INFORMATION SECURITY SERIES

“十二五”国家重点图书

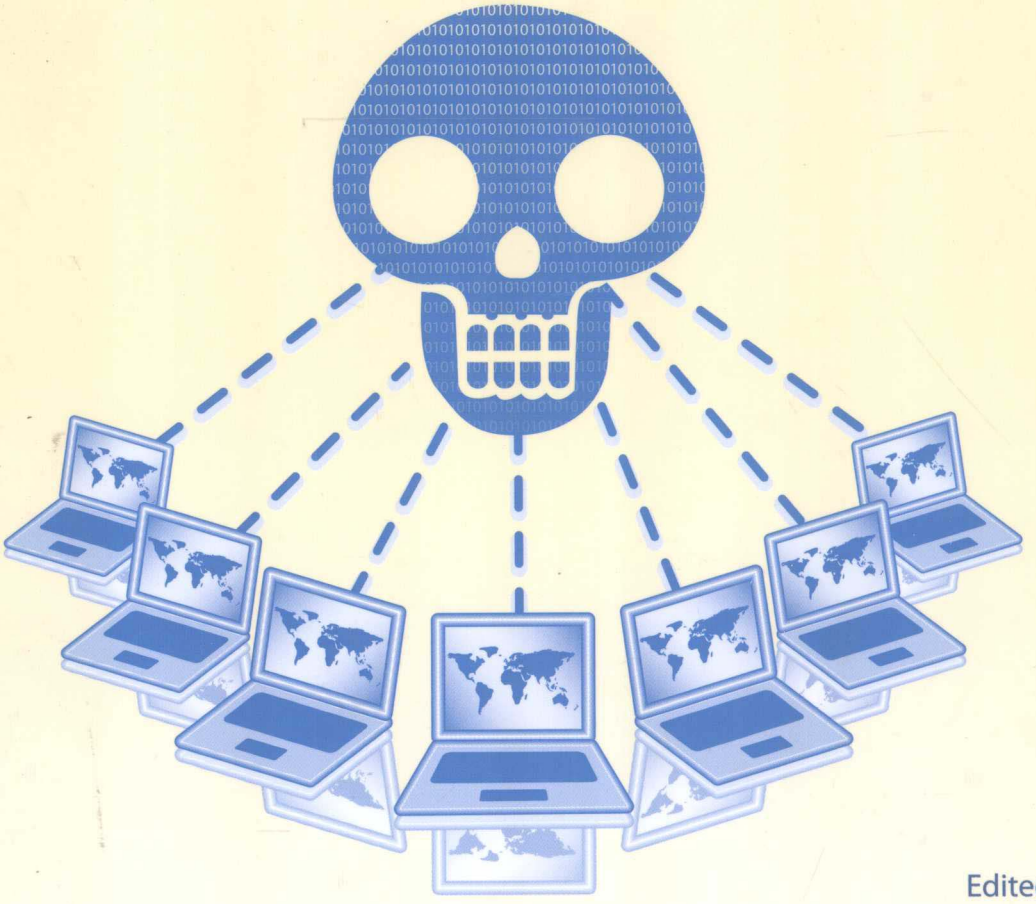

Edited by

Markus Jakobsson

The Death of the Internet

因特网死亡

高等教育出版社
HIGHER EDUCATION PRESS

“十二五”国家重点图书

INFORMATION SECURITY SERIES

The Death of the Internet

因特网死亡

YINTEWANG SIWANG

Edited by

Markus Jakobsson

高等教育出版社·北京
HIGHER EDUCATION PRESS BEIJING

图书在版编目(CIP)数据

因特网死亡 = The Death of the Internet : 英文 / (美) 雅各布森 (Jakobsson, M.) 主编. — 北京 : 高等教育出版社, 2012.10
(信息安全系列)
ISBN 978-7-04-030117-5

Ⅰ. ①因… Ⅱ. ①雅… Ⅲ. ①互联网络-信息安全-研究-英文 Ⅳ. ①TP393.4

中国版本图书馆 CIP 数据核字 (2012) 第 224351 号

策划编辑 陈红英　责任编辑 陈红英　封面设计 张 楠　版式设计 杜微言
责任校对 张小镝　责任印制 朱学忠

出版发行	高等教育出版社	咨询电话	400-810-0598
社　　址	北京市西城区德外大街 4 号	网　　址	http://www.hep.edu.cn
邮政编码	100120		http://www.hep.com.cn
印　　刷	涿州市星河印刷有限公司	网上订购	http://www.landraco.com
开　　本	787mm × 1092mm　1/16		http://www.landraco.com.cn
印　　张	25	版　　次	2012 年 10 月第 1 版
字　　数	660 千字	印　　次	2012 年 10 月第 1 次印刷
购书热线	010-58581118	定　　价	79.00 元

物 料 号　30117-00

“十二五”国家重点图书

INFORMATION SECURITY SERIES

INFORMATION SECURITY SERIES

Information Security Series systematically introduces the fundamentals of information security design and application. The goals of the Series are:

- to provide fundamental and emerging theories and techniques to stimulate more research in cryptology, algorithms, protocols, and architectures
- to inspire professionals to understand the issues behind important security problems and the ideas behind the solutions
- to give references and suggestions for additional reading and further study

Publications consist of advanced textbooks for graduate students as well as researcher and practitioner references covering the key areas, including but not limited to:

— Modern Cryptography
— Cryptographic Protocols and Network Security Protocols
— Computer Architecture and Security
— Database Security
— Multimedia Security
— Computer Forensics
— Intrusion Detection

LEAD EDITORS

EDITORIAL BOARD

For A and Art.

Foreword

It is tempting to believe that Internet security is somebody else's problem, or that it is a problem that eventually will vanish, as technology improves. This is a dangerous belief; burglary has not vanished because of improvements in door and window locks, and Internet security is similarly unlikely to change as technology gets better. In considering the recent past, I think it is easy to say that increased awareness of Internet security has not had much impact on the rate of victimization of consumers. Given how many Internet users there now are—significantly over 1 billion people, and as high as 2 billion by some counts—consciously improving the rate of awareness of populations at this scale is incredibly difficult and time consuming. And technology leaps often make things worse, quite as much as they make things better.

This is not a Chicken Little situation of "the sky is falling, the sky is falling," but we will find ourselves in an increasingly difficult situation soon, unless we start to pay more attention to Internet problems than we have done in the last few years. To begin with, we need to ensure that we not only understand the problem—and not just its manifestation—but also its underlying reasons for being. Then, we need to start designing new technology to save ourselves—and our users—from greed and crime, from the very things that have made the Internet so successful: how efforts scale, how everybody can participate, and from the low costs of entry. We also need to consider possible regulation, as it is unlikely that the technology industry's call for self-regulation will be heeded any more than that of the road or aviation industries in the 1910s and 1920s.

Let me tell you a little bit about myself first, to give you some perspective on my viewpoint. I have been a technology strategist for nearly three decades, honing my craft first within the confines of a corporate environment, but in recent years increasingly looking outside that mothership. In the early 2000s, I spent quite a bit of time in the identity space—I was President of the Liberty Alliance, which was an open standards consortium that developed the first meaningful identity federation protocol, SAML 2.0. Since 2006, I have been CISO at PayPal. Given PayPal's global reach, the size of our user base (in mid-2011, over 100 million active customers) and the nature of our systems, which move money from any arbitrary point A to point B on the planet, we tend to find ourselves at the leading edge of new classes of criminal attacks. Willy nilly, we find ourselves having to craft solutions to problems that the rest of the industry barely recognizes as problems, let alone admits there are solutions for.

Here is what I believe we must do. We must begin by understanding our vulnerabilities, whether they are social or technical. We must then instrument our systems, both technical and societal, to collect metrics about everything relevant. After all, how can we argue about reality objectively without having hard data about it? And then we must take the next step, and measure what is not yet reality—trying to predict behaviors and vulnerabilities, in other words. With these hard won insights, we must then create plans for the future. You cannot design a system—especially not a security system—without understanding what affects security and *everything* that affects security.

That is what this book is about. It describes the Internet, and the mobile Internet, in a crisp and convincing manner. It is infused with the anticipation of trends and describes how these will affect us, for good and bad. And it gives examples of novel approaches that we can take to change the course of the future, and avoid what otherwise may become what the title of the book states—The Death of the Internet.

I encourage anyone with an interest in the Internet, in technology, in online commerce, or indeed in a fair and open society to read this book. These are important topics and this book does an excellent job in provoking alternative ways of thinking about them.

MICHAEL BARRETT

San Jose, CA
May 2012
Chief Information Security Officer, Paypal

Preface

Imagine life without electricity. Although most of humankind has managed just well without electricity, its loss would certainly impact society most profoundly. Now imagine that hundreds of thousands of criminals all over the world could make a quick profit by doing something that—little by little—killed the electric infrastructure. And that politically adversarial individuals and governments could speed up this looming catastrophe if they wanted to. It is terrible to imagine. And yet, it is a very real threat—although to the *Internet.*

Criminals have an array of ways to abuse the Internet. Most commonly in order to make money: Internet crime is both profitable and safe for criminals to engage in. Online crime scales exceptionally well, and is fast to perpetrate. It is often difficult to identify abuse, and almost always difficult to track down the criminals. And most of the time, offenders who are detected and blocked simply vanish, only to resurface with a new pseudonym shortly thereafter. While financial abuses are difficult to block and track, *politically* motivated abuse is yet harder to control. This is since politically motivated attacks do not involve taking money out of the system—which is often the hardest and riskiest part of online crime.

Any disruptions to the Internet would send shock waves through society. It would affect telephony; banking; how corporations do business; how the energy grid is controlled; and how many of us make a living. It would disrupt government, media, and military. It would impact our entire infrastructure—including our *trust* infrastructure.

Like a bridge that may tolerate increasing strains until it comes crashing down, the Internet may hold up well until the tipping point is reached. We must not wait for that moment. We must understand the problems, and defend against them—before they develop, if possible. We must understand how things can go wrong, and how we can engineer things better.

This book describes the problems the Internet is facing, and gives examples of some possible solutions. You do not need a deep technical background to understand the *general nature* of these. At the same time, each chapter has in-depth material for readers who do not want to stop at understanding the general concepts, but who want to know *exactly* how things work.

I hope that the insights that you will gain by reading this book will help you make decisions or designs—depending on who you are—that will help rescuing the Internet from the assault it is under.

MARKUS JAKOBSSON, PHD

Mountain View, CA
May 2012
Principal Scientist of Consumer Security, PayPal

Is the Title of this Book a Joke?

Maybe you thought the title of this book was simply chosen to demand attention, or a silly joke, and that the Internet cannot realistically be killed. If that is so, I want to start off by convincing you that it just is not so. *Killed* sounds drastic. Let us for a moment say "rendered useless" or "more or less abandoned." Is *that* possible?

You may ask: What would render the Internet useless? And what would make people abandon it? Let me start off my explanation with an analogy.

Think about traveling by air. We all know that some flights get delayed. But most are not, and those that are delayed are only *reasonably* delayed. A few hours at most, but most often, they are delayed much less than that. We also know that people die in airplane accidents most every year. But most travelers arrive safe.

Imagine that most flights arrived late, and often quite drastically delayed. Maybe a week, maybe two. And imagine further that airplane accidents became dramatically more common. Maybe half of all flights would not arrive *at all*, but everybody on the flight would end up on the bottom of the ocean.

Nobody in their right mind would fly if this were so. It would render air travel useless from a practical perspective, and people would abandon it and take the train, or even walk rather than setting foot in an airplane. In other words, these increases of inconveniences and risks would *kill* aviation.

Now, let us talk about the Internet. The first commercial spam message was sent on March 5, 1994 by an Arizona-based law firm. In the years to follow, spam became more and more prevalent. Still, it shocked a lot of people when the amount of spam overtook that of legitimate emails. Many did not think that it could *ever* become that bad. In spite of impressive advancements in spam blocking technologies, *less than 5%* of email is legitimate at the time of writing, and most of us receive one or two spam messages every day. But a spam message that manages to sneak by the filters typically only wastes a few seconds of our time, with no further consequences to the typical recipient. Things *could* be worse.

What would happen if, in spite of our best efforts to keep the Internet secure, less than 5% of websites were secure, and the rest were hosting malware? Defenses may improve, but what if our normal activities still resulted in malware slipping through once or twice a day? What would be the consequences to online commerce if only 5% of advertisements were honest, and the rest attempted to defraud buyers? Looking from the other side, what would advertisers do if 95% of users had infected computers that constantly were committing clickfraud? How would we be affected if less than 5% of the information we find was correct?

The likely answer is that there would be a *drastic* change in how we use the Internet, and what we *dare* to use it for. Following our aviation analogy, these increases of inconveniences and risks could and *would* kill the Internet.

That is what this book is about. This book explains what might kill the Internet by making it useless and dangerous. And how we are inching toward a tipping point where the result would be the death of the Internet. Where is that point? Nobody knows.

This book also investigates what can be done to stop that from happening, given a thorough understanding of what the problem is. It does not contain an exhaustive list of all the dangers, and certainly not a complete list of meaningful solutions. But it *does* explain how to think about the problems and the solutions in a way that helps you—and others like you—start thinking about how we can prevent the death of the Internet. We depend too much on it to let it go.

So, no, the title of the book is not a cheap attention grabber or a joke. *I* am serious when I say that the *situation* is too serious.

Acknowledgments

Internet security—and *in*security—is both a compelling and terrifying topic to write about. Even more than other aspects of the Internet, it is amorphous and under constant evolution—fueled by both the introduction of new features and services and the criminal realization that these offer new opportunities. It is a vast topic. It is technical, legal, and social. It requires an understanding of the markets, computing, and psychology. You may feel that reading this book is much like drinking from a fire hose. This is also how writing it has been.

We would like to thank BITS and the BITS Security Steering Committee for permitting the reuse of portions of the BITS Malware Risks and Remediation Report. The full report, developed by members of the BITS Security Working Group, is publicly available at http://www.bits.org. BITS addresses issues at the intersection of financial services, technology, and public policy, where industry cooperation serves the public good, such as critical infrastructure protection, fraud prevention, and the safety of financial services. BITS is the technology policy division of the Financial Services Roundtable, which represents 100 of the largest integrated financial services companies providing banking, insurance, and investment products and services to the American consumer.

I could not have pulled this off on my own, and I am indebted to my many contributors, all of whom have invested their time and passion in making this book fantastic. In particular, I want to thank Ruj Akavipat, Adam Barth, Dan Boneh, Garth Bruen, Igor Bulavko, Elie Bursztein, Juan Caballero, Richard Chow, Michael Conover, Mayank Dhiman, Ori Eisen, Bruno Gonçalves, Baptiste Gourdin, Mark Grandcolas, Jeff Hodges, Mohammad Hossein Manshaei, Jean-Pierre Hubaux, Nathaniel Husted, Hampus Jakobsson, William Leddy, Debin Liu, Filippo Menczer, Steven Myers, Dimitar Nikolov, Yuan Niu, Adrienne Porter Felt, Ariel Rabkin, Emilee Rader, Gustav Rydstedt, Elaine Shi, Christopher Soghoian, Dawn Song, Sid Stamm, Andy Steingruebl, Nevena Vratonjic, David Wagner, Rick Wash, Ruilin Zhu, and Members of the BITS Security Working Group and staff leads Greg Rattray and Andrew Kennedy.

I know few people as hardworking and dedicated as Liu Yang, and I owe my sanity to him for helping me with all practical issues surrounding conversions to LaTeX, hacking the template to make things look nice, and ensuring that everything was complete.

I am also thankful for the administrative help I have received from Wiley and HEP. Finally, thanks to all my wonderful colleagues, many of whom also contributed to this book.

Contributors

Ruj Akavipat, Department of Computer Engineering, Mahidol University, Bangkok, Thailand
Adam Barth, IMDEA Software Institute, Google, Inc., University of California, Berkeley, San Francisco, CA, USA
Dan Boneh, Department of Computer Science and Electrical Engineering, Stanford University, Stanford, CA, USA
Garth Bruen, KnujOn.com LLC, Brookline, MA, USA
Igor Bulavko, PayPal, Inc., San Jose, CA, USA
Elie Bursztein, Security Laboratory, Computer Science Department, Stanford University, Stanford, CA, USA
Juan Caballero, IMDEA Software Institute, Madrid, Spain
Richard Chow, Palo Alto Research Center (PARC), Palo Alto, CA, USA
Michael Conover, School of Informatics and Computing, Indiana University, Bloomington, IN, USA
Mayank Dhiman, PEC University of Technology, Ambala, Chandigarh, India
Ori Eisen, The 41st Parameter, Inc., Scottsdale, AZ, USA
Adrienne Porter Felt, Computer Science Division, University of California, Berkeley, CA, USA
Aurélien Francillon, EURECOM, Sophia Antipolis, France
Philippe Golle, Google, Inc., Mountain View, CA, USA
Bruno Gonçalves, Northeastern University, Boston, MA, USA
Nathan Good, Principal Good Research LLC, Berkeley, CA USA
Baptiste Gourdin, LSV, INRIA & ENS-Cachan, Paris, France
Mark Grandcolas, FatSkunk, Inc., Mountain View, CA, USA
Jeff Hodges, PayPal, Inc., San Jose, CA, USA
Jean-Pierre Hubaux, School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland
Nathaniel Husted, School of Informatics and Computing, Indiana University, Bloomington, IN, USA
Hampus Jakobsson, Independent Researcher, Limhamn, Sweden
Markus Jakobsson, PayPal, Inc., San Jose, CA, USA
Andrew Kennedy, BITS/The Financial Services Roundtable, Washington, DC, USA
William Leddy, PayPal, Inc., Georgetown, Washington, DC, USA
Debin Liu, Information Risk Management, PayPal, Inc., Austin, TX, USA
Mohammad Hossein Manshaei, Department of Electrical and Computer Engineering, Isfahan University of Technology, Isfahan, Iran
Ryusuke Masuoka, Fujitsu, Sunnyvale, CA, USA
Filippo Menczer, School of Informatics and Computing, Indiana University, Bloomington, IN, USA

Jesus Molina, Fujitsu, Sunnyvale, CA, USA
Steven Myers, School of Informatics and Computing, Indiana University, Bloomington, IN, USA
Dimitar Nikolov, School of Informatics and Computing, Indiana University, Bloomington, IN, USA
Yuan Niu, Yahoo!, Sunnyvale; University of California, Davis, CA, USA
Adrian Perrig, Cybersecurity Laboratory (CyLab), Department of Electrical and Computer Engineering, Department of Engineering and Public Policy, and School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA
Ariel Rabkin, Electrical Engineering and Computer Science Department, University of California, Berkeley, CA, USA
Emilee Rader, Department of Telecommunication, Information Studies and Media, College of Communication Arts and Sciences, Michigan State University, East Lansing, MI, USA
Greg Rattray, Delta Risk LLC, Washington, DC, USA
Gustav Rydstedt, Blizzard Entertainment, Huntington Beach, CA, USA
Elaine Shi, Palo Alto Research Center (PARC), University of California, Berkeley, CA, USA
Christopher Soghoian, Center for Applied Cybersecurity Research, Indiana University, Bloomington, IN, USA
Dawn Song, IMDEA Software Institute, Google, Inc., Computer Science Department, University of California, Berkeley, CA, USA
Jeff Song, Fujitsu, Sunnyvale, CA, USA
Sid Stamm, Independent Security and Privacy Researcher, Santa Clara, CA, USA
Andy Steingruebl, Information Risk Management, PayPal, Inc., San Jose, CA, USA
Dahn Tamir, Entropy Management Services, Techlist, Las Vegas, NV, USA
Nevena Vratonjic, School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland
David Wagner, Computer Science Division, University of California, Berkeley, CA, USA
Rick Wash, School of Journalism, and Department of Telecommunication, Information Studies and Media, Michigan State University, East Lansing, MI, USA
Ruilin Zhu, Peking University, Beijing, China

Contents