UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Sophos Ltd. and Sophos Inc.
Petitioners,

v.

Webroot Inc.,
Patent Owner.

Case No. IPR2024-00252

U.S. Patent No. 8,201,243

Declaration of Dr. Mary K. Bolin
In Support of Petitioners in the
Petition for *inter Partes* Review of U.S. Patent No. 8,201,243

Mail Stop **PATENT BOARD**
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

I, Mary K. Bolin, declare as follows:

1. I am Emeritus Professor at the University of Nebraska, Lincoln Libraries. I have held the following positions: Professor and Catalog and Metadata Librarian, Digital Initiatives and Special Collections Department, University of Nebraska—Lincoln Libraries (2015–2020); Professor and Chair, Discovery and Resource Management and Coordinator of Collection Development, University of Nebraska, Lincoln Libraries (2004–2015); Lecturer, San José State University School of Information: Lecturer, (2008–present); Head of the Technical Services Department, University of Idaho Library (1993–2003; promoted to Professor in 1999); Head of the Cataloging Department, University of Idaho Library (1986–1993; promoted to Associate Professor and tenured in 1991); Head of Records Maintenance Section, Cataloging Department, University of Georgia Libraries (1981–1986).

2. Over the course of my work as a librarian, professor, researcher, student research supervisor, and author of numerous scholarly papers and other publications, I have had extensive experience with cataloging records and online library management systems built around Machine-Readable Cataloging (MARC) standards. I also have substantial experience in authenticating printed documents and establishing the date when they were accessible to ordinarily skilled researchers.

3. Attachment A is my full resume.

4. I have been retained by counsel for Sophos Ltd. and Sophos Inc. (collectively, "Petitioners") to evaluate, authenticate, and establish the dates of public availability of certain documents in the above-captioned cases. For my time performing this service, I am being paid my usual hourly fee. I have no financial interest in this

Sophos Ex. 1009
Page 2 of 318

litigation, or affiliation with Petitioners. My compensation is not dependent upon the outcome of, or my testimony in, this litigation or any other proceeding, including before the U.S. Patent Office.

A. Documents of Interest

5. The documents that are the subject of this declaration are the following:

- **Document 1**. Li, Z. et al., *UCLog: A Unified, Correlated Logging Architecture for Intrusion Detection*. 12th Int'l Conf. on Telecomm. Sys.: Modeling and Analysis (ICTSM), 330-343 (2004).

- **Document 2.** Salvador Mandujano Vergara, *A Multiagent Approach to Outbound Intrusion Detection* (2004) (published Ph.D. dissertation, Instituto Tecnológico y de Estudios Superiores de Monterrey).

B. Methodology

6. I searched each of these documents in the OCLC database, Google Scholar, and one or more subscription databases such as IEEE Xplore and ACM Digital Library. I also searched the library catalogs of institutions such as the University of Nebraska—Lincoln and other large research libraries. My goal was to find a record in one of these databases that contained a date indicating when the document was available to the public.

7. *Use of MARC Records*. The "Machine Readable Cataloging" standard, also known as "MARC," is an industry-wide standard method of organizing library catalog information. The Library of Congress first developed MARC in the 1960s. A MARC-compatible library is one that has a catalog consisting of individual MARC records for each of its works. Today, MARC is the primary communication protocol for the transfer and storage of bibliographic metadata in libraries.

8. OCLC is a database of MARC records used by libraries worldwide. When an OCLC-participating institution acquires a work, it creates a MARC record for this work

in its computer catalog system, and OCLC automatically supplies the date of creation. The MARC record creation date reflects the date on which (or shortly after) the item was first acquired and catalogued. Once the MARC record is created by a cataloger at an OCLC-participating member institution, it is instantaneously available to any OCLC-participating members, and, therefore available to the public.

9. OCLC was created "to establish, maintain and operate a computerized library network and to promote the evolution of library use, of libraries themselves, and of librarianship, and to provide processes and products for the benefit of library users and libraries, including such objectives as increasing availability of library resources to individual library patrons and reducing the rate of rise of library per-unit costs, all for the fundamental public purpose of furthering ease of access to and use of the ever-expanding body of worldwide scientific, literary and educational knowledge and information."[1] Among other services, OCLC and its members are responsible for maintaining the WorldCat database (http://www.worldcat.org/), which is used by independent and institutional libraries throughout the world. OCLC is the international cooperative to which libraries catalog records and others pay to access MARC records. The MARC records referenced in this declaration, extracted from the OCLC database, were generated in the ordinary course of business. Further, I personally located and downloaded each of these records associated with the documents on which my declaration focuses.

10. OCLC provides its members online access to MARC records through its OCLC Connexion database. When an OCLC member institution acquires a work, it

[1] Third Article, Amended Articles of Incorporation of OCLC Online Computer Library Center, Incorporated (available at http://www.oclc.org/en-US/membership/councils.html) (Attachment B).

creates a MARC record for this work in its computer catalog system in the ordinary course of its business. MARC records created at the Library of Congress are batch-loaded into the OCLC database through a subscription to MARC Distribution Services daily or weekly. Once the MARC record is created by a cataloger at an OCLC member institution or is tape-loaded from the Library of Congress, it is then made available to any other OCLC members online, and, therefore, made available to the public.

11. A MARC record comprises several fields, each of which contains specific data about the work. Each field is identified by a standardized, unique, three-digit code corresponding to the type of data that follows. For example, a work's title is recorded in field 245, the primary author of the work is recorded in field 100, an item's International Standard Book Number ("ISBN") is recorded in field 020, an item's Library of Congress call number is recorded in field 050, and the publication date is recorded in field 260 or 264 under the subfield "c." If a work is a periodical, then its publication frequency is recorded in field 310, and the publication dates (*e.g.*, the first and last publication) are recorded in field 362, which is also referred to as the enumeration/chronology field. Field 650 can include terms related to the topical subject of the work. The library that created the record is recorded in field 040 in subfield "a" with a unique library code. When viewing the MARC record online via OCLC's Connexion database, hovering over this code with the mouse reveals the name of the library. I used this method of "mousing over" the library codes in the OCLC database to identify the originating library for all of the MARC records in this declaration, as well as the "display holdings" command to see a list of libraries that own the item. In the MARC records in the OCLC Connexion database,

MARC field 040 contains the OCLC symbols of institutions transcribing, creating, and modifying the record.

12. When an OCLC member institution creates a new record, OCLC automatically supplies the date of creation for that record. The MARC record creation date reflects the date on which, or shortly after which, the item was first acquired or cataloged.

13. The date of creation for the MARC record appears in the fixed field (008), characters 00 through 05. Field 005 is a 16 digit field that records the Date characters yyyymmdd including 4 for the year, 2 for the month, and 2 for the day. The Date entered on file never changes. These dates are automatically supplied by the OCLC system and cannot be changed by the user. When the MARC record is created in OCLC, it is available throughout the world to other libraries and in WorldCat, the public version of OCLC Connexion. Thus, the date of creation of the MARC record by a cataloger at an OCLC member institution reflects when the underlying item is accessible to the public.

14. For each of Documents 1-2, I included a search of OCLC, www.oclc.org, which is a large shared database used by libraries around the world to create and exchange cataloging data, using MARC format, which is a universal communications format for that data.

15. For those documents for which there is a record in OCLC, I determined when the record was created and which libraries own the item. I then visited the website of the library or libraries that own the item, searched their library catalog, and determined when the item was added to the library's collection.

16. I am familiar with many online catalogs, databases, and search engines. In preparing for this declaration, I used authoritative information systems such as the above-mentioned WorldCat (https://www.worldcat.org), the online catalog of the Library of Congress (https://catalog.loc.gov), shared research databases such as ResearchGate and CiteSeerX, the online catalog of a number of large research libraries such as the University of Nebraska—Lincoln, IEEExplore, and Google Scholar search for records. These records are identified and discussed in this declaration.

C. Document 1

- Li, Z. et al., *UCLog: A Unified, Correlated Logging Architecture for Intrusion Detection*. 12th Int'l Conf. on Telecomm. Sys.: Modeling and Analysis (ICTSM), 330-343 (2004).

17. ***Authentication of Document 1***

18. Document 1 is an article published in the proceedings of the *12th International Conference on Telecommunication Systems-Modeling and Analysis (ICTSM)* ("proceedings"). I searched OCLC for any MARC records related to Document 1 and/or the proceedings. I found two relevant MARC records in OCLC.

19. First MARC Record: Attachment D is the MARC record from OCLC created by Leibniz Information Centre for Science and Technology University Library in Hannover, Germany ("TIB"). As can be seen in Attachment D, the MARC record has an "Entered" date of 20050531, which means it was created on May 31, 2005 by TIB.

20. Attachment D indicates ("attached holdings") that TIB has the proceedings in its collection, as it is common practice to create a MARC record for each publication added to a library's collection.[2]

[2] TIB has declined to provide a scanned copy of Document 1 from their records, indicating that they do not provide assistance to librarians in the United States.

21. Second MARC Record: Attachment E is the MARC record from OCLC created by the National Research Council Canada ("NRC"). As shown in Attachment E, the MARC record has an "Entered" date of "20050422", which means it was created on April 22, 2005 by NRC.

22. Attachment E indicates ("attached holdings") that NRC has the proceedings in its collections.

23. I was able to obtain a copy of the proceedings from NRC. Attachment F is the PDF of Document 1 from the collection of NRC. There is a label on the back cover that indicates that the proceedings were added to the NRC collection on "5-04-21", i.e. April 21, 2005.[3] The label is viewable at Attachment F, pg. 2. This is consistent with the MARC record, as the MARC record was created the day after the label on Attachment F indicates the NRC received the proceedings.

24. The Library of Congress's call number is recorded in field 050. Attachment E indicates that the call number for the proceedings in field 050 is "TK5101.A1 I619 2004." In addition to being classified under the Library of Congress's classification system, field 650 contains several terms, both in English and French, related to the subject matter of the proceedings. Thus, as of its cataloging, the proceedings, which contained Document 1, was indexed according to its subject matter by the Library of Congress's classification system and the field 650 terms.

25. ***Public Availability of Document 1.***

26. Based on the fact that there are two MARC records in OCLC for Document 1, one created by NRC and one by TIB, created on April 22, 2005 and May 31, 2005,

[3] Canada uses the Y-M-D date format: Canada uses YMD https://en.wikipedia.org/wiki/Date_format_by_country. *See* Attachment C.

Sophos Ex. 1009
Page 8 of 318

respectively, and that National Research Council Canada added it to their collection on 4-22-2005, as it would have been customary, shortly after a document was received from NRC, to have it be publicly available within a month of receipt, it is my opinion that Document 1 was available by June 2005 at the latest and could have been located by a skilled researcher.

D. Document 2.

- Salvador Mandujano Vergara, *A Multiagent Approach to Outbound Intrusion Detection* (2004) (published Ph.D. dissertation, Instituto Tecnológico y de Estudios Superiores de Monterrey).

27. ***Authentication of Document 2.***

28. Document 2 is found in Google Scholar, and a PDF, Attachment G, was retrieved from the Institutional Repository (archive of theses, dissertations, and other scholarly output from one institution) of the Instituto Tecnologico y de Estudios Superiores de Monterrey at https://repositorio.tec.mx/handle/11285/572223.

29. Attachment H is the MARC record from OCLC created by Instituto Tecnologico y de Estudios Superiores. As shown in Attachment H, the MARC record has an "Entered" date of "20041209", which means it was created on December 9, 2004 by Instituto Tecnologico y de Estudios Superiores.

30. The Library of Congress's call number is recorded in field 050. Attachment H indicates that the call number for the proceedings in field 050 is "TK5105.59 2004." In addition to being classified under the Library of Congress's classification system, field 650 contains the descriptor "Seguridad de computadores." Thus, as of its cataloging, the proceedings, which contained Document 1, was indexed according to its subject matter by the Library of Congress's classification system and the field 650 term.

Sophos Ex. 1009
Page 9 of 318

31. Attachment I is the record for Document 2 from the online library catalog of the Instituto Tecnologico y de Estudios Superiores de Monterrey, showing a publication date of 2004, and a date in the Description field of December 2004. (https://www.babbel.com/en/magazine/how-to-write-the-date-in-spanish#:~:text=In%20pretty%20much%20every%20country,be%2007%2F09%2F2014). This is the record that is shown if someone finds this document using the library online catalog of the Instituto Tecnologico y de Estudios Superiores. Attachment I also indicates that there are various terms associated with Document 2, such as "Detection" and "Intrusion."

32. Attachment J is the record for Document 2 from the institutional repository of the Instituto Tecnologico y de Estudios Superiores de Monterrey, showing a publication date of 2004-1-12, *i.e.*, December 1, 2004. *See* Attachment K (https://www.babbel.com/en/magazine/how-to-write-the-date-in-spanish#:~:text=In%20pretty%20much%20every%20country,be%2007%2F09%2F2014). This is the record that is shown if someone finds this document in Google Scholar.

33. ***Public Availability of Document 2***

34. Based on the fact that Document 2 is found in Google Scholar, and a PDF was retrieved from the Institutional Repository (archive of theses, dissertations, and other scholarly output from one institution) of the Instituto Tecnologico y de Estudios Superiores de Monterrey with a publication date of December 1, 2004, in the library catalog with a publication date of December 2004, and that the MARC record shows a creation date of December 9, 2004, and that these documents showed that Document 2

was indexed and/or catalogued, Document 2 was available in December 2004, and could have been located by a skilled researcher.

* * *

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements are made with the knowledge that willful false statements and the like are punishable by fine, imprisonment, or both, under Section 1001 of Title 18 of the U.S. Code. I state under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed this December 11, 2023, in Lincoln, Nebraska.



Mary K. Bolin

Sophos Ex. 1009
Page 11 of 318

ATTACHMENT A

Mary K. Bolin, PhD
2433 U Street
Lincoln, Nebraska 68503
mary@marykbolin.com

Mary K. Bolin, PhD

Professor Emeritus
University Libraries
University of Nebraska—Lincoln
mbolin2@unl.edu

Lecturer,
School of Information
San José State University
mary.bolin@sjsu.edu

EDUCATION

PhD, Educational Studies (Educational Leadership and Higher Education), 2007, University of Nebraska

Dissertation: *A Typology and Discourse Analysis of the Status and Appointments of Librarians at Land Grant Universities* (available at digitalcommons.unl.edu)

MA, English, 1999, University of Idaho
Thesis: *Grace: A Contrastive Analysis of a Biblical Semantic Field* (available at digitalcommons.unl.edu)

MSLS, 1981, University of Kentucky

BA, Linguistics, 1976, University of Nebraska—Lincoln

WORK EXPERIENCE

Center for Digital Research in the Humanities. Fellow, 2019- Emeritus, 2020-

University of Nebraska--Lincoln Libraries: Professor and Catalog and Metadata Librarian, Digital Initiatives and Special Collections, 2015-2020. Emeritus, 2020-

University of Nebraska—Lincoln Libraries: Professor and Chair, Discovery and Resource Management, 2004-2015; Coordinator of Collection Development, 2012-2015

University of Idaho Library: Head, Technical Services Department, 1993-2003. Promoted to Professor in 1999

University of Idaho Library: Head, Cataloging Department, 1986-1993. Promoted to Associate Professor and tenured, 1991

University of Georgia Libraries: Head, Records Maintenance Section, Cataloging Department, 1981-1986

Teaching:
San José State University School of Information:

LIBR 281, Metadata. 2009-
LIBR 248, Introduction to Cataloging and Classification. 2013-
LIBR 202, Information Retrieval. 2008-2013

Associate Coordinator, Gateway PhD program, San Jose State University, 2021-2022

American Library Association

Cataloging for Non-Catalogers, Online Continuing Education Course, 2017-
Intermediate Cataloging for Non-Catalogers, Online Continuing Education Course, 2018-

Lyrasis

Webinar on BibFrame, taught twice a year, 2020-
Quick and Clean Cataloging, Parts 1, 2, and 3. 2022-

Independent Study in Idaho

Library Science C415, Technical Services for Small Libraries, 2004-2013
Library Science C416, Library Services to Rural Libraries, 2004-2013
Library Science C422, Use of the School Library, 1999-2013

Library Science C425/525, Organization and Management of Small Libraries, 1988-2013
Library Science C427, Library or Media Center Practicum, 1995-2013

Gateway PhD Program (San Jose State University and Queensland University of Technology). SJSU supervisor for two PhD students, 2015-

Gateway PhD Program (San Jose State University and Manchester Metropolitan University). Associate Coordinator, 2021-2022

Consulting
Expert witness for multiple attorneys employed by Wolf, Greenfield & Sacks, P.C., Hughes Hubbard & Reed LLP, and other firms. 2017-

Editing
Language and academic editing for EditSprings, a service for international scholars, 2020-

SCHOLARSHIP

Refereed Publications:

(Publications available from digitalcommons.unl.edu)

Pytlik Zillig, Brian, and Mary K. Bolin. (Forthcoming 2023).“Visualization of Semantic Primes in Texts from Early Modern English Drama.” *Journal of the Japanese Association for the Digital Humanities.*

Bolin, Mary K. (2022). R*efocusing Academic Libraries through Learning and Discourse: The Idea of a Library.* Chandos/Elsevier.

Bolin, Mary K. (2022). “Academic Libraries and Writing Centers: Collaborations at US Public Research Universities.” *School of Information Student Research Journal, 11*(2). Retrieved from https://scholarworks.sjsu.edu/ischoolsrj/vol11/iss2/2

Bolin, Mary K. (2022). *Refocusing Academic Libraries through Learning and Discourse: The Idea of a Library*. Oxford, UK: Chandos/Elsevier.

Bolin, Mary K. (2022). "Metadata, Catalog Systems, Linked Data, and the Evolving ILS." In: Sandra Hirsh (Ed.), *Information Services Today: An Introduction.* 3rd ed. New York: Rowman and Littlefield.

Wintermute, Harriet E., and Bolin, Mary K. (2018). "Metadata Generators for Backlog Reduction: Metadata Maker Streamlines Cataloging and Facilitates Transition. *Technical Services Quarterly 35*(2) doi: https://doi.org/10.1080/07317131.2018.1422879

Bolin, Mary K. (2018). "Metadata, Catalog Systems, Linked Data, and the Evolving ILS." In: Sandra Hirsh (Ed.), *Information Services Today: An Introduction.* 2nd ed. New York: Rowman and Littlefield.

Bolin Mary K. (2017). *The 21-century Academic Library: Global Patterns of Organization and Discourse.* Oxford: Chandos Publishing. (Chandos Information Professional Series)

Bolin, Mary K. (2017). Linguistics and LIS: A Research Agenda." *SLIS Student Research Journal,* 7:1.
Retrieved from http://scholarworks.sjsu.edu/slissrj/vol7/iss1/2

Bolin, Mary K. (2014). "Open Access Web Resources for Library Continuing Education and Training." In S. Hines (Ed.), *Revolutionizing the Development of Library and Information Professionals: Planning for the Future* (pp. 107- 124). Hershey, PA: .
doi:10.4018/978-1-4666-4675-9.ch008.

Bolin, Mary K. (2014). "The Language of Academic Librarianship: The Discourse of Promotion and Tenure. "*Advances in Library Administration and Organization 32*.

Bolin, Mary K. "Librarian Status at US Research Universities: Extending the Typology." *Journal of Academic Librarianship* 34:4 (September 2008)

Bolin, Mary K. "A Typology of Librarian Status at Land Grant Universities." *Journal of Academic Librarianship*, 34: 3 (May 2008)220:230.

Bolin, Mary K. "The Library and the Computer Center: Organizational Patterns at Land Grant Universities." *Journal of Academic Librarianship*, 31:1 (2005)

Eckwright, Gail Z. and Mary K. Bolin. "The Hybrid Librarian: The Affinity of Collection Management with Technical Services and the Organizational Benefits of an Individualized Assignment." *Journal of Academic Librarianship*, 27:6 (2001)452:457.

Bolin, Mary K. "Catalog Design, Catalog Maintenance, Catalog Governance," L*ibrary Collections, Acquisitions, and Technical Services.* 24: (2000)53-63.

Bolin, Mary K. "The Collegial Environment and the Functional Organization," *Journal of Library Administration*. 29:2(1999):49-61.

Bolin, Mary K. "Cataloging Environmental Impact Statements and Related Documents," *Resource Sharing and Library Networks*. 14:1(1999): 21-40.

Bolin, Mary K. "Waiting for the Electrician or Someone like Him: Client-Centered Technical Services" Proceedings of the Association of College and Research Libraries National Conference, Nashville, April 1997.
< http://www.ala.org/Content/ContentGroups/ACRL1/Nashville_1997_Papers/Bolin.htm >

Grants:

Electronic Theses and Dissertations. Seed Grant from University of Idaho Office of Research, 2002. Co PI with Gail Z. Eckwright. $6500.

Interoperability of Metadata for Thematic Research Collections. Grant to the Center for Digital Research in the Humanities from the Institute of Museum and Library Services 2005-2008. Project Team Member.
http://unlcms.unl.edu/cas/center-for-digital-research-in-the-humanities/projects/pages/iom_files/IOM_Final_Narrative_Report_July_2008.pdf

Other Publications:

Library Philosophy & Practice. Peer-reviewed electronic library science journal. Mary K. Bolin and Gail Z. Eckwright, founders and editors. http://digitalcommons.unl.edu/libphilprac

Library Science C415, Technical Services for Small Libraries, through Independent Study in Idaho. 2004.

Library Science C416, Library Services to Rural Libraries, through Independent Study in Idaho. 2004.

Library Science C422, Use of the School Library, through Independent Study in Idaho. Revised 1999.

Library Science C425/525, Organization and Management of Small Libraries through Independent Study in Idaho. Revised 2005.

Library Science C427, Library or Media Center Practicum, Independent Study in Idaho. Revised 1999.

Presentations:

Bolin, Mary K. “BibFrame and the MARC Cataloging Environment: an Introduction.” Lyrasis Webinar, June 16, 2020, August 17, 2021.. Recording available http://lyrasis.adobeconnect.com/paufo9e0dm30/

Pytlik Zillig, Brian, and Mary K. Bolin. “A Little More than Kin, a Little Less than Kind”: the Semantic Field Grace in Early Modern English" Poster Session, Text Encoding Initiative Meeting, Graz, Austria, September 18, 2019

Pytlik Zillig, Brian, and Mary K. Bolin. “Shakespeare and the Enumeration of Semantic Primes.” Research Paper Presentation, Text Encoding Initiative Meeting, Tokyo, Japan, September 13, 2018

"Library Philosophy and Practice: Creating and Maintaining an Open-Access E-Journal," Presentation to UNCL meeting (libraries at campuses of the University of Nebraska), Lincoln, Nebraska, March 20, 2008.

Presentation to "Thinking outside the Borders" participants on the open-access electronic journal Library Philosophy and Practice and its role in the international library community. University of Nebraska—Lincoln Libraries, October 24, 2007

Panel discussion on open access electronic journals. GWLA directors' meeting. Linda K. Hall Library, Kansas City, Missouri. April 2007.

"Waiting for the Electrician or Someone Like Him: Client-Centered Technical Services." Association of College and Research Libraries National Conference, Nashville, April 5, 1997.

" 'Some Problems of an Editor': the Idaho Librarian from Manuscript to Published Issue," Joint Annual Conference of the Idaho Library Association and the Washington Library Association, Spokane, Washington, April 6, 1995.

"Automating Idaho's Libraries: a Historical Perspective." Idaho Library Association Annual Conference, Idaho Falls, October 4, 1990.

"Small Library Management." Idaho Library Association Spring Conference, Lewiston, Idaho, May 13, 1989.

"No More 'Good Workers': Excellence and Academic Library Classified Staff." Association of College and Research Libraries 5th Annual Conference, Cincinnati, April 7, 1989.

"Communication to and from Cataloging Departments." Heads of Cataloging Discussion Group, American Library Association Midwinter Meeting, San Antonio, January 1988.

Discussion Leader at the LAMA (Library Administration and Management) Management Exchange. American Library Association Annual Meeting, San Francisco, June 1987.

"Staff Turnover." Heads of Cataloging Discussion Group, American Library Association Midwinter Meeting, Chicago, January 1987.

SELECTED PROFESSIONAL SERVICE

Reviewer. *Resources for College Libraries*. Linguistics bibliography, Spring 2020.

Peer reviewer for ALISE National Conference presentations, Spring 2020.

Peer reviewer for *Journal of Academic Librarianship,* 2012-

Peer reviewer for refereed journal, *Western Nigerian Journal of Medical Sciences* (http://www.wnjms.org/), a publication of Ladoke Akintola University of Technology (LAUTECH) Teaching Hospital, 2008-

PhD Reviewer for students from Aligarh Muslim University, India, 2010-

Outside reviewer for numerous library faculty promotion and tenure cases in US universities.

University of Nebraska--Lincoln

Chair, University Libraries Faculty, 2018

Faculty Senate Academic Rights and Responsibilities Committee, 2016- Chair, 2018-19

Faculty Senate – 2004-2008; 2008-2011

Faculty Senate Executive Committee 2004-2008

Academic Planning Committee (APC), 2005-2008; Vice Chair, 2006-2007; Chair, 2007-2008

Aesthetic Review Committee 2006-2007

APC monitor for academic program reviews: Department of Political Science, 2005; Department of Educational Psychology 2006

Pacific Northwest Library Association:

Editor, *PNLA Quarterly*, December 2002-December 2013

Idaho Library Association:

Editor, *The Idaho Librarian*, September 1990-January 1996

ATTACHMENT B

OCLC GLOBAL & REGIONAL COUNCILS

Member-focused, member-driven

A key benefit of OCLC membership is the opportunity to elect delegates who reflect your local and regional views on the issues facing libraries today. Member libraries have a voice through their elected representatives, and are encouraged to actively participate in elections and member meetings, and through direct dialogue with their elected delegates.

Regional Councils facilitate connections in three world regions: Europe, the Middle East, and Africa (EMEA); Asia Pacific (APRC); and the Americas (ARC). Members from each region elect delegates to serve on the OCLC Global Council. Global Council brings worldwide viewpoints together, informing and guiding the cooperative from their unique perspectives.

Global Council | Americas Regional Council | EMEA Regional Council | Asia Pacific Regional Council

Global Council

The Global Council is comprised of 48 delegates, each representing one of three regions (ARC, EMEA, APRC). The delegates work on behalf of the global membership to reflect the needs of member institutions. Members can contact any Global Council delegate to share ideas, offer feedback, and ask questions. Global Council committees complete much of the group's work. The working groups include:

- Global Council Executive Committee

A message from Evi Tramantza, Global Council Chair



Welcome from your Global Council!

- Global Council Bylaws Committee
- Global Council Nominating Committee
- Global Council Program Committee

Get involved with OCLC

See your Global Council delegates (https://www.oclc.org/en/membership/councils/directory.ht

l Council delegates, we serve to direction, research, products, and lecting the needs of the many ommunities that comprise OCLC's round the world.

Finding common ground is of critical importance to addressing the challenges we face today, from supporting environmental sustainability to encouraging shared understanding in pursuit of equity, openness, progress, prosperity, and peace around the world. By working together, we can accelerate the process of achieving our shared goals and move our members worldwide further, and faster, together.

Please join us in this dialogue in support of global libraries. We invite you to partner with your Global Council delegates in pursuit of the strategies and solutions that will promote our shared values, inform our work, and sustain our impact.

Recent news

Check back soon for new Global and Regional Councils news.

Past events

15 June 2023

Redefining the library experience: Insights and inspiration (https://www.oclc.org/en/events/2023/redefining-library-experience-insights-inspiration.html)

Join Lynn Silipigni Connaway, Executive Director, OCLC Research, and library leaders from OCLC's Global Council as they share results from the 2023 global survey on "Redefining the library experience."

24 January 2023

Redefining the library experience: Putting ideas into action (https://www.oclc.org/en/events/2023/redefining-library-experience-ideas-into-action.html)

Join Lesley Langa, Associate Research Scientist, OCLC Research, and library leaders from OCLC's Global Council for a discussion on how libraries around the world are rethinking the library experience.

Upcoming events

There are no events currently scheduled.

More Global and Regional Council events » (https://www.oclc.org/en/events/councils.html)

Interested in learning more or getting involved?

If you are interested in participating in Americas, Asia Pacific, or EMEA Regional Council activities, or in representing your region on the OCLC Global Council, please contact Member Relations (/content/oclc-forms/en_us/membership/contact.html).

Sophos Ex. 1009
Page 23 of 318

Get involved with OCLC

Stay in the know

Get the latest product updates, research, events, and much more—right to your inbox.

Subscribe now (https://www.oclc.org/en/email.html?utm_campaign=email-subscription-web-organic&utm_medium=digital&utm_source=webpage-email-sub&utm_content=email-subscription-footer-button&utm_term=en)

About

About OCLC (https://www.oclc.org/en/about.html)
Latest news (https://www.oclc.org/en/news.html)
Leadership (https://www.oclc.org/en/about/leadership.html)
Finance (https://www.oclc.org/en/about/finance.html)
OCLC Trust Center (https://www.oclc.org/en/trust.html)
Advancing racial equity (https://www.oclc.org/en/about/diversity-and-advancing-racial-equity.html)
Next blog (https://blog.oclc.org/next/)
Contact Us (https://www.oclc.org/en/contacts.html)

Events

Upcoming events (https://www.oclc.org/en/events.html)
Upcoming webinars (https://www.oclc.org/en/events/webinars.html)

Information for

Academic libraries (https://www.oclc.org/en/academic.html)
Public libraries (https://www.oclc.org/en/public.html)
Research libraries (https://www.oclc.org/en/rlibs.html)
Special libraries (https://www.oclc.org/en/special.html)
Groups and consortia (https://www.oclc.org/en/groups.html)
Partners (https://www.oclc.org/en/partnerships.html)

Membership

Overview (https://www.oclc.org/en/membership.html)
Libraries and groups (https://www.oclc.org/en/membership/institutions.html)
Individuals (https://www.oclc.org/en/membership/individuals.html)
Councils (https://www.oclc.org/en/membership/councils.html)
Resources (https://www.oclc.org/en/membership/member-resources.html)

Products

Wise (https://www.oclc.org/en/wise.html)
WorldShare Management Services (https://www.oclc.org/en/worldshare-management-services.html)
WorldShare Interlibrary Loan (https://www.oclc.org/en/worldshare-ill.html)
WorldCat Discovery (https://www.oclc.org/en/worldcat-discovery.html)
OCLC Cataloging Subscription (https://www.oclc.org/en/cataloging-subscription.html)
EZproxy (https://www.oclc.org/en/ezproxy.html)
Dewey Services (https://www.oclc.org/en/dewey.html)
CONTENTdm (https://www.oclc.org/en/contentdm.html)
All products and services » (https://www.oclc.org/en/services.html)

Visit related sites

Support & Training (https://help.oclc.org/)
OCLC Community Center (https://community.oclc.org)
OCLC Research (https://www.oclc.org/research/home.html)
Developer Network (https://www.oclc.org/developer/home.en.html)
WebJunction (https://www.webjunction.org/home.html)
President's Leadership Blog (http://www.skipprichard.com/)



(https://www.oclc.org/en/careers.html)

Careers

About Careers (https://www.oclc.org/en/careers.html)
View open positions in the Americas (https://oclc.wd1.myworkdayjobs.com/OCLC_
View open positions in Europe and Asia Pacific (https://oclc.wd1.myworkdayjobs.com/OCLC_

(http://www.facebook.com/pages/OCLC/20530435726) (https://www.youtube.com/OCLCvideo) (https://www.linkedin.com/company/oclc) (https://www.instagram.com/oclc_global) (https://blog.oclc.org/next/) (https://www.oclc.org/en/rss.html)

© 2023 OCLC (https://policies.oclc.org/en/copyright.html)
Domestic and international trademarks and/or service marks of OCLC, Inc. and its affiliates (https://policies.oclc.org/en/copyright/trademarks.html)

ATTACHMENT C

WIKIPEDIA
The Free Encyclopedia

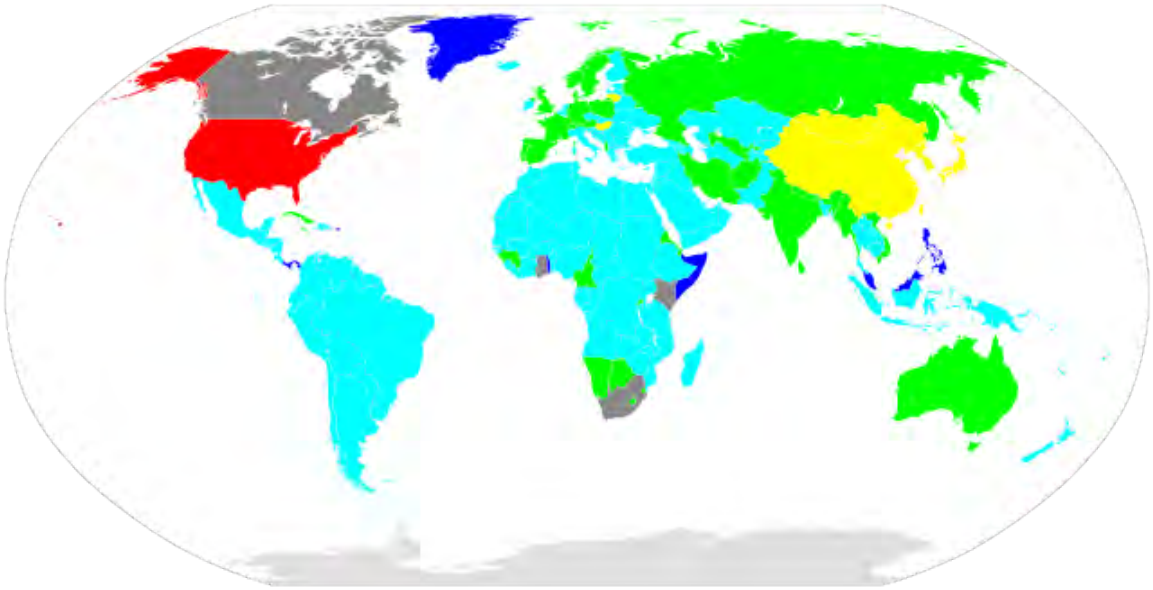
Date format by country

The legal and cultural expectations for date and time representation vary between countries, and it is important to be aware of the forms of all-numeric calendar dates used in a particular country to know what date is intended.

Writers have traditionally written abbreviated dates according to their local custom, creating all-numeric equivalents to day–month formats such as "9 October 2023" (09/10/23, 09/10/2023, 09-10-2023 or 09.10.2023) and month–day formats such as "October 9, 2023" (10/09/23 or 10/09/2023). This can result in dates that are impossible to understand correctly without knowing the context. For instance, depending on the order style, the abbreviated date "01/11/06" can be interpreted as "1 November 2006" for DMY, "January 11, 2006" for MDY, and "2001 November 6" for YMD.

The ISO 8601 format YYYY-MM-DD (2023-10-09) is intended to harmonize these formats and ensure accuracy in all situations. Many countries have adopted it as their sole official date format, though even in these areas writers may adopt abbreviated formats that are no longer recommended.

Usage map



Sophos Ex. 1009
Page 26 of 318

Colour	Order styles	Main regions and countries (population of each region in millions)	Total population (millions)
Cyan	DMY	**Europe**: Italy (60), Ukraine (42), Romania (19), Netherlands (17), others (65) **North America**: Mexico (127), various Caribbean islands (26) **Central America**: Guatemala (18), Honduras (9.2), others (19) **South America**: Brazil (210), Colombia (51), Chile (19), Argentina (45), Peru (32), Venezuela (32), others (24) **North Africa**: Egypt (99), Algeria (43), Morocco (35), Tunisia (12), others (11) **East Africa**: Somalia (16) **West, Central, and Southern Africa**: Nigeria (193), Ethiopia (99), DRC (87), Tanzania (56), Sudan (41), Uganda (40), others (323) **West Asia**: Turkey (82), Iraq (40), Saudi Arabia (33), Yemen (30), others (107) **Central Asia**: Tajikistan (8.9), Kyrgyzstan (6.4), Turkmenistan (5.9) **East and Southeast Asia**: Malaysia (33), Indonesia (268), Thailand (66), Cambodia (16), others (8.9) **South Asia**: Pakistan (212), Bangladesh (166) **Oceania**: Papua New Guinea (8.6), New Zealand (5.0), others (5.5)	2,914
Yellow	YMD	China (1,398), Japan (126), South Korea (52), North Korea (25), Taiwan (24), Hungary (12), Mongolia (3.3), Lithuania (2.8), Bhutan (0.74).	1,643
Magenta	MDY	Some U.S. island territories (0.55)	0.55
Green	DMY, YMD	India (1,366), Russia (147), Vietnam (95), Germany (83), Iran (82), United Kingdom (67), France (66), Myanmar (54), Spain (47), Poland (38), Uzbekistan (33), Afghanistan (32), Nepal (30), Australia (26), Cameroon (24), Sri Lanka (22), others (131)	2,317
Blue	DMY, MDY	Philippines (107), Togo (7.5), Panama (4.2), Puerto Rico (3.2), Cayman Islands (0.63), Greenland (0.056)	122.6
Red	**MDY**, YMD	United States (328)	339
Grey	MDY, DMY, YMD	South Africa (60), Kenya (52), Canada (38), Ghana (30)	180

Listing

Table coding

All examples use example date 2021-03-31 / 2021 March 31 / 31 March 2021 / March 31, 2021 – except where a single-digit day is illustrated.

Basic components of a calendar date for the most common calendar systems:

D – day
M – month
Y – year

Specific formats for the basic components:

yy – two-digit year, *e.g. 21*

yyyy – four-digit year, *e.g. 2021*
m – one-digit month for months below 10, *e.g. 3*
mm – two-digit month, *e.g. 03*
mmm – three-letter abbreviation for month, *e.g. Mar*
mmmm – month spelled out in full, *e.g. March*
d – one-digit day of the month for days below 10, *e.g. 2*
dd – two-digit day of the month, *e.g. 02*
ddd – three-letter abbreviation for day of the week, *e.g. Fri*
dddd – day of the week spelled out in full, *e.g. Friday*

Separators of the components:

/ – oblique stroke (slash)
. – full stop, dot or point (period)
- – hyphen (dash)
– space

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Afghanistan	Yes	Yes	No	Short format: d/m/yyyy (Year first, month, and day in right-to-left writing direction) Long format: yyyy mmmm d (Day first, full month name, and year in right-to-left writing direction)	
Åland	Yes	Yes	No	Short format: yyyy-mm-dd Long format: d mmmm yyyy	
Albania	Yes	Yes	No	dd/mm/yyyy Some YMD[1][2][3]	
Algeria	No	Yes	No	[4] (dd/mm/yyyy)[5]	
American Samoa	No	No	Yes	(mm/dd/yy)	
Andorra	No	Yes	No		
Angola	No	Yes	No		
Anguilla	No	Yes	No		
Antigua and Barbuda	No	Yes	No		
Argentina	Sometimes	Yes	No	Numeric format: yyyyMMdd (Example: 20030613) Short format: dd/mm/yy (Example: 13/06/03) Medium format: dd/mm/yyyy (Example: 13/06/2003) Long format: d' de 'mmmm' de 'yyyy (Example: 13 de junio de 2003) Full format: dddd d' de 'mmmm' de 'yyyy (Example: viernes 13 de junio de 2003).[6]	
Armenia	No	Yes	No	(dd.mm.yyyy)[7][8]	
Aruba	No	Yes	No	[9]	
Australia	Yes	Yes	No	mmmm d, yyyy is sometimes used, usually informally in the mastheads of magazines and newspapers,[10][11] and in advertisements, video games, news, and TV shows, especially those emanating from the United States. MDY in numeric-only form is never used. The ISO 8601 date format (2023-10-9) is the recommended short date format for government publications.[12]	AS ISO 8601-2007
Austria	Yes	Yes	No	(Using dots (which denote ordinal numbering) as in d.m. (yy)yy or sometimes d. month (yy)yy).[13][14]	ÖNORM ISO 8601
Azerbaijan	No	Yes	No	(dd.mm.yyyy)[15]	
Bahamas	No	Yes	No		
Bahrain	No	Yes	No	[16]	

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Bangladesh	No	Yes	No	Not officially standardised. Bengali calendar dates are also used: দদ-মম-বববব	
Barbados	No	Yes	No		BNS 50:2000[17]
Belarus	No	Yes	No	(dd.mm.yyyy)[18][19]	
Belgium	No	Yes	No	(dd/mm/yyyy)[20] or (dd.mm.yyyy)[21][22]	NBN Z 01-002
Belize	No	Yes	No	[23][24]	
Benin	No	Yes	No		
Bermuda	No	Yes	No		
Bhutan	Yes	No	No		
Bolivia	No	Yes	No	[25]	
Bonaire	No	Yes	No		
Bosnia and Herzegovina	No	Yes	No	(d. m. yyyy. or d. mmmm yyyy.)	
Botswana	Yes	Yes	No	yyyy-mm-dd for Setswana and dd/mm/yyyy for English	
Brazil	No	Yes	No	(dd/mm/yyyy)[26][27]	
British Indian Ocean Territory	No	Yes	No		
British Virgin Islands	No	Yes	No		
Brunei	No	Yes	No	[28]	
Bulgaria	No	Yes	No	(dd.mm.yyyy)[29][30]	
Burkina Faso	No	Yes	No		
Burundi	No	Yes	No		
Cambodia	No	Yes	No	Short format: dd/mm/yy Long format: d mmmm yyyy	
Cameroon	Yes	Yes	No	(d)d/(m)m/yyyy or d mmmm yyyy for Aghem, Bafia, Basaa, Duala, English, Ewondo, French, Fula, Kako, Kwasio, Mundang, Ngiemboon and Yangben yyyy-mm-dd for Meta' and Ngomba	
Canada	Yes	Yes	Yes	ISO 8601 is the only format that the Government of Canada and Standards Council of Canada officially recommend for all-numeric dates.[31][32][33] However, usage differs with context.[34][35] All three long forms are used in Canada. For English speakers, MDY (mmmm-dd-yyyy) (example: April 9, 2019) is used by many English-language publications and media company products as well as the majority of government documents written in English.[36]	CAN/CSA-Z234.4-89 (R2007)[37]

Sophos Ex. 1009
Page 30 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				For French and English speakers, DMY (dd-mmmm-yyyy) is used (example: 9 April 2019/le 9 avril 2019). This form is used in formal letters, academic papers, military, many media companies and some government documents, particularly in French-language ones. Federal regulations for shelf life dates on perishable goods mandate a year/month/day format, but allow the month to be written in full, in both official languages, or with a set of standardized two-letter bilingual codes such as 2019 JA 07 or 19 JA 07.	
Cape Verde	No	Yes	No		
Cayman Islands	No	Yes	Yes	DMY and MDY are used interchangeably. Official forms generally tend towards DMY. Month is often spelled out to avoid confusion.	
Central African Republic	No	Yes	No		
Chad	No	Yes	No		
Chile	No	Yes	No	[38] In Chile the format dd/mm/yyyy is used only, or you can also say "3 June 2023" or in Spanish "3 de junio del 2023"You can also use the short format, example "03/06/23".	
China	Yes	No	No	National standard format is yyyy-mm-dd (with leading zeroes) and (yy)yy年(m)m月(d)d日 (with or without leading zeroes)[39] Uyghur languages in Xinjiang usually give date examples in the form يىل 18-ئاۋغۇست-2017 or 2017-8-18 (i.e. yyyy-d-mmm) but this form is never used when writing in Chinese;[40] casually many people use (yy)yy/(m)m/(d)d or (yy)yy.(m)m.(d)d (with or without leading zeroes). See Dates in Chinese.	GB/T 7408-2005
Christmas Island	Yes	Yes	No		
Cocos (Keeling) Islands	Yes	Yes	No		
Colombia	No	Yes	No	[41]	
Comoros	No	Yes	No		
Congo (East and West)	No	Yes	No		
Cook Islands	No	Yes	No		
Costa Rica	No	Yes	No	[42]	
Croatia	No	Yes	No	(d. m. yyyy. or d. mmmm yyyy.)[43][44] See Date and	

Sophos Ex. 1009
Page 31 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				time notation in Croatia for details on cases used.	
Cuba	Yes	Yes	No	[45]	
Curaçao	No	Yes	No		
Cyprus	No	Yes	No	dd/mm/yyyy [46]	
Czech Republic	Yes	Yes	No	(d. m. yyyy or d. month yyyy)[47][48]	ČSN ISO 8601
Denmark	Yes	Yes	No	Examples: Long date: 7. juni 1994. Long date with weekday: onsdag(,) den 21. december 1994. Numeric date: 1994-06-07[49] (The format dd.mm.(yy)yy is the traditional Danish date format.[50] The international format yyyy-mm-dd or yyyymmdd is also accepted, though this format is not commonly used. The formats d. 'month name' yyyy and in handwriting d/m-yy or d/m yyyy are also acceptable.[51])	DS/ISO 8601:2005[52]
Djibouti	Yes	Yes	No	Short format: dd/mm/yyyy (Day first, month number and year in left-to-right writing direction) in Afar, French and Somali ("d/m/yy" is a common alternative). Gregorian dates follow the same rules but tend to be written in the yyyy/m/d format (Day first, month number, and year in right-to-left writing direction) in Arabic language. Long format: d mmmm yyyy or mmmm dd, yyyy (Day first, full month name, and year or first full month name, day, and year, in left-to-right writing direction) in Afar, French and Somali and yyyy ،mmmm d (Day first, full month name, and year in right-to-left writing direction) in Arabic	
Dominica	No	Yes	No		
Dominican Republic	No	Yes	No	[53]	
East Timor	No	Yes	No		
Ecuador	No	Yes	No	[54]	
Egypt	No	Yes	No	[55][56][57]	
El Salvador	No	Yes	No	[58]	
Equatorial Guinea	No	Yes	No	(dd/mm/yyyy or d mmmm yyyy) for French and Spanish	
Eritrea	Yes	Yes	Sometimes	Short format: dd/mm/yyyy for Afar, Bilen, English, Saho, Tigre and Tigrinya. Gregorian dates follow the same rules but tend to be written in the yyyy/m/d (Day first, month number and year in right-to-left writing direction) format in Arabic language. Long format: D MMMM YYYY (Day first, full month name, and year in left-to-right writing direction) for Bilen, English, Tigre and Tigrinya,	

Sophos Ex. 1009
Page 32 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				YYYY ،MMMM D (Day first, full month name, and year in right-to-left writing direction) for Arabic and MMMM DD, YYYY (First full month name, day and year in left-to-right writing direction) for Afar and Saho	
Estonia	Sometimes	Yes	No	dd.mm.yyyy, d.m.(yy)yy or d. mmmm yyyy (mmmm may be substituted by Roman numerals). In more formal, international contexts yyyy-mm-dd is the preferred allowed format.[59]	
Eswatini	Yes	Yes	No	YMD (in Swati), DMY (in English)	
Ethiopia	No	Yes	Sometimes	(dd/mm/yyyy or dd mmmm yyyy) for Amharic, Tigrinya and Wolaytta (dd/mm/yyyy or mmmm dd, yyyy) for Afar, Oromo and Somali[60]	
Falkland Islands	No	Yes	No		
Faroe Islands	No	Yes	No		
Federated States of Micronesia	No	No	Yes	[61]	
Finland	No	Yes	Sometimes	Finnish: d.m.yyyy[62] or in long format d. mmmm yyyy Inari Sami: mmmm d. *p.* yyyy Northern Sami: mmmm d. *b.* yyyy Skolt Sami: mmmm d. *p.* yyyy Swedish: d mmmm yyyy (Note: Month and year can be shortened)	
Fiji	No	Yes	No	[63]	
France	Yes	Yes	No	(dd/mm/yyyy) for Alsatian, Catalan, Corsican, French and Occitan[64][65] (yyyy-mm-dd) for Breton, Basque and Interlingua	NF Z69-200
French Guiana	No	Yes	No		
French Polynesia	No	Yes	No		
Gabon	No	Yes	No		
Gambia	No	Yes	No		
Georgia	No	Yes	No	(dd.mm.yyyy) (In Georgian calendar dates, century digits may be omitted, e.g., dd-mm-yy.)	
Germany	Yes	Yes	No	The format dd.mm.yyyy using dots (which denote ordinal numbering) is the traditional German date format.[66] Since 1996-05-01, the international format yyyy-mm-dd has become the official standard date format, but the handwritten form *d. mmmm yyyy* is also accepted (see DIN 5008). Standardisation applies to all applications in the scope of the standard including uses in government, education, engineering and sciences. Since 2006, the old format (d)d.(m)m.(yy)yy is allowed again as alternative to the yyyy-mm-dd format in areas	DIN ISO 8601:2006-09, used in DIN 5008:2011-04[67]

Sophos Ex. 1009
Page 33 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				where there is no risk of ambiguation. See Date and time notation in Europe.	
Ghana	Yes	Yes	Yes	(yyyy/mm/dd) for Akan (dd/mm/yyyy) (m/d/yyyy) for Ewe	
Gibraltar	No	Yes	No		
Greece	No	Yes	No	[68][69]	ELOT EN 28601
Greenland	No	Yes	Yes	Danish: d. mmmm yyyy Greenlandic: mmmm d.-*at*, yyyy[70]	
Grenada	No	Yes	No		
Guadeloupe	No	Yes	No		
Guam	No	No	Yes		
Guatemala	No	Yes	No	Short format: dd/mm/yyyy Long format: d de mmmm de yyyy or dddd, d de mmmm de yyyy[71]	
Guernsey	No	Yes	No		
Guinea	Yes	Yes	Sometimes	Short format: dd/mm/yyyy (Day first, month and year in left-to-right writing direction) in French and Fulah. Gregorian dates follow the same rules but tend to be written in yyyy/mm/dd (Day first, month number, and year in right-to-left writing direction) format in N'ko language. Long format: D MMMM YYYY (Day first, month and year in left-to-right writing direction) for French and Fulah and YYYY, DD MMMM (First full month name, day, and year in right-to-left writing direction) for N'ko	
Guinea-Bissau	No	Yes	No		
Guyana	No	Yes	No		
Haiti	No	Yes	No		
Hong Kong	Yes	Yes	Rarely	(yy)yy年(m)m月(d)d日 (if without leading zeros) for Chinese[72] and in British English, (d)d/(m)m/(yy)yy in short format. d mmmm yyyy (Casually many people use with commas: d mmmm, yyyy) in long format. Both expanded forms dd-mmmm-yyyy and mmmm-dd-yyyy are used interchangeably in Hong Kong, except the latter was more frequently used in media publications and commercial purpose, such as *The Standard*.	
Honduras	No	Yes	No	[73]	
Hungary	Yes	No	No	yyyy. mm. (d)d.	MSZ ISO 8601:2003

Sophos Ex. 1009
Page 34 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				The year is written in Arabic numerals. The name of the month can be written out in full or abbreviated, or it can be indicated by Roman numerals or Arabic numerals. The day is written in Arabic numerals.[74][75][76][77]	
Iceland	No	Yes	No	(dd.mm.yyyy)[78][79]	IST EN 28601:1992
India	Yes	Yes	Sometimes	In India, the dd-mm-yyyy is the predominant short form of the numeric date usage. Almost all government documents need to be filled up in the dd-mm-yyyy format. An example of dd-mm-yyyy usage is the passport application form.[80][81][82] Though not yet a common practice, the BIS (Bureau of Indian Standards) of the Government of India introduced the standard named "IS 7900:2001 (Revised in 2006) Data Elements And Interchange Formats – Information Interchange – Representation Of Dates And Times" which officially recommends use of the date format yyyy-mm-dd; for example, 2013-09-10, 20130910, or 2013 09 10 for the date 10 September 2013. Dates in the Bodo language are in mm/dd/yyyy. The majority of English-language newspapers and media publications in India use mmmm dd, yyyy.	IS 7900:2001
Indonesia	No	Yes	Rarely	On English-written materials, Indonesians tend to use the M-D-Y but was more widely used in non-governmental contexts. English-language governmental and academic documents use DMY.	
Iran, Islamic Republic of	Yes	Yes	No	Short format: yyyy/mm/dd[83] in Persian Calendar system ("yy/m/d" is a common alternative). Gregorian dates follow the same rules in Persian literature but tend to be written in the dd/mm/yyyy format in official English documents.[84] Long format: YYYY MMMM D (Day first, full month name, and year in right-to-left writing direction)[83]	
Iraq	No	Yes	No	Short format: (dd/mm/yyyy)[85]	
Ireland	No	Yes	No	(dd-mm-yyyy). dd/mm/yyyy is also in common use[86][87]	IS/EN 28601:1993
Isle of Man	No	Yes	No		
Israel	No	Yes	No	The format dd.mm.yyyy using dots is the common format. dd/mm/yyyy is also in common use. The Jewish calendar is in limited use, mainly for Jewish holidays, and follows the DMY format.[88][89][90]	
Italy	No	Yes	No	(dd/mm/yyyy)[91]	UNI EN 28601
Ivory Coast	No	Yes	No		
Jamaica	Yes	Yes	No	[92]	
Jan Mayen	No	Yes	No		
Japan	Yes	No	No	Often in the form yyyy年mm月dd日;[93] sometimes Japanese era year is used, e.g. 平成18年12月30日.[94]	JIS X 0301:2002

Sophos Ex. 1009
Page 35 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Jersey	No	Yes	No		
Jordan	No	Yes	No	[95][96]	
Kazakhstan	Sometimes	Yes	No	Short format: (yyyy.dd.mm) in Kazakh[97] and (dd.mm.(yy)yy) in Russian[98] Long format: yyyy 'ж'. d mmmm in Kazakh;[99] d MMMM yyyy in Russian Full format in Kazakh: yyyy 'ж'. dd mmmm	
Kenya	Yes	Yes	Yes	(yy/mm/dd)[100] (dd/mm/yyyy) (m/d/yyyy) for Swahili[101]	
Kiribati	No	Yes	No		
North Korea	Yes	No	No	[102]	
South Korea	Yes	No	No	National standard format is yyyy-mm-dd (with leading zeroes) and (yy)yy년 (m)m월 (d)d일 (with or without leading zeroes)[39][103] casually many people use (yy)yy.(m)m.(d)d(.) (with or without leading zeroes, with or without the last full stop).	KS X ISO 8601
Kosovo	No	Yes	No		
Kuwait	No	Yes	No	[104]	
Kyrgyz Republic	No	Yes	No	(dd.mm.yyyy)[105]	
Lao People's Democratic Republic	No	Yes	No	[106][107]	
Latvia	No	Yes	No	Short format: dd.mm.yyyy.[108] Long format: yyyy. *gada* d. mmmm	
Lebanon	No	Yes	No	[109]	
Lesotho	Yes	Yes	No	yyyy-mm-dd for Sesotho and dd/mm/yyyy for English	
Liberia	No	Yes	No		
Libya	No	Yes	No	[110]	
Liechtenstein	No	Yes	No	(dd.mm.yyyy)[111]	
Lithuania	Yes	No	No	(yyyy-mm-dd)[112] yyyy <m.> <month in genitive> d <d.>	LST ISO 8601:1997 (obsolete) LST ISO 8601:2006 (current)[113]
Luxembourg	No	Yes	No	(dd.mm.yyyy)[114]	ITM-EN 28601
Macau	Yes	Yes	No	YMD（年月日）	

Sophos Ex. 1009
Page 36 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				(same as Hong Kong)[115] DMY (in Portuguese and British English)	
Madagascar	No	Yes	No		
Malawi	No	Yes	No		
Malaysia	No	Yes	No	dd-mm-yyyy[116]	
Maldives	Yes	Yes	No	Short format: yy/mm/dd (Day first, month next and year last in right-to-left writing direction) Long format: dd mmmm yyyy (Year first, full month name and day last in right-to-left writing direction)	
Mali	No	Yes	No		
Malta	No	Yes	No		
Marshall Islands	No	No	Yes	[117]	
Martinique	No	Yes	No		
Mauritania	No	Yes	No		
Mauritius	No	Yes	No		
Mayotte	No	Yes	No		
Mexico	No	Yes	No	[118]	NOM-008-SCFI-2002
Moldova	No	Yes	No		
Monaco	No	Yes	No	[119]	
Mongolia	Yes	No	No	National standard format is yyyy-mm-dd (with leading zeroes) and yyyy оны (m)m сарын (d)d (with or without leading zeroes) Traditional Mongolian languages in Mongolia usually give date examples in the form 2017ᠣᠨ ᠬᠣᠶᠠᠳᠤᠭᠠᠷ ᠰᠠᠷ᠎ᠠ 2ᠨᠢ but this form is never used when writing in Mongolian Cyrillic; casually many people use yyyy/(m)m/(d)d or yyyy.(m)m.(d)d (with or without leading zeroes).[120]	MNS-ISO 8601
Montenegro	No	Yes	No	Both d.m.yyyy. and dd.mm.yyyy. are accepted. A period is used as a separator and after the year because the Montenegrin language writes these numbers as ordinal numbers that are written as the corresponding cardinal number, with a period at the end.[121]	
Montserrat	No	Yes	No		
Morocco	No	Yes	No	[122]	
Mozambique	No	Yes	No		
Myanmar	Yes	Yes	No	YMD for Burmese calendar. DMY for Gregorian calendar.	

Sophos Ex. 1009
Page 37 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Nagorno-Karabakh Republic	No	Yes	No	(dd.mm.yyyy)[15][8]	
Namibia	Yes	Yes	No	DMY[123]	
Nauru	No	Yes	No		
Nepal	Yes	Yes	Sometimes	DMY, YMD in official Nepali Vikram Samvat calendar (also see Nepal Sambat which is also in use); MDY in Gregorian dates are used for newspapers (English language) and PCs[124]	
Netherlands	No	Yes	No	Using hyphens as in "dd-mm-yyyy".[125]	NEN ISO 8601, NEN EN 28601, NEN 2772
New Caledonia	No	Yes	No		
New Zealand	No	Yes	No	[126]	
Nicaragua	No	Yes	No	[127]	
Niger	No	Yes	No		
Nigeria	No	Yes	Sometimes	Short format: (d)d/(m)m/(yy)yy for Edo, English, Fulani, Hausa, Ibibio, Igbo, Kanuri and Yoruba language[128] Long format: d mmmm yyyy for English, Hausa and Igbo and mmmm dd, yyyy for Edo, Fulani, Ibibio, Kanuri and Yoruba language	
Niue	No	Yes	No	dd/mm/yyyy	
Norfolk Island	No	Yes	No		
North Macedonia	No	Yes	No	(dd.mm.yyyy)[129]	
Northern Mariana Islands	No	No	Yes	[130]	
Norway	Yes	Yes	Rarely	dd.mm.yyyy; leading zeroes and century digits may be omitted, e.g., 10.02.16; ddmmyy (six figures, no century digits, no delimiters) allowed in tables. ISO dates yyyy-mm-dd can be used for "technical" purposes. The fraction form $^{d}/_{m}$-y is incorrect, but is common and considered passable in handwriting. Lule Sami and Southern Sami dates mmmm d. *b.* yyyy.[131]	NS-ISO 8601[132]
Oman	No	Yes	No	[133]	
Pakistan	No	Yes	No		
Palestine (Palestinian Authority, West Bank and Gaza Strip)	No	Yes	No	(dd/mm/yyyy)	
Palau	No	Yes	Rarely	[134] Formerly including: (m)m/(d)d/(yy)yy in English and (yy)yy/m(m)/(d)d in Japanese	
Panama	No	Yes	Yes	Short format: mm/dd/yyyy Long format: d *de* mmmm *de* yyyy[135]	
Papua New Guinea	No	Yes	No		
Paraguay	No	Yes	No	[136]	

Sophos Ex. 1009
Page 38 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Peru	No	Yes	No	[137]	
Philippines	No	Yes	Yes	**Long formats**: English: mmmm d, yyyy DMY dates are also used occasionally, primarily by, but not limited to, government institutions such as on the data page of passports, and immigration and customs forms. Filipino: *ika*-d *ng* mmmm(,) yyyy[138] (Note: Month and year can be shortened. Filipino dates may also be written in mmmm d, yyyy format in civil use but still pronounced as above.) **Short/numerical format**: mm/dd/yyyy for both languages.	
Pitcairn Islands	No	Yes	No		
Poland	Yes	Yes	No	Traditional format (DMY): (dd.mm.yyyy,[139] often with dots as separators; more official is d <month in genitive> yyyy, or, less frequently, d <month in Roman numerals> yyyy)[140][141] Official format (YMD): The ISO 8601 YYYY-MM-DD format is used in official documents, banks, computer systems and the internet in Poland.	PN-90/N-01204
Portugal	Yes	Yes	No	Mostly (dd/mm/yyyy) and (dd-mm-yyyy); some newer documents use (yyyy-mm-dd).[142]	NP EN 28601
Puerto Rico	No	Yes	Yes	English: mmmm d, yyyy Spanish: d *de* mmmm *de* yyyy	
Qatar	No	Yes	No	[143]	
Réunion	No	Yes	No		
Romania	No	Yes	No	(dd.mm.yyyy)[144][145] Also widely used: (d)d-mmm-yyyy (3 letters of month name with the notable exception of Nov for November, which would otherwise be *noiembrie*) and (d)d-XII-yyyy (month number as a Roman numeral with lines above AND below, slowly deprecating)	
Russian Federation	Yes	Yes	No	yyyy-mm-dd dd.mm.yyyy (dd.mm.(yy)yy);[146] more official is d <month in genitive> yyyy г. (= *g.*, short for *goda*, i.e. *year* in genitive) Bashkir, Ossetian, Sakha and Tatar languages in Russia usually give date examples in the form 22 май 2017 й, 22 майы, 2017 аз, ыам ыйын 22 күнэ 2017 с., 22 май 2017 ел but this form is never used when writing in Russian.	GOST R 7.0.64-2018 GOST R 7.0.97-2016
Rwanda	Yes	Yes	No	(yyyy/mm/dd or yyyy mmmm dd) for Kinyarwanda (dd/mm/yyyy or d mmmm yyyy) for English and French	

Sophos Ex. 1009
Page 39 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Saint Barthélemy	No	Yes	No		
Saint Helena, Ascension and Tristan da Cunha	No	Yes	No	[147]	
Saint Kitts and Nevis	No	Yes	No		
Saint Lucia	No	Yes	No		
Saint Martin	No	Yes	No	[148]	
Saint Pierre and Miquelon	No	Yes	No		
Saint Vincent and the Grenadines	No	Yes	No		
Samoa	No	Yes	No		
San Marino	No	Yes	No		
São Tomé and Príncipe	No	Yes	No		
Saudi Arabia	No	Yes	No	(dd/mm/yyyy in Islamic and Gregorian calendar systems,[149][150]	
Senegal	No	Yes	No		
Serbia	No	Yes	No	(d.m.yyyy. or d. mmmm yyyy.)[151][152][153][154]	
Seychelles	No	Yes	No		
Sierra Leone	No	Yes	No		
Singapore	Yes	Yes	Sometimes	(Chinese representation: yyyy年m月d日, no leading zeroes)[155] DMY in English, Malay and Tamil languages[156] MDY (in long format) also sometimes used, especially in media publications, commercial usage, and some governmental websites.	
Sint Eustatius	No	Yes	No		
Sint Maarten	No	Yes	No		
Slovakia	No	Yes	No	(d. m. yyyy)[157]	
Slovenia	No	Yes	No	(d. m. yyyy or d. mmmm yyyy)[158]	
Solomon Islands	No	Yes	No		
Somalia	No	Yes	No	Short format: dd/mm/yyyy	
South Africa	Yes	Yes	Sometimes	(yyyy/mm/dd and dd mmmm yyyy) in English (yyyy-mm-dd and dd mmmm yyyy) in Afrikaans[159] (yyyy-mm-dd and yyyy mmmm d) in Xhosa[160] MDY in Zulu[161]	SANS 8601:2009[162]

Sophos Ex. 1009
Page 40 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Spain	Yes	Yes	No	(dd/mm/yyyy) for Asturian, Catalan, Galician, Spanish and Valencian[163] (yyyy/mm/dd) for Basque[164]	UNE EN 28601
Sri Lanka	Yes	Yes	Rarely	(yyyy-mm-dd) for Sinhala and (d-m-yyyy) for Tamil English-language media and commercial publications use Month-day-year in long format, but only Day-month-year format (both long and short numeric) are used in governmental and other English documents of official contexts.	
Sudan	No	Yes	No		
South Sudan	No	Yes	No		
Suriname	No	Yes	No		
Svalbard	No	Yes	No		
Sweden	Yes	Yes	No	National standard format is yyyy-mm-dd.[165] dd.mm.yyyy format is used in some places where it is required by EU regulations, for example for best-before dates on food[166] and on driver's licenses. d/m format is used casually, when the year is obvious from the context, and for date ranges, e.g. 28-31/8 for 28–31 August. The textual format is "d mmmm yyyy" or "den d mmmm yyyy".	SS-ISO 8601
Switzerland	No	Yes	No	(dd.mm.yyyy or d. mmmm yyyy) for French, German, Italian and Romansh[167][168]	SN ISO 8601:2005-08
Syrian Arab Republic	No	Yes	No	[169]	
Taiwan	Yes	No	No	Short format: yyyy/(m)m/(d)d[170] or yyyy-mm-dd[171] Long format: yyyy年m月d日, in most context year is represented using ROC era system: 民國95年12月30日.[172]	CNS 7648
Tajikistan	No	Yes	No	(dd.mm.yyyy)[173]	
Tanzania	No	Yes	No		
Thailand	No	Yes	No	dd/mm/yyyy (in governmental sector with Buddhist Era years instead of Common Era)[174]	TIS 1111:2535 in 1992
Togo	No	Yes	Yes	(dd/mm/yyyy) in French and (mm/dd/(yy)yy) in Ewe	
Tokelau	No	Yes	No		
Tonga	No	Yes	No		
Trinidad and Tobago	No	Yes	No	[175]	
Tunisia	No	Yes	No	[176]	

Sophos Ex. 1009
Page 41 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
Turkey	No	Yes	No	Short format: dd.mm.yyyy[177][178] Long format: d mmmm yyyy Full format: d mmmm yyyy dddd[179]	
Turkmenistan	No	Yes	No	(dd.mm.(yy)yy *ý.*), yyyy-*nji ýylyň* d-*nji* mmmm[180][181]	
Turks and Caicos Islands	No	Yes	No		
Tuvalu	No	Yes	No		
Uganda	No	Yes	No		
Ukraine	No	Yes	No	(dd.mm.(yy)yy;[182][183] some cases of dd/mm/yyyy[184])	
United Arab Emirates	No	Yes	No	[185][186]	
United Kingdom	Yes	Yes	No	Most style guides follow the DMY convention by recommending d mmmm yyyy (sometimes written dd/mm/yyyy) format in articles (e.g. *The Guardian*'s, and the *Oxford Style Manual*).[187][188] Some newspapers use dddd mmmm d, yyyy for both the banner and articles,[189] while others stick to DMY for both.[190] In addition, YMD with four-digit year is used increasingly especially in applications associated with computers, and as per British standard BS ISO 8601:2004,[191] avoiding the ambiguity of the numerical versions of the DMY/MDY formats.	BS ISO 8601:2004
United States Minor Outlying Islands	No	No	Yes	Same as the US	
United States of America	Yes	Rarely	Yes	(Civilian vernacular: m/d/yy or m/d/yyyy;[192][193] other formats, especially d mmm(m) yyyy (but no short DMY formats) and yyyy-mm-dd (but rarely any other short YMD formats and rarely any long YMD formats), are sometimes prescribed or used—particularly in military, academic, scientific, computing, industrial, or governmental contexts. See Date and time notation in the United States.)	ANSI INCITS 30-1997 (R2008) and NIST FIPS PUB 4-2
United States Virgin Islands	No	No	Yes	[194]	
Uruguay	No	Yes	No	[195][196]	
Uzbekistan	Yes	Yes	No	(dd.mm.yyyy Cyrillic, dd/mm yyyy Latin)[197][198][199]	
Vanuatu	No	Yes	No		
Venezuela, Bolivarian Republic of	No	Yes	No	[200][201][202]	
Vietnam	Yes	Yes	Sometimes	Long format: "Ngày (d)d tháng (m)m năm yyyy" (leading zeros required by Circular No. 01/2011/TT-BNV by the Ministry of Home Affairs)[203] or ngày (d)d tháng (month in textform) năm yyyy.	

Sophos Ex. 1009
Page 42 of 318

Country	All-numeric date format			Details	ISO 8601
	YMD	DMY	MDY		
				Short format (interchangeably): (d)d/(m)m/yyyy or (d)d-(m)m-yyyy; (d)d.(m)m.yyyy is also in use.[204] In English documents: ▪ Short format: yyyy-mm-dd[205] ▪ Long format: mmmm d, yyyy[206] In historical documents: era names năm thứ _ tháng [m]m (or in textform) ngày(mồng) [d]d (or in textform).	
Wallis and Futuna	No	Yes	No		
Yemen	No	Yes	No	[207][208]	
Zambia	No	Yes	No		
Zimbabwe	Yes	No	No	[209]	

See also

- Date and time representation by country
- Common Locale Data Repository, a database that covers national date and time notations
- ISO 8601

References

1. "NLS information page – Albanian (Albania)" (https://web.archive.org/web/20090112192544/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=041C&OS=Windows%2BVista). Microsoft. Archived from the original (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=041C&OS=Windows%20Vista) on 2009-01-12. Retrieved 2008-10-21.
2. (in Albanian) Municipality of Tirana (Bashkia e Tiranës) (http://www.tirana.gov.al/) Archived (https://web.archive.org/web/20110414082050/http://www.tirana.gov.al/) 2011-04-14 at the Wayback Machine
3. "Globalization Library – Locale Data: Albania" (ftp://ftp.software.ibm.com/software/globalization/locales/Albania-Albanian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
4. "Globalization Library – Locale Data: Algeria" (ftp://ftp.software.ibm.com/software/globalization/locales/Algeria-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
5. "NLS information page – Arabic (Algeria)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=1401&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-21.
6. "Spanish (Argentina) (es-AR)" (https://www.ibm.com/support/knowledgecenter/en/SSS28S_3.0.0/com.ibm.help.forms.doc/locale_spec/i_xfdl_r_formats_es_AR.html). IBM. Archived (https://web.archive.org/web/20191031062623/https://www.ibm.com/support/knowledgecenter/en/SSS28S_3.0.0/com.ibm.help.forms.doc/locale_spec/i_xfdl_r_formats_es_AR.html) from the original on 2019-10-31. Retrieved 2019-10-21.
7. (in Armenian) National Assembly of Armenia (http://www.parliament.am/) Archived (https://web.archive.org/web/20090331131955/http://www.parliament.am/) 2009-03-31 at the Wayback Machine
8. "NLS information page – Armenian (Armenia)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=042B&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-21.

9. Overheid, Aruba. "News" (https://www.government.aw/news/news_47033/). *www.government.aw*. Archived (https://web.archive.org/web/20210510025148/https://www.government.aw/news/news_47033/) from the original on 2021-05-10. Retrieved 2021-05-10.
10. "Latest News" (http://www.newscorpaustralia.com/news). *News Corp Australia*. Archived (https://web.archive.org/web/20171104025041/http://www.newscorpaustralia.com/news) from the original on 2017-11-04. Retrieved 2017-11-04.
11. "The West Australian Demo" (http://enewspaper2.smedia.com.au/wandemo/). *The West Australian*. 2016-08-16. Archived (https://web.archive.org/web/20170901045345/http://enewspaper2.smedia.com.au/wandemo/) from the original on 2017-09-01. Retrieved 2017-11-04.
12. *Style manual for authors, editors and printers* (6 ed.). John Wiley & Sons Australia. 2002. p. 171. ISBN 978-0-7016-3647-0.
13. (in German) Federal Chancellery of Austria (Bundeskanzleramt) (http://www.bka.gv.at/site/1/default.aspx) Archived (https://web.archive.org/web/20110115071554/http://www.bka.gv.at/site/1/Default.aspx) 2011-01-15 at the Wayback Machine
14. "NLS information page – German (Austria)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0C07&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-21.
15. "NLS information page – Azeri (Latin, Azerbaijan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=042C&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-21., also NLS information page – Azeri (Cyrillic, Azerbaijan) (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=082C&OS=Windows%20Vista) Archived (https://web.archive.org/web/20220514044548/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=082C&OS=Windows+Vista) 2022-05-14 at the Wayback Machine
16. "Globalization Library – Locale Data: Bahrain" (ftp://ftp.software.ibm.com/software/globalization/locales/Bahrain-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
17. "The International system of units (SI) and its practical application including the designation of times and dates" (https://web.archive.org/web/20120425072331/http://www.bnsi.bb/?q=content%2Finternational-system-units-si-and-its-practical-application-including-designation-times-and-). *Metrication Board*. Barbados National Standard Institute (BNSI). 2000. Archived from the original (http://www.bnsi.bb/?q=content/international-system-units-si-and-its-practical-application-including-designation-times-and-) on 2012-04-25. Retrieved 2011-10-16.
18. (in Belarusian) President of the Republic of Belarus (Прэзідэнта Рэспублікі Беларусь) (http://www.president.gov.by/by/) Archived (https://web.archive.org/web/20101218230711/http://president.gov.by/by/) 2010-12-18 at the Wayback Machine
19. "Globalization Library – Locale Data: Belarus" (ftp://ftp.software.ibm.com/software/globalization/locales/Belarus-Belarusian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
20. Belgian Federal Government – Portal (http://www.belgium.be/) Archived (https://web.archive.org/web/20161107223001/http://www.belgium.be/) 2016-11-07 at the Wayback Machine
21. "Globalization Library – Locale Data: Belgium (Dutch, English, French)" (http://www.ibm.com/support/knowledgecenter/SSS28S_8.1.0/XFDL/i_xfdl_r_formats_fr_BE.dita). IBM. Archived (https://web.archive.org/web/20160320023812/http://www.ibm.com/support/knowledgecenter/SSS28S_8.1.0/XFDL/i_xfdl_r_formats_fr_BE.dita) from the original on 2016-03-20. Retrieved 2016-03-17.
22. "datumnotatie" (https://www.taaltelefoon.be/datumnotatie) (in Dutch). Belgium: Taal Telefoon. Archived (https://web.archive.org/web/20170522071431/https://www.taaltelefoon.be/datumnotatie) from the original on 2017-05-22. Retrieved 2017-06-01.
23. "National Language Support (NLS) API Reference -- Belize" (https://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=2809&OS=Windows%20Vista). Microsoft. Retrieved 2016-05-11.
24. "Her Majesty's Government of Belize" (https://web.archive.org/web/20110208193548/http://www.belize.gov.bz/). Government of Belize. Archived from the original (http://www.belize.gov.bz) on 2011-02-08. Retrieved 2017-05-18.
25. "Globalization Library – Locale Data: Bolivia" (ftp://ftp.software.ibm.com/software/globalization/locales/Bolivia-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
26. (in Portuguese) Government of Brazil – Notícias (News) (http://www.brasil.gov.br/navegue_por/noticias/) Archived (https://web.archive.org/web/20110821002142/http://www.brasil.gov.br/navegue_por/noticias) 2011-08-21 at the Wayback Machine

27. "Globalization Library – Locale Data: Brazil" (ftp://ftp.software.ibm.com/software/globalization/locales/Brazil-Portuguese_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
28. "NLS Information for Brunei Darussalam" (http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=083E&OS=Windows%207). Microsoft. Archived (https://web.archive.org/web/20160201110338/http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=083E&OS=Windows%207) from the original on 2016-02-01. Retrieved 2014-12-11.
29. (in Bulgarian) Parliament of Bulgaria (http://www.parliament.bg/) Archived (https://web.archive.org/web/20050624002434/http://www.parliament.bg/) 2005-06-24 at the Wayback Machine
30. "Globalization Library – Locale Data: Bulgaria" (ftp://ftp.software.ibm.com/software/globalization/locales/Bulgaria-Bulgarian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
31. Collishaw, Barbara (2002). "FAQs on Writing the Date" (https://www.btb.termiumplus.gc.ca/tpv2guides/guides/favart/index-eng.html?lang=eng&lettr=indx_titls&page=9geHWsOpC7ZI.html). *Terminology Update*. **35** (2): 12. Archived (https://web.archive.org/web/20180808171827/http://www.btb.termiumplus.gc.ca/tpv2guides/guides/favart/index-eng.html?lang=eng&lettr=indx_titls&page=9geHWsOpC7ZI.html) from the original on 2018-08-08. Retrieved 2018-10-30.
32. "Getting on the Same Page When It Comes to Date and Time" (https://www.scc.ca/en/news-events/news/2018/getting-same-page-when-it-comes-date-and-time). *Standards Council of Canada*. 2018-01-11. Archived (https://web.archive.org/web/20180330075920/https://www.scc.ca/en/news-events/news/2018/getting-same-page-when-it-comes-date-and-time) from the original on 2018-03-30. Retrieved 2018-03-29.
33. "TBITS 36: All-Numeric Representation of Dates and Times – Implementation Criteria" (http://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=17284&section=text). Treasury Board of Canada. 1997-12-18. Archived (https://web.archive.org/web/20120319150729/http://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=17284&section=text) from the original on 2012-03-19. Retrieved 2012-03-17.
34. "Cheque Specifications" (https://www.payments.ca/sites/default/files/standard_006_complete.pdf) (PDF). Canadian Payments Association. Archived (https://web.archive.org/web/20170925181137/https://www.payments.ca/sites/default/files/standard_006_complete.pdf) (PDF) from the original on 25 September 2017. Retrieved 17 March 2012.
35. Blaze Carlson, Kathryn (29 October 2011). "Is 02/04/12 February 4, or April 2? Bill seeks to end date confusion" (https://news.nationalpost.com/news/canada/is-020412-february-4-or-april-2-bill-seeks-to-end-date-confusion). *National Post*. Retrieved 25 September 2017.
36. Collishaw, Barbara (2002). "FAQs on Writing the Date" (https://www.noslangues-ourlanguages.gc.ca/en/favourite-articles/faqs-on-writing-the-date). *Terminology Update*. Government of Canada. p. 12.
37. "CAN/CSA-Z234.4-89 (R2007): All-Numeric Dates and Times" (https://www.scc.ca/en/standardsdb/standards/4449). Standards Council of Canada. 1989-12-31. Archived (https://web.archive.org/web/20180330075843/https://www.scc.ca/en/standardsdb/standards/4449) from the original on 2018-03-30. Retrieved 2018-03-29.
38. "Globalization Library – Locale Data: Chile" (ftp://ftp.software.ibm.com/software/globalization/locales/Chile-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
39. "Date Format in the United States | ISO" (https://iso.mit.edu/americanisms/date-format-in-the-united-states/). Archived (https://web.archive.org/web/20211121063731/https://iso.mit.edu/americanisms/date-format-in-the-united-states/) from the original on 21 November 2021. Retrieved 21 November 2021.
40. "中华人民共和国国家标准 GB/T 7408-2005/ISO 8601:2000 数据元和交换格式 信息交换 日期和时间表示法" (http://www.doc88.com/p-990319057545.html). Archived (https://web.archive.org/web/20170216062628/http://www.doc88.com/p-990319057545.html) from the original on 2017-02-16. Retrieved 2017-02-15.
41. "Globalization Library – Locale Data: Colombia" (ftp://ftp.software.ibm.com/software/globalization/locales/Colombia-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
42. "Globalization Library – Locale Data: Costa Rica" (ftp://ftp.software.ibm.com/software/globalization/locales/CostaRica-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
43. (in Croatian) Government of the Republic of Croatia (http://www.vlada.hr/hr/aktualne_teme_i_projekti) Archived (https://web.archive.org/web/20101225195720/http://vlada.hr/hr/aktualne_teme_i_projekti) 2010-12-25 at the Wayback Machine
44. "Globalization Library – Locale Data: Croatia" (ftp://ftp.software.ibm.com/software/globalization/locales/Croatia-Croatian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
45. "Inicio" (http://www.cubanews.acn.cu/). Cuban News Agency. Archived (https://web.archive.org/web/20160321231029/http://www.cubanews.acn.cu/) from the original on 2016-03-21. Retrieved 2016-03-20.

46. "tax department" (http://www.mof.gov.cy/mof/taxdep.nsf/index_en/index_en?OpenDocument/). Cyprus Tax department. Retrieved 2016-03-23.
47. "Internetová jazyková příručka Ústavu pro jazyk český" (https://prirucka.ujc.cas.cz/?id=810). Archived (https://web.archive.org/web/20210508051349/https://prirucka.ujc.cas.cz/?id=810) from the original on 2021-05-08. Retrieved 2021-05-14.
48. "Radio Praha – zprávy z České republiky" (http://www.radio.cz/cz/). *radio.cz*. Archived (https://web.archive.org/web/20110104142740/http://www.radio.cz/cz) from the original on 2011-01-04.
49. "Danish language locale for Denmark, Narrative Cultural Specification" (http://std.dkuug.dk/cultreg/registrations/narrative/da_DK,_4.3.html). *dkuug.dk*. Archived (https://web.archive.org/web/20150426151945/http://std.dkuug.dk/cultreg/registrations/narrative/da_DK,_4.3.html) from the original on 2015-04-26.
50. "Globalization Library – Locale Data: Denmark" (ftp://ftp.software.ibm.com/software/globalization/locales/Denmark-Danish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
51. "datoer" (http://sproget.dk/raad-og-regler/artikler-mv/svarbase/SV00000046). *sproget.dk*. Archived (https://web.archive.org/web/20130912191240/http://sproget.dk/raad-og-regler/artikler-mv/svarbase/SV00000046) from the original on 2013-09-12.
52. "DS/ISO 8601:2005" (https://web.archive.org/web/20161104050130/http://webshop.ds.dk/da-dk/s%C3%B8gning/01-140-30-dokumenter-inden-for-administration--handel-og-industri/ds-iso-86012005). Denmark: Dansk Standard. 2005-10-26. Archived from the original (http://webshop.ds.dk/da-dk/s%C3%B8gning/01-140-30-dokumenter-inden-for-administration--handel-og-industri/ds-iso-86012005) on 2016-11-04. Retrieved 2016-11-12.
53. "Globalization Library – Locale Data: Dominican Republic" (ftp://ftp.software.ibm.com/software/globalization/locales/DominicanRepublic-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
54. "Globalization Library – Locale Data: Ecuador" (ftp://ftp.software.ibm.com/software/globalization/locales/Ecuador-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
55. (in Arabic) [Arab Republic of Egypt: Cabinet/Office of the Prime Minister]
56. "Globalization Library – Locale Data: Egypt" (ftp://ftp.software.ibm.com/software/globalization/locales/Egypt-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
57. "ICU Demonstration - Locale Explorer (ar_EG)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ar_EG). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035820/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ar_EG) from the original on 2022-03-18. Retrieved 2022-03-18.
58. "Globalization Library – Locale Data: El Salvador" (ftp://ftp.software.ibm.com/software/globalization/locales/ElSalvador-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
59. "Õigekirjutus. Numbrite kirjutamine" (http://www.eki.ee/books/ekk09/index.php?p=2&p1=10) (in Estonian). Eesti Keele Instituut. Archived (https://web.archive.org/web/20141007202242/http://www.eki.ee/books/ekk09/index.php?p=2&p1=10) from the original on 2014-10-07. Retrieved 2013-11-10.
60. "NLS Information for Ethiopia" (http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=045E&OS=Windows%207). Microsoft. Archived (https://web.archive.org/web/20160201110338/http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=045E&OS=Windows%207) from the original on 2016-02-01. Retrieved 2014-12-11.
61. "Global Sourcebook for International Data Management – Micronesia" (https://www.grcdi.nl/gsb/micronesia.html). Archived (https://web.archive.org/web/20191212162626/https://www.grcdi.nl/gsb/micronesia.html) from the original on 12 December 2019. Retrieved 13 December 2019.
62. Ajanilmaukset (http://www.kielikello.fi/index.php?mid=2&pid=11&aid=1709) Archived (https://web.archive.org/web/20171020033758/http://www.kielikello.fi/index.php?mid=2&pid=11&aid=1709) 2017-10-20 at the Wayback Machine Kielikello 2/2006. Institute for the Languages of Finland. Retrieved 2017-10-20
63. "ICU Demonstration - Locale Explorer (en_FJ)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=en_FJ). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035821/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=en_FJ) from the original on 2022-03-18. Retrieved 2022-03-18.
64. (in French) Government of France – Prime Minister (http://www.premier-ministre.gouv.fr/fr/) Archived (https://web.archive.org/web/20110120220408/http://www.premier-ministre.gouv.fr/fr/) 2011-01-20 at the Wayback Machine
65. "Globalization Library – Locale Data: France" (ftp://ftp.software.ibm.com/software/globalization/locales/France-French_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.

Sophos Ex. 1009
Page 46 of 318

66. "Globalization Library – Locale Data: Germany" (ftp://ftp.software.ibm.com/software/globalization/locales/Germany-German_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
67. de:Datumsformat#ISO 8601 und EN 28 601, Retrieved 2010-03-10
68. (in Greek) Hellenic Parliament (Greece) (http://www.parliament.gr/default.asp) Archived (https://web.archive.org/web/20100207225851/http://www.parliament.gr/default.asp) 2010-02-07 at the Wayback Machine
69. "Globalization Library – Locale Data: Greece" (ftp://ftp.software.ibm.com/software/globalization/locales/Greece-Greek_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
70. Oqaasileriffik/Greenlandic Language Secretariat
71. "Globalization Library – Locale Data: Guatemala" (ftp://ftp.software.ibm.com/software/globalization/locales/Guatemala-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
72. "Globalization Library – Chinese (Traditional Han, Hong Kong SAR China) (zh-Hant-HK)" (http://publib.boulder.ibm.com/infocenter/forms/v3r5m1/topic/com.ibm.form.designer.locales.doc/i_xfdl_r_formats_zh_Hant_HK.html). IBM. Retrieved 2012-03-29.
73. "Globalization Library – Locale Data: Honduras" (ftp://ftp.software.ibm.com/software/globalization/locales/Honduras-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
74. "A magyar helyesírás szabályai" (https://helyesiras.mta.hu/helyesiras/default/akh12). MTA. 2015. Archived (https://web.archive.org/web/20180414080519/http://helyesiras.mta.hu/helyesiras/default/akh12) from the original on 2018-04-14. Retrieved 2018-04-14.
75. "Globalization Library – Locale Data: Hungary" (ftp://ftp.software.ibm.com/software/globalization/locales/Hungary-Hungarian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
76. "NLS information page – Hungarian (Hungary)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=040E&OS=Windows%20Vista). Microsoft. Retrieved 2009-02-03.
77. Akadémia, Magyar Tudományos (2015). *A magyar helyesírás szabályai* (12. ed.). Akadémiai Kiadó. ISBN 978-963-05-8630-6.
78. (in Icelandic) Government of Iceland (Stjórnarráðið)) (http://www.stjornarrad.is/) Archived (https://web.archive.org/web/20110722022116/http://www.stjornarrad.is/) 2011-07-22 at the Wayback Machine
79. "Globalization Library – Locale Data: Iceland" (ftp://ftp.software.ibm.com/software/globalization/locales/Iceland-Icelandic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
80. "Standards Published" (https://web.archive.org/web/20070901205604/http://bis.org.in/sf/pow1.asp?d=MS&c=005&x=DOCUMENTATION+AND+INFORMATION+SECTIONAL+COMMITTEE). Bureau of Indian Standards. Archived from the original (http://www.bis.org.in/sf/pow1.asp?d=MS&c=005&x=DOCUMENTATION+AND+INFORMATION+SECTIONAL+COMMITTEE) on 2007-09-01. Retrieved 2008-09-20.
81. "President of India" (http://india.gov.in/outerwin.htm?id=http%3A%2F%2Fwww.presidentofindia.nic.in%2F). Archived (https://web.archive.org/web/20081219231347/http://india.gov.in/outerwin.htm?id=http%3A%2F%2Fwww.presidentofindia.nic.in%2F) from the original on 2008-12-19. Retrieved 2022-05-14.
82. "Globalization Library – Locale Data (Pacific Region)" (https://web.archive.org/web/20090226232539/http://www-01.ibm.com/software/globalization/locales/locales_pr.jsp). IBM. Archived from the original (http://www-01.ibm.com/software/globalization/locales/locales_pr.jsp) on 2009-02-26. Retrieved 2008-10-12.
83. بانک مرکزی ایران (http://www.cbi.ir/default.aspx) (in Persian). The Central Bank of the Islamic Republic of Iran. Archived (https://web.archive.org/web/20090421214557/http://www.cbi.ir/default.aspx) from the original on 2009-04-21. Retrieved 2009-04-30.
84. Iranian Passport Datapage.jpg, Retrieved 2016-10-01
85. "NLS information page – Arabic (Iraq)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0801&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-21.
86. "Home" (http://www.taoiseach.gov.ie/index.asp). *taoiseach.gov.ie*. Archived (https://web.archive.org/web/20090526054501/http://www.taoiseach.gov.ie/index.asp) from the original on 2009-05-26.
87. "Globalization Library – Locale Data: Ireland" (ftp://ftp.software.ibm.com/software/globalization/locales/Ireland-English_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
88. *Jerusalem Post* (http://www.jpost.com/) Archived (https://web.archive.org/web/20080519190250/http://www.jpost.com/) 2008-05-19 at the Wayback Machine
89. *Times of Israel* (https://www.timesofisrael.com/) Archived (https://web.archive.org/web/20161023180015/http://www.timesofisrael.com/) 2016-10-23 at the Wayback Machine

90. Israel Government Portal (http://www.gov.il/firstgov) Archived (https://web.archive.org/web/20110120142636/http://www.gov.il/FirstGov/) 2011-01-20 at the Wayback Machine, English translation
91. "Globalization Library – Locale Data: Italy" (ftp://ftp.software.ibm.com/software/globalization/locales/Italy-Italian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
92. US Commercial Service (2008-12-23). "Doing Business in JAMAICA: A Country Commercial Guide for U.S. Companies" (http://ctrc.sice.oas.org/TRC/Articles/Jamaica/Commercial_Guide_2009.pdf) (PDF). Organization of American States. p. 25. Archived (https://web.archive.org/web/20140407065423/http://ctrc.sice.oas.org/TRC/Articles/Jamaica/Commercial_Guide_2009.pdf) (PDF) from the original on 2014-04-07. Retrieved 2014-04-01.
93. "ICU Demonstration - Locale Explorer (ja_JP)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ja_JP). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220514044546/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ja_JP) from the original on 2022-05-14. Retrieved 2022-03-18.
94. "Prime Minister of Japan and His Cabinet" (http://www.kantei.go.jp/). Archived (https://web.archive.org/web/20080913193311/http://www.kantei.go.jp/) from the original on 2008-09-13. Retrieved 2008-10-12.
95. "Globalization Library – Locale Data: Jordan" (ftp://ftp.software.ibm.com/software/globalization/locales/Jordan-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
96. "NLS information page – Arabic (Jordan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=2C01&OS=Windows%20Vista). Microsoft. Retrieved 2009-02-08.
97. https://adilet.zan.kz/kaz/docs/P1100001570 Archived (https://web.archive.org/web/20210709182454/https://adilet.zan.kz/kaz/docs/P1100001570) 2021-07-09 at the Wayback Machine Section 25
98. "Official rules of documenting in governmental authorities" (https://adilet.zan.kz/rus/docs/P1100001570). Government of Kazakhstan (in Kazakh and Russian). Archived (https://web.archive.org/web/20141201180727/http://adilet.zan.kz/rus/docs/P1100001570/compare) from the original on 2014-12-01.
99. "DateTime::Locale::kk_KZ - Locale data examples for the Kazakh Kazakhstan (kk-KZ) locale - metacpan.org" (https://metacpan.org/pod/distribution/DateTime-Locale/lib/DateTime/Locale/kk_KZ.pod). *metacpan.org*. Archived (https://web.archive.org/web/20210510144406/https://metacpan.org/pod/distribution/DateTime-Locale/lib/DateTime/Locale/kk_KZ.pod) from the original on 2021-05-10. Retrieved 2021-05-09.
100. "Date Time Local" (http://kobesearch.cpan.org/htdocs/DateTime-Locale/sw_KE.html#Short). Retrieved 2009-08-31.
101. "NLS Information for Windows 7 – Kiswahili (Kenya)" (https://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=0441&OS=Windows%207). *National Language Support (NLS) API Reference*. Microsoft. Archived (https://web.archive.org/web/20171027041222/https://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=0441&OS=Windows%207) from the original on 2017-10-27. Retrieved 2017-08-15.
102. "ICU Demonstration - Locale Explorer (ko_KP)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ko_KP). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220514044544/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=ko_KP) from the original on 2022-05-14. Retrieved 2022-03-18.
103. "Globalization Library – Locale Data: Korea (Simplified Chinese)" (https://web.archive.org/web/20090112210604/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0412&OS=Windows%2BVista). IBM. Archived from the original (ftp://ftp.software.ibm.com/software/globalization/locales/Korea-Korean_Date.pdf) (PDF) on 2009-01-12. Retrieved 2008-10-12.
104. "Globalization Library – Locale Data: Kuwait" (ftp://ftp.software.ibm.com/software/globalization/locales/Kuwait-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
105. "NLS information page – Kyrgyz (Kyrgyzstan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0440&OS=Windows%20Vista). Microsoft. Retrieved 2009-02-08.
106. "NLS Information for Laos" (http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=0454&OS=Windows%207). Microsoft. Archived (https://web.archive.org/web/20160201110338/http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=0454&OS=Windows%207) from the original on 2016-02-01. Retrieved 2014-12-11.
107. "ICU Demonstration - Locale Explorer (lo_LA)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=lo_LA). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035820/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=lo_LA) from the original on 2022-03-18. Retrieved 2022-03-18.

108. "Dokumentu izstrādāšanas un noformēšanas kārtība" (https://likumi.lv/doc.php?id=219491). Latvijas Vēstnesis. Archived (https://web.archive.org/web/20171230060433/https://likumi.lv/doc.php?id=219491) from the original on 2017-12-30. Retrieved 2017-12-29.
109. "Globalization Library – Locale Data: Lebanon" (ftp://ftp.software.ibm.com/software/globalization/locales/Lebanon-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
110. "NLS information page – Arabic (Libya)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=1001&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-29.
111. "NLS Information for Liechtenstein" (http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=1407&OS=Windows%207). Microsoft. Archived (https://web.archive.org/web/20160201110338/http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=1407&OS=Windows%207) from the original on 2016-02-01. Retrieved 2014-12-11.
112. "Globalization Library – Locale Data: Lithuania" (ftp://ftp.software.ibm.com/software/globalization/locales/Lithuania-Lithuanian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
113. "Lietuvos standartizacijos departamentas – el. parduotuvė" (http://www.lsd.lt/standards/catalog.php?ics=0&pid=626252). Archived (https://web.archive.org/web/20160201110338/http://www.lsd.lt/standards/catalog.php?ics=0&pid=626252) from the original on 2016-02-01. Retrieved 2015-12-15.
114. "Globalization Library – Locale Data: Luxembourg (German)" (ftp://ftp.software.ibm.com/software/globalization/locales/Luxembourg-German_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
115. "Government News" (https://web.archive.org/web/20080912124045/http://www.gov.mo/egi/Portal/rkw/public/view/area.jsp?id=21) (in Chinese). Macao SAR Government Portal. Archived from the original (http://www.gov.mo/egi/Portal/rkw/public/view/area.jsp?id=21) on 2008-09-12. Retrieved 2008-10-13.
116. *Malaysia: Doing Business, Investing in Malaysia Guide Volume 1 Strategic, Practical Information, Regulations, Contacts* (https://books.google.com/books?id=ftqtDwAAQBAJ&dq=ddmmyyyy+date+format+malaysia&pg=PA11). Washington D.C.: International Business Publications. 2017. p. 11. ISBN 9781438713168.
117. "RMI" (http://www.rmiembassyus.org). Embassy of the Republic of the Marshall Islands Washington, DC. Archived (https://web.archive.org/web/20170611024125/http://rmiembassyus.org/) from the original on 2017-06-11. Retrieved 2017-06-12.
118. "Globalization Library – Locale Data: Mexico" (ftp://ftp.software.ibm.com/software/globalization/locales/Mexico-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
119. "NLS Information for Monaco" (http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=180C&OS=Windows%207). Microsoft. Archived (https://web.archive.org/web/20160201110338/http://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=180C&OS=Windows%207) from the original on 2016-02-01. Retrieved 2014-12-11.
120. "Монгол Улсын Ерөнхийлөгч" (https://web.archive.org/web/20081014005530/http://www.president.mn/index.php) (in Mongolian). President of Mongolia. Archived from the original (http://www.president.mn/index.php) on 2008-10-14. Retrieved 2008-10-12.
121. "Vlada Crne Gore (Government of Montenegro)" (http://www.gov.me/) (in Montenegrin). Archived (https://web.archive.org/web/20090204025310/http://www.gov.me/) from the original on 2009-02-04. Retrieved 2009-02-08.
122. "Globalization Library – Locale Data: Morocco" (ftp://ftp.software.ibm.com/software/globalization/locales/Morocco-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
123. "Archived copy" (http://www.gov.na/documents/10181/12843/Namibian_Citizenship_Act_No.pdf/d5541c06-16de-4410-8c5a-4b7e2c945a0f). Archived (https://web.archive.org/web/20150402142207/http://www.gov.na/documents/10181/12843/Namibian_Citizenship_Act_No.pdf/d5541c06-16de-4410-8c5a-4b7e2c945a0f) from the original on 2015-04-02. Retrieved 2015-03-11.
124. "Public opinion necessary on contempt of court bill: Nembang | Top Stories" (http://www.ekantipur.com/2014/09/08/top-story/public-opinion-necessary-on-contempt-of-court-bill-nembang/394759.html). Archived (https://web.archive.org/web/20140908223236/http://www.ekantipur.com/2014/09/08/top-story/public-opinion-necessary-on-contempt-of-court-bill-nembang/394759.html) from the original on 2014-09-08. Retrieved 2014-09-08.
125. "Globalization Library – Locale Data: Netherlands" (ftp://ftp.software.ibm.com/software/globalization/locales/Netherlands-Dutch_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
126. "Globalization Library – Locale Data: New Zealand" (ftp://ftp.software.ibm.com/software/globalization/locales/NewZealand-English_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.

127. "Globalization Library – Locale Data: Nicaragua" (ftp://ftp.software.ibm.com/software/globalization/locales/Nicaragua-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
128. "OTHER LINKS ON NIGERIA" (http://www.motherlandnigeria.com/other_links.html). Archived (https://web.archive.org/web/20140226021828/http://www.motherlandnigeria.com/other_links.html) from the original on 2014-02-26. Retrieved 2014-02-14.
129. "Globalization Library – Locale Data: Macedonia" (ftp://ftp.software.ibm.com/software/globalization/locales/Macedonia-Macedonian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
130. "Commonwealth of the Northern Mariana Islands" (https://web.archive.org/web/20170505204247/http://gov.mp/). Archived from the original (http://gov.mp) on 2017-05-05. Retrieved 2017-05-06.
131. "Tall, tid og dato" (http://www.sprakradet.no/sprakhjelp/Skriveregler/Dato/) [Numbers, time and date] (in Norwegian). Norway: Språkrådet [Language Council]. 2015-06-30. Archived (https://web.archive.org/web/20161204064702/http://www.sprakradet.no/sprakhjelp/Skriveregler/Dato) from the original on 2016-12-04. Retrieved 2017-01-12.
132. "Internasjonal datostandard" (https://www.standard.no/nyheter/nyhetsarkiv/ikt/2013/internasjonal-datostandard/). Archived (https://web.archive.org/web/20160215105632/https://www.standard.no/nyheter/nyhetsarkiv/ikt/2013/internasjonal-datostandard/) from the original on 2016-02-15. Retrieved 2016-02-10.
133. "Globalization Library – Locale Data: Oman" (ftp://ftp.software.ibm.com/software/globalization/locales/Oman-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
134. Republic of Palau National Government (https://web.archive.org/web/20130615141726/http://palaugov.org/)
135. "National Language Support (NLS) API Reference -- Panama" (https://www.microsoft.com/resources/msdn/goglobal/default.mspx?submitted=180A&OS=Windows%20Vista). Microsoft. Retrieved 2017-06-17.
136. "Globalization Library – Locale Data: Paraguay" (ftp://ftp.software.ibm.com/software/globalization/locales/Paraguay-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
137. "Globalization Library – Locale Data: Peru" (ftp://ftp.software.ibm.com/software/globalization/locales/Peru-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
138. "Globalization Library – Locale Data: Philippines" (ftp://ftp.software.ibm.com/software/globalization/locales/Philippines-English_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
139. "Globalization Library – Locale Data: Poland" (ftp://ftp.software.ibm.com/software/globalization/locales/Poland-Polish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
140. (in Polish) Kancelaria Prezesa Rady Ministrów (Republic of Poland – The Chancellery of the Prime Minister) (http://www.kprm.gov.pl/) Archived (https://web.archive.org/web/20101216115814/http://kprm.gov.pl/) 2010-12-16 at the Wayback Machine
141. Zagórska Brooks, Maria (1975). *Polish Reference Grammar*. Walter de Gruyter. p. 35. ISBN 90-279-3313-8.
142. "Globalization Library – Locale Data: Portugal" (ftp://ftp.software.ibm.com/software/globalization/locales/Portugal-Portuguese_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
143. "Globalization Library – Locale Data: Qatar" (ftp://ftp.software.ibm.com/software/globalization/locales/Qatar-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
144. (in Romanian) Guvernul României (Government of Romania) (http://www.gov.ro/) Archived (https://web.archive.org/web/20171017201950/http://www.gov.ro/) 2017-10-17 at the Wayback Machine
145. "Globalization Library – Locale Data: Romania" (ftp://ftp.software.ibm.com/software/globalization/locales/Romania-Romanian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
146. "Globalization Library – Locale Data: Russia" (ftp://ftp.software.ibm.com/software/globalization/locales/Russia-Russian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
147. "ICU Locale "English (St. Helena)" (en_SH)" (https://www.localeplanet.com/icu/en-SH/index.html). *www.localeplanet.com*. Archived (https://web.archive.org/web/20210420112712/https://www.localeplanet.com/icu/en-SH/index.html) from the original on 2021-04-20. Retrieved 2021-05-10.
148. "ICU Locale "French (St. Martin)" (fr_MF)" (https://www.localeplanet.com/icu/fr-MF/index.html). *www.localeplanet.com*. Archived (https://web.archive.org/web/20210510025148/https://www.localeplanet.com/icu/fr-MF/index.html) from the original on 2021-05-10. Retrieved 2021-05-10.
149. "Globalization Library – Locale Data: Saudi Arabia" (ftp://ftp.software.ibm.com/software/globalization/locales/SaudiArabia-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
150. "NLS information page – Arabic (Saudi Arabia)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0401&OS=Windows%20Vista). Microsoft. Retrieved 2008-10-29.

151. "Главна страна" (https://sr.wikipedia.org/). 2016-10-22. Archived (https://web.archive.org/web/20050208190449/http://www.sr.wikipedia.org/) from the original on 2005-02-08 – via Wikipedia.
152. (in Serbian) Влада Републике Србије (Serbian Government) (http://www.srbija.gov.rs/?change_lang=cr) Archived (https://web.archive.org/web/20110511150828/http://www.srbija.gov.rs/?change_lang=cr) 2011-05-11 at the Wayback Machine
153. "NLS information page – Serbian (Cyrillic)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0C1A&OS=Windows%20Vista). Microsoft. Retrieved 2008-11-16.
154. "NLS information page – Serbian (Latin)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=081A&OS=Windows%20Vista). Microsoft. Retrieved 2008-11-16.
155. "Globalization Library – Locale Data: Singapore (Simplified Chinese)" (ftp://ftp.software.ibm.com/software/globalization/locales/Singapore-SimplifiedChinese_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
156. "Globalization Library – Locale Data: Singapore (English)" (ftp://ftp.software.ibm.com/software/globalization/locales/Singapore-English_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
157. "Slovak Grammar" (http://www.juls.savba.sk/ediela/psp2000/psp.pdf) (PDF). Veda. Archived (https://web.archive.org/web/20141011062702/http://www.juls.savba.sk/ediela/psp2000/psp.pdf) (PDF) from the original on 2014-10-11. Retrieved 2014-06-13.
158. "Pisanje datumov v slovenščini" (https://www.lektorsko-drustvo.si/jezikovni-pogovori/pisanje-datumov-v-slovenscini). Lektorsko društvo Slovenije. Archived (https://web.archive.org/web/20200520062016/https://www.lektorsko-drustvo.si/jezikovni-pogovori/pisanje-datumov-v-slovenscini) from the original on 2020-05-20. Retrieved 2020-01-19.
159. "ICU Demonstration - Locale Explorer (af_ZA)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=af_ZA). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035824/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=af_ZA) from the original on 2022-03-18. Retrieved 2022-03-18.
160. "ICU Demonstration - Locale Explorer (xh_ZA)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=xh_ZA). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035822/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=xh_ZA) from the original on 2022-03-18. Retrieved 2022-03-18.
161. "ICU Demonstration - Locale Explorer (zu_ZA)" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=zu_ZA). *icu4c-demos.unicode.org*. Archived (https://web.archive.org/web/20220318035821/https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=zu_ZA) from the original on 2022-03-18. Retrieved 2022-03-18.
162. "SANS 8601:2009 (Ed. 2.00)" (https://store.sabs.co.za/catalog/product/view/id/216987/s/sans-8601-2009-ed-2-00/). *SABS Webstore*. Archived (https://web.archive.org/web/20211124053222/https://store.sabs.co.za/catalog/product/view/id/216987/s/sans-8601-2009-ed-2-00/) from the original on 2021-11-24. Retrieved 2021-11-24.
163. "Globalization Library – Locale Data: Spain" (ftp://ftp.software.ibm.com/software/globalization/locales/Spain-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12. Catalan language locale for Spain also indicates dd/mm/yyyy for Common Date format.
164. "ICU Demonstration - Locale Explorer" (https://icu4c-demos.unicode.org/icu-bin/locexp?d_=en&_=eu_ES). *icu4c-demos.unicode.org*. Retrieved 2023-02-05.
165. "Globalization Library – Locale Data: Sweden" (ftp://ftp.software.ibm.com/software/globalization/locales/Sweden-Swedish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
166. "Regulation (EU) No 1169/2011 of the European Parliament and of the Council of 25 October 2011 on the provision of food information to consumers, amending Regulations (EC) No 1924/2006 and (EC) No 1925/2006 of the European Parliament and of the Council, and repealing Commission Directive 87/250/EEC, Council Directive 90/496/EEC, Commission Directive 1999/10/EC, Directive 2000/13/EC of the European Parliament and of the Council, Commission Directives 2002/67/EC and 2008/5/EC and Commission Regulation (EC) No 608/2004 Text with EEA relevance" (https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32011R1169&from=EN#d1e38-57-1). 2011-10-25. Annex X, 2.c). Archived (https://web.archive.org/web/20200809065209/https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32011R1169&from=EN#d1e38-57-1) from the original on 2020-08-09. Retrieved 2020-08-23.
167. Switzerland Federal Administration – Press releases and speeches (http://www.news.admin.ch/index.html?lang=en) Archived (https://web.archive.org/web/20110117172145/http://www.news.admin.ch/index.html?lang=en) 2011-01-17 at the Wayback Machine dd.mm.yyyy format seen in all languages.

168. "Globalization Library – Locale Data: Switzerland (French, German, Italian)" (https://web.archive.org/web/20090809230551/http://www-01.ibm.com/software/globalization/locales/locales_eu.jsp). IBM. Archived from the original (http://www-01.ibm.com/software/globalization/locales/locales_eu.jsp) on 2009-08-09. Retrieved 2008-10-12.
169. "Globalization Library – Locale Data: Syria" (ftp://ftp.software.ibm.com/software/globalization/locales/Syria-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
170. "台北101 官方網站" (https://www.taipei-101.com.tw/tw/). *Taipei 101*. Archived (https://web.archive.org/web/20210508052025/https://www.taipei-101.com.tw/tw) from the original on 2021-05-08. Retrieved 2021-05-09.
171. "中天公告 Archives" (https://www.ctitv.com.tw/category/announce). *中天電視* (in Chinese (Taiwan)). Archived (https://web.archive.org/web/20210510143126/https://www.ctitv.com.tw/category/announce) from the original on 2021-05-10. Retrieved 2021-05-10.
172. "Globalization Library – Locale Data: Taiwan (Simplified Chinese)" (ftp://ftp.software.ibm.com/software/globalization/locales/Taiwan-TraditionalChinese_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
173. "NLS information page – Tajik (Cyrillic, Tajikistan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0428&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044554/http://www.microsoft.com/favicon.ico) from the original on 2022-05-14. Retrieved 2008-10-29.
174. "Globalization Library – Locale Data: Thailand" (ftp://ftp.software.ibm.com/software/globalization/locales/Thailand-Thai_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
175. "NLS information page – English (Trinidad and Tobago)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=2C09&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044545/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=2C09&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2008-10-29.
176. "Globalization Library – Locale Data: Tunisia" (ftp://ftp.software.ibm.com/software/globalization/locales/Tunisia-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
177. (in Turkish) Türkiye Cumhuriyeti Cumhurbaşkanlığı (Presidency of the Republic of Turkey) (http://www.tccb.gov.tr/) Archived (https://web.archive.org/web/20140912124059/http://www.tccb.gov.tr//) 2014-09-12 at the Wayback Machine
178. "Globalization Library – Locale Data: Turkey" (ftp://ftp.software.ibm.com/software/globalization/locales/Turkey-Turkish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
179. "ICU Locale "Turkish (Turkey)" (tr_TR)" (https://www.localeplanet.com/icu/tr-TR/index.html). *www.localeplanet.com*. Archived (https://web.archive.org/web/20210510142355/https://www.localeplanet.com/icu/tr-TR/index.html) from the original on 2021-05-10. Retrieved 2021-05-10.
180. (in Turkmen) Government of Turkmenistan (http://www.turkmenistan.gov.tm/) Archived (https://web.archive.org/web/20090323035347/http://www.turkmenistan.gov.tm/) 2009-03-23 at the Wayback Machine
181. "NLS information page – Turkmen (Turkmenistan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0442&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044553/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0442&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2009-02-08.
182. (in Ukrainian) Government of the Ukraine (http://www.kmu.gov.ua/control/uk) Archived (https://web.archive.org/web/20101229044141/http://www.kmu.gov.ua/control/uk/) 2010-12-29 at the Wayback Machine
183. "NLS information page – Ukrainian (Ukraine)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0422&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044545/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0422&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2009-02-08.
184. "Globalization Library – Locale Data: Ukraine" (ftp://ftp.software.ibm.com/software/globalization/locales/Ukraine-Ukrainian_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
185. "Globalization Library – Locale Data: United Arab Emirates" (ftp://ftp.software.ibm.com/software/globalization/locales/UnitedArabEmirates-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
186. "NLS information page – Arabic (U.A.E.)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=3801&OS=Windows%20Vista). Microsoft. Retrieved 2009-02-08.
187. "The Guardian Style Guide – dates" (https://www.theguardian.com/styleguide/d). *The Guardian*. Archived (https://web.archive.org/web/20131125001247/http://www.theguardian.com/styleguide/d) from the original on 2013-11-25. Retrieved 2013-11-25.

188. Ritter, R. M., ed. (2003). "7.10.1". *The Oxford Style Manual*. Oxford: Oxford University Press. pp. 178–179. ISBN 978-0-19-860564-5.
189. "The Times frontpage" (http://www.thetimes.co.uk/tto/news). *The Times*. Archived (https://web.archive.org/web/20140211224356/http://www.thetimes.co.uk/tto/news/) from the original on 2014-02-11. Retrieved 2014-02-12.
190. "The Guardian frontpage" (https://www.theguardian.com/uk). *The Guardian*. Archived (https://web.archive.org/web/20140212015927/http://www.theguardian.com/uk) from the original on 2014-02-12. Retrieved 2014-02-12.
191. "BS ISO 8601:2004 – Data elements and interchange formats. Information interchange. Representation of dates and times" (http://shop.bsigroup.com/ProductDetail/?pid=000000000030078726). BSI Group. 2005-03-10. Archived (https://web.archive.org/web/20140222134919/http://shop.bsigroup.com/ProductDetail/?pid=000000000030078726) from the original on 2014-02-22. Retrieved 2014-02-10.
192. "Globalization Library – Locale Data: United States" (ftp://ftp.software.ibm.com/software/globalization/locales/UnitedStates-English_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
193. "NLS information page – English (United States)" (https://archive.today/20090112045306/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0409&OS=Windows+Vista). Microsoft. Archived from the original (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0409&OS=Windows%20Vista) on 2009-01-12. Retrieved 2009-02-08.
194. "Welcome to the U.S. VIRGIN ISLANDS" (https://web.archive.org/web/20150507120837/http://www.gov.vi/). United States Virgin Islands Government. Archived from the original (http://www.gov.vi/) on 2015-05-07. Retrieved 2017-06-20.
195. (in Spanish) Government of Uruguay: Documentos de Interés (documents of interest) (http://www.presidencia.gub.uy/_web/pages/docs.htm) Archived (https://web.archive.org/web/20100301080703/http://www.presidencia.gub.uy/_web/pages/docs.htm) 2010-03-01 at the Wayback Machine
196. "Globalization Library – Locale Data: Uruguay" (ftp://ftp.software.ibm.com/software/globalization/locales/Uruguay-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
197. (in Uzbek) Government of Uzbekistan (http://www.gov.uz/uz/) Archived (https://web.archive.org/web/20121123022925/http://www.gov.uz/uz/) 2012-11-23 at the Wayback Machine
198. "NLS information page – Uzbek (Cyrillic, Uzbekistan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0843&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044546/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0843&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2008-10-29.
199. "NLS information page – Uzbek (Latin, Uzbekistan)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0443&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044554/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=0443&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2008-10-29.
200. (in Spanish) Government of Venezuela: Noticias (News) (http://www.gobiernoenlinea.ve/noticias-view/ver_noticias.pag?sw=1) Archived (https://web.archive.org/web/20110205204844/http://www.gobiernoenlinea.ve/noticias-view/ver_noticias.pag?sw=1) 2011-02-05 at the Wayback Machine
201. "Globalization Library – Locale Data: Venezuela" (ftp://ftp.software.ibm.com/software/globalization/locales/Venezuela-Spanish_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
202. "NLS information page – Spanish (Venezuela)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=200A&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044556/http://www.microsoft.com/favicon.ico) from the original on 2022-05-14. Retrieved 2009-02-08.
203. "Thông tư số 01/2011/TT-BNV của Bộ Nội vụ : Hướng dẫn thể thức và kỹ thuật trình bày văn bản hành chính" (http://www.chinhphu.vn/portal/page/portal/chinhphu/hethongvanban?class_id=1&mode=detail&document_id=99777). Archived (https://web.archive.org/web/20210510031401/http://www.chinhphu.vn/portal/page/portal/chinhphu/hethongvanban?class_id=1&mode=detail&document_id=99777) from the original on 2021-05-10. Retrieved 2021-05-10.
204. "Toàn cảnh Covid-19: Tin tức, số liệu, phân tích" (https://thanhnien.vn/content/OTg0MTU5.html). *Báo Thanh Niên*. 2020-08-13. Archived (https://web.archive.org/web/20210908031322/https://thanhnien.vn/content/OTg0MTU5.html) from the original on 2021-09-08. Retrieved 2021-05-10.

205. "NLS information page – Vietnamese (Vietnam)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=042A&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044551/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=042A&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2008-10-29.
206. "VnExpress International - Latest news, business, travel and analysis from Vietnam" (https://e.vnexpress.net/). Archived (https://web.archive.org/web/20191104130751/https://e.vnexpress.net/) from the original on 2019-11-04. Retrieved 2021-05-10.
207. "Globalization Library – Locale Data: Yemen" (ftp://ftp.software.ibm.com/software/globalization/locales/Yemen-Arabic_Date.pdf) (PDF). IBM. Retrieved 2008-10-12.
208. "NLS information page – Arabic (Yemen)" (http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=2401&OS=Windows%20Vista). Microsoft. Archived (https://web.archive.org/web/20220514044624/http://www.microsoft.com/globaldev/nlsweb/default.mspx?submitted=2401&OS=Windows+Vista) from the original on 2022-05-14. Retrieved 2008-10-29.
209. "ICU Locale "English (Zimbabwe)" (en_ZW)" (http://www.localeplanet.com/icu/en-ZW/). Locale Planet. Archived (https://web.archive.org/web/20160915174014/http://www.localeplanet.com/icu/en-ZW/) from the original on 2016-09-15. Retrieved 2016-09-09.

External links

- Index of NLS information page (http://msdn.microsoft.com/en-us/goglobal/bb896001.aspx) Global Development and Computing Portal, published by Microsoft. Links on page lead to individual country date formats.

Retrieved from "https://en.wikipedia.org/w/index.php?title=Date_format_by_country&oldid=1178728233"

ATTACHMENT D

WorldCat: Proceedings of the Twelfth International Conference on Telecommun...

OCLC	254719202	**No holdings in LDL - 1 other holding**									
Books	**Rec Stat**	c		**Entered**	20050531		**Replaced**	20220515010958.9			
Type	a	**ELvl**	M	**Srce**	d		**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	1		**Biog**		**MRec**		**Ctry**	cau
		Cont		**GPub**			**LitF**	0	**Indx**	0		
Desc	c	**Ills**		**Fest**	0		**DtSt**	s	**Dates**	2004,		

040 GBVCP ‡b ger ‡e rakwb ‡c GBVCP ‡d OCLCQ ‡d OCLCO
020 0971625352
020 9780971625358
084 53.70 ‡2 bcl
090
049 LDLL
110 2_ International Conference on Telecommunication Systems ‡n 12 ‡d 2004 ‡c Monterey, Calif. ‡0 (DE-588)6044077-6
245 10 Proceedings of the Twelfth International Conference on Telecommunication Systems : ‡b modeling and analysis ; July 22-25, 2004 Monterey, California / ‡c sponsors: American Telecommunication Systems Management Association ...
260 Monterey, Calif. ‡c 2004
300 V, 352 Seiten
336 Text ‡b txt ‡2 rdacontent/ger
337 ohne Hilfsmittel zu benutzen ‡b n ‡2 rdamedia/ger
338 Band ‡b nc ‡2 rdacarrier/ger
710 2_ American Telecommunication Systems Management Association ‡0 (DE-588)5339868-3
711 2_ International Conference on Telecommunication Systems ‡n (12 ‡d 2004.07.22-25 ‡c Monterey, Calif.)
029 0_ GBVCP ‡b 487740491

Action Status Delete Holdings _ Export _ Label _ Produce _ Update Holdings _ Validate _

Sophos Ex. 1009

ATTACHMENT E

WorldCat: Proceedings of the Twelfth International Conference on Telecommun...

OCLC	60537577	**No holdings in LDL - 3 other holdings**									
Books	**Rec Stat**	c		**Entered**	20050422		**Replaced**	20220528134741.0			
Type	a	**ELvl**	M	**Srce**	d	**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	1	**Biog**		**MRec**		**Ctry**	txu
		Cont	b	**GPub**		**LitF**	0	**Indx**	0		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	2004,		

040 CAI ‡b eng ‡c CAI ‡d OCLCF ‡d OCLCO ‡d OCLCQ ‡d OCL ‡d OCLCO
020 0971625352
020 9780971625358
050 14 TK5101.A1 I619 2004
090
049 LDLL
111 2_ International Conference on Telecommunication Systems ‡n (12th : ‡d 2004 : ‡c Monterey, Calif.)
245 10 Proceedings of the Twelfth International Conference on Telecommunication Systems : ‡b modeling and analysis, July 22-25, 2004, Monterey, California / ‡c sponsors American Telecommunications Systems Management Association [and others].
260 [Dallas, Tex.] : ‡b [Edwin L. Cox School of Business, Southern Methodist University], ‡c [2004?]
300 v, 352 pages : ‡b illustrations ; ‡c 28 cm
336 text ‡b txt ‡2 rdacontent
337 unmediated ‡b n ‡2 rdamedia
338 volume ‡b nc ‡2 rdacarrier
504 Includes bibliographical references.
650 _0 Telecommunication systems ‡v Congresses.
650 _0 Telecommunication ‡x Technological innovations ‡v Congresses.
650 _0 Data transmission systems ‡v Congresses.
650 _0 Computer networks ‡v Congresses.
650 _6 Systèmes de télécommunications ‡0 (CaQQLa)201-0023227 ‡v Congrès. ‡0 (CaQQLa)201-0378219
650 _6 Télécommunications ‡0 (CaQQLa)201-0005415 ‡x Innovations ‡0 (CaQQLa)201-0379286 ‡v Congrès. ‡0 (CaQQLa)201-0378219
650 _6 Réseaux d'ordinateurs ‡0 (CaQQLa)201-0007877 ‡v Congrès. ‡0 (CaQQLa)201-0378219
650 _7 Computer networks. ‡2 fast ‡0 (OCoLC)fst00872297
650 _7 Data transmission systems. ‡2 fast ‡0 (OCoLC)fst00887993
650 _7 Telecommunication systems. ‡2 fast ‡0 (OCoLC)fst01146038
650 _7 Telecommunication ‡x Technological innovations. ‡2 fast ‡0 (OCoLC)fst01145958
655 _7 Conference papers and proceedings. ‡2 fast ‡0 (OCoLC)fst01423772
710 2_ Southern Methodist University.
710 2_ American Telecommunications Systems Management Association.

Action Status Delete Holdings _ Export _ Label _ Produce _ Update Holdings _ Validate _

Sophos Ex. 1009

ATTACHMENT F

MAIN
TK5101
.A1
I619
2004

Proceedings of the ... Inter
MONO/CONF CAT'G

Proceedings of

Twelfth International Conference on Telecommunication Systems

Modeling and Analysis

July 22-25, 2004
Monterey, California

Sponsors

- American Telecommunications Systems Management Association
- Cebrowski Institute for Information Superiority
- IFIP Working Group 7.3 "Computer system modeling and performance evaluation"
- INFORMS Technical Section on Telecommunications
- INFORMS Information Systems Society
- Naval Postgraduate School, Monterey, CA
- School of Management, The University of Texas at Dallas

Sophos Ex. 1009
Page 60 of 318

CISTI/ICIST NRC/CNRC
MAIN
Received on: 05-04-21
Proceedings of the ...
International conference on
telecommunication systems



NOT FOR LOAN
PAS DISPONIBLE POUR LE PRÊT

Proceedings of the

Twelfth International Conference on Telecommunication Systems

Modeling and Analysis

July 22-25, 2004
Monterey, California

Sponsors

- **American Telecommunications Systems Management Association**
- **Cebrowski Institute for Information Superiority**
- **IFIP Working Group 7.3 "Computer system modeling and performance evaluation"**
- **INFORMS Technical Section on Telecommunications**
- **INFORMS Information Systems Society**
- **Naval Postgraduate School, Monterey, CA**
- **School of Management, The University of Texas at Dallas**

Sophos Ex. 1009
Page 62 of 318

Table of Contents

A GRASP Heuristic for the Delay-Constrained Multicast Routing Problem
Nina Skorin-Kapov, Mladen Kos 1
A Policy Based Collaboration Infrastructure for P2P Networking
Mark R. Cornwell, Lalana Kagal, James E. Just, Tim Finn 11
A Role Model for Classification and Analysis of Internet Service Providers
Oliver Heckman, Ralf Steinmetz 21
Agents for Wireless Environments
Badr Benmammar, Francine Krief 37
An Architecture for Managing Access to and Permission for Multiple Versions of Objects in a Distributed Environment
Tom Rozenbroek 46
An Architecture for m-Learning with Wireless Networks
Michael Bartolacci, John C. Wiginton 64
An Distributed Architecture for the Publication and Management of Objects, which provides End-to-End Security
Randall Breeden, Tom Rozenbroek 67
An Optimization Approach for Connectivity in Wireless Networks
Michael Bartolacci, Abdullah Konak, Roger Whitaker 79
Analysis of the Packet Loss Process for Multimedia Traffic
György Dan, Viktória Fodor, Gunnar Karlsson 83
Army Enterprise Infostructure – Transport (AEI-T) Modeling and System (M&S) Engineering Effort
Carl Agren 95
CDN Economics and Exogenous Congestion
Eugenio de la Rosa, John H. Hartman, Terril N. Hurst 101
Coupling Network Simulators to Real-World Networks
Kevin M. McNeill, Wenji Wu 110
Dependence of Angle-to-Arrival Fluctuations on Source Coherence for Free Space Optics Links in a Turbulent Atmosphere
Yahya Baykal, Halil T. Eyyuboğlu 121
Design of an Economics-Based Software Infrastructure for Secure Utility Computing on Supercomputing Clusters
Gregory A. Koenig, William Yurcik 127
Developing IP Network Structures with Regard to Traffic Engineering and Failure Resilience
Gerhard Haβlinger, Stefan Schnitter, Martin Franzke 142

Sophos Ex. 1009
Page 63 of 318

Dynamic Adaptive Epidemic Models for Peer-to-Peer Network
Donald P. Gaver, Patricia A. Jacobs, Cheng-Chung Kao 155
FORCENet: Weve been here before
Rex Buddenberg 164
Four Dimensions of Enterprise Resource Planning (ERP) Systems: A New Classification Approach
Zinovy Radovilsky 171
Geographical NetFlows Visualization for Network Situational Awareness: NaukaNet Administrative Data Analysis System (NADAS)
Greg Cole, Natasha Bulashova, William Yurcik 177
I3MP Defense in Depth Security Levels
Steven K. Bates 185
Integration application with Public WLAN married to GPRS/3G
Wei Ye 191
Intracell and intercell interference reduction in FBWA networks through enhanced resource allocation techniques
A. Vavoulas, T. Rokkas, D. Varoutas, Th. Sphicopoulos 199
IPv6 Interoperability Network – Lessons Learned
Scott J. Baier 207
Management and Control Plane Techniques in Self-Aware Networks
Natalia Gaviria, Kevin M. McNeill, Jerzy Rozenblit 213
Mobility in the IP Networks
Mounir Frikha, Lilia Maalej 219
Network Management Using Active Probing
Marat Zhanikeev, Yoshiaki Tanaka 236
On-Time Performance and the Value of RFID in Cross-Docking
Steve Peng 246
Packet Loss and Delay Jitter Characteristics of the Internet – VoD Trial Service with Packet Recovery
Tetsuo Morita, Yutaka Majima, Toshihiro Takashima, Takahiro Kusumoto, Takashi Sasaki, Yutaka Takahashi 252
Peer-to-Peer Music Distribution
Marcin Matuszewski 259
Solving the "Last Mile" Problem for the Global Information Grid Bandwidth Expansion to Tactical Units
James E. Just, Dhananjay Phatak, Thomas Goff 270
Telecommunications in Complex Systems Failure: The Electricity Blackout of August 14, 2003
John C. Hoag 286
The Channel Characterization for Mobile Communication Employing Stratospheric Platform
Iskandar Febryn, Shigeru Shimamoto 293

Sophos Ex. 1009
Page 64 of 318

The Network Testbed Sizing Problem – A Proposed Research Framework
Kevin M. McNeill 299
Towards a First Responder Contingency Information Capability
Geoffrey G. Xie, Su Wen, John H. Gibson, Gurminder Singh 311
Towards Device-Aware Networks
Su Wen, Gurminder Singh, John H.Gibson, Arijit Das 321
UCLog: A Unified, Correlated Logging Architecture for Intrusion Detection
Zhenmin Li, Jed Taylor, Elizabeth Partridge, Yuanyuan Zhouk, William Yurcik, Cristina Abad, James J. Barlow, Jeff Rosendale 329
Virtual Path Optimization in ATM Networks
Jeffrey L. Duffany 344

Sophos Ex. 1009
Page 65 of 318

UCLog: A Unified, Correlated Logging Architecture for Intrusion Detection

Zhenmin Li[†], Jed Taylor[†], Elizabeth Partridge[†‡],
Yuanyuan Zhou[†], William Yurcik[‡],
Cristina Abad[†‡], James J. Barlow[‡] and Jeff Rosendale[‡]

[†]Department of Computer Science,
University of Illinois at Urbana-Champaign
{zli4, jtaylr2, epatrid, yyzhou}@uiuc.edu
[‡]National Center for Supercomputing Applications (NCSA)
{byurcik, cabad, jbarlow, jeffr}@ncsa.uiuc.edu

Abstract

Activity logs can be used for intrusion detection; however, most previous work on intrusion detection examines only activity logs from a single component. Doing so fails to take advantage of the naturally existing correlations among activities in different types of logs, such as network logs and system call logs.

This paper explores correlation for intrusion detection. Specifically, we propose UCLog, *a unified logging architecture that can effectively capture correlations among entries in different types of logs. UCLog enables the intrusion detection system to make some sense of the myriad of different available logs and correlate the information the logs present to enhance the intrusion detection process.*

We have evaluated UCLog by using it to detect the infection of a host with the Yaha virus. Our results show significant improvement when the information available in several logs is correlated.

Key words: *intrusion detection, auditing, logging, file I/O monitoring*

1 Introduction

Computer systems are vulnerable targets for terrorists and criminals to break-in. Intrusions can compromise the integrity, confidentiality, or availability of the systems and the information they hold [8]. For example, in January 2004 when the SQL Slammer Worm in less than seven hours the Internet traffic was overloaded to the point that many private networks were jammed and web sites stopped responding. The cost of the shut downs and cleaning up the infection was estimated to be around $1 billion [25].

An intrusion detection system (IDS) attempts to detect attacks by monitoring system or network behavior. While many existing IDSs require manual definitions of normal and abnormal behavior (intrusion signatures) [37], recent work has shown that it is possible to identify abnormalities automatically using machine learning or data mining techniques [15, 5, 7, 18, 23] to generate models or rules, which the IDS can use to detect intrusions.

However, most of the previous work on intrusion detection focuses on activities generated by a single source, resulting in many false positives and undetected intrusions. For example, some studies [23] detect intrusions by monitoring the network traffic such as tcpdump logs, whereas other studies [7, 18] monitor only system call logs. These studies usually have high false positive rates. Correlating activities from all involved components and then analyzing them together would give the IDS a better *picture* of the system, increasing the detection accuracy and reducing the false positives.

Unfortunately, few studies have been conducted on fusing audit log data from heterogeneous sources into a single framework for assessing suspicious behavior. Some existing systems, such as Windows, provide APIs to allow both device drivers or applications to log events in a centralized view [27]. Similarly, the ReVirt system [16] can record all actions below the level of the virtual machine, allowing arbitrarily detailed information about all user-level behavior at a given time. These systems provide a centralized mechanism to log activities from heterogeneous sources, but activities are not organized in a structured making it hard to capture their correlations.

Another common problem with current IDSs is that it is hard or impossible for the user to specify his/hers critical applications which should be watched for instrusions. The UCLog architecture makes it easy to add application-specific loggers and monitors, which form a middleware layer that enable the IDS to monitor different applications

Sophos Ex. 1009
Page 66 of 318

without modifying the code of the applications or the IDS itself.

This paper proposes a unified logging architecture, called UCLog, that effectively explores activity correlations for intrusion detection. UCLog is motivated from and complements previous work on secure logging such as Self-Securing Storage System (S4) [39] and ReVirt [16]. UCLog also complements previous work. More specifically, the UCLog architecture provides several benefits:

- UCLog organizes all audit logs from heterogeneous sources in a very structured way and is closely integrated into the intrusion detection system.
- UCLog effectively exploits activity correlations between different components to improve the accuracy of intrusion detection and to reduce the number of false positives.
- At an intrusion, UCLog can be used to automatically identify the source and the execution path of the intrusion, whereas most current systems require an administrator to manually examine the logs to find such information.
- UCLog can easily be extended to use information from new applications in the intrusion detection process, by adding the corresponding logger and monitor to the system.

We have implemented the UCLog architecture on a real system running Microsoft Outlook Express on Windows 2000. We evaluate this system using many normal and abnormal traces. Abnormal traces are collected when the system is infected by e-mail viruses. Using these traces, we have studied how correlating activities from the system call sequence, network activity, and file I/O operations can improve the intrusion detection accuracy. Our results show that our proposed architecture can not only provide the means to make an intrusion more obvious to an IDS, but can also help reduce the number of false positives. Compared to a baseline system that only monitors a single component, our system improves the accuracy by up to a factor of 2.5.

This paper is organized as follows. Section 2 briefly describes background material. Section 3 discusses the effectiveness of correlation for intrusion detection. Section 4 presents the UCLog architecture. Section 5 and 6 present the evaluation methodology and results. Section 7 summarizes related work. In Sect. 8 we describe what still needs to be done, and in Sect. 9 we conclude.

2 Background

This section describes some background on activity logging and how log files are used in intrusion detection systems (IDSs).

2.1 Activity Logging

Activity logging is a mechanism to collect the information about past activities that the system or administrator can use to characterize access behavior, audit user accounts, replay operations, or detect anomalies and rollback to a previous state [16, 39]. For example, the system can log every user login/logout to monitor any possible break-ins. Research issues about logging include what and how to log, how to manage and protect logs, and how to use the logged data.

Event logging in Microsoft Windows provides a standard, centralized way for applications to record important software and hardware events [27]. By default, there are three types of logs: security, system, and application. Additionally, Windows event logging provides an event log component and a dynamic-link library (DLL) that includes some application programming interfaces (APIs) for logging management such as storing, retrieving, and backing up data. MS Windows event logging organizes these events in a very simple way but does not provide any support to find correlations between events from different components, nor is it designed to easily connect with any IDS.

To avoid the risk of compromise, it is important to develop protection mechanisms for logging. Dunlap et al. propose ReVirt [16], in which the operating system is moved to a virtual machine and logging takes place below the level of the virtual machine. This protects the integrity of the logging system. Furthermore, the ReVirt system can replay the instruction-by-instruction execution of the system in order to test the effects of a particular command. The ReVirt logging architecture is designed for simple monitoring, recovery, and replay. But similar to the MS Windows event logger, it does not support correlations among activities in different components and is not connected with any IDS.

Sophos Ex. 1009
Page 67 of 318

Additionally, the systems examined in this paper would benefit from using a secure logging system such as Rio/Vista [26] or Self Securing Storage System(S4) [39]. Rio/Vista is a cache designed to survive OS crashes, coupled with a recoverable memory library which allows fast logging [26]. S4 isolates logging from user interactions, to reduce performance costs [39].

Our work complements previous work such as ReVirt, MS Windows Event logger, Rio/Vista, and S4 by providing a unified logging architecture that captures correlations and integrates with the IDS to build a system that is highly secure, reliable and self-recoverable.

2.2 Intrusion Detection Systems

Since log files contain information about past activities, the data collected by a logging system can be used by an IDS to detect signs of suspicious behavior.

A generic IDS contains three major components: sensors, analyzers, and user interface [2]. Sensors are responsible for collecting data, such as network packets and log files. The analyzer determines if an intrusion has occurred by analyzing the data from sensors. The user interface enables an administrator to view output from the system and to control the behavior of the IDS.

Analyzers can be classified in two: policy-based and anomaly-based [3]. In a policy-based detection system, known signatures of intrusive or suspicious behavior are defined by the security policy. In contrast, an anomaly-based detector reacts to deviations from previously observed normal behavior (e.g. by using data mining or machine learning techniques [23, 7, 11, 5]). These definitions of normal behavior can be automatically updated. Our system can easily combine both mechanisms. Each log can be processed by any type of analyzer, and the analyzers collaborate with each other to identify suspicious behavior.

Most previous work on IDSs analyze a single component, ignoring the correlations among different components. We first show how correlation can be used to improve the accuracy of intrusion detection. Then, we study how to build a logging architecture that can conveniently capture these correlations and integrate them with the IDS. More details on existing IDSs can be found in Sect. 7.

3 Exploring Correlations for Intrusion Detection

We have conducted several experiments to examine the correlation between logs of different system components. The motivation of our experiments is to justify that the correlation can potentially increase the accuracy of an intrusion detection system. In other words, when the network shows abnormal behavior, is the system call sequence also abnormal? And vice versa?

In our experiments, we have collected several network and system call traces. Some of these traces are collected when the system behaves normally, and the others are collected when the system is infected by e-mail viruses. We use data mining tools to analyze these traces separately and identify abnormal regions for each trace, then we examine the results of every network trace and the results of its corresponding system call trace (collected simultaneously) to check for correlations between the network activity and system calls.

Our experiments consist of three steps: collecting audit data, creating detection models by using data mining techniques, and then applying the models to detect abnormal behavior.

In order to train and test such a system, the first step is to collect multiple traces of normal and abnormal user behavior simultaneously from different subsystems. In our experiments, we collected traces from a Windows 2000 host, but the methods can be similarly applied to any other operating systems. The logging tools we used are described in Sect. 5. We used traces containing the Win32 Yaha virus [31] and normal traces when training our system.

The next step is to train the detection system by applying a data-mining algorithm, RIPPER [12, 13], to the trace data to extract useful behavior patterns. The rules generated by RIPPER are used to classify normal and abnormal system call sequences. The processing detail is described in Sect. 5.

After obtaining the rules from the training data, the next step is prediction and post-processing of test data. Similar to Lee's post-processing approach in [23], each sequence of system calls in the testing trace is predicted as "normal" or "abnormal" according to the rules. The predictions are then divided into regions of size $3l + 1$. If the number of abnormal prediction in a region is greater than a half, we predict

Sophos Ex. 1009
Page 68 of 318

the region as "abnormal". Next, we take an overlapping sliding window of size 500 and pass it over the sequence of predictions. For each system call sequence, we calculate the number of abnormal predictions in this sliding window. Results are graphed in Fig. 1 and give us an idea of the behavior of a given abnormal trace.

Figure 1(a) shows the results of the system call trace collected when the system is infected by the Yaha virus. As shown on the curve, there are some patches of increased abnormal behavior, but it is difficult to determine if any intrusion has really taken place, especially for the first abnormal spike. It is possible that either the user has just issued a command or tried a new feature that he has never tried before, or that a real intrusion occurs, i.e., the user just opened an e-mail attachment that compromises the local address book and fires off infected e-mail to other people in the address book. By examining the system call behavior alone, these types of questions can be difficult to answer because the system calls alone may not look all that abnormal.

When we examine the network traffic corresponding to the same time period, we can see an obvious increase in the network traffic as shown in Fig. 1(b). This means that the abnormal increase in the network traffic confirms the abnormal behavior in the system call sequence, therefore, the IDS can report an intrusion warning with a higher confidence.

A more detailed demonstration of how log correlation can improve the intrusion detection process was presented in [1].

4 UCLog Architecture

In this paper, we propose a unified logging architecture that captures the correlations among activities from different components and is closely integrated into the intrusion detection system. UCLog consists of several components as shown in Fig. 2: loggers, monitors, a centralized activity database, and a background rule generator. The loggers and monitors form a middleware layer that enables the IDS to use information held by different and heterogeneous logs and correlate it to improve its accuracy. Loggers and monitors should be coupled with corresponding system modules such as network, user shell, file systems, and individual applications. Components in the UCLog architecture could communicate with each other using secure messages with authentication [6, 36].

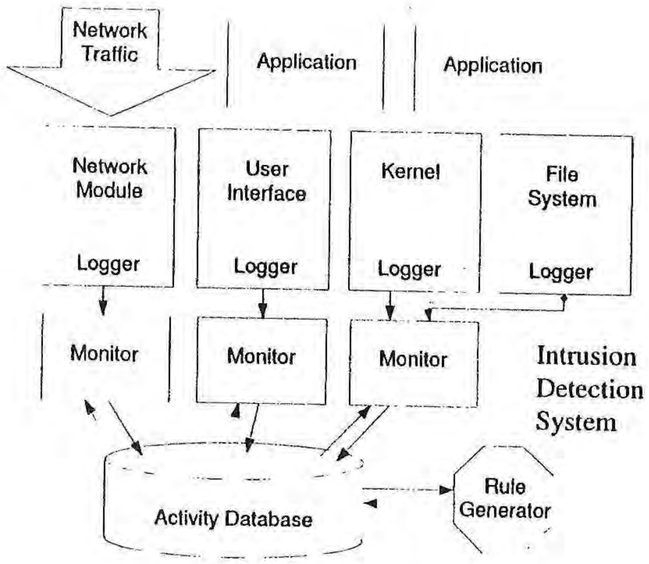
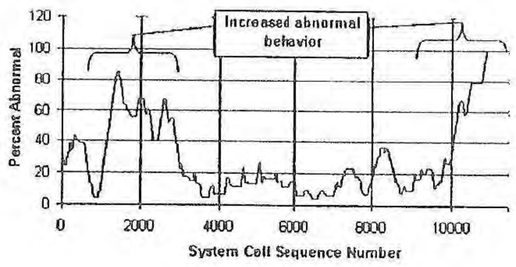


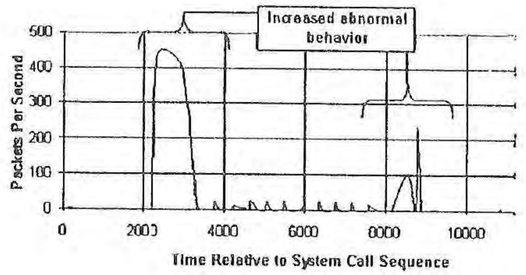
Figure 2. UCLog architecture

Loggers collect information about various aspects of system behavior and send this data to *monitors*. Monitors is watch system behavior and collaborate with each other through the centralized *activity database*. Monitors filter and compress audit data collected by loggers, and compare this data with existing models or signatures. The centralized activity database stores and manages audit data in an organized way. It also supports queries from a monitor to find out correlated activities watched by other monitors, or to trace to the specific processes or users that directly or indirectly cause the occurring intrusion. The *rule generator* produces the models for normal or abnormal behavior and signatures for previous intrusions. The rule generator can run either in the background or off-line, and either on the same or possibly several other machines. In the following subsections, we describe each component in more detail.

Sophos Ex. 1009
Page 69 of 318



(a) System call trace



(b) Network traffic trace

Figure 1. Abnormality in Yaha virus traces

4.1 Loggers

A logger records information about the behavior of a particular part of the environment and transmits that data to the corresponding monitor. For ease of deployment, a middleware layer of loggers with a common programming interface should be implemented. Thus, hiding the details of how to process each log from the other components. Each logging module must be coupled with the subsystem from which the logger is collecting audit data. For example, the network subsystem should have a logger module to collect network activities. Similarly, the user command shell also needs a logger to record user-entered commands. Several particularly useful logging modules for intrusion detection include:

- **Kernel API Logger**: A kernel API logger should record system calls, the parameters of those calls, and the time at which they are made.

- **Network Logger**: A network logger could record all packets passing through the system, all packet headers received by the system, or even record higher, summary-level information (e.g. Cisco NetFlows [10]). The actual implementation would depend on the desired circumstances.

- **File-System Logger**: A file-system logger might detect changes to significant blocks, such as the password file, or detect unusual access patterns. For example, a process which attempts to touch every block once could be harmless, such as a disk defragmentation process, or it could be some user attempting to illegally duplicate the disk.

- **User Behavior Logger**: It would be useful to have some module that records a user's interactions with the system, such as when each user logs in and out.

- **System Events Logger**: A facility that monitors the children, files, and directories created by some process would be useful in many situations.

- **Application-Specific Logger**: Additional application-specific loggers can be added as needed.

Each of the types of loggers listed can collect host-specific or system-wide information. The latter are *distributed loggers*. Distributed loggers are very useful and can be implemented to log system-wide information. For example, it is very common for organizations to have dedicated servers that collect and store logs obtained from different sources (e.g. *syslogs* obtained from each host in the organization). Obtaining a system-wide view is very helpful for intrusion detection, as related

Sophos Ex. 1009
Page 70 of 318

malicious activities may appear in several hosts. Correlating this information can help detect attacks that otherwise may not be obvious.

To reduce the performance overhead, loggers do very little processing but instead stream the data to the corresponding monitor. Some loggers may only be used when fine-granularity logging is required, e.g. an attack is suspected.

4.2 Monitors

Monitors receive audit data from one or more loggers and are responsible for filtering and classifying the stream entries, and detecting any abnormal behavior. Besides loggers, a monitor needs to interact with the centralized activity database to find correlated activities from other sources. Moreover, a monitor also has to consult with the rule generator to classify abnormal behavior. In addition, monitors also communicate with each other to confirm intrusions.

When a monitor receives audit data from loggers, it needs to reduce the data size by pruning unnecessary information such as highly variable parameters, user variable names, and any other unpredictable information. To further pack the audit data, the monitor can compress the packet statistics to connection-level information. A monitor may also combine several log entries from different loggers. For example, if several different loggers generate information from the same event (for example, a file I/O system call is recorded by both the system call logger and the file I/O logger. In this case, the monitor combines those logs to a unified format. The data the monitors store in the activity database must be stored in a unified format, so that the rule learning and processing module can understand any entry. There are several unified log formats that can be used for this purpose, for example, the one used by Snort [37], the one used by Harvester [17], the Common Intrusion Specification Language (CISL) [38], among others. Ideally the unified format should adhere to some standard, and it is believed that an XML format may become a standard [28].

Once audit data is properly filtered, a monitor compares the stream to previously generated models or signatures by using data-mining techniques such as outliers analysis [19] or machine learning techniques. More specifically, a monitor would examine the audit data sequence and check whether a subsequence matches an intrusion signature or violates a certain model that captures normal behavior. For example, a simple model could be "file activities generated by ordinary users are fewer than 5 I/Os per second during 12am to 6am". Then if the file I/O monitor detects many file accesses from an ordinary user, it can generate a warning that an intrusion is likely taking place.

Classifying information in an isolated way does not take full advantage of the semantic structure supported by the UCLog architecture. It is often useful to correlate the outputs of the monitors to strengthen the classification judgment. For example, a system call sequence may have a fraction of abnormal behavior. Before generating an intrusion alarm, it is better to consult with other components to see if correlated activities are also showing abnormal activities. In addition, when a stream is flagged as unusual, a monitor may need to query the activity database and the rule generator to determine how previous instances of this pattern were classified (if they occurred and were not part of the rule set).

4.3 Activity Database

The activity database stores and manages audit data in a very organized way. It is implemented as an efficient relational database management system in order to present monitors with an interface to query the database to find correlated activities or to trace to the root and the path of an intrusion. To prevent intrusions from breaking into the activity database, the activity database should run in a protected environment and should support only a secured query interface to monitors and the rule generator.

The activity database consists of a set of tables, each corresponding to the filtered log entries from a monitor. Each table describes some measurable and distinct aspect of system behavior. For example, in our prototype implementation, there is a table of system calls which records <*processID, processName, systemCall, date, timeStamp, etc*>, a table for network connections recording <*startTime, duration, sourceHost, destHost, bytesFromSource, bytesFromDest, processId, etc*>, and another for system behavior recording <*date, time, eventCode, pid, user, parentID*>.

Sophos Ex. 1009
Page 71 of 318

The database schema is crucial to correlate audit data from different sources. Correlations between logs exist when two log entries have the same or related values for one or more fields. For example, a network connection might be associated with a given system call if the time is within a certain distance and both belong to the same process. Two system calls are correlated if they are made by two different processes that are executed by the same user. Two file I/Os can be correlated if they are operating on the same file or the same directory. Any such correlations can be easily queried using a SQL statement.

In addition to the output of monitors, the activity database also stores "rules" (models and signatures) to describe appropriate or illegal system behavior. Since data can be retrieved in so many different ways from the activity database, it is possible to store several different granularities of a rule. In general, the most commonly used rule set would apply to any system-level behaviors; however, the activity database can also store rules describing how a particular instance (e.g., a user, an application or a file) should behave. For example, the activity database can store rules describing the network traffic generated by FTP or the SMTP daemon. Although instance-specific rules would be too expensive to use when filtering all the data streams, they would be of invaluable use in identifying (or ruling out) potential causes of anomalies.

Another key advantage of the UCLog architecture is the ability to trace back and identify the source and the execution path of the occurring intrusion. At an intrusion, a system administrator who is assigned to analyze the intrusion might ask questions like "What process was running when a large spike of network activity occurred yesterday at midnight?". In most current systems, the administrator needs to examine the logs manually in order to find out the answers to his/her questions. But in UCLog, the activity database supports various queries to provide answers to these questions. To further simplify the administrator's job, some simple tools can be built on top of the activity database with a user-friendly interface and relieve the administrator from writing SQL statements.

One of the most challenging issues of the activity database is scalability. If the system is running for a long time or on a large-scale network, the size of database may quickly increase to Giga or Tera bytes. One solution toward reducing the footprint of the activity database is to investigate algorithms to compress old events and remove obsolete events. Another solution is running the activity database on a dedicated server.

4.4 Rule Generator

The rule generator is responsible for constructing the models for normal behavior and signatures for previous intrusions. The behavior models are used to detect any anomalies, while the signatures are used to detect known intrusions [4].

To model normal behavior, the rule generator retrieves the training data from the activity database. It is convenient to query particular data from the activity database to generate instance-specific rules via SQL statements. For example, to generate rules for normal system call sequences of a specific application in a specific time period, the generator only needs to query the data of normal system call logs with the WHERE conditions of *processName* and *timeStamp* of an SQL statement. Then the training data is formatted and fed to modeling algorithms. In our current implementation, we use a data-mining algorithm, RIPPER, to generate rules for system call sequences, and statistics methods for other logs.

In contrast, since a signature-based IDS depends on the nature of various known intrusions, most intrusion signatures may come from outside of the system, possibly from an anti-virus library. The rule generator can import such signatures from other systems as well as from intrusions captured as part of training data. Data-mining techniques suitable for the latter.

The rules can be stored in the database for future use, or directly returned to the monitor that initiates the generating rule query. When applying the models to anomaly detection, the IDS can decide what percentage of the activity to flag as abnormal by verifying if the activity follows the behavior models.

4.5 Application Programming Interface

An application programming interface (API), to allow an application or a logger to communicate with a specified monitor, will be implemented in future versions. It would also be useful to incorporate some

Sophos Ex. 1009
Page 72 of 318

mechanism for specifying non-trivial correlations between two audit tables. In our current implementation, logging imposes only semantic information to the logged records; correlation occurs only at query time. Future versions may include some mechanism for the user to specify a relationship between tuples or fields.

5 Evaluation Methodology

The goal of our evaluation is to answer the following questions: "Can UCLog improve the accuracy of intrusion detection?", "Is it efficient?", "Does it impose high space overheads?".

5.1 Experimental Architecture

We have implemented the proposed architecture on two 2.56 GHz, Pentium 4 PCs with 512 MB RAM. Each machine has a 32 GB hard drive with an Intel Ethernet Adapter. Machine 1 runs Red Hat Linux 7.3 and acts as the e-mail server. Machine 2 runs Windows 2000 Service Pack 3[1], serving as a client machine. The UCLog architecture is implemented on Machine 2.

We use several different loggers to capture data from multiple vantage points. For network behavior we use WinDump [14], a Windows port of the popular tcpdump. We use TDIMonitor version 1.01 [34], a logging utility that monitors the Transport Driver Interface. The information provided by WinDump and TDIMonitor are combined into the network traffic monitor and the activity database in order to associate a particular process ID with each network connection. To capture system call behavior we use APISpy32 v3.0 [21] with some modifications that allow us to capture the system time of each system call. Due to some limitations of APISpy32 we also use FileMon v5.01 [35] to capture certain file I/O operations. We use the built-in Windows 2000 security logging features in order to capture process information (start time, end time, process ID, and parent process ID) that we later use in correlating the results from different logs.

[1] Earlier versions of Windows contain a bug that prevents us from retrieving the correct process ID for each activity (see the Microsoft Knowledge Base Article 277743).

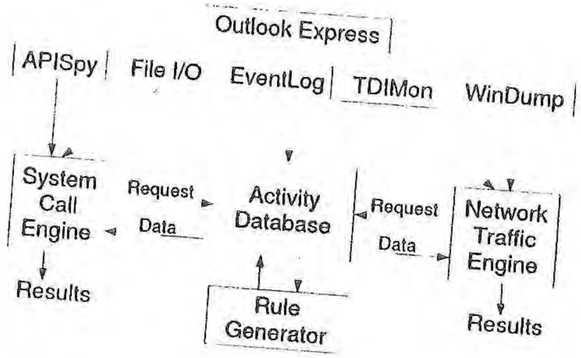


Figure 3. Experimental architecture

In Fig. 3 we give an illustration of our testing architecture. The architecture consists of five loggers, two monitors, a centralized activity database and a rule generator. The data from the five loggers is stored in the central database. While system calls can be logged without any processing, network traffic requires some pre-processing. The network traffic is stored on a connection basis. All of the traffic for a particular connection is aggregated together and stored as one entry in the table. When the data is retrieved from the database it is also necessary to perform some calculations, such as number of connections within the past 10 seconds, in order to help detect an intrusion.

5.2 Intrusion Detection Through Data Mining

While others have used data mining techniques to detect intrusions [23], to our knowledge none of them have correlated information from multiple logs in the data mining process.

When using these data mining techniques to distinguish between an intrusion and normal behavior, it is necessary to capture logs that describe both behaviors. We selected Outlook Express as the entry point for our intrusion. First, we use the loggers to capture several hours of *normal* end-user behavior to use in the training phase of our IDS. The traces record around 800,000 activities. Then, we captured several hours of *abnormal* behavior. Abnormal behavior consists of opening an infected e-mail message and executing the virus. We take roughly 80% of the data points from each set to train the data mining-based rule

Sophos Ex. 1009
Page 73 of 318

generator and use the remaining 20% of the data for testing.

The training process is performed similarly to Lee's study [23]. The system call sequence is broken up into smaller subsequences of 13[2] system calls for training, and each new subsequence is labeled as either normal or abnormal, depending upon which trace it comes from. It is necessary to understand that a normal trace is likely to contain some abnormal subsequences, but the percentage is quite small. Also, an abnormal trace will contain some normal system call subsequences since opening an infected e-mail message does contain some normal operations. These subsequences are fed into the data mining tool RIPPER [12, 13]. RIPPER takes the system call sequences and generates a set of rules that is used to classify new sequences as either normal or abnormal. Once this training has taken place we are ready to start detecting intrusions.

5.3 Log Files

For testing, we used five loggers to collect audit data. Each logger produces a lot of entries; therefore, some filtering is necessary before adding the data into the activity database.

- **System Call Log** – Generated by APISpy32, contains a list of system calls executed by all running processes on the system. We extract the system call, process ID, and time of execution.
- **WinDump Log** – Generated by WinDump. It contains the contents of the UDP and TCP packet headers of all traffic to and from the client machine.
- **TDIMonitor Log** – Generated by TDIMonitor. Captures network traffic at a different layer than WinDump. This logger is able to associate a process ID with each connection. This information is combined with the WinDump log to give us a process ID with each packet header.
- **File I/O** – Created by FileMon. It records all of the file I/O on the system. This information is combined with the system call log and stored in the database.
- **System Events** – Contains the system events generated by Windows 2000. We configured this log to record information on the creation and termination of all processes. This information includes process ID of each process and its parent process, along with the user-name that created the process.

[2]Our tests showed that size 13 provided optimal results.

5.4 Experiment Procedures

We compare the results of UCLog with a baseline IDS. UCLog and the baseline system are very similar. The primary difference between the two is that UCLog explores correlations by using the activity database, whereas the baseline system does not use correlations. The baseline system is very similar to the one used by Lee [23]. In our comparison, we mainly examine the abnormal percentage for both normal and abnormal runs.

In our first set of experiments we measure the baseline results just using the IDS without taking advantage of correlations. We use the 20% of traces that are not used in training as our testing data. The rules that RIPPER generated are applied to the traces. Each sequence of 13 system calls is classified as either normal or abnormal according to the rules. Next, a window of size 7[3] is passed over the list of sequences. If more than half of the sequences in that window are determined to be abnormal, then that region is considered abnormal. Once that window passes over all of the data we can determine what percentage of the trace is "abnormal". For a trace that does not contain an intrusion, we expect the percentage of abnormality to be quite small. If an intrusion takes place, we expect the percentage of abnormality to be noticeably larger. This procedure is similar to many other data mining-based intrusion detection systems [3, 23].

In our second set of experiments we measure the results of using the UCLog architecture to exploit correlations. Once again, we use the 20% of the traces that were not used in training the IDS. The network monitor observes the network traffic and in the meantime also queries the activity database to check if correlated system call data has shown similar behavior. The network monitor combines the system call behavior and the network behavior and makes a judgment on whether

[3]On our tests, a post processing window of size 7 provided optimal results.

Sophos Ex. 1009
Page 74 of 318

the entire system is suspicious. If only one of the network or system call monitors shows abnormal behavior, the threshold for abnormality is decremented, and vice versa.

6 Evaluation Results

Our results show that UCLog can significantly improve the accuracy of intrusion detection. In abnormal runs, UCLog can improve the confidence level of abnormal behavior. In normal runs, UCLog can decrease the number of false positives. Moreover, UCLog is relatively efficient and has reasonable space requirements.

Figure 4 compares the results without correlations (baseline) and with correlations (UCLog). The numbers are measured in terms of percentage of abnormal regions. The higher the percentage, the more likely an intrusion has occurred. If we compare two IDSs, the better IDS should give a higher abnormal percentage for an abnormal trace, and give a lower abnormal percentage for a normal trace. The results show that UCLog performs better than the baseline system and improves the intrusion detection accuracy.

For abnormal traces, Fig. 4(a) shows that the percentage of abnormality is higher when correlation is used. For example, in a30, UCLog shows 31% abnormality, versus 19% for the baseline system. This indicates that UCLog has higher confidence in reporting intrusion warnings.

For normal traces, UCLog reports lower abnormality than the baseline system (see Fig. 4(b)). For some traces, the percentage abnormality with UCLog is significantly lower. For example, with n10, the baseline system reports 1.4% abnormality, whereas UCLog reports only 0.8% abnormality. Therefore, if 1% is the cutoff, the baseline would report a false intrusion but UCLog would not.

6.1 Monitoring File I/O

We have also examined whether monitoring file I/O can be useful in improving intrusion detection accuracy. Pennington et al. [29] presented a storage-based IDS, which monitors access to specific files or directories. The administrator is able to describe which files or directories the IDS should monitor and the specific restrictions for that file or directory. We implement a similar monitor for our architecture to demonstrate how monitoring and correlating an additional source can help to improve the IDS accuracy.

We added a file I/O monitor that watches the *c:\recycler* directory. The Yaha virus used in our experiments creates several files inside this directory. This is actually a common technique that current worms and viruses often use to conceal their whereabouts. Our new file I/O monitor watches any I/O into this directory and generate warnings upon file creations.

The Windows address book is often compromised by a virus and used to find other potential victims. The Yaha virus reads the address book and sends itself to the contacts. We decided that for testing purposes, no process should be able to secretly access the address book other than Outlook Express or Explorer. In fact, in recent versions of Outlook if you try and access the Contacts information from another program you are prompted for permission. So, our new monitor also looks for any access to the address book by an unauthorized process.

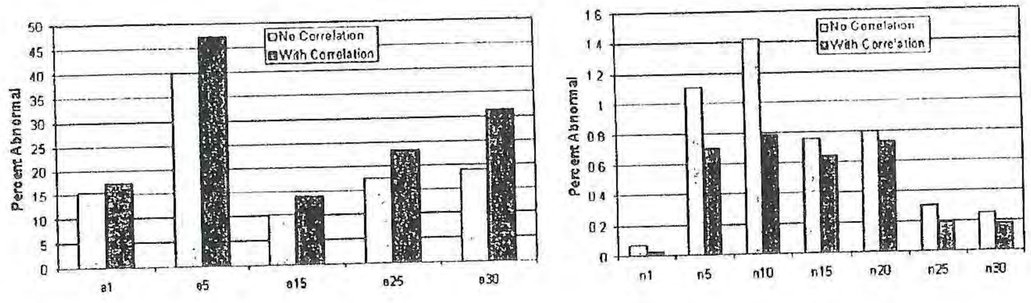
The information from our file I/O monitor is correlated with the other logs in the same way as we correlated the network traffic logs. If we detect any files created in the *c:\recycler* directory, or if we detect any unauthorized access to the Windows address book, we query the database for the network traffic information and the system call information. Then we reprocess the system call data taking into account any anomalies in the network traffic along with the new information from the file I/O monitor. If the file I/O monitor detects the expected abnormal behavior, it increases the abnormality number.

Figure 5 shows the results with the new file I/O monitor. Each abnormal trace contained the suspicious behavior that we were monitoring for. The normal traces did not contain any of the suspicious behavior. Notice that as we correlate more logs, our ability to detect intrusions increases. For example, in a30, the abnormality is 48% with correlations, whereas the abnormality without correlations only gives 19%, a factor of improvement of 2.5.

6.2 Query Performance

To evaluate the usefulness of our trace-back capabilities, we measure the average time to perform several common queries to the activity

Sophos Ex. 1009
Page 75 of 318



(a) Abnormal traces

(b) Normal traces

Figure 4. Comparison of UCLog with baseline system. Figures are not in the same scale

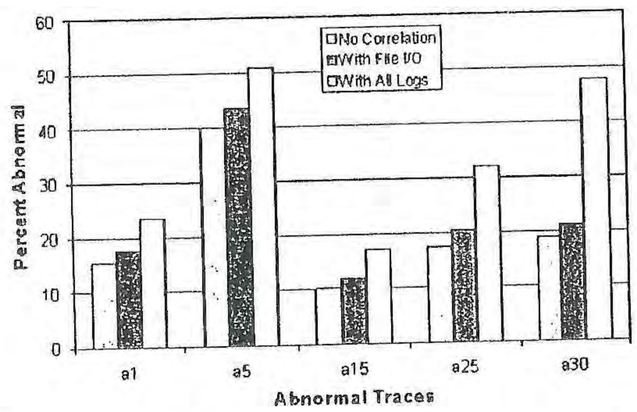


Figure 5. Percentage of abnormality when using no correlation, when correlating file I/O with system calls and when correlating all three logs

database. This was done by selecting several different parameters for each relevant query, and performing multiple runs for each set of parameters. The results are summarized in Table 1. All four simple queries (1–4) return results within 0.01 to 0.43 seconds, which indicates the UCLog architecture can be used to detect intrusions in a timely manner. Query 5 combines two of the other queries, as might be done while investigating an attack. Even for more advanced queries such as 5, the average time to run them is still small.

Table 1. Micro-benchmark results

ID	Query	Avg. Time
1	Select all system calls occurring in some time range	0.196 sec
2	Select all system calls belonging to some process in a given trace	0.425 sec
3	Select all program names belonging to some parent	0.011 sec
4	Select all process IDs receiving packets in some time range	0.082 sec
5	Query 4 and then Query 1 (with processID filtering)	0.6 sec

6.3 Space Requirement

In addition to facilitating queries and detection, the structured database format is more space efficient than storing raw logs. The sum of the three tables stored requires 90,561,664 bytes. In contrast, the total size of the unprocessed log files is 720,896,371 bytes. This 87%

Sophos Ex. 1009
Page 76 of 318

reduction in the size of data that must be stored is a result of eliminating duplicate data especially between the APISpy and FileMon logs, and as a result of removing the data packets from the WinDump log.

7 Related Work

Our work builds upon a large body of previous work. ReVirt [16] proposed by Dunlap et al. is one such example. It records all actions below the virtual machine. It also uses fault-tolerance techniques to protect the integrity of the logging system, and allow ReVirt to replay execution instruction-by-instruction even if the kernel of the operating system is compromised. However, their study does not address how an intrusion would be detected. The UCLog architecture would fill this missing part in the ReVirt system to detect intrusions.

Another related work is secure storage. The self-securing storage system (S4) [39] maintains multiple file versions for recovery and logs all operations for auditing. Based on S4, Pennington et al. [29] propose storage-based intrusion detection that allows storage systems to transparently monitor abnormal activity. In this system, the IDS embedded in a storage device continues to work even if the client systems are compromised. Their study only monitors file I/O operations, whereas UCLog monitors activities from multiple sources.

Schneier and Kelsey [36] describe a computationally cheap method to prevent log files from being read or undetectably compromised by the attackers, by using encryption, authentication, and one-way hash functions. This technique would be useful to integrate into UCLog to ensure the integrity of audit data.

Much work on OS-based intrusion detection has focused on approaches that exploit system call interposition and tracing [18, 20, 23, 22, 7, 40]. The first two studies detect anomalies by extracting patterns of system call sequences from traces and then monitoring for unknown patterns. Lee and Cabrera et al. [23, 7] employ data mining techniques to extract patterns from system call traces for intrusion detection. Warrender et al. [41] compare several methods of representing normal behavior of system calls, such as enumeration of observed sequences, methods based on relative frequencies of sequences, data mining techniques, and hidden Markov models. Ko et al. [22] integrate intrusion detection techniques with software wrapping technology to improve effectiveness. Wagner et al. [40] propose to detect intrusions via static analysis. Their idea is to constrain the system call trace of a program's execution to be consistent with its source code. When an intrusion occurs, it causes the program to behave in a manner inconsistent with the intent of its author. The biggest difference between Wagner's and our IDSs is that our intrusion detection model is built from the normal system call trace itself using logs, while Wagner's specification is constructed from the program code by static analysis. In many cases, the source code is not provided by the vendor. Even if the source code is available, the limitation of compilers such as the common aliasing problem, makes this solution not applicable to many large commercial server programs.

In [24], Lee et al. propose an architecture for a real-time based system. Such a system consists of sensors that filter and format data, detectors that analyze sensor data for attacks, a data warehouse which provides models of attacks for use by the detectors, and a module to update the data warehouse. Although these sensors have a standardized interface, correlations are not exploited.

Among others, the IDES project and the CIDF [38] standard would allow intrusion detection systems to use multiple kinds of log files. The IDES project, an early design for an intrusion-detection system, formats logged data so that it can be organized by user behavior [15]. The description given in [15], however, is more a high-level discussion of what log information would be useful for an IDS, rather than a consideration of how multiple kinds of logged data can be practically correlated. CIDF [38], the Common Intrusion Detection Framework standard, has been proposed to coordinate distinct intrusion detection systems. CIDF follows a model similar to [24]. Components share tagged data, using a Lisp-like unified format. There is evidence that an XML protocol may become standard [28]. While this standard would allow log files to be used interchangeably within a given IDS; our goal is to allow multiple logs to be used together.

Logging has also been used to improve system reliability [9, 33, 32, 30]. DejaVe [9] is a record/replay tool for Java. By logging the logical thread schedule information, DejaVe can force the exact same schedule during re-executing. Ronsse et al. [32] propose RecPlay that combines record and replay techniques so that it allows debuggers to

Sophos Ex. 1009
Page 77 of 318

debug parallel programs in a deterministic mode.

Recently, Wu et al. [42] proposed CIDS, a Collaborative Intrusion Detection System that combines alarms from multiple IDSs (Snort, Libsafe and Sysmon) and uses a Bayesian network to aggregate and combine the alarms to improve accuracy. The input data, learning method and architecture differs from UCLog, but the basic idea is the same: to improve IDS accuracy by correlating information from different and preferably disjoint sources. We plan to compare UCLog with CIDS in the future.

8 Future Work

Our work has several limitations. First, we have only evaluated our system in Windows 2000 with only the Outlook Express application. It would be interesting to evaluate other systems and other applications. Second, we have not yet integrated our system with a rollback-capable system such as S4. This remains as our immediate future work. Third, it would be interesting to see whether multi-stream data mining would be helpful to exploit correlations. Lastly, we have not compared UCLog with other intrusion detection methods such as the ones using compiler analysis. We believe both static analysis and dynamic analysis of log files would be useful for intrusion detections and can complement each other.

9 Conclusions

This paper proposed UCLog, a unified logging architecture that effectively explores activity correlations for intrusion detection. UCLog organizes audit data from heterogeneous sources into a structured way and is closely integrated into the intrusion detection system. It also exploits activity correlations between different components to improve the intrusion detection accuracy. Moreover, it can trace back to the source and the execution path of an intrusion, relieving administrators from manually examining the logs to find such information. By extending the UCLog architecture by adding additional specific loggers and monitors, the user can tailor the IDS and make it closely monitor critical applications.

We have implemented UCLog on a real system running Outlook Express on Windows 2000. Our prototype architecture uses MySQL to manage the activity database and data mining tools to generate rules and detect anomalies. We evaluated UCLog using normal and abnormal traces. Using these traces, we studied how correlating activities from the system call sequence, network activity, and file I/O can help improve the IDS accuracy. Compared to a baseline system that only monitors a single component, our system that correlates system call sequences, network traffic and file I/Os improves the accuracy by up to a factor of 2.5. UCLog also reduces false alarms when the system is perfectly normal.

References

[1] C. Abad, J. Taylor, C. Sengul, W. Yurcik, Y. Zhou, and K. Rowe. Log correlation for intrusion detection: A proof of concept. In *Proc. of the 19th Computer Security Applications Conf. (ACSAC 2003)*, Dec. 2003.

[2] J. Allen, A. Christie, W. Fithen, J. McHugh, J. Pickel, and E. Stoner. State of the practice of intrusion detection technologies. Technical Report CMU/SEI-99TR-028, School of Computer Science, Carnegie Mellon University, 2000.

[3] S. Axelsson. Research in intrusion-detection systems: A survey. Technical Report 98–17, Dep. of Computer Engineering, Chalmers University of Technology, Dec. 1998.

[4] S. Axelsson. Intrusion detection systems: A survey and taxonomy. Technical Report 99–15, Dep. of Computer Engineering, Chalmers University of Technology, Mar. 1999.

[5] D. Barbará, J. Couto, S. Jajodia, and N. Wu. ADAM: A testbed for exploring the use of data mining in intrusion detection. *ACM SIGMOD Record*, 30(4):15–24, 2001.

[6] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In *Proc. of the 16th Intl. Cryptology Conf. (CRYPTO 1996)*, Aug. 1996.

[7] J. Cabrera, L. Lewis, and R. K. Mehra. Detection and classification of intrusions and faults using sequences of system calls. *ACM SIGMOD Record*, 30(4):25–34, 2001.

Sophos Ex. 1009
Page 78 of 318

[8] CERT Coordination Center, Carnegie Mellon University. CERT/CC overview incident and vulnerability trends, May 2003. <http://www.cert.org/present/cert-overview-trends/>. (Feb. 2004).

[9] J.-D. Choi and H. Srinivasan. Deterministic replay of java multithreaded applications. In *Proc. of the 2nd SIGMETRICS Symp. on Parallel and Distributed Tools*, Aug. 1998.

[10] Cisco Systems, Inc. Netflow services and applications, a whitepaper, July 2002. <http://www.cisco.com/warp/public/cc/pd/iosw/ioft/neflct/tech/napps_w%p.htm>. (Feb. 2004).

[11] C. Clifton and G. Gengo. Developing custom intrusion detection filters using data mining. In *Proc. of the Military Communications Intl. Symp. (MILCOM 2000)*, Oct. 2000.

[12] W. W. Cohen. Fast effective rule induction. In *Proc. of the 12th Intl. Conf. on Machine Learning*, July 1995.

[13] W. W. Cohen. Learning trees and rules with set-valued features. *AAAI/IAAI*, 1:709–716, 1996.

[14] L. Degioanni, F. Risso, and P. Viano. Windump: tcpdump for Windows, Jan. 2002. <http://windump.polito.it/>. (Feb. 2004).

[15] D. E. Denning. An intrusion-detection model. *IEEE Transactions on Software Engineering*, 13(2):222–232, 1987.

[16] G. W. Dunlap, S. T. King, S. Cinar, M. A. Basrai, and P. M. Chen. ReVirt: Enabling intrusion analysis through virtual-machine logging and replay. In *Proc. of the 5th Symp. on Operating Sys. Design and Implementation (OSDI 2002)*, Dec. 2002.

[17] farm9.com, Inc. Harvester architecture overview, Aug. 2002. <http://farm9.org/Harvester/ArchitectureOverview.php>. (Feb. 2003).

[18] S. Forrest, S. A. Hofmeyr, A. Somayaji, and T. A. Longstaff. A sense of self for Unix processes. In *Proc. of the IEEE Symp. on Research in Security and Privacy*, May 1996.

[19] J. Han and M. Kamber. *Data mining: concepts and techniques.* Morgan Kaufmann Publishers Inc., 2000.

[20] S. A. Hofmeyr, S. Forrest, and A. Somayaji. Intrusion detection using sequences of system calls. *Journal of Computer Security*, 6(3):151–180, 1998.

[21] Y. Kaplan. Api spying techniques for windows 9x, nt, and 2000, Apr. 2000. <http://www.internals.com/ar-ticles/apispy/apispy.htm>. (Feb. 2003).

[22] C. Ko, T. Fraser, L. Badger, and D. Kilpatrick. Detecting and countering system intrusions using software wrappers. In *Proc. of the 9th USENIX Security Symp.*, Aug. 2000.

[23] W. Lee and S. Stolfo. Data mining approaches for intrusion detection. In *Proc. of the 7th USENIX Security Symp.*, Aug. 1998.

[24] W. Lee, S. Stolfo, P. Chan, E. Eskin, W. Fan, M. Miller, S. Hershkop, and J. Zhang. Real time data mining-based intrusion detection. In *Proc. of the DARPA Information Survivability Conf. and Exposition (DISCEX II)*, June 2001.

[25] R. Lemos. Counting the cost of slammer. In *CNET News.com*, Jan. 2003.

[26] D. E. Lowell and P. M. Chen. Free transactions with Rio Vista. In *Proc. of the 16th ACM Symp. on Operating Sys. Principles (SOSP 1997)*, Oct. 1997.

[27] Microsoft Corporation, MSDN Library. Logging application, server, and security events, 2004.

[28] S. Northcutt and J. Novak. *Network Intrusion Detection: An Analysts' Handbook.* New Riders, 2000.

[29] A. Pennington, J. Strunk, J. Griffin, C. Soules, G. Goodson, and G. Ganger. Storage-based intrusion detection: Watching storage activity for suspicious behavior. Technical Report CMU-CS-02-179, School of Computer Science, Carnegie Mellon University, Oct. 2002.

[30] M. Prvulovic, Z. Zhang, and J. Torrellas. ReVive: Cost-effective architectural support for rollback recovery in shared-memory multiprocessors. In *Proc. of the 29th Intl. Symp. on Computer Architecture (ISCA 2002)*, May 2002.

[31] L. Rohde. Yaha virus lingers into the new year, Jan. 2003. IDG News Service.

[32] M. Ronsse and K. D. Bosschere. RecPlay: A fully integrated practical record/replay system. *ACM Transactions on Computer Sys. (TOCS)*, 17(2):133–152, 1999.

[33] M. Ronsse and W. Zwaenepoel. Execution replay for treadmarks. In *Proc. of the 5th Euromicro Workshop on Parallel and Distributed Processing*, Jan. 1997.

[34] M. Russinovich. TDIMon, July 2000. <http://www.sysinternals.com/ntw2k/utilities.shtml>. (Feb. 2004).

Sophos Ex. 1009
Page 79 of 318

[35] M. Russinovich and B. Cogswell. Filemon for Windows, July 2003. <http://www.sysinternals.com/ntw2k/utilities.shtml>. (Feb. 2004).

[36] B. Schneier and J. Kelsey. Secure audit logs to support computer forensics. *ACM Transactions on Information and Sys. Security (TISSEC)*, 2(2):159–176, 1999.

[37] Snort Project, The. Snort: The open source network intrusion detection system, June 2004. <http://www.snort.org>. (Jun. 2004).

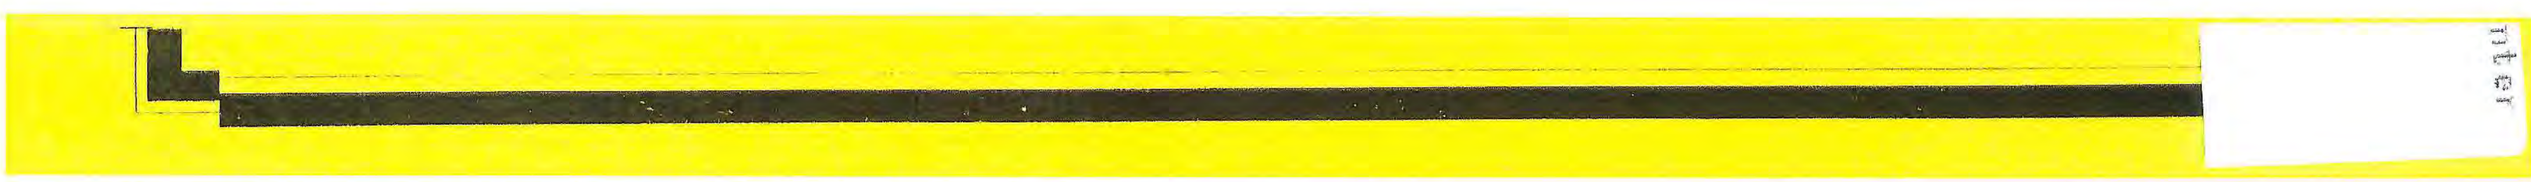
[38] S. Staniford-Chen, B. Tung, and D. Schnackenberg. The common intrusion detection framework (CIDF). In *Information Survivability Workshop*, Oct. 1998.

[39] J. D. Strunk, G. R. Goodson, M. L. Scheinholtz, C. Soules, and G. R. Ganger. Self-securing storage: Protecting data in compromised systems. In *Proc. of the 3rd Symp. on Operating Sys. Design and Implementation (OSDI 2000)*, Dec. 2000.

[40] D. Wagner and D. Dean. Intrusion detection via static analysis. In *IEEE Symp. on Security and Privacy*, May 2001.

[41] C. Warrender, S. Forrest, and B. A. Pearlmutter. Detecting intrusions using system calls: Alternative data models. In *IEEE Symp. on Security and Privacy*, May 1999.

[42] Y.-S. Wu, B. Foo, Y. Mei, and S. Bagchi. Collaborative intrusion detection system (CIDS): A framework for accurate and efficient ids. In *Proc. of the 19th Computer Security Applications Conf. (ACSAC 2003)*, Dec. 2003.



Sophos Ex. 1009
Page 80 of 318

ATTACHMENT G

Instituto Tecnológico y de Estudios Superiores de Monterrey

Monterrey Campus

Doctoral Program in Artificial Intelligence



A Multiagent Approach to Outbound Intrusion Detection

PhD thesis

by

Salvador Mandujano Vergara

December – 2004

Sophos Ex. 1009
Page 83 of 318

Instituto Tecnológico y de Estudios Superiores de Monterrey

Monterrey Campus

Graduate Program in Electronics, Computing, Information and Communications

We, the undersigned, as members of the evaluation committee, hereby recommend the thesis entitled "A Multiagent Approach to Outbound Intrusion Detection" defended by Salvador Mandujano Vergara on Friday August 6th, 2004, to be accepted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Artificial Intelligence.

Arturo Galván Rodríguez, PhD (advisor)

David Garza Salazar, PhD
Chair, Graduate Program in
Electronics, Computing,
Information and Communications

Ramón Brena Pinero, PhD

Jesús Arturo Pérez Díaz, PhD

Gerardo Ayala San Martín, PhD

Pascal Meunier, PhD

Monterrey, Mexico, December – 2004

Sophos Ex. 1009
Page 85 of 318

Acknowledgments

I must thank all the support I have received from professors, researchers and fellow students during these years of graduate school. This dissertation would have never been possible without their help.

Firstly, I acknowledge the authorities of ITESM for providing me with the fellowship that allowed me to enroll into this doctoral program and for the teaching and research opportunities that help us develop skills beyond the classroom. I also thank CONACyT for the funding I was awarded to study my Master's. This dissertation is strongly influenced by ideas and motivation originated during that period.

I very specially thank the members of my dissertation committee, Dr. Ramón Brena, Dr. Gerardo Ayala, Dr. Jesús Arturo Pérez, and Dr. Pascal Meunier, who were so kind in accepting our invitation to review this work. Their thorough feedback and suggestions helped me deliver a better document and identify research ideas I will definitely consider in future projects.

While at ITESM, I worked under the supervision of Dr. Arturo Galván, my principal advisor. I thank him for providing me with the guidance and freedom that allowed me to enjoy the exploration of new ideas in the field of computer security. His technical competence and values as a person are equally admirable and I hope many more students have the chance to work with him.

I thank Dr. Rogelio Soto for being so helpful in aligning the resources that allowed me to work with the people of the ITESM Center for Intelligent Systems. For his follow-through during my studies and for supporting my research, my sincere gratitude.

Dr. Juan Arturo Nolazco has been instrumental in many respects to my doctoral studies. He received me as an assistant in his information security group and supported the research line my advisor and I were trying to pursue. I am profoundly thankful to him for the confidence placed in me and for the opportunities I had to develop my work and to contribute to the growth of the group.

I thank all my ITESM professors for introducing me into the world of artificial intelligence whose different areas are worth a lifetime. Thanks Dr. Francisco Cantú, Dr. Hugo Terashima, Dr. Ramón Brena, Dr. José Luis Gordillo, and Dr. Manuel Valenzuela.

I must especially thank Dr. Robbert van Renesse for having me as a visiting scholar at the Department of Computer Science of Cornell University. Working at Cornell was an amazing learning experience that would have not been possible without his interest and support. I thank Robbert for believing in my proposal and for encouraging my interaction with faculty and students. I also thank Dr. Emin Gun Sirer, Dr. Rich Caruana, and Dr. Mark Bickford for their comments and for the invaluable time they so kindly devoted to me.

Some Purdue people got me really involved into computer security during my Master's and encouraged me to pursue this degree. I must thank Dr. Sam Wagstaff for being my advisor, Dr. Sunil Prabakhar and Dr. Pascal Meunier for their support and for the opportunity to work with their teams, Dr. Clay Shields for his invaluable lessons, and Dr. Mike Atallah for his incredible cryptography class and for his guidance in our research projects.

The very first thought I had of a career in research appeared when I was studying high school at UNAM's Escuela Nacional Preparatoria thanks to the encouragement of Prof. María Elena Montoya, to whom I will be forever indebted, and the Program for Early Initiation to Research (*Programa de Iniciación Temprana a la Investigación*). This motivation developed throughout college at UNAM where I met Prof. José Origel, my thesis mentor, Dr. Francisco Ugalde, and Dr. Diego Zamboni who were role models that helped me find the determination to aim for higher intellectual goals.

I must also thank Cheryl, my girlfriend, for her patience and understanding, and for the immense love she always shows for me, and all my friends and classmates for their help and for the moments we shared together. In alphabetical order, I thank Alejandra, Bonnie, Carlos, Charlie, Doris, Eduardo, Emily, Emmanuel, Evy, Fernando, Filip, Germán, Guanfeng, Héctor, Leonardo, Lingfeng, Michael, Patrick, Rossy, Saurabh , Tineke, Víctor and Yulianto.

Last but not least, I thank my family for all their love and for always supporting my decisions. I thank Felicia, my sister, for her words and strength, and Eduardo, my brother, for showing me a different way of understanding this journey called life. To my parents, Beatriz and Salvador, all my love and respect for the solid principles with which the raised us and for their everlasting support. Words will never suffice to thank them enough.

Salvador Mandujano Vergara
Monterrey, Mexico, December – 2004

To Beatriz, my mother.

To Salvador, my father.

To Felicia, my sister.

To Eduardo, my brother.

A Multiagent Approach to Outbound Intrusion Detection

A dissertation submitted by

Salvador Mandujano Vergara

to the faculty of the

Graduate Program in Electronics, Computing, Information and Communications

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Artificial Intelligence

Instituto Tecnológico y de Estudios Superiores de Monterrey

Monterrey, Mexico **December – 2004**

Sophos Ex. 1009
Page 91 of 318

Abstract

"A Multiagent Approach to Outbound Intrusion Detection". Ph.D. dissertation by Salvador Mandujano Vergara, Instituto Tecnológico y de Estudios Superiores de Monterrey. Advisor: Prof. Arturo Galván. December – 2004.

This is a dissertation on the topic of intrusion detection. It supports the philosophy of system vigilance by exploring the concept of outbound intrusion detection, which is concerned with the identification and collection of evidence that helps prove local resources are being used to compromise external systems.

We discuss the motivation behind the approach, explain the need for splitting the scope of intrusion detection into sub-problems, and present trends in computer security that reveal basic design considerations that need to be taken into account when developing modern information security tools.

We propose a multiagent architecture for outbound intrusion detection supported by an ontology. Groups of agents collectively monitor outbound network traffic and local activity in order to identify references to neighboring systems that may be indicative of a compromise attempt.

We organize agents into sub-environments called agent cells that are connected to each other in a non-hierarchical fashion. Different classes of agents and cells compose the system, which performs attack modeling by employing multiple concurrent agents. Detection cells implement independent misuse intrusion strategies whose output is systematically fed to correlation cells capable of more accurate diagnosis.

We present an attack-source-centric ontology that extends previous work in the area. It enables message interpretation and enhanced agent communication within the architecture simplifying at the same time system maintenance and facilitating the integration of new components.

We describe the implementation of the proposed architecture through the FROID prototype as a proof of concept. This is a misuse-based intrusion detection system built with agent and semantic web open-source technology whose particular focus is the identification of automated remote attack tools. It performs signature generation, matching, and correlation, and supports a signature deployment mechanism over the Internet.

We introduce a similarity matching method that improves the performance of existing algorithms by leveraging entropy and frequency properties of the input hereby reducing search time.

We link detection with incident response by procuring low false alarm rates that allow us to study local and external reaction methods to outbound intrusion events. We also present a component of the architecture that performs tracing of interactive sessions as a way of identifying the root location of a security event.

We describe the experimental design and report the results obtained with the prototype that show the feasibility of the approach as an alternate way of containing the impact of security incidents through the integration of a mesh of monitoring agents.

Copyright © Salvador Mandujano Vergara, 2004

Declaration

I hereby declare that this dissertation entitled "An Multiagent Approach to Outbound Intrusion Detection" describes my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any degree or diploma by ITESM or any other institution of higher learning, except where due acknowledgment has been explicitly made in the text.

Salvador Mandujano Vergara

Monterrey, Mexico, December – 2004

List of Figures

2.1 Anomaly-based intrusion detection. 13

2.2 Misuse-based intrusion detection. 15

2.3 A frequency system-call profile describing the behavior of a program [147]. 16

2.4 A simple Bayesian network connecting intrusion-related variables. CPT's store a-priori probabilities such as the probability of intrusion and the proportion of disk fragmentation 17

2.5 A sample neural network with window that predicts the next command once n commands have appeared in the input (window size is $n = 3$. 19

2.6 A Petri net corresponding to a partial order of events leading to a compromise state [103] 20

2.7 A parse tree corresponding to a program that can be evolved through Genetic Programming to detect a security-relevant event [41]. 21

2.8 Code vulnerable to buffer overflow before and after embedding a sensor into it. On the left side, lines 2–6 form the embedded sensor [223]. 22

2.9 Caller-callee relationships are distinctive of a program and can be used to identify abnormal behavior. An unexpected transition may be indicative of anomalous activity (see dashed line; no such transition exists in the actual program). 23

2.10 SVM splitting the feature space into two classes (they could represent normal and irregular behavior). 24

2.11 Example of DNA sequence alignment [39]. 26

2.12 The increasing sophistication of attacks and the decreasing expertise of users [2, 126]. 34

2.13 Number of security incident and number of vulnerabilities reported per year to the CERT–CC (*Coordination Center, Computer Emergency Response Team* [32]). 34

Sophos Ex. 1009
Page 99 of 318

2.14 Use of information security technologies (*2003 Global Security Survey, Deloitte & Touche* [48]) and intrusion detection technologies (*CSI/FBI Computer Crime & Security Survey 2003* [169]). The charts show the percentage of surveyed companies that report using this type of software. 35
2.15 Number of known Denial of Service (DoS) vulnerabilities and attacks per year (*PestPatrol's detection database, PestPatrol Inc. 2002* [155]). 36
2.16 (a) Replication is performed by a pool of agents which are contacted through a proxy agent. (b) Replication can also be peer-to-peer. 39
2.17 Hierarchical component of a multiagent intrusion detection architecture [224]. . . 42
2.18 Enabling high-level agent functions through an ontology layer. 42

3.1 Data collection focus of different intrusion detection approaches. 50
3.2 Attacks are detected at the originating host or at an intermediate host and can be blocked. This creates a security structure centered on the protection of others. 51
3.3 Characterization of automated attack tools through the use of locally-measurable and observable variables. 53
3.4 Under the assumption that the number of outgoing network packets an internal host generates is less than the incoming network packets a gateway sees, OID allows for more efficient packet processing. 54
3.5 A mesh of trusted OID agents that share plaintext of encrypted attack data over secure channels 56
3.6 Performance of an OID system: *Number of comparisons* vs. *time* (a comparison tries to match program files against program signatures). 57

4.1 Features of the OID environment 64
4.2 Primary OID tasks (YAWL diagram; see Appendix A). 65
4.3 Agent interconnection architectures (S = sensor, C = correlator, R = reactor; arrows indicate data exchange flow). 69
4.4 Grouping agents functionally into specific-purpose sub-environments or *agent cells*. 71
4.5 OID architecture and main cell classes. 72
4.6 Internal dataflow: (a) Traffic-signature cells, and (b) Process-signature cells (see Appendix A for YAWL symbols). 74
4.7 Internal dataflow: (a) Traffic-monitoring cells, and (b) Process-monitoring cells (see Appendix A for YAWL symbols). 75
4.8 Dataflow inside Correlation cells (see Appendix A for YAWL symbols). 76

Sophos Ex. 1009
Page 100 of 318

4.9 Deployment of attack program signatures. Repositories connected to the Internet receive signatures from registered signature generation sites. Agents download signatures and ontology updates (M represents a mirror repository). 76
4.10 Graphical notation 77
4.11 A target-centric ontology for intrusion detection [203, 202]; bold arrows represent subclasses; soft arrows represent properties. 79
4.12 A source-centric ontology for outbound intrusion detection (see Appendix B for its definition in OWL) 80
4.13 Overlapping between the target-centric intrusion detection ontology and the source-centric OID ontology. 81
4.14 OWL statements defining the Signature–Process Signature subclass predicate. . . 81
4.15 Literal attributes of the *Traffic Signature* object. 81
4.16 OWL definition of a sample Traffic Signature instance. 82
4.17 The Agent-cell class 83
4.18 The Agent class 83
4.19 The Data-source class 84
4.20 The Signature class 85
4.21 The Reaction class 85
4.22 The Process class 85
4.23 The Repository class 86
4.24 The IP-packet class 86
4.25 The Executable class 87
4.26 Loop unwinding. If a user connects from host h_s to host h_t doing a loop through hops 1, 2, 3, 4, 5 and 6 as shown above, agent A at h_t will start the tracing process and will be able to identify the loop to reduce tracing time in the future. 90

5.1 Sample TCP socket table. 93
5.2 Hierarchy of classes derived from the primary JADE-agent class. 94
5.3 Packet reading assignment for two ($n = 2$) monitoring cells. 95
5.4 Failure cascade for sensors performing packet capture. 96
5.5 Sample backup structure for $m = 3$: when A is a primary agent, B is its first backup and C is its second backup; when B is a primary agent, C is its first backup and A is its second backup; when C is a primary agent, A is its first backup and B its its second backup. 97

Sophos Ex. 1009
Page 101 of 318

5.6 Payload data of a `telnet` packet. One byte is used for sub-protocol window size negotiation. The maximum value corresponds to FF which produces a deviation between 0 and 30 in the hash value. 99

5.7 Network-traffic signatures of three different attack programs considering executions of *only one variant* per program. 99

5.8 Network-traffic signatures of three different attack programs considering *two variants*. 100

5.9 Snort's signature matching linked list [170, 37]. 101

5.10 Data structure for attack signatures. Nodes are sorted from highest to lowest frequency. 102

5.11 Splitting power per signature attribute. 104

5.12 The matching method gives more weight to attributes with high splitting power (see Table 5.4). 107

5.13 Function to compute payload hash σ_i sub-scores. 107

5.14 The text segment can be used to characterize a program. The figure shows an excerpt from the text segment of `smurf4`, one of the surveyed attack tools. 109

5.15 Process signature linked list. Every s_i node represents a program-variant signature. 110

5.16 Fusing hypothesis H_1 and H_2 into H_3. 112

6.1 Sample attack program containing root-shell code for two different target platforms: Linux and FreeBSD. 116

6.2 Experimental setup for the testing phase. 119

6.3 The Mucus intrusion detection stimulator. 121

6.4 Average hypothesis score $\sigma(x)$ corresponding to each attack tool using one traffic-monitoring cell. 123

6.5 Average hypothesis score $\sigma(x)$ corresponding to each attack tool using two traffic-monitoring cells. 124

6.6 Average hypothesis score $\sigma(x)$ corresponding to each attack tool using three traffic-monitoring cells. 125

6.7 Average hypothesis score $\pi(x)$ corresponding to each attack tool using one process-monitoring cell. 126

6.8 Average hypothesis score $\pi(x)$ corresponding to each attack tool using two process-monitoring cells. 127

Sophos Ex. 1009
Page 102 of 318

6.9 Detection time with network monitoring under different traffic levels using one, two, and three cells. 128
6.10 Detection time with process monitoring under different workload levels using one and two cells. 129
6.11 Detection accuracy using correlation with one network-monitoring cell, one process-monitoring cell, and one correlation cell. 129
6.12 Detection accuracy using correlation with two network-monitoring cells, two process-monitoring cells, and two correlation cells. 131
6.13 Left: average detection time using two monitoring cells and one correlation cell. Right: average detection time using two pairs of monitoring cells and two correlation cells. 131

7.1 TCP datagram format (RFC 793) [38]. B–143
7.2 IP v.4 datagram format (RFC 791) [38]. B–145

Sophos Ex. 1009
Page 103 of 318

Contents

List of Figures **xix**

List of Tables **xxv**

1 Introduction **3**

1.1 Context and basic concepts . . . 3

1.1.1 Information security fundamentals . . . 3

1.1.2 The critical need for vigilance . . . 5

1.2 Limitations of intrusion detection systems . . . 6

1.3 Solution approach . . . 7

1.4 Thesis statement . . . 8

1.5 Document organization . . . 9

2 Background **11**

2.1 Intrusion detection . . . 11

2.1.1 Intrusion detection systems . . . 11

2.1.2 Classification of intrusion detection systems . . . 13

2.1.3 Detection techniques . . . 15

2.1.4 Data collection and fusion . . . 27

2.1.5 Challenges of current models . . . 30

2.2 Recent trends in information security . . . 33

2.3 Connection back-tracing . . . 37

2.3.1 Limited troubleshooting visibility . . . 37

2.3.2 Related projects . . . 37

2.4 Agent-based intrusion detection . . . 38

2.4.1 Agents . . . 39

Sophos Ex. 1009
Page 105 of 318

2.4.2 Multiagent systems . . . 40
2.4.3 Multiagent components of intrusion detection architectures . . . 41
2.5 Taxonomies and ontologies in information security . . . 43

3 Problem description 49
3.1 Outbound intrusion detection . . . 49
3.1.1 Objective . . . 49
3.1.2 Main characteristics . . . 50

4 A multiagent outbound intrusion detection architecture 59
4.1 A multiagent approach to intrusion detection . . . 59
4.2 Specific objective of the architecture . . . 60
4.3 Characterization of automated attack tools . . . 60
4.3.1 Attack program signatures . . . 61
4.4 Environment . . . 63
4.5 Task definitions . . . 64
4.6 Task-to-agent mapping . . . 66
4.7 Organization of single agents . . . 68
4.8 Agent cells . . . 71
4.8.1 Agent classes and cell classes . . . 72
4.8.2 Signature and ontology deployment . . . 73
4.8.3 Operational outline . . . 74
4.9 Ontological model . . . 75
4.9.1 Supporting agent-based security technologies with ontologies . . . 76
4.9.2 Target-centric ontology for intrusion detection . . . 77
4.10 Source-centric ontology for outbound intrusion detection . . . 82
4.10.1 Use of the ontology . . . 86
4.11 Session tracing . . . 88
4.11.1 Internal design . . . 88
4.11.2 Hop identification . . . 89

5 Implementation 91
5.1 Technologies . . . 91
5.2 Triggering event . . . 92
5.3 Agent creation . . . 93

Sophos Ex. 1009
Page 106 of 318

5.4 Network data capture . . . 94
5.4.1 Buffering and coordination . . . 95
5.4.2 Packet read backup . . . 96
5.5 Process data capture . . . 97
5.6 Matching of network signatures . . . 98
5.6.1 Primary data structure . . . 100
5.6.2 Splitting power analysis . . . 102
5.6.3 Analysis of network data . . . 103
5.6.4 Matching method for network signatures . . . 106
5.7 Matching process signatures . . . 108
5.7.1 Primary data structure . . . 109
5.7.2 Matching method for process signatures . . . 110
5.7.3 Signature fusion . . . 111
5.8 Session tracing implementation . . . 113
5.9 Reaction methods . . . 113

6 Experiments and results 115
6.1 Experimental objectives . . . 115
6.2 Test data . . . 115
6.2.1 Exploit types . . . 115
6.2.2 Target platforms . . . 116
6.2.3 Programming languages . . . 117
6.2.4 Test program set . . . 117
6.3 Experimental setup . . . 119
6.4 Signature generation . . . 120
6.5 A word on compiled and interpreted code . . . 120
6.6 Experiments: Signature matching and correlation . . . 120
6.6.1 Traffic generation . . . 121
6.6.2 Workload control . . . 122
6.6.3 Cell accuracy . . . 123
6.6.4 Cell performance . . . 127
6.6.5 Correlation accuracy . . . 130
6.6.6 Correlation performance . . . 131
6.7 Interactive session tracing . . . 132

Sophos Ex. 1009
Page 107 of 318

6.8 Results 133

7 Conclusions, contributions and future work 135

7.1 Conclusions 135

7.2 Contributions 137

7.3 Future directions 138

Appendix A YAWL symbols A–141

Appendix B TCP and IP datagram formats B–143

Appendix C Outbound intrusion detection ontology in OWL C–149

Appendix D Classification Trees D–155

Appendix E Process signatures E–165

Appendix F Network signatures F–171

Sophos Ex. 1009
Page 108 of 318

List of Tables

2.1 Classification of intrusion detection systems by input type [223, 13]. 29
2.2 Classification of intrusion detection systems by their data analysis structure [223, 13]. 31
2.3 Classification of intrusion detection techniques 47

4.1 Description of the Sensor, Correlator, and Reactor agent classes 67
4.2 Description of the Name and Ontology Server agent class 68

5.2 Attributes stored on each node type. 104
5.4 Splitting power per attribute in network signatures. 106
5.6 Rules to implement automated response. 114

6.2 Automated remote-attack tools used in the experiments. 118
6.4 Matching scores for one, two, and three traffic-monitoring-cell scenarios. 125
6.5 Tracing time with and without encryption (RSA key exchange + DES chaining block ciphering); m corresponds to the number of connection hops or length of the chain. 132

Sophos Ex. 1009
Page 109 of 318

Chapter 2

Background

This chapter is divided into five main sections. They present basic concepts used and studied throughout subsequent chapters.

We first define some intrusion detection terminology. Computer security trends observed in recent surveys are presented to support the motivation behind this dissertation. The problem of connection tracing, which is closely related to intruder identification, is also introduced. Finally, the multiagent approach to intrusion detection is explained as well as some relevant projects in the area of security taxonomies and ontologies.

2.1 Intrusion detection

This section introduces intrusion detection concepts relevant to the objectives of this dissertation. It also describes the "state of the art" on detection techniques, shows recent trends in computer security and cyber-crime, and lists some of the primary research challenges in intrusion detection.

2.1.1 Intrusion detection systems

There are several ways to avoid being the target of an intrusion or to react to one [76]:

1. *Prevention.* To proactively reduce the likelihood of intrusive activity by influencing the root causes of the problem (e.g., to remove unnecessary services from a host).

2. *Preemption.* To actually counter-attack the source of the intrusion before it reaches its objective (e.g., launching a denial of service attack toward the presumed offender [208]).

3. *Deterrence.* To persuade the attacker to stop the attack (e.g., using system banners announcing the presence of the security tools installed on the system).

4. *Deflection.* To make the attacker think he has succeeded (e.g., using honey-pot boxes [193]).

5. *Use of countermeasures.* To react to the intrusion as it develops (e.g., applying patches and reconfiguring the system).

Sophos Ex. 1009
Page 111 of 318

6. *Detection*. To identify intrusive attempts and incidents in order to keep the corresponding authority informed (e.g., reporting anomalous URL strings received by the webserver).

Out of these possible reactions, intrusion detection, along with prevention, is the approach that has been explored the most [13, 4]. Twenty years ago, Anderson wrote seminal work in the area which increased the interest in computer security [6]. A few years later, the first intrusion detection model by Denning opened the field and eventually became central reference to multiple security architectures developed around the world [49].

Intrusion detection is concerned with the identification of activities that have been generated with the intention of compromising the security of local resources [4]. An attacker may be trying to gain access to a system from outside of the network, but it can also be the case that a legitimate internal user is abusing the access control mechanism in order to disclose sensitive data or to modify files without having been authorized to do so.

Intrusion detection software gathers and analyzes data from different sources within the system in order to identify activity that may be indicative of a compromise attempt [223]. An intrusion detection tool will alert the administrators who will have to take action in order to limit the damage and fix any security hole present in the system.

A related concept that needs to be mentioned is that of *intrusion tolerance* [51, 217].

DEFINITION 7

Intrusion tolerance. The ability of a system to continue working despite of intrusion [132]. Its main goals are: (1) to recognize the failure event, (2) to asses the extent of the damage, (3) to stop the event, and (4) to recover essential services during or after the incident. It is also often referred to as *survivability*.

This discipline comprises the sub-area of *fault tolerance* which studies ways to make a system resistant to failure (fault tolerance does not contemplate the case of induced failures or intrusions [210]).

DEFINITION 8

Fault tolerance. The ability to resist failures that are natural to the system (i.e., random operational faults [157]).

Intrusion tolerance also studies the problem of designing modular systems so that components can be easily isolated in the event they are compromised hereby allowing the rest of the system to keep providing all or part of its overall functionality [99, 12].

It is important to make the point that intrusion detection defers from intrusion tolerance and fault tolerance in that intrusion detection does not have to do with the design and implementation of robust systems that resist failure or attack. Intrusion detection's only goal is to identify violations to a local security policy and to report them via alert messages to the administrators [63].

Sophos Ex. 1009
Page 112 of 318

2.1.2 Classification of intrusion detection systems

There are two primary methods to implement intrusion detection: (1) *anomaly detection* (also known as *profile-based intrusion detection* and *statistical intrusion detection*), and (2) *misuse detection* (also known as *signature* or *pattern-based detection*).

DEFINITION 9

Anomaly detection. An intrusion detection technique that creates a definition of "normality" and flags as intrusive any activity that deviates from it [126].

Anomaly detection. Anomaly detection assumes intrusions constitute irregular events. It creates behavior profiles by observing the "normal" activity of users or applications (Figure 2.1). Once these profiles have been built through training or using heuristics, any activity that diverges from them will be considered as potentially intrusive. False positives and false negatives may arise due to the threshold that divides normal from abnormal behavior.

A **false positive** is the classification of legitimate activity as intrusive, whereas a **false negative** is the classification of illicit activity as benign – the latter being a much more serious problem as it represents failure to identify hazardous events [103].

Anomaly-based intrusion detection systems are computationally expensive due to the techniques they employ. They typically require the construction and tuning of behavioral models (some of these systems are described in [152, 108, 83, 175, 109, 64]). Classifying what is "normal" from what is not is a difficult task that involves not only technical factors such as compliance with protocol definitions or the understanding of the internals of an application, but also aspects such as the expertise level of the users, their preferences, and the date and time. They also contribute to the complexity of identifying what is an acceptable variation of the system's security context and they are sometimes difficult to model as input for software systems [4].

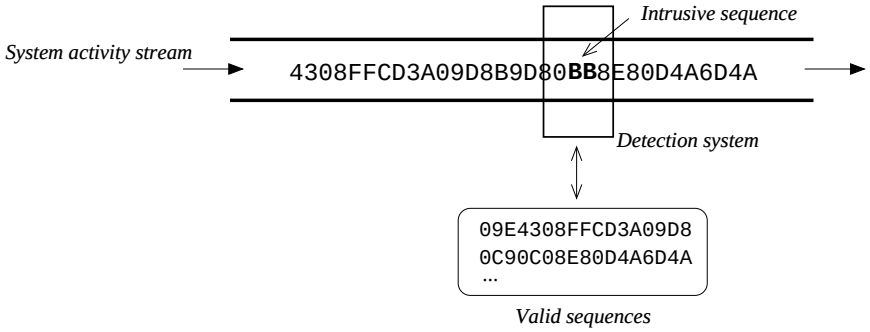


Figure 2.1: Anomaly-based intrusion detection.

– *Advantages of anomaly detection.* The main advantage of anomaly detection is that it can potentially discover attacks that have not been seen before [103]. Any evidence of activity variating from the profile will be considered for further analysis regardless of what has been previously seen. Another advantage is that the system can be left to run unattended once the

Sophos Ex. 1009
Page 113 of 318

normality profile has been created. This considerably reduces maintenance time [60].

– *Disadvantages of anomaly detection.* One of the disadvantages of this approach is that unusual behavior is not always a sign of attack. This may produce false alarms that report incorrect information to the administrator if not implemented correctly [2]. Another problem is that anomaly detection systems could be tricked into accepting anomalous behavior: an attacker could make the system consider illicit activities as normal behavior if he has access to the system during the training phase [60]. In general, the methods to create and evaluate these profiles are computationally intensive and complicate real-time detection [2, 13].

DEFINITION 10

Misuse detection. An intrusion detection technique that defines patterns of abnormal activity and then looks for occurrences of those patterns in the monitored data streams. An occurrence signifies a potential intrusion [13, 126]).

Misuse detection. The misuse detection approach assumes that intrusive activity can be described through patterns and that those patterns can be captured by means of event strings or state machines [105, 187]. These systems identify *signatures* that contain information distinctive to an attack (Figure 2.2).

DEFINITION 11

Attack signature. A pattern or fingerprint that captures the essence of an attack and which can be used to identify the attack when it occurs [142].

Monitoring programs observe the different data streams present on the system and match them against a bank of attack signatures [170, 194, 53, 72]).

Misuse detection is a pattern matching method. Once signatures are generated, the problem reduces to one of search and matching. The primary difficulty of this approach is to generate signatures that correctly represent what constitutes intrusive actions [126]. Signatures that are too specific may not catch slight variations of an attack, and signatures that are too flexible might generate false alarms by flagging as intrusive activity that is not.

– *Advantages of misuse detection.* Misuse detection techniques do not consume as much resources as anomaly detection methods do. The processing of signatures can be highly optimized as to comply with performance expectations by adjusting parameters in the matching algorithms [194, 75]. Thus, the rate of false alarms can be controlled since matching accuracy and signature specificity constitute a trade-off directly bound to the security level these systems offer [126].

– *Disadvantages of misuse detection.* Signatures need to be very precise for them to be useful and the task of signature specification is rather time consuming and requires a high level of expertise [13, 126]. In the event an attack signature overlaps with licit behavior, false positives may arise. If the signature is not specific enough the attack will go unnoticed and will compromise the security of the system. The main disadvantage of this approach is that, as opposed to anomaly-based intrusion detection, it cannot detect threats for which a signature has not been produced. Signature banks must be permanently updated and signatures need to be defined and tuned in order to detect new attacks.

Sophos Ex. 1009
Page 114 of 318

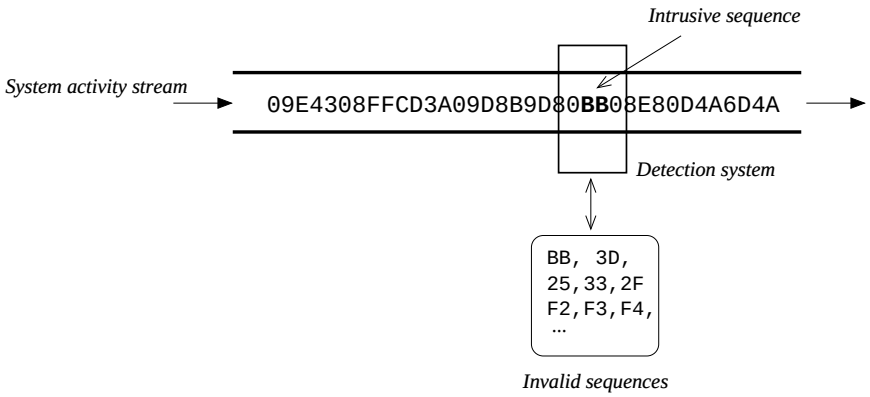


Figure 2.2: Misuse-based intrusion detection.

Hybrid methods. The two main approaches, anomaly and misuse detection, have advantages that complement each other. In fact, a breed of hybrid systems are the ones dominating the market and some more have become successful research projects (e.g., [164, 205, 77, 187, 46, 194, 140, 5].

These systems are fed with attack patterns they will detect and, at the same time, implement learning and profiling algorithms to necessary to identify invalid actions [90, 2]. The solution hybrid intrusion detection systems provide is richer detection-wise as they typically access two or more information sources which gives them the advantage of having more data elements to draw conclusions from. By correlating information of a normality profile with evidence of a particular attack, a system will be in a better position to accurately identify illicit behavior [60].

2.1.3 Detection techniques

In order to perform the actual task of detecting unusual events from evidential data, intrusion detection systems have made use of many different techniques. This section outlines the most representative methods developed to date (we include the techniques surveyed by Kumar and Spafford in [104] and extend the list with more recent detection mechanisms):

1. **Frequency-based detection**. This was one of the first anomaly detection methods developed (it first appeared in Denning's seminal paper [49]). It captures frequency patterns of users and programs through profiles that contain abnormality values $P_1, P_2, ..., P_n$ corresponding to a set of n system *measures* (i.e., representative system variables). In order to determine whether the system has encountered an unstable state, the observed frequency values f_i are combined through an expression such as:

$$\gamma = a_1 f_1^2 + a_2 f_2^2 + ... + a_n f_n^2; \quad a_i > 0 \qquad \ldots (1)$$

which captures frequency information for all the measures using a series of weights a_i. This technique detects intrusions by computing security level γ on a permanent basis and comparing it to a threshold value δ. If $\gamma > \delta$, an intrusion may be underway [182, 104].

Sophos Ex. 1009
Page 115 of 318

The different measure types that can be evaluated are diverse and can be divided into three main types: (1) ordinal measures (such as the number of files open by a certain device in a given period of time), (2) distribution measures (such as the statistical distribution of CPU utilization during the day), and (3) categorical measures (such as the number of successful transfers and the number of failed transfers requested to a DNS server).

Profiles capturing normality values could be of the form:

```
<U, file, execute; Error(0), CPU(00:01-00:04), ExpectedExec(10:00-11:30)>
```

This could indicate that whenever user U executes $file$, no errors are expected, CPU processing is likely to fall between 1 and 4 seconds, and the execution is expected to take place between 10:00AM and 11:30AM. This way it is possible to identify an intrusion if it rises the value of a measure over the limits stipulated in the profile [66]. If a program is described through the system calls it uses, a frequency profile can be built in order to describe its most common behavior (Figure 2.3):

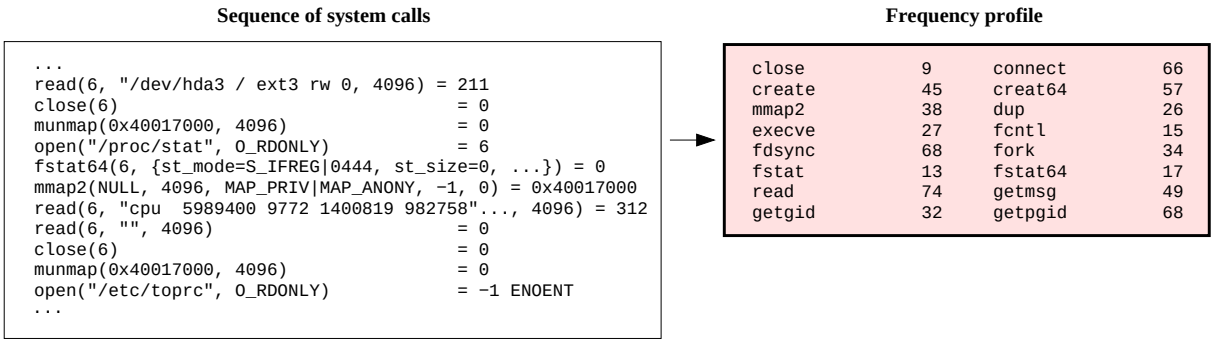


Figure 2.3: A frequency system-call profile describing the behavior of a program [147].

This profile describes how often the program requests each of the system-call it uses. Taking this profile, if we observe that, for instance, the number of `close()` system-calls is up to 30, it is likely that there is a problem on the system.

2. **Probabilistic detection models**. Let $A_1, A_2, ..., A_n$ be n different measure variables which can take one of two values: 1 to indicate the measure is anomalous, and 0 otherwise. Let I be the hypothesis that the system is undergoing intrusion. The sensitivity of each measure A_i is determined by $P(A_i = 1|I)$ and $P(A_i = 1|\neg I)$. By combining these probabilities we can determine the likelihood that the system is undergoing attack given the values of the measures:

$$P(I|A_1, A_2, ..., A_n) = P(A_1, A_2, ..., A_n|I) \times \frac{P(I)}{P(A_1, A_2, ..., A_n)} \quad \ldots (2)$$

Sophos Ex. 1009
Page 116 of 318

This requires the joint probability distribution of the measures given I and $\neg I$ [104]. Assuming probabilistic independence of the measures, obtaining $P(A_1, A_2, ..., A_n|I)$ would be reduced to:

$$P(A_1, A_2, ..., A_n|I) = \prod_{i=1}^{n} P(A_i|I) \quad \ldots (3)$$

and

$$P(A_1, A_2, ..., A_n|\neg I) = \prod_{i=1}^{n} P(A_i|\neg I) \quad \ldots (4)$$

A more realistic example would take into account the interdependence of the various measures. Lunt et al. [115] accomplished this through the use of *covariance matrices*. The compound anomaly of the system using vector $A = [A_1, A_2, ..., A_n]$ is given by

$$A^T C^{-1} A \quad \ldots (5)$$

where matrix C stores the interdependence between each pair of anomalies A_i and A_j at $C[i, j]$. Through a covariance matrix it is possible to consider the fact that the inputs that intervene in security diagnosis are interrelated [143].

3. **Bayesian networks**. Detection models using Bayesian learning take the basic structure of probabilistic detection in which a measure variable may affect others [183, 172]. This information is used to create belief networks called *Bayesian networks* in order to graphically represent causal dependencies between random variables. These networks are directed acyclic graphs and are able to compute joint distributions by specifying only a subset of probabilities corresponding to the neighbors of a node (Figure 2.4).

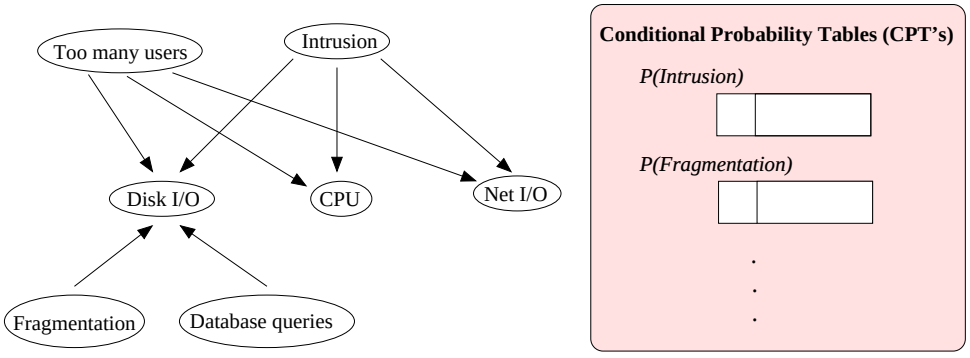


Figure 2.4: A simple Bayesian network connecting intrusion-related variables. CPT's store a-priori probabilities such as the probability of intrusion and the proportion of disk fragmentation

.

Bayesian networks use sets of *conditional probability tables* (CPT's) that store probability values corresponding to the random variables given certain preconditions. These probabilities are usually taken from the experience of past events and contribute to effectively predict the

Sophos Ex. 1009
Page 117 of 318

desired variable [94]. Intrusion is determined with a certain probability given the presence or absence of evidence.

Figure 2.4 shows a small network which can be used to answer questions such as "what is the probability of intrusion given the observed level of disk fragmentation?", or "what is the probability of intrusion given that the number of current logged-on users is high and CPU load is low?"

4. **Conditional models using event sequences**. Instead of considering system measures, this is a technique that detects irregular activity based on sequences and sub-sequences of events $E = [e_1 \prec e_2 \prec ... \prec e_n]$, where e_i is an event of type i and $\prec$ represents event precedence in time. These algorithms try to come up with a value for $P(I|E)$ where I indicates the probability of intrusion and E the probability of occurrence of an event pattern. Bayes calculus [173] can also be applied to determine how likely intrusion is given that pattern E appeared:

$$P(I|E) = \frac{P(E|I) \times P(I)}{P(E)} \quad \ldots (6)$$

An experienced system administrator could quantify approximates of a-priori probabilities such as the probability of occurrence of a given sequence of events at a certain time frame. It is then possible apply a probabilistic method that will be the one detecting intrusions and intrusion attempts. This model becomes complex when it is necessary to consider the large amount of intrusion patterns that can be observed on a system [104].

5. **Expert systems**. Expert systems have also been used in intrusion detection [181, 91, 105]. These systems embody the knowledge of an expert in order to identify unusual data and irregular actions. They are classified as anomaly-based as they explicitly specify the patterns to look for [103]. The success of these methods is directly related to two factors: (1) the expertise of the security officer that will be used as input to the detection mechanism, and (2) the effectiveness of the implementation to coherently structure the expertise of the human into a computer program.

   Expert systems have also been used to interpret system measures in intrusion detection, however, the well-known limitations of this approach for doing uncertainty reasoning have stopped the progress in the area [104].

6. **Inductive generalization**. This is a rule-based approach which characterizes the behavior of a system or user using inductive generalization under the assumption that system events can be predicted [80, 200]. A set of rules of the form:

$$E_a \rightarrow E_b \rightarrow E_c \Rightarrow (E_d = 0.95, E_e = 0.05) \quad \ldots (7)$$

expresses the probability of seeing an event appear after a sequence of other events. For instance, once events a, b, and c have appeared sequentially, event d is 95% likely to appear, and e will appear with a probability of 5%. These are both normal, acceptable sequences.

Sophos Ex. 1009
Page 118 of 318

Rules help to know beforehand the probability of being under attack given a set of observed patterns [200]. Events will be flagged as intrusive when they match the left-hand side of a rule but diverge on the probabilities that appear in right-hand side.

It is claimed that if an intruder tries to guide the rule-learning phase as to accept irregular activity as normal, the very meaning of rules can help determine something irregular has been captured [200] (i.e., rules are easily read and interpreted by administrators). This anomaly-based method has a major weakness: any unknown attack which does not match the left-hand side of a rule will not be even considered for further inspection.

7. **Neural networks**. Supervised learning methods have also been used in intrusion detection. These methods develop classifiers that predict output values based on a set of input attributes. Neural networks can be used for supervised learning and have proved effective at classifying different sorts of data when complemented with pruning, early stopping, and other heuristics [174, 27, 173]. A neural network contains a set of nodes organized into layers. Input and output layers are interconnected through one or more intermediate layers and the learning is performed by updating edge weights and adapting the parameters of the functions present on each node.

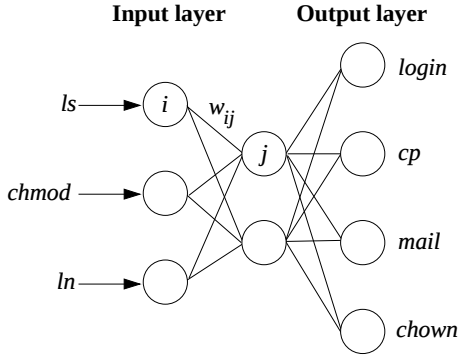


Figure 2.5: A sample neural network with window that predicts the next command once n commands have appeared in the input (window size is $n = 3$.

In the case of intrusion detection, a neural network can be used to to predict the next event to appear. For instance, the command history of a user can be used to train a neural network. Once training is done, a window of k commands will be sliding in time in order to monitor the activity of the user. If the neural network indicates that after certain k events, event e_j should follow and a different event is observed, an alarm is activated. Similarly, system calls can also be predicted after a training phase that will identify interdependencies among system-call streams hereby creating a sense of normality for the system [64].

Figure 2.5 shows a small network that predicts the next command given a three-command input. Command sequences must be divided into *windows* so that supervised learning can take place. In this case, the window will be moving one command at a time and will use three commands as input in order to determine which of the output values is more likely to

Sophos Ex. 1009
Page 119 of 318

appear next. Training time is the main restriction of the neural-network approach for the problem of intrusion detection. However, it deals very well with noisy data [174].

8. **Attack scenario models**. This misuse-detection approach describes attacks through sequences of behaviors called *scenarios* [68]. A database containing attack scenarios is accessed by the monitoring system in order to tell which of the scenarios is most likely to be present at a given time.

   These systems update the database using evidential reasoning calculus. The use of a mathematically-sound theory like evidential reasoning makes retraction of intermediate conclusions easy as evidence to the contrary accumulates. When determining the behavior to look for next, the system tries to maximize a factor α_x defined by:

$$\alpha_x = \frac{P(S_x|B)}{P(S_x|\neg B)} \quad \ldots (8)$$

   where S_x is a scenario and B corresponds to a behavior that confirms the presence of such scenario. Like other misuse detection systems, building attack scenario models require high levels of expertise and continuous updates [104].

9. **Petri nets**. This is a misuse detection technique which models attacks through Petri nets [105, 103]. Petri nets are graphs composed by states and conditions which have well defined semantics. The transitions from a start state into a final state describe the evolution of attacks. Petri nets have been used to model system events such as user commands and system calls (two popular types of input data used to detect intrusion).

   The graphical representation provided by a Petri net conveys an intuitive understanding of the events that compose an attack and allow for the representation of partial order as well – finite-state automata have also been explored in a similar way [201]. Rather than handle fixed sequences of events, Petri nets can describe series of events that are loosely connected to each other. Figure 2.6 shows a net describing an attack scenario. The steps that compose an attack may have variable precedence relations not being part of an absolute-order sequence.

```
1. cp /bin/sh /usr/spool/mail/root
2. chmod 4755 /usr/spool/mail/root
3. touch x
4. mail root < x
5. /usr/spool/mail/root
```

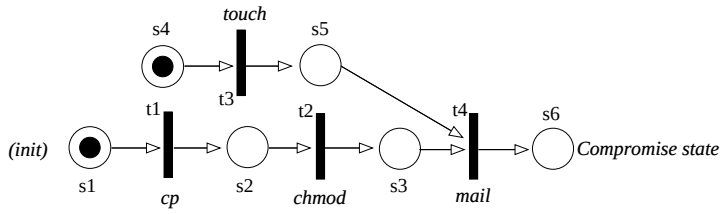


Figure 2.6: A Petri net corresponding to a partial order of events leading to a compromise state [103]

Petri nets allow us to describe and match partially ordered sequences as an alternative way of visualizing attack scenarios. In Figure 2.6, state $s5$ may be reached well before or after an

Sophos Ex. 1009
Page 120 of 318

alternate transition reaches $s3$. "A partial order state transition analysis allows more than one sequence of events in the state transition diagram" [66]. A single diagram can be used to represent several intrusion scenarios involving the same set of actions. This simplifies modeling and speeds up detection.

10. **Genetic programming**. Many security systems use monolithic programs to perform intrusion detection functions. A different approach is to have multiple small programs in charge of simple monitoring tasks. These programs can be further evolved through Genetic Programming techniques so that the best candidate for a determined task can be found [41].

    Genetic programming is a machine-learning method derived from *Genetic Algorithms* which evolves a population of programs, and not a population of sample points, to solve a problem. Possible solutions are represented as parse trees whose component nodes are manipulated through genetic operations such as crossover and mutation. A fitness function aids in the identification of the best individuals (i.e., programs). In the case of intrusion detection, normal and abnormal activity is presented to the programs in order to evaluate how well they do at distinguishing between the two.

    Figure 2.7 depicts a sample parse tree corresponding to a block of code whose purpose is to watch for packets arriving at a host but destined for another host [41] (a simplistic way of determining something wrong may be happening).

```
FOR_EACH_PACKET DO
   IF( ip-destination-address-of-packet
          NOT_EQUAL my-ip-address )
     THEN Generate(alert-broadcast)
   END_IF
END_FOR
```

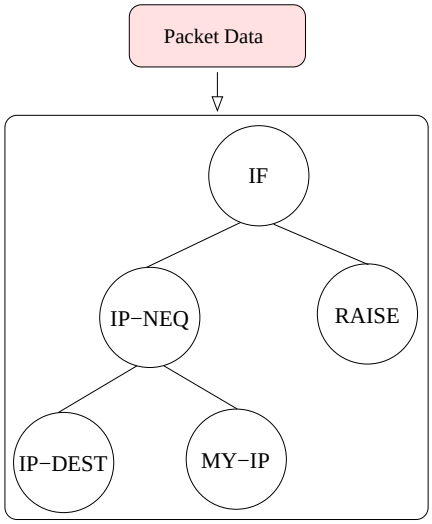


Figure 2.7: A parse tree corresponding to a program that can be evolved through Genetic Programming to detect a security-relevant event [41].

Each component of a program, which can be expressed using different languages and graph structures, will be moved from place to place connecting with other different nodes during the evolution stage. This will last until the component finds the location at which it can produce best detection results to the program as a whole. Important to notice that a good

Sophos Ex. 1009
Page 121 of 318

number of combinations will produce incorrect trees from the programming logic perspective whose fitness will be low. These will be discarded and the best ones will be selected to implement the intrusion detection strategy.

11. **Embedded sensors**. Another technique called *embedded sensors* [223, 98] proposes the use of code segments that are inserted into programs during the development phase. These code segments automatically search for abnormal conditions at execution time including well-known error or attack conditions such as stack overrun and data-type abuse (Figure 2.8).

```
1 char buf[256];
2 strcpy(buf, getenv("HOME"));
```

(a) Code before inserting sensor

```
1 char buf[256];
2 {
3    if( strlen(getenv("HOME")) > 255) {
4       log_alert("buffer overflow");
5    }
6 }
7 strcpy(buf, getenv("HOME"));
```

(b) Code after inserting sensor

Figure 2.8: Code vulnerable to buffer overflow before and after embedding a sensor into it. On the left side, lines 2–6 form the embedded sensor [223].

The type of checking these tools provide is equivalent to the implementation of good coding practices with the difference that sensors can be specific to certain attack types and are reusable. Among the disadvantages of this approach are the need to have the source code of the program to protect, the need to implement sensors in the same language as the program, and reduced portability [59].

12. **Functional dependency checking**. This is a recent misuse detection approach based on functional dependencies present in a program. In order to perform certain actions, a function executes others that provide the results it needs. At each function call, one function is the *caller* and other is the *callee*. Caller-callee relations can be captured through a graph using executable or source code: at run time, the runtime stack can be read in order to determine what are the functions being called by each function; by looking at the source code, it can be determined the same type of information but at a higher level [71].

    These dependencies are characteristic to a software program and it is possible to implement a monitor which detects graph transitions not present in the graph (system-call precedence as well as regular function-calls can be used; [211, 186, 201]).

    Figure 2.9 shows an example of this approach. The structure of a program is analyzed as to create a graph describing the relationships among its functions. This graph is built during a training phase when legitimate behavior is exercised and captured by the graph. Once in detection mode, if a security monitor observes a transition that is not present in the program's graph, it is likely that the program has been altered [186].

Sophos Ex. 1009
Page 122 of 318

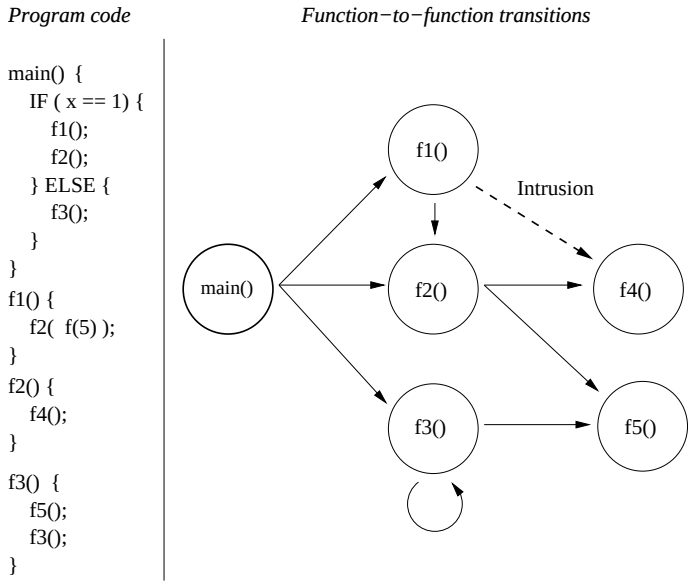


Figure 2.9: Caller-callee relationships are distinctive of a program and can be used to identify abnormal behavior. An unexpected transition may be indicative of anomalous activity (see dashed line; no such transition exists in the actual program).

13. **Data mining**. Some researchers have studied ways to apply data mining techniques to intrusion detection [114, 113]. Data mining is concerned with the discovery of new knowledge in large data collections. This discipline employs machine learning and statistical analysis methods to identify hidden patterns and trends a human would not be able to detect efficiently on large databases. Inferred patterns or rules allow the prediction of future results and, in the case of intrusion detection where the amount of data to inspect is quite large, this kind of prediction can help detect furtive attempts to circumvent the security of a system.

    For example, suppose that 10% of the packets received by a host have the SYN and ACK flags on and that 30% of the packets that have the ACK flag on also have the SYN flag on. These two variables could be associated using a rule $ACK \Rightarrow SYN$ with *support degree* $s = 0.1$ and *rule confidence* $c = 0.3$:

$$ACK \Rightarrow SYN \qquad s = 0.1, c = 0.3 \qquad \ldots (9)$$

    Rules like this can be computed for different components in order to build a normality profile that helps identify illicit activities [112].

14. **Fuzzy data mining** has also been studied [22]. The integration of fuzzy logic with data mining helps create more abstract patterns at a higher level. Rather than describing a measure through an abnormality range $[a, b]$, a security incident can be described using terms such as "high" or "low" which are more easily interpreted by humans [212]. By incorporating this concept to data mining, a value can contribute to the support of more than one fuzzy set. For instance, we can have a rule:

Sophos Ex. 1009
Page 123 of 318

$$\{SYN = LOW,\ FIN = LOW\} \Rightarrow \{RES = LOW\} \qquad s = 0.5,\ c = 0.9 \qquad \ldots (10)$$

where SYN is the number of *synchronize* flags, FIN is the number of *finalize* flags, and RES is the number of *reset* flags in a certain time period. Instead of having numeric values describing these figures, fuzzy terms are introduced to give more flexibility to the rule. Fuzzy sets describing these terms (e.g., LOW, $HIGH$) would have to be created in order to map the input before performing intrusion detection [22].

15. **Support Vector Machines**. Support vector machines (SVM) are one of the most recent supervised-learning methods and is based on the addition of dimensions that allow for regression methods to be executed [135, 219, 130]. SVM's are constructs that perform binary classification. A set of input data points are mapped into multidimensional-space *vectors* using *kernel functions* so that they can be split into two classes.

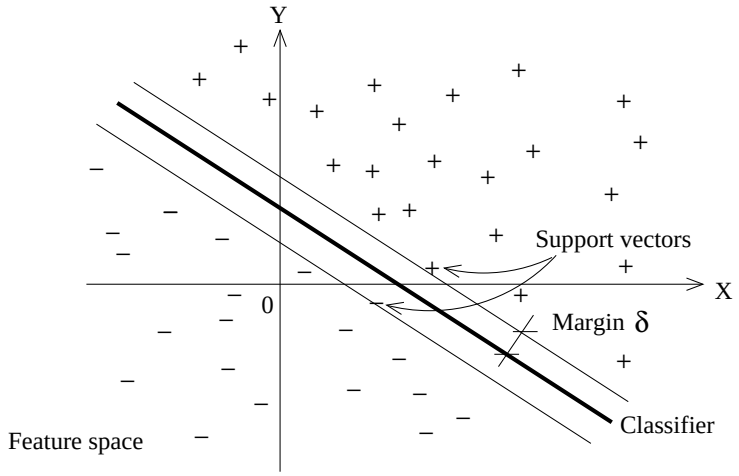


Figure 2.10: SVM splitting the feature space into two classes (they could represent normal and irregular behavior).

Some of these vectors are selected to define the boundary between the classes, and a hyperplane is computed through regression in order to split the input. It is then gradually adjusted by minimizing a *marginal-error distance* δ measured from the representative vectors (i.e. support vectors) to the hyperplane. It then becomes an optimization problem.

Just like neural networks, SVM's are trained in order to create classifier functions that help separate normal from abnormal activity. The training speed of SVM's is one of its most notable characteristics as well as its scalability (SVM's are "relatively insensitive" to the number of input data points [219]). They can learn large pattern sets and are able to scale better than neural networks.

Sophos Ex. 1009
Page 124 of 318

Figure 2.10 shows an example of a SVM on a bi-dimensional feature space. Vectors $v_i = (x_i, y_j) \in X \times Y \times \{\pm\}1$ (which could represent packet information such as packet length and checksum) are classified by a linear function which has a variable marginal distance δ determined by support vectors close to the boundary. This margin should be maximized in order for classification to be accurate.

SVM's have been used to predict security events and also to classify system behavior and fingerprints [135, 219].

16. **Planning**. An approach presented in [66] involves the construction of partial-order plans (i.e., partial-order subsequences) that limit the exponential growth of the number of states in a state-machine representing an intrusion scenario [173].

    A planner builds "plans that achieve goals which describe situations that are desirable to an intruder" [66]. A monitor will search through the set of active plans and not through the set of all possible situations, hereby reducing overhead.

    A planning program for intrusion detection will use a knowledge base which contains dependencies between actions. These dependencies will be used to create alternative plans. For example, the arrangement of preconditions P_i necessary for achieving a certain state S can be given by an expression like:

$$(P_1, P_2, ..., P_k) \wedge (P_i \prec P_k) \wedge ... \wedge (P_j) \prec P_m \quad \ldots (11)$$

    which could be used to represent a more concrete plan:

$$(cp, chmod, touch, mail) \wedge (cp \prec chmod) \wedge (chmod \prec mail) \wedge (cp \prec mail) \quad \ldots (12)$$

    That means, a set of actions need to appear and they must keep some relative precedence among them. A plan is started with a start state and further steps are added as preconditions appear. A planner uses precondition sequences to generate attack schemes that can be then validated with Petri nets or a similar tool [66, 104].

    First order logic can be also used to capture a penetration plan using the same approach. For example:

$$\begin{aligned}
&\exists\ file_1, file_2, file_3, x \\
&\quad owner(file_1) = x \wedge \\
&\quad cp(file_1, file_2) \wedge \\
&\quad chmod(file_2, 4755) \wedge \\
&\quad touch(file_3) \wedge \\
&\quad mail(root, file_3) \wedge \\
&\quad cp(file_1, file_2) \prec chmod(file_2) \wedge \\
&\quad chmod(file_2) \prec mail(root, file_3) \wedge \\
&\quad touch(file_3) \prec mail(root, file_3) \wedge
\end{aligned}$$

Sophos Ex. 1009
Page 125 of 318

$\Rightarrow attacker(x)$

which represents the conditions necessary to determine x is an attacker. Inference engines can be used to evaluate such plans with strong logic support. Each condition is evaluated and, in case all the conditions of a given plan are given, the system is able to tell intrusion is occurring.

17. **Bioinformatics sequence-alignment methods**. Sequence alignment is key to determine the similarity between two DNA sequences. Nucleotides or amino-acids are aligned in each sequence and a score is computed as to tell how well the sequences align with one another [39]. Pattern-matching algorithms originally developed for quaternary logic can be applied to determine the similarity between an input and an attack stream (Figure 2.11).

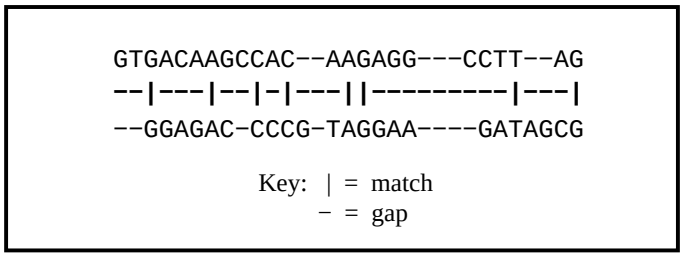


Figure 2.11: Example of DNA sequence alignment [39].

These methods are adapted to inputs consisting of more than four event types (i.e., the number of system-call types on Linux environments are over 200). In general, alignment of observed commands with long attack patterns is rewarded without heavily penalizing gaps or misalignment. The reason is that a matched subsequence is highly relevant for intrusion detection compared to misaligned subsequences. Matrices and n-dimensional data structures are built in order to compute the matching score of an input against n attack patterns. As shown in Figure 2.11, protein patterns slide over the input similarly to other matching algorithms [104, 37].

The main difference with binary matching methods is that Bioinformatics algorithms involve more than two input symbols, which allow them to be extended more easily to the case of intrusion detection sequences where the input domain has more symbols.

The techniques above are some of the most common methods utilized to detect illicit activity in computer systems. Table 2.3 at the end of the chapter lists the reviewed techniques and classify them according to three main points: (1) their basic approach (anomaly-based intrusion detection (AID), misuse-based intrusion detection (MID)), (2) their foundation (probabilistic method or learning method), and (3) the data source they can use (audit information, network packets, or a different type).

Sophos Ex. 1009
Page 126 of 318

2.1.4 Data collection and fusion

Data sources

Intrusion detection is based on information collection and processing. Multiple data sources are available to systems that try to identify suspicious symptoms. The first intrusion detection tools were *monolithic architectures* and would process data at a single location [49]. With the expansion of information systems, more data sources became available and intrusion detection systems diversified by specializing on specific data sources.

There exist two main aspects that can be taken into account for classifying intrusion detection systems [13]. By the *data* they use they can be: (1) host-based systems, or 2) network-based systems. By their *processing methods* they can be classified as: (1) centralized, or (2) distributed. We now describe these four categories which are relevant to the architecture proposed by this dissertation.

(1) Host-based intrusion detection

These systems are designed to protect a single server by scrutinizing large amounts of data contained in audit trails or **log files** as well as other system data structures. These files describe a system by keeping a record of every atomic service requested and every problem reported by the kernel and applications.

The design of audit trail modules and audit data reduction are research areas in information security that aim at improving the quality (i.e., expressiveness and regularity) of the information stored by historical records. Auditing was considered a secondary aspect of operating systems [4, 56] but now, due to the importance of computer security, there are numerous projects devoted to create consistent models that provide the operating system and intrusion detection tools with security-relevant information in a format that can be queried by automated entities [107, 90].

An audit tuple typically consists of fields such as user id, process id, log-on time, permissions mask, and even system-call numbers and system-call types. This information can be used to detect some types of attacks, but it does not suffice for catching others. Not all operating system events are recorded to audit files [56]. There are actions that actually take place but do not leave any trail of its execution. This constitutes a problem for auditing and intrusion detection systems since important data may be missing which would otherwise enable them to find an intruder.

The processing of audit data is expensive as the number of audit records generated by users and applications can be huge (in fact, for audit data to be suitable for intrusion detection it is typically necessary to enable *full-audit mode* on the host. This produces large files that are not easy to handle and prevents real-time execution forcing security checking to be performed off line. Additionally, audit data is prone to modification by an intruder as it is regularly stored on disk.

System data structures are another source of evidence available on the host. As a process executes, it modifies its environment through actions. The changes can be observed on user tables, open-file tables, disk reads and writes, the process table, etc. This sort of data is more reliable from the security perspective as it is more difficult to be tampered with (alterations to system data structures are more evident [125, 101]). The disadvantage of basing a design on this

Sophos Ex. 1009
Page 127 of 318

data is that it is volatile information and hence it is necessary to cache it or put it in a buffer so that it can be analyzed and kept as evidence [46].

(2) Network-based intrusion detection

A network intrusion detection system processes network packets. In Ethernet networks, it sets the network interface in promiscuous mode in order to read all packets [170, 143]. Unlike host-based systems, the information a network-based system uses is not typically stored on a central repository and so it needs to be quickly filtered as it is being read (buffering methods are typically implemented in order to cope with high traffic loads [180]). In general, host-based intrusion detection can detect some attacks that network-based detection cannot and vice versa.

Hybrid systems are not strange [152, 77] but, in general, network-based intrusion detection have been more successfully implemented in working, commercial products [126, 2].

Although the network-based approach has some limitations with switched networks and encrypted connections (a system cannot look into an encrypted transfer unless it has the key), it allows for a prompt detection as the network packets received often contain the first indication of attack [143].

System-call traces, audit logs, and network traffic are the most popular data sources in intrusion detection. Other sources can be incorporated but their relevance is smaller. For instance, user command history can be used to tell what a user has been doing during the day [108, 109]. It stores command names and, possibly, the parameters passed to the shell interpreter. Session variables and alias expansion complicates this approach as well as the fact that users can remove or disable their own command history preferences. Keystroke recorders can also be used [171] but the information that can be captured with them is limited as many attackers use scripts to automate their tasks and this prevents keystroke monitors from succeeding.

Table 2.1 shows some of the most popular intrusion detection systems classified into host-based and network-based systems [223, 13]. Notice that many projects have involved host-based models, however, network-based intrusion detection has been more effective and more widely accepted [2, 126].

Data fusion for intrusion detection

Historically, intrusion detection systems have been integrating more and more data sources to collect information. Data collection can be performed locally by accessing information stored on the host. Research shows, however, that by correlating information from a number of sources, detection accuracy increases [1, 185].

DEFINITION 12

Data fusion. Study of the integration of information records extracted from multiple, heterogeneous sources [16, 185].

Data warehousing tools as well as Web search engines are two applications that have proved the benefits of integrating data. Intrusion detection also deals with this problem and there have been efforts to design architectures that collect information from many distributed locations

Sophos Ex. 1009
Page 128 of 318

Network-based	Host-based
ACME (Advanced Counter Measures Environment) [28] AAFID (*Autonomous Agents for Intrusion Detection*)[224] AFJ(*Anzen Flight Jacket*) [8] BlackICE Sentry [20] Bro [152] CaptIO [29] Centrax [141] CERN NSM (*Network Security Monitor*) [133] CyberCop [138] CyberTrace [44] DefenseWorx [139] DIDS (*Distributed Intrusion Detection System*)[187] Dragon [58] EMERALD (*Event Monitoring Enabling Responses to Anomalous Live Disturbances*) [140] eTrust ID [26] GrIDS (*Graph-based Intrusion Detection System*) [194] Hummer [158] LANguard SELM [70] LISYS (*Lightweight Intrusion detection SYStem*) [82] Manhunt [198] NetProwler [199] NetRanger [36] NetSTAT (*Network-based State Transition Analysis Tool*) [209] Network Security Agent [88] NFR (*Network Flight Recorder*) [164] NID (*Network Intrusion Detector*) [52] NSM (*Network Security Monitor*) [77] OpenSnort Sensor [191] RealSecure [89] SecureNetPro [129] Shadow [145] Snort [170] T-sight [57]	AAFID (*Autonomous Agents for Intrusion Detection*) [224] ADS (*Attack Detection System*) [167] AID (*Adaptive Intrusion Detection*) [188] ALVA (*Audit Log Viewer and Analyzer*) [131] AppShield [176] ASAX (*Advanced Security Audit Trail Analysis*) [75] CARDS [218] Centrax [141] CIDS (*Cooperative Intrusion Detection*) [215] CompWatch [53] CyberCop Monitor [138] CylantSecure [216] DIDS (*Distributed Intrusion Detection System*) [187] DPEM (*Distributed Program Execution Monitor*) [102] Dragon [58] EMERALD (*Event Monitoring Enabling Responses to Anomalous Live Disturbances*) [140] Entercept [123] eTrust Audit [26] FormatGuard [40] GrIDS (*Graph-based Intrusion Detection System*) [194] Haystack [175] HP IDS/9000 [79] Hummer [158] Hyperview [46] IDA (*Intrusion Detection and Avoidance*) [189] IDES (*Intrusion Detection Expert System*) [50] IDIOT (*Intrusion Detection In Our Time*) [105] Intruder Alert [197] Janus [72] JiNao [95] Kane Security Monitor [55] LIDS (*Linux Intrusion Detection System*) [85] MIDAS (*Monitoring, Intrusion Detection, Administration System*) [182] NADIR (*Network Anomaly Detection and Intrusion Reporter*) [81] NIDES (*Next-Generation Intrusion Detection Expert System*) [5] NSTAT (*Network State Transition Analysis Tool*) [97] pH [190] POLYCENTER [47] PRCis [160] Psionic HostSentry [35] Psionic PortSentry [35] RealSecure [89] SecureNetPro [129] Shadow [145] StormWatch [148] Tripwire [100] UNICORN [34] USTAT (*Unix State Transition Analysis Tool*) [87]

Table 2.1: Classification of intrusion detection systems by input type [223, 13].

and consolidate it into a main processing unit [185, 224, 187]. These systems are known as **centralized processing systems**.

Sophos Ex. 1009
Page 129 of 318

Security-relevant evidence can be captured from sources such as network packets, audit log files, system tables, integrity checkers, vulnerability scanners, anti-virus software, and security monitors. There is already sufficient data. The challenge is in identifying the mechanisms to efficiently process this data so that attack signs can be found [18].

Some research teams have built tools that integrate host information and network data [187, 194]. A great deal of effort has been put into developing structured ways of correlating these two types of information so that attacks can be more accurately detected. It is important to mention that even when it is theoretically a good idea to use as many data sources as possible, the diversity of data formats and services significantly complicates this task [185].

Unlike centralized processing systems, there are others that collect evidence in a distributed fashion but do not centralize the processing [187, 194, 95]. These systems are known as **distributed processing systems**. Every component processes its own data and, by means of message passing, they share their findings in order to reach conclusions regarding the security state of the system. The complexity of these systems ends up impacting performance and the overhead they impose is typically unacceptable [92]. The distributed approach is more resilient to failure compared to the centralized model.

Data fusion can be performed by abstract layers. These compose a communication structure in which data-capture modules mechanically send the information they observe to data correlators [81, 5, 140]. These correlators have more intelligence built in and are responsible for the actual data processing. Many intrusion detection systems support hierarchical structures like this and consequently suffer from central-point-of-failure problems [224, 15, 78, 215].

These systems constitute the area of *distributed intrusion detection* which analyzes challenges such as *spam* and distributed denial of service [111, 187]. In recent years, viruses, worms, and distributed attacks have encouraged the design of new solution alternatives. Data fusion has represented a natural way of dealing with the problem, although some data representation and sharing issues need to be resolved in order to make distributed intrusion detection a more reliable and robust solution [112, 187].

Table 2.2 classifies intrusion detection systems into *centralized processing* and *distributed processing* systems according to the data analysis structure they have [223].

2.1.5 Challenges of current models

Intrusion detection technologies still have significant improvement areas that need to be worked out in order to create effective and efficient products. The difficulty of this security problem resides in the difficulty to accurately identify anomalous behavior given the amount, correctness, location, availability, and structure of the data that needs to be analyzed. Some of the main challenges relevant to the topics of this dissertation are described next.

1. *False alarms.* Many intrusion detection tools suffer from detection inaccuracy [2, 90, 4, 126, 13]. The high rate of false positives and false negatives that intrusion detection tools generate does not help system administrators simplify their daily tasks. On the contrary, users may eventually opt for disabling this software rather than having to interpret the stream of alarms reported by the security system that may, in fact, not be reflecting the truth. This is a major limitation of intrusion detection technologies and is the priority for multiple data fusion and model-based projects.

Sophos Ex. 1009
Page 130 of 318

Centralized processing	Distributed processing
ACME (Advanced Counter Measures Environment) [28] ADS (*Attack Detection System*) [167] AID (*Adaptive Intrusion Detection*) [188] ALVA (*Audit Log Viewer and Analyzer*) [131] AppShield [176] ASAX (*Advanced Security Audit Trail Analysis*) [75] BlackICE Sentry [20] Bro [152] CaptIO [29] CERN NSM [133] CIDS (*Cooperative Intrusion Detection System*) [215] CompWatch [53] CyberCop Monitor [138] CyberTrace [44] CylantSecure [216] DefenseWorx [139] Entercept [123] eTrust Audit [26] Haystack [175] Hyperview [46] IDA (*Intrusion Detection and Avoidance*) [189] IDES (*Intrusion Detection Expert System*) [50] IDIOT (*Intrusion Detection in Our Time*) [105] Intruder Alert [197] Janus [72] Kane Security Monitor [55] LANguard SELM [70] LIDS [85] Manhunt [198] MIDAS (*Monitoring, Intrusion Detection, Administration System*) [182] NADIR [81] NetProwler [199] NetRanger [36] Network Security Agent [88] NID [52] NIDES (*Next-Generation Intrusion Detection Expert System*) [5] NSM (*Network Security Monitor*) [77] NSTAT (*Network State Transition Analysis Tool*) [97] OpenSnort Sensor [191] pH [190] POLYCENTER [47] PRCis [160] Psionic HostSentry [35] Snort [170] T-sight [57] Tripwire [100] UNICORN [34] USTAT (*Unix State Transition Analysis Tool*) [87] Wisdom & Sense [205]	AAFID (*Autonomous Agents for Intrusion Detection*) [224] AFJ (*Anzen Flight Jacket*) [8] CARDS [218] Centrax [141] DIDS (*Distributed Intrusion Detection System*) [187] DPEM (*Distributed Program Execution Monitor*) [102] Dragon [58] EMERALD (*Event Monitoring Enabling Responses to Anomalous Live Disturbances*) [140] FormatGuard [40] GrIDS (*Graph-based Intrusion Detection System*) [194] HP IDS/9000 [79] Hummer [158] JiNao [95] LISYS (*Lightweight Intrusion detection SYStem*) [82] NetSTAT (*Network-based State Analysis Transition Tool*) [209] NFR (*Network Flight Recorder*) [164] RealSecure [89] StormWatch [148]

Table 2.2: Classification of intrusion detection systems by their data analysis structure [223, 13].

2. *Amount of data.* The amount of data available within a single host can make the intrusion detection problem intractable. In the case of network-based intrusion detection, the traffic

Sophos Ex. 1009
Page 131 of 318

going through a gateway, for instance, can be prohibitively large making it practically impossible for a security scanner to check every single packet without dropping some of them [143, 4]. Packet dropping severely damages the effectiveness of this type of protection tools.

Similarly, audit facilities can generate very large log files which limits the likelihood of performing real-time processing [4]. What this does is that audit analysis has to be moved to the end of the day for off-line checking, which is not as desirable as real-time detection.

Intrusion detection is a needle-in-the-hay type of problem: the amount of "normal" activity data is huge compared to the pieces of irregular activity that need to be found.

3. *Performance.* Data filtering and screening is a computationally demanding task that is central to intrusion detection [126]. Providing a system with real-time monitoring capabilities implies good search and pattern matching algorithms as well as efficient buffering of network traffic, audit trails, etc. which guarantees no data piece goes through without being inspected. Given the detailed checking needed by intrusion detection, a lot of resources are allocated and consumed and the performance of the hosting system ends up being impacted. As a consequence, heavy intrusion detection software and business processes cannot cohabit and the former is sometimes disabled to give full privileges to priority services [60].

4. *Switched networks.* Modern networks utilize more intelligent devices to route packets from one location into the other. Many intrusion detection prototypes were based on the ability to configure a network device in promiscuous mode (e.g., [152, 170, 5]) in a way that the monitoring mechanism could easily see and analyze all Ethernet packets. This allowed the system to correlate packets addressed to different hosts connected to the same segment, and also to protect several hosts at the same time. With modern network technology, this approach may no longer be as successful.

   Switches only forward packets to its intended destination to optimize bandwidth use [38] (i.e., broadcasting is limited). Now, the number of security software instances that need to be installed is larger and network-based systems need to be strategically placed so that they are able to protect as many hosts as possible inside a LAN. The filtering functions of a switch which represent a performance advantage from the bandwidth perspective, are in fact a challenge to the protection of networked systems [126, 2].

5. *Encryption.* Encrypted data cannot be easily interpreted unless a decryption key is available [179]. Information technology has witnessed the development of cryptographic protocols and tools that protect the privacy of data. File transfers, interactive sessions, and email can all be protected with encryption. Although this secures the information they handle, it renders further security analysis impossible – necessary in the case they transport illicit material [142]. As more companies become security-conscious the use of encryption will increase [169, 124, 48]. Intrusion detection tools are currently inadequate to deal with this problem as they are blind to encrypted communications.

6. *Security of intrusion detection technology.* Security tools are expected to be secure themselves – otherwise they could not be trusted and their results would be worthless [2, 42]. As software products, intrusion detection systems face the same security challenges other applications do (e.g., poor implementation and weak design). In fact, not much attention has been put

Sophos Ex. 1009
Page 132 of 318

into the secure design of this type of software [134]. Only a few systems can be considered moderately secure but the rest is prone to direct or evasion attacks [13, 161]. This is an aspect that cannot be ignored when designing protection systems.

7. *Reaction to incidents.* Although it is not part of the traditional definition of intrusion detection [4, 60, 126], response to incidents is a very desirable feature that a few intrusion detection systems have actually tried to automate [72, 215]. Taking action based on incorrect information can have serious consequences especially when done on production system. The problem of automated response has remained unexplored due, in great part, to the low detection-accuracy levels intrusion detection tools achieve nowadays [2, 13]. This area will develop when detection algorithms become more precise and when intrusion alarms reveal meaningful security events.

2.2 Recent trends in information security

Over the last five years, protection tools, security threats, and cyber-crime have presented behavioral patterns that reveal changes in the way computer and information security is conceived. These trends need to be studied and fully understood so that new technology can be developed to effectively mitigate the damages caused by worms, viruses, hacking, espionage, and other modern threats.

1. *Decreasing average technical expertise in attackers.* The number of security incidents has been growing steadily with the years reflecting the expansion of the Internet. Recent cyber-attacks are very serious and extremely sophisticated in nature [124]. This is in part due to the increasing complexity of software along with the fact that more people are making use of security software which makes break-ins more difficult [169, 124].

   The average technical knowledge of attackers, however, has decreased (Figure 2.12). It is no longer necessary to be a systems expert for launching an attack on the Internet. A few people create scripts and automated tools which are eventually obtained by a larger number of novice users. They may not know the details of how the exploits work, but have the curiosity of seeing what they actually do. These users represent the most dangerous group as they do not know the consequences of the attack tools they execute [42].

2. *Accelerated growth in the number of vulnerabilities and attacks.* One of the clearest trends in information security is the fact that more and more vulnerabilities are discovered everyday and, consequently, more attacks that exploit them also appear in the scene [169, 124, 48]. Figure 2.13 contains some numbers that show this trend.

   The purpose of the release of vulnerability information is to make users and administrators aware of the existing threats. However, this information is also used by attackers that devise new techniques to abuse systems.

3. *Increasing number of external attacks.* Years ago the rule was to see internal attackers abusing their privileges to compromise a local host. External attacks were not as common, but that has changed [201]. With the connectivity provided by modern networks more people

Sophos Ex. 1009
Page 133 of 318

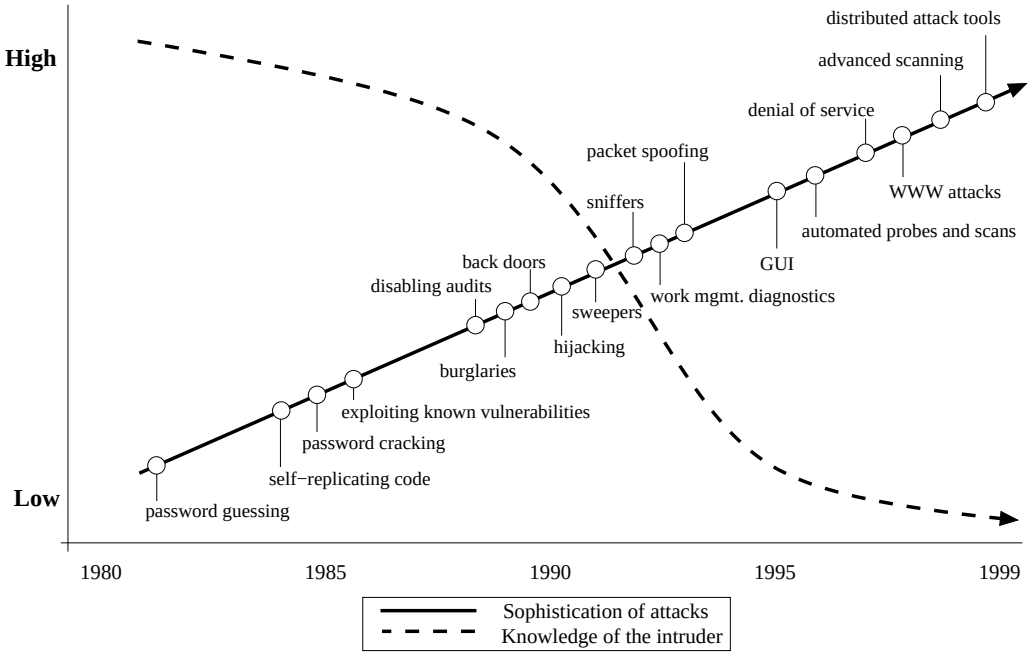


Figure 2.12: The increasing sophistication of attacks and the decreasing expertise of users [2, 126].

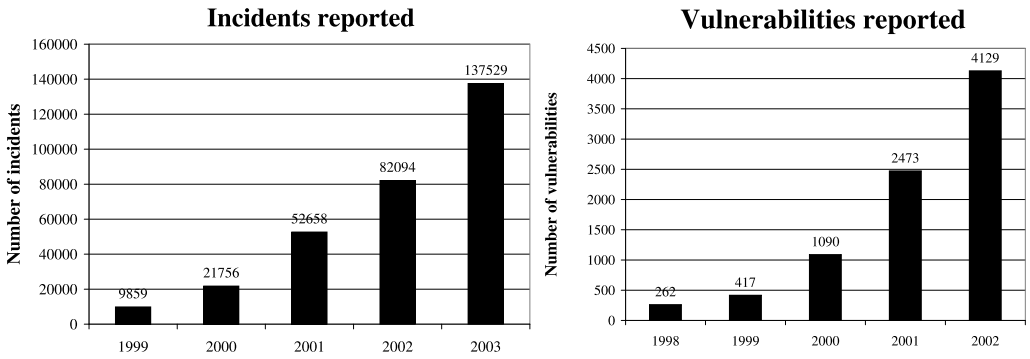


Figure 2.13: Number of security incident and number of vulnerabilities reported per year to the CERT–CC (*Coordination Center, Computer Emergency Response Team* [32]).

have access to the Internet and have the potential of becoming an attacker. As a result, the number of attacks coming from the outside, according to the information reported by many companies, has significantly increased [169, 165, 2, 42] in parallel to the discovery of remotely-exploitable vulnerabilities which the on the raise [201]. Given the visibility administrators have of their own systems, internal users are easier to catch, that is why this trend reveals an increasing number of incidents originated by outsiders who can very well hide through

Sophos Ex. 1009
Page 134 of 318

connection hopping in the net [118, 117].

4. *Increasing use of security technology.* Firewall software, authentication systems, and anti-virus technology are three common weapons security administrators usually have on their belts [169, 165]. According to recent surveys, intrusion detection programs are being added by more people as a component of the typical corporate security suite (Figure 2.14). Although this demonstrates the importance of security monitoring and encourages the creation of better tools to help identify illicit activity, vulnerabilities and exploits are on the raise using more complex methods to defeat the security mechanisms in place.

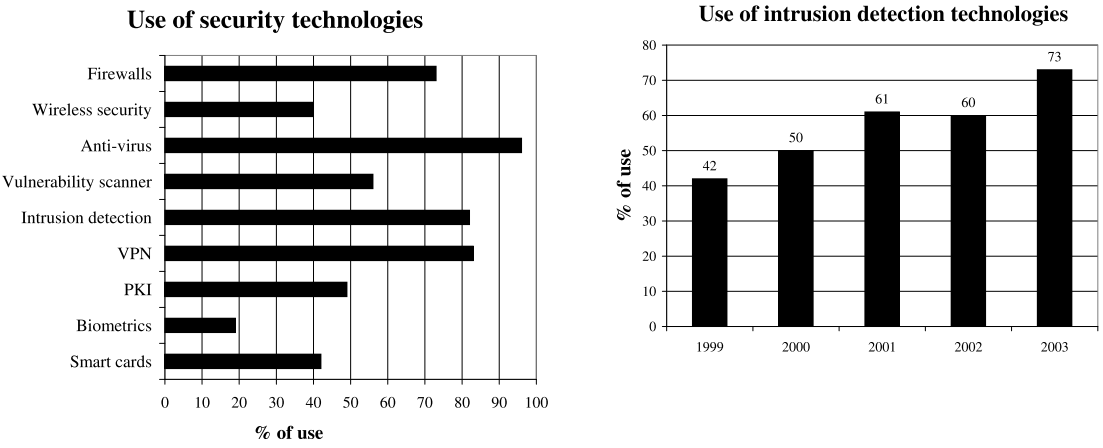


Figure 2.14: Use of information security technologies (*2003 Global Security Survey, Deloitte & Touche* [48]) and intrusion detection technologies (*CSI/FBI Computer Crime & Security Survey 2003* [169]). The charts show the percentage of surveyed companies that report using this type of software.

5. *Increasing availability of automated attack tools.* Recent statistics on software vulnerabilities and security reports coincide by indicate that an increasing majority of attacks on the Internet are being executed with the use of automated attack tools [169, 165]. Years ago there were not many websites that stored programs that exploited vulnerabilities automatically. Given the openness to report software bugs as a security strategy [163, 10] more users have access to the details of how to build programs that automate systems probing, scanning, and the actual exploitation of security wholes [103]. The increasing number of attack tools available on the Internet coupled with a low average of technical expertise from users is a combination that represents a severe challenge to information security researchers and practitioners.

6. *More Distributed Denial of Service Attacks.* The number of denial of service attacks (i.e., attacks aimed to exhaust the resources of a host by flooding it with service requests) orchestrated from remote locations keeps increasing (Figure 2.15 ;[169]). These attacks sometimes compromise a group of machines that will be later used to launch an attack using remotely-managed tools toward a common target.

Sophos Ex. 1009
Page 135 of 318

The inability of current network protocols – in particular the TCP/IP family – to preserve the integrity of packets makes of this attack one of the most difficult to defend against [111, 161]. Every field in a packet can be altered in such a way as to evade security mechanisms, which complicates the task of packet filtering and source back-tracing. The latter has been studied for several years now and it is still considered a major challenge due to the underlying network infrastructure we use to keep the Internet running [222].

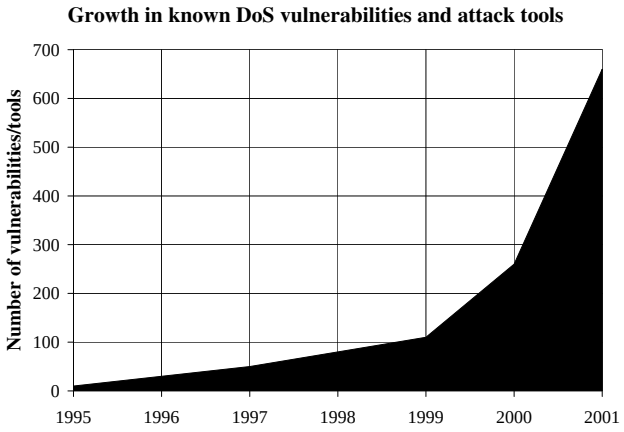


Figure 2.15: Number of known Denial of Service (DoS) vulnerabilities and attacks per year (*PestPatrol's detection database, PestPatrol Inc. 2002* [155]).

7. *Network hopping for attack relaying.* Attackers protect their identity by hopping among a series of hosts before launching an attack. This connection chain hides them from administrators and law enforcement as it is very difficult to identify the exact location from which a connection is truly coming [118, 11]. Additionally, proxies and anonymity tools can be used to complicate the task of identifying the root of an incident, an important aspect in a cyber-crime investigations [168, 73].

8. *Increasing exploitation of recent vulnerabilities.* The time between the discovery of vulnerabilities and the appearance of a number of attacks exploiting it has been decreasing [124]. In the past, two or three-year-old vulnerabilities were frequently used to compromise systems which had not been patched. The vulnerability-age range was considerably large. Now, more than 50% of the attacks involve vulnerabilities reported during the last five months. Two factors have influenced this trend. Firstly, vulnerability information is more readily available on the Internet and an attacker planning to compromise a system can look for the most recent vulnerabilities to increase the odds of spotting vulnerable systems somewhere. Secondly, more companies are becoming security-aware and this has reduced the number of systems that have old bugs or misconfiguration problems [169].

Sophos Ex. 1009
Page 136 of 318

2.3 Connection back-tracing

Given a set of hosts $H = \{h_1, h_2,..., h_n\}$, with $n > 2$, a user could create a connection chain of m hops as follows: he could connect from host h_i into h_{i+1}; from h_{i+1} into h_{i+2}; from h_{i+2} into h_{i+3}; and so on, until finally connecting to target host h_{i+m}. The problem of finding the hosts that compose the connection chain and identifying the root host h_i given H, h_{i+m}, and user name u on h_{i+m}, is known as the **connection back-tracing problem**.

For a small $n = |H|$ this problem is straightforward. Just a few hosts need to be analyzed in order to identify the source of a connection sequence. However, as n becomes larger (e.g. $n \gg 2$), the complexity of the problem increases since a larger number of hosts in H multiplies the possibilities of creating a more confusing connection path.

2.3.1 Limited troubleshooting visibility

System administrators have limited access to information regarding who is connecting to one of their servers. On a LAN, administrators typically have login access to all servers, which makes the tracing process easier. However, for connections coming from the outside of the LAN, they are not able to freely access all other servers for scrutiny. The visibility of this type of connections reduces to one link in the connection chain. That is, an administrator can only see the terminal number or the IP address/host-name of the immediate host from where users are logging on. He is unable to tell, however, whether that host is in fact h_i (the workstation of the attacker). In order to find out whether it is the root host or not, he will have to contact the administrator of the remote system who, in turn, will have to perform a similar checking.

This job is cumbersome and requires time, communication, and trust from the parties involved. For the case of a connection that has been closed already, audit information will have to be examined in order to find the nodes of the chain that lead to the origin of the connection. This off-line sort of tracing has an important flaw: it counts on the existence and integrity of audit log files. For clean systems this would not be a problem but for systems that have been compromised, and in the case of a connection chain created to perpetrate an attack this could probably be the case, these files could have been tampered with during a previous break-in [222]. On the other hand, if a connection is still alive, i.e. there is a packet flow between source and destination, information contained on system tables and network packets can be used to trace the user while he is on-line.

2.3.2 Related projects

There are some projects that touch on the problem of connection-route tracing. One is the method proposed in [222]. This model implements network traffic monitors that perform session tracing once a server has been compromised. A system that computes deviations from the traffic observed at two different hosts helps determine if these hosts were used in the same connection chain. Since checking network traffic at a node involves huge amounts of data, multiple packet monitors permanently filter and record specific system activity to build their own logs. These files are analyzed by a data correlator to identify the hosts of a connection chain. It is important to notice that this system deals with the problem of deleted logs by creating its own records. These records, however, may be the target of the same type of attack.

Sophos Ex. 1009
Page 137 of 318

IDA (Intrusion Detection Agent system) [11] is another system that detects the origin of an information exchange related to an incident. It is primarily an intrusion detection system that employs mobile agents to detect local attacks. From a main host, investigation agents are sent over to the requester that needs an integrity check. These agents use the MLSI (*Marks Left by a Suspected Intruder*) strategy to reduce the amount of data needed to flag system activity as anomalous. In order to collect further information on a break-in, IDA does passive route tracing using audit log data. The structure of the tracing method is not specified by the authors but its agents broadcast connection evidence to their peers as a way of identifying possible hosts involved in the attack.

The *TBAIR* system (*Tracing-based Active Intrusion Response*) [213] is a network-based intrusion detection tool that, unlike most traditionally-passive intrusion detection systems, adds active response to security events using a technique called *Based on Sleepy Watermark Tracing*. TBAIR tries to attack the root of the vulnerability exploiting problem by locating the originators of the incident to hold them accountable for their intrusions and their method is based on the analysis of evidence collected from the network.

Thumbprinting [195] is a technique that places several processes across the network in order to capture activity signatures or "thumbprints". This method does not require having processes on all hosts and, like in the case of [222], it replicates packet information for future review. This system is based on the fact that the packet content of a session passing though a set of hosts is invariant and can be compared in order to identify a host in the chain. This system is capable of detecting the root of a connection but it is not designed to identify the exact hosts that are part of the chain.

2.4 Agent-based intrusion detection

Multiple artificial intelligence techniques have been applied to information security problems. There are authentication systems that use artificial neural networks to perform fingerprint and retinal-pattern recognition in order to identify the users of a system with a high precision level [18, 137, 156]. Fuzzy logic, decision trees, and support vector machines have also made possible the creation of face recognition and keystroke profiling tools to enhance access control, a critical security component of a system [171, 120, 219]. Redundancy through replication using autonomous agents has been also studied in order to strengthen the robustness of high-availability platforms [63, 122] (Figure 2.16).

Agent-based software in many different areas have helped understand the security requirements of services and applications [14, 21]. The field of intrusion detection has not been the exception and artificial intelligence has also contributed with some solutions [224].

Data mining software, for instance, has been built to monitor the security of a system and has been capable of successfully identifying electronic fraud and, in some cases, back-trace intruders [114, 113, 11]. A number of anomaly detection techniques have been implemented using neural networks and support vector machines [135, 174, 27, 46]. Their performance is still a challenge but further research in this young area is likely to produce better solutions.

Inside a host, different data sources need to be consolidated in order to discover distributed attacks. Outside of the host, networking has allowed the integration of diverse platforms with the purpose of supporting electronic services and resource sharing. This integration has expanded

Sophos Ex. 1009
Page 138 of 318

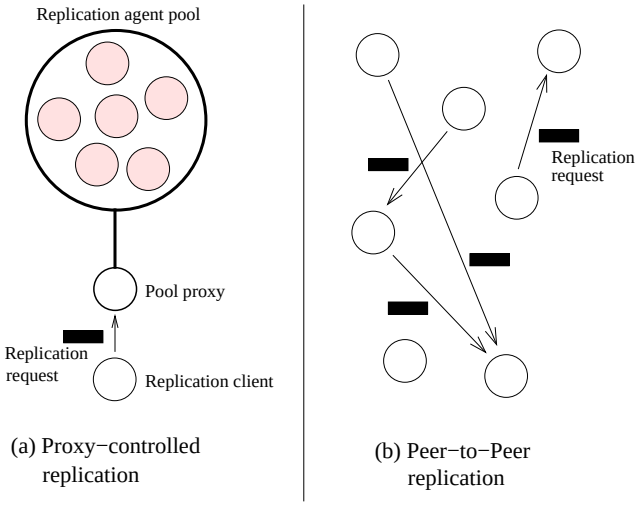


Figure 2.16: (a) Replication is performed by a pool of agents which are contacted through a proxy agent. (b) Replication can also be peer-to-peer.

the coverage area for which a monitoring module is responsible. It is no longer a single server or application what needs to be protected. Every single component can be vulnerable to attack and needs to be secured. Even if one of these components has not been apparently compromised, the activity observed by its neighbors may indicate the opposite – a victim does not always have enough visibility as to infer it is being attacked [60].

Intrusion detection systems need to be able to correlate input from multiple local or distributed sources in order to prevent attackers from gaining unauthorized access to it. The data integration nature of intrusion detection has made researchers look at the agent and multiagent approach in the search for more and better solutions [224].

2.4.1 Agents

The definition of agent we use throughout this dissertation is the following:

DEFINITION 13

Agent. A computational entity with a certain degree of autonomy that perceives its environment through sensors and modifies it through effectors [173].

Some authors suggest agents need to be independent enough as not to require intervention from humans (otherwise they could be considered mere software applications [214, 3]). In order to reach that independence or autonomy, it is necessary for agents to be able to decide what action to take next based on their input perceptions. This decision may be influenced by a utility or goal function that gives meaning to the agent (in the case of security monitoring agents, rarely are these functions explicitly defined [41, 224]). If all of the agent's actions are selected as to reach this objective based on the available information, the agent is said to be rational [172].

Sophos Ex. 1009
Page 139 of 318

Rationality is another desirable characteristic of agents and recent publications emphasize the issue of how to determine whether or not an agent is trying to achieve its goal, and how well it is doing in the process [173]. Performance functions are employed to evaluate the agent's performance and a rational agent will always try to maximize this function regardless of the uncertainty of the environment and just based on its perceptions and prior knowledge.

Sensorial information (i.e., perceptions) is key to an agent's success. It allows it to make decisions and acquire knowledge. Learning is a process through which an agent will use perceptions from the past to make well informed decisions at present predicting the future [159, 173]. The knowledge-base of an agent captures information meaningful to the utility function and that helps an agent achieve its objectives.

There are other characteristics that, depending on the application, are desirable in an agent. Such is the case of mobility, persistence, initiative, communication, coordination, and sociability [173]. The latter three represent important considerations of what researchers have called *distributed artificial intelligence* [214].

2.4.2 Multiagent systems

An agent can be designed to work alone but in many practical scenarios multiple agents are executed to accomplish a single task. Sociability is the ability an agent has to communicate and interact with its peers in order to accomplish such task and **Distributed Artificial Intelligence** (DAI) is a discipline that studies the use of collections of agents to perform joint tasks [172, 151].

DAI comprises two broad issues: (1) finding ways to split a certain problem into subproblems so that it can be solved more quickly by having several entities working concurrently, and (2) utilizing multiple existing agents to solve a problem. The slim line that divides these domains has in fact created two different research areas, namely, *distributed problem solving* and *multiagent systems* [146].

Distributed Problem Solving (DPS) has to do with the analysis of the structure of problems and their potential solutions. Some problems can be divided into into several smaller problems whose solutions can be obtained separately and then combined to resolve the main problem [116]. The problem could also be transformed into an equivalent representation that allows parallel processing. Distributed problem solving has an algorithmic perspective that differs from the Multiagent Systems approach to artificial intelligence [54, 146].

Multiagent Systems (MAS) are primarily concerned with the behavior of a collection of agents. It focuses on developing mechanisms to get groups of agents to cooperatively accomplish a goal they could not solve individually. This includes problems like collaboration, negotiation, coordination, and knowledge sharing where agents need to use their skill-set to communicate and solve, as a team, a common problem [14, 146, 25, 86].

From a different perspective, we can also define DAI as a subarea of MAS. Agents in DAI are mostly cooperative and friendly. MAS studies this basic behavior plus more complex cases in which agents negotiate, complete, and can even become selfish..

Information security researchers have been interested in agents for two main reasons. Firstly, agents are implemented in software and, as such, are prone to attack. This constitutes an interesting research problem as the mobility and communication features of agents complicate

Sophos Ex. 1009
Page 140 of 318

the task of securing information and guaranteeing agent availability [74, 122, 157, 45, 62]. The problem of agents trying to compromise their peers has been studied. However, the fact that the actual execution environment can be hostile represents a threat to the integrity of agent-based software and is a topic that has proved challenging and more interesting [177, 178, 19]).

Secondly, and being the aspect that occupies this dissertation, a multiagent system can serve the purpose of carrying out security-related tasks [78, 93]. System administrators take advantage of their experience in order to identify actions that could uncover an intrusion. That experience coupled with their ability to correlate information and draw conclusions, is something that needs to be built in computer intrusion detection systems. The automation of this process through monolithic intrusion detection programs [49, 182, 50, 75] has due to challenges like the ones we presented in previous sections (see Section 2.1.5). For that reason, there have been different architectures that integrate agents to accomplish the goal of providing a host with monitoring capabilities. Following a description of the components of some of these architectures.

2.4.3 Multiagent components of intrusion detection architectures

Some agent-based architectures have been implemented in the past for the problem of intrusion detection [25, 224, 158, 15]. They all share structural aspects and component types that are mentioned next and that constitute the base of the system presented here.

Agent-based intrusion detection architectures

The term "agent" has been overused in many information technology areas [21] including computer security. The aim and characteristics of agents have not been yet fully understood by, and integrated with neighbor areas which are often attracted by what this promising computing approach has to offer.

Intelligent capabilities such as learning, autonomy, and pro-activity are rarely seen in agent-based intrusion detection implementations. The architecture by Balasubramaniyan et al. [15], for instance, displays communication between processing entities limited to the structure of a message-passing mechanism. It does not enable agents to do inference with the information contained in the messages they receive, and the module interdependence of the system prevents even the most basic levels of agent autonomy.

Crosbie and Spafford [41] also designed an architecture to implement system monitoring. Some of their agents perform data collection and analysis but they were not designed to use seen events or to evaluate the success/failure of their re/actions (i.e., they do not accumulate any sort of experience) to improve on their own performance. The Hummer project [158] and the AAFID architecture [224] are agent-based intrusion detection environments where data is captured, inspected, and correlated to perform intrusion detection (Figure 2.17). Some of the hierarchical structures used to integrate the components of these models, however, severely limit agent autonomy and constitute potential attack points and/or bottlenecks that cannot be present in robust security software.

An element required for some agent functions is a layer which provides a common interpretation of the environment. This layer is missing in all of the agent-based intrusion detection projects reviewed in the literature. An ontology could provide security agents with a better

Sophos Ex. 1009
Page 141 of 318

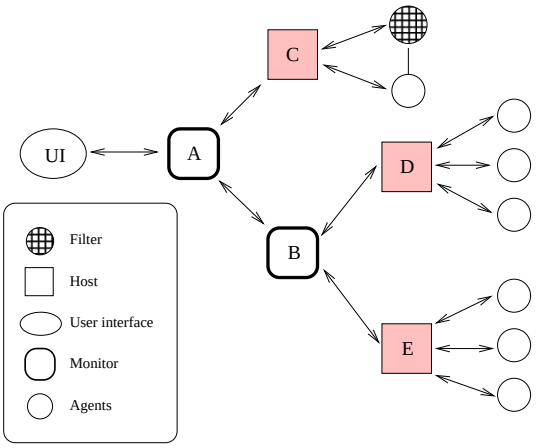


Figure 2.17: Hierarchical component of a multiagent intrusion detection architecture [224].

foundation for more naturally integrating components and sharing information while enabling the development of more adaptable and useful intelligent software [23].

DEFINITION 14

Ontology. The specification of the objects, concepts, and relationships in an area of interest [214].

By having an ontological model as the underlying element of a multiagent platform,

Agent communication and development would benefit from an underlying ontological model that gives the agent platform solid ground for building more ellaborate functionality for intrusion detection tasks. Figure 2.18 depicts a logical composition of blocks representing fundamental components of an intrusion detection system based on intelligent software. When built on top on an ontology, agents can be more easily maintained as the overall functioning of the system is guided by the definitions and relationships stored on the ontology.

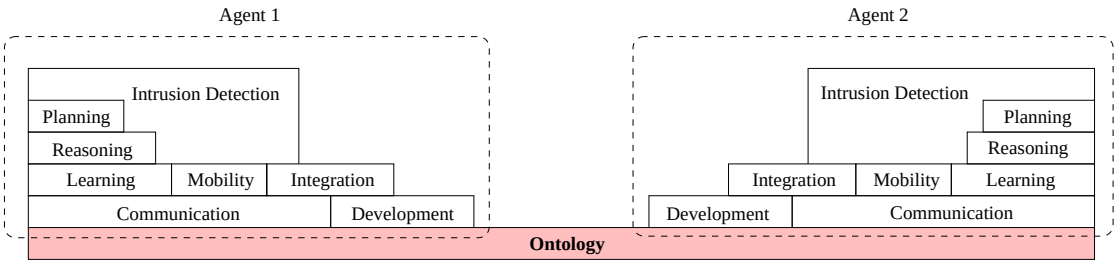


Figure 2.18: Enabling high-level agent functions through an ontology layer.

Sophos Ex. 1009
Page 142 of 318

Agent-based intrusion detection architectures

There is a series of tasks that are necessary to perform intrusion detection and which can be distributed among a group of agents. They can be summarized in three major tasks: *data collection*, *integration*, and *correlation* or *reasoning* [4, 192, 103].

- **Sensors**. Data collectors are usually referred to as *sensors* in intrusion detection [215, 78] (notice that sensors are also pieces of code used to detect anomalies when embedded into a program; see [223] for details). Sensors can be placed at different locations within a network in order to capture information that may contain signs of intrusion. They could be agents that simply read and forward to a central authority all the information they receive, or they can be enabled with filtering capabilities that allow them to pre-process the data before sending it to other agents. Further more, sensors can have intelligent functions built in – through learning and reasoning techniques, e.g. – so that they not only reduce and filter, but also classify and check for consistency the data they make available to its peers [41].

- **Integrators**. Integration agents receive information from others, condense it, and re-write it using formats higher-level agents can understand. The role of integrators is to relieve other layers from structuring and sanitizing the data, as well as to take over the role of querying every sensor that may have acquired interesting information security-wise. Interrupt-based or polling methods are equally used for the communication of these systems [192, 224].

- **Correlators**. The agents that implement the core of intrusion detection algorithms are the *correlators*. Their objective is to use the data they are sent in order to detect signs of intrusion. In order to do this, they need to be equipped with mechanisms that allow them to create relationships between the pieces of information they have. Decision trees, inference models, and neural networks are just some of the techniques that have been utilized by this type of agents to accomplish their goal [174, 27, 215, 45].

There are several ways to organize these agent classes. Each agent can be responsible for reaching a specific subgoal whose results are to be sent to other agents in charge of processing low-level information. Most existing intrusion detection models use a hierarchical structure like this [224, 15, 106, 78, 41] (Figure 2.17). However, recent architectures have proposed peer-to-peer communication where a self-contained agent is capable of performing several task by itself [92]. A hierarchical structure is more efficient from the performance point of view but, since the number of agents that compose a security architecture is typically small [15, 224, 41], a peer-to-peer system provides desirable advantages from the availability point of view that make it a good alternative.

2.5 Taxonomies and ontologies in information security

Systems security has struggled with the fast-paced expansion of the Internet. As a result, the formalization of many concepts is still missing as it is the case with intrusion detection which lacks structured representation, communication, and classification models. This section describes some relevant efforts in these areas related to the objectives and motivation of the present dissertation.

Sophos Ex. 1009
Page 143 of 318

- **Anderson's attacker classification**. One of the earliest classifications to appear in the area of intrusion detection was the definition of attacker types [6]. It divided attackers into *external penetrators* (i.e., outsiders), *masqueraders* (i.e., users who steal identities from others), *misfeasors* (i.e., those who abuse the privileges they have been granted), and *clandestine users* (i.e., those who are given access by another user but are not known by system administrators and audit records). This classification is important since it was one of the initial efforts to understand the phenomenon of illicit system use.

- **Kumar's taxonomy of computer intrusions**. This work states that intrusions are similar among one another and can be grouped into generic classes described by patterns [103].

  This project presents a hierarchy for classifying intrusion signatures based on the structural interrelationships among the events that compose them. For instance, a file-write event could be preceded by a permission-change event at the beginning of a particular intrusion scenario and the sequence of these two events – with their corresponding parameters and the existence of some execution variables – could be common to more than one intrusion type. A signature can be created as to represent this precedence and all intrusions that matched the signature will be put together in the classification tree.

  The events defined by this taxonomy are composed of lower-level elements such as system calls and system constructs (e.g., i-nodes, permission masks). This taxonomy shows how colored Petri nets can be used to model the classification and matching system (similar signatures share the same Petri net; [105]). Even though this taxonomy was not designed to classify vulnerabilities but the signatures to detect when they are exploited, it became the foundation of a vulnerability database [128].

- **CVE classification of vulnerability types**. The *Common Vulnerabilities and Exposures* initiative (CVE) is a list of standardized names for vulnerabilities and other information security exposures [121, 33]. The goal of CVE is to make it easier to share data across separate vulnerability databases and security tools [127]. The content of CVE is a result of a collaborative effort of an editorial board which includes representatives from numerous security-related organizations such as security tool vendors, academic institutions, and government as well as other prominent security experts.

- **Axelsson's taxonomy of intrusion detection systems**. This project classifies intrusion detection systems according to nine different functionality, structure, and security attributes [13]. These attributes are: (1) detection principle (anomaly, policy, hybrid), (2) detection time (real-time, non-real-time), (3) granularity (batch, continuous), (4) data source (host, network, hybrid), (5) type of response (passive, active), (6) data processing (centralized, distributed), (7) data collection (centralized, distributed), (8) security (low, moderate, high), (9) interoperability (low, moderate, high).

  The major merit of this classification system is that it employs frequently-used concepts that appear in the literature describing security systems and methods. This helps compare intrusion detection tools from the more fundamental perspectives.

- **CIDF: A framework for the cooperation of intrusion detection systems**. The *Common Intrusion Detection Framework* (CIDF) was an effort to develop specifications that

Sophos Ex. 1009
Page 144 of 318

enabled different intrusion detection components to interoperate and share information as well as to allow for the re-usability of modules [196].

The CIDF is based on the assumption that intrusion detection systems (at least the ones developed by projects funded by DARPA) can be modeled through standard components such as sensors, analysis engines, and response elements. The project was motivated by the lack of standardization. The group defined a language to describe objects and relationships through LISP-like expressions that conveyed semantic information. (This language is in fact similar to XML [162] and the proposed structure could be actually implemented with modern ontology-representation languages such as DAML+OIL and OWL).

The project did not have the impact the security community was expecting but it did leave a precedence on how to define and interleave intrusion detection concepts into a single common, meaningful structure. The CIDF motivated the creation of the IDMEF project [43] and the terminology it defined has been used to develop portions of the communication systems of some research projects [112].

- **IDMEF: A standard format for intrusion detection communications**. Motivated by the fact that there is no standardized way for intrusion detection systems to communicate, a group funded by the Internet Engineering Task Force developed the *Intrusion Detection Message Exchange Format* (IDMEF) [43].

  This is a specification for a standard alert format and a standard transport protocol. In fact, it also defines the requirements and structure of intrusion detection systems. These specifications use XML to encode alerts which contain attributes a large subset of intrusion detection products use to report incidents. It defines data types, message classes, alert classes, among other concepts.

  The benefits of using XML certainly give the IDMEF the portability and flexibility intrusion detection systems require to exchange and store critical information, and represents a good way to link intrusion detection tools with other technologies through a standardized language.

  The IDMEF format is fairly new but it has been already implemented in a Snort [170] plug-in to report alarms.

- **A target-centric intrusion detection ontology**. Most terminology used in intrusion detection is described from the point of view of the attacker [103, 126]. Undercoffer and his team developed an ontology to define intrusion detection terms and relationships based on the analysis of over 4,000 classes of computer attacks [203, 202].

  The core structure of the ontology contains types of system targets, mechanism of attack, consequences of attack, and location of the attacker. It is argued that the taxonomic characteristics used to describe attacks should be observable and measurable by the target system. This ontology has been implemented in DAML+OIL and can be used to exchange information between security tools that share the resources this model defines (this ontology is intersected by the outbound intrusion detection ontology proposed in this dissertation; see Chapter 4 for details).

Other related projects have proposed alternate ways to use ontologies in information security problems. Such is the case of an extension to a software engineering ontology that is being

Sophos Ex. 1009
Page 145 of 318

developed to include security considerations while software agents are being defined (like other applications, agent-software is not typically designed with security in mind; [134]). There are new projects in the area of knowledge management and representation that are still in their early stages. They promise to contribute important ontological constructs in the future to the security and intrusion detection areas, as pointed out by Raskin et al. [166].

Sophos Ex. 1009
Page 146 of 318

		Approach		Foundation		Data source		
	Technique name	**AID**	**MID**	**Probabilistic method**	**Learning method**	**Audit data**	**Network data**	**Other**
1	Frequency analysis	X				X	X	X
2	Probabilistic detection	X		X		X	X	
3	Bayesian networks			X	X			
4	Conditional event sequences	X		X		X		
5	Expert systems	X	X			X	X	X
6	Inductive generalization	X		X		X		
7	Neural networks	X			X	X		
8	Attack scenarios		X	X		X		
9	Petri nets		X			X		
10	Genetic programming		X				X	X
11	Embedded sensors	X	X			X	X	
12	Functional dependency	X						X
13	Data mining	X			X	X		
14	SVM's	X			X	X	X	
15	Planning		X			X		
16	Bioinformatics	X				X	X	X

Table 2.3: Classification of intrusion detection techniques

Sophos Ex. 1009
Page 147 of 318

Chapter 1

Introduction

Existing computer security tools are focused on the protection of local resources. A different, unselfish approach would be to protect a system by making sure that others are not able to launch attacks toward it. By monitoring local activity and outbound traffic, a system should be able to identify compromise attempts originated locally so that they are immediately stopped. This dissertation explores this idea. The design and implementation of a prototype architecture serves as a proof-of-concept that also provides experimental results regarding the use of agent and semantic web technology in the development of security software. This first chapter describes the overall aim and structure of the dissertation.

1.1 Context and basic concepts

This section presents basic concepts pertaining the area of computer security and discusses the need for having monitoring systems that complement secure software design and systems maintenance practices in the task of guaranteeing the security of a computing platform.

1.1.1 Information security fundamentals

Information security has been a very active research area over the last two decades. High connectivity and massive access to network resources combined with the widespread use of vulnerable software has resulted in hundreds of security breaches that leave considerable losses behind [169] putting at risk computer systems used to conduct the electronic operations of corporations, governments, and individuals worldwide [126, 124].

Security is a complex term involving technical, social, and economic aspects. For the purposes of this dissertation, which concerns the security of computer systems, the concept of information security or simply **security**, can be summarized as follows:

DEFINITION 1

Security (1). The ability a system has to prevent failure and to achieve its design objectives in spite of failure – either it be an intentionally induced fault or a fault caused by an operational error.

Sophos Ex. 1009
Page 149 of 318

A secure system should behave as it is supposed to while resisting attempts to trick it into performing actions that are not part of its expected functioning [18].

The objectives that information security pursues can be condensed in five main goals that help understand the specific purpose of existing security software [179, 156]:

1. *Confidentiality.* To protect information from unauthorized disclosure (e.g., using encryption).
2. *Integrity.* To guarantee that an information resource is not unexpectedly modified so that it remains trustworthy (e.g., protecting documents using secure hash algorithms).
3. *Availability.* To ensure that a resource is accessible at the time it is required (e.g., implementing controlled-redundancy strategies).
4. *Authentication.* To effectively validate the identity of a party (e.g., using password-based access control).
5. *Non-repudiation.* To guarantee that no party is able to deny actions it actually performed (e.g., implementing digital signatures).

From these five points we can derive a functional definition of computer security that some experts and practitioners in the area use:

DEFINITION 2

Security (2). A concept that integrates the confidentiality, integrity, and availability of a system as well as its ability to guarantee authentication and non-repudiation to all users and applications.

Each of these aspects has different priority levels depending on the protected resource. The value of a resource determines what aspects should be protected and how much it is worth investing in its protection. A database table, for instance, may contain strategic information and sufficient resources should be allocated for its protection at all times. Preserving the confidentiality of a message, on the other hand, may be necessary for just a short period of time after which the message becomes public and it is no longer necessary to keep it private, or it may not matter who reads it as long as it remains unchanged. These guidelines or rules that describe the set of acceptable actions are the components of a **security policy** [156, 126, 18].

DEFINITION 3

Security policy. Set of guidelines that indicate what and how resources should be protected by explicitly specifying all the valid ways in which the components of a system are allowed to interact.

Security policies are a basic construct in information security. They are used to differentiate abnormal from normal behavior. Since software is used to handle and store electronic information, it should be designed as to guarantee compliance with the security policies at the site and must be robust enough as to resist attack attempts.

The growth of the Internet has multiplied the number of inexpert users. Many of them are not familiar with even the more basic concepts of security and are often in charge of computer

Sophos Ex. 1009
Page 150 of 318

systems and data banks [2, 4]. As a result, security policies do not always exist and the way users try to protect their systems is rather empirical, insufficient, and incorrect [18]. Living in a market-driven economy where often what matters is to deliver a product by the deadline regardless of quality, security is not always a priority for software firms. They overlook this aspect and release products that may be very good from the functionality point of view, but that are weak at the security level. If these products are not properly configured and maintained, they will become vulnerable systems that could be the target of attacks or be infected by viruses. In either case, the productivity of a company could be impacted [169, 124].

Computer security is a real problem and also a very fertile research area in computer science. New threats appear every day challenging the quality and configuration of all kinds of platforms ranging from web farms to even wireless telephony. Viruses and worms such as Sasser, Nimda, Code Red, Blaster, and Slammer [225, 154] have hit thousands of computers resulting, at the very least, in millions of people losing a considerable number of effective work-hours. When a popular website is compromised by hackers or when a denial of service attack is underway, the media will not fail to report the incident [111, 84]. In cases like these, the reputation of entire institutions may be hurt and the consequences may cause a considerable loss of market share.

A system that has not been found to be vulnerable today could be brought down tomorrow. Software must be permanently maintained with patches and updates, which constitutes the most effective procedure for systematically enhancing the security of a computer system [60, 180]. This, however, does not suffice to counter the effects of new and existing security threats [2].

1.1.2 The critical need for vigilance

Software is developed to meet design objectives. Even though a program may cover all these objectives in an actual implementation, it may also reach unexpected states at execution time that escaped the mind of its designers [207, 60]. The conditions necessary for going into one of these states may seem so unlikely that systems analysts ignore them without explicitly indicating what the program must do in those special cases. If a user carefully modifies the code, execution environment, input, or configuration of a program, the conditions for running into one of those states may suddenly appear and the behavior of the system can become unpredictable and/or incorrect.

Undesirable results may reduce the usefulness of the program as well as the trust users have with the system as a whole [18]. Some of those unexpected states may actually allow users to access information they were not supposed to, or to execute code without being authorized. The security policy will be violated by these states compromising the confidentiality of data, the integrity of its resources, or the availability of its services.

Secure coding practices can be used to build more robust software, however, the possibility of human error is always present. Fundamental software such as operating systems, network protocols, and database management systems are complex creation of the human kind. This complexity extends to the applications that interact with them which represents a challenge to the construction of programs that meet their design criteria. As the number of components of a computer systems increases, programs grow larger requiring more lines of code to develop both basic and sophisticated tasks as well as to keep backward compatibility. Unfortunately, the likelihood of their containing programming errors increases as well [207, 163].

Sophos Ex. 1009
Page 151 of 318

The existence of an exposure window between the discovery of a new **vulnerability** and the generation of a patch that fixes it, gives attackers the opportunity to infiltrate by taking advantage of security holes before they are removed [10, 163].

DEFINITION 4
Software vulnerability. A coding bug, configuration error, or design flaw that can be exploited to compromise the security of a system.

Additionally, system misconfiguration by unexperienced administrators is a problem that increases every year as the speed of growth of the Web outperforms the speed at which security education and training is being deployed [169]. Technology is developed to solve these issues but the effort seems not to suffice either. Even when firewalls, strong cryptography and other tools are available to protect the integrity and confidentiality of data, users keep sharing their passwords, insiders abuse their privileges, and backups are considered a cumbersome task by many system administrators [124, 126].

The imperfection of software and configuration management makes necessary the use of mechanisms that protect systems when their security fails. Before, during, and after the exposure window, systems must be guaranteed a protection level that limits the damage someone could cause by breaking into them. A number of solutions have been developed to help administrators deal with this issue. From among all solution strategies, system monitoring, also known as **vigilance**, has proved effective at controlling the impact of security incidents [4, 18].

Monitoring tools such as firewalls, anti-virus software, and **intrusion detection systems** scrutinize the activity of a system and constitute a protection layer that aims at identifying security breaches that escape the operating system and applications.

DEFINITION 5
Intrusion. Any violation attempt or actual violation to the security policy committed by an internal or external user [126].

DEFINITION 6
Intrusion detection. The set of techniques used to identify intrusions on computer systems. Intrusion detection is used to find attacks that cannot be prevented [60].

Assuming that code bugs, misconfiguration, privilege misuse, and lack of security training will continue to exist, it is critical to develop security vigilance solutions that complement good software engineering in order to protect the millions of computer systems connected to the Internet and used daily around the globe.

1.2 Limitations of intrusion detection systems

Intrusion detection technologies have been developed since the late 80's [49]. Many techniques have been implemented into commercial products and research prototypes following different

Sophos Ex. 1009
Page 152 of 318

approaches. The results obtained from them have exposed clear opportunity areas that need to be considered in order to develop more effective and secure intrusion detection tools.

Among the most important challenges in the area, we can name the following four which are addressed in this dissertation:

- Low detection accuracy reflected by high rates of false positives and false negatives that limit the usefulness of this type of security software.
- Immense amounts of data that need to be inspected in order to detect signs of intrusive activity.
- High overhead generated by the amount of resources that need to be allocated by intrusion detection tools in order to achieve real-time operation.
- Reduced portability and lack of standards that impede intrusion detection systems to share information (e.g., attack signatures) and to be integrated with other technologies.
- Limited connection between intrusion detection and automated response.

Most intrusion detection systems suffer from these limitations [126, 2]. One of the most recent models proposed for solving the problem of intrusion detection has been the use of agent-based systems. These systems take advantage of the distributed nature of the problem and leverage some of the basic and most attractive features of software agents. These systems also have some common design problems that need to be taken into consideration should we want to develop robust agent-based information security solutions:

- Frequent use of hierarchical structures for the arrangement of agents which results in potential bottle-necks and single-point-of-failure flaws.
- Limited autonomy of the agents composing this type of systems.
- Negligible use of standard communication and knowledge management mechanisms that may allow intrusion detection systems to integrate with other technologies and to facilitate their own maintenance.
- Lack of prompt generation and deployment of attack signatures that help quickly identify cyber-attacks.
- Limited security of the intrusion detection systems themselves.

1.3 Solution approach

In this dissertation, the goal of intrusion detection is complemented with an additional monitoring layer that concentrates on aspects rarely considered by existing intrusion detection systems: the inspection of outbound network traffic and the identification of attack attempts originated locally to compromise other systems in the outside.

Sophos Ex. 1009
Page 153 of 318

This approach is studied as a way of limiting the extension of the damage an intruder may cause and constitutes the foundation of a type of security tools whose objective is to protect neighboring systems by blocking intrusion attempts launched locally toward them and by alerting them about the fact that they are the target of an attack. Therefore, for the sake of its own security, a system could be interested in other hosts having such a monitoring system in place.

In particular, instead of modeling the intruder's behavior (i.e., human behavior) as the way of identifying abnormal activity, our solution approach focuses on the modeling of surveillance and attack tools (i.e., program behavior.) Attack software constitutes a repository of security information readily available on the Internet which has remained greatly unexplored and that could help improve the development of security tools. By detecting a rogue program in execution we expect to be able to protect other systems from being victims of a break-in. Further incidents could be prevented by countering the attack at the root or at an intermediate point instead of just assuming the defensive role of the target.

Given the nature of the problem, agent teams are used to implement this approach to intrusion detection. An evidence correlation mechanism is utilized to increase detection accuracy at a level that enables the integration of the detection task with automated response to security incidents.

The perspective of this dissertation regarding intrusion detection allows for the definition of a broader security context, poses new research questions, and provides an alternate way of approaching the security monitoring problem from a cooperative, less selfish stand-point.

1.4 Thesis statement

This dissertation has been developed to validate the following main hypotheses:

1. Preventing compromise attempts launched locally toward external targets is an effective way of limiting the effects of a security threat.

2. The data bank comprised of all the remote attack tools available on the Internet constitutes a valuable source of security information that can be used to develop better protection mechanisms.

3. The use of more than one software agent for the modeling of a rogue program produces more accurate detection results.

4. The generation and deployment of remote-attack program signatures or *fingerprints* can be automated.

5. A multiagent intrusion detection architecture supported by an ontological model requires less maintenance than non-agent software.

6. The use of virtual machines in the construction of trusted sub-environments helps agents more easily validate the integrity of the receiving (potentially malicious) host.

Sophos Ex. 1009
Page 154 of 318

1.5 Document organization

This document is organized as follows. Chapter 1 has presented an overview of the topics central to this dissertation as well as basic terms, the thesis statement, and the specific security aspects to address. Chapter 2 reviews related projects and background subjects in depth. It describes detection techniques, outlines the multiagent paradigm, and presents recent trends in computer security that motivate some of the main ideas of this work. Chapter 3 introduces the concept of outbound intrusion detection emphasizing its differences with the traditional intrusion detection approach. Chapter 4 is a description of the proposed architecture. It contains the definition of the ontology that supports the model as well as the organization of its components and their functions. Chapter 5 explains the implementation of the architecture, justifies some important design ideas, and describes the algorithms used to perform signature generation and matching. Chapter 6 describes the experimental design, presents results, and discusses findings from the execution of the proof-of-concept prototype under different scenarios. Chapter 7 lists the conclusions of the dissertation, summarizes its contributions, and identifies possible directions for future research. All cited references and the appendices appear at the end of the document.

Sophos Ex. 1009
Page 155 of 318

Chapter 3

Problem description

One of the primary topics of this dissertation is a computer security problem derived from intrusion detection for which a multiagent architecture is being proposed. This chapter identifies and defines outbound intrusion detection. It discusses its main features and compares it to intrusion detection techniques making special emphasis on the key advantages this approach has regarding security monitoring. The proposed architecture is described in the next chapter.

3.1 Outbound intrusion detection

At an airport it is critical to verify the identity of those who are going to access the gate area. It should be equally important to make sure no one on board a plane can take control of it with the intention of attacking public buildings or other facilities. Similarly, intrusion detection programs must guarantee that only legitimate traffic and users are let into a system, but it is also necessary to have security measures that prevent users from utilizing the system to compromise others – as typically occurs with DDoS attacks, e.g. [126].

Outbound intrusion detection is a variant of the intrusion detection problem. Its objectives vary from those of traditional intrusion systems but the two approaches share some fundamental aspects such as their monitoring nature and the type of security threats they handle. This section describes the outbound intrusion detection model.

3.1.1 Objective

Outbound intrusion detection (OID) focuses, not on protecting local resources from being compromised as intrusion detection techniques typically do, but on preventing those resources from being used to compromise its neighbors or any other systems connected to the Internet. OID aims at containing the impact of an incident by monitoring outbound activity at the host level so that the system cannot be utilized as an attack launcher or intrusion relayer. Figure 3.1 shows the data collection priority of three different intrusion detection models.

The function of OID is to identify system activity that contains attack signs and which is to be sent through a network channel. The major benefit of this approach is that, unlike incoming traffic whose origin and integrity is difficult to verify, a host is able to tell how outgoing traffic is being generated given the visibility it has of itself. It has access to detailed information such

Sophos Ex. 1009
Page 157 of 318

as the process that generated the traffic, the program attached to that process, the user who started up the process, and even the system calls that were requested to the kernel for execution.

DEFINITION 15

Outbound intrusion detection (OID). Set of techniques used to identify and report attempts to compromise remote systems using network resources and being originated at a local host.

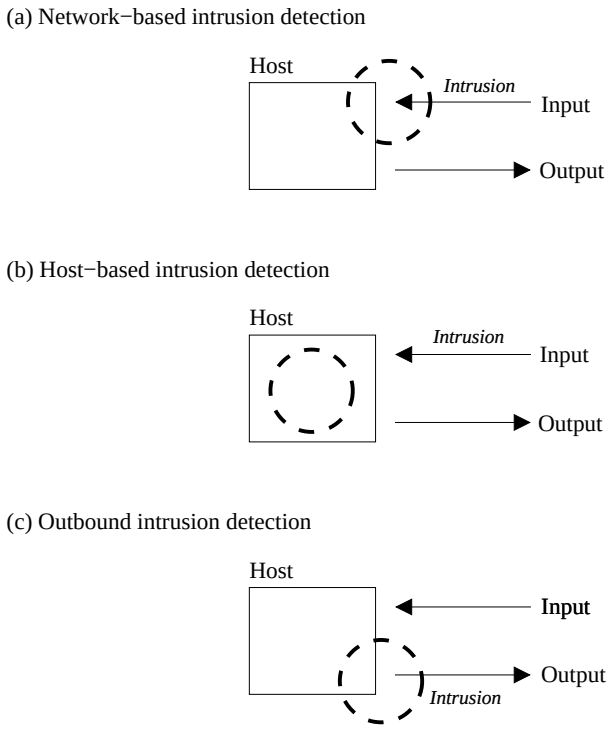


Figure 3.1: Data collection focus of different intrusion detection approaches.

3.1.2 Main characteristics

The OID approach has some distinctive characteristics that are worth exploring from a research perspective. The design of information security systems has been traditionally centered on the protection of critical components and their perimeter [156]. OID allows for the development of alternate security models by moving the spotlight from the target onto the source of an incident. Attack of neighbors is a sign of compromise that can be used to create a security context if the attacks are analyzed at the originating system. The following are some of the most significant features of this method which promotes a less selfish protection mechanism.

Sophos Ex. 1009
Page 158 of 318

1. **A less selfish approach to intrusion detection**. Intrusion detection systems protect resources local to a computer network by monitoring what happens inside the network and by inspecting the traffic addressed to internal hosts. These systems are not concerned about the security of other systems in the outside, and outbound traffic is rarely analyzed. The OID approach aims at developing a net of collective monitoring. The security of a system H is protected by others that make sure no attack is launched from them toward H. H does the same by making sure that no attack is launched locally toward external systems. Although the ultimate purpose is achieving H's own security, each system is not actively pursing its own protection. This creates a monitoring net that provides security to its members that, in turn, would be interested in other systems to become part of the net.

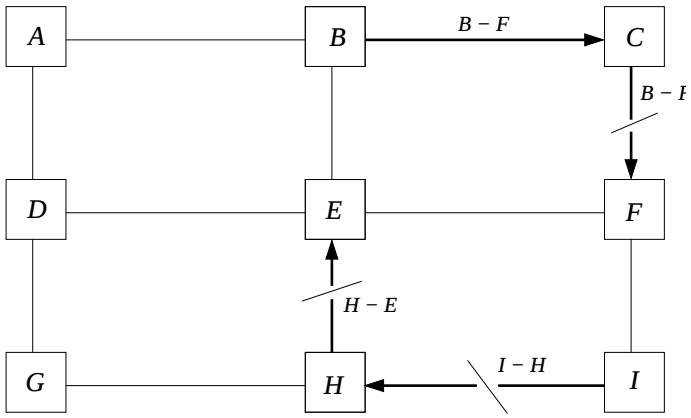


Figure 3.2: Attacks are detected at the originating host or at an intermediate host and can be blocked. This creates a security structure centered on the protection of others.

In Figure 3.5 an attack aimed at H could be detected and blocked by the source system (provided it supports OID). Similarly, if H is being used to compromise E, H itself will detect the incident, block it, and inform E about it. If host B does not support OID and is being used to compromise F through C, the OID system at C could prevent this from happening.

2. **Higher-quality evidence and detection**. "An intrusion detection system can only be as good as the data on which it bases its decisions" [223]. The change in perspective proposed by OID provides increased quality of data.

   A network-based intrusion detection system receives traffic which may or may not come from legitimate hosts or applications – this is left to the security system to decide. The design of the most popular network protocols we use did not take into account some basic security considerations and, as a result, it is possible to modify the fields in a network packet before releasing it into the network. A user can inject packets after having changed, for instance, the source address, the source port number, its network flags, and the segment size in an attempt to evade the security of the remote system he is trying to compromise [161, 143].

Sophos Ex. 1009
Page 159 of 318

The victim system has no way of knowing how these packets were created. It needs to rely exclusively on the information contained in the packets it receives to make decisions, and there are no guarantees about the integrity of these packets – especially if they arrive from a public network like the Internet. An advantage of the OID approach in this respect is broader visibility: a host has total access and control over its own resources.

Some network attacks modify flags and IP addresses on packets as to circumvent the security policy of a firewall [161, 143]. Firewalls that have not been properly configured may be tricked into accepting outside packets they were "supposed" to reject. If the security policy indicates certain packet types are allowed inside the LAN, but they cannot go into the Internet, an attacker could send packets whose source IP address corresponds to an inner host [143]. If no firewall rule verifies that packets were truly generated by internal hosts, the security mechanisms will be bypassed.

A system implementing OID could monitor this sort of traffic by making sure that all outgoing packets have the right source address hereby protecting other systems from being compromised by users at the local host (unlike an outbound firewall that is in charge of checking network packets only, OID complements this kind of checking with the analysis of the process generating such packets in order to detect events more accurately).

The execution of stimulators of intrusion detection systems produces traffic that may contain attack signatures [136]. An intrusion detection tool will react to this input if its detection algorithm is based solely on network packets. OID tools will not react to stimulators. They count on context information in addition to the network packets that enable them to identify what kind of program is being executed. An OID system looks at the process originating the traffic. Traditional intrusion detection tools do not.

3. **Characterization of attack tools**. Signature generation is a task that requires a high level of expertise [126]. In order to extract the essence of an attack, or an attack type, it is necessary to be familiar with the operational details of network protocols, operating systems, applications, as well as the way systems react to certain inputs [142]. Capturing what constitutes irregular activity into a pattern or a signature that can be used by a detection algorithm is not easy. Signatures that try to identify an attack type that covers multiple individual attacks could be too broad and would generate false positives. If signatures are too tight, they will miss small variations of the same attack.

   Given the increasing number of automated attack tools available on the Web [169], it would be desirable to characterize their execution using anomaly-detection techniques like those implemented by other security systems to model the normal behavior of services and applications (Figure 3.3). With the quality and amount of information available locally to OID systems, programs could be characterized and then distinguished from other applications at runtime. Root-kits could be identified, not by the static footprint they leave on a file-system – where they are easily hidden [125]–, but at execution time when they are actually harmful.

   Attack programs that aim at compromising remote systems through the Internet could be modeled through the network traffic they create by taking advantage of the access the host has to the IP stack and kernel data structures of the source system.

Sophos Ex. 1009
Page 160 of 318

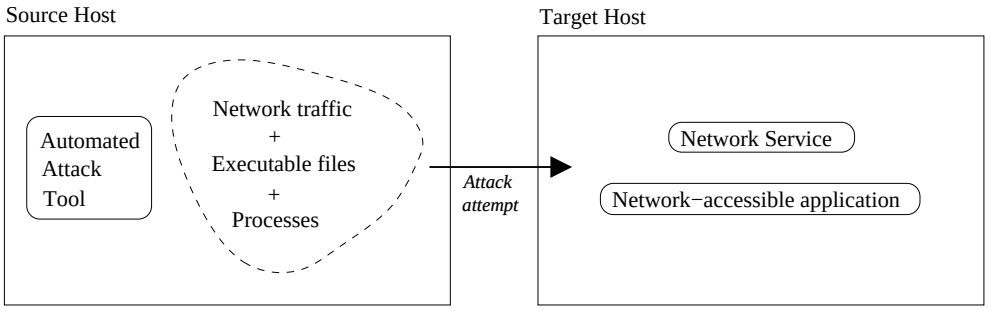


Figure 3.3: Characterization of automated attack tools through the use of locally-measurable and observable variables.

4. **More effective incident response**. Some intrusion detection systems have implemented response to incidents in the past but most of them leave this function to the administrator [126, 13]. If a system is under attack, there are several ways in which to react: firewall tables can be updated to close some ports or to filter out specific packet types; services can be temporarily disabled; passwords can be changed; security patches can be applied; the system can even be brought off-line for post-mortem, forensics work. All these security measures do not prevent the attacking system from targeting a new host should the current victim close all its entry points. This sort of reaction typically helps strengthen the security of the victim system [2] but, besides the effects of the event perceived by the target, there is no easy way to collect further information about the attack, or to react to it by disabling the actual processes responsible for launching the attack at the intruder's machine.

   The type of response an OID-based mechanism can implement is more definitive as access to the source of the incident is possible. There is no need to block local services or ports that could significantly damage the functionality of the system. OID systems are in the position to implement a more desirable type of reaction such as stopping user and system processes, restricting sessions, collecting evidence including the actual attack program and commands used (similar to the information a honey-pot system can capture), and sharing this evidence with the pretended victims so that they are able to heighten their security accordingly. If the system has been rooted, the results of an OID system are uncertain. OID therefore focuses on the majority of users executing attack tools, who are assumed not to have the expertise to modify the system in an attempt to hide their actions.

5. **Early detection of attack and surveillance tools**. Intrusion detection systems raise alarms as soon as they identify attack signs on the inspected evidence. Depending on the type of attack and detection method used, this could take a variable amount of time. An anomaly detector, for instance, could identify in seconds an event that deviates from the security profile of an application, whereas a stateful intrusion detection system could take several minutes to spot evidence of attack on a packet stream.

   By identifying attack tools through signatures, a system may be able to react as soon as the attack program is uploaded into memory, or when it puts the first packets onto the network. The development of response strategies based on an early detection capabilities may allow security systems to limit the extension and consequences of security breaches, not after the

Sophos Ex. 1009
Page 161 of 318

fact, but as soon as the incident arises.

An attack tool may be stored in an encrypted or compressed format (e.g., uuencode, zip, or ciphered with a cryptographic algorithm). This will prevent some security mechanisms that look into the file-system's contents for the presence of well-known attack programs from succeeding. However, attack tools need to adopt executable form in order for them to be run. It is then when they can be more easily identified.

The OID approach supports the idea of detecting automated tools which are used by users to compromise systems connected to the Internet and which account for a significant number of security incidents around the world [124, 169]. It is important to notice that if the OID program is disabled or put to sleep, it will be unable to succeed. That is why external monitoring of OID tools is necessary in order to guarantee they are working as expected.

6. **Improved packet processing without enabling promiscuous mode**. One of the challenges of network intrusion detection systems is the amount of data they need to inspect (see section 2.1.5). These tools operate in promiscuous mode and are generally placed at choke points where they can make the most out of their functionality by looking at a greater percentage of the traffic that goes into the network. However, the high traffic rates of intermediate or interconnection systems force them to drop packets in certain situations. In order to avoid this problem, OID is focused on protecting leaf nodes (e.g., servers and desktop computers) within a LAN by exclusively processing outgoing traffic emanated from them.

   Programs also exist that can detect hosts connected to a segment that are operating in promiscuous mode [2]). They will alert attackers about the potential presence of an intrusion detection system which could make them try to evade detection. An application that does not set a network device in promiscuous mode is more stealthy and may increase the chances of finding illicit activity. Since the objective of OID is to protect systems by observing outgoing traffic only there is no need to enable promiscuous mode.

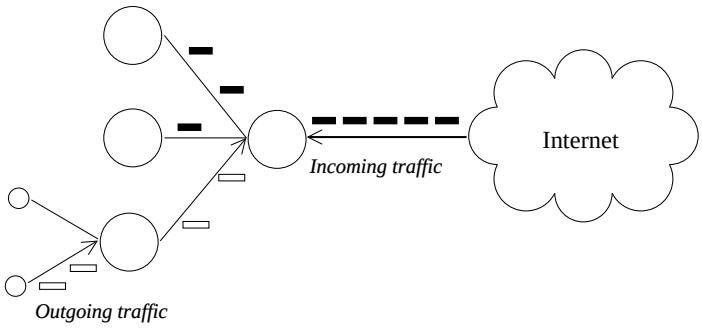


Figure 3.4: Under the assumption that the number of outgoing network packets an internal host generates is less than the incoming network packets a gateway sees, OID allows for more efficient packet processing.

Assuming that the rate of outgoing traffic from a switched-network host is less than the

Sophos Ex. 1009
Page 162 of 318

incoming traffic a gateway or a router receives, this approach allows for better packet handling by bounding the amount of data that needs to be analyzed (Figure 3.4).

OID does not contemplate traffic filtering capabilities as most firewalls and several network-based intrusion detection systems do. These systems identify irregular traffic in order to dynamically block packets preventing them to make it to their ultimate destination. Our approach does not hold a packet and looks into it to decide whether or not it can be allowed to get through. Packets are scrutinized after copying some of them (i.e., the first packets generated by an application) into an analysis buffer and so no packet is retained. This approach does not suffer from the need of having to evaluate the sanity of every single packet, which could constitute a bottleneck problem if the retention-and-checking algorithm is inappropriate for high traffic loads.

OID is a hybrid system that integrates network-based and host-based data sources. It distributes the detection task among all the hosts of a LAN instead of centralizing it on a firewall or router. Every system is responsible for monitoring its own outbound traffic, which constitutes a more scalable processing strategy.

7. **No need to keep stateful transfers for the reassembly of fragmented packets**. Many network-based intrusion detection systems work with complete packets only and do not provide reassembly support for fragmented packets [90, 143]. Since network-based intrusion detection systems need to evaluate packets as quickly as possible as to avoid flow congestion, some of them do not keep track of the fragmentation state of packet streams due to the overhead generated by having to delay a packet's evaluation until all its fragments are put together so that the detection algorithm can be applied. Evasion exploits based on fragmentation split a packet into pieces that individually bypass firewall controls and other filtering mechanisms hence achieving undesirable results on the target [143].

   OID handles packet reconstruction differently. Since it has access to complete packets at the sending location via the IP stack, packets can be analyzed before reaching the fragmentation phase – in case they have a fragmentation flag turned on. The reassembly overhead does not appear when capturing packets at this point of the transmission. This significantly expands detection coverage to include local attack attempts that may circumvent the security mechanisms of a remote system by dividing the payload of a packet through the use fragmentation. An attack could generate raw packets to evade OID but, since the OID tool is not only looking at network traffic but also at the actual processes in execution, detection is still possible.

8. **Possibility of detecting attacks over encrypted channels**. A major disadvantage of intrusion detection based on the analysis of network traffic is encryption. An encrypted data transfer will render a detection system ineffective if it captures packets and applies matching algorithms to the payload portion. Cryptography effectively scrambles data to protect its confidentiality and integrity [178]. Frequency analysis and many other cryptanalysis techniques are not effective with strong cryptographic algorithms which could allow an attacker to hide information detectable by an intrusion detection system. The use of encryption in virtual private networks (VPN's), network protocols, email applications, as well as terminal

Sophos Ex. 1009
Page 163 of 318

emulators and network-connection software, generates packets that cannot be handled by intrusion detection programs.

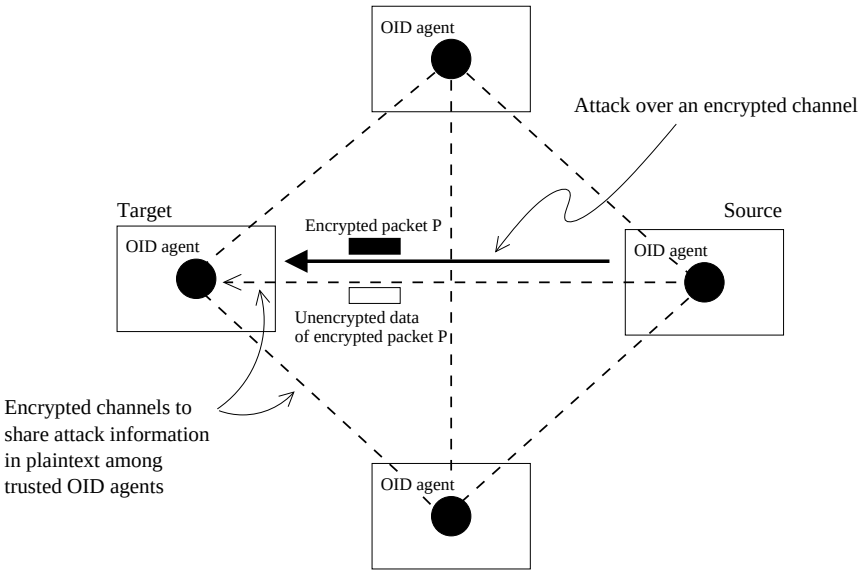


Figure 3.5: A mesh of trusted OID agents that share plaintext of encrypted attack data over secure channels

With OID it is potentially possible to look at certain data before it is actually encrypted. For instance, it can be determined whether an encrypted session contains commands crafted to defeat a target system. That is because the commands can be looked into in the clear by capturing the command history of the user. A mesh of trusted, certified intrusion detection agents could be created to alert systems about the presence of attackers like these who are using encrypted sessions (Figure 3.5; if the information is encrypted at the application level, this is not possible). A network-based intrusion detection system installed at the target machine is blind to this kind of traffic, but can receive information from an OID system at the attacking location which has access to the execution details of the attacker's processes, including plaintext.

9. **Improved chances to detect stealthy attacks expanding over long periods of time**. "Intrusions require the accumulation of substantial state information for their identification" [2]. An effective way to escape detection by some intrusion detection systems is to extend the attack period by inserting delays between subsequent actions. These delays will give time for the system to receive other network traffic and, given the cost of maintaining an intrusion's state, the detection system may drop a pending analysis line if the next set of actions does not arrive before a certain threshold [126]. In an attack modeled by a state-machine or a Petri net, for instance, the initial transitions leading toward a compromise state may quickly appear, but the detection process might not complete if future actions take sufficient time to appear as to exceed the intrusion detection's threshold.

Sophos Ex. 1009
Page 164 of 318

If this sort of attack is embedded into an attack tool, an OID system at the source location could identify it at execution time. System-call traces, function-call dependencies, as well as resource usage are distinctive features of a program that can be used to identify them. This can help defeat attacks that expand its activity over several minutes or hours.

10. **Performance improvements over time by signing safe programs**. Intrusion detection systems continuously monitor network and host activity. They either compare the data they observe against normality profiles, or against attack signatures.

    Similarly, an OID system can check the execution of programs against a data structure containing attack-program signatures. Once a program has been verified, it can be signed as to indicate it is safe to execute. It will not be necessary to verify it again unless new signatures are added or unless its integrity has been compromised (remember that a program is deemed as dangerous only if it matches a network and/or a process signature; programs for which no signature exists will go unnoticed; the halting problem prevents a security system from determining whether any given program will produce risky results). In the long run, this will benefit the performance of the OID security mechanism which will have to run only when new programs or when new program signatures appear (Figure 3.6).

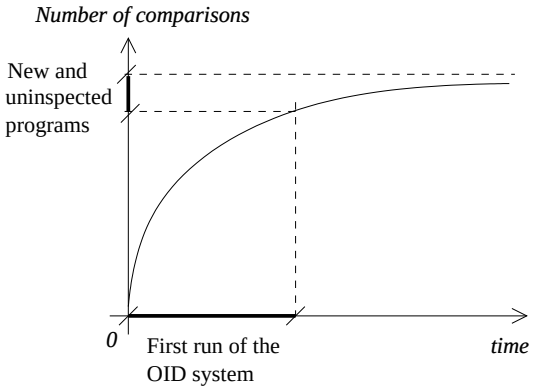


Figure 3.6: Performance of an OID system: *Number of comparisons* **vs.** *time* **(a comparison tries to match program files against program signatures).**

11. **Resilience to denial of service attempts**. If an OID system was protecting a host, a user might try to launch a denial of service attack in order to bring it down locally. Having an OID access to the source of such attack (it executes on the same machine as the attack program) it is able, not only to realize it is under attack, but also to react to the incident by suspending sessions or any other measures to make the threat disappear. This makes of this approach a more robust security monitoring strategy.

    Additionally, distributed denial of service attacks are usually executed using remotely managed tools. At one of the compromised servers on which an instance of the attack tool has been downloaded, an OID program could be executing. If a signature exist for the tool,

Sophos Ex. 1009
Page 165 of 318

the security system will detect it and will be able to alert the remote target(s) about the existence of a latent threat.

Sophos Ex. 1009
Page 166 of 318

Chapter 4

A multiagent outbound intrusion detection architecture

This chapter describes a multiagent architecture proposed for the problem of outbound intrusion detection. It discusses main tasks and agent classes, the characteristics of the environment, introduces the underlying ontology that allows the interaction of agents and describes the functioning of the architecture.

4.1 A multiagent approach to intrusion detection

Given the very specific tasks that need to be developed by an OID system, the information exchange necessary to integrate these tasks, and the intelligent functions involved in intrusion detection, the multiagent approach appears as a good option.

Intrusion detection is a process requiring human skills that researchers have tried to build into computer programs [65]. An administrator needs to analyze different sources of dynamic data (e.g., kernel data structures, audit log files, and network traffic) which may contain a few, small pieces representing security information. In order to do that, he will need to correlate events very precisely based on attack patterns he knows, or by identifying and interpreting irregular signs during the operation of the system [67]. If any security incident has occurred in the past, the administrator will use that previous experience to handle the data and detect attack attempts more accurately.

This kind of monitoring not aided with intrusion detection tools could be performed in real time for a single host or a small network. Monitoring a WAN or a large corporate LAN would definitely be a tough task for a single administrator or a small team of security analysts given the dimensions, dynamics, and heterogeneity of computing platforms used nowadays [101]. Training a security team or hiring remote administration is an expensive solution but, even when security issues can arise in any kind and size of organization, the costs can prevent the management team from investing in the protection of its information assets. Intrusion detection tools must be used to automate vigilance, reduce costs, and increase effectiveness. The development of intelligent software which helps simplify systems administration is necessary.

Agent-based architectures are an option to implement the functions needed to perform in-

Sophos Ex. 1009
Page 167 of 318

trusion detection [224, 15, 223]. These functions can be assigned to specific agents that will be responsible for performing certain tasks autonomously. In order to enable inference processes, meaningful communication, and smooth integration with other security mechanisms, it is necessary for agents to have a common understanding of the different objects and relations present in their environment. This can be achieved with ontology support as mentioned later in this chapter.

4.2 Specific objective of the architecture

The two main approaches to intrusion detection are to model what is normal, and to model what is abnormal. The creation of normality profiles that describe licit and regular operations is not an easy goal to accomplish [108]. The number of kernel components, devices, and applications that compose a system, along with the execution states, variables, events, and messages it handles constitute a complex model. A system like this can be described, but the task of characterizing it through the definition of a sense of "system self" which captures stability and safety conditions has proved challenging [83, 64].

A simpler problem is to characterize what constitutes an abnormality – assuming that the majority of the events within a system are normal [64]. In the case of OID, characterizing outbound activity, such as network traffic, poses similar problems to the ones that appear when modeling the normal functioning of a system.

"Regular traffic is not well-behaved" [143]. Due to the number and diversity of the systems that can connect to the Internet as well as the number of protocols they utilize, it is not trivial to generalize what normal traffic is to a router or a firewall. Creating behavior profiles of the outgoing traffic of a host might seem a simpler issue, however, the execution conditions of internal applications as well as the service requests that need to be processed, generate packet streams for which it is be difficult to define a single, precise normality profile [143].

Our particular approach to OID will be to define signatures of illicit activity. We leverage the quality of information available to the host that describes its internal structures and execution state in order to precisely detect patterns that may indicate someone is trying to compromise the system by launching attacks from a local host.

Given the trends observed in information security over the last years (Section 2.2), we will concentrate on the more specific problem of *characterizing remote attack tools that generate IP packets detecting them at runtime and reacting to these incidents in an automated way.* IP is chosen because it is the most widely used protocol on the Internet and it is also frequently used to perpetrate cyber-attacks [142, 161]. The detection mechanism will be connected to response methods that will counter compromise attempts locally, will alert neighbor systems, and will trace the intruder back in case he is connecting through an interactive shell session to the local server [118].

4.3 Characterization of automated attack tools

OID's objective is unselfish: to achieve the system's own security through the protection of others [119]. If the system watches for the security of its neighbors, and these neighbors follow

Sophos Ex. 1009
Page 168 of 318

the same approach, a higher, cooperative security level will eventually appear.

One way of protecting other systems is to identify users who may be trying to bring them down. A common way to analyze and attack network servers and workstations is through the use of automated attack tools [126]. Users download these programs from the Internet with the intention of compromising all sorts of systems – although sometimes they just want to play a prank on someone's workstation.

In order to identify these tools, one could just look for the static file image of the program as integrity checkers and root-kit detectors normally do [100]. Unfortunately, compression and encryption make this approach unfeasible. The program, however, will have to adopt executable form so that it can be run. It is then when attack tools can be detected at the hosting system and administrators can react to the incident more definitely. They can also warn the attempted victim about the possibility of being the target of a wider, orchestrated attack (e.g., distributed denial of service).

Many aspects may be considered for the characterization of programs [184]. We consider two groups of them: (1) the peculiarities of the network traffic generated by an attack tool, and (2) the resources accessed by its process(es) when in execution (note that the OID architecture presented in this document concentrates on remote attack tools only). The former is an approach widely used by network-based intrusion detection tools. The difference is that an OID system has access to details such as how packets are being actually generated (including how the source IP address and/or port is being spoofed) and this information is not accessible to traditional intrusion detection systems which are located on the victim's side.

Analyzing the resources accessed by an attack tool also helps identify its execution. This is because features like the data and text segments uploaded into RAM, the libraries and other files open by the program, as well as the sequence of system calls requested to the kernel are, as a whole, distinctive enough to identify a program from the rest [184]. Similarly, this type of attack software frequently crafts network packets with very distinctive characteristics that can be easily identified as anomalous or intentionally-wrong. If packets that belong to the same exchange are captured and put together for analysis, it may be possible to identify the attack program that originated them. This is one of the hypothesis of this project.

4.3.1 Attack program signatures

In order to capture execution and communication data to model a program, two signature types are considered. They are used independently by detection agents which in turn supply their matching results to correlators that generate a composed diagnosis. Each signature should contain attributes that help uniquely identify an attack program.

1. **Network traffic signature**. An OID traffic signature must contain information extracted from the sequence of network packets generated by an automated attack tool. The TCP/IP traffic signature presented here considers the first k packets of an exchange (outbound packets only). They are captured and condensed in order to build a single signature that describes them all. Some attributes are more distinctive than others for purposes of packet stream analysis [143]. The proposed structure is as follows:

```
{Source Port, Source Address, Destination Port, Destination Address,
 Fragmentation, Window Size, {Flags 1, Flags 2, ..., Flags_ k}, {Hash 1,
```

Sophos Ex. 1009
Page 169 of 318

```
Hash 2, ..., Hash k}}
```

`Source Port` data is usually modified in evasion attacks in order to circumvent filtering rules. Clients typically use non-privileged ports to contact a server listening on privileged ports; although some servers contact other servers using privileged ports, as in the case of DNS's port 53, seeing packets with a privileged port number as source port going out of a workstation may constitute a sign of anomalous activity.

The `Source Address` field may also suffer from the same type of problems as source port numbers do. An attacker may be inserting addresses of the internal hosts of a network in order to bypass a misconfigured bastion firewall that accepts incoming packets from the outside as long as they have a local source IP address.

`Destination Port` and `Destination Address` can be both used to characterize tools that aim at a particular target or service.

`Fragmentation` flags may indicate the attacker is forcing packet fragmentation in order to avoid network-based intrusion detection systems that look for suspicious strings in the payload but which do not support the expensive task of packet reassembly. (Notice that this "intention" information is rarely accessed by intrusion detection systems which only see packets but have no way to tell how they are being generated).

`Window Size` is an attribute used in surveillance operations to fingerprint an operating system by seeing how it reacts to malformed packets (e.g., those with a '0' window size; different protocol implementations react differently to exception and error conditions and this can be used to identify the type or version of the services a host is running).

TCP `Flags` are one of the most evident signs of furtive and malicious activity. They can be altered in order to influence the synchronization of an exchange (e.g., to do session hijacking [143]) or to launch denial of service attacks (e.g., a SYN-flood attempt).

The above fields can be used to distinguish abnormal sessions in an anomaly-based fashion. An additional set of attributes that can be used to more accurately identify the originating program is packet payloads. Many attack tools use root-shell code[1] which represents a very distinctive characteristic of a program and its programmer [211].

Other fields such as packet checksums, time-to-live values, and sequence numbers [142] (Appendix B provides a description of the TCP and IP packet formats) can be also considered but we decided to include some of the most descritive ones only.

2. **Process signature**. Profiling a program in execution involves aspects which are operating-system dependent or that depend on the peculiarities of the executing environment (e.g., configuration and set-up conditions caused by the software being executed [184]). Some of the attributes that a process signature may contain are captured in the following signature structure:

```
{Program Name, Execution Parameters, Text Segment Checksum,
 Number of Files Open, {File Open 1, File Open 2, ..., File Open k},
 Number of Dynamic Libraries Open, {DLibrary 1, DLibrary 2, ..., DLibrary k},
 Memory Allocated at Upload, Hash of the File Image}
```

[1] Root-shell. A sequence of commands precompiled for a particular architecture and stored as a character array on an attack program that will be inserted into memory during a buffer-overflow attempt

Sophos Ex. 1009
Page 170 of 318

With a different weight, these attributes will build a platform-dependent signature that describes part of the behavior of attack tools at runtime.

Unlike the overall structure of a program, `Program Name` is an attribute easily modifiable. However, some attackers may download and execute the tool as-is and file names would help identify the type program being used.

`Execution Parameters` are particular to each program. Sometimes the options passed in the command line need to follow a fixed order, and options are typically limited and very specific.

The `Text Segment` contains code that will be executed by the processor. It is dependent on the source code and the compiler used and reflects the way the program has been structured. A representation of the segment (for a determined architecture) can be used. (In the case of scripts, this attribute is not being considered as it is not straightforward to extract the code in executable format from a file image. The actual code of the script is used instead).

The `Number of Files Open` is a characteristic of a program that can be used to complement the other attributes without having to do process tracing to monitor system calls. The specific names of these files (`File Open` i) as well as the order in which they are open can also contribute to the description of programs in execution.

Many programs link `Dynamic Libraries` made available by the operating system. The number and name of these libraries, as well as the order in which they are accessed is information that reveals the internals of the monitored program.

The amount of RAM `Memory Allocated at Upload` time, including text and data segments, can also be captured as a descriptive attribute.

The `File Image` stored on hard disk is a feature used by integrity checkers in order to detect illicit modifications. Compilation, compression, and encryption can render this attribute invalid, however, using the right weights, it can be helpful in the detection of tools that have been downloaded and executed without being prior modification.

4.4 Environment

The environment OID agents inhabit is a network host. All observable and measurable information within the host constitutes a potential perception to the sensing agents of a multiagent architecture. Agents that implement incident response will in turn modify this environment based upon security-relevant events about which they have been warned. This environment can be described according to different *dimensions* that help understand it and determine the appropriate agent architecture to use [173] (Figure 4.1).

By the capability of an agent to acquire meaningful information, the environment is *fully observable*. The OID approach allows agents to access complete state of the environment. The information captured by the sensing tasks is enough to decide what to do at any given time.

By the sequence of events observed by an agent, this environment is non-deterministic or *stochastic*. An action performed by an agent cannot always have the same outcome. Random factors affect the result of an action such as a change in a program's state or an error condition. The outcome of an OID agent action can be any of a set of possible results whose probability is variable (a response agent, for instance, could have sent a termination signal to a process but a parent-child binding could have prevented this from happening as the agent was expecting). Furthermore, there are elements whose behavior cannot be known a-priori such as users and

Sophos Ex. 1009
Page 171 of 318

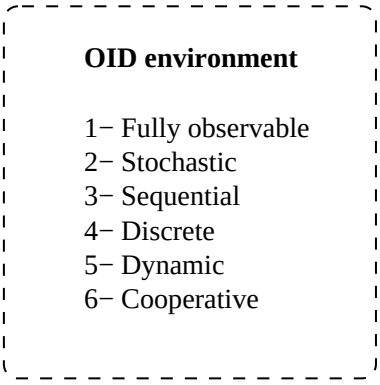


Figure 4.1: Features of the OID environment

the reception of network packets from clients. These elements have also a direct impact on the result of actions and the status of the system.

This environment is *sequential* as the choice of actions by an agent at a certain point in time depends not only on the current state but also on previous events that influenced the current state. The detection of intrusion actions is highly dependent upon events observed in the past which helps interpret the contents of messages at present. Additionally, an action performed by an agent can and will affect other actions and decisions at subsequent times. The state of an intrusion event needs to be followed in order to determine the point at which illegal actions are being performed so that evidence can be obtained.

The host environment is *discrete*. The number of percepts types and actions relevant to OID agents is limited. The possible execution states of an application form a finite set as well – error conditions can be many, but finite at the end–, and the host's timeline upon which events appear is discrete too.

The OID environment is *dynamic* since it can change while agents are making decisions. This dynamics is part of any multiuser or multitask computer environment which adds difficulty to the modeling of its components and their interaction, the inspection of data streams, and the reaction incidents.

The approach to OID we are taking defines a *cooperative multiagent environment* in which a number of diverse agents contribute to the goal of identifying attack programs in execution. All the agent subgoals are defined toward the achievement of that common objective and no competitive function is part of the approach.

4.5 Task definitions

The task of characterizing, detecting, and reacting to automated attack and surveillance programs constitutes a phase of misuse detection that integrates direct monitoring of host-based data sources.

From the literature on agent-based intrusion detection systems, there are three high-level tasks common to all of them [92, 41, 224]: (1) data collection, (2) data integration, (3) data correlation. In addition to these, signature deployment, evidence sharing, and other functions

Sophos Ex. 1009
Page 172 of 318

are also implemented.

In the case of OID, we can identify four high-level tasks very similar to the ones mentioned above which are composed of a number of more specific, atomic sub-tasks. The four major tasks are: (1) sensing, (2) signature generation, (3) signature detection, and (4) reaction or response.

1. **Sensing**. Intrusion detection systems model rogue processes and attacks by using different variables such as system-calls invoked by a process (sometimes including parameters) [83, 211], the set events generated by a process [108, 109], modifications made to kernel data structures [9], the functional structure of programs [186, 211], and the impact a program has on resources usage when it is executed [50, 5].

   Data collection, also known as *sensing* when discussing agent systems, is a basic function needed to obtain all that information. Data must be screened in order to optimize communications with subsequent phases. There are two subgoals of data collection for OID: (1) to observe an attack tool in execution in order to extract data to be used in signature generation, and (2) to monitor the system during regular operation in order to supply the detection task with analyzable data.

   The better the quality of the data (e.g., complete, compact, trustworthy, properly-formatted data), the faster the matching may be and the better the chances of providing a suitable security level to the resources being protected.

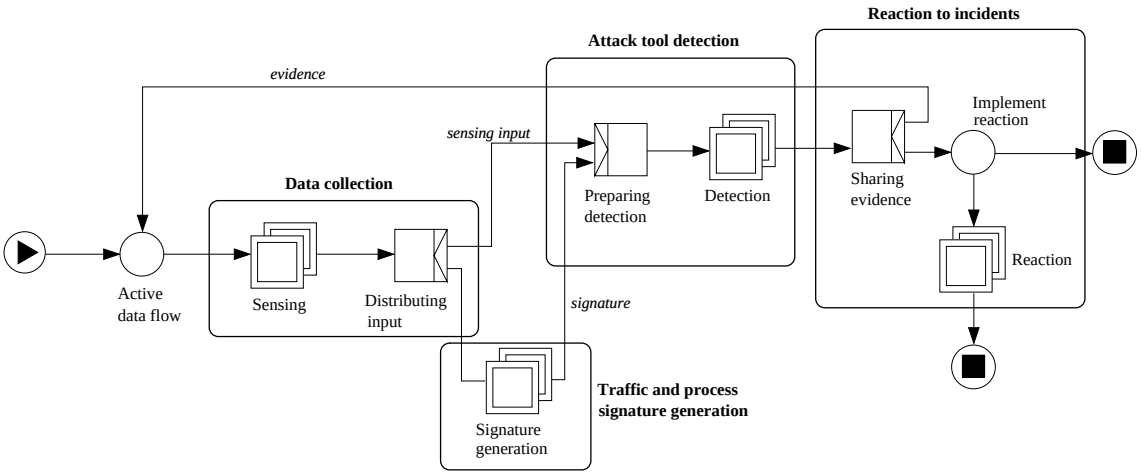


Figure 4.2: Primary OID tasks (YAWL diagram; see Appendix A).

2. **Signature generation**. The functions of signature generation are to integrate data emanated from the program to model, and to condense it into detectable patterns. The critical component of an automatic signature generation algorithm is to determine which aspects of a program need to be taken into account. These have to be distinctive enough for distinguishing a program from the rest and, at the same time, flexible enough to effectively detect small variations of the program.

Sophos Ex. 1009
Page 173 of 318

Signature design also needs to consider the suitability of the signature format being used [142, 43]. Signatures are supplied to the detection module and, if the format has not been properly designed, or if the signature contains redundant or irrelevant information, detection will be unnecessarily delayed.

The signature generation task is also in charge of accumulating all the signatures on a repository that will be accessed by the detection mechanism at execution time, and it will also update the repository as new, better, or incorrect signatures are reported.

3. **Signature detection**. The detection task is a data correlation function. Information received from the sensing modules needs to be integrated and interpreted in order to tell whether or not there is something anomalous in the system. For that purpose, the signature repository needs to be open and all pertinent signatures retrieved. Once system activity data has been properly structured and the signature bank uploaded into memory, the detection algorithm can be applied.

   Given that the OID focus of this architecture is on the detection of attack tools used to compromise remote systems through network channels, the event that will trigger the signature detection algorithm is the request to access a network device for initiating a data transfer (e.g., a *socket()* system-call requested to the kernel). When this sort of event is detected, it will be necessary to inspect the execution of the program in order to determine whether it matches or not any of the signatures.

4. **Response**. Automated response relies on accurate detection [126, 60]. There are two types of response: (1) passive and (2) active. Passive response to incidents includes session logging, increased monitoring of a service or process, and sending alert messages to the administrator. Active response has higher detection-accuracy requirements as it reacts to incidents by drastically modifying the environment with the intention of containing attack attempts.

   The consequences of inappropriate response to misdiagnosis can be severe especially if reconfiguration is being performed, processes are being suspended or killed, sessions are being terminated, files are being deleted, access is being denied, or if an emergency shutdown is scheduled [2].

   An OID system can effectively perform both types. Through active reaction it could, for example, suspend a hostile process in execution, and through passive actions it could alert neighbor systems about the existence of a security threat.

4.6 Task-to-agent mapping

The four high-level tasks of OID can be mapped onto specific agents that will implement, along with other support-function agents, the mechanisms to monitor the execution of programs. Tables 4.1 and 4.2 describe the three main agent classes that compose this architecture, namely, *Sensors, Correlators*, and *Reactors* (agent subclasses corresponding to each of these major classes are defined by the ontology; see Section 4.10).

Sophos Ex. 1009
Page 174 of 318

Class name	**Sensor (S agent)**
Purpose	To provide data collection capabilities for the multiagent architecture. This agent class is in charge of monitoring data sources and to extract meaningful information from them that serves the purpose of generating and detecting signatures.
Perceptions	(1) Data streams observable within the host (e.g., environment variables, network packets, system-call sequences) (2) Messages from other agents
Actions	(1) To listen for system activity that indicates signs of intrusion attempt (2) To screen data based on the relevance of the information it carries (3) To decide what information to send to peer agents (4) To send data to other agents

Class name	**Correlator (C agent)**
Purpose	(1) To integrate data from multiple sources in order to create program signatures (2) To integrate data from multiple sources in order to detect program signatures (3) To trigger the response mechanism by alerting reactor agents about intrusion events
Perceptions	Messages from other agents
Actions	(1) To create program signatures (2) To store signatures on the repository (3) To detect program signatures

Class name	**Reactor (R agent)**
Purpose	To implement passive and active response to security incidents as identified by correlators
Perceptions	(1) Messages from other agents (2) Status of system elements related to incident response (e.g., files opened by a process, process status) (3) Status of security mechanisms related to incident response (e.g., active firewall rules, file permission masks)
Actions	(1) To decide on the response type to implement according to the detected event (2) To develop active and/or passive response to incidents (3) To store evidence on the repositories

Table 4.1: Description of the Sensor, Correlator, and Reactor agent classes

Sophos Ex. 1009
Page 175 of 318

Class name	**Name and Ontology server (NOS agent)**
Purpose	(1) To keep track of the location of agents (2) To supply location information to agents (3) To update the ontology image (4) To supply the ontology definition to new agents (5) To report changes on the ontology to all agents
Perceptions	(1) Creation of new agents (2) Destruction of existing agents (3) Changes in the status or location of agents (4) Ontology update events (5) Information request messages
Actions	(1) To register agents (2) To remove agents (3) To update agents location and status (4) To supply location information (5) To update the ontology (6) To send ontology updates to all agents

Table 4.2: Description of the Name and Ontology Server agent class

4.7 Organization of single agents

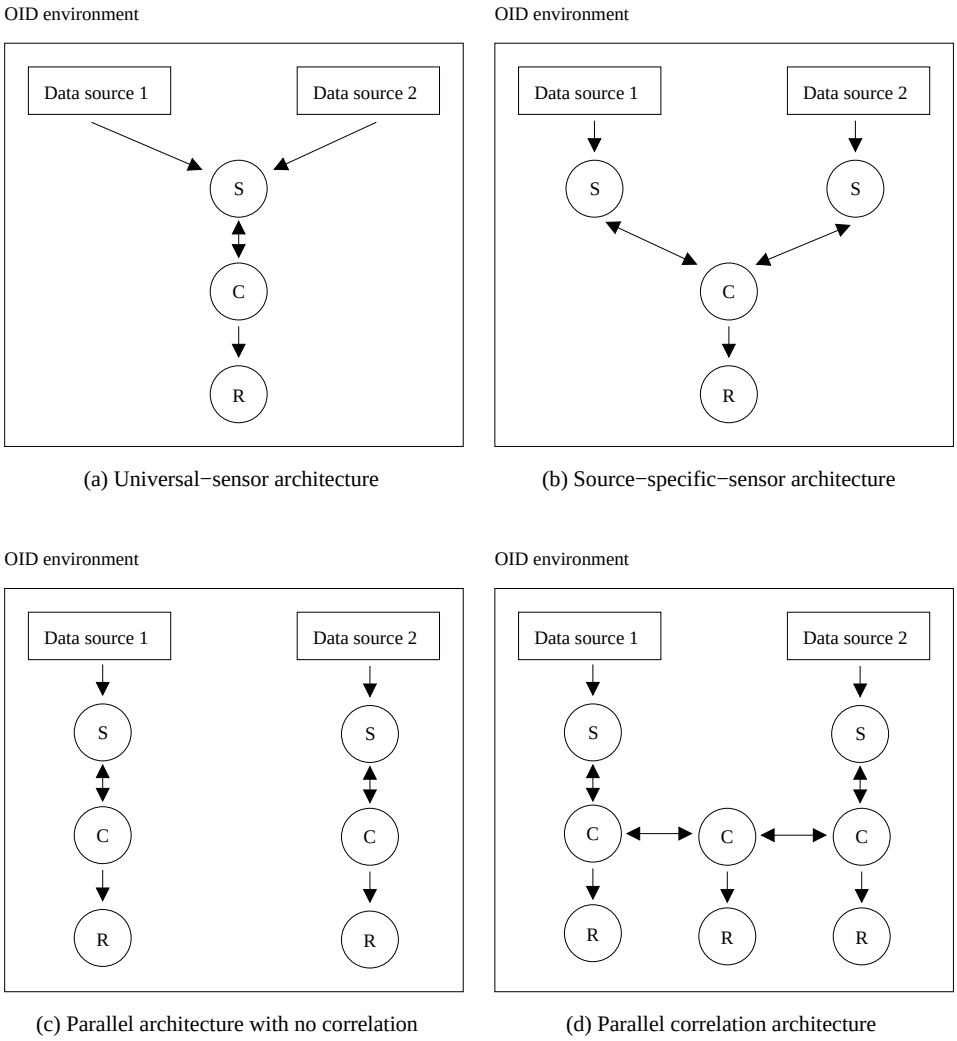
All the architectures described in the intrusion detection literature employ one or more processing modules – called agents by some implementations – arranged in many different ways. Seven architectures were analyzed in order to identify the most common inter-agent connections ([92, 224, 15, 25, 41, 78, 221, 158]). As mentioned earlier, these architectures have some functional commonalities that make them comparable.

Taking three basic agent classes described in the previous section, we describe four connection architectures and discuss their advantages and disadvantages (agents S, C, and R agents are used; since no architecture uses an ontology and only two architectures give notion of how agents are located by others [92, 221], the NOS name and ontology agent-type is not included). The R reactor agent in this analysis is responsible for any activity suggested by correlators as a result of data analysis.

Figure 4.3 depicts the four models. Both, peer-to-peer and tree structures were considered. We focus on the connection of the agents to the data sources and on how they exchange information.

1. **Universal-sensor architecture**. This is a simplistic architecture that could be considered a monolithic structure from the functionality perspective. It has a hierarchical structure containing a sensor, a correlator, and a reactor. The sensor is in charge of monitoring all the

Sophos Ex. 1009
Page 176 of 318



(a) Universal−sensor architecture

(b) Source−specific−sensor architecture

(c) Parallel architecture with no correlation

(d) Parallel correlation architecture

Figure 4.3: Agent interconnection architectures (S = sensor, C = correlator, R = reactor; arrows indicate data exchange flow).

available data sources. There is a unique correlator that generates and detects signatures, and one reactor that takes care of incident response.

– *Advantages*. It is a model simple to implement and the consumption of resources is kept to a minimum.

– *Disadvantages*. Integrating the functionality to access different data sources into a single agent might not be the best alternative. It may increase complexity unnecessarily as the agent should carry the code to handle all data sources available. This model limits portability and does not scale well to the addition of more data sources (basically, the agent needs to be changed). A major disadvantage of this hierarchical structure is that, if the sensor agent is not working properly or not working at all, the whole mechanism gets stuck.

2. **Source-specific-sensor architecture**. This architecture has independent, specialized sen-

Sophos Ex. 1009
Page 177 of 318

sors for each data source. The sensors send the data to a central correlation agent that triggers incident response by contacting a single reactor.

– *Advantages.* The task of monitoring data sources is split into several sensors. This allows for dynamically adding and removing sensors depending on the need to monitor more data sources at the same time. The correlation mechanism still receives evidence from different data sources which increases the likelihood of identifying illicit programs. If one of the sensors fails, other sensors are still able to provide the correlator with data and to detect intrusion attempts.

– *Disadvantages.* In case the correlator gets disabled or presents any functional problem, there is no other agent with the capability of detecting and generating signatures. The system, in that case, would fail and no reaction would be carried out.

3. **Parallel architecture with no correlation**. This model includes independent, parallel mechanisms for detecting intrusions on each data source..

   – *Advantages.*It is a more robust model in terms availability. If one or more of the agents in charge of monitoring a certain data source goes off, the system is still able to provide protection – including reaction to incidents – for one portion of the system.

   – *Disadvantages.* Intrusion detection systems that integrate multiple data sources have proved more effective at detecting attacks [1]. The main weakness of this architecture is that it does not take advantage of the availability of multiple data sources. This might increase the rate of false alarms as correlators are not working with as much information as if they were receiving data from all evidence sources.

4. **Parallel correlation architecture**. This architecture integrates the best features of the other architectures presented before. There are independent agents in charge of the monitoring of each data source and they supply, at the same time, evidence to a central correlator that fuses the data and has the potential to identify intrusion attempts more accurately.

   – *Advantages.* This system has a central correlator that receives and fuses information from different sensors in order to determine the best way to react to incidents. In case this correlator fails, the system is still able to provide full protection by detecting attack signatures at every data source. In the event a sensor fails, the rest of the mechanism is still operational and keeps providing security monitoring to the system. Since the sensing functionality is stored on independent sensors, this model allows for portability and scalability.

   – *Disadvantages.* A disadvantage of this architecture is the number of agents it requires. The first architecture (see (*a*) in Figure 4.3) used three agents whereas this architecture uses $3n + 2$ agents, with n being the number of data sources to monitor. This is computationally much more expensive if agents execute separately – e.g., agents are stand-alone processes; it is a different story for agents implemented as threads [110].

Out of the four models, the last one presents clear advantages regarding availability. The isolation provided by grouping agents by data source and then connecting these "stand-alone" groups through correlation elements increases the robustness of the structure while still providing evidence fusion.

Sophos Ex. 1009
Page 178 of 318

This agent groupings are exploited in the proposed OID architecture described next in order to develop a model more resilient to the malicious-host problem[2].

4.8 Agent cells

The proposed architecture is not composed only by agents that operate individually and co-operate with their peers [224, 41]. We group agents functionally into teams that cohabit sub-environments called **agent cells** (Figure 4.4).

DEFINITION 16

Agent cell. A functional software construct that groups several agents for the purpose of achieving a common goal.

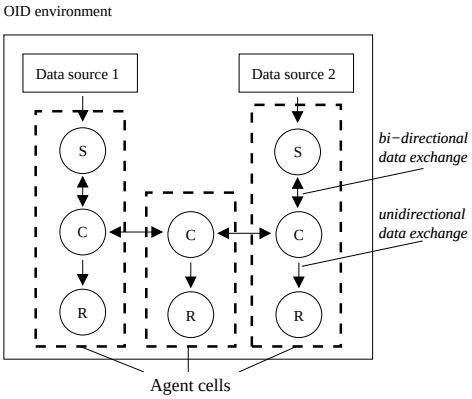


Figure 4.4: Grouping agents functionally into specific-purpose sub-environments or *agent cells***.**

An *agent cell* is a software environment which agents can trust. Unlike a host, whose overall hardware/software integrity can be difficult to asses, virtual machines and interpreters are software constructs whose operational safety is easier to verify mechanically by the agent themselves through digital signatures or other cryptographic means. An agent implemented as a process would execute on the operating system directly [15, 158], but a cell agent executes inside the cell which is the one in touch with the underlying operating system. The task of securing execution is left to the cell and not to each of the agents.

Agents that share a common goal are put into the sub-environment of the same cell. The tasks necessary to develop intrusion detection are actually implemented by the agents but cells integrate them functionally for security and communication purposes. Figure 4.5 shows the agent cells that compose the architecture.

[2]Malicious-host problem. The execution of agents on a rogue environment that gives no guarantees regarding execution correctness or agent integrity [153, 178]

Sophos Ex. 1009
Page 179 of 318

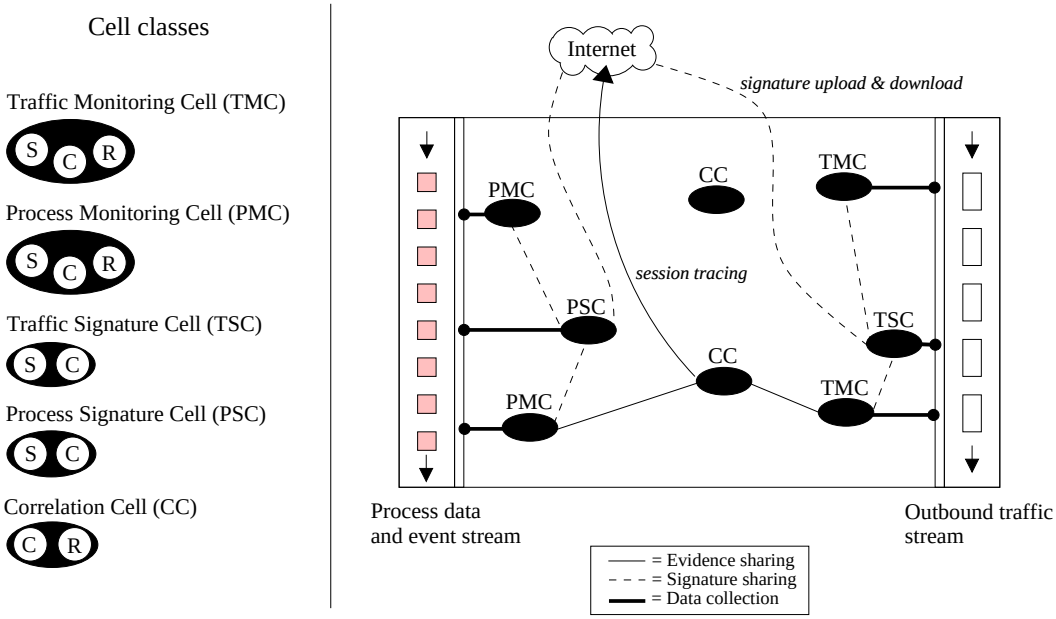


Figure 4.5: OID architecture and main cell classes.

The creation of cells allows agents to perform independent detection functions. A cell can be focused on network-level intrusion detection while another one works on process-level detection. Then, both can be integrated by means of correlation agents. Cells are self-contained in the sense that sensing, detection, and reaction capabilities are all available inside the cell. Agent cells are arranged non-hierarchically as a way to provide increased robustness to the OID system (Janakiraman et al. study peer-to-peer intrusion detection in [92]). If one of the cells is down, others will be available.

In Figure 4.5 for instance, if a TMC cell is down (see cell-type descriptions below), another TMC cell or a PMC could be up, and vice versa. Similarly, if a CC correlation cell is down, the rest of the monitoring cells will continue to provide protection to the platform even when the correlation functionality may be off.

4.8.1 Agent classes and cell classes

There are several agent and agent-cell classes composing the architecture. Agent types were introduced before and here we describe their sub-classes (for an overall description of the classes involved in the architecture, see the ontology in Figure 4.12 later in this chapter):

- *Sensors*. Sensors are responsible for collecting data. They are further classified into *Traffic Sensors* and *Process Sensors* depending on the data source they access.

- *Reactors*. Reactors implement reaction procedures to incidents. They are further classified into *Guards*, which execute local response, and *Tracers*, which connect with other agents located at remote hosts in order to back-trace suspicious sessions identified in a security incident [118].

Sophos Ex. 1009
Page 180 of 318

- *Correlators*. Correlators receive data from *Sensors* and inspect it through detection mechanisms. They are also responsible for triggering response from *Reactors*. Correlators are further classified into *Traffic Detectors*, *Process Detectors*, and *Hybrid Detectors* – hybrid detectors that fuse evidence from the other two detector types.

A number of instances of these classes are the ones performing signature generation and detection, however, in order to achieve higher-level goals such as implementing network-based OID, individual agents are put into agent cells, which can be of five different types:

1. *Traffic Signature Cell (TSC)*. Composed of agents S and C. S provides traffic data to the correlator C (a *Signer*) in order to generate traffic signatures corresponding to the input attack tool (Figure 4.6).

2. *Process Signature Cell (PSC)*. Composed of agents S and C (notice that agent types are generic; specific sensor and correlator sub-types exist for the tasks of handling process and network data; see ontology description in the next section). S provides process data to the correlator C (a *Signer*) in order to generate a process signature corresponding to a certain attack tool (Figure 4.6).

3. *Traffic Monitoring Cell (TMC)*. Composed of agents S, C, and R. It implements attack tool detection by processing network packets. S provides traffic data to the correlator C (a *Detector*) which in turn may command R to react upon an event (Figure 4.7).

4. *Process Monitoring Cell (PMC)*. Composed of agents S, C, and R. It implements attack tool detection by inspecting process data. S provides process data to the correlator C a *Detector* which in turn may command R to react upon an event (Figure 4.7).

5. *Correlation Cell*. Composed of agents C and R. It combines the output received from traffic and process monitoring cells (TMC and PMC) in order to more accurately detect the execution of a rogue program. Once C receives information from monitoring cells, it may command R to react upon an incident (Figure 4.8).

4.8.2 Signature and ontology deployment

There are two signature-related processes that compose the architecture, one to generate signatures, and one to match them against observed events. Signature generators are called upon by a user who wishes to generate signatures for an attack tool. These signatures will be utilized by the monitoring agents in order to automatically detect attack attempts launched by a local (legitimate or intruding) user.

Signatures can be generated and used by local agents only, however, they can also be shared with other systems. Alternatively, they can be generated at a central location and then deployed to agents. Many anti-virus software systems use a signature deployment mechanism via the Internet which supplies signature updates to subscribers [2]. A similar model could be implemented for attack program signatures.

Figure 4.9 shows a distribution scheme. A central repository serves as storage for the signatures produced by signature-generation sites. This repository is mirrored by other Internet

Sophos Ex. 1009
Page 181 of 318

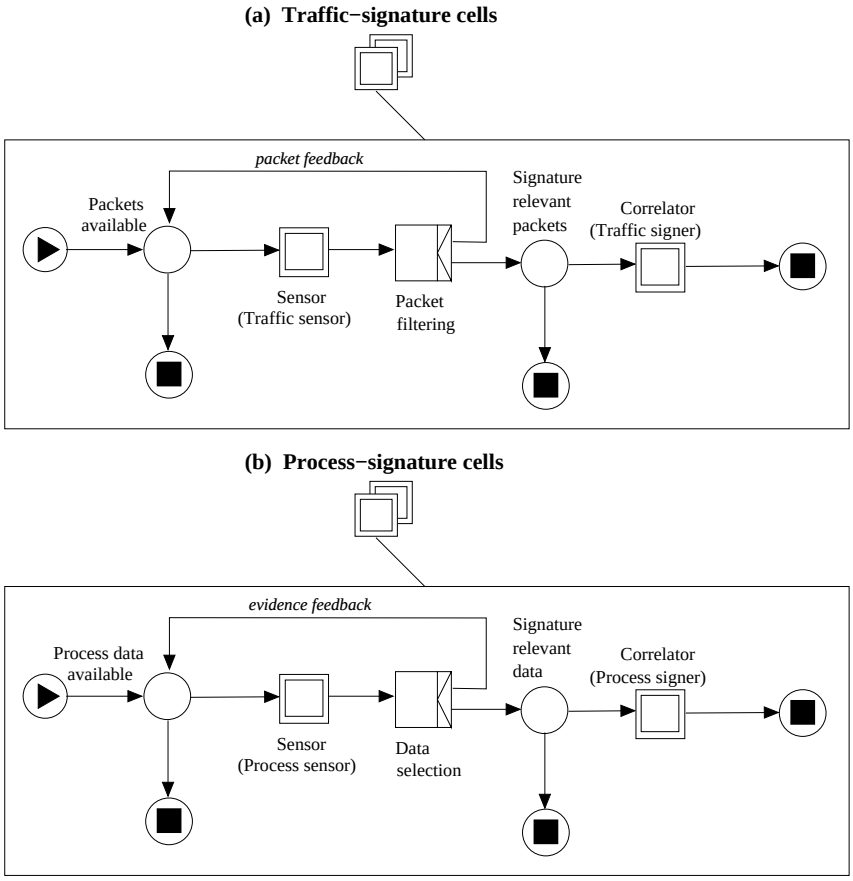


Figure 4.6: Internal dataflow: (a) Traffic-signature cells, and (b) Process-signature cells (see Appendix A for YAWL symbols).

servers which are accessed periodically by agents in order to obtain the most up-to-date version of the signatures or the ontology (an ontology for intrusion detection is rather static; there is no need to generate an administration agent that answers queries or automatically maintains the ontology [31]). Cryptographic certificates can be used to sign the bundles containing the signatures so that agents are able to verify their integrity and authenticity.

4.8.3 Operational outline

Summarizing, two data streams are the input to the architecture: outbound network traffic and process execution data (Figure 4.2). For signature generation, signature cells access these streams and extract information at the time the attack tool is executed. They share these signatures with local monitoring agents and alternatively download signature updates from an Internet signature repository (the repository is maintained by registered signature generation sites;see Figures 4.5 and 4.9). Traffic monitoring cells permanently screen network packets looking for a match with the existing signatures. In parallel, process monitoring cells do the same accessing the process data stream. Either of these two monitoring cells may identify an

Sophos Ex. 1009
Page 182 of 318

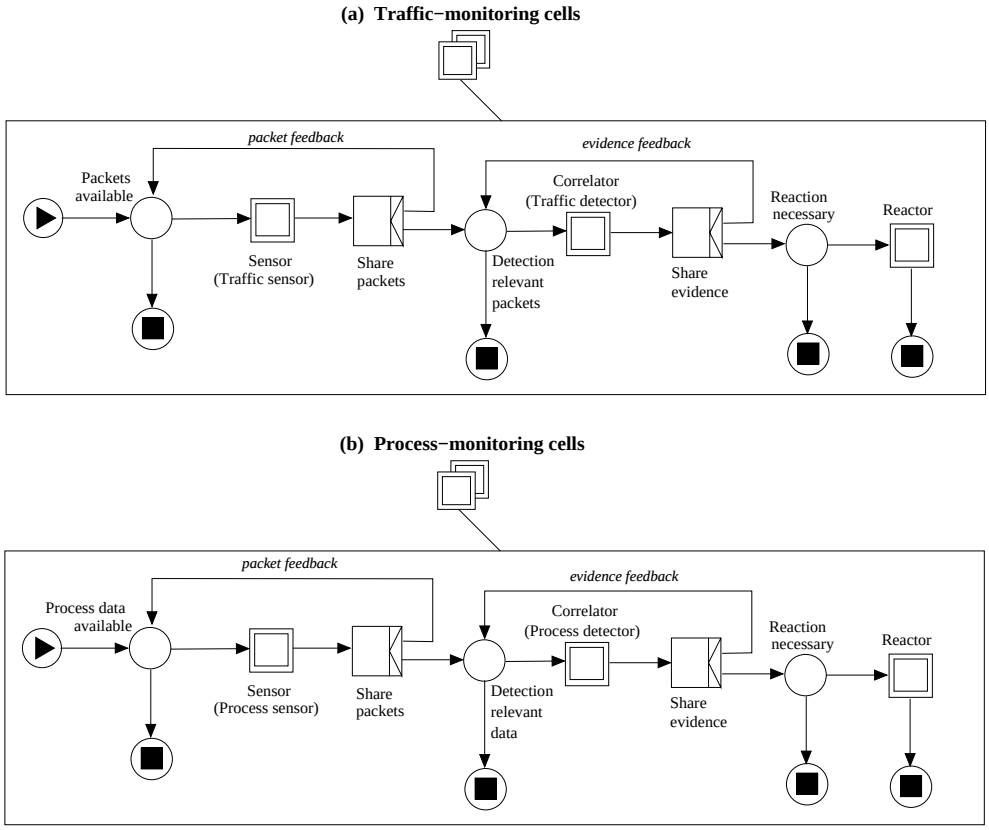


Figure 4.7: Internal dataflow: (a) Traffic-monitoring cells, and (b) Process-monitoring cells (see Appendix A for YAWL symbols).

attack tool and react accordingly. An additional correlation cell receives evidence from the monitoring cells, which it integrates in order to more precisely identify the attack program that has been detected.

4.9 Ontological model

The OID problem approached from a multiagent perspective presents the need for a common understanding of concepts and relationships, and for a more simple and meaningful integration with other security technologies such as firewalls and vulnerability scanners. Given the amount of data that needs to be analyzed, it is also necessary to filter and optimize communication contents in order to speed up all tasks. An ontology also helps reduce the size of the messages exchanged, limits redundancy, and allows for a more intelligent information analysis [204]. (When describing ontology components graphically, bold arrows will represent subclass relationship and soft arrows will represent attributes; see figure 4.10.)

Sophos Ex. 1009
Page 183 of 318

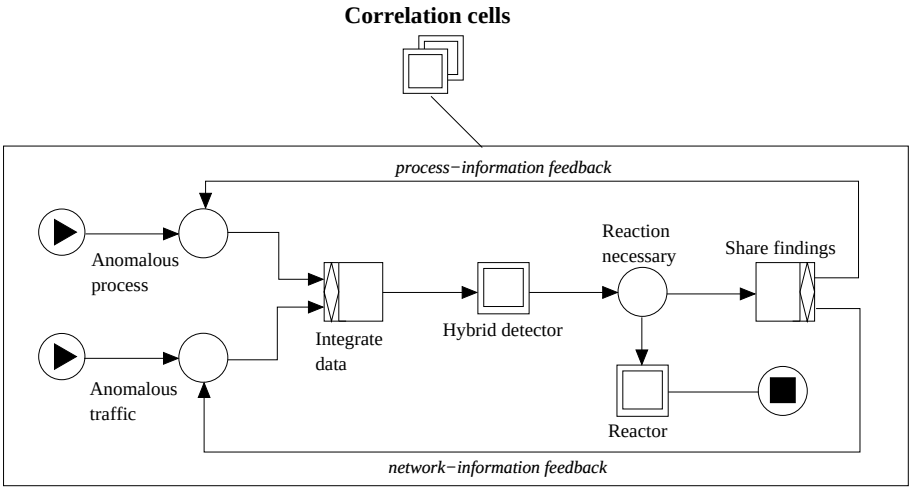


Figure 4.8: Dataflow inside Correlation cells (see Appendix A for YAWL symbols).

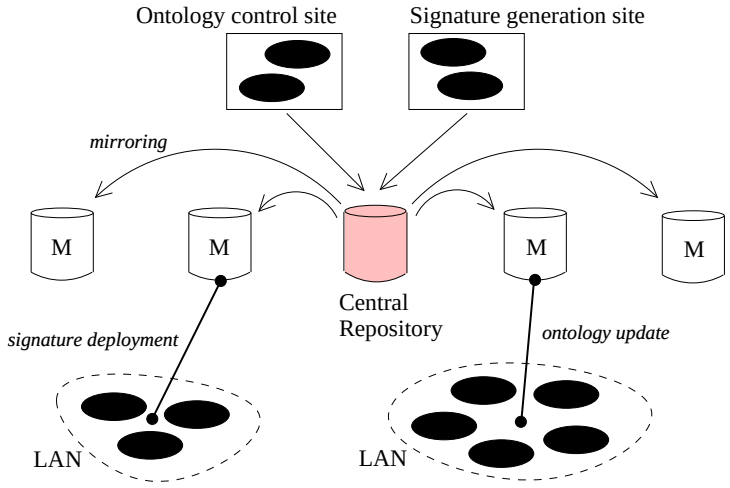


Figure 4.9: Deployment of attack program signatures. Repositories connected to the Internet receive signatures from registered signature generation sites. Agents download signatures and ontology updates (M represents a mirror repository).

4.9.1 Supporting agent-based security technologies with ontologies

Having a central specification of the structure of a data set helps automate the verification of the input for compliance with the structure. This is a basic benefit of using ontologies, however, its main advantage appears when logic processes need to be implemented and when data is being exchanged among different system components.

Most research in intrusion detection has been in taxonomic models. A taxonomy is always part of an ontology but, in addition to the domain classification benefits it offers, ontologies

Sophos Ex. 1009
Page 184 of 318

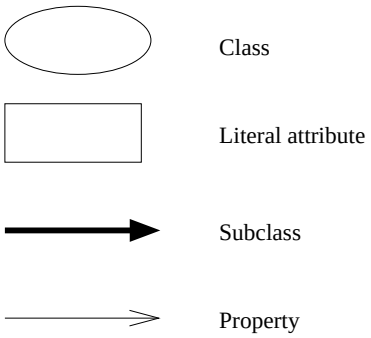


Figure 4.10: Graphical notation

provide constructs that allow a system, in this case an OID system, to reason over existing instances of the taxonomy [203]. They constitute a meta-data framework of the objects of the environment and describes them with all their features and relationships.

Relationships and attributes allow to infer information not explicitly specified when communicating a message. For instance, if the ontology establishes that (in RDF notation [203]): `(Agent, subclass, Reactor)` and also that `(Reactor, subclass, Tracer)`, an agent cell may correctly infer that a given `tracer` agent will inherit the functionality of `Agent` and `Reactor` as well even when this has not been explicitly said about `Tracer`. This information needs not be communicated, but only the basic facts.

An ontology for an intrusion detection architecture is necessary to enable more intelligent behavior in agents, to optimize communication contents and interpretation, and to give formalism to the way the components of an architecture interact (an aspect that is usually ignored). "... An ontology decouples the data model representing an intrusion from the logic of the intrusion detection system"[203]. From a software engineering perspective, this facilitates the integration with other technologies by sharing data in a standard format that can be equally conceptualized using the machine-interpretable constructs of the ontology.

4.9.2 Target-centric ontology for intrusion detection

In the area of intrusion detection there have not been very many efforts aimed at defining ontological schemes that enhance and simplify the interaction of operating systems and security applications [196, 43]. The most relevant project is the one by Undercoffer et al. [202, 203]. They defined a target-centric ontology for intrusion detection which is based on the analysis of over 4,000 intrusion instances. Figure 4.11 shows a representation of this ontology and Figure 4.12 shows the ontology for OID which is proposed in this dissertation and which is explained in the next sections.

The target-centric ontology project defines intrusion detection concepts from the point of view of the victim claiming that intrusion detection systems should exclusively handle information that can be measured and observed locally [202]. This same perspective is used in OID but, in this case, the position of the attacker is at a local host: the perspective of the attacking party in OID is equivalent to the perspective of a victim in intrusion detection models. Only locally observable information is used in OID to characterize a program's behavior and the network

Sophos Ex. 1009
Page 185 of 318

traffic it generates.

The taxonomic structure of this ontology defines three main classes:

- *target of attack* = {*network, kernel space, application, other*}
- *means of attack* = {*input validation, exploits, configuration*}
- *consequences of attack* = {*DoS, user access, root access*}.

that are used to describe how an attacker compromises a target through a certain means and this produces some consequences. All the relations among these classes are based on the way the set of analyzed intrusion instances behaves and covers local-to-local as well as remote-to-local scenarios.

There is a point where this intrusion detection ontology and an OID ontology intersect. This point is the *Network* resource (Figure 4.13). From the intrusion detection perspective, *Network* is a subclass of *Location* and represents a potential point from which a host can be attacked. From the OID point of view, *Network* constitutes a *data source* that needs to be monitored by a sensor agent in order to collect relevant information (see Figure 4.12).

Subject–predicate–object structure

An ontology can be understood as a vocabulary that defines the terminology of concepts used by agents and the nomenclature of the relationships between them. It defines the semantics and structure of an integrated agent-based system and allows for inference processes to access known facts in order to create new knowledge from them.

The description of the OID ontology that appears next follows the $(Subject, Predicate, Object)$ RDF statement convention [203]. Let R be a set of system *Resources*, $P \subseteq R$ a subset representing resource *Properties*, and L a set of literal values (e.g., not subjects or classes, but simple values like integers or strings). Then $Subject \in R$, $Predicate \in P$, and $Object \in (R \cup L)$. This relationships can be also described graphically.

This ontology has been defined in the OWL ontology language [7]. OWL extends RDF with richer constructs and allows to specify the meaning of classes and properties in various ways, so that intelligent programs can reason about it. OWL has been chosen because is one of the most recently developed ontology languages, it has undergone a more rigorous review process than other languages such as DAML+OIL and even RDF, and because it "has received significant input to its design from researchers with experience in building inference engines" [7].

OWL documents are typically stored in terms of an XML/RDF file that contains: (1) *classes* to define concepts from the domain, (2) *individuals* to represent specific instances of the classes, and (3) *properties* define the values each individual can take.

For instance, the statement `(Signature, subclass, Process_Signature)` which defines a subclass predicate would be represented in OWL as shown in Figure 4.14. (The document preamble defines some name-spaces such as RDF, RDFS, and OWL itself used to correctly reference objects from those domains).

In the case of an individual of the *Traffic Signature* class, which stores the representation of attack-traffic signatures introduced in the next section, its attributes could be graphically described as shown in Figure 4.15.

Sophos Ex. 1009
Page 186 of 318

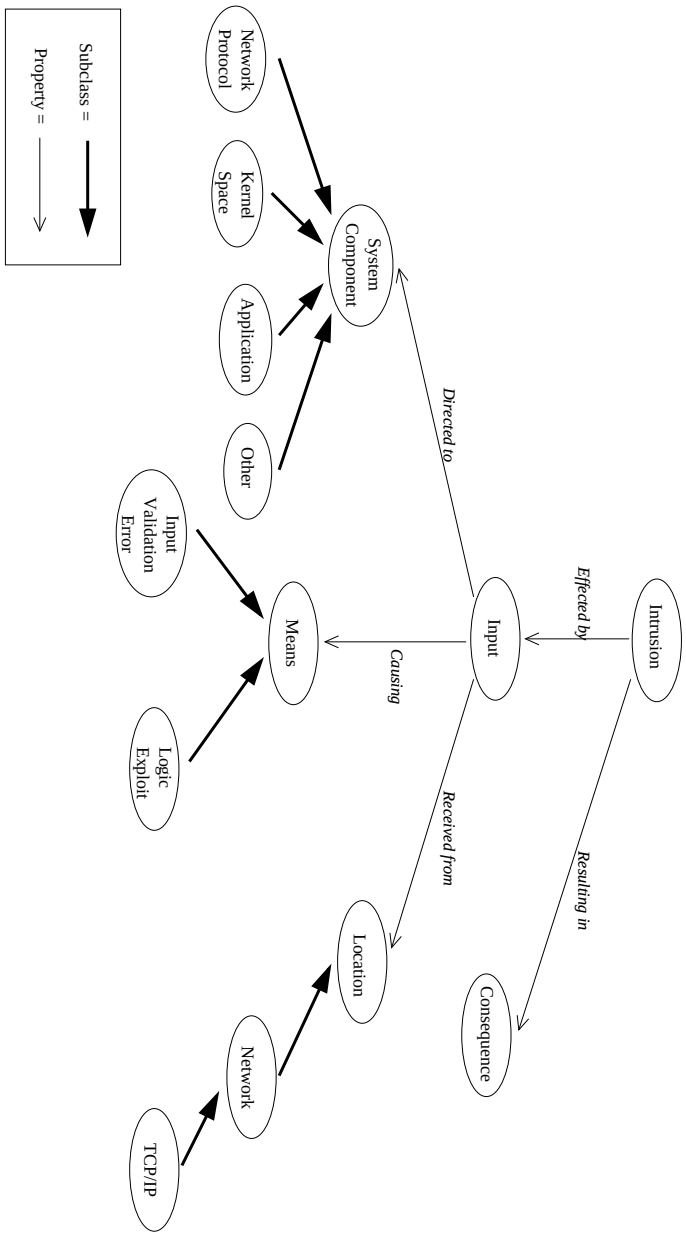


Figure 4.11: A target-centric ontology for intrusion detection [203, 202]; bold arrows represent subclasses; soft arrows represent properties.

Literal attributes qualify the subjects of a class. They represent atomic features that need no further organization or classification and that are directly attached to the subject in question (e.g., the size of a packet payload; it contains an integer number which does not have any

Sophos Ex. 1009
Page 187 of 318

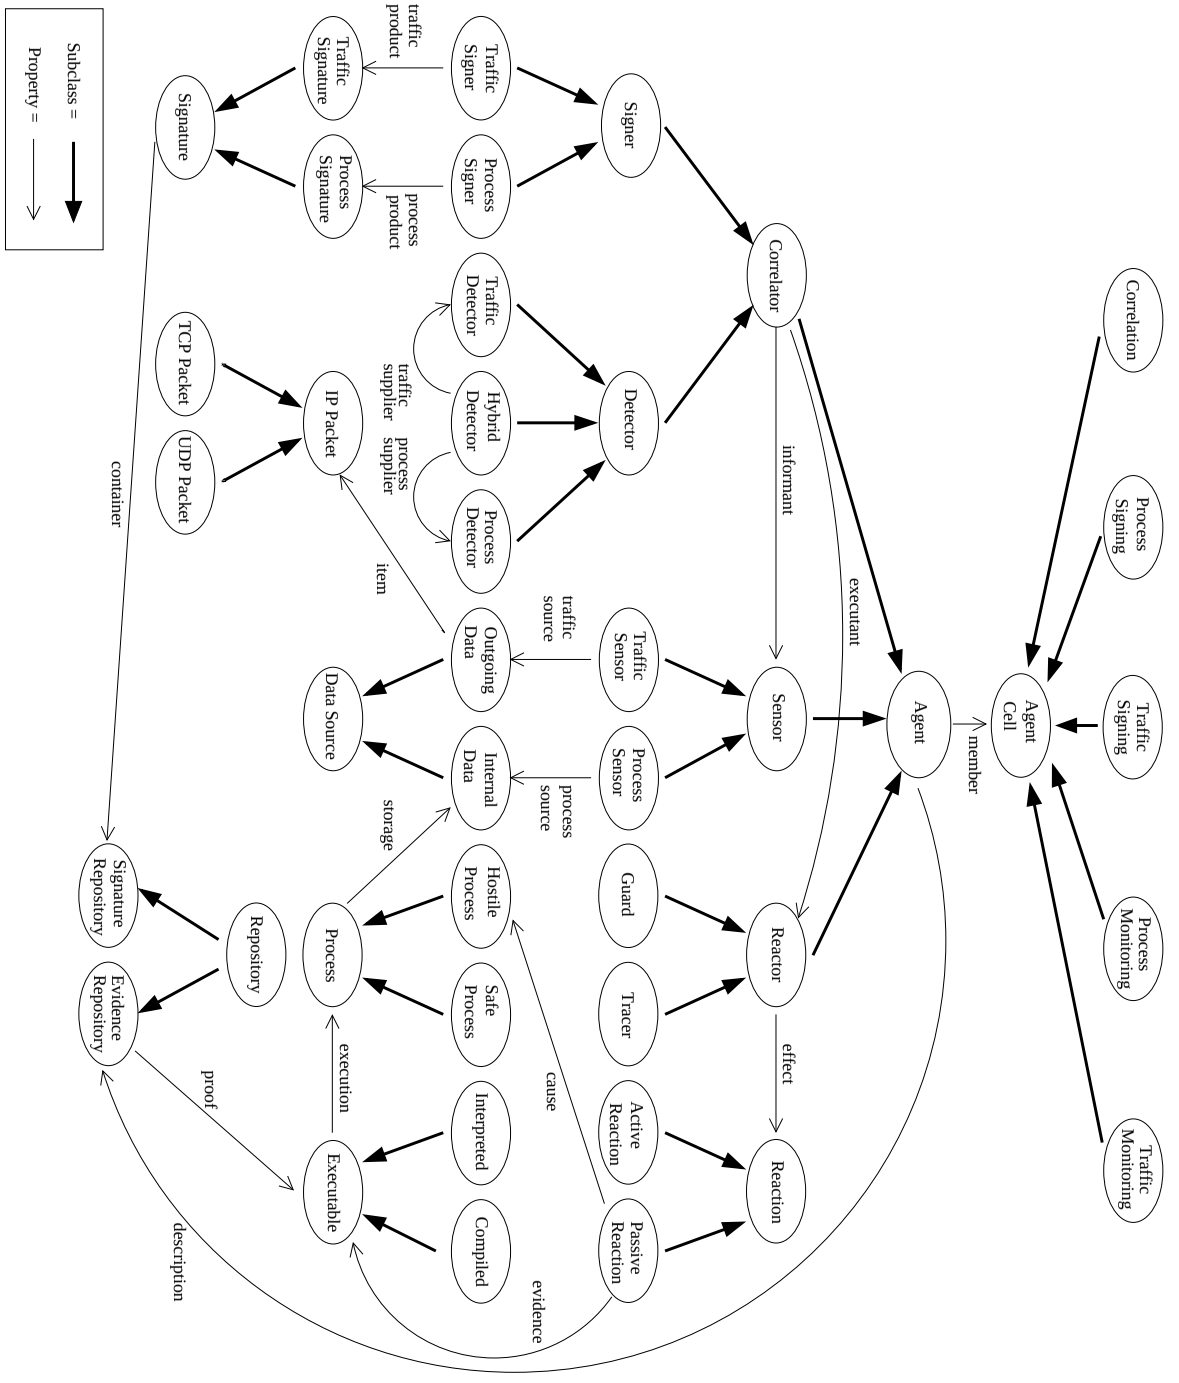


Figure 4.12: A source-centric ontology for outbound intrusion detection (see Appendix B for its definition in OWL)

Sophos Ex. 1009
Page 188 of 318

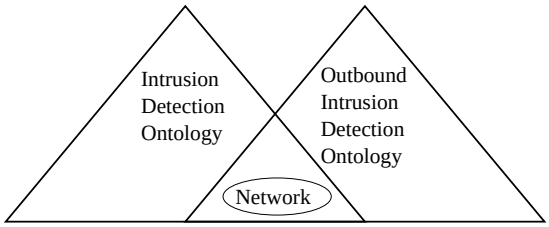


Figure 4.13: Overlapping between the target-centric intrusion detection ontology and the source-centric OID ontology.

```
<rdf:RDF
    xmlns:protege="http://protege.stanford.edu/plugins/owl/protege#"
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
    xmlns:owl="http://www.w3.org/2002/07/owl#"
    <owl:Ontology rdf:about="OID Ontology">
        <owl:imports rdf:resource="http://protege.stanford.edu/plugins/owl/protege"/>
    </owl:Ontology>

    <owl:Class rdf:ID="ProcessSignature">
        <rdfs:subClassOf>
            <owl:Class rdf:ID="Signature"/>
        </rdfs:subClassOf>
    </owl:Class>
...
</rdf:RDF>
```

Figure 4.14: OWL statements defining the Signature–Process Signature subclass predicate.

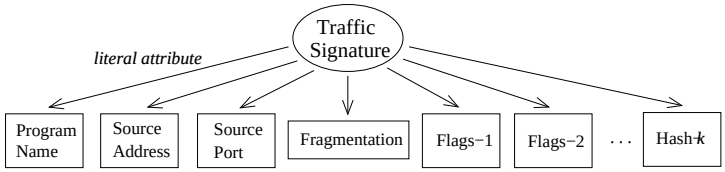


Figure 4.15: Literal attributes of the *Traffic Signature* object.

sub-categories or attributes itself).

If the parameters defining a traffic signature changed, the ontology would have to be updated and the agents would receive this structural modification through the signature deployment system. For instance, if we wished to include a new *programID* attribute into *Traffic Signature*, the class would change and an individual, "tesoIISb" in this case, would look like in Figure 4.16.

These structures allow all architectural components to equally interpret what happens in the environment. The relationships between an exploit program and a file, a signature-generation agent and a data source, a signature and a repository, etc., can be all clearly defined. This

Sophos Ex. 1009
Page 189 of 318

```
<TrafficSignature rdf:ID="tesoIISb">
   <programName>tesoIISb</programName>
   <fragmentation>1</fragmentation>
   <sourceAdddress>127.0.0.1</sourceAdddress>
   <sourcePort>0</sourcePort>
   <destinationAddress>132.243.46.41</destinationAddress>
   <destinationPort>80</destinationPort>
   <windowSize>5480</windowSize>
   <flags>0,8,1,10,1,10,1,1,1,1,1,0,0,0,0,0</flags>
   <hashes>0,0,0,0,1528,1820,0,4104,0,1739,0,0,0,0,0,0</hashes>
   <programID>50</programID>
</TrafficSignature>
```

Figure 4.16: OWL definition of a sample Traffic Signature instance.

reduces ambiguity, makes communication simpler, and creates an interface other agents and software applications can access for communicating with the OID architecture.

4.10 Source-centric ontology for outbound intrusion detection

The first part of this section describes the taxonomic structure of the proposed ontology including object classes, subclasses, and instances. Instances are shown only for subclasses with three or less instances given that some subclasses can potentially have hundreds of instances, like in the case of the IP packet class. The second part of this section outlines the relations among objects of the ontology. The whole contents of the ontology appear in Appendix C.

Taxonomic structure

The components of this ontology are based based on the commonalities found among the different intrusion detection architectures surveyed ([224, 41, 158, 92, 221]). These commonalities were extended with the particular points of the OID approach. The ontology contains nine primary classes that cover all the objects that appear on an OID scenario: *agent class, agent-cell class, data-source class, signature class, reaction class, process class, repository class, IP-packet class*, and *executable class*.

1. **Agent-cell class**. An agent-cell is composed of one or more agents to perform a certain task. There are five types of agent class: correlation cells, (2) process-signing cells, (3) traffic-signing cells, (4) process-monitoring cells ,and (5) traffic-monitoring cells.

2. *Correlation cells* develop integration and interpretation functions. They are composed of a sensor, a correlator, and a reactor agent.

3. *Process-monitoring cells* are responsible for observing the behavior of processes in order to extract data for the detection and signature generation cells. They are composed of a sensor, a correlator, and a reactor agent.

4. *Process-signature cells* receive process information in order to define process signatures. They are composed of a sensor and a correlator agent.

Sophos Ex. 1009
Page 190 of 318

5. *Traffic-monitoring cells* are responsible for observing outbound traffic generated by suspicious programs and supply this data to the detection and signature generation cells. They are composed of a sensor, a correlator, and a reactor agent.

6. *Traffic-signature cells* receive network packets in order to define traffic signatures. They are composed of a sensor and a correlator agent.

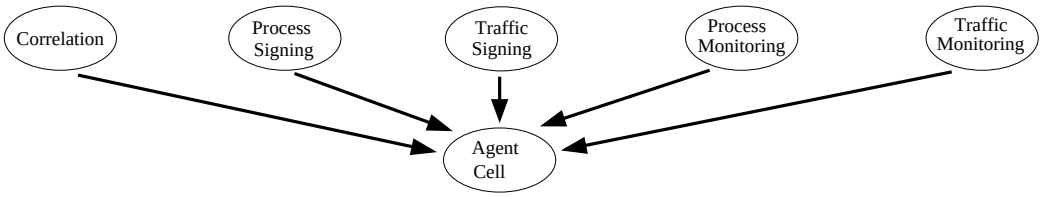


Figure 4.17: The Agent-cell class

7. **Agent class**. There are three main agent types: (1) sensors, (2) correlators, and (3) reactors. These, in turn, have two subclasses each (Figure 4.18).

   Sensors can be: (1) *process monitors*, or (2) *traffic monitors*. This is because a program will be modeled according to its execution as a process and also through the network packets it generates.

   Correlators can be: (1)*detectors*, or (2)*signature generators*. While signature generators receive input from sensors with the objective of defining program signatures, detectors receive input to detect the occurrence of such signatures.

   A *reactor* agent can be: a (1) *tracer*, or a (2) *guard*. Guards implement local reaction to incidents by updating the environment or by passively increasing the monitoring level over a certain resource. Tracers cover the function of back-tracing the connection of an intruder in case an interactive session has been established with a local host.

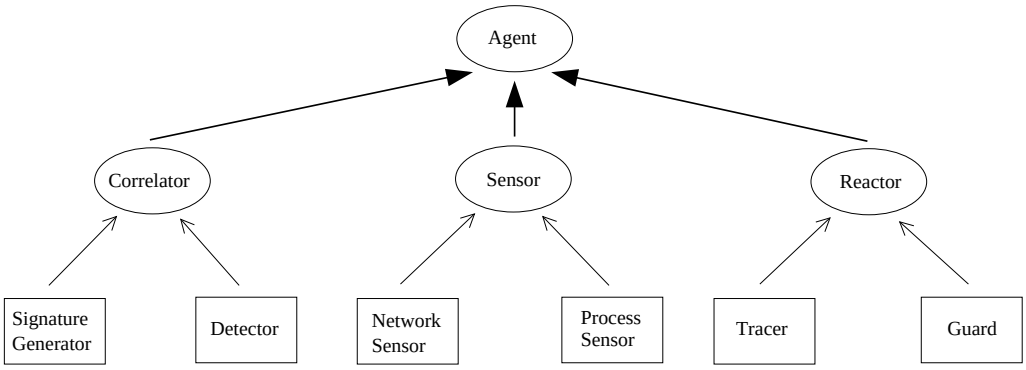


Figure 4.18: The Agent class

Sophos Ex. 1009
Page 191 of 318

8. **Data-source class**. Data sources are classified into two subclasses: (1) data sources for external use, and (2) data sources for internal use (Figure 4.19). It is the type of the ultimate information consumer what determines the essence of the data from a security-monitoring perspective.

   Sources that produce *data for internal use* (e.g., the kernel process table) can be locally accessed. Although a remote party could also access a local system to submit a query, these sources are not intended for external use.

   There are sources that generate *data for external use* (e.g., outgoing network packets). This information has no internal meaning unless it is used for debugging purposes (incoming traffic is not being considered as it does not have the significance to OID outgoing traffic does). In the case of OID, this is a basic information resource.

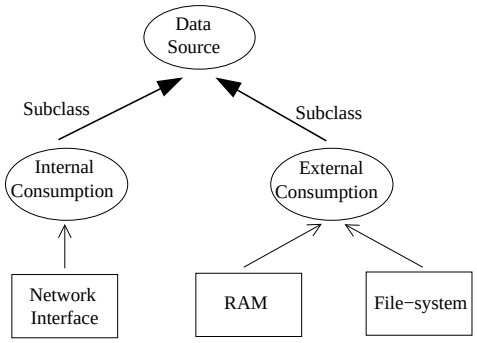


Figure 4.19: The Data-source class

9. **Signature class**. There are two signature types utilized within the OID architecture: (1) traffic signatures, and (2) process signatures (Figure 4.20). These two types of pattern can be used independently or jointly to detect an attack tool.

   A *process signature* contains the execution properties of a program, in particular, regarding the effects it has on local resources while in execution.

   A *traffic signature* characterizes the type of network traffic generated by an attack tool when it is executed. Notice that both signatures are used to characterize the same object (i.e., an attack program).

10. **Reaction class**. There are two types of reaction to incidents: (1) passive reactions, and (2) active reactions (Figure 4.21).

    Passive reactions do not take any steps toward stopping an attack attempt. They focus on collecting additional information that can be used as evidence in the future (back-tracing is an example of passive reaction).

    The goal of *active reactions* is to counter an intrusion attempt by a user who is executing an attack program locally toward an external victim. This class of reaction has more impact on the environment than the mere information-oriented goal of passive reaction.

Sophos Ex. 1009
Page 192 of 318

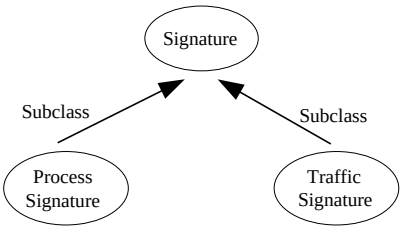


Figure 4.20: The Signature class

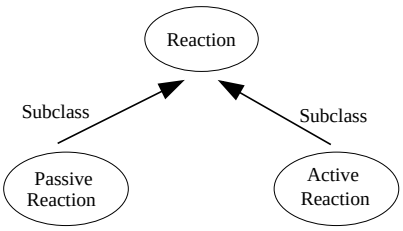


Figure 4.21: The Reaction class

11. **Process class**. A process is a key object for OID. It represents the execution of an attack program and contains critical information for the detection algorithms. There exist two different types of process in this ontology: (1) legitimate processes, and (2) hostile processes (Figure 4.22).

    This categorization is based on matches found by a detectors. *Hostile processes* are subject to further inspection, whereas *legitimate processes* do not have to be monitored once it has been verified they do not pose any threat to the security of other systems.

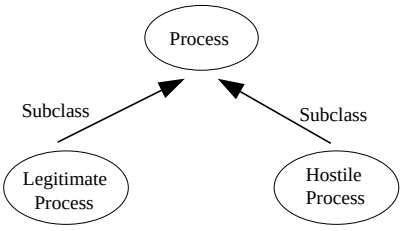


Figure 4.22: The Process class

12. **Repository class**. There are two different types of repository: (1) signature repositories, and (2) evidence repositories (Figure 4.23).

    The purpose of an *evidence repository* is to store information that helped correlators to reach intrusion-related conclusions.

    A *signature repository* is in charge of storing the program signatures generated by the signers. A repository is a passive entity, and as such, its administration is not autonomous. Agents

Sophos Ex. 1009
Page 193 of 318

are in charge of keeping these repositories up to date (alternatively, a security analyst can be in charge of maintaining this type of repository).

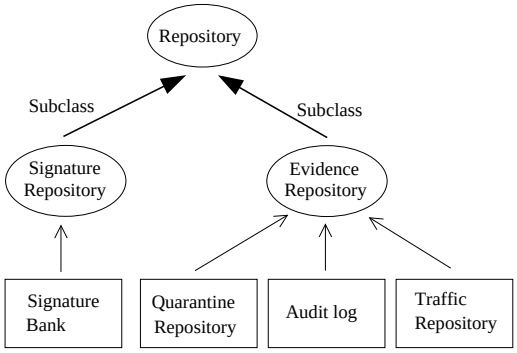


Figure 4.23: The Repository class

13. **IP Packet class**. Network packets are key to signature definition and detection. This ontology supports two subclasses of IP traffic: (1) TCP packets, and (2) UDP packets (Figure 4.24). IP traffic counts for a big percentage of the attacks performed over the Internet [143] and the ontology could be extended in case other protocols needed to be supported.

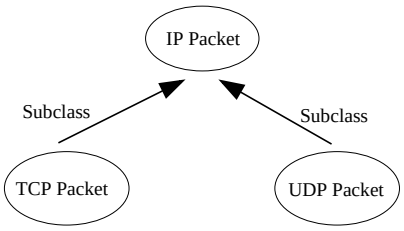


Figure 4.24: The IP-packet class

14. **Executable class**. The proposed OID approach is based on the execution of programs. There are two main categories we are considering: (1) attack programs which execute via an interpreter, and (2) attack programs that are compiled into binary format and can be understood by the host (Figure 4.25).

    It is necessary to split the Executable class into these two categories because the analysis that needs to be performed for each of them is different and because the information exchanged among agents is dependent on the type of executable that is being inspected.

4.10.1 Use of the ontology

The ontology is used by agents to process messages and to derive information not explicitly conveyed to the listener during a communication. There are two main tasks: (1) message interpretation and (2) consistency verification.

Sophos Ex. 1009
Page 194 of 318

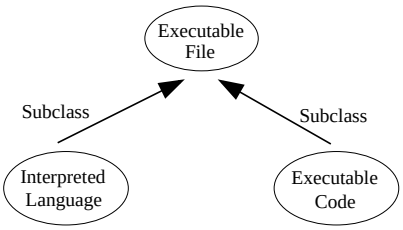


Figure 4.25: The Executable class

- **Message interpretation**. All messages exchanged between agents are based on the OID ontology. The sender builds the message using ontology classes in no particular order and filling out attribute values. The receiver expects to see certain message types depending of its own agent type and therefore looks for statements that are meaningful to itself. A major advantage of this approach is that the knowledge of the agent (i.e., the ontology and its knowledge base) is used to reduce the amount of data that needs to be exchanged.

  For instance, when preparing a message the sender simply specifies its agent ID and the listener will be able to fetch its agent type from its own knowledge base – updated every time a cell is spawned. From this single piece of data it is possible to infer parent classes and relationships the sender keeps with other system components.

  When a traffic signature is open by a *Traffic Detector*, the agent may require additional information regarding the corresponding process signature. Specifically, it will have to know the originating *Process Signer* that created the signature. Even when the *Signer ID* is not explicitly indicated on the message, the ontology indicates a traffic signature is generated by a certain *Traffic Signer* and the agent will have to first query its ontology and then its knowledge base in order to find the specific agent that will in turn point it at the right direction to find the process signature of the program.

- **Consistency verification**. When a message is received by an agent, the ontology is used to perform automatic consistency verification. Not only does it allow the receiver to check the integrity of the classes involved in the message (i.e., whether or not subject-to-property relationships are maintained), but also it can verify that the attribute values are valid and reflect the current state of the system.

  In its simplest way, the verification mechanism will check every message for compliance with the ontology by analyzing the structure of each of the classes included. For some classes this is a basic requirement as any inconsistency can and will affect the performance of some agents (e.g., the contents of a traffic signature; it must be complete and correct).

  Before going ahead with passive reaction, a *Reactor* may verify if the process that triggered the detection mechanism is still running. In order to do this, the *Reactor* will check the ontology and will know that a reaction must have been generated by an executable file. This executable may have a process attached to it and the agent will check whether the process is still alive in order to proceed with process suspension or simply with a change of execution privileges of the file.

Sophos Ex. 1009
Page 195 of 318

4.11 Session tracing

The tracing component of the architecture aims at identifying the hops (i.e., IP addresses or host names) of a connection chain that ends at our host. Connection chains are used by intruders to hide the origin of their connection. From an OID perspective, identifying the components of a connection chain helps warn other hosts about potential security threats.

Our system is composed of a number of tracing agents that communicate with each other sharing information on user connections [3]. There is an agent running at each host h_i and the ideal situation is having an agent running on every host of the network. Every agent has client and server capabilities that enable it to request information from other agents and to supply information to them. An agent is not useful by itself. It needs to collect information from the others in order to identify the links of a connection chain. The agent at the target host will eventually collect information from a number of its peers and will be in a position to inform to the local administrator about the possible origin of the connection.

The perceptions of a tracing agent are received from two main sources: (1) the state of the system regarding user sessions and connections, and (2) messages from other agents requesting or providing information.

The environment an agent inhabits is highly dynamic and provides the communication channels necessary for an agent to stay in touch with its peers over the network – the environment can also represent a threat to the tracing system if malicious users gain access to it.

As a result of the changes perceived by an agent, it can perform a number of different actions: (a) send messages to an agent, (b) request messages from an agent, (c) broadcast messages to all tracing agents, (d) generate encryption keys, (e) encrypt messages, (f) decrypt messages, (g) identify connection loops, (h) retrieve system status information, and (h) save its knowledge to a safe location.

4.11.1 Internal design

The knowledge of an agent is stored on a repository known as **knowledge base** (KB). This KB includes tracing host information and user data. A tree structure stores the information on the hosts that are part of a connection chain. Information regarding user names, login times, and dates is stored on a separate structure. These two elements represent all the knowledge an agent has about the state of its environment. Should this knowledge be lost, the agent would lose track of previous events and would have to be updated by its peers.

The actions an agent takes depend on the information of the state of the system and its perceptions. Embodied into the agent are a set of rules that enable the agent to work autonomously deciding what to do. As new knowledge is fed into the KB, more rules from the rule base need to be analyzed before making any conclusions and proceeding upon them. The type of perceptions and actions do not change but the state of the environment does evolve hereby modifying the behavior of agents. The type of rules that make up the rule base are of the form:

[3] This system was implemented in C and has been successfully tested separately from the Java agents that implement OID [118]; the integration of the two modules is still pending (see Section 7.3 in the last chapter)

Sophos Ex. 1009
Page 196 of 318

```
R1:
IF ((HOW-CONNECTED(User)=Console) AND (WHAT-DOING(User)=Ssh-to(Target)))
⇒ TRACING-DONE();


R2:
IF ((MESSAGE-TYPE(Message)=Tracing-request) AND (VALID-PUBLIC-KEY(P-key)=False))
⇒ DENY-CONNECTION() AND BROADCAST-TREE() AND ALERT();
```

In order to protect their information exchanges, all agents have an RSA key-pair that is used to negotiate symmetric session keys before encrypting their messages [179]. Encryption at this level guarantees that the information will not be visible at any point before reaching the top of the IP stack of the recipient party. This will deter sniffing attacks that would otherwise be possible with agent communication in plain text [93].

The tracing process starts when a connection needs to be monitored. Agent A at target host h_t receives local information from the system and checks for remote connections. If there is a connection from a remote machine h_r, A will contact agent B at host h_r in order to research the connection. A private channel will be opened using the asymmetric encryption keys of the agents. A symmetric session key will be agreed upon and all messages will be encrypted with it for that exchange. The same procedure will be followed if there are relevant connections from outside of the host of h_r coming into the server. A set of secure channels will be setup by the agents in order to share the information they have regarding connections. One hop at a time, the links composing a connection chain will be found.

4.11.2 Hop identification

The tracing information obtained by an agent is shared with the rest of the agents through broadcast communication. This guarantees that the KB's of the agents are consistent with the state of the environment and that, in the event of a security incident, the knowledge acquired by an agent will not be lost as it has been replicated by other agents.

If we have a set H of n hosts, and a user creates a connection chain C of m hops, the system will behave well – with a *O(n)* worst case scenario. Suppose $m \gg n$. If the number q $(0 < q < m)$ of connections that go from host h_x into host h_y – or vice versa – form a loop, the arrangement of the agents will find at most two relevant connections between the hosts hereby unwinding the loop (e.g., a connection from A into B and a connection from B into A involving the same two user ID's). This is due to the fact that an agent is capable of identifying repeated connections toward itself by storing information on the connected users. In the case of loops passing through a total of k hosts, the chain will be reduced to at most $k-1$ hops. This makes a worst case scenario of *O(n)* with $n = |H|$.

The unwinding method is based on previously observed activity. An agent first stores user name and host-name information on its KB. When a tracing process is started, an agent will engage in communication with several neighbor agents in order to collect data that allow them to draw conclusions. Each time agent A receives new information, it checks its KB for loops. This is done as follows.

All tracing requests have an ID number w. Suppose agent A initiates a tracing round. If

Sophos Ex. 1009
Page 197 of 318

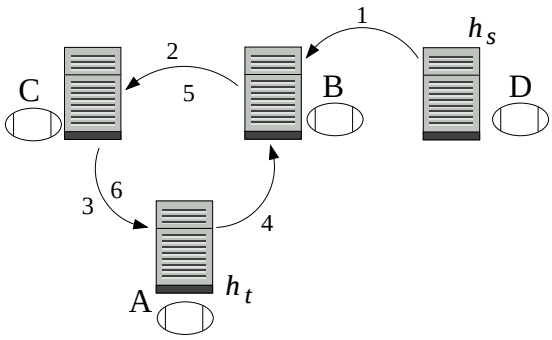


Figure 4.26: Loop unwinding. If a user connects from host h_s to host h_t doing a loop through hops 1, 2, 3, 4, 5 and 6 as shown above, agent A at h_t will start the tracing process and will be able to identify the loop to reduce tracing time in the future.

while waiting for news on a particular w the agent is being asked by a cooperating agent B for information to complete tracing round w as well, the agent knows it is part of a loop since it is receiving information on a tracing round that it itself started (all agents, including initiators have this property). The trace is stored on its KB tree T and, if the user being traced repeats a loop that passes through the same host where agent A is located, this agent will reduce the tracing time by omitting a further requests to follow a loop.

Once an agent F reports a local connection (e.g., connection at the console) when investigating a determined chain, the originating agent A concludes the search and reports the components of the $A \rightarrow F$ chain.

Sophos Ex. 1009
Page 198 of 318

Chapter 6

Experiments and results

This chapter presents a set of experiments performed to test the implementation of the OID architecture. It describes the test set and experimental setup, the design of the experiments, and the results obtained from the evaluation of the accuracy and performance of the model.

6.1 Experimental objectives

The goal of the test phase is to analyze the feasibility of the OID approach and to evaluate the accuracy and performance of the FROID implementation under different traffic and workload conditions. The specific objectives are: (1) to evaluate the detection accuracy of individual cells, (2) to evaluate the detection accuracy of cells when detecting intrusion jointly, (3) to evaluate the performance of individual cells, and (4) to evaluate the performance of cells when detecting intrusion jointly.

6.2 Test data

The OID architecture has been tested using a sample group of automated remote-attack programs downloaded from the Internet. Some of these programs are proof-of-concept scripts that were developed to exemplify how a certain vulnerability can be exploited on a determined platform. They were obtained from vulnerability repositories and databases that store this kind of software and often provide references to security patches and workarounds [24, 149, 61].

The rest of the programs was obtained from the websites of different black-hat communities [150]. These programs have been specifically developed to exploit security bugs and they typically make reference to vulnerabilities already reported on vulnerability repositories.

6.2.1 Exploit types

The types of attack program being considered are four: (1) buffer-overflow exploits, (2) denial-of-service tools (DoS), (3) surveillance tools, and (4) misconfiguration exploits. Our search for experimental data confirmed the fact that buffer overflows are by far the most frequently exploited type of vulnerability and many of the tools surveyed belong to that type [124] (Table

Sophos Ex. 1009
Page 199 of 318

6.2). (Since we are evaluating the OID approach, we have only considered programs that target remotely-exploitable vulnerabilities).

DoS is a type of attack which is often automated through programs due to the need of submitting multiple requests or network packets to the target system in an attempt to rendering it inoperable [111]. Similarly, surveillance software often provides the possibility of analyzing a range of services, IP addresses, and port numbers. This is easily accomplished by means of scripts and programs that allow the user to repeat the execution of a series of operations. Misconfiguration errors are system administration flaws present at hosts with poor security management. The exploitation of this type of vulnerability has also been built into attack programs and therefore we have included some of them in our testing.

6.2.2 Target platforms

The programs that compose the test set can compromise the security of specific services and applications running on the target. In the case of buffer overflows, the pre-compiled code that has to be inserted into the RAM of the target system is platform-specific and some tools contain code for multiple platforms the intruder can choose from (Figure 6.1).

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include <sys/types.h>
#include <sys/socket.h>
#include <netinet/in.h>
#include <arpa/inet.h>

#define DUMMY            0x41414141
#define BUFSIZE          512
#define IP_INDEX         5
#define PORT_INDEX       11

#define FMTSTR           0x1
#define OVERFLOW         0x2

#define JMPCODE          "\xeb\x0c"

static char linux_code[] =
    "\x31\xc0\x50\x50\x68\xc0\xa8\x01\x01\x66\x68\x30\x39\xb0\x02"
    "\x66\x50\x89\xe6\x6a\x06\x6a\x01\x6a\x02\x89\xe1\x31\xdb\x43"
    "\x30\xe4\xb0\x66\xcd\x80\x89\xc5\x6a\x10\x56\x55\x89\xe1\xb3"
    "\x03\xb0\x66\xcd\x80\x89\xeb\x31\xc9\x31\xc0\xb0\x3f\xcd\x80"
    "\x41\xb0\x3f\xcd\x80\x41\xb0\x3f\xcd\x80\x31\xd2\x52\x68\x2f"
    "\x2f\x73\x68\x68\x2f\x62\x69\x6e\x89\xe3\x52\x53\x89\xe1\xb0"
    "\x0b\xcd\x80\x31\xc0\x40\xcd\x80";

static char freebsd_code[] =
    "\x31\xc0\x50\x50\x68\xc0\xa8\x01\x01\x66\x68\x30\x39\xb4\x02"
    "\x66\x50\x89\xe2\x66\x31\xc0\x50\x40\x50\x40\x50\x50\x30\xe4"
    "\xb0\x61\xcd\x80\x89\xc7\x31\xc0\xb0\x10\x50\x52\x57\x50\xb0"
    "\x62\xcd\x80\x50\x57\x50\xb0\x5a\xcd\x80\xb0\x01\x50\x57\x50"
    "\x83\xc0\x59\xcd\x80\xb0\x02\x50\x57\x50\x83\xc0\x58\xcd\x80"
    "\x31\xc0\x50\x68\x2f\x2f\x73\x68\x68\x2f\x62\x69\x6e\x89\xe3"
    "\x50\x53\x89\xe2\x50\x52\x53\x50\xb0\x3b\xcd\x80\x31\xc0\x40"
    "\x50\x50\xcd\x80";

struct target {

...
```

Figure 6.1: Sample attack program containing root-shell code for two different target platforms: Linux and FreeBSD.

Sophos Ex. 1009
Page 200 of 318

Given their analysis purpose, network surveillance tools can be used to probe many types of system. A service may have been implemented for different platforms so the objective of surveillance programs will be identifying which of those implementations and what operating system the target system is running.

In order to reflect platform dependency, there are three platforms that are being considered: (1) Windows* systems (which includes all the MS Windows operating systems), (2) Unix/Linux (covering Unix and Linux variants), and (3) CISCO (includes CISCO 677/678, CISCO IOS, and Catalyst).

6.2.3 Programming languages

The bast majority of attack tools we found are written in any of two programming languages: C and Perl. Considering its support for pointer arithmetic, its predominant use in the development of Unix components and applications, as well as the functionality it provides for accessing system resources, C is the language of choice for many programmers that develop attack tools.

Unlike C, Perl is an interpreted language frequently used for creating shell and CGI scripts. It provides powerful data and text manipulation functions that make it suitable for writing vulnerability exploits for many services such as `HTTP`, DNS, and `Sendmail`.

We have considered these two languages in our testing to reflect the fact that they are the two most widely used languages for writing attack code.

6.2.4 Test program set

The experimental phase of the project considered fifty different remote attack programs. Their names appear listed in Table 6.2 and they cover the attack types, platforms, and languages mentioned in the previous subsections. Four of these tools require root access to be executed (`exceed, nxt, remo`, and `ss`). Again, OID covers the majority of the incidents under the assumption that inexpert users are not capable of hiding their activities. These four tools have been included into the testing set given the type of traffic they generate without considering the access level they require.

	Exploit name	**Type**	**Platform**	**Language**
1	305monit	Buffer overflow	Windows*	C
2	akill2	Buffer overflow	Windows*	Perl
3	anubisexp	Buffer overflow	Unix/Linux	C
4	apache-nose	Buffer overflow	Unix/Linux	C
5	apache-scalp	Buffer overflow	Unix/Linux	C
6	arksink	Buffer overflow	Unix/Linux	C
7	atphttp	Buffer overflow	Unix/Linux	C
8	bazooka	Buffer overflow	Windows*	Perl
9	bysin	DoS	Unix/Linux	C
10	cge	DoS	CISCO	Perl
11	cobain	Buffer overflow	Unix/Linux	C
12	delefate	Buffer overflow	Windows*	C
13	emule	DoS	Windows*	Perl

Sophos Ex. 1009
Page 201 of 318

14	esignal	Buffer overflow	Windows*	Perl
15	ethboom	DoS	Windows*-Unix/Linux	C
16	exceed	DoS	Unix/Linux	C
15	ethboom	DoS	Windows*-Unix/Linux	C
16	exceed	DoS	Unix/Linux	C
17	eXtremail2	DoS	Unix/Linux	C
18	ezploit	Misconfiguration	Windows*	C
19	fping	Surveillance	Windows*-Unix/Linux	C
20	freezer	DoS	Unix/Linux	C
21	fspscan	Surveillance	Windows*-Unix/Linux	C
22	gtkftpd	Buffer overflow	Unix/Linux	C
23	iawebmail	Buffer overflow	Windows*	Perl
24	ircd	DoS	Windows*	C
25	isspam	Buffer overflow	Windows*	C
26	lsasrv	Buffer overflow	Windows*	C
27	namedsploit	Buffer overflow	Unix/Linux	C
28	niprint	Buffer overflow	Windows*	C
29	nptd	Buffer overflow	Unix/Linux	C
30	numby	Surveillance	Unix/Linux	C
31	nxt	Buffer overflow	Windows*-Unix/Linux	C
23	pam	Buffer overflow	Unix/Linux	C
24	phorum	DoS	Windows*	Perl
25	picobof	Buffer overflow	Windows*-Unix/Linux	C
26	priv8	Buffer overflow	Unix/Linux	Perl
36	proftpd	Buffer overflow	Unix/Linux	C
37	qex	Buffer overflow	Unix/Linux	C
38	realown	Buffer overflow	Windows*	C
39	remo	Buffer overflow	Windows*	C
40	rootdown	Misconfiguration	Unix/Linux	C
41	rpcexec	Buffer overflow	Windows*	C
42	smurf4	DoS	Windows*-Unix/Linux	C
43	spxeneo	DoS	Windows*	Perl
44	ss	Surveillance	Windows*-Unix/Linux	Perl
45	sslbomb	DoS	Windows*	C
46	symanhoddos	DoS	Windows*	C
47	telexp	Buffer overflow	Windows*	C
48	tlsdos	DoS	Windows*	C
49	vwexploit	Buffer overflow	Unix/Linux	C
50	yapping	Surveillance	Windows*-Unix/Linux	C

Table 6.2: Automated remote-attack tools used in the experiments.

Sophos Ex. 1009
Page 202 of 318

6.3 Experimental setup

In order to run the experiments we used a setup that included two source systems (*attackers*) and three target systems (*victims*). These two groups were connected to separate segments of a LAN that supports hundreds of users every day [1].

Name	Function	OS	RAM	Processor	NIC
System 1	Source	RedHat Linux 9.0	512Mb	Intel Celeron @ 800MHz	10/100 Ethernet
System 2	Source	RedHat Linux 9.0	1Gb	Intel Pentium @ 1.8GHz	10/100 Ethernet
System 3	Target	MS Windows 95	128Mb	Intel Pentium @ 700MHz	10/100 Ethernet
System 4	Target	MS Windows 98	512Mb	Intel Pentium @ 1.2GHz	10/100 Ethernet
System 5	Target	RedHat Linux 7.0	256Mb	Intel Pentium @ 650MHz	10/100 Ethernet

Figure 6.2 depicts our experimental setup. The two source machines run RedHat Linux 9.0 and each of the target machines supports a different operating system (Windows 95, Windows 98, and RedHat Linux 7.0). The targets were all configured with default settings from the installation CD's and no service packs or patches were applied. The three target systems run different operating systems that have services vulnerable to a good number of existing attacks. The selected programs exploit flaws on TCP/IP-based network services that are installed along with the operating systems (for instance, service daemons for HTTP, DNS, News, e-mail, and network printing; all these are services reachable through the network).

The OID approach works with one or more machines that try to defeat the security of a target using the Internet and a suite of attack and surveillance tools. Although instances of FROID can be installed on all the machines, OID focuses on detection at the source of the event and there is where FROID agents were installed.

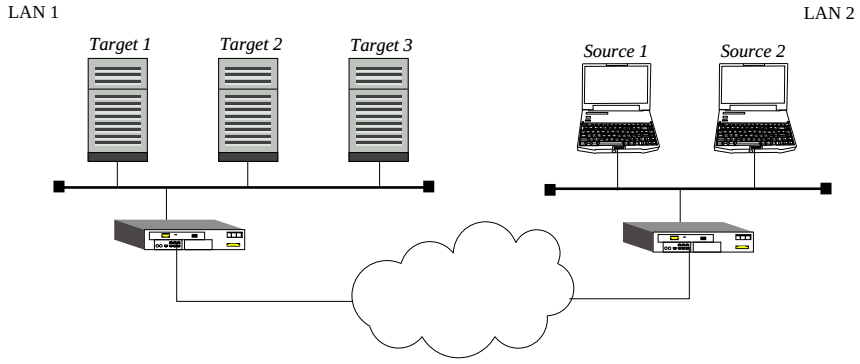


Figure 6.2: Experimental setup for the testing phase.

[1] The CSF network of the Department of Computer Science of Cornell University as well as resources from the Center for Intelligent Systems and the Department of Computer Science of ITESM were used in the development and testing of FROID.

Sophos Ex. 1009
Page 203 of 318

6.4 Signature generation

Source system 1 was used to generate signature pairs for each of the 50 attack programs (process signatures appear in Appendix E and network signatures in Appendix F). Two cells need to be active when a program is to be modeled: a network signature cell and a process signature cell.

The system is executed via the Java compiler which receives a call to the JADE environment class:

```
% java jade.Boot N:NetworkSignatureCell T:TrafficSignatureCell
```

Once active, both cells will download the ontology and will start polling the connections file for new entries that indicate a program is sending data through the network interface. They will separately generate signatures that will be appended to the signature repositories of processes and network traffic. For the purpose of having a controlled environment, no network activity other than the one created by the program in evaluation was present at signature-generation time.

For each program at least one execution variant was considered. In order to further classify variants, we divided the behavior of the target system into *vulnerable* and *not vulnerable*. From a networking stand-point, the reaction the source sees from its target differs between these two categories and for that reason, each of the signatures belongs in one of these two groups. As a result, attack attempts as well as successful attacks can be modeled and detected.

The systems takes in average 2.44 sec to compute a network signature and 1.20 sec to compute a process signature. Network signatures take more time since in some cases packets are delayed and the sensor of the monitoring cells has a behavior that makes it wait until a certain number of packets have been received or until the 6 sec timeout threshold is reached; whatever comes first.

After signature generation, signature matching experiments took place.

6.5 A word on compiled and interpreted code

Our test set included 10 interpreted-code programs (PERL programs) and 40 compiled-code programs (C programs). Network-level detection works the same in both cases: a stream of network datagrams is generated by the executable or by the interpreter and it is analyzed by monitoring cells for signs of outbound intrusion.

For process-level detection, the text segment of the interpreter includes the code of the program to execute. This limits the task of computing a hash of a text segment for characterizing the program. However, arguments, file size, etc. are available for evaluating the similarity of the input program with the signatures in the repository. For keeping the scoring system consistent, all PERL programs received a π_i sub-score of $0.2/2 \times 1$ for the MD5 attribute – otherwise all PERL programs would have gotten lower process-signature scores.

6.6 Experiments: Signature matching and correlation

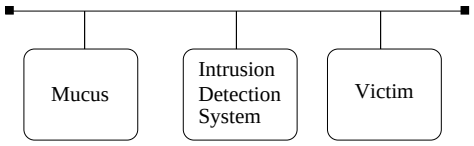
The signature matching method was tested with accuracy and performance experiments. The two main factors that influence these evaluations are the system's internal workload and the

Sophos Ex. 1009
Page 204 of 318

amount of traffic that goes through the network interface [184]. These variables should be controlled in order to effectively measure the precision and response time of the system.

6.6.1 Traffic generation

In order to control and measure the amount of traffic in our tests, we used Mucus [136]. Mucus is an open-source traffic generation tool designed to evaluate network-based intrusion detection systems. It translates Snort signatures [170] into network packet streams and allows the user to vary transmission rates, to modify their contents and type, and to randomize payloads (Figure 6.3). Mucus can be run on any computer that has a network connection.



(a) Topology

```
[smv@localhost mucus-1.0]# mucus
usage: mucus
  [-n <number-of-attack-instances>]
  [-r <random-seed>]
  [-S] (suppress-network-traffic)
  [-d <destination-ip-address>]
  [-s <source-netmask>]
  [-D <fixed-delay-in-seconds>]
  [-v] (turn on verbosity)
  [-R] (random 0.1 - 1.0 sec delay
  [-A] (random rule selection)
  [-N <rule-number>] specific rule
  [-P] (randomize packet payloads, default)
  [-p] (don't randomize packet payloads)
  [-E] (establish TCP sessions for TCP rules)
    [-G <fake-gateway-ip>] (required by -E)
    [-V <fake-victim-ip>]  (required by -E)
  [-e] (don't establish TCP sessions, default)
  ** snort rules are expected on the standard input
```

(b) Options

```
[smv@localhost src]# ./mucus -n 10 -A -r 1234 -d 192.168.1.20 -s 128.0.0.0/8 < ./snort-rules-186
--[ MUCUS options ]---------------
number of attack instances: 10
random seed: 1234
no fixed delay specified
using fixed destination address: 192.168.1.20
randomizing packet payloads.
not establishing TCP sessions.
  ** generating network traffic **
---------------------------------
HOME_NET == 192.168.1.0
EXTERNAL_NET == 0.0.0.0
SMTP == 192.168.1.0
HTTP_SERVERS == 192.168.1.0
SQL_SERVERS == 192.168.1.0
DNS_SERVERS == 192.168.1.0
initializing RNG with seed: 1234
1225 rules parsed.
10 instances of these rules will be generated.
[Wed Mar 16 12:11:08 2004 378] TCP - "WEB-CGI AT-admin.cgi access" - 128.197.61.87:38350 -> 192.168.1.20:80
[Wed Mar 16 12:11:08 2004 381] TCP - "EXPLOIT NNTP Cassandra Overflow" - 128.232.59.235:56184 -> 192.168.1.20:119
[Wed Mar 16 12:11:08 2004 383] TCP - "INFO - Possible Squid Scan" - 128.219.40.5:237 -> 192.168.1.20:3128
[Wed Mar 16 12:11:08 2004 388] ICMP - "ICMP Source Quench (Undefined Code!)" - 128.104.184.71:51294 -> 192.168.1.
...
```

(c) Execution example

Figure 6.3: The Mucus intrusion detection stimulator.

Mucus version 1.0 had to be modified for it to accept shorter packet delays. By default, delays can only go from 0.1 to 1 sec, which is not a good granularity for simulating high traffic loads. The `usleep()` function used in the program receives an argument in microseconds which is multiplied by 1×10E6. We changed this in order to allow the traffic generator to receive delays in microseconds as well. The change in `mucus.cpp` (lines 498–502), was :

Sophos Ex. 1009
Page 205 of 318

```
// before returning, sleep some time as specified by the user
if (_fixedDelaySpecified) {
  //usleep((unsigned int) ((float) _fixedDelay * (float) 1E6));
  usleep((unsigned int) ((float) _fixedDelay * (float) 1));
}
```

Three traffic workload levels have been defined for our testing (the effect of the system load on Mucus is considered part of the normal functioning of the system under those conditions and has not been separately analyzed; the generation of packets for purposes of OID cannot be isolated to a different box). They were used to create different execution conditions at the system:

1. **Low traffic volume**. No network packets are artificially generated. The only traffic present is the one produced by the attack tool itself.

2. **Moderate traffic volume**. An incoming packet is generated every t_i microseconds where $t_i \in [1, 500]$ is a pseudo-random integer, and an outgoing packet is generated after t_o microseconds of the incoming packet where $t_o \in [1, 500]$ is a pseudo-random integer.

3. **High traffic volume**. An incoming packet is generated every t_i microseconds where $t_i \in [1, 10]$ is a pseudo-random integer, and an outgoing packet is generated after t_o microseconds of the incoming packet where $t_o \in [1, 10]$ is a pseudo-random integer.

Sensor agents see all the packets the network interface handles – that includes incoming as well as outgoing traffic. They select outgoing packets and pass them on to correlation agents within the same cell. For that reason, the experiments consider both, artificially-generated incoming and outgoing traffic.

6.6.2 Workload control

The workload at the source system has an impact on the performance of agent cells. When many applications are executing concurrently, cells will have to compete for resources through the scheduling mechanism.

In order to test our application under workload conditions we could replicate, we define two levels:

1. **Low workload**. At low workload the source system does not run any application other than the operating system, the monitoring cells (i.e., Java virtual machines), and the program that is to be observed.

2. **Moderate workload**. In order to generate a moderately high level of system activity consistent throughout the testing, we executed a set of applications on the source machine. These programs consume a high volume of processing cycles per execution and were run in parallel to our JADE cells.

   Three programs were selected and they were executed twice each in order to spawn a total of six processes. All programs, including JADE cells, were executed with root privileges. The programs were: (1) `find` (searches the hard disk for a file; the system takes around 17 sec to

Sophos Ex. 1009
Page 206 of 318

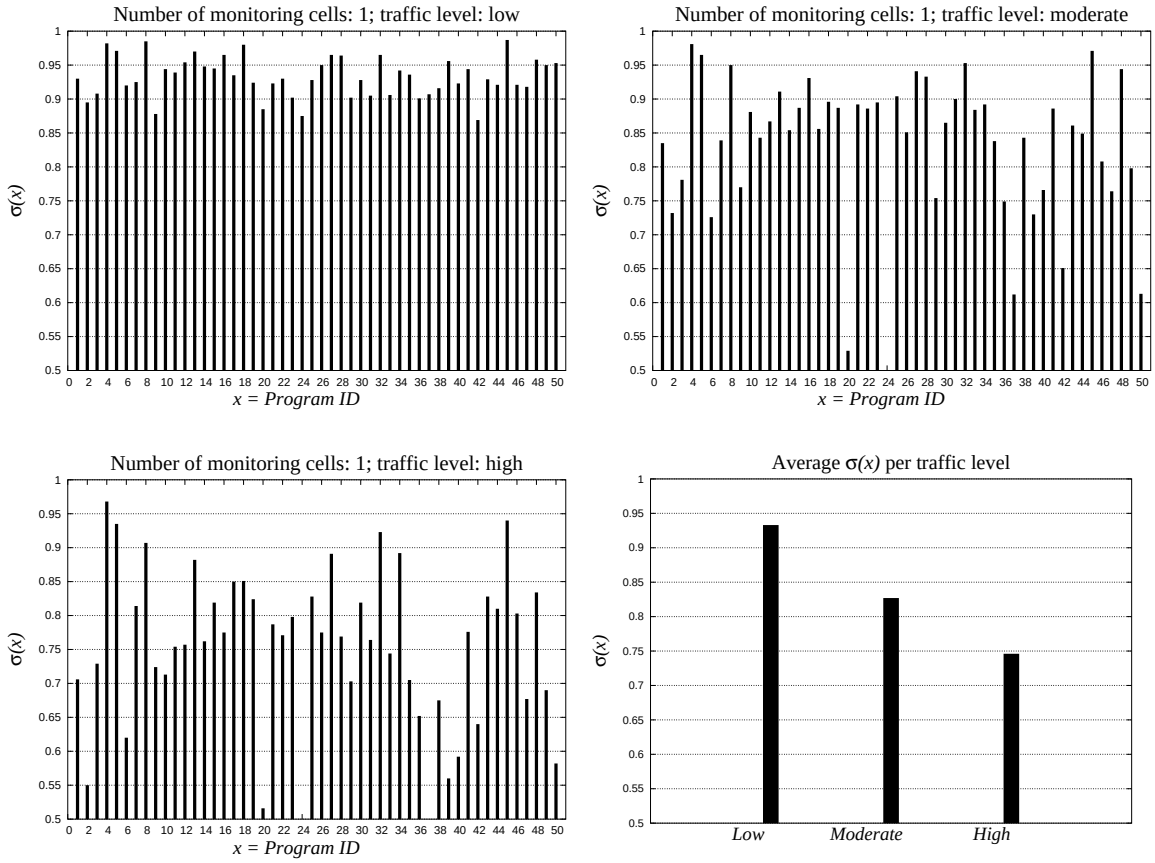


Figure 6.4: Average hypothesis score $\sigma(x)$ corresponding to each attack tool using one traffic-monitoring cell.

locate the file), (2) `top` (is a permanent loop that lists the processes executing on the server in real time), and (3) `cp` (copies a 94 Mb file from one disk partition into another; the copy takes around 12 sec to complete).

6.6.3 Cell accuracy

In order to evaluate the accuracy of the system several experiments were performed using *source system 1* and the three targets:

1. Accuracy of traffic-monitoring cells operating independently. We evaluated the ability of traffic-monitoring cells to detect attack programs in execution.

We run three experiments to include one, two, and three traffic-monitoring cells active at the same time. For each case we considered three sub-cases corresponding to each of the three traffic levels defined in previous sections.

The 50 attack programs were executed three times each. All the signatures generated during the signature generation phase were present on the signature repositories of the cells.

Sophos Ex. 1009
Page 207 of 318

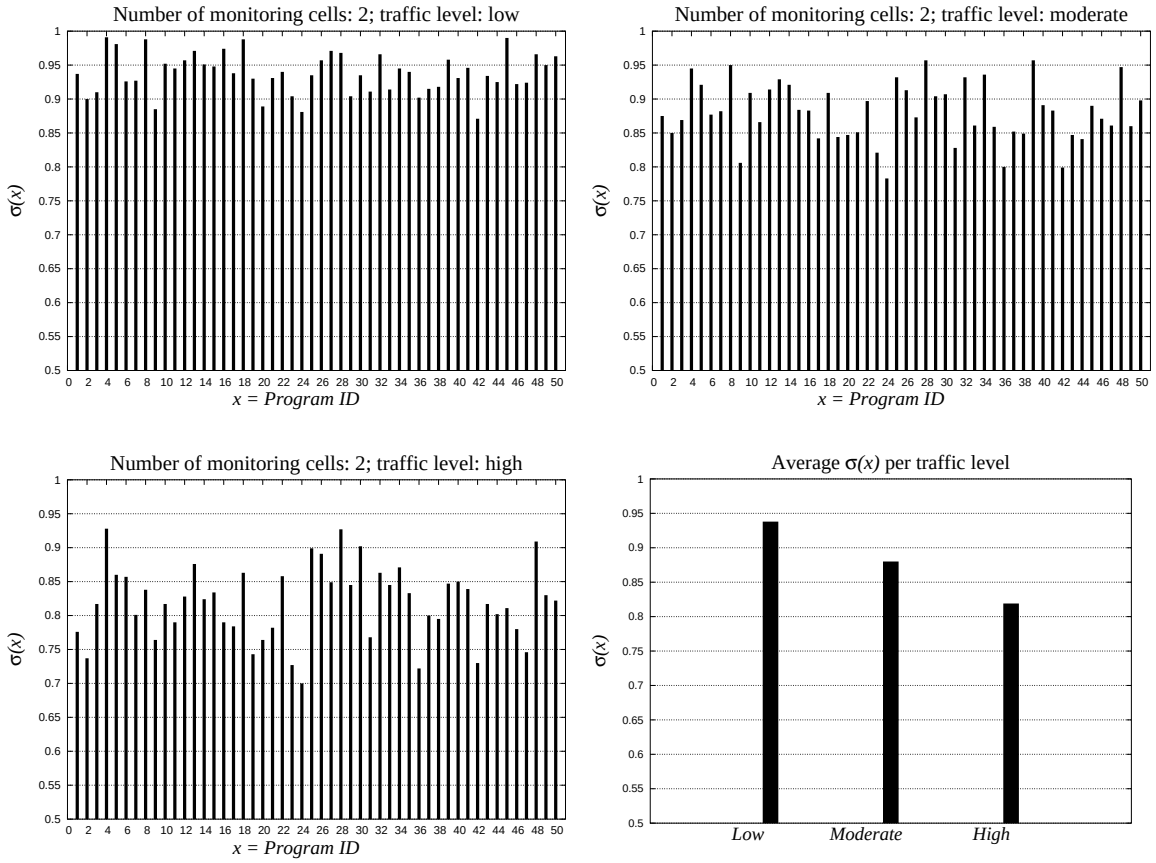


Figure 6.5: Average hypothesis score $\sigma(x)$ corresponding to each attack tool using two traffic-monitoring cells.

We computed the average hypothesis score σ from the three executions of each program. The results for one, two, and three monitoring cells are shown in Figures 6.4, 6.5, and 6.6.

The precision of the system varies with the number of active cells. When only one cell is active all matching scores lay above 0.8. However, even with moderate traffic load, the system experiences a clear decrease in its ability to positively detect the programs in execution (from an average of 0.931, scores drop to 0.827 for a 11.1% accuracy loss). A single cell has no backup peers that support its execution. A missed packet in a transfer may not be recovered and this makes the matching method produce a low score as it is unable to verify all the packets necessary to match certain signatures on the linked list. At high traffic volume, a lowest score 0.479 appears and the average σ drops 9.69% to 0.748. This average could be acceptable but the system fails to detect eight attack programs with scores below 0.6 where automated response could not safely take place.

When one more cell is added, accuracy increases. Accuracy under low traffic volume remains stable and high, but there is an improvement at moderate and high traffic volumes. When two cells are active, each one functions as the read peer of the other. Packets are evenly split between them and more resources are available for the processing. At moderate volume, an average score of 0.883 is achieved and 0.820 at high volume traffic. These represent an improvement of 6.77%

Sophos Ex. 1009
Page 208 of 318

and 9.65% for moderate and high volume traffic respectively, from the execution with only one network-monitoring cell.

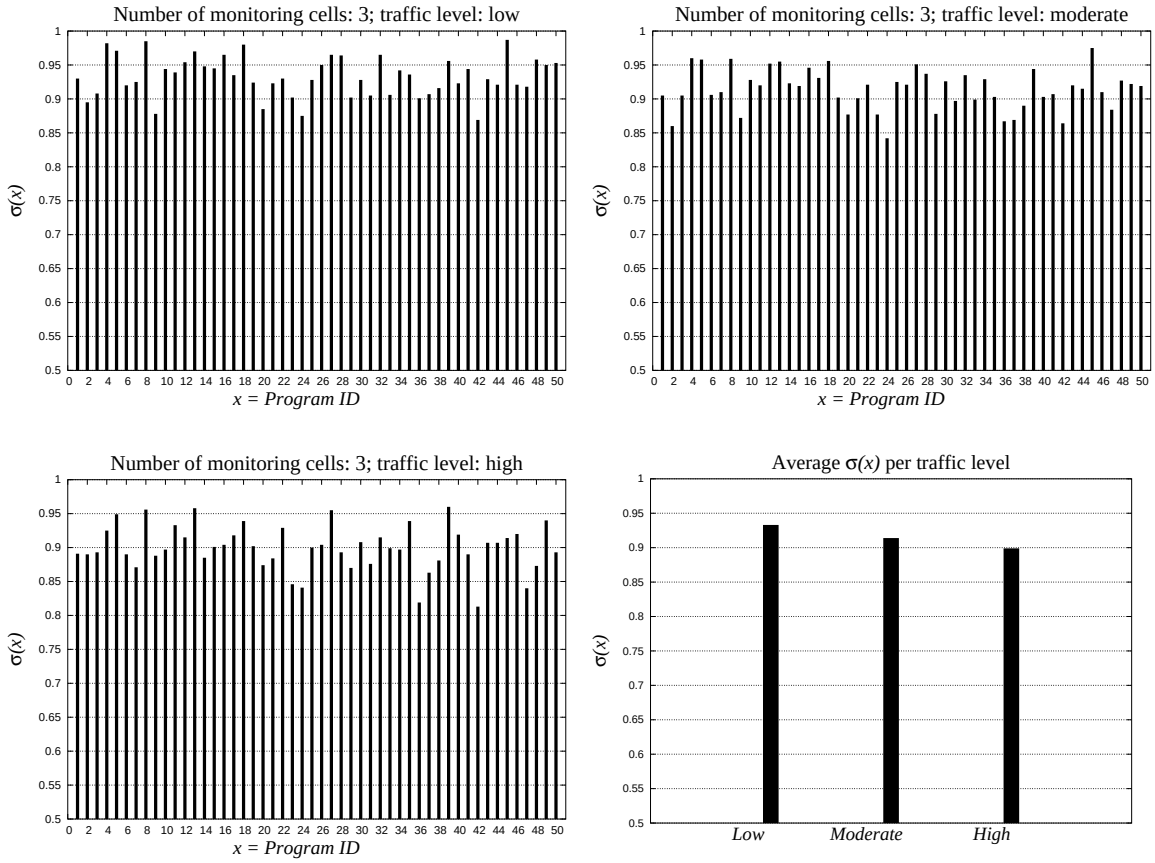


Figure 6.6: Average hypothesis score $\sigma(x)$ corresponding to each attack tool using three traffic-monitoring cells.

		Number of cells		
		1	2	3
Traffic	Low	0.931	0.941	0.933
	Moderate	0.827	0.883	0.921
	High	0.748	0.820	0.895

Table 6.4: Matching scores for one, two, and three traffic-monitoring-cell scenarios.

The execution with three monitoring cells renders the system slow but detection accuracy clearly improves. Low traffic still yields high scores and accuracy at moderate and high traffic conditions goes up to 0.921 and 0.895 respectively. This accounts for a 4.30% increase and 9.14%

Sophos Ex. 1009
Page 209 of 318

from the two-cell case, and 11.36% and 19.65% from the one-cell case for moderate and high volume traffic, respectively.

Monitoring network traffic with more than one cell gives better detection results at the cost of performance (see next section for performance evaluation results).

2. Accuracy of process-monitoring cells operating independently. We evaluated the accuracy of process-monitoring cells when they operate independently (i.e., no correlation). We considered scenarios with one and two cells. A three-cell scenario was not considered given that, unlike network packets that need to be buffered for analysis, program execution information is always placed on more stable storage (the `/proc` pseudo-filesystem, e.g.) and this helps cells that missed to capture information pieces belonging to their assigned process. The cell backup system is still in place in case case the cell is unable to monitor a process.

Each of the test cases has two variants corresponding to different workload levels. Figures 6.7 and 6.8 show average scores from three executions per program.

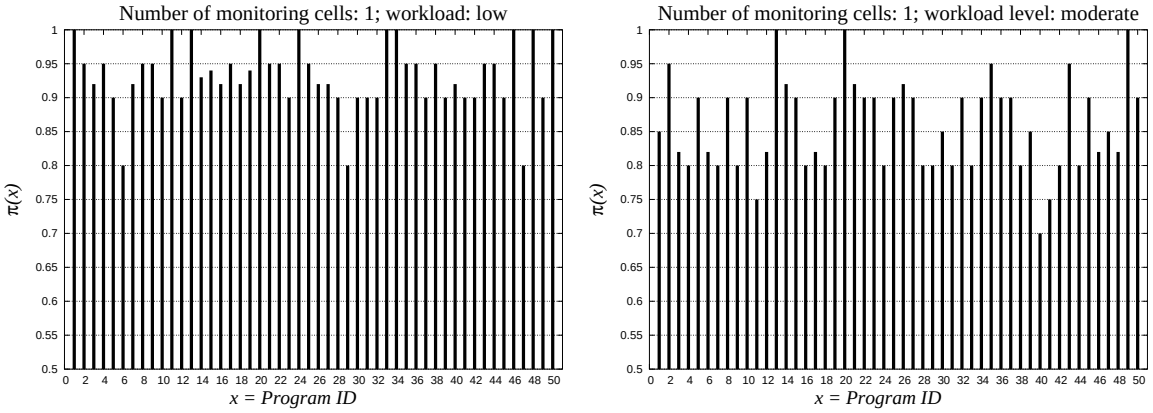


Figure 6.7: Average hypothesis score $\pi(x)$ corresponding to each attack tool using one process-monitoring cell.

When a single traffic-monitoring cell is active, detection accuracy is over 0.75 for the 50 programs when they execute under low workload (remember that scores are dependent upon the individual variants of the program). A 0.8 lowest hypothesis score is achieved for program ID's 6, 29, and 47 (see Table 6.2) and many other programs are identified with a 1.0 score. Score differences are larger than the those obtained with network cells because there are less attributes and most of them return a boolean value that is multiplied by a weight $\omega \in \{0.1, 0.2\}$.

When we execute the six processes for generating a heavier workload, accuracy decreases – although we still have a few programs scoring 1.0. If the program is not successful at compromising its target, execution will complete in a few milliseconds and its reference data on `/proc` will disappear. In that case, not all the attributes of the signature will be checked and the score will be low.

Having two process-monitoring cells does not significantly improve the accuracy of the system when operating under low workload. The fact that programs may fail to compromise

Sophos Ex. 1009
Page 210 of 318

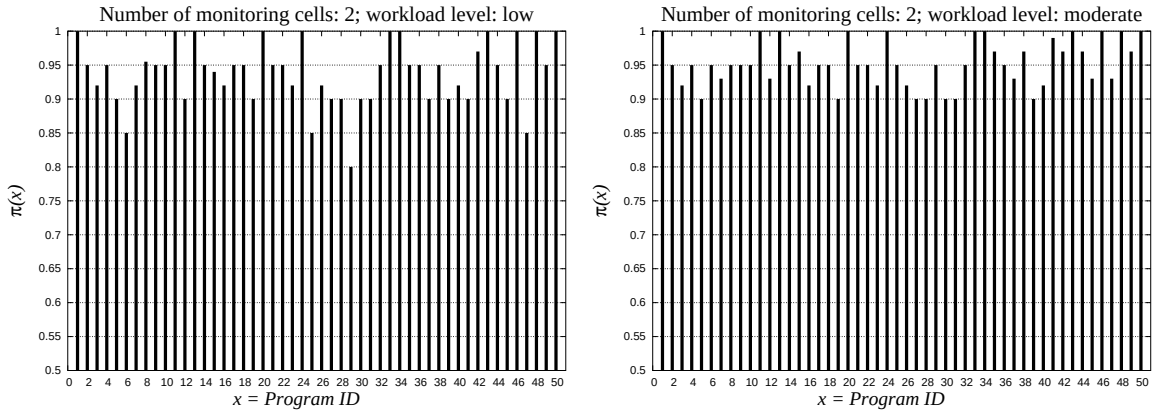


Figure 6.8: Average hypothesis score $\pi(x)$ corresponding to each attack tool using two process-monitoring cells.

a target and will therefore execute very quickly equally affects one or two cells. Two cells are able to handle moderate workload better, though. This is most likely because, while the cell to which the program was assigned is busy catching and analyzing data, the backup cell is ready to jump and take over the checking. The smallest score obtained with low workload is 0.8 whereas with two parallel cells the lowest score is much higher at 0.9.

6.6.4 Cell performance

The time FROID cells take to detect the execution of a program varies with the number of active cells, the workload at the source system, and the amount of network traffic. Average detection times from three executions of each remote automated-attack tool were computed and are shown next.

1. Detection time with network monitoring cells operating independently. Figure 6.9 contains the detection times of network monitoring cells. Each program was executed three times and an overall average of all of them was computed for the three levels of network traffic previously defined.

The execution of cells using an interpreter or, in this case, a virtual machine, is very expensive performance-wise. The benefits of portability and isolation the cell concept provides have a cost in performance. This is important as an critical aspect for intrusion detection tools is real-time detection. In a matter of milliseconds a system may be compromised and that is why fast, effective detection and correlation algorithms are necessary.

When one cell executes under low traffic conditions, detection takes 1.427 sec in average. This time is high compared to other intrusion detection systems that report 0.6 sec or less to detect an intrusion [103, 224]. A comparison like this, however, lacks validity as the intrusion detection approach is different from the one followed by OID. OID will detect programs in execution on our host. The evidence captured by cells gives the system more time to evaluate hypotheses and the fact that no packets or programs are blocked until the detection mechanism yields an outcome – as some intrusion detection tools do [8, 188, 164] – allows FROID to take

Sophos Ex. 1009
Page 211 of 318

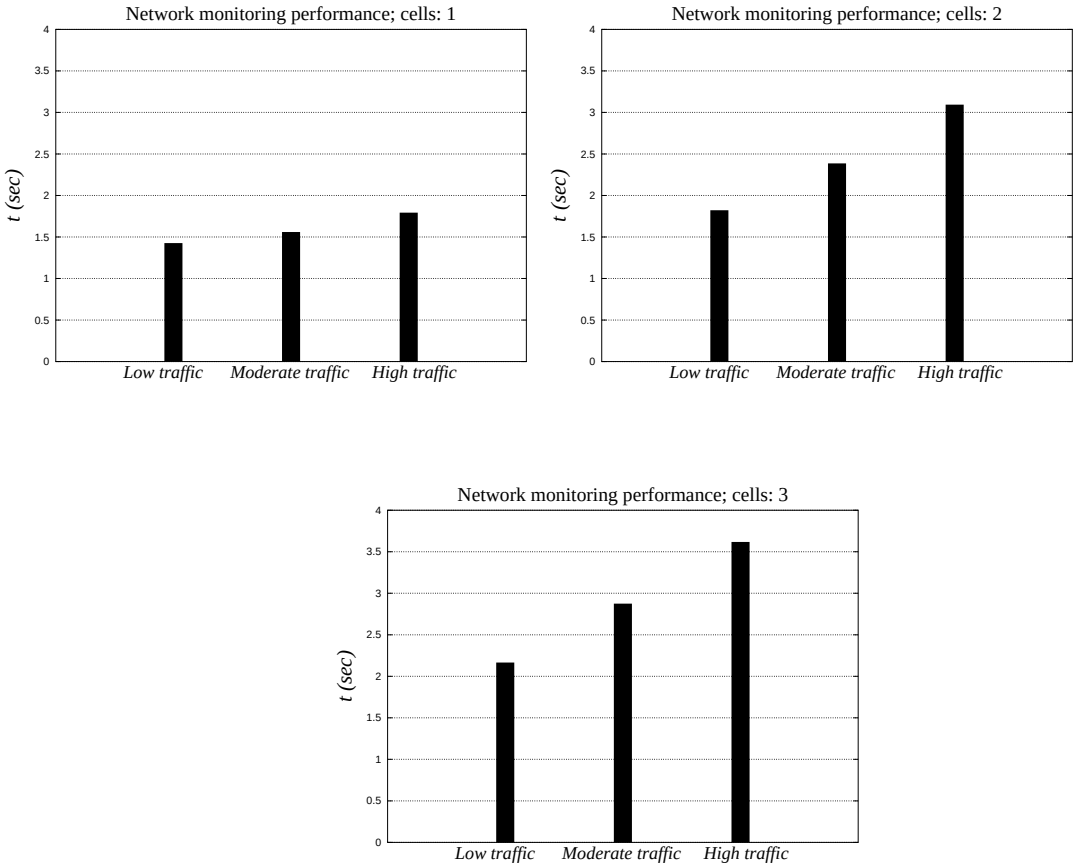


Figure 6.9: Detection time with network monitoring under different traffic levels using one, two, and three cells.

more time and still be able to identify the tools.

With low traffic, a network cell takes 1.560 sec in average to detect an attack tool and 1.797 sec when it operates under high traffic volume. These increments are acceptable. When working with low traffic, packets arrive one after the other with no major delay. Under heavier traffic conditions, the system will have to actually do filtering and this will increase detection time.

The execution of a second cell implies an additional virtual machine which means slower detection time (for accuracy figures please refer to the previous section). Time almost doubles from low traffic with one cell (1.427 sec) to high traffic with two cells (3.084 sec). The issue with such a detection time is that an attack tool or the attacker himself could have the time to remove evidence from the filesystem, memory, and even `/proc` if the security mechanism takes even few seconds to respond.

The last case with three cells considerably increases detection time to 2.166 sec with low traffic, 2.875 sec with moderate traffic, and 3.618 sec under high traffic conditions.

2. Detection time with process monitoring cells operating independently. A similar increase is observed when more than one process monitoring cell is active. Figure 6.10 contains

Sophos Ex. 1009
Page 212 of 318

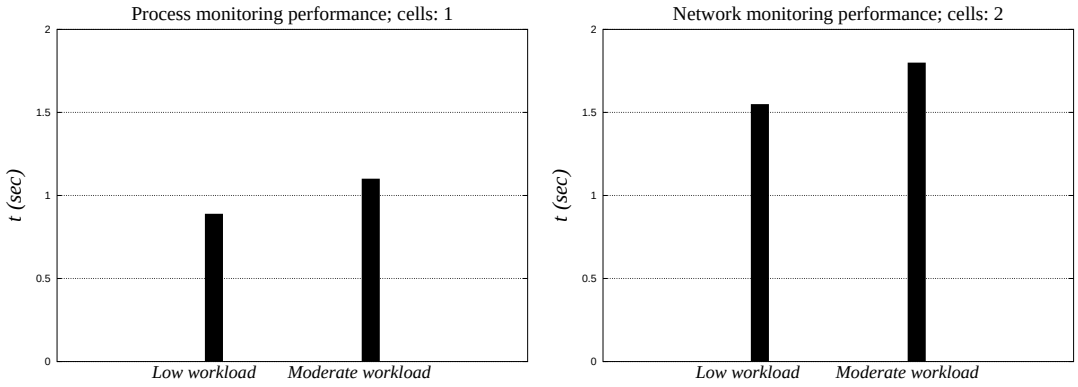


Figure 6.10: Detection time with process monitoring under different workload levels using one and two cells.

average detection times for all the 50 attack programs executed three times each and checked by one and two cells.

The impact of adding a second cell is a 79.64% increase in detection time under low workload conditions (from 0.894 sec it goes to 1.606 sec) and 60.69% under moderate workload (from 1.183 to 1.901 sec).

Agents are executed as threads within cells but their performance is not as relevant as observing the behavior of the system when multiple virtual machines are running carrying out intrusion detection tasks. This is because the proposed architecture involves software sub-environments and a way to implement them is through virtual machines and interpreters.

All the above charts expose the high cost of implementing agent cells through Java virtual machines. The constructs JADE provides are a convenient way to implement a multiagent platform but using more than one environment at a time has an prohibitively high cost that affects the performance of the source system where rogue programs are supposed to be detected.

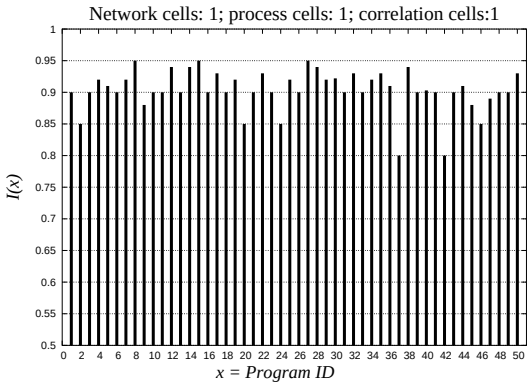


Figure 6.11: Detection accuracy using correlation with one network-monitoring cell, one process-monitoring cell, and one correlation cell.

Sophos Ex. 1009
Page 213 of 318

6.6.5 Correlation accuracy

The correlation method described in the previous chapter was tested for performance and accuracy. The latter involved combining the outcome from each monitoring cell in charge of measuring the similarity of the program's behavior with the signatures.

In order to compute hypothesis scores (I scores) two experiments were executed. The first one involved one network-monitoring cell and one process monitoring cell executing on the same source system with low traffic and workload conditions (see Sections 6.6.1 and 6.6.2 for details). The second one employed two concurrent network-monitoring cells and two process-monitoring cells working under high network traffic conditions.

1. One network-monitoring cell, one process-monitoring cell, and one correlation cell; low network traffic; low workload. All process and network signatures were uploaded into both cells. Mucus was not used and no other applications were executing on the source system. Both cells were reporting an outcome to the correlation cell.

The two cells were spawned one after the other and each of the 50 attack programs was executed once. Figure 6.11 shows the I scores reported by the correlation cell.

Since the two monitoring cells are reporting a hypothesis and a score (i.e., no cell failures) we are able to integrate the scores from the two matching systems without having to consider the cases where one of the hypothesis is missing. In order to trigger reaction, an I score must have to be at least 0.7 to suspend the process corresponding to the PID the cells received, and both σ and π scores must exist and be 0.5 or greater. The certainty that two cells are reporting the same hypothesis allows for the cell to implement automated response based on sound evidence.

Average correlation scores are smaller than network-monitoring scores but they have the big advantage of having been generated using independent entities that evaluate the input program separately.

2. Two network-monitoring cells, two process-monitoring cells, and two correlation cells; high network traffic; moderate workload. Mucus was set up to generate a high amount of traffic that kept the network interface rather busy and two instances of each of the three application selected for producing a moderately high system workload were launched. These six processes were accompanied by six FROID cells: two for network monitoring, two for process monitoring, and two for correlation. Figure 6.12 shows the average I scores computed between the two correlation cells for each attack program.

Average scores are much lower than when single cells were used. This is most likely due to the amount of traffic being generated and the time each cell has for evaluating each program. All I scores range from over 0.7 to 0.971 which is enough for triggering automated reaction.

An important point to make is that the workload and traffic conditions of this experiment really appeared as an abnormally heavy scenario at the source system that allowed us to test the functioning of the architecture. Despite of this, the accuracy level permitted the system to positively identify the programs in execution that had a signature on any of the repositories.

Sophos Ex. 1009
Page 214 of 318

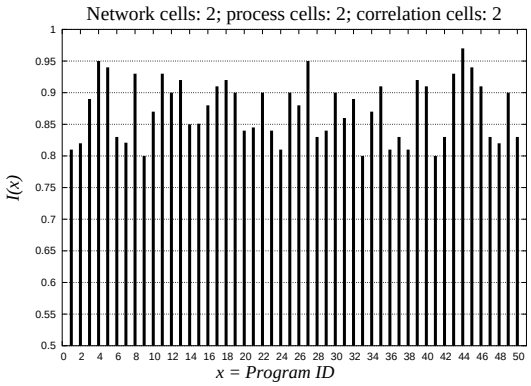


Figure 6.12: Detection accuracy using correlation with two network-monitoring cells, two process-monitoring cells, and two correlation cells.

6.6.6 Correlation performance

The two experiments involving correlation exhibited poor performance. The use of several virtual machines (up to six) proves not to be an effective way of implementing agent cells. Figure 6.13 shows the average time that correlation cells took to produce an I score for the input – this represents detection time for single executions of all the programs corresponding to the two experiment scenarios exposed in the previous section.

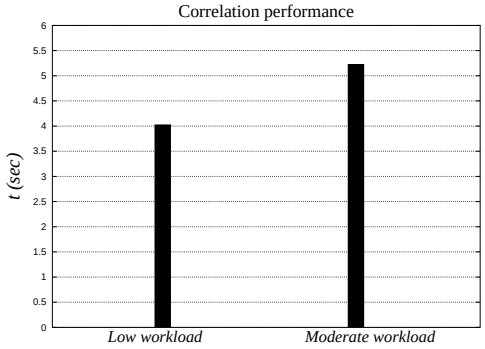


Figure 6.13: Left: average detection time using two monitoring cells and one correlation cell. Right: average detection time using two pairs of monitoring cells and two correlation cells.

A worst case of 5.23 sec detection time occurs when six cells are executed. This time does not represent an issue from the accuracy perspective, but giving a rogue program five execution seconds could have severe consequences on the target. Furthermore, attack systems could be small or big servers. A big, corporate server could probably host several instances of the FROID cells to provide continuous checking, however, a small system running the software would have trouble delivering fair performance to the user, which is one of the important challenges of intrusion detection technology.

Sophos Ex. 1009
Page 215 of 318

6.7 Interactive session tracing

In order to evaluate the efficiency of the tracing method, some experiments were prepared. The testing environment consisted of set of seven Linux machines running kernels 2.1 and above on RedHat 7.2, RedHat 7.3, RedHat 8.0 and Mandrake 8.0 operating system installations.

The general structure of the experiments was the creation of connection chains of length $m \in \{2, 4, 8, 16, 32\}$ using n hosts where $n \in \{3, 5, 7\}$. This number of test boxes was selected since our experience tells us it is not likely that, for session performance purposes, an attacker – especially beginners – can use more than seven stepping-stone hosts.

All cases where there is no loop in the chain, that is, where $m = n - 1$, are easily resolved so we prepared connections that include a series of loops along the chain. For instance, in the three-server scenario, the connection goes in circles from host h_1 into h_2, from host h_2 into host h_3, from host h_3 back into host h_1 and so on for a total of m hops. The experiments show that the algorithm effectively unwinds the loops. When an agent finds the same loop repeated several times, it will not try to solve it over and over again. Its knowledge base will contain user and connection information that allows it to conclude that a user is creating a loop.

Part of the contribution of this method is the fact that, by using loop unwinding, the performance of the algorithm has an upper threshold which is linear on the number of hosts n through which the user connects. A set of n agents will be in charge of tracing the connection and even for a very long chain, an agent located on each hosts guarantees that loops will not delay the tracing process in any way.

m	Time without encryption (sec)	Time with encryption (sec)
2	0.03	0.84
4	0.10	1.13
8	0.62	3.09
16	0.69	4.20
32	0.73	3.99

Table 6.5: Tracing time with and without encryption (RSA key exchange + DES chaining block ciphering); m corresponds to the number of connection hops or length of the chain.

Table 6.5 includes tracing times with and without encryption. Both columns indicate that the growth of the tracing time is below linear with respect to the number of hops m (this makes us think this model should escalate well to large networks). The time difference from tracing with and without encryption is considerable. This only reminds us of the performance cost of using cryptographic algorithms to protect our data. It is not likely that a connection chain extends over more than thirty hosts, for instance (latency would render the connection inoperable). It follows that, even with 1024-bit keys, the use of encrypted communication is the way to go regarding agent communication for this problem.

In all cases the system was able to locate the origin of the connection. However, no experiments were performed as to evaluate the robustness of the model once one or more agents stop

Sophos Ex. 1009
Page 216 of 318

cooperating.

6.8 Results

The results of the test phase indicate the multiagent approach to OID is feasible when using agent and ontology technologies. The performance of the FROID prototype, however, reveals clear performance limitations when more than one cell executes on the host system. Remote-attack programs can be detected by agents that characterize them as processes and also as network traffic generators. By integrating both detections an agent is able to provide a more accurate identification of the executing program. Grouping agents into teams provides the protection mechanism with increased robustness measured by the accuracy of the detection when workload and traffic conditions increase and the backup mechanism gets activated (misses remain low). The definition of program signatures allows the system to have an ontology-supported signature repository that can be deployed over the network in order to minimize exposure while speeding up reaction time.

Sophos Ex. 1009
Page 217 of 318

Chapter 5

Implementation

The First Resource for Outbound Intrusion Detection system (**FROID**) is a proof-of-concept implementation of the OID architecture proposed in previous chapters. It supports signature generation, detection, and correlation and has been developed using agent and semantic web technologies. This chapter discusses its design details and functioning. Experimental results are presented in Chapter 6.

5.1 Technologies

Multiple agent frameworks, knowledge representation languages, and automatic reasoning libraries exist. They cover specific functions that need to be combined with other tools in order to build a multiagent system capable of more intelligent behavior. The following is a brief description of the tools that were used in the implementation of this project.

- **JADE** (Java Agent DEvelopment framework) [17] is an open-source Java-based environment to build and host agent-based systems. It follows FIPA standards and contains libraries that provide different levels of communication and control functions that can be used to define the behavior of agents. JADE supports an ACL (Agent Communication Language) message exchange layer as well as a basic ontology that can be extended for specific applications.

  JADE lacks support for intelligent agent functions but, given that it is Java-based software, it can interact with other systems that do provide specialized functionality (e.g., Jena). JADE models can be distributed across several machines running different operating systems as support for agent location and mobility is also available.

  A GUI allows users to visualize agents, message exchanges, and the status of other system components. This simplifies debugging and allows for a more intuitive understanding of the agent platform.

- **Jena** is a framework for developing semantic Web applications [30]. Its API supports RDF-based languages such as DAML+OIL and OWL and contains a number of Java constructs for querying and manipulating data models. Search and navigation functions as well as set operations like union, intersection, and difference can be applied to the input knowledge base in order to locate objects and attributes or to answer other types of query.

Sophos Ex. 1009
Page 219 of 318

Its class methods for accessing RDF repositories make of Jena a valuable tool when integrating knowledge management capabilities into agents.

- **JPCAP** (Java Packet CAPture) is a network packet capture library [96]. The PCAP API that is available in C was ported to Java and became JPCAP providing almost the same functionality for packet capture and analysis. JPCAP supports several protocols through a set of Java classes that allow the user to read and modify any datagram. It also provides a graphical interface for doing traffic analysis in real time via `tcpdump` filters that taken as the basis for developing more elaborate traffic screening mechanisms.

- **Protégé** is an open-source ontology editor and knowledge acquisition system [144]. It allows the user do define an ontology, to instantiate classes, and to view knowledge bases through other applications. Protégé was written in Java and has multiple plug-ins for different ontology languages. It has been designed to guide developers and domain experts in the process of knowledge-based systems development through a set of modules that can be separately used.

  Protégé also allows the user to create, run, and save queries that return instances or object attributes from a knowledge base. It can verify instances through consistency-checking routines and identify immediate dependencies and transitivity between classes.

- **OWL** is a Web ontology language used to process information contained in documents rather than just presenting it to the user [7]. Like other ontology representation languages, OWL helps express the meaning of the objects that are part of an environment as well as represent their attributes and existing relationships. OWL provides more modeling primitives for expressing semantic contents and its design is based on lessons learned from the implementation and revision of the DAML+OIL ontology language.

5.2 Triggering event

The event that triggers the detection mechanism and the signature generation sequence is system call 117 on the Linux kernel 2.4.20-8: `opensock(int, int, char*)`. Whenever the kernel sees this type of request, monitoring cells are alerted and received the *PID* of the caller in order to begin inspection or attribute extraction.

`opensock` maintains the TCP socket table stored under the `/proc` pseudo-filesystem. A sample TCP socket table is shown in Figure 5.1. This table is permanently accessed by monitors in order to detect whether a new exchange is in progress – this includes both monitor types, network and process monitors. Each sensor agent reads up the packets it will review as well as the packets corresponding to the agents it is backing up (see description of the agent backup system below).

As mentioned in previous chapters, the FROID implementation focuses on TCP packets only as this protocol is widely used by attack and surveillance tools.

Sophos Ex. 1009
Page 220 of 318

```
sl  local_address rem_address   st tx_queue rx_queue tr tm->when retrnsmt   uid t inode
 0: 00000000:8000 00000000:0000 0A 00000000:00000000 00:00000000 00000000   29  0 1635 1 d2d8fa80 300 0 0 2 -1
 1: 0100007F:8001 00000000:0000 0A 00000000:00000000 00:00000000 00000000    0  0 2096 1 d362b080 300 0 0 2 -1
 2: 00000000:0203 00000000:0000 0A 00000000:00000000 00:00000000 00000000    0  0 2095 1 d362ba80 300 0 0 2 -1
 7: 8584110A:8108 E712020A:006E 06 00000000:00000000 03:00001639 00000000    0  0 0 2 c8c4e200
 8: 8584110A:8104 E612020A:006E 06 00000000:00000000 03:00000697 00000000    0  0 0 2 c8c4e280
 9: 8584110A:8106 EB06D894:0050 01 00000000:00000000 00:00000000 00000000  500  0 25559 1 c0c35a80 27 4 0 2 -1
10: 8584110A:8107 EB06D894:0050 01 00000000:00000000 00:00000000 00000000  500  0 25568 1 c0c35580 29 4 8 2 -1
    ...
```

Figure 5.1: Sample TCP socket table.

5.3 Agent creation

The core of the agents implemented for this project is based on JADE [17].

JADE provides agent name location and communication primitives that have been used to speed up the development of the system. Each agent cell is implemented as an JADE environment also called *container* which is in turn a Java virtual machine. Each environment executes agent-support classes that provide all basic methods for the functioning of the agent environment.

Each container executes multiple concurrent agents using threads and can communicate with other environments through message passing. An OID cell is equivalent to a single environment and contains a combination of agents as indicated in the previous chapter. Figure 5.2 shows the class inheritance structure utilized for deriving all agents:

A set of JADE *behaviours* define the functioning of each agent. *Behaviours* are routines used to group into modules the different methods an agent is able to execute depending on its state and perceptions. For instance, a $TrafficSensor$ agent contains a permanent-loop behaviour to monitor outgoing traffic as well as a timer behaviour that helps the agent determine whether it already has enough packets in the buffer in order to model a particular attack program:

```
public class TrafficSensor extends Sensor
{
   protected void setup()
   {
      addBehaviour( new LoopBehavior( this ) );
      addBehaviour( new TimerBehavior( this ) );
   }
...
```

Every agent obtains a copy of the ontology at start-up from a webserver known beforehand – no automation has been implemented to update the ontology given the static nature of the objects and relationships in the design of the OID environment; for our purpose, security is a priority over eficiency. The *model* variable stores the ontology and is used to interpret and compose messages at runtime:

```
   private final static String ontologyURL = "url:http://OIDhost/oidOntology.owl";

   public void loadOntology() {
```

Sophos Ex. 1009
Page 221 of 318

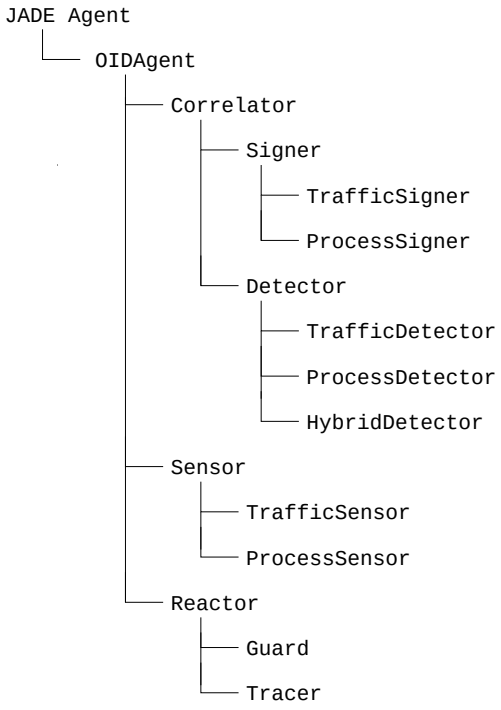


Figure 5.2: Hierarchy of classes derived from the primary JADE-agent class.

```
      model = ModelFactory.createOntologyModel( OntModelSpec.OWL_MEM, null );
      model.read( ontologyURL );
   }
...
```

5.4 Network data capture

Network sensors read outgoing network packets from the network interface via the JPCAP library [96]. There is no need for them to operate in promiscuous mode as only outgoing packets are being captured and they are available at the network device without the need of further privileges. For the FROID implementation all sensor agents use a `tcpdump` filter that screens TCP outgoing network packets:

```
filter = "proto TCP and dst not " + localAddress.getHostAddress();
```

In the above example, a variable stores a filter string in `tcpdump`-syntax and packets corresponding to the TCP protocol whose destination address is not the local IP address will be captured. Preliminary screening like this performed by network sensors is intended for correlation and reaction agents to receive only the packets that will mean something to them without having to spend time doing basic filtering by themselves.

During signature generation a sensor will capture up to 16 outgoing TCP packets. These packets come from an application that has requested a socket to send data to the outside. When an application generates less than 16 packets, the `packet counter` variable is updated as to

Sophos Ex. 1009
Page 222 of 318

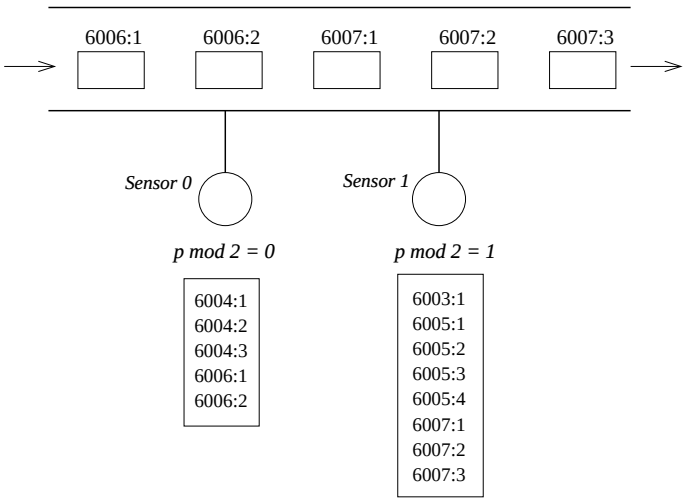


Figure 5.3: Packet reading assignment for two ($n = 2$) monitoring cells.

show this peculiarity. We have considered 16 packets in the design after having experimented with the execution of multiple attack tools. Most of them would send more than five but less than 17 packets toward their targets at launch time. In order to consider as many packets as possible without compromising performance by having to buffer a large number of packets once a socket is open, we use up to 16 packets and use a packet counter to describe how many packets the application generates.

Sensors capture packets and immediately forward them to correlation agents. Many concurrent sensors may be running and reading packets at the same time. When multiple monitoring cells are activated, monitoring is coordinated through a read backup mechanism described next.

5.4.1 Buffering and coordination

When a single sensor is in operation it may be the case that it misses some packets due to software failure or increased workload. In order to avoid this, the FROID system supports the execution of multiple, parallel monitoring cells.

Let n be the number of sensors available for network traffic monitoring. The packets each sensor sees are determined by the source port number p assigned to the socket being used. This is a consecutive integer number that increases as new sockets are requested [38]. In the case of $n = 2$, *Sensor* 0 and *Sensor* 1 will be in charge of reading up all outgoing TCP packets as a team. Modulo arithmetic is used to coordinate agents as follows:

$$p \ mod \ n \ = \ i \qquad \ldots (1)$$

where i indicates *Sensor* i will be in charge of reading the packet stream corresponding to source port p when n sensors are in operation. In Figure 5.3, *Sensor* 0 gets all packets corresponding to even port numbers whereas *Sensor* 1 is in charge of odd port numbers. As new sensors are added, n is updated accordingly so that there is always a primary agent responsible for a given socket.

Sophos Ex. 1009
Page 223 of 318

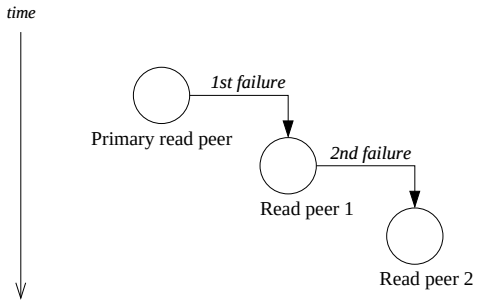


Figure 5.4: Failure cascade for sensors performing packet capture.

This decentralized structure of cells requires additional communication resources for agent location and message passing including a copy of the agent name table which, as the ontology definition, needs to be stored on every cell. In order to become operational as soon as possible, new cells are spawned from an existing cell of the same type from which they will inherit basic runtime data.

The addition of a new monitoring cells is communicated through broadcast. Although this represents an expensive communication approach to large-scale systems, the avoidance of central repositories it provides allows for increased robustness at the cost of more communication exchanges. Given the number of cells the FROID architecture is designed to support, this increased message exchanging is small (i.e., $m - 1$ where m is the number of active cells).

5.4.2 Packet read backup

A cascade mechanism allows a sensor to miss packets by designating backup agents that will support it in case of failure (Figure 5.4). The same mechanism is used for process sensors that, rather than handling network packets, extract data from processes. We now take the network sensor example to now explain the method.

Once a cell assumes responsibility for analyzing a certain socket number, it lets its read peers know so. A *read peer* is a backup monitoring cell that provides support to another similar cell in case it fails to read or scrutinize one or more packets. Each sensor is assigned two read peers making a total of three agents that will aid in the filtering of a single network transfer (if only one monitoring cell is active, no backup exists; in the case of two cells, they will function as backup of each other). Thus every outbound packet is always temporarily buffered by three different agents: the primary agent and its two read peers.

It is important to remember two design considerations made for this project: (1) the architecture does not prevent any outbound packet from passing through even when it may contain signs of attack; and (2) all buffered packets are discarded once a correlator receives them for detection.

Fault-tolerance principles suggest the use of three redundant entities to support stochastic failure [210]. Even when our entities are not redundant in the sense that they are not mirror copies of the primary agent but instead fully operational agents that provide data capture services to others, the principle still holds. For the sake of simplicity we have assumed probabilistically-independent failures in our implementation, although there may actually be

Sophos Ex. 1009
Page 224 of 318

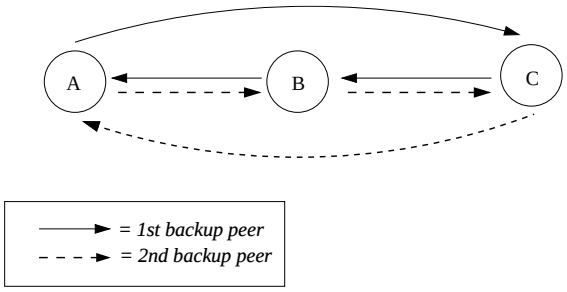


Figure 5.5: Sample backup structure for $m = 3$: when A is a primary agent, B is its first backup and C is its second backup; when B is a primary agent, C is its first backup and A is its second backup; when C is a primary agent, A is its first backup and B its its second backup.

some correlation between them (see Section 7.3 in the last chapter).

In case an agent fails to notify its intention to evaluate a transfer, its peers will not receive any message indicating the socket is already being taken care of. If a read peer receives such packet and has not received confirmation from the primary agent after two seconds, it will proceed to do the filtering and its cell will do the matching as it owns a full copy of the packets (Figure 5.5). The agent will announce this to the other read peer and, if it is itself processing a packet sequence, it will first finish with it before proceeding with its backup work. The second read peer is activated when it has received a packet corresponding to a primary agent but does not receives confirmation from the first neither the second agent after four seconds (two seconds for the primary agent failing to report, and two seconds for the second peer failing to report).

Only one pending transfer is allowed per backup peer (the agent works as first or second read peer only; cannot be both at the same time). A peer supports a primary agent when it fails but it is unable to support a second failing agent when it is already performing any backup function. At any point, an agent can only be assigned a maximum of three roles: primary agent, first read backup peer, and second read backup peer.

Similarly, once the primary agent is done doing signature matching with a stream of packets, it must announce this to its two read peers. The reason behind this consideration is the case in which an agent experiences trouble after it announced it will be processing a certain transfer but cannot complete the task. In that situation the same cascade procedure applies. The first backup peer will take over the operation and the second backup peer will do so in case the other two fail to report the matching. When an agent cannot report due to communication and not operational problems, it will complete its task but will not let the other agents know about it. In this case a packet stream will be checked more than once but that is more desirable than missing to do the checking.

5.5 Process data capture

Process data is derived from two main sources: (1) executable files and (2) the `/proc` pseudo-filesystem. Once a program has been identified as the originator of an outgoing packet sequence, the executable is used to extract the attribute values that will go into the process signature.

Sophos Ex. 1009
Page 225 of 318

This information is first extracted from the executable file itself and confirmed with the data stored on the `/proc` file system.

The MD5 message digest of the text segment of the file needs to be computed in the first place. The text segment is extracted with `size`, then converted into a `String` and then passed to an instance of the MD5 class (no key is used with MD5 to create a MAC [179]; signatures must be verifiable by anyone):

```
import org.w3c.tools.crypt.*;
...
     Md5 txtSegmentMd5 = new Md5(txtSegment);
     String digestMd5 = txtSegmentMd5.getStringDigest();
```

The size of the text and data segments are obtained through `size` and then verified under the `/proc` directory corresponding to the `PID` of the process. If there is a discrepancy between the two values, the one stored under `/proc` is taken first as it reflects what the system is seeing. A discrepancy between the two reduces the weight of the attribute when matching is applied (see the signature matching section ahead).

The possible arguments to the program are then stored on the signature including the minimum and maximum number of arguments that can appear in the command line when the program is invoked. At detection time, this information is extracted from `/proc` as well.

The size of the executable, a strong attribute when other attributes match but a weak attribute by itself, is also recorded using the actual file of the attack tool which is located by the *Signature generation* and the *Monitoring* agents using the execution directory and file name derived from the PID.

5.6 Matching of network signatures

An attack tool can be executed using different arguments in the command line. Each combination of arguments makes the program behave differently by following a certain path inside the text segment and by varying the number and contents of the network packets sent to the target.

We took a test set of attack tools and executed them using different combination of parameters. Each combination or *variant* was executed three times so that process and network signatures could be generated for each variant. The objective was to evaluate the similarity of signatures corresponding to the same attack tool when using different parameters, and also to evaluate the similarity of signatures corresponding to different attack tools.

In the case of network traffic, we executed the same variant three times. For instance, we took attack tool A_{10} and compared the signatures produced by every execution. They were consistent except for hash values.

Hash values are computed by adding up the payload bytes of a packet. These values are checksums of the payload whose length is limited with $mod\ 10E6$ in order to preserve a fixed checksum size of six decimal digits, thus we refer to them as hash values.

The payload of the packet is not consistent from execution to execution as TCP wraps other packet types that contain variable fields. Such is the case of `telnet`, where the size of the payload is the same but its contents vary due to session negotiation flags. For that reason, a *deviation* field to represent the variation of the hash value due to sub-protocol flags is used. In

Sophos Ex. 1009
Page 226 of 318

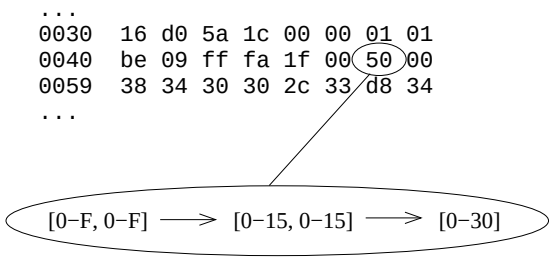


Figure 5.6: Payload data of a `telnet` packet. One byte is used for sub-protocol window size negotiation. The maximum value corresponds to FF which produces a deviation between 0 and 30 in the hash value.

the case of `telnet`, for instance, the deviation is an integer $d_{telnet} \in [0, 30]$. Figure 5.6 depicts this variation.

Network signatures corresponding to executions of the same variant are very similar. The *deviation* field helps capture the possible variations that may arise when transporting subprotocols.

We compared signatures corresponding to one single variant of different attack programs. Figure 5.7 is a representation of the signatures of three different attack programs using three attributes. Executions of the same program remain close to each other as explained before. The signatures corresponding to different programs are easily distinguishable as they form clusters that reflect the fact that the programs behave differently.

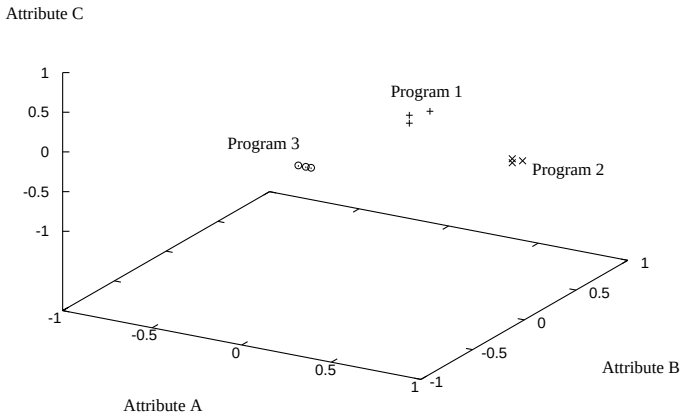


Figure 5.7: Network-traffic signatures of three different attack programs considering executions of *only one variant* per program.

Sophos Ex. 1009
Page 227 of 318

When different variants for each program are exercised, the situation is different. They are not easily distinguishable for the same program and not even for the same variant. In Figure 5.8, for instance, a variant of *Program* 2 appears separate from the rest. However, the signatures of a second variant of *Program* 2 appear close to the cluster of a variant of *Program* 3. Even in a 3D representation, it is not easy to determine a single cluster that groups all the different variants of the same program. If there are over 100 programs and more than one variant per program, the situation is more complicated.

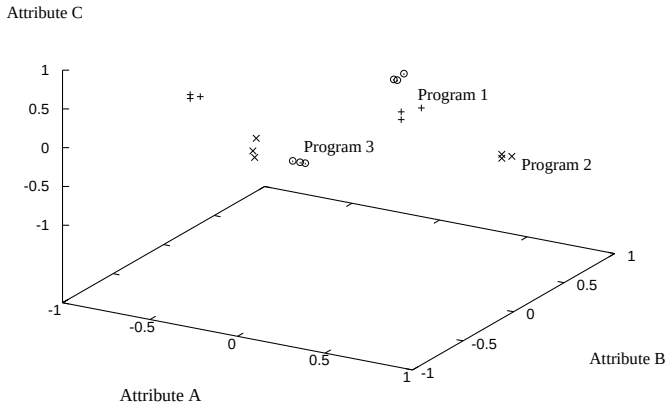


Figure 5.8: Network-traffic signatures of three different attack programs considering *two variants*.

Consequently, our design contemplates program-variant signatures in order to capture the specifics of the of network traffic a program generates when using a different set of arguments rather than creating a single signature to group all variants together.

5.6.1 Primary data structure

Our network-signature matching algorithm is based on the work of Roesch [170] which was enhanced by Coit et al. in [37] using trees. We have modified the basic algorithm by considering the amount of information each field conveys. The idea is to reduce the time an agent spends matching its perceptions against its knowledge base in order to speed up response time.

The algorithm used by Snort groups attributes into two main node categories: (1)high-gain nodes and (2)low-gain nodes – the categorization is based on the experience of the authors in monitoring network activity. Our approach is to develop this categorization using experimental results from entropy and splitting power analysis applied to attack tool traffic so to use a data

Sophos Ex. 1009
Page 228 of 318

structure that takes into consideration the different weights of its attributes and the frequency of its values – since read requests are much more frequent than read requests in this type of pattern-matching systems, it is possible to support a permanently-sorted data structure.

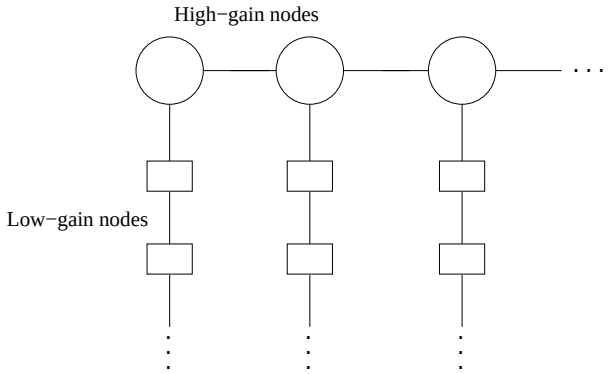


Figure 5.9: Snort's signature matching linked list [170, 37].

Figure 5.9 shows the bi-dimensional data structure used to store signatures in Snort. High-gain nodes contain a small group of heavy attributes. They are checked first for a match and, if there is one, the matching proceeds down the linked list where low-gain nodes are. These nodes contain the rest of the attributes that eventually lead to a classification of the input and each of them represents a signature. Once a new attack pattern needs to be added, new low-gain nodes are added and high-gain nodes updated.

Given the number of attributes we have, the data organization we used in the FROID implementation considers a third node type. In Figure 5.10 different node sizes reflect the fact that some nodes are more frequent than others. Node α_1, for instance, contains values that are common to a larger number of signatures than the values stored by node $alpha_3$. Nodes are placed in the linked list according to its weight, ω_i. Even though all α nodes contain the same attributes, the values they store give more relevance to some of them (i.e., if a certain combination of TCP flags helps easily classify packets, its splitting power will likely be high and it will be part of an *alpha* node at the front of the list).

β nodes are also sorted by frequency but their attributes contain attributes with less splitting power than α-node attributes. Unlike Snort, β nodes does not represent individual signatures. They group attribute values common to several signatures. If a match with an α node and a match with a β node have been found, the algorithm now tries to match γ nodes with the input. γ nodes do represent actual signatures and they are added to the corresponding β node list once a new signature needs to be considered. Important to recall that this method classifies packet sequences and not individual packets as Snort do.

In order to determine the order of attributes to build our three-dimensional linked list we have analyzed data obtained from attack tools.

Sophos Ex. 1009
Page 229 of 318

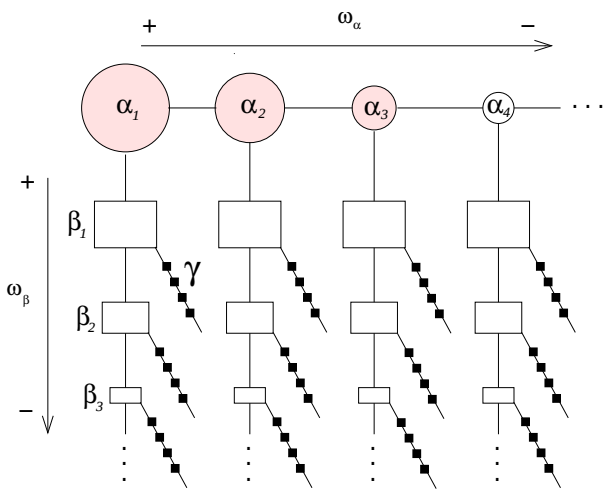


Figure 5.10: Data structure for attack signatures. Nodes are sorted from highest to lowest frequency.

5.6.2 Splitting power analysis

There are three splitting functions commonly used in classification trees [130]. These functions determine the amount of information a certain attribute conveys so that the most compact tree can be created.

Regardless of whether the tree creation technique is a greedy or multi-pass refinement one, it can be used to extract information about the weight of attributes for classification purposes. The three functions we used to compute the relevance of each attribute classification-wise are:

- Entropy. Represents the expected number of bits needed to encode the class of a randomly drawn sample using information theory coding:

$$Entropy(S_i) = \sum_{v \in Values(S_i)}^{k} -p(v)\ log_k(p(v)) \qquad \ldots (2)$$

- Information Gain. Computes the reduction in entropy due to the splitting on an attribute).

$$InfoGain(S_i) = Entropy(S_i) - \sum_{v \in Values(S_i)} \frac{\mid S_{i,v} \mid}{\mid S_i \mid} Entropy(S_{i,v}) \qquad \ldots (3)$$

- Gini index. Measures the impurity of a node considering how many samples go though each of its branches.

$$GINI(S_i) = 1 - \sum_{v \in Values(S_i)} p(v)^2 \qquad \ldots (4)$$

Sophos Ex. 1009
Page 230 of 318

where S_i represents attribute i; $v \in Values(S_i)$ is a class for attribute i (continuous values are discretized in order to define a finite number of attribute classes); $|S_i|$ is the cardinality of the attribute defined by the number of input samples it receives; $|S_{i,v}|$ is the cardinality of each attribute class defined by the number of input samples that go into that class; and $p(v)$ is the probability of a random input sample being classified in class v as per the evaluation of its i-th attribute.

Let σ_j represent a splitting function. Different σ_j functions generate different trees, however, they all identify the most relevant attributes so that they can be placed at a shallow depth in the tree hereby helping more rapidly predict the classification of an input sample and reducing, at the same time, the tree's depth.

The use of more than one splitting function allows to reduce the selection bias of each technique and to identify high-value attributes with increased accuracy [130]. In order to integrate the three entropy values we use the following expression with domain in $\mathbb{R}$ and codomain in the $[0, 1]$ closed interval:

$$\lambda_{\sigma_j}(S_i) = \frac{\mid \sigma_j(S_i) \mid}{Max_{\sigma_j} - Min_{\sigma_j}} \quad \ldots (5)$$

where S_i represents the i-th attribute, σ_j the j-th splitting function, $j \in \{1, 2, 3, ..., k\}$, k is the number of splitting functions being used, and Max_j and Min_j are the maximum and minimum values generated by the i-th function, respectively. Notice that the functions we use are either asymptotic or have an absolute maximum. Their values can be normalized to 1 so that they can be integrated into a final weight value Λ:

$$A(S_i) = (\sum_{j=1}^{k} \sigma_j(S_i))/k \quad \ldots (6)$$

5.6.3 Analysis of network data

The decision trees produced with each function appear in Appendix D. They were generated using the decision tree module of the Gnome Data Miner tool [69] which implements the selected splitting functions. Table 5.4 shows the splitting-power values obtained for each of the attributes (Figure 5.11 plots the averages computed from the three splitting functions).

α ($A > 0.8$)	β ($A > 0.5$)	γ ($A \leq 0.5$)
fragmentation	destinationPort1	l7
windowSize	f5	l11
packetCounter	f6	h9
l4	sourceAddress	f16
h4	f8	f13
h5	sourcePort	h16
h6	h7	f12
deviation	h8	f15

Sophos Ex. 1009
Page 231 of 318

l5	destinationAddress	f14
l6	f4	h12
f7	destinationPort2	l12
	f9	h12
	f10	l15
	f3	l10
	h14	f2
	h11	h10
	l9	l16
	l8	h2
	l14	h13
	f11	h3
		l2
		l16
		l3
		l13
		h15
		f14

Table 5.2: Attributes stored on each node type.

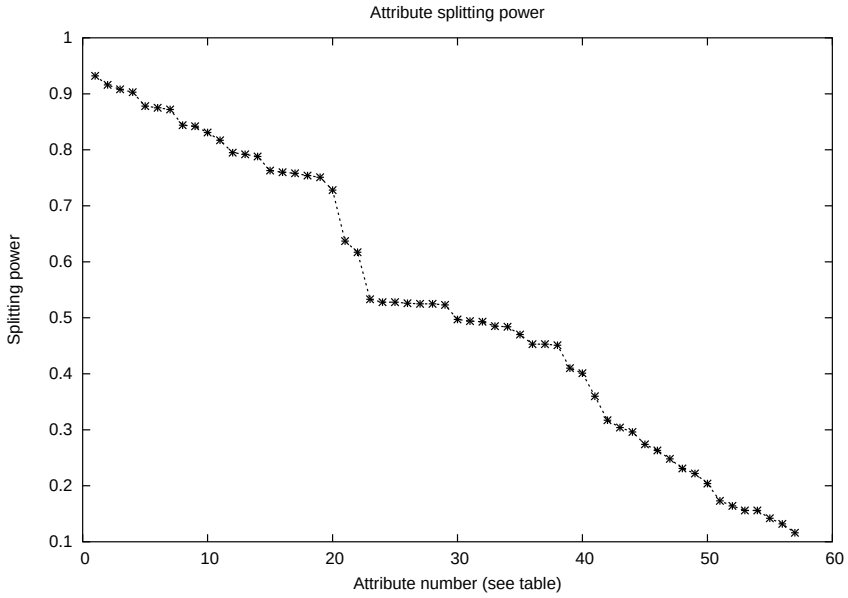


Figure 5.11: Splitting power per signature attribute.

	Attribute	*Entropy*	*InfoGain*	*GINI*	*Average* (*A*)
1	fragmentation	0.972	0.926	0.899	0.932
2	windowSize	0.902	0.922	0.923	0.916
3	packetCounter	0.874	0.886	0.965	0.908
4	l4	0.922	0.934	0.852	0.903
5	h4	0.920	0.822	0.893	0.878
6	h5	0.846	0.920	0.860	0.875
7	h6	0.824	0.924	0.869	0.872
8	deviation	0.869	0.899	0.764	0.844
9	l5	0.868	0.892	0.765	0.842
10	l6	0.837	0.829	0.826	0.831
11	f7	0.712	0.823	0.917	0.817
12	dport1	0.867	0.723	0.794	0.795
13	f5	0.768	0.783	0.825	0.792
14	f6	0.790	0.823	0.751	0.788
15	sAddress	0.817	0.725	0.748	0.763
16	f8	0.865	0.691	0.723	0.760
17	sPort	0.786	0.823	0.666	0.758
18	h7	0.743	0.725	0.794	0.754
19	h8	0.734	0.726	0.792	0.751
20	dAddress	0.753	0.689	0.743	0.728
21	f4	0.530	0.699	0.681	0.637
22	dport2	0.526	0.491	0.834	0.617
23	f9	0.634	0.509	0.455	0.533
24	f10	0.578	0.509	0.497	0.528
25	f3	0.578	0.483	0.523	0.528
26	h14	0.426	0.529	0.623	0.526
27	h11	0.590	0.560	0.425	0.525
28	l9	0.574	0.523	0.478	0.525
29	l8	0.593	0.447	0.530	0.523
30	l14	0.547	0.519	0.425	0.497
31	f11	0.535	0.425	0.523	0.494
32	l7	0.498	0.490	0.491	0.493
33	l11	0.425	0.500	0.530	0.485
34	h9	0.547	0.405	0.501	0.484
35	f16	0.464	0.426	0.521	0.470
36	f13	0.463	0.397	0.498	0.453
37	h16	0.434	0.225	0.699	0.453
38	f12	0.408	0.550	0.396	0.451
39	f15	0.400	0.356	0.474	0.410
40	f14	0.446	0.534	0.223	0.401
41	h12	0.385	0.285	0.409	0.360
42	l12	0.427	0.252	0.271	0.317
43	h1	0.366	0.489	0.057	0.304
44	l15	0.377	0.303	0.208	0.296

Sophos Ex. 1009
Page 233 of 318

45	l10	0.328	0.125	0.370	0.274
46	f2	0.355	0.216	0.217	0.263
47	h10	0.019	0.037	0.688	0.248
48	l16	0.367	0.217	0.110	0.231
49	h2	0.185	0.293	0.188	0.222
50	h13	0.207	0.058	0.346	0.204
51	h3	0.186	0.205	0.129	0.173
52	l2	0.102	0.197	0.193	0.164
53	l1	0.157	0.190	0.123	0.156
54	l3	0.144	0.194	0.131	0.156
55	l13	0.128	0.144	0.155	0.142
56	h15	0.107	0.166	0.124	0.132
57	f1	0.132	0.108	0.108	0.116

Table 5.4: Splitting power per attribute in network signatures.

According to the results, we have selected the attributes with splitting power 0.8 or above to go into α_i nodes. Attributes whose splitting power is 0.5 or greater are stored on β nodes, and attributes with splitting power less than 0.4 are stored on γ nodes. Table 5.2 condenses these selection criteria.

5.6.4 Matching method for network signatures

The matching method is a score-based system. An overall score *sigma* represents the similarity between the input packets and a determined signature. This score is composed of a linear combination of sub-scores $sigma_i$ that correspond to a match between the i-th attribute of the signature and the input.

Each attribute contributes differently to σ scores (Figure 5.12). The weight ω_i of an attribute is determined by its splitting power which is also indicative of its relevance as a similarity measure:

$$\omega_i = ln(k) - ln(i) \qquad \ldots (7)$$

where k is the total number of attributes. Weights are combined into σ and then normalized to generate an output in the range $[0, 1]$:

$$\sigma \;=\; \sum_{i=1}^{k} \frac{\sigma_i \omega_i}{\sum_{j=1}^{k} \omega_i} \qquad \ldots (8)$$

All σ_i sub-scores return a value in the $[0, 1]$ range as well. In fact, all attributes return a boolean value, $\sigma_i \in 0, 1$ except for the *hash* attributes which are not checked for identical match as the others. As mentioned earlier, TCP sub-protocols may change the payload of a sequence and that is why the *deviation* attribute is used.

We use a triangular function to evaluate similarity of hash values. The maximum of the function corresponds to a match of hash values with zero deviation (i.e., perfect match; Figure

Sophos Ex. 1009
Page 234 of 318

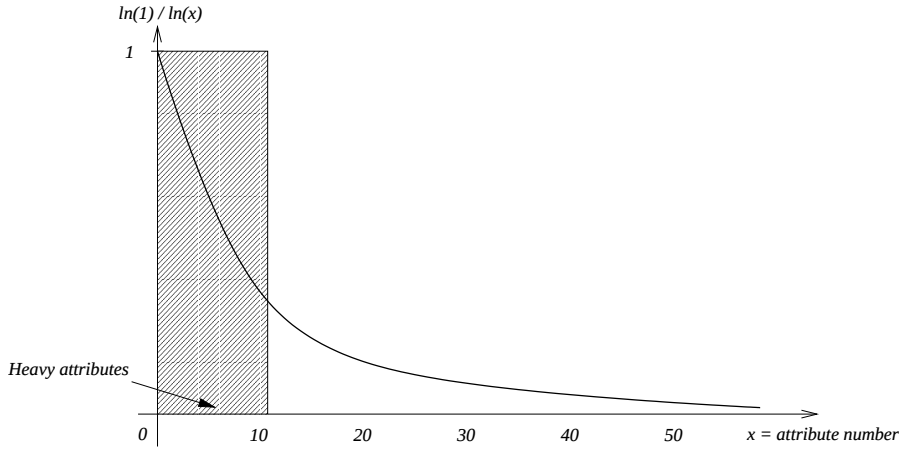


Figure 5.12: The matching method gives more weight to attributes with high splitting power (see Table 5.4).

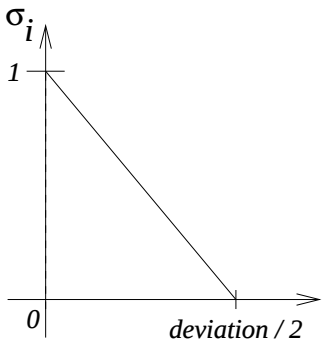


Figure 5.13: Function to compute payload hash σ_i sub-scores.

5.13). Any difference between the signature and input hash will be penalized with a lower σ_i as the deviation increases. Those matches whose deviation is greater than the one stored on the signature will get a zero similarity value.

The value of σ_i for a hash attribute is computed as follows:

$$\sigma_i = 1 - \frac{|h_i - h_{input}|}{deviation} \quad \ldots (9)$$

where h_{input} is the hash of the input packet's payload, h_i represents the i-th hash in the signature, and *deviation* is the deviation attribute.

All other attributes are checked for exact matches. IP addresses, port numbers, window size, etc. are features for which different degrees of similarity do not work properly. If, for instance, a signature corresponds to an attack program that always targets a determined port number and the input uses a port number close to the port number in the signature (for instance,

Sophos Ex. 1009
Page 235 of 318

port 22 and port 23), this does not mean the input in fact resembles the signature [103]. Different protocols may be run on different port numbers and so any similarity between signatures would be merely coincidental.

These attributes yield a 0 or a 1 that will enable or disable the attribute's weight in the calculation of σ. When the search reaches γ nodes, the input will be matched against multiple signatures under a determined β_i node. The highest σ score will be picked to denote to which signature the input resembles the most. The first-fit flaw [1] some intrusion detection systems suffer is avoided by comparing the input against a number of signatures and taking the closest match.

The following pseudocode provides a high-level perspective of how the matching algorithm for network signatures works:

```
# Network signature matching.  High-level pseudocode.
IF input MATCHES α_i node THEN
    ACCESS α_i list
    IF input MATCHES β_j node THEN
        ACCESS β_j list
        max = 0
        FOR all γ_k nodes IN β_j DO
            IF MATCH(input, γ_k) > MATCH(input, γ_max_i THEN
                max = k
            END_IF
        END_FOR
        PRINT The input corresponds to signature:  γ_max
    END_IF
ELSE
    PRINT The transfer is safe
END_IF
```

5.7 Matching process signatures

In the case of process signatures, we removed the *dynamic libraries* field considered in the design phase after reviewing the data. The number of dynamic libraries attack programs use is very small. This is perhaps due to the fact that these programs need to be portable and self-contained. A way of reducing platform dependency is to minimize the number of system libraries a program links at run time (malicious dynamic libraries are not included into the scope of this dissertation, however, they are used by some tools to compromise vulnerable systems [225, 154]). We found that more than 50% of the surveyed programs use at most two dynamic libraries. From the program modeling perspective this does not convey meaningful information to the signature and so has been omitted in the implementation.

The rest of the fields has been integrated into the signature. The text segment of Linux

[1] First-fit flaw. Returning the first match encountered during the search without considering subsequent, possibly better matches.

Sophos Ex. 1009
Page 236 of 318

```
smurf4:     file format elf32-i386

Contents of section .text:

804885c 31ed5e89 e183e4f0 50545268 008f0408  1.^.....PTRh....
804886c 68d08e04 08515668 0c890408 e8fffeff  h....QVh........
804887c fff49090 5589e553 50e80000 00005b81  ....U..SP.....[.
804888c c3aa1a00 008b8378 00000085 c07402ff  .......x.....t..
804889c d08b5dfc c9c39090 5589e583 ec08803d  ..]......U......=
80488ac b8a30408 007529a1 08a20408 8b1085d2  .....u).........
80488bc 741789f6 83c004a3 08a20408 ffd2a108  t...............
80488cc a204088b 1085d275 ebc605b8 a3040801  .......u........
80488dc c9c389f6 5589e583 ec08a130 a3040885  ....U......0....
80488ec c07419b8 00000000 85c07410 83ec0c68  .t........t....h
80488fc 30a30408 e8fb76fb f783c410 c9c39090  0.....v.........
804890c 5589e556 5383c480 83e4f0b8 00000000  U..VS...........
804891c 29c4c745 c4000000 0083ec0c 68002000  )..E........h. .
804892c 00e8eafd ffff83c4 10894594 e8a60300  ..........E.....
804893c 0083ec08 683f8e04 086a02e8 a0fdffff  ....h?...j......
804894c 83c41083 7d08057f 1083ec0c 8b450cff  ....}........E..
804895c 30e89903 000083c4 1083ec0c 8b450c83  0............E..
...
```

Figure 5.14: The text segment can be used to characterize a program. The figure shows an excerpt from the text segment of `smurf4`, one of the surveyed attack tools.

programs was chosen over the use of the heap and stack segments given its unchanging contents. Executable code is a valuable element for characterizing programs that is readily available with executables [211] (Figure 5.14). A change on this segment implies recompilation of the source program which can be limited by removing compilers from the system and enabling filesystem integrity checkers.

The number and order of arguments passed to a program are distinctive features highly dependent on the coding habits and preferences of the developer. File size is an attribute that completes the signature. It is used by integrity checkers in order to detect changes to critical files. To us it serves the purpose of confirming a hypothesis regarding signature matching.

When put together, all these information pieces help identify the type of program that is being executed. Even when some of these attributes can individually reveal the presence of an attack tool, there exist programs that remove themselves from the filesystem after execution [125] hereby eliminating the possibility of counting on the executable file resource to identify the attack program. It is for that reason that multiple attributes are considered in order to more effectively tell whether a certain rogue process is exploiting or surveying a system in the outside.

5.7.1 Primary data structure

For the matching of process signatures we use a linked list sorted on the *text segment size* attribute. The number of attributes composing process signatures is significantly less than the attributes used for network signatures. Each node in the process signature linked list represents an actual signature – notice that signatures corresponding to the same program may appear in nodes that do not share a link (Figure 5.15).

This list is permanently sorted and the matching algorithm checks first the node that has the closest text segment size to the input. It then proceeds to match the input against the predecessor and the successor of such node as explained in the next section.

Sophos Ex. 1009
Page 237 of 318

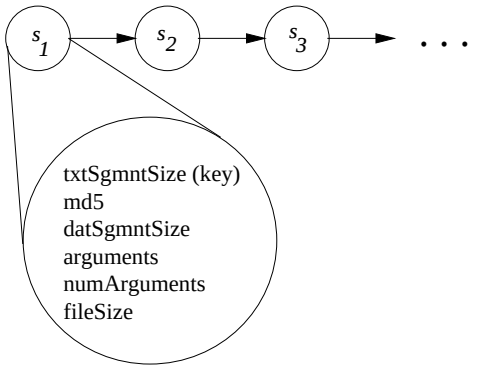


Figure 5.15: Process signature linked list. Every s_i node represents a program-variant signature.

The backup system used for network monitoring cells is also used for monitoring processes. Each process sensor announces to two backup peers – assigned in a circular fashion – its intention to analyze a certain program in execution. Unlike network sensors, no data capture is necessary. The only reference the monitor needs is the PID of the process. All other attributes can be derived from this reference. Should a sensor fail to report to its peers, the cascade sequence will activate one of its two peers which will take over the analysis of the program (see Figure 5.4).

5.7.2 Matching method for process signatures

A score system is also used for process signatures. An overall score π captures the similarity of the input with a signature node. Each signature attribute has a different weight ω_i that multiplies attribute sub-scores π_i:

$$\pi = \sum_{i=1}^{k} \pi_i \omega_i \qquad \ldots (10)$$

The following weights have been assigned to each attribute [184]:

- MD5: $\omega_1 = 0.2$. The matching function for the MD5 attribute returns a value $r_1 \in \{0, 1\}$ where 1 corresponds to an exact match and 0 to different hash values.
- Text segment size: $\omega_2 = 0.2$. The matching function returns a value $r_2 \in \{0, 1\}$ where 1 corresponds to an exact match and 0 to different text segments.
- File size: $\omega_3 = 0.2$. The matching function returns a value $r_3 \in \{0, 1\}$ where 1 corresponds to an exact match and 0 to different file sizes.
- Arguments: $\omega_4 = 0.2$. Depending on the number of arguments that match the signature, the matching function for this attribute returns a positive real value $r_4 \in [0, 1]$. Each argument contributes with $1/p$ to the value of π_4 where p is the maximum number of arguments that can be matched.

Sophos Ex. 1009
Page 238 of 318

- Number of arguments: $\omega_5 = 0.1$. The number of arguments passed to the program at execution time contributes with $1/q$ where q is the maximum number of arguments the attack tool may receive.
- Data segment size: $\omega_6 = 0.1$. The matching function returns a value $r_6 \in \{0, 1\}$ where 1 corresponds to an exact match and 0 to different data segments.

The following is the pseudocode of the algorithm. It locates the most likely signature node using the text segment size field as an index and computes match score π with that node as well as with its predecessor and/or successor. The highest score of the three is chosen and is reported as the most likely match.

```
# Process signature matching.  High-level pseudocode.
tss = TEXT_SEGMENT_SIZE(input)
i = LOCATE_INDEX(tss, SignatureList)
IF i ≠ 0 THEN
    π1 = MATCH(input, SignatureList[i])
    IF SignatureList HAS PREDECESSOR
        π2 = MATCH(input, SignatureList[i].PREDECESSOR)
    ELSE
        σ2 = 0
    END_IF
    IF SignatureList HAS SUCCESSOR
        π3 = MATCH(input, SignatureList[i].SUCCESSOR)
    ELSE
        π3 = 0
    END_IF
    πmax = MAX(π1, π2, π3)
    PRINT The input corresponds to signature:  πmax
ELSE
    PRINT The process is safe
END_IF
```

5.7.3 Signature fusion

Monitoring cells process signatures separately. The results reported by a process monitoring cell are independent of the ones a network monitoring cell reports. If a network cell goes down and is unable to perform detection, or if an attack tools has been executed but the network interface is down and no packets have being generated, a process monitoring cell could still identify the program being executed in order for the security system to take action. Similarly, upon failure of a process monitoring cell, network cells can provide full detection support to the system by analyzing outbound traffic.

As proved by previous research, the integration of evidence sources is necessary to increase the accuracy of an intrusion detection system [1, 220]. This is especially important when automated response is being implemented. If what the system reports as illicit activity is really

Sophos Ex. 1009
Page 239 of 318

violating the security policy in place, it is indeed possible to incorporate automated response into the security system. When detection results are not fully reliable, no effective automated response can be built on top of the intrusion detection system.

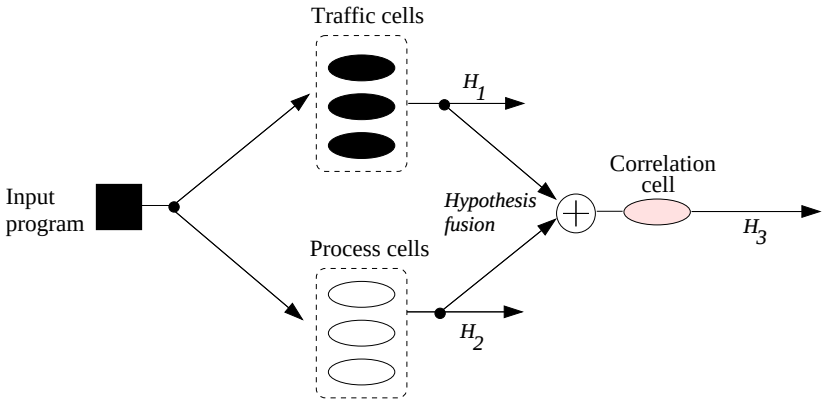


Figure 5.16: Fusing hypothesis H_1 and H_2 into H_3.

For that reason, the two detection systems implemented into FROID are integrated by correlation cells that determine a single score I which involves two hypotheses regarding which signature the observed input corresponds to. A **hypothesis** is the signature that the matching mechanism reports as most likely. In our implementation, hypothesis H_1 points at a certain network signature with similarity score σ, and hypothesis H_2 points at a process signature with similarity score π (Figure 5.16).

The integrated I similarity score is computed by a correlation cell that receives input from two monitoring cells as follows:

$$(1)\ If\ \{\exists\ H_1\ \wedge\ \nexists\ H_2\} \Rightarrow \{H_3 = H_1 \wedge I = \sigma\}$$
$$(2)\ If\ \{\nexists\ H_1\ \wedge\ \exists\ H_2\} \Rightarrow \{H_3 = H_2 \wedge I = \pi\}$$
$$(3)\ If\ \{\nexists\ H_1\ \wedge\ \nexists\ H_2\} \Rightarrow \{H_3 = I = NULL\}$$
$$(4)\ If\ \{\exists\ H_1, H_2 \mid H_1 = H_2\} \Rightarrow \{H_3 = H_1 = H_2 \wedge I = \tfrac{\sigma+\pi}{2}\}$$
$$(5)\ If\ \{\exists\ H_1, H_2 \mid H_1 \neq H_2\} \Rightarrow \begin{cases} \sigma > \pi & \Rightarrow & H_3 = H_1 \wedge I = \sigma \\ \pi > \sigma & \Rightarrow & H_3 = H_2 \wedge I = \pi \\ \sigma = \pi & \Rightarrow & H_3 = I = NULL \end{cases}$$

This cell will then implement automated response based on the information received from two monitoring cells and its reaction will be more definitive relying on the evidence found by independent detection mechanisms. The last section of this chapter describes response mechanisms.

Sophos Ex. 1009
Page 240 of 318

5.8 Session tracing implementation

The session tracing model was implemented separately from the JADE agents at an earlier stage of the project [2]. Every agent is a stand-alone process that uses sockets to provide client and server functionality. It can request information from others and it can also share information.

A 1024-bit RSA key-pair $\{K_{Pub}, K_{Priv}\}$ is generated before the agent is launched. Once the agent is started up, it is ready to collaborate into the tracing of a session. When contacting a peer agent, it exchanges public keys and verifies the signature on the received key. The certification must come from the master key whose public key is accessible to all agents. In order to speed up message encryption, the sender follows a well known encryption procedure [179]: (1) it generates a session key K_{DES} for symmetric encryption; (2) it then encrypts the message using K_{DES}; (3) it encrypts K_{DES} using the public key of the recipient; and finally (4) it sends the encrypted session key and the encrypted message to the destination. The implementations of the RSA and DES algorithms we used are the ones provided by OpenSSL.

The prototype was developed on a RedHat Linux 8.0 box. Although agent libraries can be easily obtained for different compilers, they typically put unnecessary functionality into the agent that produces heavy processes. For this reason and with the intention of providing the best performance possible, the implementation was carried out in C minimizing the number of libraries linked to the executables.

Considering that programs and data files may get damaged in a security incident and that a security tool should not rely on them, our agents do not use the output from programs such as `finger`, `who` and `ps` to read the information stored on system tables like `utmp` and `wtmp`. They instead incorporate this functionality into their own body which gives them extended independence by making them self-contained entities. The size of each agent is around 20 kilobytes to which we add two encryption keys that need to be uploaded (this is just five additional kilobytes) along with the knowledge and rule bases.

The knowledge base of each agent includes a tree structure T where all connection chains are stored. Whenever a host has received all the information regarding a particular session, it stores the connection chain $c = [h_1 \rightarrow h_2 \rightarrow ... \rightarrow h_n]$ into T and sends it to all available agents. Every agent reads c and locates the hosts in c that match its own tree. It then creates new branches to keep its knowledge base up to date.

5.9 Reaction methods

Depending on the similarity scores reported, different response actions can be implemented. When the cells report high similarity scores, the type of reaction that can be performed can be more definitive. In case low scores are reported, it will be better to proceed safely by triggering actions that have not irreversible consequences.

Three cell types are capable of reaction: (1) network monitoring cells can react upon the packet sequences they observe, (2) process monitoring cells can react upon the execution of a process they consider as suspicious, and (3) correlation cells may act upon fused input to request session tracing and to alert neighboring systems.

[2] This prototype system was implemented and successfully tested; the integration with FROID cells is still pending (see Section 7.3).

Sophos Ex. 1009
Page 241 of 318

The types of reaction implemented in FROID are a used to prove the need of having high detection-accuracy levels in order to be able to let a system automatically respond to activity classified illicit.

Response types

```
A = send matching data to evidence repository
B = move executable to evidence repository
C = copy command history to evidence repository
D = suspend process
E = change file execution privileges
F = send kill signal to process
G = start tracing
H = alert target system
```

Traffic Monitoring Cell			
	$\sigma \geq 0.5$	$\Rightarrow$	A
	$\sigma \geq 0.7$	$\Rightarrow$	B
	$\sigma \geq 0.8$	$\Rightarrow$	C
Process Monitoring Cell			
	$\pi \geq 0.5$	$\Rightarrow$	A
	$\pi \geq 0.7$	$\Rightarrow$	B
	$\pi \geq 0.8$	$\Rightarrow$	C
Correlation Cell			
	$\sigma \geq 0.5,\ \pi \geq 0.5,\ I\ \geq 0.7$	$\Rightarrow$	D
	$\sigma \geq 0.5,\ \pi \geq 0.5,\ I\ \geq 0.8$	$\Rightarrow$	E + G
	$\sigma \geq 0.5,\ \pi \geq 0.5,\ I\ \geq 0.9$	$\Rightarrow$	F + H

Table 5.6: Rules to implement automated response.

Notice that a single score can trigger more than one reaction. In the case of high similarity scores from monitoring cells, a correlation cell will integrate the hypotheses and will perform actions to effectively stop the attack launched toward an external system.

Sophos Ex. 1009
Page 242 of 318

Chapter 7

Conclusions, contributions and future work

This chapter condenses final remarks and conclusions, lists the contributions of the dissertation, and describes potential lines for future research.

7.1 Conclusions

- Monitoring for compromise attempts launched locally toward external targets is an effective way of countering the expansion of security threats.
- Outbound intrusion detection is a feasible approach to systems monitoring that complements traditional intrusion detection systems by analyzing outgoing network traffic and by identifying rogue programs in execution. A reduction in the intrusion detection search space contributes increased accuracy to the model.
- Repositories of automated attack tools on the Internet are a valuable security resource that can be analyzed and organized as to produce signatures that characterize the behavior of these programs and to detect them at runtime.
- The use of multiple software agents in the generation and matching of signatures increases detection accuracy by enabling evidence correlation from various data sources and creates a protection mechanism capable of graceful monitoring degradation in the event of system failure or compromise attempt.
- Remote attack programs in execution can be successfully identified using two basic operational aspects: the network traffic they generate, and the characteristics of process that executes them. Separate signatures can be automatically derived and a correlation method can integrate them in order to generate a more reliable matching score.

Sophos Ex. 1009
Page 243 of 318

- Attack signatures supported by an ongology can be easily deployed to intrusion detection agents connected to the Internet as soon as updates become available.

- The use of an ontological model not only allows a multiagent intrusion detection system to have a clean and organized way of handling its knowledge base and its security repositories but it also aids the design and implementation phases of the system. It reduces maintenance time by guiding the definition of the dataflow and even the internal structure of agents.

- The use of virtual machines for the implementation of trusted sub-environments helps reduce the impact a malicious host has on agents by simplifying the process of verifying the integrity of the receiving environment.

- Entropy and frequency analysis of the type of data that is to be handled by an intrusion detection system allows for faster matching algorithms by exploiting the splitting power of the attributes and the locality of reference of the signatures which help speed up the classification of the input.

- The different execution variants that remote automated attack tools may have do not have significant similarities between them as to derive a single fingerprint or signature per attack program. A better approach to follow is the generation of signatures that identify each variant separately and accurately.

- The use of agent and semantic web technologies in the development of security software offers development speed advantages but, due to the functionality that comes with these tools and which is necessary to perform even the most basic functions, it imposes significant overhead to the system that results in poor overall performance of the security system and the executing platform.

- Increased detection accuracy effectively enables the development of automated response methods that allow for various types of reaction either when the input event is classified as illicit, or simply when it is abnormal at a certain degree (determined by a signature matching scoring system).

- The deployment of an agent network in charge of tracing subsequent connections of an intruder across systems connected to the Internet is a feasible way to identify the origin of an interactive connection whose main challenge is the protection of the agents against malicious hosts.

Sophos Ex. 1009
Page 244 of 318

7.2 Contributions

As for contributions, we can say that this dissertation:

- Introduced the notion of outbound intrusion detection as an unselfish, cooperative way of approaching intrusion detection and discussed its advantages and disadvantages.
- Proved that outbound traffic must be inspected for signs of intrusion as much as inbound traffic is in order to detect modern security threats and attack programs.
- Improved the performance of a signature matching algorithm by using experimental data to rank signature attributes by the amount of information they convey.
- Showed how repositories of attack tools can be used as a valuable resource of security information that can be used to improve the design and expand the scope of intrusion detection software.
- Described the behavior of public-domain automated attack tools for the purpose of signature generation and detection.
- Showed that the use of multiple agents for modeling a rogue program contributes increased detection accuracy and robustness to the security mechanism.
- Proposed an intrusion detection architecture based on agent cells organized non-hierarchically and suggested the use of software sub-environments as a viable way to guaranteeing safe execution to agents.
- Proposed and implemented a source-centric ontology for an outbound intrusion detection multiagent architecture that extends ontology work in the area of intrusion detection systems.
- Showed that automated signature generation and deployment over the Internet is possible for an outbound intrusion detection system.
- Implemented a proof-of-concept prototype system for outbound intrusion detection using semantic web and agent technology exposing the advantages and disadvantages of utilizing these tools.
- Explored the problem of session tracing, proposed an agent-based solution, and described its implementation and results. borrowed from bioinformatics.

Sophos Ex. 1009
Page 245 of 318

7.3 Future directions

The following are some ideas to do research derived from the results of this dissertation:

- To design OID models in which monitoring is performed within societies of trusted hosts with a well-defined membership structure. The flow of network traffic could be restricted to members of a group or society that implement OID to protect one another. This would allow us to study the benefits of security-related group memberships as well as the different incident response methods that can be used when defining groups of computers that trust one another and monitor outbound activity among them.

- To develop mechanisms that allow agents to automatically verify the integrity of the receiving sub-environment (e.g., virtual machines and interpreters) in order to minimize the risks posed by malicious hosts. This represents a more feasible way to go about host validation than trying to determine whether the system as a whole is trustworthy.

- To experiment with longer and shorter signatures in order to identify possible performance and precision gains considering variable-size and variable-attribute signatures. The fixed signature structure used in our implementation leaves room for improvement especially if signatures can be adapted to the runtime features of the programs. This could be done by selecting a smaller number of high-value attributes per program variant rather than using a set of attributes common to all programs to characterize them.

- To analyze the potential of ontological models for audit trail reduction. In order to optimize the performance of intrusion detection tools, audit trails must have a regular structure, must contain security-relevant data, and must be as compact as possible. The use of an ontology could help develop a standard for reduced audits that relied on pre-established relationships among attributes in order to infer data. That way, a shorter input would be necessary and information could be derived at the time is needed rather than having it stored without being used.

- To develop an outbound intrusion detection system utilizing supervised-learning techniques that eliminate the need for program signatures. Our research showed the dissimilarity of attack programs while in execution. This aspect can be further explored in order to determine whether a classifier is able to generalize this type of input and to detect unseen attack programs.

- To build and analyze a larger signature repository in order to identify attack-tool groups or clusters. The definition of attack-class signatures could reduce the number of total signatures the system needs to check at detection time hereby improving its responsiveness. Data mining

Sophos Ex. 1009
Page 246 of 318

and machine learning techniques could be used to extract patterns that group programs together, and to detect such patterns in real time.

- To extend the system to other protocols such as UDP and to create more specific, sub-protocol signatures (e.g., ICMP signatures). The first step has been to generate per-program-variant signatures. This only contemplated a small group of TCP/IP-based attack tools but the approach could be extended to UDP and other protocols in order to discover operational differences that help design better signature layouts and algorithms.

- To develop sound, automatic ways of protecting intrusion detection evidence stored on the knowledge bases of software agents. This could be achieved through the use of encryption in such a way as to be able to use this information before a jury in security-related cases. The protection of agent data is an important challenge since agents are vulnerable to misuse by the hosting system. Methods for safeguarding agent-collected evidence in a legally-valid format are of paramount importance to the success of distributed intrusion detection.

- To integrate intrusion and outbound intrusion detection systems so that they are able to share security information and interact with each other (e.g., letting an intrusion detection tool trigger the outbound intrusion detection system when a change in the local security context appears). A more comprehensive security model might appear by interleaving the focus on local resource protection intrusion detection systems have, with the collective monitoring approach proposed by outbound intrusion detection.

- To test the outbound intrusion detection system under extreme, real workload conditions using a very large number of attack signatures. This could help discover new ways of organizing and assigning tasks to agents, and to evaluate the functioning and design of the ontological model.

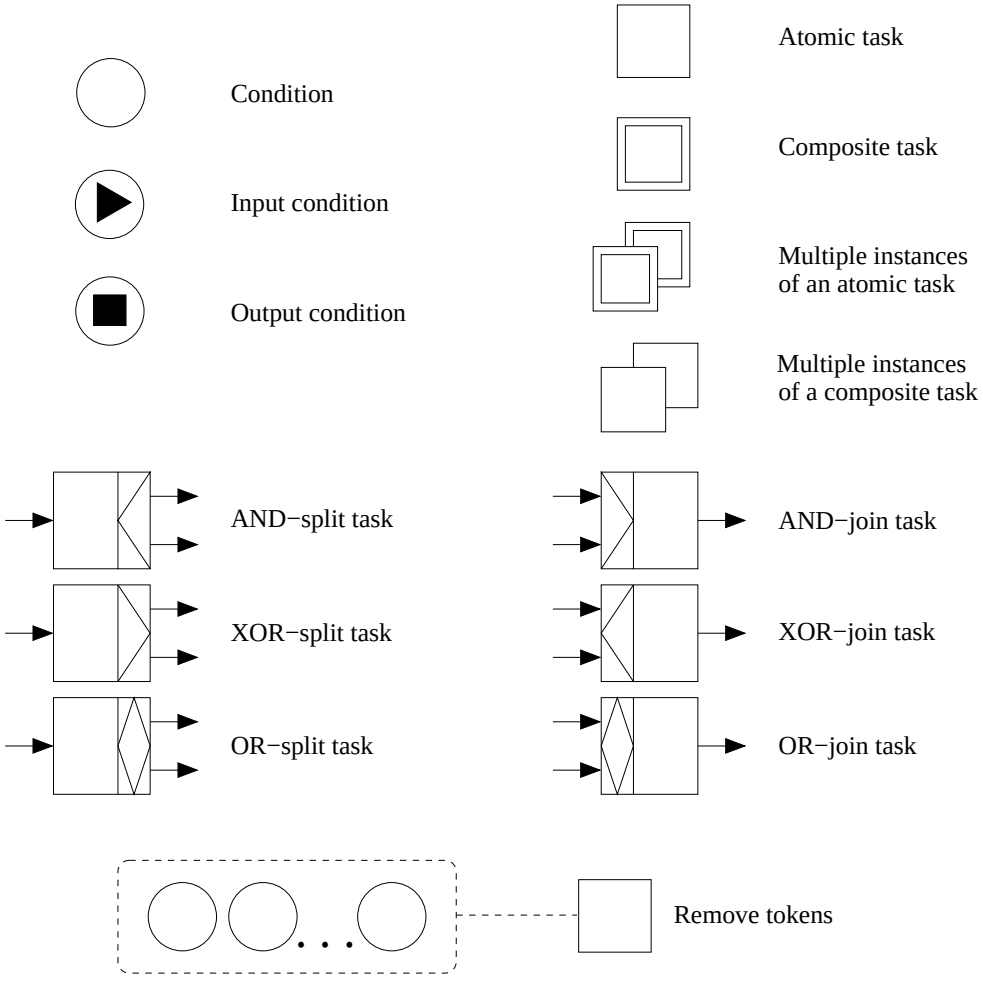
- To extend the functionality of connection-tracing agents to the exchange of encrypted session information among trusted hosts so that suspicious activity can also be detected when encrypted channels are used to compromise a system. Currently, network-based intrusion detection systems are blind to encrypted traffic. If agents placed at different hosts are able to share the contents of encrypted sessions originated at their sites, it could be possible to discover (either, with the traditional intrusion detection approach, or with outbound intrusion detection) an attack in progress that is using ciphered channels.

- To test the FROID system with distributed denial of service tools (DDoS) so that signatures for this type of attack can be generated and deployed as soon as a copy of the tool becomes available. This could help shed some light on the automatic generation of execution signatures for worms and not only for humanly-operated programs.

Sophos Ex. 1009
Page 247 of 318

Appendix A: YAWL symbols

YAWL (*Yet Another Workflow Language*) is a very powerful workflow language based on a rigorous survey of the best of breed solutions on the market [206]. Its constructs are inspired by an advanced concurrency control language, Petri Nets. The basic symbols used to create workflow diagrams in YAWL are the following:



Sophos Ex. 1009
Page 249 of 318

Appendix B: TCP and IP datagram formats

B.1 TCP datagram format

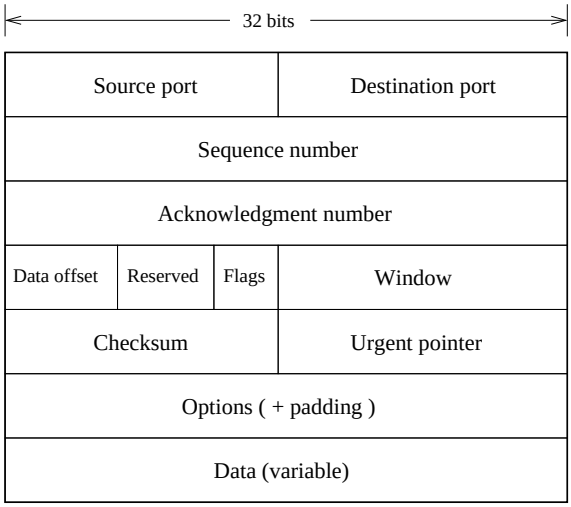


Figure 7.1: TCP datagram format (RFC 793) [38].

Together with Internet Protocol (IP), TCP is one of the core protocols underlying the Internet. They are usually referred to as a group by the term "TCP/IP". TCP enables two computers to establish a connection and exchange information. TCP guarantees delivery of data and also guarantees that information packets will be delivered in the same order in which they were sent. The fields of the TCP packet are described here [38]:

- **Source port and destination port**. Identify the points (sockets) at which upper-layer source and destination processes receive TCP services.
- **Sequence number**. Usually specifies the number assigned to the first byte of data in the current message. Under certain circumstances, it can also be used to identify an initial sequence number to be used in the upcoming transmission.

Sophos Ex. 1009
Page 251 of 318

- **Acknowledgment number**. Contains the sequence number of the next byte of data that the sender of the packet expects to receive.
- **Data offset**. Indicates the number of 32-bit words in the TCP header.
- **Reserved**. Is reserved for future use.
- **Flags**. Carries a variety of control information.
- **Window**. Specifies the size of the senderś receive window (buffer space available for incoming data).
- **Checksum**. Provides information used to determine whether the header was damaged in transit.
- **Urgent pointer**. Points to the first urgent data byte in the packet.
- **Options**. Specifies various TCP options.
- **Data**. Contains upper-layer information.

Example

TCP HEADER FORMAT

```
OCTET 1,2 Source Port (SRC_PORT)
OCTET 3,4 Destination Port (DEST_PORT)
OCTET 5,6,7,8 Sequence Number (SEQ)
OCTET 9,10,11,12 Acknowledgement Number (ACK)
OCTET 13,14 Data Offset (4 bit)+Reserved (6 bit)+
Control Flags(6 bit) (DTO, FLG)
OCTET 15,16 Window (WIN)
OCTET 17,18 Checksum (TCP_SUM)
OCTET 19,20  Urgent Pointer (URP)
OCTET 21,22,23  Options (OPT)
OCTET 24 Padding
OCTET 25,26... Data
```

TCP/IP Packet Example:

```
7E 21 45 00 00 4B 57 49 40 00 FA 06 85 77 C7 B6 78 0E CE D6 95 50 00 6E 04 9F 74 5B EE A2
59 9A 00 0E 50 18 24 00 E3 2A 00 00 2B 4F 4B 20 50 61 73 73 77 6F 72 64 20 72 65 71 75 69
72 65 64 20 66 6F 72 20 61 6C 65 78 75 72 2E 0D 0A 67 B2 7E

Start 7E
SEP 21
IP Header 45 00 00 4B 57 49 40 00 FA 06 85 77 C7 B6 78 0E CE D6 95 50
TCP Header 00 6E 04 9F 74 5B EE A2 59 9A 00 0E 50 18 24 00 E3 2A 00 00
Data 2B 4F 4B 20 50 61 73 73 77 6F 72 64 20 72 65 71 75 69 72 65 64 20 66 6F 72
20 61 6C 65 78 75 72 2E 0D 0A
FCS 67 B2
Stop 7E
```

TCP Header:

```
SRC_PORT=110  DEST_PORT=1183   SEQ=745BEEA2  ACK=599A000E   DTO=5  FLG=18  WIND=9216
TCP_SUM=E32A  URP=0000   (No Options)
```

Data:

```
+OK Password required for alexur\r\n

Control Flags (FLG=18):

FLG=00011000 Urgent Pointer  URG=0
Acknowledgment ACK=1
Push Function PSH=1
Reset connection RST=0
Synchronization SYN=0
Finished data FIN=0

State = ACK-PSH
```

B.2 IPv4 datagram format

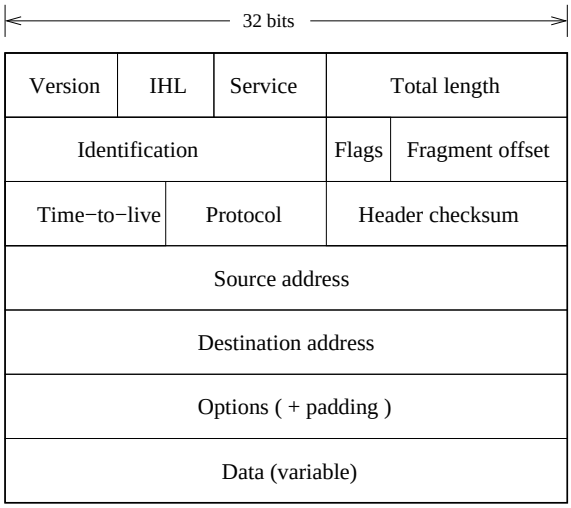


Figure 7.2: IP v.4 datagram format (RFC 791) [38].

The Internet Protocol (IP) provides a frame for encapsulating other protocols like TCP and UDP. The IP header informs the recipient among other things of: the destination and source addresses of the packet, number of octets in the packet, whether the packet can be fragmented or not, how many hops can the packet traverse, the protocol that the packet carries, etc. The IP version currently utilizied is 4. The fields of the IP packet are as follows [38]:

- **Version**. Indicates the version of this IP datagram.

Sophos Ex. 1009
Page 253 of 318

- **IP Header Length (IHL)**. Indicates the datagram header length in 32-bit words.
- **Service or Type of Service**. Specifies how a particular upper-layer protocol would like the current datagram to be handled. Datagrams can be assigned various levels of importance using this field.

  Today this field is used primarily to provide quality of service (QoS) capabilities to TCP/IP for applications requiring predictable bandwidth or delay. RFC 2474 describes a method by which the TOS field is replaced by a DS field that is used to provide differentiated services (DiffServ) on networks. This field is split into two parts. The first 6 bits are used for the DSCP codepoint, which is used to differentiate traffic. The last 2 bits, or CU, are ignored by DiffServ-compliant nodes.
- **Total Length**. Specifies the length of the entire IP packet, including data and header, in bytes.
- **Identification**. Consists of an integer identifying this datagram. This field is used to help piece together datagram fragments.
- **Flags**. Consists of 3 bits, of which the low-order 2 bits control fragmentation. One bit specifies whether the packet can be fragmented; the second bit specifies whether the packet is the last fragment in a series of fragmented packets.
- **Time-to-Live**. Maintains a counter that gradually decrements down to zero, at which point the datagram is discarded. This keeps packets from looping endlessly.
- **Protocol**. Indicates which upper-layer protocol receives incoming packets after IP processing is complete.
- **Header Checksum**. Helps ensure IP header integrity.
- **Source Address**. Specifies the sending node.
- **Destination Address**. Specifies the receiving node.
- **Options**. Allows IP to support various options, such as security.
- **Data**. Contains upper-layer information.

Example

IP HEADER FORMAT

```
OCTET 1 Version (4 bit)+IHL (4 bit) (VER, IHL)
OCTET 2 Type of service (TOS)
OCTET 3,4 Total Length (TOL)
OCTET 5,6  Identification (ID)
OCTET 7,8 Flags (3 bit)+Fragment Offset (13 bit) (FLG, FRO)
OCTET 9 Time to Live (TTL)
OCTET 10 Protocol (PRO)
OCTET 11,12  Header Checksum (IP_SUM)
```

```
OCTET 13,14,15,16  Source Address (SRC)
OCTET 17,18,19,20 Destination Address (DEST)
OCTET 21,22,23 Options (OPT)
OCTET 24 Padding
OCTET 25, 26 ... Data
```

In the example shown the IP packet is encapsulated in a PPP frame including the flag sequence (7E), address (FF 03), protocol (00 21) and FCS (0B 81). In turn, the IP packet encapsulates a UDP message.

Example:

```
7E FF 03 00 21 45 00 00 40 00 01 00 00 3C 11 E0 31 CE D9 8F 1F C7 B6 78 CB 04 63 00
35 00 2C AB DA 00 01 01 00 00 01 00 00 00 00 00 00 04 70 6F 70 64 02 69 78 06 6E 65
74 63 6F 6D 03 63 6F 6D 00 00 01 00 01 0B 81 7E

Start 7E
Address FF 03
SEP 00 21
IP Header 45 00 00 40 00 01 00 00 3C 11 E0 31 CE D9 8F 1F C7 B6 78 CB
Data 04 89 00 35 00 2C AB B4 00 01 01 00 00 01 00 00 00 00 00 00 04 70 6F
70 64 02 69 78 06 6E 65 74 63 6F 6D 03 63 6F 6D 00 00 01 00 01
FCS 0B 81
Stop 7E
```

IP Header:

```
VER=4 IHL=5 TOS=0 TOL=64 ID=1 FLG=00 FRO=00 TTL=60 PRO=17 IP_SUM=E031
SRC=206.217.143.31. DEST=199.182.120.203. OPT=00000000
```

Appendix C: Outbound intrusion detection ontology in OWL

This is the top level definition of the source-centric OID ontology that supports the proposed architecture (Figure 4.12). It has been created in OWL [7] using the Protégé ontology definition tool [144].

```
<?xml version="1.0"?>
<rdf:RDF
    xmlns="http://localhost/OIDontology#"
    xmlns:protege="http://protege.stanford.edu/plugins/owl/protege#"
    xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
    xmlns:owl="http://www.w3.org/2002/07/owl#"
  xml:base="http://localhost/OIDontology">
  <owl:Ontology rdf:about="OIDontology">
    <owl:imports rdf:resource="http://protege.stanford.edu/plugins/owl/protege"/>
  </owl:Ontology>
  <owl:Class>
    <owl:unionOf rdf:parseType="Collection">
      <rdf:Description rdf:about="http://www.w3.org/2002/07/owl#Thing"/>
      <owl:Class rdf:about="#TrafficSignature"/>
    </owl:unionOf>
  </owl:Class>
  <owl:Class rdf:ID="Tracer">
    <rdfs:subClassOf>
      <owl:Class rdf:about="#Reactor"/>
    </rdfs:subClassOf>
  </owl:Class>
  <owl:Class rdf:ID="ProcessMonitoringCell">
    <rdfs:subClassOf>
      <owl:Class rdf:ID="Cell"/>
    </rdfs:subClassOf>
  </owl:Class>
  <owl:Class rdf:ID="IPpacket"/>
  <owl:Class rdf:ID="ProcessSignature">
    <rdfs:subClassOf>
      <owl:Class rdf:ID="Signature"/>
    </rdfs:subClassOf>
  </owl:Class>
  <owl:Class rdf:ID="CompiledExecutable">
    <rdfs:subClassOf>
      <owl:Class rdf:ID="Executable"/>
```


Sophos Ex. 1009
Page 257 of 318

```
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="ProcessSigner">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Signer"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="TrafficSigner">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Signer"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="HostileProcess">
  <rdfs:subClassOf>
    <owl:Class rdf:ID="Process"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="Reaction"/>
<owl:Class rdf:ID="ProcessSensor">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Sensor"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="EvidenceRepository">
  <rdfs:subClassOf>
    <owl:Class rdf:ID="Repository"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="TrafficDetector">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Detector"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="ActiveReaction">
  <rdfs:subClassOf rdf:resource="#Reaction"/>
</owl:Class>
<owl:Class rdf:ID="ProcessSigningCell">
  <rdfs:subClassOf rdf:resource="#Cell"/>
</owl:Class>
<owl:Class rdf:ID="TrafficSensor">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Sensor"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="Reactor">
  <rdfs:subClassOf>
    <owl:Class rdf:ID="Agent"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="ProcessDetector">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Detector"/>
  </rdfs:subClassOf>
</owl:Class>
```

```
<owl:Class rdf:ID="SignatureRepository">
  <rdfs:subClassOf rdf:resource="#Repository"/>
</owl:Class>
<owl:Class rdf:ID="SafeProcess">
  <rdfs:subClassOf rdf:resource="#Process"/>
</owl:Class>
<owl:Class rdf:ID="TCPpacket">
  <rdfs:subClassOf rdf:resource="#IPpacket"/>
</owl:Class>
<owl:Class rdf:ID="DataSource"/>
<owl:Class rdf:ID="InterpretedExecutable">
  <rdfs:subClassOf rdf:resource="#Executable"/>
</owl:Class>
<owl:Class rdf:ID="InternalData">
  <rdfs:subClassOf rdf:resource="#DataSource"/>
</owl:Class>
<owl:Class rdf:ID="TrafficSignature">
  <rdfs:subClassOf rdf:resource="#Signature"/>
</owl:Class>
<owl:Class rdf:ID="Correlator">
  <rdfs:subClassOf rdf:resource="#Agent"/>
</owl:Class>
<owl:Class rdf:ID="HybridDetector">
  <rdfs:subClassOf>
    <owl:Class rdf:about="#Detector"/>
  </rdfs:subClassOf>
</owl:Class>
<owl:Class rdf:ID="Signer">
  <rdfs:subClassOf rdf:resource="#Correlator"/>
</owl:Class>
<owl:Class rdf:ID="TrafficMonitoringCell">
  <rdfs:subClassOf rdf:resource="#Cell"/>
</owl:Class>
<owl:Class rdf:ID="UDPpacket">
  <rdfs:subClassOf rdf:resource="#IPpacket"/>
</owl:Class>
<owl:Class rdf:ID="Sensor">
  <rdfs:subClassOf rdf:resource="#Agent"/>
</owl:Class>
<owl:Class rdf:ID="Guard">
  <rdfs:subClassOf rdf:resource="#Reactor"/>
</owl:Class>
<owl:Class rdf:ID="Detector">
  <rdfs:subClassOf rdf:resource="#Correlator"/>
</owl:Class>
<owl:Class rdf:ID="CorrelationCell">
  <rdfs:subClassOf rdf:resource="#Cell"/>
</owl:Class>
<owl:Class rdf:ID="OutgoingData">
  <rdfs:subClassOf rdf:resource="#DataSource"/>
</owl:Class>
<owl:Class rdf:ID="TrafficSigningCell">
  <rdfs:subClassOf rdf:resource="#Cell"/>
</owl:Class>
```

Sophos Ex. 1009
Page 259 of 318

```
<owl:Class rdf:ID="PassiveReaction">
  <rdfs:subClassOf rdf:resource="#Reaction"/>
</owl:Class>
<owl:ObjectProperty rdf:ID="container">
  <rdfs:domain rdf:resource="#Signature"/>
  <rdfs:range rdf:resource="#SignatureRepository"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="member">
  <rdfs:domain rdf:resource="#Agent"/>
  <rdfs:range rdf:resource="#Cell"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="trafficSource">
  <rdfs:domain rdf:resource="#TrafficSensor"/>
  <rdfs:range rdf:resource="#OutgoingData"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="trafficSupplier">
  <rdfs:domain rdf:resource="#HybridDetector"/>
  <rdfs:range rdf:resource="#TrafficDetector"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="executant">
  <rdfs:range rdf:resource="#Reactor"/>
  <rdfs:domain rdf:resource="#Correlator"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="processSource">
  <rdfs:range rdf:resource="#InternalData"/>
  <rdfs:domain rdf:resource="#ProcessSensor"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="effect">
  <rdfs:range rdf:resource="#Reaction"/>
  <rdfs:domain rdf:resource="#Reactor"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="evidence">
  <rdfs:range rdf:resource="#Executable"/>
  <rdfs:domain rdf:resource="#PassiveReaction"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="item">
  <rdfs:range rdf:resource="#IPpacket"/>
  <rdfs:domain rdf:resource="#OutgoingData"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="processSupplier">
  <rdfs:domain rdf:resource="#HybridDetector"/>
  <rdfs:range rdf:resource="#ProcessDetector"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="cause">
  <rdfs:range rdf:resource="#HostileProcess"/>
  <rdfs:domain rdf:resource="#PassiveReaction"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="processProduct">
  <rdfs:domain rdf:resource="#ProcessSigner"/>
  <rdfs:range rdf:resource="#ProcessSignature"/>
</owl:ObjectProperty>
<owl:ObjectProperty rdf:ID="storage">
  <rdfs:range rdf:resource="#InternalData"/>
  <rdfs:domain rdf:resource="#Process"/>
```

```
  </owl:ObjectProperty>
  <owl:ObjectProperty rdf:ID="execution">
    <rdfs:range rdf:resource="#Process"/>
    <rdfs:domain rdf:resource="#Executable"/>
  </owl:ObjectProperty>
  <owl:ObjectProperty rdf:ID="proof">
    <rdfs:domain rdf:resource="#EvidenceRepository"/>
    <rdfs:range rdf:resource="#Executable"/>
  </owl:ObjectProperty>
  <owl:ObjectProperty rdf:ID="informant">
    <rdfs:range rdf:resource="#Sensor"/>
    <rdfs:domain rdf:resource="#Correlator"/>
  </owl:ObjectProperty>
  <owl:ObjectProperty rdf:ID="trafficProduct">
    <rdfs:range rdf:resource="#TrafficSignature"/>
    <rdfs:domain rdf:resource="#TrafficSigner"/>
  </owl:ObjectProperty>
  <owl:ObjectProperty rdf:ID="description">
    <rdfs:range rdf:resource="#EvidenceRepository"/>
    <rdfs:domain rdf:resource="#Agent"/>
  </owl:ObjectProperty>
</rdf:RDF>

<!-- Created with Protege (with OWL Plugin 1.2 beta, Build 139)  http://protege.stanford.edu -->
```

Appendix D: Classification Trees

These trees were generated using Gnome Data Miner's decision trees module [69].

Entropy tree

```
Splitting function: Entropy

/*-------------------------------------------------------------------
  domains
-------------------------------------------------------------------*/
dom(packetCounter) = ZZ [2, 16];
dom(fragmentation) = { 0, 1 };
dom(sourceAddress) = IR [-8205, -92];
dom(destinationAddress) = ZZ [-9720, -2322];
dom(destinationPort1) = ZZ [1, 8080];
dom(destinationport2) = ZZ [0, 65535];
dom(windowSize) = { 0, 1 };
dom(h8) = ZZ [0, 3360];
dom(h4) = ZZ [0, 3256];
dom(h5) = ZZ [0, 3360];
dom(f3) = ZZ [0, 8];
dom(f4) = ZZ [0, 9];
dom(f7) = ZZ [0, 8];
dom(f8) = ZZ [0, 8];
dom(f10) = ZZ [0, 8];
dom(f15) = ZZ [0, 8];
dom(l4) = ZZ [0, 1448];
dom(l6) = ZZ [0, 1448];
dom(programID) = { 1a, 1b, 2a, 2b, 2c, 3a, 4a, 4b, 5a, 5b, 6a, 6b,
  7a, 7b, 8a, 8b, 9a, 9b, 10a, 10b, 11a, 12a, 12b, 13a, 14a, 15a,
  16a, 17a, 18a, 18b, 19a, 20a, 21a, 22a, 23a, 24a, 25a, 26a, 27a,
  27b, 28a, 29a, 30a, 31a, 32a, 33a, 34a, 35a, 36a, 36b, 37a, 38a,
  39a, 40a, 41a, 42a, 43a, 44a, 45a, 46a, 47a, 48a, 49a, 50a };

/*-------------------------------------------------------------------
  decision tree
-------------------------------------------------------------------*/
dtree(programID) =
{ (l4|29)
  <:{ (f7|0.5)
      <:{ (packetCounter|4.5)
```


Sophos Ex. 1009
Page 263 of 318

```
        <:{ (h8|1075)
            <:{ (packetCounter|2.5)
                <:{ (destinationPort1|37)
                    <:{ 37a: 6 },
                    >:{ 34a: 5 }},
                >:{ (packetCounter|3.5)
                    <:{ 1a: 5 },
                    >:{ 36a: 5 }}},
            >:{ (destinationPort1|480)
                <:{ (sourceAddress|-4148.5)
                    <:{ 39a: 5 },
                    >:{ 29a: 5 }},
                >:{ (packetCounter|2.5)
                    <:{ 17a: 5 },
                    >:{ 31a: 5 }}}},
        >:{ (packetCounter|5.5)
            <:{ (f3|6)
                <:{ (destinationAddress|-3716)
                    <:{ 25a: 5 },
                    >:{ 11a: 5 }},
                >:{ (h5|4905)
                    <:{ 8b: 5 },
                    >:{ 4a: 5 }}},
            >:{ (packetCounter|9.5)
                <:{ (packetCounter|7)
                    <:{ 23a: 5 },
                    >:{ 24a: 5 }},
                >:{ (packetCounter|12)
                    <:{ 15a: 5 },
                    >:{ 20a: 5 }}}}},
    >:{ (h5|1035)
        <:{ (destinationPort1|24)
            <:{ (packetCounter|9)
                <:{ (destinationPort1|22.5)
                    <:{ 9a: 5 },
                    >:{ 1b: 5 }},
                >:{ (destinationport2|32767.5)
                    <:{ 27a: 5 },
                    >:{ 35a: 5 }}},
            >:{ (destinationPort1|3066)
                <:{ (packetCounter|8.5)
                    <:{ 10b: 5 },
                    >:{ 49a: 5 }},
                >:{ (packetCounter|7.5)
                    <:{ 7b: 5 },
                    >:{ 44a: 5 }}}},
        >:{ (packetCounter|14)
            <:{ (destinationPort1|51.5)
                <:{ (packetCounter|11.5)
                    <:{ 36b: 5 },
                    >:{ 14a: 5 }},
                >:{ (destinationPort1|111.5)
                    <:{ 32a: 5 },
                    >:{ 21a: 5 }}},
```

```
            >:{ (destinationPort1|348.5)
                <:{ (destinationPort1|40.5)
                    <:{ 41a: 5 },
                    >:{ 48a: 5 }},
                >:{ (destinationPort1|3320)
                    <:{ 46a: 5 },
                    >:{ 7a: 5 }}}}},
>:{ (14|641)
    <:{ (packetCounter|13)
        <:{ (packetCounter|7.5)
            <:{ (packetCounter|6.5)
                <:{ (packetCounter|5.5)
                    <:{ 6b: 5 },
                    >:{ 16a: 5 }},
                >:{ (fragmentation)
                    0:{ 38a: 5 },
                    1:{ 42a: 5 }}},
            >:{ (windowSize)
                0:{ (destinationPort1|51.5)
                    <:{ 47a: 5 },
                    >:{ 13a: 5 }},
                1:{ (packetCounter|10.5)
                    <:{ 43a: 5 },
                    >:{ 40a: 5 }}}},
        >:{ (f8|4)
            <:{ (windowSize)
                0:{ (packetCounter|15)
                    <:{ 28a: 5 },
                    >:{ 8a: 5 }},
                1:{ (packetCounter|15.5)
                    <:{ 45a: 5 },
                    >:{ 30a: 5 }}},
            >:{ (f15|4)
                <:{ (destinationport2|55)
                    <:{ 18b: 5 },
                    >:{ 6a: 5 }},
                >:{ (packetCounter|15)
                    <:{ 33a: 5 },
                    >:{ 18a: 5 }}}}},
    >:{ (h4|2433)
        <:{ (f10|0.5)
            <:{ (destinationPort1|51)
                <:{ (l6|724)
                    <:{ 10a: 5 },
                    >:{ 2b: 5 }},
                >:{ (packetCounter|9)
                    <:{ 4b: 5 },
                    >:{ 50a: 5 }}},
            >:{ (packetCounter|15)
                <:{ (packetCounter|12)
                    <:{ 3a: 5 },
                    >:{ 9b: 5 }},
                >:{ (f4|3)
                    <:{ 2a: 5 },
```

```
                    >:{ 2c: 5 }}}},
          >:{ (packetCounter|7.5)
              <:{ (windowSize)
                  0:{ (destinationPort1|51.5)
                      <:{ 12a: 5 },
                      >:{ 12b: 5 }},
                  1:{ (packetCounter|6.5)
                      <:{ 19a: 5 },
                      >:{ 26a: 5 }}},
              >:{ (f7|4)
                  <:{ (packetCounter|12)
                      <:{ 5b: 5 },
                      >:{ 5a: 5 }},
                  >:{ (l6|1317)
                      <:{ 27b: 5 },
                      >:{ 22a: 5 }}}}}};

/*-----------------------------------------------------------------------
  number of attributes: 19
  tree height         : 7
  number of nodes     : 127
  number of tuples    : 321
-----------------------------------------------------------------------*/
```

Information Gain tree

```
Splitting function: Information Gain

/*-----------------------------------------------------------------------
  domains
-----------------------------------------------------------------------*/
dom(packetCounter) = ZZ [2, 16];
dom(fragmentation) = { 0, 1 };
dom(sourceAddress) = IR [-8205, -92];
dom(destinationAddress) = ZZ [-9720, -2322];
dom(destinationPort1) = ZZ [1, 8080];
dom(destinationport2) = ZZ [0, 65535];
dom(windowSize) = { 0, 1 };
dom(stdDev) = { 0, 30 };
dom(h5) = ZZ [0, 3360];
dom(h8) = ZZ [0, 3360];
dom(h4) = ZZ [0, 3256];
dom(f4) = ZZ [0, 9];
dom(f7) = ZZ [0, 8];
dom(f8) = ZZ [0, 8];
dom(f15) = ZZ [0, 8];
dom(l6) = ZZ [0, 1448];
dom(programID) = { 1a, 1b, 2a, 2b, 2c, 3a, 4a, 4b, 5a, 5b, 6a, 6b,
  7a, 7b, 8a, 8b, 9a, 9b, 10a, 10b, 11a, 12a, 12b, 13a, 14a, 15a,
  16a, 17a, 18a, 18b, 19a, 20a, 21a, 22a, 23a, 24a, 25a, 26a, 27a,
  27b, 28a, 29a, 30a, 31a, 32a, 33a, 34a, 35a, 36a, 36b, 37a, 38a,
```

```
  39a, 40a, 41a, 42a, 43a, 44a, 45a, 46a, 47a, 48a, 49a, 50a };

/*-----------------------------------------------------------------------
  decision tree
-----------------------------------------------------------------------*/
dtree(programID) =
{ (packetCounter|2.5)
  <:{ (destinationPort1|37)
      <:{ 37a: 6 },
      >:{ (sourceAddress|-4148.5)
          <:{ (destinationPort1|1473.5)
              <:{ (h8|934)
                  <:{ 34a: 5 },
                  >:{ 39a: 5 }},
              >:{ 17a: 5 }},
          >:{ 29a: 5 }}},
  >:{ (stdDev)
      0 :{ (windowSize)
           0:{ (fragmentation)
               0:{ (packetCounter|15)
                   <:{ (packetCounter|4)
                       <:{ 1a: 5 },
                       >:{ (h5|1292)
                           <:{ (packetCounter|10.5)
                               <:{ (packetCounter|5.5)
                                   <:{ (destinationAddress|-3716)
                                       <:{ (destinationPort1|3042)
                                           <:{ (destinationPort1|51)
                                               <:{ (destinationPort1|20.5)
                                                   <:{ 8b: 5 },
                                                   >:{ 6b: 5 }},
                                               >:{ 4a: 5 }},
                                           >:{ 25a: 5 }},
                                       >:{ 11a: 5 }},
                                   >:{ (packetCounter|6.5)
                                       <:{ (destinationAddress|-7415)
                                           <:{ 16a: 5 },
                                           >:{ 23a: 5 }},
                                       >:{ (h8|2325)
                                           <:{ (packetCounter|8.5)
                                               <:{ (packetCounter|7.5)
                                                   <:{ (destinationport2|32767.5)
                                                       <:{ 12b: 5 },
                                                       >:{ 38a: 5 }},
                                                   >:{ (destinationPort1|7080)
                                                       <:{ (destinationPort1|3095)
                                                           <:{ (destinationPort1|22.5)
                                                               <:{ 9a: 5 },
                                                               >:{ (destinationPort1|95)
                                                                   <:{ (h8|5135)
                                                                       <:{ 47a: 5 },
                                                                       >:{ (h5|2232)
                                                                           <:{ 13a: 5 },
                                                                           >:{ 4b: 5 }
```

```
                                                        }
                                                    },
                                                >:{ 10b: 5 }
                                            }},
                                        >:{ 44a: 5 }
                                    },
                                >:{ 24a: 5 }}},
                            >:{ (destinationPort1|23.5)
                                <:{ (l6|724)
                                    <:{ 10a: 5 },
                                    >:{ 2b: 5 }},
                                >:{ (destinationPort1|52.5)
                                    <:{ 49a: 5 },
                                    >:{ (l6|1317)
                                        <:{ 27b: 5 },
                                        >:{ 22a: 5 }
                                    }}}},
                        >:{ (destinationPort1|51)
                            <:{ (destinationport2|32767.5)
                                <:{ 27a: 5 },
                                >:{ 35a: 5 }},
                            >:{ (destinationport2|32767.5)
                                <:{ 50a: 5 },
                                >:{ 3a: 5 }}}}},
                >:{ (destinationPort1|126.5)
                    <:{ (destinationPort1|95)
                        <:{ (destinationPort1|51)
                            <:{ 36b: 5 },
                            >:{ 32a: 5 }},
                        >:{ 15a: 5 }},
                    >:{ 21a: 5 }}},
            >:{ (destinationPort1|50.5)
                <:{ 28a: 5 },
                >:{ (destinationport2|55)
                    <:{ 33a: 5 },
                    >:{ 9b: 5 }}}}},
    >:{ (h8|1135)
        <:{ 41a: 5 },
        >:{ (destinationPort1|348.5)
            <:{ (destinationPort1|51)
                <:{ 5a: 5 },
                >:{ (destinationport2|66)
                    <:{ (destinationport2|11)
                        <:{ (f4|3)
                            <:{ (f7|4.5)
                                <:{ (f8|4)
                                    <:{ 2a: 5 },
                                    >:{ (f15|4)
                                        <:{ 18b: 5 },
                                        >:{ 18a: 5 }}},
                                >:{ 8a: 5 }},
                            >:{ 2c: 5 }},
                        >:{ 48a: 5 }},
                    >:{ 6a: 5 }}},
```

```
                        >:{ 46a: 5 }}}},
            1:{ (packetCounter|7.5)
                <:{ (packetCounter|5.5)
                    <:{ 36a: 5 },
                    >:{ 42a: 5 }},
                >:{ 5b: 5 }}},
        1:{ (packetCounter|15.5)
            <:{ (packetCounter|14)
                <:{ (packetCounter|11)
                    <:{ (packetCounter|4.5)
                        <:{ 31a: 5 },
                        >:{ (packetCounter|8)
                            <:{ (packetCounter|6.5)
                                <:{ 19a: 5 },
                                >:{ (destinationPort1|3022)
                                    <:{ 26a: 5 },
                                    >:{ 7b: 5 }}},
                            >:{ 43a: 5 }}},
                    >:{ 20a: 5 }},
                >:{ 45a: 5 }},
            >:{ (destinationPort1|3022)
                <:{ 30a: 5 },
                >:{ 7a: 5 }}}},
    30:{ (packetCounter|10)
        <:{ (packetCounter|7.5)
            <:{ 12a: 5 },
            >:{ 1b: 5 }},
        >:{ (windowSize)
            0:{ 14a: 5 },
            1:{ 40a: 5 }}}}};

/*------------------------------------------------------------------------
  number of attributes: 17
  tree height         : 20
  number of nodes     : 127
  number of tuples    : 321
------------------------------------------------------------------------*/
```

GINI tree

```
Splitting function: GINI

/*------------------------------------------------------------------------
  domains
------------------------------------------------------------------------*/
dom(packetCounter) = ZZ [2, 16];
dom(fragmentation) = { 0, 1 };
dom(sourceAddress) = IR [-8205, -92];
dom(destinationAddress) = ZZ [-9720, -2322];
dom(destinationPort1) = ZZ [1, 8080];
dom(destinationport2) = ZZ [0, 65535];
dom(windowSize) = { 0, 1 };
```

```
dom(h4) = ZZ [0, 3256];
dom(h5) = ZZ [0, 3360];
dom(h8) = ZZ [0, 3360];
dom(f4) = ZZ [0, 9];
dom(f7) = ZZ [0, 8];
dom(f8) = ZZ [0, 8];
dom(f15) = ZZ [0, 8];
dom(l6) = ZZ [0, 1448];
dom(programID) = { 1a, 1b, 2a, 2b, 2c, 3a, 4a, 4b, 5a, 5b, 6a, 6b,
  7a, 7b, 8a, 8b, 9a, 9b, 10a, 10b, 11a, 12a, 12b, 13a, 14a, 15a,
  16a, 17a, 18a, 18b, 19a, 20a, 21a, 22a, 23a, 24a, 25a, 26a, 27a,
  27b, 28a, 29a, 30a, 31a, 32a, 33a, 34a, 35a, 36a, 36b, 37a, 38a,
  39a, 40a, 41a, 42a, 43a, 44a, 45a, 46a, 47a, 48a, 49a, 50a };

/*-----------------------------------------------------------------
  decision tree
-----------------------------------------------------------------*/
dtree(programID) =
{ (packetCounter|15.5)
  <:{ (windowSize)
      0:{ (packetCounter|13)
          <:{ (packetCounter|11.5)
              <:{ (f7|4)
                  <:{ (h4|2338)
                      <:{ (packetCounter|6.5)
                          <:{ (packetCounter|5.5)
                              <:{ (packetCounter|2.5)
                                  <:{ (sourceAddress|-4148.5)
                                      <:{ (destinationPort1|1473.5)
                                          <:{ (destinationPort1|94)
                                              <:{ (destinationPort1|37)
                                                  <:{ 37a: 6 },
                                                  >:{ 34a: 5 }},
                                              >:{ 39a: 5 }},
                                          >:{ 17a: 5 }},
                                      >:{ 29a: 5 }},
                                  >:{ (packetCounter|4.5)
                                      <:{ (packetCounter|3.5)
                                          <:{ 1a: 5 },
                                          >:{ 36a: 5 }},
                                      >:{ (destinationAddress|-3716)
                                          <:{ (destinationPort1|3042)
                                              <:{ (destinationPort1|51)
                                                  <:{ (destinationPort1|20.5)
                                                      <:{ 8b: 5 },
                                                      >:{ 6b: 5 }},
                                                  >:{ 4a: 5 }},
                                              >:{ 25a: 5 }},
                                          >:{ 11a: 5 }}}},
                              >:{ (destinationAddress|-7415)
                                  <:{ 16a: 5 },
                                  >:{ 23a: 5 }}},
                          >:{ (packetCounter|7.5)
                              <:{ (fragmentation)
```

```
                        0:{ (destinationPort1|51.5)
                            <:{ 12a: 5 },
                            >:{ (destinationport2|32767.5)
                                <:{ 12b: 5 },
                                >:{ 38a: 5 }}},
                        1:{ 42a: 5 }},
                    >:{ (fragmentation)
                        0:{ (destinationPort1|7080)
                            <:{ (destinationPort1|3095)
                                <:{ (destinationPort1|95)
                                    <:{ (destinationPort1|51.5)
                                        <:{ (destinationPort1|22.5)
                                            <:{ 9a: 5 },
                                            >:{ (destinationport2|4595.5)
                                                <:{ 47a: 5 },
                                                >:{ 1b: 5 }}
                                          },
                                        >:{ (destinationport2|32767.5)
                                            <:{ 13a: 5 },
                                            >:{ 4b: 5 }}},
                                    >:{ 10b: 5 }},
                                >:{ 44a: 5 }},
                            >:{ 24a: 5 }},
                        1:{ 5b: 5 }}}},
                >:{ (packetCounter|9.5)
                    <:{ (destinationPort1|52.5)
                        <:{ (destinationPort1|23.5)
                            <:{ (16|724)
                                <:{ 10a: 5 },
                                >:{ 2b: 5 }},
                            >:{ 49a: 5 }},
                        >:{ (16|1317)
                            <:{ 27b: 5 },
                            >:{ 22a: 5 }}},
                    >:{ (destinationPort1|51)
                        <:{ (destinationport2|32767.5)
                            <:{ 27a: 5 },
                            >:{ 35a: 5 }},
                        >:{ (destinationport2|32767.5)
                            <:{ 50a: 5 },
                            >:{ 3a: 5 }}}}},
            >:{ (destinationPort1|126.5)
                <:{ (destinationPort1|95)
                    <:{ (destinationPort1|51)
                        <:{ 36b: 5 },
                        >:{ 32a: 5 }},
                    >:{ 15a: 5 }},
                >:{ 21a: 5 }}},
        >:{ 14a: 5 }},
    >:{ (destinationPort1|50.5)
        <:{ 28a: 5 },
        >:{ (destinationport2|55)
            <:{ 33a: 5 },
            >:{ 9b: 5 }}}},
```

```
     1:{ (packetCounter|14)
         <:{ (packetCounter|12.5)
             <:{ (packetCounter|10.5)
                 <:{ (packetCounter|8)
                     <:{ (packetCounter|6.5)
                         <:{ (packetCounter|4.5)
                             <:{ 31a: 5 },
                             >:{ 19a: 5 }},
                         >:{ (destinationPort1|3022)
                             <:{ 26a: 5 },
                             >:{ 7b: 5 }}},
                     >:{ 43a: 5 }},
                 >:{ 40a: 5 }},
             >:{ 20a: 5 }},
         >:{ 45a: 5 }}},
  >:{ (destinationPort1|3320)
     <:{ (f4|3)
         <:{ (h5|52341)
             <:{ (destinationPort1|21.5)
                 <:{ (destinationPort1|11)
                     <:{ 41a: 5 },
                     >:{ 30a: 5 }},
                 >:{ 5a: 5 }},
             >:{ (destinationport2|66)
                 <:{ (h8|321)
                     <:{ (f4|3)
                         <:{ (f7|4.5)
                             <:{ (f8|4)
                                 <:{ 2a: 5 },
                                 >:{ (f15|4)
                                     <:{ 18b: 5 },
                                     >:{ 18a: 5 }}},
                             >:{ 8a: 5 }},
                         >:{ 2c: 5 }},
                     >:{ 48a: 5 }},
                 >:{ 6a: 5 }}},
         >:{ 46a: 5 }},
     >:{ 7a: 5 }}};

/*--------------------------------------------------------------------
  number of attributes: 16
  tree height         : 16
  number of nodes     : 127
  number of tuples    : 321
--------------------------------------------------------------------*/
```

Appendix E: Process signatures

The table expanding the following pages corresponds to the process signatures of the programs surveyed in this project.

Sophos Ex. 1009
Page 273 of 318

```
<ProcessSignature rdf:ID=305monit>
       <md5> ead28b1b7133a311dcb2c0ab8c2b5c4a </md5>
       <textSegmentSize> 4103 </textSegmentSize>
       <dataSegmentSize> 508 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 15819 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=akill2>
       <md5> 7c57fda886a90ab8d3201a557a79ab20 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 1796 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=anubisexp>
       <md5> 4e0aa24f43a8b13ac33e06eb162dbaea </md5>
       <textSegmentSize> 10678 </textSegmentSize>
       <dataSegmentSize> 868 </dataSegmentSize>
       <arguments> -b-l-r-w </arguments>
       <numArguments> 2-6 </numArguments>
       <size> 23806 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=apache-nosejob>
       <md5> 340ffbab2f45b95c33cd3cf082282834 </md5>
       <textSegmentSize> 9821 </textSegmentSize>
       <dataSegmentSize> 972 </dataSegmentSize>
       <arguments> -h-t-o-b-d-z-r-w-c </arguments>
       <numArguments> 1-5 </numArguments>
       <size> 22523 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=apache-scalp>
       <md5> 9d4ad46076f207c90a5438734efb5632 </md5>
       <textSegmentSize> 6391 </textSegmentSize>
       <dataSegmentSize> 620 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 18769 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=arksink>
       <md5> 318472fdcad36d111831d6880ee83610 </md5>
       <textSegmentSize> 3759 </textSegmentSize>
       <dataSegmentSize> 376 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 15450 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=atphttp>
       <md5> fb654b077f6d67ce87779b15a4663e27 </md5>
       <textSegmentSize> 4095 </textSegmentSize>
       <dataSegmentSize> 420 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1-3 </numArguments>
       <size> 15742 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=bazooka>
       <md5> db390f1e415a0b3d1d2770c892a15420 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 5-6 </numArguments>
       <size> 4773 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=bysin>
       <md5> c6b4738762ecc12d41420474ef4561af </md5>
       <textSegmentSize> 3885 </textSegmentSize>
       <dataSegmentSize> 336 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 15493 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=cge>
       <md5> 21297199a69c4b86a57ce4170dda4482 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 18530 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=cobain>
       <md5> f0cc271a17471fcd8555602dc84764d5 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 1528 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=delefate>
       <md5> c100a9d4f27fc5f4e5160f21d68e680f </md5>
       <textSegmentSize> 4396 </textSegmentSize>
       <dataSegmentSize> 864 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 16664 </size>
</ProcessSignature>
```


Sophos Ex. 1009
Page 274 of 318

```
<ProcessSignature rdf:ID=emule>
       <md5> f97ec52de0388c043e7f5547f5f4b5aa </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 3362 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=esignal>
       <md5> 87ef0810b9c74733065752ff2e9feb8b </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 3207 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=ethboom>
       <md5> decbb70da04e8a22d3baf6f4e7938cf7 </md5>
       <textSegmentSize> 3584 </textSegmentSize>
       <dataSegmentSize> 2560 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1-2 </numArguments>
       <size> 6656 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=exceed>
       <md5> 43f3d018d0922e363781a7982db1f359 </md5>
       <textSegmentSize> 3176 </textSegmentSize>
       <dataSegmentSize> 300 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 14301 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=extremail2>
       <md5> a6896905340d841524db18566c94dedb </md5>
       <textSegmentSize> 2745 </textSegmentSize>
       <dataSegmentSize> 312 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 14133 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=ezploit>
       <md5> 3113afa9e6b095b1818f188542b5d56e </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 3993 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=fping>
       <md5> 8f7ed0db109379a6547c21c26d1e00e5 </md5>
       <textSegmentSize> 17809 </textSegmentSize>
       <dataSegmentSize> 704 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 66152 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=freezer>
       <md5> 262490c35aaa4933cb0b5b2554456cfe </md5>
       <textSegmentSize> 2816 </textSegmentSize>
       <dataSegmentSize> 312 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 5 </numArguments>
       <size> 14215 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=fspscan>
       <md5> e248a6e47086d5a3ef4238f395aa44ca </md5>
       <textSegmentSize> 5237 </textSegmentSize>
       <dataSegmentSize> 396 </dataSegmentSize>
       <arguments> o-l-p-r-t-i </arguments>
       <numArguments> 1-5 </numArguments>
       <size> 17622 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=gtkftpd>
       <md5> 47081c74be9f704d7d317aad1af1ca1e </md5>
       <textSegmentSize> 9403 </textSegmentSize>
       <dataSegmentSize> 516 </dataSegmentSize>
       <arguments> -h-P-s-u-p-c-b-a-n-r-d </arguments>
       <numArguments> 1-11 </numArguments>
       <size> 22220 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=iawebmail>
       <md5> a87b4ffcd85a199911affa76df9d3d18 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 6316 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=ircd>
       <md5> a12e35c8195d88e27d52e4e9328b5a6f </md5>
       <textSegmentSize> 2459 </textSegmentSize>
       <dataSegmentSize> 316 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 13786 </size>
</ProcessSignature>
```


Sophos Ex. 1009
Page 275 of 318

```
<ProcessSignature rdf:ID=isspam>
       <md5> 420d323eb9c0cee5d575e9fd6557507e </md5>
       <textSegmentSize> 3384 </textSegmentSize>
       <dataSegmentSize> 916 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 16093 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=lsasrv>
       <md5> fff16d7ca96e0d7c1c3a734a6c0f23ea </md5>
       <textSegmentSize> 5653 </textSegmentSize>
       <dataSegmentSize> 2980 </dataSegmentSize>
       <arguments> -t </arguments>
       <numArguments> 1-5 </numArguments>
       <size> 22290 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=namedsploit>
       <md5> 98a2ff0deedb9750d7ae12ccf7137b53 </md5>
       <textSegmentSize> 6425 </textSegmentSize>
       <dataSegmentSize> 1296 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1-3 </numArguments>
       <size> 19027 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=niprint>
       <md5> 68f0946727b02511b5f921d32cb11561 </md5>
       <textSegmentSize> 2441 </textSegmentSize>
       <dataSegmentSize> 1140 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 14642 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=ntpd>
       <md5> 063e963f3d5cec5daf7e6936e6fd3740 </md5>
       <textSegmentSize> 4140 </textSegmentSize>
       <dataSegmentSize> 688 </dataSegmentSize>
       <arguments> -t </arguments>
       <numArguments> 2 </numArguments>
       <size> 16305 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=numby>
       <md5> 9fb7a5ebb06c48a284b9c2f1bfe3ec82 </md5>
       <textSegmentSize> 17308 </textSegmentSize>
       <dataSegmentSize> 488 </dataSegmentSize>
       <arguments> -m-v-analize-azconnect-azpost </arguments>
       <numArguments> 1-4 </numArguments>
       <size> 95188 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=nxt>
       <md5> 024c02fda0b3ef72975fe7f4d600ea75 </md5>
       <textSegmentSize> 5917 </textSegmentSize>
       <dataSegmentSize> 1104 </dataSegmentSize>
       <arguments> -t-s </arguments>
       <numArguments> 2-4 </numArguments>
       <size> 19458 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=pam>
       <md5> f12ce5a9fc714c6e8c2454cb1a274c37 </md5>
       <textSegmentSize> 4703 </textSegmentSize>
       <dataSegmentSize> 884 </dataSegmentSize>
       <arguments> -h-p </arguments>
       <numArguments> 1-2 </numArguments>
       <size> 17081 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=phorum>
       <md5> 6c1c7939bd91e17b8534abccc4f2c9d1 </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 0 </numArguments>
       <size> 2907 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=picobof>
       <md5> 3df483d4809936e885e7fd9dc6a54166 </md5>
       <textSegmentSize> 5058 </textSegmentSize>
       <dataSegmentSize> 308 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2-3 </numArguments>
       <size> 16449 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=priv8>
       <md5> 7196f50ff2b0c806cf549add4c7278db </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> -h-p-t-o </arguments>
       <numArguments> 4 </numArguments>
       <size> 3717 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=proftpd>
       <md5> e4a71ee5629a744abafc1b7e101b90bd </md5>
       <textSegmentSize> 11436 </textSegmentSize>
       <dataSegmentSize> 884 </dataSegmentSize>
       <arguments> -d-u-p-l-h-o-t </arguments>
       <numArguments> 7 </numArguments>
       <size> 25837 </size>
</ProcessSignature>
```


Sophos Ex. 1009
Page 276 of 318

```
<ProcessSignature rdf:ID=qex>
       <md5> 4ff02f4c30707d56e88799bbe44e9ec6 </md5>
       <textSegmentSize> 5631 </textSegmentSize>
       <dataSegmentSize> 400 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 3 </numArguments>
       <size> 17496 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=realown>
       <md5> 338fe8fa41a7cd1fbbb5cdf72ee66f1d </md5>
       <textSegmentSize> 3185 </textSegmentSize>
       <dataSegmentSize> 6800 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 21791 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=remo>
       <md5> 3e505342cdf4892297e405c61cb574f5 </md5>
       <textSegmentSize> 4116 </textSegmentSize>
       <dataSegmentSize> 2248 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 17741 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=rootdown>
       <md5> 97ace86f8075e71c2e306e393012020f </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> -h-c-i-p-r-v </arguments>
       <numArguments> 2-6 </numArguments>
       <size> 12416 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=rpcexec>
       <md5> 1ae30ab01b2cd2ba7b7eb3e0cf37adf3 </md5>
       <textSegmentSize> 4032 </textSegmentSize>
       <dataSegmentSize> 2397 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 17744 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=smurf4>
       <md5> 5e62a17140d5198587044c90c79d5b88 </md5>
       <textSegmentSize> 4346 </textSegmentSize>
       <dataSegmentSize> 432 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 5 </numArguments>
       <size> 16356 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=spxeneo>
       <md5> 9752ef6c1df794be31bc985d5faca28c </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 1399 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=ss>
       <md5> 1034ee8bece356c5370d39673a3fef1a </md5>
       <textSegmentSize> 0 </textSegmentSize>
       <dataSegmentSize> 0 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 1 </numArguments>
       <size> 17215 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=sslbomb>
       <md5> cc7bdc36853ee5fae400cde398a22a8c </md5>
       <textSegmentSize> 4142 </textSegmentSize>
       <dataSegmentSize> 1720 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2-3 </numArguments>
       <size> 4076 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=symanhoddos>
       <md5> 7703f81e3af39355c943f5917a9f6b58 </md5>
       <textSegmentSize> 4288 </textSegmentSize>
       <dataSegmentSize> 324 </dataSegmentSize>
       <arguments> -fi-tp-ti-n </arguments>
       <numArguments> 2-4 </numArguments>
       <size> 15852 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=telexp>
       <md5> 64c9619b45ef671a765ac9da6bc373d3 </md5>
       <textSegmentSize> 3573 </textSegmentSize>
       <dataSegmentSize> 956 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 3 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=tlsdos>
       <md5> 79c150d47e2bd0485552f710c588d899 </md5>
       <textSegmentSize> 4110 </textSegmentSize>
       <dataSegmentSize> 1720 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 2 </numArguments>
       <size> 3 </size>
</ProcessSignature>
```


Sophos Ex. 1009
Page 277 of 318

```
<ProcessSignature rdf:ID=vwexploit>
       <md5> fc0e8ed87f104ae8c93d7a886d752d70 </md5>
       <textSegmentSize> 3514 </textSegmentSize>
       <dataSegmentSize> 2516 </dataSegmentSize>
       <arguments> null </arguments>
       <numArguments> 4 </numArguments>
       <size> 17686 </size>
</ProcessSignature>

<ProcessSignature rdf:ID=yapping>
       <md5> cd9cb37c44a832d0661ec61dcdedab23 </md5>
       <textSegmentSize> 2496 </textSegmentSize>
       <dataSegmentSize> 312 </dataSegmentSize>
       <arguments> -h-s-c </arguments>
       <numArguments> 3 </numArguments>
       <size> 14137 </size>
</ProcessSignature>
```


Sophos Ex. 1009
Page 278 of 318

Appendix F: Network signatures

The table expanding the following pages corresponds to the network signatures of the programs surveyed in this project.

Sophos Ex. 1009
Page 279 of 318

```
<TrafficSignature rdf:ID=apache-nose>
    <packetCounter> 3 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 0 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 30 </deviation>
    <flags> "0,8,1,5,4,5,4,0,0,0,0,0,0,0,0" </flags>
    <lengths> "0,0,142,298,0,0,0,0,0,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,1132,2924,0,0,0,0,0,0,0,0,0,0,0" </hashes>
    <execVariant> 1a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=apache-nose>
    <packetCounter> 8 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 23 </destinationPort1>
    <destinationPort2> 9191 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 30 </deviation>
    <flags> "0,8,1,1,1,1,8,1,0,0,0,0,0,0,0" </flags>
    <lengths> "0,0,0,0,1032,0,0,0,0,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,0,0,1950,0,0,0,0,0,0,0,0,0,0" </hashes>
    <execVariant> 1b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=apache-scalp>
    <packetCounter> 16 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 0 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,1,1,1,1,1,1,1,1,1,1,1,1,1" </flags>
    <lengths> "0,0,0,0,1448,1448,1448,1448,1448,1448,1448,1448,1448,1448,1448" </lengths>
    <hashes> "0,0,0,0,2138,1950,1950,1950,1950,2013,1264,1950,1950,1950,1950,1950" </hashes>
    <execVariant> 2a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=apache-scalp>
    <packetCounter> 9 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 22 </destinationPort1>
    <destinationPort2> 0 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,1,1,1,1,1,1,1,0,0,0,0,0,0" </flags>
    <lengths> "0,0,0,1448,1448,1448,0,1448,1448,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,0,2138,1950,1950,0,1950,1950,0,0,0,0,0,0" </hashes>
    <execVariant> 2b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=apache-scalp>
    <packetCounter> 16 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 0 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,1,5,1,1,1,1,1,1,1,1,1,1,1" </flags>
    <lengths> "0,0,0,1448,1448,1448,1448,1448,1448,1448,1448,1448,1448,1448,1448" </lengths>
    <hashes> "0,0,0,2143,1950,1950,1950,1166,1950,1950,1950,1950,1950,1950,1950" </hashes>
    <execVariant> 2c </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=atphttp>
    <packetCounter> 10 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 65535 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,1,1,7,8,0,0,1,1,0,0,0,0,0" </flags>
    <lengths> "0,0,0,820,0,0,0,0,96,1232,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,0,1928,0,0,0,0,1562,1674,0,0,0,0,0,0" </hashes>
    <execVariant> 3a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=bazooka>
    <packetCounter> 5 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 0 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,8,8,8,0,0,0,0,0,0,0,0,0,0" </flags>
    <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
    <execVariant> 4a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=bazooka>
    <packetCounter> 8 </packetCounter>
    <fragmentation> 0 </fragmentation>
    <sourceAddress> -8205 </sourceAddress>
    <sourcePort> 0 </sourcePort>
    <destinationAddress> -5110 </destinationAddress>
    <destinationPort1> 80 </destinationPort1>
    <destinationPort2> 65535 </destinationPort2>
    <windowSize> 0 </windowSize>
    <deviation> 0 </deviation>
    <flags> "0,8,1,1,7,8,0,0,0,0,0,0,0,0,0" </flags>
    <lengths> "0,0,0,820,0,0,0,0,0,0,0,0,0,0,0" </lengths>
    <hashes> "0,0,0,1928,0,0,0,0,0,0,0,0,0,0,0" </hashes>
    <execVariant> 4b </execVariant>
</TrafficSignature>
```

```
<TrafficSignature rdf:ID=exceed>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 110 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,0,8,1,1,1,1,1,1" </flags>
       <lengths> "0,0,0,1448,1448,89,0,0,0,0,0,0,6,100,182,0" </lengths>
       <hashes> "0,0,0,2576,1950,2312,0,0,0,0,0,0,349,1512,977,977" </hashes>
       <execVariant> 5a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=exceed>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 1 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,7,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,1448,89,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2576,1950,2312,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 5b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=numby>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 110 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,7,8,1,1,1,1,1,1,1" </flags>
       <lengths> "0,0,0,266,0,0,190,0,0,0,0,6,100,182,0,33" </lengths>
       <hashes> "0,0,0,2838,0,0,1334,0,0,0,0,349,1306,1483,0,0" </hashes>
       <execVariant> 6a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=numby>
       <packetCounter> 5 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,4,1,1,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,266,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2838,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 6b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=pam>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 6023 </destinationPort1>
       <destinationPort2> 110 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,8,8,8,8,8,1,1,1,1,1,1,1,1" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,6,100,182,0,33,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,349,1223,1617,0,1689,1689" </hashes>
       <execVariant> 7a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=pam>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 6022 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,8,8,8,8,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 7b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=realown>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,8,1,1,1,1,8,1,1,1,1" </flags>
       <lengths> "0,0,0,34,0,0,0,0,39,0,0,0,0,48,0,0" </lengths>
       <hashes> "0,0,0,1759,0,0,0,0,1868,0,0,0,0,1814,0,0" </hashes>
       <execVariant> 8a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=realown>
       <packetCounter> 5 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 19 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,8,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 8b </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 281 of 318

```
<TrafficSignature rdf:ID=qex>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,27,0,538,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,286,0,839,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 9a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=qex>
       <packetCounter> 14 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 110 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,1,1,1,1,1,0,0" </flags>
       <lengths> "0,0,0,1448,1448,1448,1448,664,0,0,27,0,23,0,0,0" </lengths>
       <hashes> "0,0,0,2432,2484,2484,2484,2651,0,0,1720,0,1105,0,0,0" </hashes>
       <execVariant> 9b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=telexp>
       <packetCounter> 9 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 30 </deviation>
       <flags> "0,8,1,1,1,1,1,1,1,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,1448,0,1448,1448,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2432,2484,0,2484,2484,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 10a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=telexp>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 110 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 30 </deviation>
       <flags> "0,8,1,1,1,1,7,1,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,7,7,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,458,451,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 10b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=vwexploit>
       <packetCounter> 5 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -2322 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,4,4,8,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 11a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=305monit>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 23 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1314,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,3256,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 12a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=305monit>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1414,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,3256,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 12b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=akill2>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,7,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,462,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,1771,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 13a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 282 of 318

```
<TrafficSignature rdf:ID=anubisexp>
       <packetCounter> 12 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 23 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 30 </deviation>
       <flags> "0,8,5,1,5,1,5,1,1,1,1,0,0,0,0,0" </flags>
       <lengths> "0,0,0,27,538,3,3,5,0,0,163,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,386,23,23,32,83,0,0,213,0,0,0,0,0" </hashes>
       <execVariant> 14a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=arksink>
       <packetCounter> 11 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 110 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,7,0,0,8,1,7,1,0,0,0,0,0" </flags>
       <lengths> "0,0,0,18,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,939,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 15a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=bysin>
       <packetCounter> 6 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -9720 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,7,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,33,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2104,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 16a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=cge>
       <packetCounter> 2 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 2812 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 17a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=cobain>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,7,8,1,1,1,1,1,7,8" </flags>
       <lengths> "0,0,0,90,47,0,1448,133,0,0,90,0,47,1448,133,0" </lengths>
       <hashes> "0,0,0,2363,2343,0,2344,2823,0,0,2363,0,2343,2344,2823,2823" </hashes>
       <execVariant> 18a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=delefate>
       <packetCounter> 6 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -2897 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,490,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2434,2655,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 19a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=emule>
       <packetCounter> 13 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,2,2,2,2,2,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,23,234,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,134,2234,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 20a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=esignal>
       <packetCounter> 11 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 143 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,8,1,1,1,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,50,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,2239,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 21a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 283 of 318

```
<TrafficSignature rdf:ID=ethboom>
       <packetCounter> 9 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,7,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,1448,1448,381,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2773,3360,3360,3360,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 22a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=extremail2>
       <packetCounter> 6 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,8,8,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,558,550,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 23a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=ezploit>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 8080 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,2,7,7,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 24a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=fping>
       <packetCounter> 5 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 6004 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,0,0,1,0,1,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 25a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=freezer>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1308,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2728,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 26a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=fspscan>
       <packetCounter> 10 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,1448,1448,1186,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,2748,3330,3330,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 27a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=fspscan>
       <packetCounter> 9 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,7,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,1448,1186,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2748,3330,3330,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 27b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=gtkftpd>
       <packetCounter> 14 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 21 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,8,8,8,1,1,1,7,8,1,1,1,7" </flags>
       <lengths> "0,0,0,32,0,0,0,0,32,0,0,0,0,32,0,0" </lengths>
       <hashes> "0,0,0,210,0,0,0,0,210,0,0,0,0,213,0,0" </hashes>
       <execVariant> 28a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 284 of 318

```
<TrafficSignature rdf:ID=iawebmail>
       <packetCounter> 2 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -092 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 445 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 29a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=ircd>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 21 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,0,9,9,9,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,255,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2523,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 30a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=isspam>
       <packetCounter> 3 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 515 </destinationPort1>
       <destinationPort2> 7788 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,0,1,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 31a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=lsasrv>
       <packetCounter> 11 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,1,1,1,0,0,0,0,0" </flags>
       <lengths> "0,0,0,27,538,3,3,5,0,0,163,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,286,839,6,5,349,0,0,1713,0,0,0,0,0" </hashes>
       <execVariant> 32a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=namedsploit>
       <packetCounter> 14 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,7,8,1,1,1,1,1,7,8" </flags>
       <lengths> "0,0,0,90,47,0,1448,133,0,0,90,1448,180,0,0,0" </lengths>
       <hashes> "0,0,0,2363,2343,0,2344,2823,0,0,2363,2343,2910,0,0,0" </hashes>
       <execVariant> 33a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=niprint>
       <packetCounter> 2 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 53 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,7,0,0,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 34a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=nptd>
       <packetCounter> 10 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 65535 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,1,1,8,1,0,0,0,0" </flags>
       <lengths> "0,0,0,5,0,1448,1448,1448,252,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,532,0,3188,3360,3360,1520,0,0,0,0,0,0,0" </hashes>
       <execVariant> 35a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=nxt>
       <packetCounter> 4 </packetCounter>
       <fragmentation> 1 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,3,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,5,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,532,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 36a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 285 of 318

```
<TrafficSignature rdf:ID=nxt>
       <packetCounter> 11 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 65535 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,1,1,8,1,0,0,0,0,0" </flags>
       <lengths> "0,0,0,5,0,1448,1448,1200,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,532,0,3188,3360,3360,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 36b </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=phorum>
       <packetCounter> 2 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 21 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,0,0,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 37a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=picobof>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 65535 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,7,7,0,7,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,234,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,4023,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 38a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=priv8>
       <packetCounter> 2 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 135 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 39a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=rootdown>
       <packetCounter> 12 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 23 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 30 </deviation>
       <flags> "0,8,1,1,1,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,201,0,802,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2352,0,234,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 40a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=rpcexec>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 1 </destinationPort1>
       <destinationPort2> 2 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,8,8,8,8,8,8,8,8,8,8,8,8" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 41a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=smurf4>
       <packetCounter> 7 </packetCounter>
       <fragmentation> 1 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,3,1,1,1,6,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,132,0,1308,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,2145,0,913,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 42a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=spxeneo>
       <packetCounter> 9 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 22 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,0,1,0,1,1,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,308,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,925,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 43a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 286 of 318

```
<TrafficSignature rdf:ID=ss>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 6080 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,8,8,8,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 44a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=sslbomb>
       <packetCounter> 15 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 1 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,7,8,1,1,1,7,8,1,1,1,7" </flags>
       <lengths> "0,0,0,31,0,0,0,0,31,0,0,0,0,31,0,0" </lengths>
       <hashes> "0,0,0,2101,0,0,0,0,2102,0,0,0,0,2103,0,0" </hashes>
       <execVariant> 45a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=symanhoddos>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 617 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,8,1,1,1,1,1,1,1,7,1,0,0" </flags>
       <lengths> "0,0,0,0,0,1448,1448,1448,1448,1448,1448,1441,0,8,0,0" </lengths>
       <hashes> "0,0,0,0,0,3360,3360,3360,3360,3360,3360,3360,0,103,0,0" </hashes>
       <execVariant> 46a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=tlsdos>
       <packetCounter> 8 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 23 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,8,8,1,1,1,1,1,1,1,1,1,1,1,1" </flags>
       <lengths> "0,0,0,203,203,203,203,203,203,203,203,203,203,203,203,203" </lengths>
       <hashes> "0,0,0,3138,3950,3950,5950,5950,3950,1012,1265,1221,1221,1221,1221,1221" </hashes>
       <execVariant> 47a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=yapping>
       <packetCounter> 16 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 22 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,8,1,1,1,1,1,1,1,1,1,1,1,1" </flags>
       <lengths> "0,0,0,0,0,0,22,544,24,0,416,0,16,48,0,64" </lengths>
       <hashes> "0,0,0,0,0,0,0,1528,1728,104,0,1323,0,43,1774,0,0" </hashes>
       <execVariant> 48a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=proftpd>
       <packetCounter> 9 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 25 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,1,7,1,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,0,16,26,1309,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,0,1260,1936,1963,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 49a </execVariant>
</TrafficSignature>

<TrafficSignature rdf:ID=remo>
       <packetCounter> 10 </packetCounter>
       <fragmentation> 0 </fragmentation>
       <sourceAddress> -8205 </sourceAddress>
       <sourcePort> 0 </sourcePort>
       <destinationAddress> -5110 </destinationAddress>
       <destinationPort1> 80 </destinationPort1>
       <destinationPort2> 0 </destinationPort2>
       <windowSize> 0 </windowSize>
       <deviation> 0 </deviation>
       <flags> "0,8,1,1,1,1,7,0,0,0,0,0,0,0,0,0" </flags>
       <lengths> "0,0,0,1448,1448,1218,0,0,0,0,0,0,0,0,0,0" </lengths>
       <hashes> "0,0,0,1878,1890,1890,0,0,0,0,0,0,0,0,0,0" </hashes>
       <execVariant> 50a </execVariant>
</TrafficSignature>
```


Sophos Ex. 1009
Page 287 of 318

Bibliography

[1] Abad, C., Taylor, J., Sengul, C., and Yurcik, W. *Log Correlation for Intrusion Detection: A Proof of Concept.* 19th Annual Computer Security Applications Conference, vol. 1, ACSAC, pp. 255–265, Las Vegas, NV, December 2003.

[2] Allen, J., Christie, A., Fithen, W., McHugh, J., Pickel, J., and Stoner, E. *State of the Practice of Intrusion Detection Technologies.* Tech. Rep. 99-TR-028, ESC-99-028, Carnegie Mellon University, Software Engineering Institute, CMU/SEI, Pittsburgh, PA, 2000.

[3] Alonso, E. *AI and Agents: State of the Art.* The AI Magazine, pp. 25–29, vol. 23, num. 3, 2000.

[4] Amoroso, E. G. *Intrusion Detection: An Introduction to Internet Surveillance, Correlation, Trace Back, and Reponse.* Intrusion.net Books. February 1999.

[5] Anderson, D., Frivold, T., Tamaru, A., and Valdes, A. *Next-generation Intrusion Detection Expert System, NIDES.* Tech. Rep. SRI–CSL–95–07, Computer Science Laboratory, SRI International, Menlo Park, CA, May 1994.

[6] Anderson, J. P. *Computer Security threat Monitoring and Surveillance.* Tech. rep., James P. Anderson Company, Fort Washington, PA, Abril 1980.

[7] Antoniou, G., and van Harmelen, F. *Web Ontology Language: OWL.* Tech. rep., Department of Computer Science, University of Crete, and Department of Artificial Intelligence, Vrije Universitiet, 2003.

[8] Anzen. *AFJ Intrusion Detector*, Anzen Corp., June 2001. URL: http://www.anzen.comp/products/afj (on-line resource).

[9] Apap, F., Honig, A., Hershkop, S., Eskin, E., and Stolfo, S., Eds. *Detecting Maliciuos Software by Monitoring Anomalous Windows Registry Accesses* (October 2002), LNCS 2516, Recent Advances in Intrusion Detection, Springer.

[10] Arbaugh, W. A., Fithen, W. L., and McHugh, J. *Windows of Vulnerability: A Case Study Analysis.* IEEE Computer, pp. 52–59, num. 1, 2000.

[11] Asaka, M., Tsuchiya, M., Onabuta, T., Okasawa, S., and Goto, S. *Local Attack Detection and Intrusion Route Tracing.* IEICE Trans. Commun. New Paradigms in Network Management, pp. 1826–1833, vol. E82-B, num. 11, 1999.

[12] Avizienis, A., Laprie, J., and Randell, B. *Fundamental Concepts of Dependability*. Tech. Rep. N01145, LAAS-CNRS, April 2001.

[13] Axelsson, S. *Research in Intrusion-Detection Systems: A Survey*. Tech. Rep. TR: 98-17, Department of Computer Engineering, Chalmers University of Technology, Goteborg, Sweeden, 1998.

[14] Ayala, G. *Intelligent Agents for a Lifelong Learning Environment in Information Technologies*. International Journal of Continuing Engineering Education and Life-Long Learning, pp. 189–200, vol. 12, num. 1-4, 2000.

[15] Balasubramaniyan, J. S., Garcia-Fernandez, J. O., Isacoff, D., Spafford, E. H., and Zamboni, D. *An Architecture for Intrusion Detection using Autonomous Agents*. 14th Anual Computer Security Applications Conference, pp. 13–24, 1998.

[16] Bass, T. *Intrusion Detection Systems and Multisensor Data Fusion*. Communications of the ACM, pp. 99–105, vol. 1, num. 1, 1999.

[17] Bellifemine, F., Poggi, A., and Rimassa, G. *JADE - A FIPA-compliant Agent Framework*. PAAM '99, pp. 97–108, London, UK, 1999.

[18] Bishop, M. *Computer Security: Art and Science*, 1st ed. Addison-Wesley Pub. Co. December 2002.

[19] Borselius, N. *Mobile Agent Security*. Electronics & Communication Engineering Journal, pp. 211–218, vol. 14, num. 5, 2002.

[20] Breiling, S., Plato, A., and Shulak, B. *BlackICE Sentry Signature Reference Guide*. Internet Security Systems, ISS, June 2001.

[21] Brena, R., Aguirre, J. L., and Trevino, A. C. *Just-in-time Information and Knowledge: Agent Technology for KM Bussiness Processes*. 2001 IEEE Systems, Man, and Cybernetics Conference, September 2001.

[22] Bridges, S. M., and Vaughn, R. B. *Intrusion Detection via Fuzzy Data Mining*. 12th Annual Canadian Information Technology Security Symposium, pp. 109–122, Ottawa, Canada, June 2000.

[23] Brugali, D., and Sycara, K. *Towards Agent Oriented Application Frameworks*. ACM Computing Surveys, pp. 21–26, vol. 32, num. 1, 2000.

[24] Bugtraq. *Bugtraq Vulnerability Repository*, Securityfocus.com, 2003. URL: http://www.securityfocus.com (on-line resource).

[25] Bussmann, S., and Müller, J. *A Communication Architecture for Cooperating Agents*. Computers and Artificial Intelligence, pp. 37–53, vol. 12, num. 1, 1993.

[26] CA. *E-Trust Intrusion Detection*, Computer Associates, Inc., December 2003. URL: http://www3.ca.com/ (on-line resource).

[27] Cannady, J. *Artificial Neural Networks for Misuse Detection*. National Information Systems Security Conference, NISSC, no. 1, pp. 443–456, Arlington, VA, October 1998.

[28] Cansian, A. M., de Souza, A. Z. T., Filho, S. L., and Moreira, E. S. *A Password-capture System for the Security of Computer Networks*. Department of Computer Science and Statistics, UNESP, Brazil, 1999.

[29] Captus Networks. *CaptIO and CaptIO-G*, Captus Networks, Inc., May 2004. URL: http://www.captusnetworks.com/ (on-line resource).

[30] Carroll, J. J., Dickinson, I., Dollin, C., Reynolds, D., Seaborne, A., and Wilkinson, K. *Jena: Implementing the Semantic Web Recommendations*. Tech. Rep. HPL-2003-146, HP Laboratories Bristol, Bristol, UK, December 2003.

[31] Ceballos-Cancino, H. G. *Ontology Administration in Multiagent Environments using an Ontology Agent applied to JITIK (in Spanish)*. Master's thesis, ITESM–Monterrey. Monterrey, Mexico, 2003.

[32] CERT. *Statistics of the Computer Emergency Response Team Coordination Center, CERT–CC*, Carnegie Mellon University, 2004. URL: http://www.cert.org/stats/ (on-line resource).

[33] Christey, S. M., and Martin, R. A. *A Progress Report on the CVE Initiative*. 14th Annual Computer Security Incident Handling Conference, Kona, HA, June 2002.

[34] Christoph, G. G., Jackson, K. A., Neuman, M. C., Siciliano, C. L., Simmonds, D. D., Stallings, C. A., and Thopson, J. L. *UNICORN: Misuse Detection for UNICOS*. ACM/IEEE Supercomputing Conference, no. 56, ACM Press and IEEE Computer Society Press, December 1995.

[35] CISCO. *The Abacus Project*, CISCO Systems, Inc., June 2001. URL: http://www.digital.com/info/security/id.htm (on-line resource).

[36] CISCO. *Net Ranger*, CISCO Systems, Inc., January 2004. URL: http://www.cisco.com/univercd/cc/td/doc/product/iaabu/netrangr/ (on-line resource).

[37] Coit, C. J., Staniford, S., and McAlerney, J. *Towards Faster String Matching for Intrusion Detection or Exceeding the Speed of Snort*. DARPA Information Survivability Conference and Exposition (DISCEX II'01), vol. 1, Anaheim, CA, June 2001.

[38] Comer, D. E., and Stevens, D. L. *Internetworking with TCP/IP series*, 1st ed. Prentice Hall. Upper Saddle River, NJ, 1997.

[39] Coull, S., Branch, J., Szymanski, B., and Breimer, E. *Intrusion Detect: A Bioinformatics Approach*. 19th Annual Computer Security Applications Conference, pp. 24–33, Las Vegas, NV, December 2003.

[40] Cowan, C., Barringer, M., Beattie, S., Kroah-Hartman, G., Frantzen, M., and Lokier, J. *Formatguard: Automatic Protection from Printf Format String Vulnerabilities*. USENIX Security Symposium, pp. 191–199, Washington, DC, August 2001.

[41] Crosbie, M., and Spafford, E. H. *Defending a Computer System using Autonomous Agents*. 18th National Information Systems Security Conference, pp. 549–558, October 1995.

[42] Curry, D., Klaus, C., Ranum, M., Spafford, E., and Sutterfieldi, L. *CSI Roundtable: Experts Discuss Present and Future of Intrusion Detection Systems*. Computer Security Journal, vol. 14, num. 1, 1998.

[43] Curry, D. A., Debar, H., and Feinstein, B. S. *Intrusion Detection Message Exchange Format*. Intrusion Detection Working Group – Internet Engineering Task Force, Internet Draft, January 2004.

[44] CyberTrace. *CyberTrace Network Security Management*, CyberTrace, Inc., February 2004. URL: http://www.cybertrace.com/ctids.html (on-line resource).

[45] Dasgupta, P., Moser, L. E., and Melliar-Smith, P. M. *The Security Architecture for MAgNET: A Mobile Agent E-commerce System*. 3rd International Conference on Telecommunications and E-commerce, pp. 289–298, Dallas, TX, November 2000.

[46] Debar, H., Becker, M., and Siboni, D. *A Neural Network Component for an Intrusion Detection System*. IEEE Computer Society Symposium on Research in Security and Privacy, Los Alamitos, CA, pp. 240–250, Oakland, CA, May 1992.

[47] DEC. *POLICENTER Security Compliance Manager*, Digital Equipment Corporation, Inc., August 1994. URL: http://www.digital.com/info/security/id.htm (on-line resource).

[48] Deloitte and Touche. *2003 Global Security Survey*, Deloitte and Touche, Inc., May 2003. URL: http://www.deloitte.com/dtt/cda/doc/content (on-line resource).

[49] Denning, D. E. *An Intrusion Detection Model*. IEEE Transactions on Software Engineering, no. 2 in 13, pp. 222–232, February 1987.

[50] Denning, D. E. *A Prototype IDES: A Real-time Intrusion Detection Expert System*. Tech. rep., Computer Science Laboratory, SRI International, Menlo Park, CA, August 1987.

[51] Deswarte, Y., Blain, L., and Fabre, J.-C. *Intrusion Tolerance in Distributed Computing Systems*. IEEE Symposium on Security and Privacy, pp. 110–121, 1991.

[52] DOE-CIAC. *DOE-CIAC's Network Intrusion Detector (NID)*, National Nuclear Security Administration, April 2004. URL: http://ciac.llnl.gov/cstc/nid/nid.html (on-line resource).

[53] Dowel C, R. *The Computer Watch Data Reduction Tool*. 13th National Computer Security Conference, pp. 99–108, Washington, DC, October 1990.

[54] Durfee, E. H., and Rosenschein, J. *Distributed Problem Solving and Multiagent Systems: Comparisons and Examples*. 13th International Workshop on DAI, M. Klein, Ed., pp. 94–104, Lake Quinalt, WA, 1994.

[55] Edwards, M. J. *Kane Security Monitor 3.03*. Windows Network and .NET Magazine, num. 1, 1998 (on-line resource.

[56] Endler, D. *Intrusion Detection Applying Machine Learning to Solaris Audit Data*. Annual Computer Security Applications Conference, ACSAC, IEEE Computer Society Press, pp. 268–279, Scottsdale, AZ, 1998.

[57] EnGarde. *T-sight: On Target Security*, ENGARDE Systems, April 2004. URL: http://www.engarde.com/software/t-sight/ (on-line resource).

[58] Enterasys. *Dragon Intrusion Detection Defense System*, Enterasys Intrusion Defense, Inc., April 2004. URL: http://www.enterasys.com/products/ids/ (on-line resource).

[59] Erlingsson, and Schneider. *SASI Enforcement of Security Policies: A Retrospective*. WNSP: New Security Paradigms Workshop, ACM Press, pp. 87–95, 2000.

[60] Escamilla, T. *Intrusion Detection: Network Security Beyond the Firewall*. John Wiley & Sons, Inc. New York, NY, 1998.

[61] Exploit Database. *Full-disclosure Exploit Database*, Black Sheep Networks, Inc., 2003. URL: http://www.exploitdatabase.com (on-line resource).

[62] Farmer, W. M., Guttman, J. D., and Swarup, V. *Security for Mobile Agents: Authentication and State Appraisal*. 4th European Symposium on Research in Computer Security, pp. 118–130, Rome, Italy, 1996.

[63] Fedoruk, A., and Deters, R. *Improving Fault-Tolerance by Replicating Agents*. 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems, AAMAS, pp. 737–744, Bologna, Italy, July 15–19 2002.

[64] Forrest, S., Hofmer, S. A., Somayaji, A., and Longstaff, T. A. *A Sense of Self for UNIX Processes*. IEEE Symposium on Security and Privacy, IEEE Press, Oakland, CA, May 1996.

[65] Frank, J. *Artificial Intelligence and Intrusion Detection: Current and Future Directions*. 17th National Computer Security Conference, pp. 22–33, October 1994.

[66] Frincke, D., and Ho, Y. *Planning, Petri Nets, and Intrusion Detection*. 21st National Information Systems Security Conference, pp. 346–361, October 1998.

[67] Garfinkel, S., and Spafford, E. H. *Practical UNIX Security*, 2nd ed. O'Reilly. April 1996.

[68] Garvey, T. D., and Lunt, T. F. *Model-based Intrusion Detection*. 14th National Computer Security Conference, NIST, National Institute of Standards and Technology/National Computer Security Center, pp. 372–385, Baltimore, MD, October 1998.

[69] GDataMine. *Gnome Data Mining Tool*, GNOME, August 2003. URL: http://www.togaware.com/datamining/gdatamine (on-line resource).

[70] GFI. *LANguard*, GFI Security, March 2004. URL: http://www.gfi.com/languard/ (on-line resource).

[71] Giffin, J. T., Jha, S., and Miller, B. P. *Efficient Context-Sensitive Intrusion Detection*. 11th Annual Network and Distributed System Security Symposium (NDSS), San Diego, CA, February 2004.

[72] Goldberg, I., Wagner, D., Thomas, R., and Brewer, E. A. *A Secure Environment for Untrusted Helper Applications*. 6th USENIX Security Symposium, pp. 1–13, San Jose, CA, July 1996.

[73] Goldschlag, D., Reed, M., and Syverson, P. *Onion Routing for Anonymous and Private Internet Connections*. Communications of the ACM, pp. 39–41, vol. 42, num. 2, 1999.

[74] Guan, X., Yang, Y., and You, J. *POM – A Mobile Agent Security Model against Malicious Hosts*. 4th International Conference on High-Performance Computing in the Asia-Pacific Region, vol. 2, pp. 1165–1166, May 2000, Beijing, China.

[75] Habra, N., Charlier, B. L., Mounji, A., and Mathieu, I. *ASAX: Software Arquitecture and Rule-based Language for Universal Audit Trail Analysis*. European Symposium on Research in Computer Security, ESORICS, pp. 435–450, num. 1, 1992.

[76] Halme, L. R., and Bauer, K. R. *A Taxonomy of Anti-intrusion Techniques*. 18th National Information Systems Security Conference, pp. 163–172, Baltimore, MD, October 1995.

[77] Heberlein, T., Dias, G., Levitt, K., Wood, B. M. J., and Wolber, D. *A Network Security Monitor*. IEEE Symposium on Research in Security and Privacy, IEEE Press, pp. 296–304, Los Alamitos, CA, May 1990.

[78] Helmer, G. G., Wong, J., Honovar, V., and Miller, L. *An Architecture for Intrusion Detection using Autonomous Agents*. Tech. rep., Department of Computer Science, Iowa State University, Des Moines, IA, 1999.

[79] Hewlett Packard. *HP IDS 9000*, Hewlett Packard, Inc., June 2001. URL: http://www.hp.com/security/products/ids (on-line resource).

[80] Hirata, K. *A Classification of Abduction: Abduction for Logic Programming*. Machine Intelligence, Oxford University Press, pp. 397–424, vol. 1, num. 14, 1995.

[81] Hochberg, J., Jackson, K., Stallings, C., McClary, J. F., DuBois, D., and Ford, J. *NADIR: An Automated System for Detecting Network Intrusion and Misuse*. Elsevier Advanced Technology Publications, pp. 235–248, vol. 12, num. 3, 1993.

[82] Hofmeyr, S. A., and Forrest, S. *Architecture for an Artificial Immune System.* Evolutionary Computation, pp. 443–473, vol. 8, num. 4, 2000.

[83] Hofmeyr, S. A., Forrest, S., and Somayaji, A. *Intrusion Detection using Sequences of System Calls.* Journal of Computer Security, pp. 151–180, vol. 6, num. 1, 1998.

[84] Houle, K. H., and Weaver, G. M. *Trends in Denial of Service Attack Technology.* Tech. rep., CERT Coordination Ceter, Pittsburgh, PA, October 2001.

[85] Huagang, X. *Build a Secure System with LIDS.* LinuxSecurity.com, num. 1, 2000 (on-line resource).

[86] Huhns, M. N., and Singh, M. P. *Agents and Multiagent systems: Themes, Approaches, and Challenges.* Readings in Agents, pp. 1–23, num. 1, 1998.

[87] Ilgun, K. *USTAT: A Real-Time Intrusion Detection System for UNIX.* IEEE Symposium on Security and Privacy, pp. 16–28, Oakland, CA, May 1993.

[88] Intouch INSA. *Network Security Agent*, Intouch INSA, February 2004. URL: http://www.ttinet.com (on-line resource).

[89] ISS. Internet Security Systems, Inc., July 2003. URL: http://www.iss.net (on-line resource).

[90] Jackson, K. A. *Intrusion Detection System (IDS) Product Survey.* Tech. Rep. LA-UR-99-3883, Los Alamos National Laboratory, Los Alamos, NM, June 1999.

[91] Jackson, K. A., DuBois, D. H., and Stallings, C. A. *An Expert System Application for Network Intrusion Detection.* 14th National Computer Security Conference, National Institute of Standards and Technology/National Computer Security Center, pp. 215–225, Washington, DC, October 1991.

[92] Janakiraman, R., Waldvogel, M., and Zhang, Q. *INDRA: A Peer-to-Peer Approach to Network Intrusion Detection and Prevention.* IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, WETICE, June 2003.

[93] Jansen, W., Mell, P., Karygiannis, T., and Marks, D. *Mobile Agents in Intrusion Detection and Response.* 12th Annual Canadian Information Technology Security Symposium, Ottawa, Canada, June 2000.

[94] Jensen, F. V. *Bayesian Networks and Decision Graphs.* Springer Verlag Ed. Berlin, Germany, 2001.

[95] Jou, Y. F., Gong, F., Sargor, C., Wu, S. F., and Cleaveland, W. R. *Architecture Design of a Scalable Intrusion Detection System for the Emerging Network Infrastructure.* Tech. Rep. CDRLA-005, Department of Computer Science, North Carolina State University, Raleigh, NC, April 1997.

[96] JPCAP. *Java Packet Capture library*, Java Foundry, February 2003. URL: http://sourceforge.net/projects/jpcap/ (on-line resource).

[97] Kemmerer, R. A. *NSTAT: A Model-based Real-time Network Intrusion Detection System*. Tech. Rep. TR-CS-97-18, Reliable Software Group, Department of Computer Science, University of California-Santa Barbara, Santa Barbara, CA, November 1997.

[98] Kerschbaum, F., Spafford, E. H., and Zamboni, D. M. *Using Internal Sensors and Embedded Detectors for Intrusion Detection*. Journal of Computer Security, vol. 10, pp. 23–70, 2002.

[99] Kihlstrom, K. P., and Narasimhan, P. *The Starfish System: Providing Intrusion Detection and Intrusion Tolerance for Middleware Systems*. IEEE Workshop on Object-oriented Real-time Dependable Systems, Guadalajara, Mexico, January 2003.

[100] Kim, E. H., and Spafford, E. H. *The Design and Implementation of TRIPWIRE: A File System Integrity Checker*. 2nd ACM Conference on Computers and Communications Security, pp. 18–29, Fairfax, VA, November 1994.

[101] Klevinsky, T. J., Laliberte, S., and Gupta, A. *Hack I.T. – Security through Penetration Testing*, 1st ed. Addison-Wesley Longman. February 2002.

[102] Ko, C., Fink, G., and Levitt, K. *Automated Detection of Vulnerabilities in Privileged Programs by Execution Monitoring*. 10th Annual Computer Security Applications Conference, IEEE Computer Society Press, pp. 134–144, Orlando, FL, 1994.

[103] Kumar, S. *Classification and Detection of Computer Intrusions*. PhD thesis, Department of Computer Sciences, Purdue University. West Lafayette, IN, 1995.

[104] Kumar, S., and Spafford, E. H. *An Application of Pattern Matching in Intrusion Detection*. Tech. Rep. CSD–TR–94–013, Department of Comptuter Sciences, Purdue University, West Lafayette, IN, June 1994.

[105] Kumar, S., and Spafford, E. H. *A Pattern Matching Model for Misuse Intrusion Detection*. 17th National Computer Security Conference, pp. 11–21, Baltimore, MD, October 1994.

[106] Kumar, S., and Spafford, E. H. *A Software Architecture to Support Misuse Intrusion Detection*. 18th National Information Security Conference, pp. 194–204, October 1995.

[107] Kuri, J., Navarro, G., M'e, L., and Heye, L. *A Pattern Matching Based Filter for Audit Reduction and Fast Detection of Potential Intrusions*. Recent Advances in Intrusion Detections conference (RAID '00), pp. 17–27, Berlin, Germany, 2000.

[108] Lane, T., and Brodley, C. E. *Temporal Sequence Learning for Anomaly Detection*. 5th ACM Conference on Computer and Communications Security, pp. 150–158, San Francisco, CA, 1998.

[109] Lane, T., and Brodley, C. E. *Temporal Sequence Learning and Data Reduction for Anomaly Detection.* ACM Transactions on Information Systems Security, pp. 295–331, vol. 2, num. 3, 1999.

[110] Lange, D. B., and Oshima, M. *Programming and Deploying Java Mobile Agents with Aglets.* Addison–Wesley Press. Menlo Park, CA, 1998.

[111] Lau, F., Rubin, S. H., Smith, M. H., and Trajovic, L. *Distributed Denial of Service Attacks.* IEEE International Conference on Systems, Man, and Cybernetics, pp. 2275–2280, Nashville, TN, October 2000.

[112] Lee, W., Nimbalkar, R. A., Yee, K. K., Patil, S. B., Desai, P. H., Tran, T. T., and Stolfo, S. J. *A Data Mining and CIDF Based Approach for Detecting Novel and Distributed Intrusions.* Lecture Notes in Computer Science, pp. 49–65, vol. 1907, num. 1, 2000.

[113] Lee, W., and Stolfo, S. *Data Mining Approaches for Intrusion Detection.* 7th USENIX Security Symposium, pp. 79–93, San Antonio, TX, 1998.

[114] Lee, W., Stolfo, S. J., and Mok, K. W. *A Data Mining Framework for Building Intrusion Detection Models.* IEEE Symposium on Security and Privacy, pp. 120–132, 1999.

[115] Lunt, T., Tamaru, A., Gilham, F., Jagannathan, R., Neumann, P., Javitz, H., Valdes, A., and Garvey, T. *A Real-time Intrusion Detection Expert System (IDES).* Tech. rep., SRI International, Computer Science Laboratory, February 1992.

[116] Lynch, N. *Distributed Algorithms.* Morgan Kaufmann Publishers. San Mateo, CA, 1996.

[117] Mandujano, S. *Toward the Preservation of a Key Feature of the Internet: Policy and Technology for Internet Anonymity.* International Conference on Technology Policy and Innovation, ICTPI '02, Monterrey, Mexico, 2002.

[118] Mandujano, S., and Galván, A. *ODISET: Online Distributed Session Tracing using Agents.* 18th International Joint Conference on Artificial Intelligence, IJCAI '03, pp. 749–754, Acapulco, Mexico, August 2003.

[119] Mandujano, S., and Galván, A. *Outbound Intrusion Detection.* International Conference on Computing, Communications and Control Technologies, CCCT '04, vol. 1, pp. 68–73, Austin, TX, August 2004.

[120] Mandujano, S., and Soto, R. *Deterring Password Sharing: User Authentication via Fuzzy c-Means applied to Keystroke Profiles.* 5th ENC International Conference in Computer Science, ed. IEEE Computer Society, pp. 181–187, Colima, Mexico, September 2004.

[121] Mann, D. E., and Christey, S. M. *Towards a Common Enumeration of Vulnerabilities.* 2nd Workshop on Research with Security Vulnerability Databases, Purdue University, West Lafayette, IN, January 1999.

[122] Marin, O., Sens, P., Briot, J. P., and Guessoum, Z. *Towards Adaptive Fault Tolerance for Distributed Multiagent Systems*. ERSADS 2001, pp. 195–201, Bertinoro, Italy, 2001.

[123] McAfee. *Entercept Security: Advanced e-Server Protection*, McAfee Security, Inc., February 2004. URL: http://www.nai.com/us/index.asp (on-line resource).

[124] McCarthy, L. *Symantec Internet Security Threat Report*. Tech. rep., Symantec, Inc., September 2003.

[125] McClure, S., Scambray, J., and Kurtz, G. *Hacking Exposed: Network Security Secrets & Solutions*, 4th ed. McGraw-Hill. February 2003.

[126] McHugh, J. *Intrusion and Intrusion Detection*. International Journal of Information Security, pp. 14–35, vol. 1, num. 1, 2001.

[127] Mell, P. *ICAT Metabase and CVE Vulnerability Search Engine*, Computer Security Division Information Technology Laboratory National Institute of Standards and Technology, NIST, 2000.

[128] Meunier, P., Mandujano, S., Song, W. F., and Ma, L. *Cooperativy Vulnerability Database*. Tech. rep., Center for Education and Research in Information Assurance and Security, CERIAS, Purdue University, West Lafayette, IN, February 2001.

[129] MimeStar. *SecureNet Pro*, MimeStar Intrusion Detection Co., March 2004. URL: http://www.mimestar.com/products/ (on-line resource).

[130] Mitchell, T. M. *Machine Learning*. McGraw Hill. 1st ed., 1997.

[131] Moitra, A. *Real-time Audit Log Viewer and Analyzer*. 4th Workshop on Computer Security Incident Handling, (Forum of Incident Response and Security Teams, FIRST), Denver, CO, August 1992.

[132] Moore, A. P., Ellison, R. J., and Linger, R. C. *Attack Modeling for Information Security and Survivability*. Tech. Rep. SEI-2001-TN-001, Survivable Systems, Carnegie Mellon University, Software Engineering Institute (SEI), Pittsburgh, PA, March 2001.

[133] Morini, P. *CERN Network Security Monitor*. 1st Workshop on Recent Advances in Intrusion Detection, Louvain la Neuve, Belgium, September 1998.

[134] Mouratidis, H., Giorgini, P., and Manson, G. *An Ontology for Modeling Security: The Tropos Approach*. Tech. rep., Department of Computer Science, University of Sheffield, Sheffield, UK, 2003.

[135] Mukkamala, S., Janowsky, G., and Sung, A. H. *Intrusion Detection using Neural Networks and Support Vector Machines*. IEEE International Joint Conference on Neural Networks 2002, pp. 1702–1707, Hawaii, May 2002.

[136] Mutz, D., Vigna, G., and Kemmerer, R. *An Experience Developing an IDS Stimulator for the Black-Box Testing of Network Intrusion Detection Systems*. 19th Annual Computer Security Applications Conference, ACSAC, pp. 374–383, Las Vegas, Nevada, December 2003.

[137] Neto, H. V., and Borges, D. L. *Fingerprint classification with neural networks*. 4th Brazilian Symposium on Neural Networks, SBRN, IEEE CS Press, pp. 66–72, Campos do Jordao, Brazil, December 1997.

[138] NETWORK ASSOCIATES. *Cybercop Monitor*, Network Associates, Inc., February 2004. URL: http://www.tlic.com/security/cybercopmonitor.cfm (on-line resource).

[139] Network Solutions. *DefenseWorx*, Network Solutions, Inc., March 2004. URL: http://www.defenseworx.com (on-line resource).

[140] Neumann, P. G., and Porras, P. A. *Experience with EMERALD to Date*. 1st USENIX Workshop on Intrusion Detection and Network Monitoring, pp. 73–80, Santa Clara, CA, April 1999.

[141] NFR. *Centrax: Hybrid Security Solution*, NFR Security Corp., March 2001. URL: http://www.nfr.com/ (on-line resource).

[142] Northcutt, S., Cooper, M., Fearnow, M., and Frederick, K. *Intrusion Signatures and Analysis*, 1st ed. New Riders, SANS GIAC. Indianapolis, IN, January 2001.

[143] Northcutt, S., and Novak, J. *Network Intrusion Detection*, 3rd ed. New Riders. Indianapolis, IN, September 2002.

[144] Noy, N. F., Sintek, M., Decker, S., Crubezy, M., Fergerson, R. W., and Musen, M. A. *Creating Semantic Web Contents with Protege-2000*. IEEE Intelligent Systems, vol. 16(2), pp. 60–71, 2000.

[145] NSWC. *Shadow ID*, Naval Surface Warfare Center, December 2003. URL: http://www.nswc.navy.mil/ISSEC/CID/ (on-line resource).

[146] O'Hare, G. M. P., and Jennings, N. R. *Foundations of Distributed Artificial Intelligence*. John Wiley & Sons. New York, NY, 1996.

[147] Okazaki, Y., and Sato, I. *A New Intrusion Detection Method based on Process Profiling*. IEEE Symposium on Applications and the Internet (SAINT), pp. 82–91, Nara, Japan, January 2002.

[148] OKENA. *Stormwatch, white paper*, Okena Systems Co., October 2000. URL: http://www.okena.com/products/literature.htm (on-line resource).

[149] OSVDB. *Open Source Vulnerability Database*, May 2004. URL: http://www.osvdb.org (on-line resource).

[150] PacketStorm. *Packet Storm Security*, June 2003. URL: http://www.packetstormsecurity.org (on-line resource).

[151] Parunak, H. V. *An Overview of Distributed Artificial Intelligence.* Foundations of Distributed Artifical Intelligence, 6th Generation Computer Technology Series, G. OHgare and N. Jennings, Eds., John Wiley and Sons, pp. 3–55, June 1996.

[152] Paxson, V. *Bro: A System for Detecting Network Intruders in Real-time.* Elsevier Computer Networks, pp. 2435–2463, vol. 31, num. 23–24, 1998.

[153] Pérez, J. A., Álvarez, D., and Sobrado, I. *A Fast Data Protection Technique for Mobile Agents against Malicious Hosts.* Electronic Notes in Theoretical Computer Science, K. Fischer and D. Hutter, Eds., vol. 63, Elsevier, 2002.

[154] Pescatore, J. *Nimda Worm Shows You Can't Patch Fast Enough*, Gartner Group, September 2001. URL: http://www.emailsolutions.com/gwnews/gartner_iis.pdf (on-line resource).

[155] PestPatrol. *Denial of Service Attacks on the Rise, whitepaper*, PestPatrol, Inc., July 2002. URL: http://www.pestpatrol.com/WhitePapers (on-line resource).

[156] Pfleeger, C. P. *Security in Computing*, 2nd ed. Prentice Hall, Inc. Upper Saddle River, NJ, 1997.

[157] Pleisch, S., and Schiper, A. *Modeling Fault-Tolerant Mobile Agent Execution as a Sequence of Agreement Problems.* 9th IEEE Symposium on Reliable Distributed Systems (SRDS 2000), IEEE Computer Society, pp. 11–20, Nurnberg, Germany, 2000.

[158] Polla, D., McConnell, J., Johnson, T., Marconi, J., Tobin, D., and Frincke, D. *A FrameWork for Cooperative Intrusion Detection.* 21st National Information Systems Security Conference, pp. 361–373, October 1998.

[159] Poole, D. L., MackWorth, A. L., and Goebel, R. G. *Computational Intelligence: A Logical Approach.* Oxford University Press. New York, NY, 1998.

[160] PRC. *PRC-Precis whitepaper*, PRC Systems, 1999. URL: http://www.bellevue.prc.com/precis/solution.pdf (on-line resource).

[161] Ptacek, T. H., and Newsham, T. N. *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection.* Tech. rep., Secure Networks, Inc., January 1998.

[162] Quin, L. *Extensible Markup Language (XML).* Tech. rep., World Wide Web Consortium (W3C), 2002.

[163] Ranum, M. J. *Full Disclosure and Open Source*, Black Hat Briefings, keynote speech, 2000.

[164] Ranum, M. J., Landfield, K., Stolarchuk, M., Sienkiewicz, M., Lambeth, A., and Wall, E. *Implementing a Generalized Tool For Network Monitoring.* 11th Systems Administration Conference (LISA), USENIX, pp. 1–8, San Diego, CA, October 1997.

[165] Rapalus, P. *Computer Crime & Security Survey '02*, Computer Security Institute (CSI) and Federal Business of Investigations (FBI), April 2002.

[166] Raskin, V., Hempelmann, C. F., Triezenberg, K. E., and Nirenburg, S. *Ontology in Information Security: a Useful Theoretical Foundation and Methodological Tool*. New Security Paradigms Workshop, ACM Press, pp. 53–59, Cloudcroft, NM, 2001.

[167] Rathmell, A., and Valeri, L. *Information Warfare and the Asymmetric Threat: An approach to early warning*. Tech. rep., International Centre for Security Analysis, King's College London, London, UK, 1997.

[168] Reiter, M., and Rubin, A. *Anonymous Web Transactions with Crowds*. Communications of the ACM, pp. 32–38, vol. 42, num. 2, 1999.

[169] Richardson, R. *Computer Crime & Security Survey '03*, Computer Security Institute (CSI) and Federal Business of Investigations (FBI), April 2003.

[170] Roesch, M. *Snort: Lightweight Intrusion Detection for Networks*. 13th Conference on Systems Administration, pp. 229–238, Berkeley, CA, 1999.

[171] Ru, W. G., and Eloff, J. H. *Enhanced Password Authentication through Fuzzy Logic*. IEEE Expert, vol. 12–6, pp. 38–45, Nov/Dec 1997.

[172] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*, 1st ed. Prentice Hall. Upper Saddle River, NJ, January 1995.

[173] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*, 2nd ed. Prentice Hall. Upper Saddle River, NJ, 2003.

[174] Ryan, J., Lin, M.-J., and Miikkulainen, R. *Intrusion Detection with Neural Networks*. Advances in Neural Information Processing Systems, M. I. Jordan, M. J. Kearns, and S. A. Solla, Eds., vol. 10, The MIT Press, pp. 943–949, Denver, CO, 1998.

[175] Samaha, S. E. *Haystack: An Intrusion Detection System*. 4th Aerospace Computer Security Applications Conference, vol. XII(440), IEEE Computer Society Press, pp. 37–44, Washington, DC, 1988.

[176] Sanctum. *AppShield: Automatic Web Application Control and Security*, Sanctum, Inc., January 2004. URL: http://www.sanctuminc.com/solutions/appshield/ (on-line resource).

[177] Sander, T., and Tschudin, C. F. *Protecting Mobile Agents Against Malicious Hosts*. Lecture Notes in Computer Science, LNCS, pp. 44–49, vol. 1419, num. 1, 1998.

[178] Schneider, F. B. *Towards Fault-tolerant and Secure Agentry*. 11th International Workshop on Distributed Algorithms, WDAG, M. Mavronicolas and P. Tsigas, Eds., Springer Verlag, Berlin, Germany, 1997.

[179] Schneier, B. *Applied Cryptography*, 2nd ed. John Wiley and Sons. New York, NY, 1996.

[180] Schneier, B. *Managed Security Monitoring: Network security for the 21st century.* Computers and Security, pp. 491–503, vol. 20, num. 6, 2001.

[181] Sebring, M. M., Shellhouse, E., Hanna, M. E., and Whitehurst, R. A. *Expert systems in intrusion detection: A case study.* National Institute of Standards and Technology/National Computer Security Center, 11th National Computer Security Conference, pp. 74–81, Baltimore, MD, October 1988.

[182] Sebring, M. M., Shellhouse, E., Hanna, M. E., and Whitehurst, R. A. *MIDAS: Multics Intrusion Detection and Alerting System.* Tech. rep., National Computer Security Center, SRI International, Ft. Meade, MA, 1988.

[183] Sebyala, A. A., Olukemi, T., and Sacks, L. *Active Platform Security through Intrusion Detection using Naive Bayesian Networks for Anomaly Detection.* London Communications Symposium, 2002.

[184] Sherwood, T., Perelman, E., Hamerly, G., and Calder, B. *Automatically Characterizing Large Scale Program Behavior.* ACM APLOS X, pp. 45–57, San Jose, CA, October 2002.

[185] Siaterlis, C., and Maglaris, B. *Towards Multisensor Data Fusion for DoS Detection.* ACM Symposium on Applied Computing, March 2004.

[186] Sirer, E. G., Renesse, R. V., Rodríguez, R., and Mandujano, S. *Dynamic Analysis of Executable Code using Function-call Graphs for Intrusion Detection.* Tech. rep., Department of Computer Science, Cornell University, Ithaca, NY, December 2003.

[187] Snapp, S. R., Smaha, S. E., Teal, D. M., and Grance, T. *The DIDS (Distributed Intrusion Detection System) Prototype.* Summer 1992 USENIX Conference, pp. 227–233, Oakland, CA, June 1992.

[188] Sobirey, M. *The Intrusion Detection System AID*, University of South Carolina, March 2004. URL: http://www.cse.sc.edu/research/isl/mirrorSobireys.shtml (on-line resource).

[189] Sobirey, M., Fischer-Hübner, S., and Rannenberg, K. *Pseudonymous Audit for Privacy Enhanced Intrusion Detection.* 13th International Information Security Conference, pp. 151–163, Copenhagen, Denmark, May 1997.

[190] Somayaji, A., and Forrest, S. *Automated Response Using System-Call Delays.* 9th USENIX Security Symposium, ACM Press, pp. 185–198, Denver, CO, August 2000.

[191] SourceFire. *OpenSnort*, Sunfire, Inc., February 2004. URL: http://www.sourcefire.com (on-line resource).

[192] Spafford, E. H., and Zamboni, D. *Data Collection Mechanisms for Intrusion Detection Systems.* Tech. rep., Center for Education and Research in Information Assurance and Security, CERIAS, Purdue University, West Lafayette, IN, June 2000.

[193] Spitzner, L. *Honeypots: Tracking Hackers*, 1st ed. Addison-Wesley Professional. September 2002.

[194] Staniford-Chen, S., Cheung, S., Crawford, R., Dilger, M., Frank, J., Hoagland, J., Levitt, K., Wee, C., Yip, R., and Zerkle, D. *GrIDS – A Graph-based Intrusion Detection System for Large Networks.* 19th National InfIntrusion Detection Working Group – ormation Systems Security Conference, vol. 1, pp. 361–370, October 1996.

[195] Staniford-Chen, S., and Heberlein, L. *Holding Intruders Accountable on the Internet.* IEEE Symposium on Security and Privacy, pp. 39–49, Oakland, CA, 1995.

[196] Staniford-Chen, S., Tung, B., and Schnackenberg, D. *The Common Intrusion Detection Framework (CIDF).* Information Survivability Workshop, Orlando, FL, October 1998.

[197] Symantec. *Intruder Alert*, Symantec, Inc., March 2004. URL: http://securityresp.symantec.com/IA.html/ (on-line resource).

[198] Symantec. *ManHunt*, Symantec, Inc., April 2004. URL: http://securityresp.symantec.com/MH.html (on-line resource).

[199] Symantec. *NetProwler*, Symantec, Inc., March 2004. URL: http://securityresp.symantec.com/NP.html (on-line resource).

[200] Teng, H. S., Chen, K., and Lu, S. C. *Security Audit Trail Analysis Using Inductively Generated Predictive Rules.* 6th Conference on Artificial Intelligence Applications, IEEE Service Center, pp. 24–29, Piscataway, NJ, March 1990.

[201] Tinnagonsutibout, C., and Watanapongse, P. *A Novel Approach to Process-based Intrusion Detection System using Read-sequence Finite State Automata with Inbound Byte Profiler.* Information and Computer Engineering Postgraduate Workshop, 2003.

[202] Undercoffer, J. L., Joshi, A., Finin, T., and Pinkston, J. *A Target-Centric Ontology for Intrusion Detection.* 18th International Joint Conference on Artificial Intelligence, IJCAI, pp. 749–754, Acapulco, Mexico, July 2003.

[203] Undercoffer, J. L., Joshi, A., Finin, T., and Pinkston, J. *A Target Centric Ontology for Intrusion Detection: Using DAML+OIL to Classify Intrusive Behaviors. Knowledge Engineering Review, Cambridge University Press*, 2004.

[204] Uschold, M., and Grüninger, M. *Ontologies: Principles, Methods, and Applications.* Knowledge Engineering Review, pp. 93–155, vol. 11, num. 2, 1996.

[205] Vaccaro, H. S., and Liepins, G. E. *Detection of Anomalous Computer Session Activity.* IEEE Computer Society Symposium on Security and Privacy, pp. 280–289, May 1988.

[206] VanderAalst, W., Aldred, L., Dumas, M., and TerHofstede, A. *Design and Implementation of the YAWL System*. Tech. rep., Department of Technology Management–Eindhoven University of Technology and Centre for IT Innovation–Quensland University of Technology, Queensland, Australia, 2004.

[207] Viega, J., and Messier, M. *Secure Programming Cookbook for C and C++*, 2nd ed. O'Reilly & Associates. July 2003.

[208] Vigna, G., and Jonnson, E., Eds. *6th International Symposium, Recent Advances in Intrusion Detection 2003* (September 2003), Lecture Notes in Computer Science, LNCS, Springer.

[209] Vigna, G., and Kemmerer, R. *NetSTAT: A Network-based Intrusion Detection System*. Journal of Computer Security, pp. 37–71, vol. 7, num. 1, 1999.

[210] Vytopil, J. *Formal Techniques in Real-time and Fault-tolerant Systems*. Springer-Verlag. Berlin, Germany, December 1991.

[211] Wagner, D., and Dean, D. *Intrusion Detection via Static Analysis*. IEEE Symposium on Security and Privacy, pp. 156–169, Washington, DC, 2001.

[212] Wang, L.-X. *A Course in Fuzzy Systems and Control*. Prentice Hall, Inc. Upper Saddle River, NJ, 1997.

[213] Wang, X., Reeves, D. S., and Wu, S. F. *Tracing-based Intrusion Response*. Journal of Information Warfare, vol. 1, num. 1, 2001.

[214] Weiss, G. *Multiagent Systems: A Modern Introduction to Distributed Artificial Intelligence*, 1st ed. MIT Press. Cambridge, MA, 1999.

[215] White, G. B., Fisch, E. A., and Pooch, U. W. *Cooperating Security Managers: A Peer-Based Intrusion Detection System*. IEEE Network, pp. 20–23, vol. 10, num. 1, 1996.

[216] Wimer, S. M. *Cylantsecure: A Scientific Approach to Security, whitepaper*, Cylant Corp., February 2004. URL: http://www.cylant.com/products/cylantsecure.htm (on-line resource).

[217] Wylie, J. J., Bigrigg, M. W., Strunk, J. D., Ganger, G. R., Kiliççöte, H., and Khosla, P. K. *Survivable Information Storage Systems*. Computer, pp. 60–68, vol. 33, num. 8, 2000.

[218] Yang, J., Ning, P., Wang, X. S., and Jajodia, S. *CARDS: A Distributed System for Detecting Coordinated Attacks*. IFIP TC11 Sixteenth Annual Working Conference on Information Security (SEC 2000), pp. 171–180, August 2000.

[219] Yao, Y., Frasconi, P., and Pontil, M. *Fingerprint Classification with Combinations of Support Vector Machines*. Lecture Notes in Computer Science, LNCS, pp. 253–258, vol. 2091, num. 1, 2001.

[220] Ye, N., Li, X., Chen, Q., Emran, S. M., and Xu, M. *Probabilistic Techniques for Intrusion Detection Based on Computer Audit Data.* IEEE Transactions on Systems, Mand, and Cybernetics, vol. 31(4), pp. 266–274, July 2001.

[221] Yee, B. S. *A Sanctuary for Mobile Agents.* Secure Internet Programming: Security Issues for Distributed and Mobile Objects, Springer Verlag, Lecture Notes in Computer Science, LNCS, pp. 261–273, Berlin, Germany, 1999.

[222] Yoda, K., and Etoh, H. *Finding a Connection Chain for Tracing Intruders.* 6th European Symposium on Research in Computer Security, ESORICS, pp. 191–205, October 2000.

[223] Zamboni, D. *Doing Intrusion Detection using Embedded Sensors.* PhD thesis, Purdue University. West Lafayette, IN, 2000.

[224] Zamboni, D. M., and Spafford, E. H. *Intrusion Detection using Autonomous Agents.* Computer Networks – Elsevier, pp. 547–570, vol. 34, num. 04, 2000.

[225] Zou, C., Gong, W., and Towsley, D. *Code Red Worm Propagation Modeling and Analysis.* 9th ACM Conference on Computer and Communications Security, pp. 138–147, Washington, DC, 2002.

ATTACHMENT H

WorldCat: A multiagent approach to outbound intrusion detection

OCLC	970596494	**No holdings in LDL - 1 other holding**									
Books	**Rec Stat** c			**Entered** 20041209		**Replaced** 20220831093048.3					
Type	t	**ELvl**	M	**Srce**	d	**Audn**		**Ctrl**		**Lang**	eng
BLvl	m	**Form**		**Conf**	0	**Biog**		**MRec**		**Ctry**	xx
		Cont	bm	**GPub**		**LitF**	0	**Indx**	0		
Desc	a	**Ills**	a	**Fest**	0	**DtSt**	s	**Dates**	2004,		

040 I8M ‡b spa ‡c I8M ‡d OCLCO ‡d OCLCF ‡d OCLCQ ‡d OCLCO ‡d OCLCQ
050 _4 TK5105.59 ‡b .M35 2004
090
049 LDLL
100 1_ Mandujano Vergara, Salvador.
245 12 A multiagent approach to outbound intrusion detection / ‡c Salvador Mandujano Vergara.
260 Monterrey, N.L., ‡c 2004.
300 xxiii, 179 p. : ‡b il. ; ‡c 28 cm
336 texto ‡b txt ‡2 rdacontent/spa
337 sin mediación ‡b n ‡2 rdamedia/spa
520 This is a dissertation on the topic of intrusión detection. It supports the philosophy of system vigilance by exploring the concept of outbound intrusión detection, which is concerned with the identification and collectiou of evidence that helps prove local resources are being used to compromise external systems. We discuss the motivation behind the approach, explain the need for splitting the scope of intrusión detection into sub-problems, and present trends in computer security that revcal basic design considcrations that need to be taken into account when developing modern informatioii security tools. We propose a multiagent architecture for outbound intrusión detection supported by an ontology. Groups of agents collectively monitor outbound network trame and local activity in order to identify references to neighboring systems that may be indicative of a compromise attempt. We organize agents into sub-environments called agent cells that are connected to cach other in a non-hierarchical fashion. Different classes of agents and cells compose the system, which performs attack modeling by employing múltiple concurrent agents. Detection cells impleinent independent misuse intrusión strategies whose output is systematically fed to correlation cells capable of more accurate diagnosis. We present an attack-source-centric ontology that extends previous work in the área. It enables mcssage interpretation and enhanced agent communication within the architecture sirn-plifying at the same time system maintenance and facilitating the integration of new conipouents. We describe the implementation of the proposed architecture through the FROID prototype as a proof of concept. This is a misuse-based intrusión detection system built with agent and semantic web open-source technology whose particular focus is the identification of automated remote attack tools. It performs signature generation, matching, and correlation. and supports a siguature deploymeut mechanism over the Internet. Wc introduce a similarity matching method that improves the performance of existing algorithms by leveraging entropy and frequency properties of the input hereby reducing search time. We link detection with incident response by procuring low false alarm rates that allow us to study local and external reaction methods to outbound intrusión events. We also present a component of the architecture that performs tracing of interactivo sessions as a way of identifying the root location of a security event. We describe the experimental design and report the results obtained with the prototype that show the feasibility of the approach as an altérnate way of containing the impact of security incidents through the integration of a mesh of monitoring agents.
502 Tesis (Doctor of Philosophy in Artificial Intelligence)--ITESM Campus Monterrey.
504 Bibliografía: p. [181-197].
500 Arturo Galván Rodríguez, Asesor.
650 _4 Seguridad de computadoras.
655 _4 Tesis académicas.
655 _7 Academic theses. ‡2 fast ‡0 (OCoLC)fst01726453

Sophos Ex. 1009

655 _7 Academic theses. ‡2 lcgft
655 _7 Thèses et écrits académiques. ‡2 rvmgf ‡0 (CaQQLa)RVMGF-000001173
700 1_ Galván Rodríguez, Arturo, ‡c PhD.

Action Status	Delete Holdings _	Export _	Label _	Produce _	Update Holdings _	Validate _

Sophos Ex. 1009

ATTACHMENT I

Instituto T de Monterrey

ce

ion

PhD thesis

by

Salvador Mandujano Vergara

Document properties

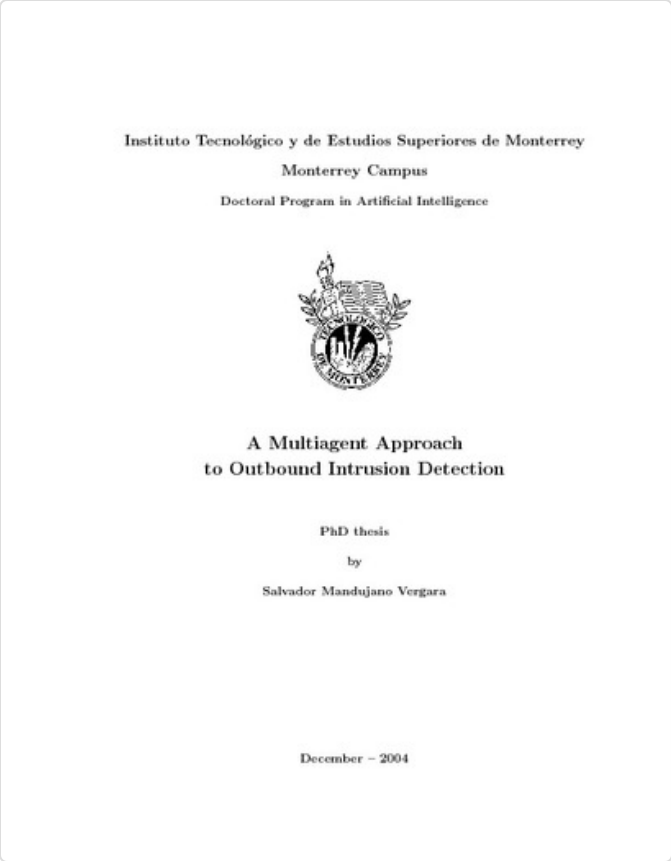
File name:	dissertationMexico.pdf
File size:	4.8 MB
Title:	-
Author:	-
Subject:	-
Keywords:	-
Created:	12/10/04, 10:50:38 AM
Modified:	12/10/04, 10:51:00 AM
Application:	TeX
PDF producer:	pdfTeX-0.14h
PDF version:	1.4
Page count:	224
Page size:	Varies
Fast web view:	No

Close

Sophos Ex. 1009
Page 311 of 318

ATTACHMENT J

A Multiagent Approach to Outbound Intrusion Detection



Instituto Tecnológico y de Estudios Superiores de Monterrey

Monterrey Campus

Doctoral Program in Artificial Intelligence

A Multiagent Approach
to Outbound Intrusion Detection

PhD thesis

by

Salvador Mandujano Vergara

December – 2004

View/Open

DocsTec_1973.pdf Size (4.817Mb)

DocsTec_1973_1.pdf Size (35.74Kb)

Date

2004-01-12

Author

Mandujano Vergara, Salvador

Metadata

Show full item record

Export citation

Abstract

A Multiagent Approach to Outbound Intrusion Detection. Ph.D. dissertation by Salvador Mandujano Vergara, Instituto Tecnológico y de Estudios Superiores de Monterrey. Advisor: Prof. Arturo Galván. December 2004. This is a dissertation on the topic of intru

URI

http://hdl.handle.net/11285/572223

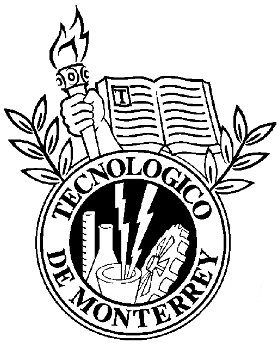
Collections

Ciencias Exactas y Ciencias de la Salud **294**

Instituto Tecnológico y de Estudios Superiores de Monter

Monterrey Campus

Doctoral Program in Artificial Intelligence



A Multiagent Approach to Outbound Intrusion Detection

Instituto Tecnológico y de Estudios Superiores de Monterrey
Campus Monterrey

Monterrey, Nuevo León a 09 de Diciembre de 2004.

Lic. Arturo Azuara Flores:
Director de Asesoría Legal del Sistema

Por medio de la presente hago constar que soy autor y titular de la obra titulada "A Multiagent Approach to Outbound Intrusion Detection", en los sucesivo LA OBRA, en virtud de lo cual autorizo a el Instituto Tecnológico y de Estudios Superiores de Monterrey (EL INSTITUTO) para que efectúe la divulgación, publicación, comunicación pública, distribución y reproducción, así como la digitalización de la misma, con fines académicos o propios al objeto de EL INSTITUTO.

El Instituto se compromete a respetar en todo momento mi autoría y a otorgarme el crédito correspondiente en todas las actividades mencionadas

Except where otherwise noted, this item's license is described as http://creativecommons.org/licenses/by-nc-nd/4.0

DSpace software copyright © 2002-2019 DuraSpace

ATTACHMENT K

Instituto T de Monterrey

Document properties

File name:	dissertationMexico.pdf
File size:	4.8 MB
Title:	-
Author:	-
Subject:	-
Keywords:	-
Created:	12/10/04, 10:50:38 AM
Modified:	12/10/04, 10:51:00 AM
Application:	TeX
PDF producer:	pdfTeX-0.14h
PDF version:	1.4
Page count:	224
Page size:	Varies
Fast web view:	No

Close

PhD thesis

by

Salvador Mandujano Vergara

Sophos Ex. 1009
Page 318 of 318