Arosha K. Bandara
Mark Burgess (Eds.)

LNCS 4543

Inter-Domain Management

First International Conference on Autonomous Infrastructure, Management and Security, AIMS 2007
Oslo, Norway, June 2007, Proceedings

ASSOCIATION FOR
COMPUTING MACHINERY

Arosha K. Bandara Mark Burgess (Eds.)

Inter-Domain Management

First International Conference on Autonomous Infrastructure, Management and Security, AIMS 2007
Oslo, Norway, June 21-22, 2007
Proceedings

江苏工业学院图书馆
藏书章

Springer

Volume Editors

Arosha K. Bandara
The Open University
Wlton Hall Campus, Milton Keynes, MK 76AA, UK
E-mail: a.k.bandara@open.ac.uk

Mark Burgess
Oslo University College
PO Box 4, St Olavs Plass, 0130 Oslo Norway
E-mail: mark@iu.hio.no

Library of Congress Control Number: 2007928329

CR Subject Classification (1998): C.2, D.4.4, H.3, H.4

LNCS Sublibrary: SL 5 – Computer Communication Networks and Telecommunications

ISSN 0302-9743
ISBN-10 3-540-72985-2 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-72985-3 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 2007
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12074792 06/3180 5 4 3 2 1 0

Lecture Notes in Computer Science 4543

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

Lecture Notes in Computer Science

For information about Vols. 1–4438

please contact your bookseller or Springer

Vol. 4543: A.K. Bandara, M. Burgess (Eds.), Inter-Domain Management. XII, 237 pages. 2007.

Vol. 4542: P. Sawyer, B. Paech, P. Heymans (Eds.), Requirements Engineering: Foundation for Software Quality. IX, 384 pages. 2007.

Vol. 4539: N.H. Bshouty, C. Gentile (Eds.), Learning Theory. XII, 634 pages. 2007. (Sublibrary LNAI).

Vol. 4538: F. Escolano, M. Vento (Eds.), Graph-Based Representations in Pattern Recognition. XII, 416 pages. 2007.

Vol. 4537: K.C.-C. Chang, W. Wang, L. Chen, C.A. Ellis, C.-H. Hsu, A.C. Tsoi, H. Wang (Eds.), Advances in Web and Network Technologies, and Information Management. XXIII, 707 pages. 2007.

Vol. 4534: I. Tomkos, F. Neri, J. Solé Pareta, X. Masip Bruin, S. Sánchez Lopez (Eds.), Optical Network Design and Modeling. XI, 460 pages. 2007.

Vol. 4531: J. Indulska, K. Raymond (Eds.), Distributed Applications and Interoperable Systems. XI, 337 pages. 2007.

Vol. 4530: D.H. Akehurst, R. Vogel, R.F. Paige (Eds.), Model Driven Architecture- Foundations and Applications. X, 219 pages. 2007.

Vol. 4529: P. Melin, O. Castillo, L.T. Aguilar, J. Kacprzyk, W. Pedrycz (Eds.), Foundations of Fuzzy Logic and Soft Computing. XIX, 830 pages. 2007. (Sublibrary LNAI).

Vol. 4527: J. Mira, J.R. Álvarez (Eds.), Bio-inspired Modeling of Cognitive Tasks, Part I. XXII, 630 pages. 2007.

Vol. 4526: M. Malek, M. Reitenspieß, A. van Moorsel (Eds.), Service Availability. X, 155 pages. 2007.

Vol. 4525: C. Demetrescu (Ed.), Experimental Algorithms. XIII, 448 pages. 2007.

Vol. 4524: M. Marchiori, J.Z. Pan, C.d.S. Marie (Eds.), Web Reasoning and Rule Systems. XI, 382 pages. 2007.

Vol. 4523: Y.-H. Lee, H.-N. Kim, J. Kim, Y. Park, L.T. Yang, S.W. Kim (Eds.), Embedded Software and Systems. XIX, 829 pages. 2007.

Vol. 4522: B.K. Ersbøll, K.S. Pedersen (Eds.), Image Analysis. XVIII, 989 pages. 2007.

Vol. 4521: J. Katz, M. Yung (Eds.), Applied Cryptography and Network Security. XIII, 498 pages. 2007.

Vol. 4519: E. Franconi, M. Kifer, W. May (Eds.), The Semantic Web: Research and Applications. XVIII, 830 pages. 2007.

Vol. 4517: F. Boavida, E. Monteiro, S. Mascolo, Y. Koucheryavy (Eds.), Wired/Wireless Internet Communications. XIV, 382 pages. 2007.

Vol. 4516: L. Mason, T. Drwiega, J. Yan (Eds.), Traffic Performance in Converged Networks. XXIII, 1191 pages. 2007.

Vol. 4515: M. Naor (Ed.), Advances in Cryptology - EUROCRYPT 2007. XIII, 591 pages. 2007.

Vol. 4514: S.N. Artemov, A. Nerode (Eds.), Logical Foundations of Computer Science. XI, 513 pages. 2007.

Vol. 4513: M. Fischetti, D.P. Williamson (Eds.), Integer Programming and Combinatorial Optimization. IX, 500 pages. 2007.

Vol. 4510: P. Van Hentenryck, L. Wolsey (Eds.), Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems. X, 391 pages. 2007.

Vol. 4509: Z. Kobti, D. Wu (Eds.), Advances in Artificial Intelligence. XII, 552 pages. 2007. (Sublibrary LNAI).

Vol. 4508: M.-Y. Kao, X.-Y. Li (Eds.), Algorithmic Aspects in Information and Management. VIII, 428 pages. 2007.

Vol. 4506: D. Zeng, I. Gotham, K. Komatsu, C. Lynch, M. Thurmond, D. Madigan, B. Lober, J. Kvach, H. Chen (Eds.), Intelligence and Security Informatics: Biosurveillance. XI, 234 pages. 2007.

Vol. 4505: G. Dong, X. Lin, W. Wang, Y. Yang, J.X. Yu (Eds.), Advances in Data and Web Management. XXII, 896 pages. 2007.

Vol. 4504: J. Huang, R. Kowalczyk, Z. Maamar, D. Martin, I. Müller, S. Stoutenburg, K.P. Sycara (Eds.), Service-Oriented Computing: Agents, Semantics, and Engineering. X, 175 pages. 2007.

Vol. 4501: J. Marques-Silva, K.A. Sakallah (Eds.), Theory and Applications of Satisfiability Testing – SAT 2007. XI, 384 pages. 2007.

Vol. 4500: N. Streitz, A. Kameas, I. Mavrommati (Eds.), The Disappearing Computer. XVIII, 304 pages. 2007.

Vol. 4496: N.T. Nguyen, A. Grzech, R.J. Howlett, L.C. Jain (Eds.), Agent and Multi-Agent Systems: Technologies and Applications. XXI, 1046 pages. 2007. (Sublibrary LNAI).

Vol. 4495: J. Krogstie, A. Opdahl, G. Sindre (Eds.), Advanced Information Systems Engineering. XVI, 606 pages. 2007.

Vol. 4494: H. Jin, O.F. Rana, Y. Pan, V.K. Prasanna (Eds.), Algorithms and Architectures for Parallel Processing. XIV, 508 pages. 2007.

Vol. 4493: D. Liu, S. Fei, Z. Hou, H. Zhang, C. Sun (Eds.), Advances in Neural Networks – ISNN 2007, Part III. XXVI, 1215 pages. 2007.

此为试读，需要完整PDF请访问：www.ertongbook.com

Vol. 4492: D. Liu, S. Fei, Z. Hou, H. Zhang, C. Sun (Eds.), Advances in Neural Networks – ISNN 2007, Part II. XXVII, 1321 pages. 2007.

Vol. 4491: D. Liu, S. Fei, Z.-G. Hou, H. Zhang, C. Sun (Eds.), Advances in Neural Networks – ISNN 2007, Part I. LIV, 1365 pages. 2007.

Vol. 4490: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), Computational Science – ICCS 2007, Part IV. XXXVII, 1211 pages. 2007.

Vol. 4489: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), Computational Science – ICCS 2007, Part III. XXXVII, 1257 pages. 2007.

Vol. 4488: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), Computational Science – ICCS 2007, Part II. XXXV, 1251 pages. 2007.

Vol. 4487: Y. Shi, G.D. van Albada, J. Dongarra, P.M.A. Sloot (Eds.), Computational Science – ICCS 2007, Part I. LXXXI, 1275 pages. 2007.

Vol. 4486: M. Bernardo, J. Hillston (Eds.), Formal Methods for Performance Evaluation. VII, 469 pages. 2007.

Vol. 4485: F. Sgallari, A. Murli, N. Paragios (Eds.), Scale Space and Variational Methods in Computer Vision. XV, 931 pages. 2007.

Vol. 4484: J.-Y. Cai, S.B. Cooper, H. Zhu (Eds.), Theory and Applications of Models of Computation. XIII, 772 pages. 2007.

Vol. 4483: C. Baral, G. Brewka, J. Schlipf (Eds.), Logic Programming and Nonmonotonic Reasoning. IX, 327 pages. 2007. (Sublibrary LNAI).

Vol. 4482: A. An, J. Stefanowski, S. Ramanna, C.J. Butz, W. Pedrycz, G. Wang (Eds.), Rough Sets, Fuzzy Sets, Data Mining and Granular Computing. XIV, 585 pages. 2007. (Sublibrary LNAI).

Vol. 4481: J. Yao, P. Lingras, W.-Z. Wu, M. Szczuka, N.J. Cercone, D. Ślęzak (Eds.), Rough Sets and Knowledge Technology. XIV, 576 pages. 2007. (Sublibrary LNAI).

Vol. 4480: A. LaMarca, M. Langheinrich, K.N. Truong (Eds.), Pervasive Computing. XIII, 369 pages. 2007.

Vol. 4479: I.F. Akyildiz, R. Sivakumar, E. Ekici, J.C.d. Oliveira, J. McNair (Eds.), NETWORKING 2007. Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet. XXVII, 1252 pages. 2007.

Vol. 4478: J. Martí, J.M. Benedí, A.M. Mendonça, J. Serrat (Eds.), Pattern Recognition and Image Analysis, Part II. XXVII, 657 pages. 2007.

Vol. 4477: J. Martí, J.M. Benedí, A.M. Mendonça, J. Serrat (Eds.), Pattern Recognition and Image Analysis, Part I. XXVII, 625 pages. 2007.

Vol. 4476: V. Gorodetsky, C. Zhang, V.A. Skormin, L. Cao (Eds.), Autonomous Intelligent Systems: Multi-Agents and Data Mining. XIII, 323 pages. 2007. (Sublibrary LNAI).

Vol. 4475: P. Crescenzi, G. Prencipe, G. Pucci (Eds.), Fun with Algorithms. X, 273 pages. 2007.

Vol. 4474: G. Prencipe, S. Zaks (Eds.), Structural Information and Communication Complexity. XI, 342 pages. 2007.

Vol. 4472: M. Haindl, J. Kittler, F. Roli (Eds.), Multiple Classifier Systems. XI, 524 pages. 2007.

Vol. 4471: P. Cesar, K. Chorianopoulos, J.F. Jensen (Eds.), Interactive TV: a Shared Experience. XIII, 236 pages. 2007.

Vol. 4470: Q. Wang, D. Pfahl, D.M. Raffo (Eds.), Software Process Dynamics and Agility. XI, 346 pages. 2007.

Vol. 4468: M.M. Bonsangue, E.B. Johnsen (Eds.), Formal Methods for Open Object-Based Distributed Systems. X, 317 pages. 2007.

Vol. 4467: A.L. Murphy, J. Vitek (Eds.), Coordination Models and Languages. X, 325 pages. 2007.

Vol. 4466: F.B. Sachse, G. Seemann (Eds.), Functional Imaging and Modeling of the Heart. XV, 486 pages. 2007.

Vol. 4465: T. Chahed, B. Tuffin (Eds.), Network Control and Optimization. XIII, 305 pages. 2007.

Vol. 4464: E. Dawson, D.S. Wong (Eds.), Information Security Practice and Experience. XIII, 361 pages. 2007.

Vol. 4463: I. Măndoiu, A. Zelikovsky (Eds.), Bioinformatics Research and Applications. XV, 653 pages. 2007. (Sublibrary LNBI).

Vol. 4462: D. Sauveron, K. Markantonakis, A. Bilas, J.-J. Quisquater (Eds.), Information Security Theory and Practices. XII, 255 pages. 2007.

Vol. 4459: C. Cérin, K.-C. Li (Eds.), Advances in Grid and Pervasive Computing. XVI, 759 pages. 2007.

Vol. 4453: T. Speed, H. Huang (Eds.), Research in Computational Molecular Biology. XVI, 550 pages. 2007. (Sublibrary LNBI).

Vol. 4452: M. Fasli, O. Shehory (Eds.), Agent-Mediated Electronic Commerce. VIII, 249 pages. 2007. (Sublibrary LNAI).

Vol. 4451: T.S. Huang, A. Nijholt, M. Pantic, A. Pentland (Eds.), Artifical Intelligence for Human Computing. XVI, 359 pages. 2007. (Sublibrary LNAI).

Vol. 4450: T. Okamoto, X. Wang (Eds.), Public Key Cryptography – PKC 2007. XIII, 491 pages. 2007.

Vol. 4448: M. Giacobini et al. (Ed.), Applications of Evolutionary Computing. XXIII, 755 pages. 2007.

Vol. 4447: E. Marchiori, J.H. Moore, J.C. Rajapakse (Eds.), Evolutionary Computation,Machine Learning and Data Mining in Bioinformatics. XI, 302 pages. 2007.

Vol. 4446: C. Cotta, J. van Hemert (Eds.), Evolutionary Computation in Combinatorial Optimization. XII, 241 pages. 2007.

Vol. 4445: M. Ebner, M. O'Neill, A. Ekárt, L. Vanneschi, A.I. Esparcia-Alcázar (Eds.), Genetic Programming. XI, 382 pages. 2007.

Vol. 4444: T. Reps, M. Sagiv, J. Bauer (Eds.), Program Analysis and Compilation, Theory and Practice. X, 361 pages. 2007.

Vol. 4443: R. Kotagiri, P.R. Krishna, M. Mohania, E. Nantajeewarawat (Eds.), Advances in Databases: Concepts, Systems and Applications. XXI, 1126 pages. 2007.

Vol. 4440: B. Liblit, Cooperative Bug Isolation. XV, 101 pages. 2007.

Vol. 4439: W. Abramowicz (Ed.), Business Information Systems. XV, 654 pages. 2007.

Preface

Research needs ideas, discourse and experimentation in order to thrive, but more than ever we are expected to make research immediately 'relevant' and available to society and the world of commerce. Of these three poles (ideas, discourse and experimentation), ideas lie farthest from a finished product, and it is therefore ideas that are most easily left behind in the rush to catch the gravy train. The pressure to prioritize applications rather than understanding hinders researchers from thinking deeply about problems, and in the worst case prevents us from truly understanding and innovating.

The first Autonomous Infrastructure Management and Security conference (AIMS 2007) was proposed as an act of optimism by the leaders of the EMANICS Network of Excellence in Network and Service Management. It was a proposal aimed at avoiding the tar-pit of "apply existing knowledge only, " to reach out for new ideas that might expand our network of concepts and solutions.

There are already many excellent conferences in the field of Network of System Management : LISA, IM, NOMS, DSOM, Policy Workshop, etc. Although there is an overlap, both in attendance and ideas, AIMS does not compete with any of these. Rather we have sought a strong cross-disciplinary forum, in which novelty and discussion are made paramount. An additional objective of AIMS is to provide a forum for doctoral students, the future leaders of our research, to discuss their research with a wider audience and receive training to help make their research careers successful. To this end, AIMS incorporates a European PhD Student Symposium and a tutorial programme that covers a broad range of topics.

We have sought sometimes bold or ambitious ideas, even those that are unfinished, and this naturally invites controversy. We have ensured nevertheless the highest standard, and each paper in this volume has received not merely acceptance, but at least one enthusiastic endorsement from a referee. For the main track, 58 submissions were received and 14 were accepted. For the European PhD symposium, 31 submissions were received and 18 accepted. Of the 15 tutorials proposed, ten were provided.

The AIMS conference was arranged and sponsored by the IST EMANICS Network of Excellence (#26854), in cooperation with the ACM. This yielded an established network of experts in the field of network and service management. Networks are not static of course, they grow as meaningful dialogue pushes at the borders of the established connections, spreading into other fields. Where should one draw the borders of a network? This is a question that some of our contributors have asked in their work. Canright and Engø-Monsen tell us that the natural borders in a network can be defined where dialogue stops, and one-way communication begins. So we have fashioned the AIMS conference around dialogue, not just presentation. With short presentations and extensive, guided

discussion of each paper, it is our goal to capture the spirit of scientific discourse, and to inspire and propel the PhD students to be the next generation of experts in our field.

We would like to thank everyone who gave their time to the project, either as a contributor or as an organizer. We are grateful to the Programme Committee, reviewers and to the organizers of students symposium and tutorial tracks for their efforts. Most of all we thank the authors of the papers for allowing us to realize the goal of our conference.

During the conference, we also benefited from student volunteers and behind-the-scenes administrators who mobilized the conference smoothly and seamlessly. The conference was supported by the EMANICS Network of Excellence (http://www.emanics.org).

June 2007

Arosha Bandara
Mark Burgess

Organization

AIMS 2007 was organized by the EC IST-EMANICS Network of Excellence (#26854) in cooperation with ACM SIGAPP, SIGMIS and Oslo University College.

Main Conference Chairs

Arosha Bandara	Open University, UK
Mark Burgess	Oslo University College, Norway

Student Symposium Chair

Rolf Stadler	Royal Institute of Technology, Sweden

Tutorial Coordinators

Isabelle Chrisment	LORIA, France
David Hausheer	University of Zurich, Switzerland

Student Symposium Technical Programme Committee

Rolf Stadler	Royal Institute of Technology, Sweden
Gyorgy Dan	Royal Institute of Technology, Sweden
Emil Lupu	Imperial College London, UK
Joan Serrat	Universitat Politècnica de Catalunya, Spain
Burkhard Stiller	University of Zurich, Switzerland
Felix Wu	University of California Davis, USA

Local Arrangements

Conference Coordinator	Mark Burgess
Finance Coordinator	Marianne Kjær Bakke
Travel and Finance Coordinator	Gro Birgitte Ruste
Marketing Liason	Kjersti Hilden Smørvik
Conference Registration	Jon Ørstavik
Conference Design	Mark Burgess

Technical Programme Committee

Arosha Bandara	The Open University, UK
Jan Bergstra	University of Amsterdam, The Netherlands
Claudio Bartolini	HP Labs, USA
Artur Binczewski	Poznan Supercomputing and Networking Center, Poland
Mark Burgess	Oslo University College, Norway
Geoffrey Canright	Telenor Research, Norway
Isabelle Chrisment	LORIA, Nancy University, France
Alva Couch	Tufts University, USA
Gabi Dreo	University of Federal Armed Forces Munich, Germany
Olivier Festor	INRIA, France
David Hausheer	University of Zurich, Switzerland
Heinz-Gerd Hegering	Leibniz Supercomputing Center, Germany
Alexander Keller	IBM, USA
Karst Koymans	University of Amsterdam, The Netherlands
Jorge Lobo	IBM Research, USA
Emil Lupu	Imperial College London, UK
Hanan Lutfiyya	University of Western Ontario, Canada
George Pavlou	University of Surrey, UK
Aiko Pras	University of Twente, The Netherlands
Helmut Reiser	Leibniz Supercomputing Center, Germany
Frode Eika Sandnes	Oslo University College, Norway
Jacques Sauve	UFCG, Brazil
Jürgen Schönwälder	Jacobs University Bremen, Germany
Joan Serrat	Universitat Politècnica de Catalunya, Spain
Morris Sloman	Imperial College London, UK
Rolf Stadler	Royal Institute of Technology, Sweden
Burkhard Stiller	University of Zurich, Switzerland
John Strassner	Motorola Inc., USA

Additional Reviewers

Cristian Morariu	University of Zurich, Switzerland
Edgar Magana	Cisco Systems, USA
Ha Manh Tran	Jacobs University Bremen, Germany
Helmut Reiser	Leibniz-Supercomputing Centre, Germany
Isabelle Chrisment	LORIA - University of Nancy, France
Marco Aurlio Spohn	Universidade Federal de Campina Grande, Brazil
Siri Fagernes	Oslo University College, Norway
Thomas Bocek	University of Zurich, Switzerland
Tiago Fioreze	University of Twente, The Netherlands

Table of Contents

Scalable Network Management

On the Impact of Management Instrumentation Models on Web Server Performance: A JMX Case Study 1
Abdelkader Lahmadi, Anca Ghitescu, Laurent Andrey, and Olivier Festor

RAQNet: A Topology-Aware Overlay Network 13
Seyed Iman Mirrezaei, Javad Shahparian, and Mohammad Ghodsi

IBGP Confederation Provisioning 25
M. Nassar, R. State, and O. Festor

Inter-domain Concepts

Ontology-Based Management for Context Integration in Pervasive Services Operations 35
J. Martín Serrano, Joan Serrat, Sven van der Meer, and Mícheál Ó Foghlú

A Terminology for Control Models at Optical Exchanges 49
Freek Dijkstra, Bas van Oudenaarde, Bert Andree, Leon Gommans, Paola Grosso, Jeroen van der Ham, Karst Koymans, and Cees de Laat

Self-forming Network Management Topologies in the Madeira Management System 61
Liam Fallon, Daryl Parker, Martin Zach, Markus Leitner, and Sandra Collins

Activity-Based Scheduling of IT Changes 73
David Trastour, Maher Rahmouni, and Claudio Bartolini

Promises and Ubiquitous Management

Estimating Reliability of Conditional Promises 85
Alva L. Couch, Hengky Susanto, and Marc Chiarini

Modeling Change Without Breaking Promises 97
Alva L. Couch, Hengky Susanto, and Marc Chiarini

Norms and Swarms 109
Mark Burgess and Siri Fagernes

此为试读，需要完整PDF请访问：www.ertongbook.com

Autonomous Infrastructure and Security

Providing Seamless Mobility in Wireless Networks Using Autonomic Mechanisms 121
John Strassner, Barry Menich, and Walter Johnson

A 'Pumping' Model for the Spreading of Computer Viruses 133
Geoffrey Canright and Kenth Engø-Monsen

Improving Anomaly Detection Event Analysis Using the EventRank Algorithm 145
Kyrre Begnum and Mark Burgess

Management Models

Modeling and Performance Evaluation of the Network and Service Management Plane 156
Abdelkader Lahmadi, Laurent Andrey, and Olivier Festor

Abstractions to Support Interactions Between Self-Managed Cells 160
Alberto Schaeffer-Filho and Emil Lupu

Self-management Framework for Unmanned Autonomous Vehicles 164
Eskindir Asmare and Morris Sloman

Towards a Toolkit for the Analysis and Design of Systems with Self-Management Capabilities 168
Ralf Koenig and Heinz-Gerd Hegering

Policy Interactions

Estimating Local Cardinalities in a Multidimensional Multiset 172
Patrick Truong and Fabrice Guillemin

Harnessing Models for Policy Conflict Analysis 176
Steven Davy and Brendan Jennings

Distributed End-to-End QoS Contract Negotiation 180
Hélia Pouyllau and Stefan Haar

Distributed and Heuristic Policy-Based Resource Management System for Large-Scale Grids 184
Edgar Magaña and Joan Serrat

Security Management

Risk Based Authorisation for Mobile Ad Hoc Networks 188
Nilufer Tuptuk and Emil Lupu

Malware Models for Network and Service Management 192
Jérôme François, Radu State, and Olivier Festor

A Survey of the High-Speed Self-learning Intrusion Detection Research Area 196
Anna Sperotto and Remco van de Meent

Logic and Validation

Distributed Case-Based Reasoning for Fault Management 200
Ha Manh Tran and Jürgen Schönwälder

Understanding Promise Theory Using Rewriting Logic 204
Kyrre Begnum and Mark Burgess

Graph Models of Critical Infrastructure Interdependencies 208
Nils Kalstad Svendsen and Stephen D. Wolthusen

Networks

Self-management of Lambda-Connections in Optical Networks 212
Tiago Fioreze and Aiko Pras

A Distributed Architecture for IP Traffic Analysis 216
Cristian Morariu and Burkhard Stiller

Iterative Key Based Routing for Web Services Addressing and Discovery 221
Marco Milanesio and Giancarlo Ruffo

Peer-to-Peer Large-Scale Collaborative Storage Networks 225
Thomas Bocek and Burkhard Stiller

Tutorials

Xen Virtualization and Multi-host Management Using MLN 229
Kyrre Begnum

Ponder2 - Policy-Based Self Managed Cells 230
Kevin Twidle and Emil Lupu

From Charging for QoS to Charging for QoE: Internet Economics in the Era of Next Generation Multimedia Networks 231
Peter Reichl

Next Generation Semantic Business Process Management 232
Armin Haller and Marin Dimitrov

Programmability Models for Sensor Networks 233
Torsten Braun

Scalable Routing for Large Self-organizing Networks 234
Thomas Fuhrmann

The IT Infrastructure Library (ITIL) – An Introduction for Practitioners and Researchers 235
Thomas Schaaf

Peer-to-Peer Market Places: Technical Issues and Revenue Models 236
Giancarlo Ruffo

Author Index 237

On the Impact of Management Instrumentation Models on Web Server Performance: A JMX Case Study*

Abdelkader Lahmadi, Anca Ghitescu, Laurent Andrey, and Olivier Festor

LORIA - INRIA Lorraine - Université de Nancy 2
615 rue du Jardin Botanique
F-54602 Villers-lès-Nancy, France
{Abdelkader.Lahmadi,Anca.Ghitescu,Laurent.Andrey,Olivier.Festor}@loria.fr

Abstract. JMX (Java Management eXtension) is a Java framework that allows any Java technology-based application or accessible resource to become easily manageable. This standard begins to be widely used within different managed systems which vary from large mainframes to small mobile devices, limited in both resource and computing capacity. Today, little is known about the costs associated with the manageability of a system. In this paper, we analyse the impact of various instrumentation models on the behavior of both the functional and the management plane. We show on a JMX instrumented web server that the service is highly affected by the management activity in driver and component models while a daemon approach limits the management impact on the functional service.

Keywords: JMX, Agent, Daemon, Driver, Component, Benchmarking, Management Performance.

1 Introduction

The Java technology deployment varies from small devices to huge data centers with a considerable number of servers. The functionality that controls these applications work is split into two main planes:(i) a value-added plane or *functional plane* that handles the users data ; (ii) the management plane that monitors and configures the functional plane. While the original functional plane was designed to be independent from the management plane, today's applications and services are far more integrated and more complex than before. The functional plane needs to expose both client's services and management interfaces.

Another important trend over the past couple of years is the emergence of the JMX standard for managing Java based applications, mainly the J2EE applications [3]. This standard aims to provide a management architecture and an

* Some of the authors of this paper are supported in part by the IST-EMANICS Network of Excellence project.

A.K. Bandara and M. Burgess (Eds.): AIMS 2007, LNCS 4543, pp. 1–12, 2007.
© Springer-Verlag Berlin Heidelberg 2007

API set that allows any Java technology-based or accessible resource to be inherently manageable. As the number of resources being managed grows and the systems become more distributed and more dynamic, the behavior of application management technologies such as JMX needs to be studied. The overhead of management activities could be important on the user perceived performance of a JMX based managed applications such as a web server where delays and throughput are the key performance metrics for quality of service guarantee [11].

In the past few years, several works [11,1,5] have looked at the performance of multi-tier Internet services which are the base of many businesses, such as retailers, auctioneers and banks. Most of them focus on studying their performance independently from the existence of the management tier. Therefore there is a need to study the performance of such Internet applications accounting for the behavior of the management tier.

Little is known about the cost associated with JMX based management activities. To assess these costs, it is necessary to collect data about the performance and operations of this management system. Furthermore, it is important to collect this data under various configurations and management assumptions. One aspect of these management configurations is the integration model of a JMX agent within a managed system. In the literature, three integrations models are proposed: *daemon*, *driver* and *component* [8]. Overhead associated with management activities of those three models on a managed system performance is unavoidable apart from switching off any instrumentation. However, basic questions we are trying to answer arise: *Does the three models impact differently a managed system performance ? Does it also impact the management part's performance ? Which model is more appropriate and in which context ?*

The main contribution of this paper is to present an experience report on the design and implementation of a simple benchmark to evaluate the three integration models and their impact on the user perceived delays and throughput of the managed web server. In complement to the work in [9], we compared the three models against the same managed application which is a small Java based web server (Tiny Java Web Server[1]) to derive credible performance assessment within comparable experimental environment. This could be helpful for JMX based management systems designers to select one of them regarding performance metrics guarantee. The remainder of this paper is organized as follows. Section 2 gives an overview of the three integration models of a JMX agent. Section 3 describes our benchmarking methodology to assess the impact of the three models. Section 4 analyses the obtained empirical results. Section 5 presents concluding remarks.

2 JMX Agent Integration Models

A common paradigm for traditional management systems and even for the autonomic management framework [7] is the manager-agent model [8]. The JMX management framework [10] is based on this model. In this model, agents mediate between managed devices or services, and managers. An agent is responsible

[1] see: http://sourceforge.net/projects/tjws/

for processing requests from managers, gathering responses, returning them to the requester and issuing notifications when registered events occur. The JMX agent, or MBean server, satisfies the definition and requirements of a management agent. Its functionality is to control the resources and make them available to remote management applications. According to the JMX specification, the JMX agent is composed of the MBean server, MBeans mapped on managed resources, agent services and adaptors or connectors. Basically the MBean server is a registry of objects that provide to the managers the attributes and operations through which the resources are monitored. On the other side, the manager is responsible for collecting management information from the agents using one of the two modes: polling or notification, and then takes any defined automated actions or solicits a human decision. The JMX API does not provide a specification for a manager, but it supplies all necessary requirements to design and develop a full management system. We are mostly interested in the deployment of the management agent in the managed resource. The main question here is: *how can a new or an existing application be designed in order to become manageable.* Therefore we implemented the three agent integration models identified by the authors of [8]: *daemon*, *driver* and *component.*

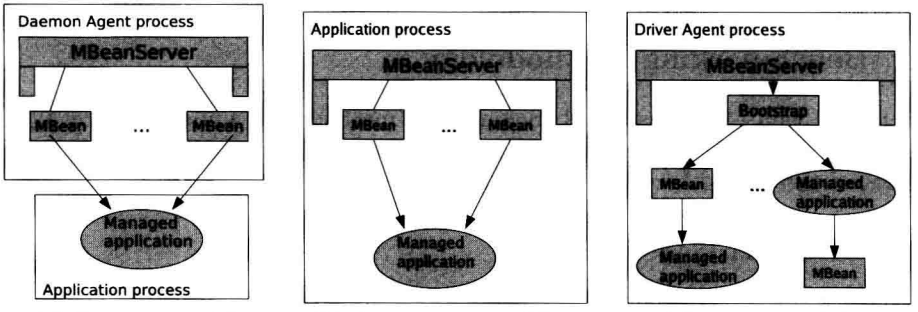

(a) Daemon model (b) Component model (c) Driver model

Fig. 1. JMX agent integration models

2.1 Daemon Model

In the daemon model as depicted in figure 1(a), the JMX agent and the managed application are running in two separate processes. The managed application is not influenced by the overhead imposed by the JMX implementation. Therefore in case of low utilization by the JMX agent, the performance of default functionality should not be affected. One advantage of using the daemon model is that the system can have a single MBean server available for multiple applications. Due to its existence outside the scope of managed applications, the MBean server can be used to control the life-cycle of the application or possibly many applications. On the other hand, problems may arise when the applications try to find or use the MBean server.

此为试读,需要完整PDF请访问: www.ertongbook.com