

LNCS 4189

Dieter Gollmann
Jan Meier
Andrei Sabelfeld (Eds.)

Computer Security – ESORICS 2006

11th European Symposium on Research in Computer Security
Hamburg, Germany, September 2006
Proceedings

Dieter Gollmann Jan Meier
Andrei Sabelfeld (Eds.)

Computer Security – ESORICS 2006

11th European Symposium on
Research in Computer Security
Hamburg, Germany, September 18-20, 2006
Proceedings

Volume Editors

Dieter Gollmann

Jan Meier

TU Hamburg-Harburg

Harburger Schlossstr. 20, 21079 Hamburg-Harburg, Germany

E-mail: diego@tu-harburg.de, j.meier@tuhh.de

Andrei Sabelfeld

Chalmers University of Technology

Dept. of Computer Science and Engineering

EDIT Bldg., Rännvägen 6B, 41296 Gothenborg, Sweden

E-mail: andrei@cs.chalmers.se

Library of Congress Control Number: 2006932266

CR Subject Classification (1998): E.3, D.4.6, C.2.0, H.2.0, K.6.5, K.4.4

LNCS Sublibrary: SL 4 – Security and Cryptology

ISSN 0302-9743

ISBN-10 3-540-44601-X Springer Berlin Heidelberg New York

ISBN-13 978-3-540-44601-9 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper SPIN: 11863908 06/3142 5 4 3 2 1 0

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Foreword

These proceedings contain the papers selected for presentation at the 11th European Symposium on Research in Computer Security – ESORICS, held in Hamburg, Germany, September 18-20, 2006.

The 160 papers submitted to ESORICS were each reviewed by at least three members of the program committee. A two-week discussion phase was then held electronically, where committee members could comment on all papers and all reviews. Finally, 32 papers were selected for presentation at ESORICS, giving an acceptance rate of about 21%.

In 2005, three specialized security workshops were organized in affiliation with ESORICS. This trend has continued. In addition to RAID, which is already a well established event in its own right, there were four more workshops this year, ESAS 2006, EuDiRights 06, STM 06, and FEE2, further strengthening the rôle of ESORICS as the major European conference on security research.

There were many volunteers who offered their time and energy to put together the symposium and who deserve our acknowledgment. We want to thank all the members of the program committee and the external reviewers for their hard work in evaluating and discussing the submissions. We are also very grateful to all those people whose work ensured a smooth organization: Joachim Posegga, who served as General Chair; Andreas Günter and his team at HITeC for taking on the conference management; Klaus-Peter Kossakowski for his efforts as Sponsorship Chair; Jan Meier for managing the ESORICS Web site, and Joachim Stehmann for the Web design; and Martin Johns for dealing with the growing number of affiliated workshops.

Last, but certainly not least, our thanks go to all the authors who submitted papers and all the attendees. We hope you found the program stimulating.

July 2006

Dieter Gollmann and Andrei Sabelfeld

Organization

General Chair

Joachim Posegga

University of Hamburg, Germany

Program Committee

Tuomas Aura

Michael Backes

Gilles Barthe

Lynn Batten

Giampaolo Bella

Joachim Biskup

Jan Camenisch

Jason Crampton

Frederic Cuppens

Marc Dacier

George Danezis

Sabrina De Capitani di Vimercati

Robert Deng

Ulfar Erlingsson

Simon Foley

Philippe Golle

Dieter Gollmann (co-chair)

Pieter Hartel

Jaap-Henk Hoepman

Sushil Jajodia

Alan Jeffrey

Audun Jøsang

Jan Jürjens

Markus Kuhn

Xuejia Lai

Kwok-Yan Lam

Volkmar Lotz

Heiko Mantel

Vashek Matyas

Flemming Nielson

Microsoft Research, UK

Saarland University, Germany

INRIA Sophia-Antipolis, France

Deakin University, Australia

University of Catania, Italy

University of Dortmund, Germany

IBM Zurich Research Laboratory,
Switzerland

Royal Holloway, UK

ENST Bretagne, France

Institut Eurécom, France

University of Leuven, Belgium

University of Milan, Italy

Singapore Management University,
Singapore

Microsoft Research, USA

University College Cork, Ireland

Palo Alto Research Center, USA

Hamburg University of Technology,
Germany

Twente University, Netherlands

Radboud University Nijmegen,
Netherlands

George Mason University, USA

Bell Labs, USA

QUT, Australia

TU Munich, Germany

University of Cambridge, UK

Shanghai Jiaotong University, PR China

Tsinghua University, PR China

SAP, France

RWTH Aachen, Germany

Masaryk University Brno, Czech Republic

DTU, Denmark

VIII Organization

Peter Ryan	University of Newcastle upon Tyne, UK
Andrei Sabelfeld (co-chair)	Chalmers University, Sweden
Babak Sadighi	SICS, Sweden
Kazue Sako	NEC Corporation, Japan
Andre Scedrov	U. Pennsylvania, USA
Einar Snekkenes	Gjøvik University College, Norway
Eijiro Sumii	Tohoku University, Japan
Paul Syverson	Naval Research Laboratory, USA
Mariemma I. Yagiüe	University of Malaga, Spain
Alf Zugenmaier	DoCoMo Labs Europe, Germany

Additional Reviewers

Imad Aad, Pedro Adão, Ana Almeida Matos, Toshinori Araki, Olav Bandmann, Vicente Benjumea, Gustavo Betarte, Stefano Bistarelli, Bruno Blanchet, Damiano Bolzoni, Ahmed Bouabdallah, Jeremy Bryans, Marzia Buscemi, Jan Cederrquist, Shiping Chen, Lukasz Chmielewski, Daniel J.T. Chong, Morshed Chowdhury, Siu-Leung Chung, Jolyon Clulow, Céline Coma, Luca Compagna, Ricardo Corin, Veronique Cortier, Nora Cuppens-Boulahia, Marcin Czenko, Mads Dam, Marnix Dekker, Markus Dürmuth, Ulrich Flegel, Jun Furukawa, David Galindo, Han Gao, Flavio D. Garcia, Joaquin Garcia, Meng Ge, Pablo Giambiagi, Hidehito Gomi, Maria Isabel González Vasco, Thomas Gross, Toshiyuki Isshiki, Ana Jancic, Thomas Jensen, Florian Kerschbaum, Naoki Kobayashi, Steve Kremer, Jan Krhovjak, Marek Kumpost, Sven Lachmund, Julien Laganier, Yassine Lakhnech, Peeter Laud, Corrado Leita, Anyi Liu, Vaclav Lorenc, Henning Makholm, Matteo Maffei, Michael Meier, Dale Miller, Kengo Mori, Antonio Muñoz Gallego, Steven Murdoch, Thanh Son Nguyen, Christoffer Rosenkilde Nielsen, Satoshi Obana, Yutaka Oiwa, Joseph Pamula, Maura Paterson, Jean-Christophe Pazzaglia, Thea Peacock, Van Hau Pham, François Pottier, Christian W. Probst, Tony Ramard, Sankardas Roy, Pierangela Samarati, Thierry Sans, Andreas Schaad, Ludwig Seitz, Fred Spiessens, Henning Sudbrock, Hongwei Sun, Petr Svenda, Isamu Teranishi, Tachio Terauchi, Olivier Thonnard, Terkel K. Tolstrup, Laurent Vigneron, Jan Vitek, Lingyu Wang, Stephen D. Wolthusen, Konrad Wrona, Hirosuke Yamamoto, Yanjiang Yang, Chao Yao, Stefano Zanero, Ye Zhang, Xibin Zhao, Hongbin Zhou, Anna Zych

Local Organization

Wiebke Frauen, Andreas Günter, Martin Johns (workshop chair), Klaus-Peter Kossakowski (sponsorship chair), Jan Meier, Joachim Stehmann (Web design), Margit Wichmann

Lecture Notes in Computer Science

For information about Vols. 1–4091

please contact your bookseller or Springer

- Vol. 4228: D.E. Lightfoot, C.A. Szyperski (Eds.), *Modular Programming Languages*. X, 415 pages. 2006.
- Vol. 4219: D. Zamboni, C. Kruegel (Eds.), *Recent Advances in Intrusion Detection*. XII, 323 pages. 2006.
- Vol. 4208: M. Gerndt, D. Kranzlmüller (Eds.), *High Performance Computing and Communications*. XXII, 938 pages. 2006.
- Vol. 4206: P. Dourish, A. Friday (Eds.), *UbiComp 2006: Ubiquitous Computing*. XIX, 526 pages. 2006.
- Vol. 4193: T.P. Runarsson, H.-G. Beyer, E. Burke, J.J. Merelo-Guervós, L. D. Whitley, X. Yao (Eds.), *Parallel Problem Solving from Nature - PPSN IX*. XIX, 1061 pages. 2006.
- Vol. 4192: B. Mohr, J.L. Träff, J. Worringer, J. Dongarra (Eds.), *Recent Advances in Parallel Virtual Machine and Message Passing Interface*. XVI, 414 pages. 2006.
- Vol. 4189: D. Gollmann, J. Meier, A. Sabelfeld (Eds.), *Computer Security - ESORICS 2006*. XI, 548 pages. 2006.
- Vol. 4188: P. Sojka, I. Kopeček, K. Pala (Eds.), *Text, Speech and Dialogue*. XIV, 721 pages. 2006. (Sublibrary LNAI).
- Vol. 4187: J.J. Alferes, J. Bailey, W. May, U. Schwertel (Eds.), *Principles and Practice of Semantic Web Reasoning*. XI, 277 pages. 2006.
- Vol. 4186: C. Jesshope, C. Egan (Eds.), *Advances in Computer Systems Architecture*. XIV, 605 pages. 2006.
- Vol. 4185: R. Mizoguchi, Z. Shi, F. Giunchiglia (Eds.), *The Semantic Web - ASWC 2006*. XX, 778 pages. 2006.
- Vol. 4184: M. Bravetti, M. Núñez, G. Zavattaro (Eds.), *Web Services and Formal Methods*. X, 289 pages. 2006.
- Vol. 4183: J. Euzenat, J. Domingue (Eds.), *Artificial Intelligence: Methodology, Systems, and Applications*. XIII, 291 pages. 2006. (Sublibrary LNAI).
- Vol. 4180: M. Kohlhase, *OMDoc - An Open Markup Format for Mathematical Documents [version 1.2]*. XIX, 428 pages. 2006. (Sublibrary LNAI).
- Vol. 4178: A. Corradini, H. Ehrig, U. Montanari, L. Ribeiro, G. Rozenberg (Eds.), *Graph Transformations*. XII, 473 pages. 2006.
- Vol. 4176: S.K. Katsikas, J. Lopez, M. Backes, S. Gritzalis, B. Preneel (Eds.), *Information Security*. XIV, 548 pages. 2006.
- Vol. 4175: P. Bücher, B.M.E. Moret (Eds.), *Algorithms in Bioinformatics*. XII, 402 pages. 2006. (Sublibrary LNBI).
- Vol. 4174: K. Franke, K.-R. Müller, B. Nickolay, R. Schäfer (Eds.), *Pattern Recognition*. XX, 773 pages. 2006.
- Vol. 4172: J. Gonzalo, C. Thanos, M. F. Verdejo, R.C. Carrasco (Eds.), *Research and Advanced Technology for Digital Libraries*. XVII, 569 pages. 2006.
- Vol. 4169: H.L. Bodlaender, M.A. Langston (Eds.), *Parameterized and Exact Computation*. XI, 279 pages. 2006.
- Vol. 4168: Y. Azar, T. Erlebach (Eds.), *Algorithms - ESA 2006*. XVIII, 843 pages. 2006.
- Vol. 4165: W. Jonker, M. Petković (Eds.), *Secure, Data Management*. X, 185 pages. 2006.
- Vol. 4163: H. Bersini, J. Carneiro (Eds.), *Artificial Immune Systems*. XII, 460 pages. 2006.
- Vol. 4162: R. Kráľovič, P. Urzyczyn (Eds.), *Mathematical Foundations of Computer Science 2006*. XV, 814 pages. 2006.
- Vol. 4160: M. Fisher, W.v.d. Hoek, B. Konev, A. Lisitsa (Eds.), *Logics in Artificial Intelligence*. XII, 516 pages. 2006. (Sublibrary LNAI).
- Vol. 4159: J. Ma, H. Jin, L.T. Yang, J.J.-P. Tsai (Eds.), *Ubiquitous Intelligence and Computing*. XXII, 1190 pages. 2006.
- Vol. 4158: L.T. Yang, H. Jin, J. Ma, T. Ungerer (Eds.), *Autonomic and Trusted Computing*. XIV, 613 pages. 2006.
- Vol. 4156: S. Amer-Yahia, Z. Bellahsene, E. Hunt, R. Unland, J.X. Yu (Eds.), *Database and XML Technologies*. IX, 123 pages. 2006.
- Vol. 4155: O. Stock, M. Schaerf (Eds.), *Reasoning, Action and Interaction in AI Theories and Systems*. XVIII, 343 pages. 2006. (Sublibrary LNAI).
- Vol. 4154: Y.A. Dimitriadis, I. Zigurs, E. Gómez-Sánchez (Eds.), *Groupware: Design, Implementation, and Use*. XIV, 438 pages. 2006.
- Vol. 4153: N. Zheng, X. Jiang, X. Lan (Eds.), *Advances in Machine Vision, Image Processing, and Pattern Analysis*. XIII, 506 pages. 2006.
- Vol. 4152: Y. Manolopoulos, J. Pokorný, T. Sellis (Eds.), *Advances in Databases and Information Systems*. XV, 448 pages. 2006.
- Vol. 4151: A. Iglesias, N. Takayama (Eds.), *Mathematical Software - ICMIS 2006*. XVII, 452 pages. 2006.
- Vol. 4150: M. Dorigo, L.M. Gambardella, M. Birattari, A. Martinoli, R. Poli, T. Stützle (Eds.), *Ant Colony Optimization and Swarm Intelligence*. XVI, 526 pages. 2006.
- Vol. 4149: M. Klusch, M. Rovatsos, T.R. Payne (Eds.), *Cooperative Information Agents X*. XII, 477 pages. 2006. (Sublibrary LNAI).
- Vol. 4148: J. Vounckx, N. Azemard, P. Maurine (Eds.), *Integrated Circuit and System Design*. XVI, 677 pages. 2006.

- Vol. 4146: J.C. Rajapakse, L. Wong, R. Acharya (Eds.), *Pattern Recognition in Bioinformatics*. XIV, 186 pages. 2006. (Sublibrary LNBI).
- Vol. 4144: T. Ball, R.B. Jones (Eds.), *Computer Aided Verification*. XV, 564 pages. 2006.
- Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pahikkala, *Advances in Natural Language Processing*. XVI, 771 pages. 2006. (Sublibrary LNAI).
- Vol. 4138: X. Cheng, W. Li, T. Znati (Eds.), *Wireless Algorithms, Systems, and Applications*. XVI, 709 pages. 2006.
- Vol. 4137: C. Baier, H. Hermanns (Eds.), *CONCUR 2006 – Concurrency Theory*. XIII, 525 pages. 2006.
- Vol. 4136: R.A. Schmidt (Ed.), *Relations and Kleene Algebra in Computer Science*. XI, 433 pages. 2006.
- Vol. 4135: C.S. Calude, M.J. Dinneen, G. Păun, G. Rozenberg, S. Stepney (Eds.), *Unconventional Computation*. X, 267 pages. 2006.
- Vol. 4134: K. Yi (Ed.), *Static Analysis*. XIII, 443 pages. 2006.
- Vol. 4133: J. Gratch, M. Young, R. Aylett, D. Ballin, P. Olivier (Eds.), *Intelligent Virtual Agents*. XIV, 472 pages. 2006. (Sublibrary LNAI).
- Vol. 4132: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006*, Part II. XXXIV, 1028 pages. 2006.
- Vol. 4131: S. Kollias, A. Stafylopatis, W. Duch, E. Oja (Eds.), *Artificial Neural Networks – ICANN 2006*, Part I. XXXIV, 1008 pages. 2006.
- Vol. 4130: U. Furbach, N. Shankar (Eds.), *Automated Reasoning*. XV, 680 pages. 2006. (Sublibrary LNAI).
- Vol. 4129: D. McGookin, S. Brewster (Eds.), *Haptic and Audio Interaction Design*. XII, 167 pages. 2006.
- Vol. 4128: W.E. Nagel, W.V. Walter, W. Lehner (Eds.), *Euro-Par 2006 Parallel Processing*. XXXIII, 1221 pages. 2006.
- Vol. 4127: E. Damiani, P. Liu (Eds.), *Data and Applications Security*. XX, X, 319 pages. 2006.
- Vol. 4126: P. Barahona, F. Bry, E. Franconi, N. Henze, U. Sattler, *Reasoning Web*. XII, 269 pages. 2006.
- Vol. 4124: H. de Meer, J.P. G. Sterbenz (Eds.), *Self-Organizing Systems*. XIV, 261 pages. 2006.
- Vol. 4121: A. Biere, C.P. Gomes (Eds.), *Theory and Applications of Satisfiability Testing – SAT 2006*. XII, 438 pages. 2006.
- Vol. 4120: J. Calmet, T. Ida, D. Wang (Eds.), *Artificial Intelligence and Symbolic Computation*. XIII, 269 pages. 2006. (Sublibrary LNAI).
- Vol. 4119: C. Dony, J.L. Knudsen, A. Romanovsky, A. Tripathi (Eds.), *Advanced Topics in Exception Handling Components*. X, 302 pages. 2006.
- Vol. 4117: C. Dwork (Ed.), *Advances in Cryptology – CRYPTO 2006*. XIII, 621 pages. 2006.
- Vol. 4116: R. De Prisco, M. Yung (Eds.), *Security and Cryptography for Networks*. XI, 366 pages. 2006.
- Vol. 4115: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence and Bioinformatics*, Part III. XXI, 803 pages. 2006. (Sublibrary LNBI).
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Computational Intelligence*, Part II. XXVII, 1337 pages. 2006. (Sublibrary LNAI).
- Vol. 4113: D.-S. Huang, K. Li, G.W. Irwin (Eds.), *Intelligent Computing*, Part I. XXVII, 1331 pages. 2006.
- Vol. 4112: D.Z. Chen, D. T. Lee (Eds.), *Computing and Combinatorics*. XIV, 528 pages. 2006.
- Vol. 4111: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roeer (Eds.), *Formal Methods for Components and Objects*. VIII, 447 pages. 2006.
- Vol. 4110: J. Díaz, K. Jansen, J.D.P. Rolim, U. Zwick (Eds.), *Approximation, Randomization, and Combinatorial Optimization*. XII, 522 pages. 2006.
- Vol. 4109: D.-Y. Yeung, J.T. Kwok, A. Fred, F. Roli, D. de Ridder (Eds.), *Structural, Syntactic, and Statistical Pattern Recognition*. XXI, 939 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), *Mathematical Knowledge Management*. VIII, 295 pages. 2006. (Sublibrary LNAI).
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H. A. Güvenir (Eds.), *Advances in Case-Based Reasoning*. XIV, 566 pages. 2006. (Sublibrary LNAI).
- Vol. 4105: B. Gunsels, A.K. Jain, A. M. Tekalp, B. Sankur (Eds.), *Multimedia, Content Representation, Classification and Security*. XIX, 804 pages. 2006.
- Vol. 4104: T. Kunz, S.S. Ravi (Eds.), *Ad-Hoc, Mobile, and Wireless Networks*. XII, 474 pages. 2006.
- Vol. 4103: J. Eder, S. Dustdar (Eds.), *Business Process Management Workshops*. XI, 508 pages. 2006.
- Vol. 4102: S. Dustdar, J.L. Fiadeiro, A. Sheth (Eds.), *Business Process Management*. XV, 486 pages. 2006.
- Vol. 4101: Y. Luo (Ed.), *Cooperative Design, Visualization, and Engineering*. X, 338 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), *PRICAI 2006: Trends in Artificial Intelligence*. XXVIII, 1263 pages. 2006. (Sublibrary LNAI).
- Vol. 4098: F. Pfenning (Ed.), *Term Rewriting and Applications*. XIII, 415 pages. 2006.
- Vol. 4097: X. Zhou, O. Sokolsky, L. Yan, E.-S. Jung, Z. Shao, Y. Mu, D.C. Lee, D. Kim, Y.-S. Jeong, C.-Z. Xu (Eds.), *Emerging Directions in Embedded and Ubiquitous Computing*. XXVII, 1034 pages. 2006.
- Vol. 4096: E. Sha, S.-K. Han, C.-Z. Xu, M.H. Kim, L.T. Yang, B. Xiao (Eds.), *Embedded and Ubiquitous Computing*. XXIV, 1170 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassarre, R. Calabretta, J.C. T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, D. Parisi (Eds.), *From Animals to Animats 9*. XV, 869 pages. 2006. (Sublibrary LNAI).
- Vol. 4094: O. H. Ibarra, H.-C. Yen (Eds.), *Implementation and Application of Automata*. XIII, 291 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaiane, Z. Li (Eds.), *Advanced Data Mining and Applications*. XXI, 1110 pages. 2006. (Sublibrary LNAI).
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), *Knowledge Science, Engineering and Management*. XV, 664 pages. 2006. (Sublibrary LNAI).

Table of Contents

Finding Peer-to-Peer File-Sharing Using Coarse Network Behaviors	1
<i>M.P. Collins, M.K. Reiter</i>	
Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses	18
<i>V. Shmatikov, M.-H. Wang</i>	
TrustedPals: Secure Multiparty Computation Implemented with Smart Cards	34
<i>M. Fort, F. Freiling, L. Draque Penso, Z. Benenson, D. Kesdogan</i>	
Private Information Retrieval Using Trusted Hardware	49
<i>S.H. Wang, X.H. Ding, R.H. Deng, F. Bao</i>	
Bridging the Gap Between Inter-communication Boundary and Internal Trusted Components	65
<i>Y. Watanabe, S. Yoshihama, T. Mishina, M. Kudo, H. Maruyama</i>	
License Transfer in OMA-DRM	81
<i>C.N. Chong, S. Iacob, P. Koster, J. Montaner, R. van Buuren</i>	
Enhanced Security Architecture for Music Distribution on Mobile	97
<i>A. Benjelloun-Touimi, J.-B. Fischer, C. Fontaine, C. Giraud, M. Milhau</i>	
A Formal Model of Access Control for Mobile Interactive Devices	110
<i>F. Besson, G. Dufay, T. Jensen</i>	
Discretionary Capability Confinement	127
<i>P.W.L. Fong</i>	
Minimal Threshold Closure	145
<i>X.-B. Zhao, K.-Y. Lam, G. Luo, S.-L. Chung, M. Gu</i>	
Reducing the Dependence of SPKI/SDSI on PKI	156
<i>H. Wang, S. Jha, T. Reps, S. Schwoon, S. Stubblebine</i>	
Delegation in Role-Based Access Control	174
<i>J. Crampton, H. Khambhammettu</i>	

Applying a Security Requirements Engineering Process	192
<i>D. Mellado, E. Fernández-Medina, M. Piattini</i>	
Modeling and Evaluating the Survivability of an Intrusion Tolerant Database System	207
<i>H. Wang, P. Liu</i>	
A Formal Framework for Confidentiality-Preserving Refinement	225
<i>T. Santen</i>	
Timing-Sensitive Information Flow Analysis for Synchronous Systems	243
<i>B. Köpf, D. Basin</i>	
HBAC: A Model for History-Based Access Control and Its Model Checking	263
<i>J. Wang, Y. Takata, H. Seki</i>	
From Coupling Relations to Mated Invariants for Checking Information Flow	279
<i>D.A. Naumann</i>	
A Linear Logic of Authorization and Knowledge	297
<i>D. Garg, L. Bauer, K.D. Bowers, F. Pfennig, M.K. Reiter</i>	
Prêt à Voter with Re-encryption Mixes	313
<i>P.Y.A. Ryan, S.A. Schneider</i>	
Secure Key-Updating for Lazy Revocation	327
<i>M. Backes, C. Cachin, A. Oprea</i>	
Key Derivation Algorithms for Monotone Access Structures in Cryptographic File Systems	347
<i>M. Srivatsa, L. Liu</i>	
Cryptographically Sound Security Proofs for Basic and Public-Key Kerberos	362
<i>M. Backes, I. Cervesato, A.D. Jaggard, A. Scedrov, J.-K. Tsay</i>	
Deriving Secrecy in Key Establishment Protocols	384
<i>D. Pavlovic, C. Meadows</i>	

Limits of the BRSIM/UC Soundness of Dolev-Yao Models with Hashes	404
<i>M. Backes, B. Pfitzmann, M. Waidner</i>	
Conditional Reactive Simulatability	424
<i>M. Backes, M. Dürmuth, D. Hofheinz, R. Küsters</i>	
SessionSafe: Implementing XSS Immune Session Handling	444
<i>M. Johns</i>	
Policy-Driven Memory Protection for Reconfigurable Hardware	461
<i>T. Huffmire, S. Prasad, T. Sherwood, R. Kastner</i>	
Privacy-Preserving Queries on Encrypted Data	479
<i>Z. Yang, S. Zhong, R.N. Wright</i>	
Analysis of Policy Anomalies on Distributed Network Security Setups	496
<i>J.G. Alfaro, F. Cuppens, N. Cuppens-Boulahia</i>	
Assessment of a Vulnerability in Iterative Servers Enabling Low-Rate DoS Attacks	512
<i>G. Maciá-Fernández, J.E. Díaz-Verdejo, P. García-Teodoro</i>	
Towards an Information-Theoretic Framework for Analyzing Intrusion Detection Systems	527
<i>G. Gu, P. Fogla, D. Dagon, W. Lee, B. Skoric</i>	
Author Index	547

Finding Peer-to-Peer File-Sharing Using Coarse Network Behaviors*

Michael P. Collins¹ and Michael K. Reiter²

¹ CERT/Network Situational Awareness, Software Engineering Institute,
Carnegie Mellon University

`mcollins@cert.org`

² Electrical & Computer Engineering Department, Computer Science Department,
and CyLab, Carnegie Mellon University

`reiter@cmu.edu`

Abstract. A user who wants to use a service forbidden by their site's usage policy can masquerade their packets in order to evade detection. One masquerade technique sends prohibited traffic on TCP ports commonly used by permitted services, such as port 80. Users who hide their traffic in this way pose a special challenge, since filtering by port number risks interfering with legitimate services using the same port. We propose a set of tests for identifying masqueraded peer-to-peer file-sharing based on traffic summaries (flows). Our approach is based on the hypothesis that these applications have observable behavior that can be differentiated without relying on deep packet examination. We develop tests for these behaviors that, when combined, provide an accurate method for identifying these masqueraded services without relying on payload or port number. We test this approach by demonstrating that our integrated detection mechanism can identify BitTorrent with a 72% true positive rate and virtually no observed false positives in control services (FTP-Data, HTTP, SMTP).

1 Introduction

Peer-to-peer file-sharing services are often constrained by organizations due to their widespread use for disseminating copyrighted content illegally, their significant bandwidth consumption for (typically) non-work-related uses, and/or the risk that they may introduce new security vulnerabilities to the organization. Karagiannis et al. [5] have shown that instead of obeying site bans on file-sharing, however, users hide their file-sharing activity. Moreover, file-sharing tools themselves are being updated to circumvent attempts to filter these services; e.g., BitTorrent developers now incorporate encryption into their products in order to evade traffic shaping.¹

* This work was partially supported by NSF award CNS-0433540, and by KISA and MIC of Korea.

¹ "Encrypting BitTorrent to Take Out Traffic Shapers", TorrentFreak Weblog, <http://torrentfreak.com/encrypting-BitTorrent-to-take-out-traffic-shapers/>

While encryption makes filtering based on traffic content difficult, filtering packets by port number (as would typically be implemented in router ACLs, for example) remains an obstacle to peer-to-peer file-sharing. As such, common hiding methods also involve changing the port number used by the service to something that is not filtered. In networks that implement a “deny-than-allow” policy, the service traffic may be sent on a common service port, in particular 80/TCP (HTTP).

In such cases, ports do not reliably convey the services using them, while deep packet examination is viable only as long as packet payload is unencrypted. Analysts therefore need alternative methods to characterize and filter traffic. In this paper, we propose an alternative service detection and identification method that characterizes services behaviorally. We hypothesize that TCP services have quantifiable behaviors that can be used to identify them without relying on payload or port numbers. For example, we expect that the majority of HTTP sessions begin with a small initial request followed by a larger response, and then terminate. If a presumed HTTP client and HTTP server were communicating using symmetric short bursts of traffic in a single long-lived session, then we would have reason to consider an alternative hypothesis, such as a chat service.

Within this paper, we focus on a specific problem that motivated this line of research: demonstrating that a user who claims to be using a common service on its standard port (such as HTTP) is using another service, specifically BitTorrent. To do so, we implement tests that characterize traffic and show how they can be used together to effectively differentiate BitTorrent traffic from common services. The goal of our research is a collection of tests which can be used by analysts or automated systems to classify traffic. Given the increasing sophistication of evasion strategies, we seek to find behaviors that can be effective with as few assumptions as possible. For example, these tests do not use deep packet examination, and are therefore applicable to encrypted and unencrypted traffic.

We calibrate and validate our approach using logs of traffic crossing a large network. From these logs, we select traffic records describing BitTorrent and major services, specifically HTTP, FTP data channel and SMTP. The log data consists of NetFlow, a traffic summarization standard developed by CISCO systems². Flow data is a compact representation of an approximate TCP session, but does not contain payload. In addition, we do not trust port numbers, making our tests port- and payload-agnostic. Despite this, we show that by classifying flows based on several behaviors we can effectively differentiate source-destination pairs engaged in BitTorrent communication from those involved in HTTP, FTP or SMTP exchanges. Specifically, our integrated test identifies BitTorrent with a 72% true positive rate and virtually no observed false positives in control services (FTP-Data, HTTP, SMTP).

The rest of this paper is structured as follows. Section 2 describes previous work done in service detection. Section 3 describes the behaviors with which we characterize flows, and that we use to distinguish certain file-sharing traffic from

² CISCO Corporation, *Netflow Services and Applications*, <http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/nappswp.htm>

other types of traffic. Section 4 describes our experiments using these classifications both individually and in aggregate, to identify BitTorrent activity. We conclude in Section 5.

2 Previous Work

Prior work in identifying file-sharing traffic varies primarily in the information used to do so. Several file-sharing detection tools have analyzed packet payload (e.g., [9,20]), a method which will not survive encryption of packet contents or might simply be infeasible due to performance or other limitations. Other approaches utilize aggregate packet attributes, such as interstitial arrival times or the presence of specific packet sequences (e.g., [22,8,2,24,3,10,11,12]). However, in sufficiently large and busy networks, even this degree of packet analysis can be problematic.

As a result, flows are increasingly used for various types of security analysis (e.g., [18,13]). Flows were originally specified by Partridge [15] for traffic summarization, and have since been adopted by CISCO for traffic reporting. NetFlow, the CISCO standard, uses timeouts to approximate TCP sessions, an approach originally developed by Claffy et al. [1]. Since flow records do not contain payload information, they are generally used for large-scale and statistical analysis. Notably, Soule et al. [21] developed a classification method to cluster flows, though they stopped short of mapping them to existing applications (or types of applications).

Since their development, peer-to-peer file-sharing systems have become targets of filtering and detection efforts. Karagiannis et al. [5] showed that peer-to-peer users increasingly evade detection by moving their traffic to alternate port numbers. Studies conducted on BitTorrent and other peer-to-peer file-sharing applications have examined the behavior of individual nodes (e.g., [4,23,7]) and application networks (e.g., [19,17]), but have not compared the behaviors observed to the behavior of more traditional services. Ohzahata et al. [14] developed a method for detecting hosts participating in the Winny file-sharing application by inserting monitoring hosts within the file-sharing network itself. Karagiannis et al. [6] developed a general method for identifying applications called Blinc, which uses various heuristics and interconnection patterns exhibited by groups of nodes to identify services. In contrast, we focus on the flow characteristics between a pair of nodes in isolation to identify the service in which they are participating, and as such our approach is complementary. Nevertheless, we believe there is potential in combining our point-to-point analyses with Blinc's on interconnection patterns, and hope to investigate this in future work.

3 Application Classification

In this section, we describe the behaviors used to differentiate BitTorrent traffic from other services. In Section 3.1 we describe a classification tree that we will

use to classify flows into different types, and in Section 3.2 we describe the intuition and formulation of our tests.

3.1 Simple Taxonomy

Our analyses use flow records; a *flow* is a sequence of packets with the same addressing information (source and destination addresses, source and destination ports, and protocol) which occur within a short time of each other [1]. A *flow record* is a summary consisting of addressing, size and timing information about the flow, but no payload. We will refer to fields of a flow record f with “dot” notation (e.g., $f.duration$ or $f.bytes$).

We restrict our data to TCP flows. Flow collection systems such as CISCO NetFlow record TCP flags by ORing the flags of every packet. As a result, flag distributions cannot be derived from multi-packet flow records, and certain behaviors—notably whether an endpoint is the initiator or responder of the TCP connection of which the flow represents one direction—are not discernible.

We divide flows into three categories: **Short Flows**, comprising three packets or less; **Messages**, which are 4–10 packets but less than 2 kB in size; and **File Transfers**, which are any flows longer than a Message. Figure 1 represents our taxonomy as a decision tree and the categories that this tree covers.

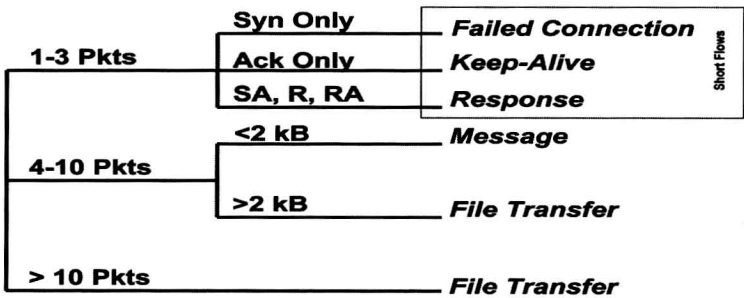


Fig. 1. Flow Classification Tree: The rules used on this tree assign each flow to a class

A **Short Flow** consists of three or fewer packets; since a complete TCP session will require at least three packets, Short Flows indicate some error in communication or anomaly in flow collection. Within Short Flows, we can acquire more information by examining the TCP flags of the flow; we use the flags to create three sub categories. A **Failed Connection** has a SYN flag and no ACKs. A **Keep-Alive** has ACK flags only. Since flow recording is timeout-based, Keep-Alives are recorded by the flow collector during long-lived sessions but currently do not have any substantial impact on analysis. A **Response** consists of any Short Flow whose flags imply a response from the TCP connection responder to the initiator: a SYN-ACK, a RST-ACK or a RST. We do not consider the other

TCP flags (e.g., PSH) significant in this analysis. As noted above, TCP flags are OR'ed in flow records, and as a result we only use flags in the Short Flow case, where the results are least ambiguous.

We define a **Message** as a flow consisting of 4–10 packets and with a total size less than 2 kB. We assume that Messages represent the structured exchange of service data between the source and destination. Example Messages include HTTP requests and the control messages sent by BitTorrent. We assume that Messages contain structured communication, as opposed to data intended for the application's users. Consequently, we expect that Messages will have fixed sizes and that certain Messages (with specific sizes) will appear more often than other ones.

We label any flow longer than 2 kB or 10 packets a **File Transfer**. We assume that a File Transfer is the exchange of non-service-specific information between two sites. We expect that certain services will tend to send shorter files than others. For example, we expect that HTTP servers will transfer less data than BitTorrent peers typically, since HTTP clients interact with users and therefore need a rapid response time.

Table 1 is a log showing BitTorrent flows; in this log, we have labeled each flow with its corresponding category. Of particular interest is the presence of repeated Failed Connections (the 144-byte SYN-only packets) and the 276-byte Message packets. Both of these behaviors will be used to construct tests in Section 3.2.

Table 1. Log of traffic and associated classification

Source Port	Destination Port	Packets	Bytes	Flags				Start Time	Class
				F	S	A	R		
3584	6881	1637	1270926	x	x			11/04/2005 21:09:33	File Transfer
3586	6881	5	276	x	x	x		11/04/2005 21:09:36	Message
3619	6881	5	276	x	x	x		11/04/2005 21:10:18	Message
3651	6881	5	276	x	x	x		11/04/2005 21:11:01	Message
3701	6881	5	276	x	x	x		11/04/2005 21:12:04	Message
1290	6881	3	144	x				11/04/2005 21:53:56	Failed Connection
2856	6881	5	636	x	x			11/04/2005 22:33:11	Message
3916	6881	5	276	x	x	x		11/04/2005 23:03:44	Message
4178	6881	5	636	x	x			11/04/2005 23:12:01	Message
4884	6881	3	144	x				11/04/2005 23:32:05	Failed Connection

3.2 Tests

In this section, we describe four tests for characterizing the flows generated by various services. Each test is performed on a log of flow records bearing the same source and destination, and hence are unidirectional. We rely on unidirectional flows for three reasons. First, CISCO NetFlow is reported unidirectionally; i.e., each direction of a connection is reported in a different flow. Second, on a network with multiple access points, there is no guarantee that entry and exit traffic