any significant deployment. They followed the old principle of the IETF by focusing on running code as opposed to theoretical results or committee agreements. Pekka Nikander, later helped by Thomas Henderson, Andrei Gurtov, a number of implementation teams, and a fair number of other people, started a long process that has now finally been completed.

We are fortunate that Andrei and the rest of the team have decided to write this book. The book describes the HIP architecture and protocols in detail. It also goes beyond the basics, explaining many of the advanced issues and applications, such as middlebox interactions, application programming interfaces, privacy issues, application of distributed hash tables to identifier resolution, and network mobility. Even if you are not deploying or writing code for HIP, the concepts borne out by this protocol design and presented in this book are applicable to future evolution of the Internet architecture in general.

The book comes out at an appropriate time. Official specifications for HIP are coming out of the IETF as RFCs, designs have been verified in many implementations and tests, a number of significant real-world deployments are being discussed, and some are already up and running. A growing body of research on Internet architecture employs ideas from

HIP. Within the set of many identifier–locator separation designs for the Internet, HIP has progressed clearly further than anything else we have so far. It is time to see what HIP can do in larger scale in the real world. In order to make that happen, the world needs a HIP book, and now we have it.

Jari Arkko
Internet Area Director, IETF

Foreword

David Hutchison

Electronic communication, whether using fixed-line or wireless networks, is now a common feature of many activities in everyday life. Enterprises and individuals alike have embraced communication as an indispensable tool. The Internet, the Web and the mobile telephone are three technologies that have made a particular impression. Used in conjunction, these form a powerful combination of technologies that are transforming the ways in which we interact with each other and with information, whether it is for work or leisure. However, these can also be quite literally addictive. There are now reports in the press that many people experience withdrawal symptoms if they are denied access – even temporarily – to their habitual or favourite means of communication. But even more significant is our growing dependency on networks for business use, including for example government, banking, travel, healthcare, monitoring and control: when the network fails – for whatever reason – the consequences may be severe. The working assumption is that the network will act as a utility, i.e. that it will provide an essentially perfect service all of the time. It seems that there is a considerable lack of awareness amongst enterprise owners of the vulnerabilities of networks and a lack of foresight about the impact that failures can have on their operations. Perhaps this is often due to complacency, although ignorance and sometimes unwillingness to invest in appropriate measures cannot be ruled out.

Security and dependability are, then, crucially important aspects of communications and networks that need very much more exposure, as well as considerably more investment by enterprise owners. The Host Identity Protocol (HIP) offers the dual prospects of more secure and more dependable communications through the separation of identity and location, as well as a number of related potential benefits. Although the first HIP draft was submitted in 1999 at the IETF, HIP remained unknown to a wider audience until recently.

This latest – and very welcome – book in the Wiley CNDS Series is written by a co-chair of the HIP Research Group at IRTF. It gives a comprehensive coverage of the subject and should appeal to researchers and practitioners alike in the field of communications and computer networks, not just to those interested in security and mobility.

David Hutchison
Lancaster University

Preface

The main goal of this book is to present a well-structured, readable and compact overview of the Host Identity Protocol with relevant extensions to the Internet architecture and infrastructure. HIP is a new protocol developed by the Internet Engineering Task Force (IETF) to address shortcomings of the current Internet in supporting secure host mobility and multihoming.

As an alternative to Mobile IP, HIP helps to solve security and Denial-of-Service issues that hindered the deployment of IP mobility in an architecturally clean way. In the TCP/IP stack, HIP is positioned between the network (IPv4, IPv6) and transport (TCP, UDP, SCTP, DCCP) layers. The transport protocols use a cryptographic host identity instead of IP addresses.

Since 2000, HIP had been specified in the IETF. Many major telecommunication vendors and operators have started internal activities involving HIP. At the time of writing, an effort is ongoing to create an experimental deployment of HIP in the Internet. HIP is used by several large research projects and often cited in the networking articles. Several public open-source interoperating HIP implementations are available for various platforms, including Linux, BSD, Windows, and Mac OS.

The need for HIP

The Internet has grown tremendously over the past twenty years and become a part of life for millions of people. The basic TCP/IP technology has served us very well. However, important issues such as mobility of Internet hosts over separate IP networks and simultaneous connections to several networks were not a part of the original Internet design. Furthermore, when the Internet grew from a small university network up to a global communication infrastructure, many security issues became apparent. The lack of reliable host authentication has prevented deployment of existing IP mobility extensions. Often, public Internet servers face Denial-of-Service (DoS) attacks that make the service unavailable to other users.

HIP is developed to address these issues in an integrated approach that fits well within the TCP/IP architecture. The original ideas on the separating of host identity and location in the Internet date back to Saltzer in “RFC 1498 On the Naming and Binding of Network Destinations”. However, only recent advances in public key cryptography and new requirements of portable terminals have made the actual design and implementation possible.

It is true that HIP is only one of many proposals developed recently in the IETF in the area of security and mobility. Compared with other proposals that often solve only a small part of the problem, HIP integrates host mobility and multihoming in a simple and elegant way.

Independently of which proposal will be deployed in the Internet in the future, we believe that ideas stemming from design and experiments with HIP provide indispensable knowledge for anybody interested in next-generation networking.

Intended audience

We hope the book will be interesting for engineers implementing new Internet protocols, researchers working in the area of next-generation Internet, and students working on Master or Doctoral theses in the area of Internet security, mobility, and multihoming.

For industry engineers, the book includes detailed information on protocol messages and packet headers. Although only the IETF specifications offer complete information necessary for the implementation, the specifications might be difficult to comprehend at first without an overall picture of HIP architecture. The specifications are divided into several RFCs with many cross-references and often repeating parts of the text. Therefore, we hope that the book will be useful to be read before the specifications to obtain a general understanding of the area. In addition to published specifications, the book also includes several experimental or historic proposals from expired Internet drafts that are not easily found otherwise. The ASCII diagrams from the specifications are redrawn to be easily understandable.

For researchers, the book includes an overview of several research articles on HIP infrastructure and design alternatives. The future Internet architecture is a subject of many research projects in the USA such as Future Internet Network Design (FIND) and in the EU (Framework Program 7). The concept of identifier–locator split forms the core of HIP and is often applied in other protocol areas. Several large projects, such as the FP6 Ambient Networks, have adopted HIP as an architectural backbone. Therefore, even the researchers not directly working with HIP may find it useful to understand the concepts and design decisions behind HIP.

University students participating in networking courses and research seminars may find the book useful in covering the Internet security, mobility, and multihoming issues comprehensively. The book can be also used as a text book on a course focusing on networking protocol design, where HIP is taken as a possible example.

The prerequisites for the book include knowledge of TCP/IP networking. Experience with UNIX socket programming is helpful to understand certain parts of the book, such as the socket API interface. Readers would benefit from knowledge of basic cryptography, although an overview of symmetric and asymmetric cryptography, IP security protocols, and DNS is given at the beginning of the book.

Coverage

The book covers all IETF specifications of HIP, several experimental extensions proposed in the IETF, and research results relevant to the development and deployment of HIP. The HIP specifications produced by the HIP Working Group in the IETF include the base specification, IPsec ESP encapsulation, mobility and multihoming extensions, rendezvous server, registration extensions, NAT traversal, and API specification. The experimental proposals from the HIP Research group in IRTF include the DHT interface, opportunistic mode, HIP-aware middleboxes, service discovery and simultaneous multiaccess, and lightweight HIP.

Research results cover the use of overlay networks for routing HIP control packets, HIP privacy extensions, multicast, and integration of HIP with other protocols such as SIP and Mobile IP.

We have considered including a CD-ROM with HIP specifications and implementations together with the book. Instead, for a number of reasons, we decided to create a web page with the book that can be kept up-to-date. First, few HIP specifications are still in the final standardization stages and can change after publication of the book. IETF and IRTF documents are freely available from the Internet. The HIP implementations are constantly improving and supplying a deprecated version on the CD-ROM would only generate obsolete feedback. An overview and setup instructions for HIP on Linux (HIPL) implementation are given in Appendix A towards the end of the book.

The book includes several practical examples of HIP packet exchanges captured with the freeware network analyzer Wireshark. Its latest version, modified to support the HIP packet formats, is available from the book web page (<http://www.hipbook.net>).

Acknowledgments

This book focuses on the Host Identity Protocol, which has been created by the joint effort of many people. Robert Moskowitz proposed HIP in IETF and served as a coordinator of development efforts in the early days of standardization. Pekka Nikander, Gonzalo Camarillo, and Tom Henderson served as main proponents of HIP architecture and specifications in the following years. Many people have contributed to creating HIP specifications in IETF/IRTF and developing HIP implementations. Many thanks to all of you for doing a great job!

This book would not have been possible without the help of many people. I thank Lars Eggert for helping to draft the book proposal and serving as a co-initiator of the book project. Thanks to anonymous reviewers who have evaluated the proposal and given plenty of useful suggestions for improving it. Tobias Heer contributed two valuable chapters to the book, covering background material and lightweight HIP. Jari Arkko wrote a nice foreword for the book. Samu Varjonen, Abhinav Pathak, Petri Jokela, and René Hummen reviewed various parts of the book. However, any remaining mistakes are those of the author. Assel Mukhametzhanova helped to maintain the project bibliography database. Thanks to Birgit Gruber, Anna Smart, and Sarah Hinton from John Wiley and Sons for help and support in getting the book published.

I am grateful to the Helsinki Institute for Information Technology and Helsinki University of Technology for hosting HIP-related research projects. I thank the Finnish Funding Agency for Technology and Innovation (Tekes), Academy of Finland, and companies that funded research work at HIIT: Elisa, Ericsson, Finnish Defence Forces, Nokia, Secgo, and TeliaSonera.

The book covers a wide range of engineering and research efforts related to HIP. I would especially like to mention a few people whose publications are referenced in the book: Jeff Ahrenholz, Petri Jokela, Andrey Khurri, Miika Komu, Teemu Koponen, Dmitry Korzun, Julien Laganier, Richard Paine, Abhinav Pathak, Laura Takkinen, Hannes Tschofenig, Samu Varjonen, Essi Vehmersalo, Rolland Vida, Ekaterina Vorobyeva, and Jukka Ylitalo. This list is by no means complete and extends to everyone else in the HIP community.

Last but not least, I would like to thank my core and extended family for love and support in getting the book written.

Andrei Gurtov
Helsinki

Abbreviations

ACK	Acknowledgment packet
ACL	Access Control List
ADSL	Asynchronous Digital Subscriber Line
AES	Advanced Encryption Standard
AH	Authentication Header
API	Application Programming Interface
AR	Access Router
ARP	Address Resolution Protocol
ASM	Any Source Multicast
BE	Base Exchange
BEET	Bound End-to-End Tunnel
BOS	Bootstrap packet
BSD	Berkeley Software Distribution
CA	Certificate Authority
CBA	Credit-Based Authorization
CBC	Cipher Block Chaining
CPU	Central Processing Unit
CRL	Certificate Revocation List
CS	Cryptographic Session
DCCP	Datagram Congestion Control Protocol
DH	Diffie–Hellman
DHCP	Dynamic Host Configuration Protocol
DHT	Distributed Hash Table
DNSSEC	Domain Name System with Security
DoS	Denial-of-Service
DR	Designated Router
DSA	Digital Signature Algorithm
ED	Endpoint Descriptor
ESP	Encapsulated Security Payload
FQDN	Fully Qualified Domain Name
FTP	File Transfer Protocol
GGSN	Gateway GPRS Support Node
GNU	GNU is not UNIX
GPL	General Public License
GPRS	General Packet Radio Service

GRUU	Globally Routable UA URI
GSM	Global System for Mobile communications
GUI	Graphical User Interface
HCVF	Hash Chain Value Parameter
Hi3	Host Identity Indirection Infrastructure
HIP	Host Identity Protocol
HIPD	HIP Daemon
HIPL	HIP for Linux
HISM	Host Identity Specific Multicast
HIT	Host Identity Tag
HMAC	Hash Message Authentication Code
HMIP	Hierarchical Mobile IP
HTTP	Hyper Text Transfer Protocol
i3	Internet Indirection Infrastructure
IANA	Internet Assigned Numbers Authority
ICE	Interactive Connectivity Establishment
ICMP	Internet Control Message Protocol
ID	Identifier
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IHC	Interactive Hash Chain
IKE	Internet Key Exchange
IMS	IP Multimedia Subsystem
IPv6	Internet Protocol version 6
IRTF	Internet Research Task Force
ISP	Internet Service Provider
L2TP	Layer 2 Tunneling Protocol
LHIP	Lightweight HIP
LRVS	Local Rendezvous Server
LSI	Local Scope Identifier
MAC	Message Authentication Code
MIME	Multipurpose Internet Mail Extensions
MIP	Mobile IP
MITM	Man-In-The-Middle
MKI	Master Key Identifier
MLD	Multicast Listener Discovery
MOBIKE	Mobile Key Exchange
MODP	More Modular Exponential
MR	Mobile Router
MTU	Maximum Transmission Unit
NACK	Negative Acknowledgment
NAT	Network Address Translator
NEMO	Network Mobility
NIC	Network Interface Card
OCALA	Overlay Convergence Architecture for Legacy Applications

ORCHID	Overlay Routable Cryptographic Hash Identifier
OS	Operating System
P2P	Peer-to-Peer
PACK	Pre-Acknowledgment
PDA	Personal Digital Assistant
PGP	Pretty Good Privacy
PIDF	Presence Information Data Format
PISA	P2P Internet Sharing Architecture
POSIX	Portable Operating System Interface
PSIG	Pre-Signature Packet
PSP	Pre-Signature Parameter
RADIUS	Remote Authentication Dial In User Service
RFC	Request for Comments
RFID	Radio-Frequency Identification
PKI	Public Key Infrastructure
ROC	Rollover Counter
RPC	Remote Procedure Call
RR	Resource Record for DNS
RSA	Rivest–Shamir–Adleman algorithm
RTO	Retransmission Timeout
RTP	Real-time Transmission Protocol
RTT	Round-Trip Time
RVA	Rendezvous agent
RVS	Rendezvous server
SA	Security Association
SACK	Selective Acknowledgment
SAD	Security Association Database
SCTP	Stream Control Transmission Protocol
SD	Service Discovery
SDP	Service Discovery Protocol
SHA	Secure Hash Algorithm
SIGMA	Signature and MAC
SIM	Subscriber Identity Module
SIMA	Simultaneous Multi Access
SIP	Session Initiation Protocol
SMTP	Simple Mail Transfer Protocol
SPD	Security Parameter Database
SPI	Security Parameter Index
SRTP	Secure Real-time Transmission Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
SSM	Source Specific Multicast
SSO	Single Sign-On
SSRC	Synchronization Source
STUN	Simple Traversal of UDP through NATs
SYN	Synchronization packet for TCP

TCP	Transmission Control Protocol
TESLA	Timed Efficient Stream Loss-tolerant Authentication
TLS	Transport Layer Security
TLV	Type-Length-Value
TRIG	Trigger packet
TTL	Time to Live
UA	User Agent
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunication System
URI	Universal Resource Identifier
UUID	Universally Unique Identifier
VIGMP	Version-Independent Group Management Protocol
VM	Virtual Machine
VMM	Virtual Machine Monitor
VoWLAN	Voice over WLAN
VPN	Virtual Private Network
Wi-Fi	Wireless Fidelity
WIMP	Weak Identifier Multihoming Protocol
WLAN	Wireless Local Area Network
WPA	Wi-Fi Protected Access
XFRM	Linux IPsec Transforms
XML	Extensible Markup Language