Sachar Paulus
Norbert Pohlmann
Helmut Reimer

Securing Electronic Business Processes

Highlights of the Information Security Solutions Europe 2006 Conference

Contents

Preface _____ **xi**

About this Book _____ **xiii**

ISCOM: On the Way for ICT Security in Italy _____ **xv**

RFID, e-ID Cards, Trusted Computing, Interoperability _____ **1**

Radio Frequency Identification (RFID) and Data Protection Legal Issues
Zoi Talido _____ 3

e-ID and Smartcards – Current Status, Hopeful Developments and Best Practices
Graham Williamson _____ 17

European Citizen Card Combined with Travel Document Function, Convergence or Divergence?
Detlef Houdeau _____ 25

Physical Unclonable Functions for enhanced security of tokens and tags
Pim Tuyls, Boris Škorić _____ 30

Hardware Security Features for Secure Embedded Devices
Helena Handschuh, Elena Trichina _____ 38

Security in Next Generation Consumer Electronic Devices
Tom Kan, Tim Kerins, Klaus Kursawe _____ 45

Security Architecture for Device Encryption and VPN
Ammar Alkassar, Michael Scheibel, Christian Stüble, Ahmad-Reza Sadeghi, Marcel Winandy _____ 54

TPM Enterprise Key Management requires centralized Hardware-based Security
Bernhard Weiss _____ 64

Implementation of DRM Systems under the EU Legal Framework
Pius Alexander Benczek 72

IT-Grundschutz: Two-Tier Risk Assessment for a Higher Efficiency in IT Security Management
Angelika Jaschob, Lydia Tsintsifa 95

ISO/IEC 24727 – A Future Standard for Smart Card Middleware
Stephan Spitz, Jens Urmann, Gisela Meister 102

Information Security Standardization – the ETSI Perspective
Charles Brookson, Dionisio Zumerle 108

Digital Signatures without the Headaches
Nick Pope, Juan Carlos Cruellas 119

Could Test Standards Help on the Way to Achieve Global e-Passport Interoperability?
Andreas M. Wolf 129

A New Standard Based Road to Interoperable Strong Authentication
Philip Hoyer 139

Identity Management, Biometrics, PKI-Solutions, Network Security 149

Identifying Patterns of Federation Adoption
Heather Hinton, Mark Vandenwauver 151

Fidelity: Federated Identity Management Security based on Liberty Alliance on European Ambit
Manel Medina, Miquel Colomer, Sandra García Polo, Antoine de Poorter 161

Deflecting Active Directory Attacks
Jan De Clercq 168

Implementing role based access control – How we can do it better!
Marko Vogel 176

Identity and Access Control – Demonstrating Compliance
Marc Sel, Bart Van Rompay 186

Robust and Secure Biometrics: Some Application Examples
T. Kevenaar, G.J. Schrijen, A. Akkermans, M. Damstra, P. Tuyls, M. van der Veen 196

Selecting the Optimal Biometric 2-factor Authentication Method – a User's Viewpoint
Gunter Bitz 204

A Face Recognition System for Mobile Phones
Paolo Abeni, Madalina Baltatu, Rosalia D'Alessandro 211

Advanced certificate validation service for secure Service-Oriented Architectures
Antonio Ruiz-Martínez, Daniel Sánchez-Martínez, C. Inmaculada Marín-López, Antonio F. Gómez-Skarmeta 218

An Introduction to Validation for Federated PKIs
Robert Dulude, David Engberg, Seth Hitchings 228

MADSig: Enhancing Digital Signature to Capture Secure Document Processing Requirements
Jean-Christophe Pazzaglia, Stefano Crosta 241

PKI Consolidation Project and Multiapplicative Smart Payment Cards
Milan Marković, Miloš Kilibarda, Aleksandar Milošević 249

Security Analysis and Configuration of Large Networks
Antonio Lioy 259

S-VPN Policy: Access List Conflict Automatic Analysis and Resolution
Simone Ferraresi, Stefano Pesic, Livia Trazza, Andrea Baiocchi 266

Lock-Keeper: A New Implementation of Physical Separation Technology
Feng Cheng, Christoph Meinel 275

SPEECH: Secure Personal End-to-End Communication with Handheld
A. Castiglione, G. Cattaneo, A. De Santis, F. Petagna, U. Ferraro Petrillo 287

Finding the Mobile Trusted Element
Fabio Ricciato, Maura Turolla, Antonio Varriale 298

Security Management, Applications 309

Centrally Administered COIs Using Cross-Organizational Trust
Kevin Foltz, Coimbatore Chandersekaran 311

Improving Assurance of Information Security RoI
Michael D. Barwise 318

Modelling the Economics of Free and Open Source Software Security
Anas Tawileh, Jeremy Hilton, Steve McIntosh 326

Securing service-oriented applications
Anthony Nadalin, Nataraj Nagaratnam, Maryann Hondo 336

A Service Oriented Trust Development Platform
Helena Rifà, Francisco Jordán 344

A Trust Label for Secure and Compliant e-ID Applications: The Belgian Experience
Geert Somers, Jos Dumortier 356

Electronic signature in Italy after ten years of “running in”
Giovanni Manca 363

Awareness Raising, Compliance, Data Protection, Cyberspace Regulation 375

Internet Early Warning System: The Global View
Norbert Pohlmann, Marcus Proest 377

IT Security Vulnerability and Incident Response Management
Wim Hafkamp 387

Blending Corporate Governance with Information Security
Yves Le Roux 396

On Privacy-aware Information Lifecycle Management in Enterprises: Setting the Context
Marco Casassa Mont 405

Regulation of State Surveillance of the Internet
Murdoch Watney 415

How Can NRA Contribute to the Improvement of IT Security?
Rytis Rainys 426

Information Security Regulation: Tomorrow Never Dies?
Andreas Mitrakas 433

Introducing Regulatory Compliance Requirements Engineering
Shahbaz Ali, Jon Hall 439

Legal Issues in Secure Grid Computing Environments
Irene Kafeza, Eleanna Kafeza, Felix Wai-Hon Chan 448

The Impact of Monitoring Technology on the Law
Pieter Kleve, Richard De Mulder, Kees van Noortwijk 455

Index 467

Preface

ENISA is proud to be working with eema, TeleTrusT, ISCOM (the Italian Institute for Communications and Information Technologies) and the German Federal Ministry of the Interior as well as the German Federal Office for Information Security for this year's 8th annual Information Security Solutions Europe Conference.

The aim of ISSE has always been to support the development of a European information security culture. ENISA is committed to this goal, in our work to assist and advise the European Commission, Member States as well as business community on network, information security and legislative requirements and we are delighted to support ISSE again this year.

The security of communication networks and information systems is of increasing concern. In order to face today's complex information security challenges it is clear that working collaboratively with one another is the key to generating new strategies to address these problems. It has been an exciting opportunity to facilitate this collaboration at ISSE 2006, and pull together the wealth of industry knowledge, information and research that we hold in Europe, and across the globe.

The success of this event in generating ideas and frank, lively debate around the complex topic of IT security is due also to the independent, varied nature of the programme, which was selected by world-wide industry specialists.

Some of the key topics explored at this year's conference have been chosen as the basis for this book, which is an invaluable reference point for anyone involved in the IT security industry.

We hope that you will find it a thought-provoking and informative read.

Andrea Pirotti, Executive Director, ENISA

About this Book

The Information Security Solutions Europe Conference (ISSE) was started in 1999 by eema and TeleTrusT with the support of the European Commission and the German Federal Ministry of Technology and Economics. Today the annual conference is a fixed event in every IT security professional's calendar.

The integration of security in IT applications was initially driven only by the actual security issues considered important by experts in the field; currently, however, the economic aspects of the corresponding solutions are the most important factor in deciding their success. ISSE offers a suitable podium for the discussion of the relationship between these considerations and for the presentation of the practical implementation of concepts with their technical, organisational and economic parameters.

From the beginning ISSE has been carefully prepared. The organisers succeeded in giving the conference a profile that combines a scientifically sophisticated and interdisciplinary discussion of IT security solutions while presenting pragmatic approaches for overcoming current IT security problems.

An enduring documentation of the presentations given at the conference which is available to every interested person thus became important. This year sees the publication of the third ISSE book – another mark of the event's success – and with about 50 carefully edited papers it bears witness to the quality of the conference.

An international programme committee is responsible for the selection of the conference contributions and the composition of the programme:

- **Ronny Bjones,** Microsoft (Belgium)
- **Alfred Büllesbach,** Daimler Chrysler (Germany)
- **Lucas Cardholm,** Ernst&Young (Sweden)
- **Roger Dean,** eema (UK)
- **Marijke De Soete,** Security4Biz (Belgium)
- **Jos Dumortier,** KU Leuven (Belgium)
- **Walter Fumy,** Siemens (Germany)
- **Boaz Gelbord,** ENISA (Greece)
- **David Goodman,** eema (UK)
- **Michael Hange,** Federal Office for Information Security (Germany)
- **John Hermans,** KPMG (Netherlands)
- **Jeremy Hilton,** Cardiff University (UK)
- **Alison James,** eema (UK)
- **Frank Jorissen,** SafeBoot (Belgium)
- **Matt Landrock,** Cryptomathic (Denmark)
- **Tim Mertens,** ENISA (Greece)
- **Andreas Mitrakas,** ENISA (Greece)
- **David Naccache,** ENS (France)
- **Sachar Paulus,** SAP (Germany)

- **Daniele Perucchini,** Fondazione Ugo Bordoni (Italy)
- **Attila Péterfalvi,** Parliamentary Commissioner for Data Protection and Freedom of Information (Hungary)
- **Norbert Pohlmann,** University of Applied Sciences Gelsenkirchen (Germany)
- **Bart Preneel,** KU Leuven (Belgium)
- **Helmut Reimer,** TeleTrusT (Germany)
- **Paolo Rossini,** Telsy Italia (Italy)
- **Wolfgang Schneider,** Fraunhofer SIT (Germany)
- **Robert Temple,** BT (UK)

The editors have endeavoured to allocate the contributions in these proceedings – which differ from the structure of the conference programme – to topic areas which cover the interests of the readers.

Sachar Paulus *Norbert Pohlmann* *Helmut Reimer*

eema (www.eema.org):

Established in 1987, eema is an independent association of IT professionals, businesses and governments providing business and technical networking opportunities at both local and regional levels in the broad areas associated with digital identity and its applications, such as security. Our mission is to stimulate the growth and effectiveness of our members' business in these areas through increased market awareness, cooperation and opportunity creation.

We aim to bring over 1,500 member representatives together in a neutral environment for education and networking puposes. We enable members to share experiences and best practice by holding meetings and conferences, by facilitating working groups who produce reports on topical subjects, and by helping members to connect with the right person to help them solve business issues or develop beneficial business relationships. All work produced by members is available free to other members, and previous papers include: Towards Understanding Identity, Role Based Access Control – a Users Guide, Secure e-mail within a Corporate Environment and Secure e-mail between Organisations.

For more information contact:
alison.james@eema.org.

TeleTrusT (www.teletrust.de):

In the 16 years of its existence TeleTrusT has evolved into a competence network for applied Cryptography and Biometrics with over 90 institutional members.

The TeleTrusT working groups produce results which create an advantageous framework for trustworthy solutions of daily business processes as well as contributing to their acceptance.

TeleTrusT brings together the interests of users and vendors. Thus vendors can satisfy the users' demands more effectively with marketable products and services, in which scalable security mechanisms are implemented.

TeleTrusT seeks and cultivates the cooperation with other organisations with similar objectives – in Germany and internationally. Thus ISSE has been organised in cooperation with EEMA, ENISA and ISCOM in Rome this year.

For further information contact:
sophie.hellmann@teletrust.de

ISCOM:
On the Way for ICT Security in Italy

The Istituto Superiore delle Comunicazioni e delle Tecnologie dell'Informazione (ISCOM) was established in 1907 as a technical-scientific department belonging to the Italian Communication Ministry. Considering its role as a nonpartisan public institution, the Institute's value added in terms of reliability and expertise is the aspect which characterizes the technical support and consultancy services it provides to businesses and entities in the TLC sector. The role of ISCOM in providing services to ICT Companies, government agencies and users is manifold, spanning from experimental and research activities to specialized training and education in the TLC field.

One of ISCOM's main missions is its proactive role in national and international law-making activities, in order to ensure greater transparency and better access to services for users, manufacturers and TLC network administrators alike.

As far as research is concerned, ISCOM is essentially focused on developing and improving TLC and IT related services. Hence, activities involve almost all areas in these fields, from telephony to television, to signal processing and treatment, from network architecture to service implementation.

ISCOM runs the Post-Graduate Specialization School in TLC (which began its activity in 1923), which provides higher education in electronic communication and information technologies; it also provides technical training and updating courses on electronic communications and information technologies, security, multimedia applications, and Quality of Service to both Ministry and government staff in general, to enhance their technical know-how and skills.

ISCOM works with several Certification Bodies to verify and control Corporate Quality System compliance with UNI EN ISO 9000 standards, is involved in monitoring Accredited Laboratory compliance with UNI CEI EN ISO/IEC 17025 rules and is a Notified Body for activities envisaged by Legislative Decree n. 269 of May 9, 2001. It is also a Notified Body under the EU Directive on radio equipment and telecommunications terminal equipment as well as a Competent Body and Notified Body on electromagnetic compatibility. In 2002, the Institute became the International Certification Body for the TETRA MoU.

Among all the numerous ISCOM fields of activity, ICT security is getting an increasing relevance. Here, ISCOM plays a leading role in various contexts, some of which are briefly summarized below:

- Due to his widely recognized non-partisan role, a government decree dated october 30, 2003 appointed ISCOM the Certification Body within the Italian certification scheme for commercial security systems and products. The Certification Body supervises all the

activities carried out within the certification scheme, which operates according to the international evaluation criteria ITSEC and Common Criteria.

- ISCOM is an Evaluation Center (Ce.Va.) for ICT systems and products dealing with classified data. The center, the only one belonging to the Italian Public Administration which has been accredited by the Autorità Nazionale per la Sicurezza (ANS), carries out evaluation activities according to ITSEC and Common Criteria.
- ISCOM runs the Training Center on ICT Security for Public Administration personnel. The Training Center provides training and raises awareness amongst government employees on ICT security, through the development of a centralized and coordinated Training and Awareness-Raising Plan aimed at disseminating security principles and methodologies throughout the Administration.
- The Institute acts as promoter and leader of several initiatives aimed at raising the national level of ICT security, by gathering the expertise of the major subjects operating in the ICT field. Among these initiatives we can recall the redaction of three guidelines, in English and Italian, on "*The quality of service in ICT networks*", "*Risk analysis and protection strategies for network security*" and "*Network security in critical infrastructures*", carried out with the contribution of experts from institutions and industry. Six more guidelines are being released; these will be focused on deepenings on risk analysis, on the outsourcing of security services, on QoS in UMTS, on QoS in broadband networks, on local emergency handling and on security certification. Moreover, ISCOM has promoted the creation of ISAC on network security, currently involving all the major Italian network operating companies.

ISCOM hosting of ISSE 2006 is a further prove of our desire to play a role in fostering the European information security debate. We look forward to a great opportunity for the exchange of ideas and experiences.

Luisa Franchina,

PhD, General Director of Istituto Superiore delle Comunicazioni
E delle Tecnologie dell'Informazione

RFID
e-ID Cards
Trusted Computing
Interoperability

Radio Frequency Identification (RFID) and Data Protection Legal Issues*

Zoi Talidou

Hellenic Data Protection Authority
Legal Auditor
Kifisias 1-3, Athens
ztalidou@dpa.gr

Abstract

Radio Frequency Identification (RFID) Technology uses radio waves to identify automatically, wirelessly, contact less and without visibility objects which, or people who have an RFID tag attached. It is being used in many sectors but raises data-protection concerns. The reasons for that are the world-wide unique identifier, the possibility of unnoticed remote reading, and the profiling through sporadic surveillance. For these reasons RFID-technology introduces new legal issues that have to be discussed: what is personal data, who is responsible for the data processing, whether the data-transmission is telecommunication, whether it presents a new way of direct marketing or if it constitutes an automatic decision.

In the early 1970s fears about loss of privacy and worries concerning data protection were focused on large, centrally held data-bases containing files about named or numbered individuals processed by huge computers situated in big rooms. As the Web, its attendant search engines and the inter-link ability of many databases in various networks have developed, the concept of "files" became trivial. Now the emerging RFID technology contributes to the realisation of the Ambient Intelligence Environment, where intelligent objects communicate with each other by exchanging information and taking decisions. That introduces us to the next step of the "Internet of the things". Technology innovation and the impact of its usage stress a rethinking and re-examining of the traditional legal principles and legal instruments in the field of data protection.

1 What RFIDs are all about

Radio Frequency Identification (RFID) Technology belongs to the broad category of automatic identification technologies[1] and uses radio waves to automatically identify wirelessly, contact less and without visibility[2] objects which, or people who have an RFID tag attached. It consists of two parts: **a tag** that contains an identification number and **a reader** who works as a scanner. This number usually acts as an input to further data processing[3]. A typical RFID tag consists of a small integrated circuit attached to a radio antenna, capable of transmitting a unique serial number. The tag can easily be embedded onto or into (textile-) products, onto their packages or even direct implanted beneath human's skin. RFID tags can be active, semi-

* This paper is based on a report conducted for LEGAL-IST

[1] What is RFID?, RFID Journal, available at: http://www.rfidjournal.com/article/articleprint/1339/-1/129/

[2] See http://en.wikipedia.org/wiki/Rfid.

[3] *Hennig, Ladkin, Sieker,* Privacy Enhancing Technology Concepts for RFID Technology Scrutinised, p.1.

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2006), 3-16

active or passive. *Passive Tags* do not have a power source; they simply reflect back energy coming from the reader antenna[4]. *Active RFID* tags on the other hand, have their own internal power source that allows them having longer range and larger memories than passive tags, as well as the ability to store additional information sent by the transceiver. A typical **reader** is a device that has one or more antennas that emit radio waves and receive signals back from the tag. This RFID reader is a data-collection instrument, and a transmitter or broadcaster of information, as it sends its data through the information network. The databases connected to these networks hold, use and disclose the gathered information.

The innovation of RFID tags is that they provide for unique identification of each tagged unit whereas bar codes are identical for every unit of the same product[5]. Prices of RFID are dropping. Many postulate that they will be the essential drivers of ubiquitous computing and will introduce the so-called "Internet of the things".

2 Use of RFID technology

2.1 Retail/Consumer Goods Sector

Companies across the retail and consumer packaged goods supply chains have been among the early adopters of RFID and Electronic Product Code (EPC) technologies. The use of this new technology is connected to the EPC Discovery Service, an aggregate database of tag "sightings" collected from independent readers. Anyone with access EPC Discovery can monitor or track the movement of a particular RFID-tagged item. The retail industry is using passive tags that implement no protection against unauthorised access to the information held. Hence the EPC can be read out directly by any RFID-reader from a six to eight meters distance[6].

2.2 Manufacturing Sector

RFID technology can increase productivity and reduce costs by enabling to track inventory, reusable containers, work in process and finished products: they can manage parts inventory with active RFID, improve the tracking of work in process, reduce parts defects, and increase factory productivity by using active RFID tags. In some cases, RFIDs aim in such seemingly simple tasks as ensuring that the right label goes on a product or that a box contains everything it should. In other cases, RFID is put through more complex uses as tracking an item through every workstation and recording every tool that performed an operation on it. This information can be used to quickly identify potential problems and correct them before they show up in the product. RFID can furthermore save companies a great amount of money spent on replacing lost tools, that can be easily traced through the tags.

---

[4] The basic of RFID Technology, RFID Journal, available at: http://www.rfidjournal.com/article/articleprint/1337/-1/129/

[5] See *International Conference of Data Protection & Privacy Commissioners,* Resolution on Radio-Frequency Identification, (Nov. 20, 2003) p. 2, available at: http://www.privacyconference2003org/resolutions/res5.DOC.

[6] Auto-ID Centre (2003): Technical report 860MHz-930MHz Class I Radio Frequency Identification Tag Radio Frequency & Logical Communication Interface Specification Candidate Recommendation, Version 1.0.1., MIT, USA, available at: http://interval.hu-berlin.de/downloads/rfid/chipklassen/4_candidate_recommendation_1_0_1.pdf

2.3 Recycling & waste management

The EPC tags may be used to automatically sort recyclable material and will also identify manufacturer, type and weight of disposable material (the manufacturer of a product that will eventually constitute hazardous waste may ultimate have to pay for its safe disposal).

2.4 Transportation/Logistics Sector

Transportation and logistic companies are already tagging product for their customers. Some of them are still examining how they can benefit internally, by improving the utilization of containers and chassis with RFID tracking. Logistics hubs can benefit from a real-time locating system, and they can improve the visibility of cargo in transit and cargo security with electronic seals.

2.5 Libraries

Libraries began using RFID systems to replace their electro-magnetic and bar code systems in the late 1990s. RFID technology in libraries promises to relieve repetitive strain injury, speed patron self-checkout, make possible comprehensive inventory and automated sorting, retrieve hidden items and support security. Many libraries (more than 130 in North America and the Stadtbibliothek of the city Wien[7]) are starting to tag every item in their collections with RFID tags. But current library RFID tags do not prevent unauthorised reading of tag data[8].

2.6 Tracking of animals (dogs, cows and sheep)

Pets can be implanted with small chips so that they may be returned to their owners if lost. They can also be used to satisfy the need to track herds and to be able to recognize when an animal is missing and, if the animal has died, locate its body[9]. Beside that, request on safe handling with animals as a result of repeated outbreaks of epidemics is pointing out electronic animal tracking through RFID as a significant solution. Following successful animal tracking trials[10], the European Council of Ministers (ECM) has adopted a law[11] throughout Europe requiring the individual electronic tagging of sheep and goats using RFID technology. Besides RFID tags are used for to identify big pets, such as dogs over 20 kilograms. Several laws at the European level make the wear of such a tag compulsory, that will have to contain at least following data: unique number for the chip, data of the pet and data of the owner of the pet.

---

[7] http://www.ekz.de/2110.html

[8] See *Molnar, Wagner,* Privacy and Security in Library RFID issues, practices and architectures, CCS'04, October 25-29 2004, Washington, DC, USA, p. 218, available at: http://www.cs.berkeley.edu/molnar/library.pdf

[9] See http://www.rfidgazette.org/asset_tracking/.

[10] See *Balch, Feldman, Wilson,* Assessment of a RFID System for Animal Tracking, The BORG Lab, Georgia Institute of Technology, Atlanta, 1.10.2004, available at: http://www.cc.gatech.edu/~storm/Feldman2004TR.pdf

[11] Council Regulation (EC) No 644/2005 of 27 April 2005 authorising a special identification system for bovine animals kept for cultural and historical purposes on approved premises as provided for in Regulation (EC) No 1760/2000 of the European Parliament and of the Council, available at: http://europa.eu.int/eur-lex/lex/LexUriServ/site/de/oj/2005/l_107/l_10720050428de00180019.pdf

2.7 Health Care Sector

Hospitals plan to deploy RFID to identify patients, call up records, reduce medical errors and improve overall productivity. A pilot project has started in July 2005 in clinical centre of Saarbrücken, where thousand of patients receive by admission a bracelet with an RFID tag on which the patient identifier is stored. Physicians and nurses may access the patient identifier and data stored on a database through a wireless network. The project is based on a solution already deployed in Jacobi Medical Centre, New York[12].

2.8 Tracking of people (schools, prisons, VIP clubs)

A group of children in Yokohama City in Japan wears active tags to keep them safe on their way to and from school[13]. Each child participating to the programme wears a bracelet with a RFID tag. Existing Wi-Fi access points used by the city for wireless Internet access work as RFID readers that receive signals send by the tags. The system can also be set up to notify parents or guardians automatically via e-mail on a cell phone or PC if a child passes a specific Wi-Fi access point on the way to or from the school. The VIP Baja Beach Club in Barcelona offers it's VIP clients the opportunity to have a syringe-injected RFID microchip implanted in their upper arms: this chip gives them special access to VIP lounges, but also acts as a debit account, from which they can pay for drinks[14]. A new tracking system has been developed which provides real-time identification and tracking of inmates and officers[15]. It handles common prison complexities such as a multi-floor, mixed indoor/outdoor environment, as well as the need for cell-level accuracy. The tag immediately detects any attempt to remove or tamper with it. The Los Angeles County jail system has reportedly engaged in a pilot project to use RFID technology to track inmates at the Pitchess Detention Centre in Castaic[16].

2.9 Passports and Ids

In May 2004 the International Civil Aviation Organisation (ICAO) adopted specifications for machine readable travel documents (MRTD) which demands for digital storage of the pass photo[17]. In compliance with the recommendations of the ICAO the Council of the European Union adopted on 13/12/2004 a regulation[18] mandating the inclusion of both facial image and fingerprints in future passports and travel documents issued by EU Member States. The new regulation aims at better protecting EU passports against falsification, at enabling better identification of passport holders and at harmonising security standard features used in the production of passports and travel documents issued by Member States[19]. As a result in November 2005 Germany introduced the first e-passport[20], equipped with biometric data stored on a

---

[12] Computer mit Augen und Ohren, at: Frankfurter Allgemeine Zeitung, 14.01.2006, p. 18.

[13] http://www.rfidjournal.com/article/articleprint/2050/-1/1/

[14] See http://news.bbc.co.uk/2/hi/technology/3697940.stm; http://www.heise.de/newsticker/meldung/53789

[15] See http://www.technologynewsdaily.com/node/1900.

[16] See http://www.socaltech.com/fullstory/0001952.html.

[17] Available at: http://www.icao.int/cgi/goto_m.pl?/icao/en/strategic_objectives.htm.

[18] Council Regulation 2252/2004 on standards for security features and biometrics in passports and travel documents issued by Member States, OJ L 385, available at: http://europa.eu.int/eur-lex/lex/LexUriServ/site/en/oj/2004/l_385/l_38520041229en00010006.pdf

[19] See e-government of the European Union news available at: http://europa.eu.int/idabc/en/document/3669/330.

[20] http://www.epass.de/

RFID tag. In Italy the Foreign Affairs Ministry issued on 17th January 2006 a decree concerning the introduction of a new electronic passport that will include biometric data contained in RFID chips[21]. The European Central Bank was moving forward with plans to embed RFID tags as thin as a human hair into the fibres of Euro bank notes by 2005[22]. Hitachi Ltd. has developed a RFID chip that requires no external antenna and makes possible the embedding of tracking and identification chips in bank notes, tickets and other paper products[23]. But now it is still uncertain whether they will force this plan or not because, according to new statements, RFID technology is not safe enough to combat monetary counterfeit[24].

2.10 Transportation: e-pass, e-plate, e-ticket

Many countries, including Greece, have developed RFID-based Electronic Toll Collection systems for a variety of highways and bridges. As a vehicle equipped with a RFID transponder enters a toll plaza equipped to accept RFID toll collection the radio frequency emitted by the electronic reader will activate the transponder. The transponder then sends out account or identification information pertaining to the vehicle. The information is received by the reader and through the antenna sent to the host computer system. The toll is then deducted from the account associated to that vehicle and the driver is signaled to proceed. The tags can be read at a speed of 100 miles per hour. Their use is simplifying the toll-collection procedure and so cutting traffic jams and the resulting levels of smog at toll booths. It is definitely clear, that these systems, once they are not designed to function anonymously, create a huge database recording the precise time and location of every toll crossing by every tagged car. For instance, the Greek "Taxes-Code for Books" poses the obligation of collecting and retaining for 6 years following data: name, residence, taxation-number, taxation authority, date of entrance, hour and exact point of entrance of the highway/bridge user. The purpose of this data processing is limited to the performance of the contract between the toll collectors and their subscribers. Nevertheless of great importance is to establish policies that will prevent toll-crossing information from being used for purposes unrelated to traffic management. So that ETC databases are not routinely used by law enforcement agencies to track the movement of suspect cars and by both divorce lawyers and labor lawyers to track the movements of people under investigation.

The British government is preparing to test new high-tech license plates containing microchips capable of transmitting unique vehicle identification numbers and other data to readers more than 300 feet away. United States are initiating their own tests of the plates, which incorporate radio frequency identification to make vehicles trackable. Greece is in the very beginning of creating working groups with representatives of both governmental and private sector/university actors for planning their developement and eventually their deployment.

---

[21] See www.statewatch.org/news/2006/feb/08italy-biometric-passports.htm. Very critical: *Juels / Molnar / Wagner,* Security and Privacy Issues in E-Passports, IEE SecureComm 2005, available at: www.cs.berkeley.edu /~dmolnar/papers/papers.html; *Rieback, Crispo, Tanenbaum,* Is your cat infected with a computer virus?, 2006, available at: www.rfidvirus.org/papers/percom.06.pdf; *Schulzki / Haddouti,* Neue Reisepässe: Mit Sicherheit teuer, available at: http://www.sicherheit-heute.de/index.php?cccpage=Verkehr

[22] See *Yoshida,* Euro Bank Notes to Embed RFID Chips by 2005, EETimes, 19.12.2001, available at: http://www.eetimes.com/story/OEG20011219S0016

[23] See http://www.computerworld.com/mobiletopics/mobile/story/0,10801,84543,00.html.

[24] See http://www.zeit.de/zeit-wissen/2006/01/Falschgeld.xml.

The public transportation network of big cities like London, Helsinki, Peking are already using e-ticket. We are talking about a chip-card, used as a recharchable ticket, which will permit the passangers of easier and faster entrance of the transportation means and the public transport companies to avoid fare dodger and to use easier and faster the system of dynamic prices.

3 Legal Implications

From the applications of RFID technology, as described above, following categories of RFID-tags arrear: We have the tags that contain only an item number. Their use is in giving information for the identification of an item. Through the linking of the RFID tag number with a products database one can find out what kind of item this is. Supposing the item information is linked to the purchaser during the payment procedure and further stored to a customers' database one may create customers' purchase profiles. Supposing the item information can be associated to a person either because this person is currently visible or this person is identifiable by other means, for instance with its RFID identification card (i.e. passport) or employee's card, this all may lead to a person's identification for various purposes (customers' profiling, surveillance of workers at workplace). The second category concerns tags that contain an identification number which reveals the identity of a person after the matching of the information contained on the tag with a backend data-base, which holds the information concerning the identity of the person. However the stronger relation to a person is to be found in the RFID tags of the third category. On these tags personal data are directly stored. They are normally active tags and contain information like name, age, nationality and so on. According to that following legal implications may arise.

3.1 Infringement of the right to privacy and data protection

RFIDs tag may be related to personal information. Data protection and the information self-determination is a precious fundamental right that should be protected from the technical development, if this proceeds without taking into account the conformity to main constitutional values and rights. It should be assured that the right to privacy and to data protection will not turn into a caprice of the individual but will still remain an obligation of the democratic society.

3.1.1 3.1.1 Identification and profiling of a person

RFID tags consist of a unique identification number. The use of the tag is to enable identifying and tracking every single item. Everyone who carries at least one so-tagged item is possible to get allocated and tracked. RFID tags function as a unique identifier and the growing interoperability of the system makes allocating and tracking possible worldwide. Beyond that, the link-ability of RFID technology to other databases and their supersets-archives can facilitate the identification process. RFID information can be used independent from information of other sources. But the facileness of the combination of both turns it into a main threat to privacy. As we saw in the application of RFID technology in the retail sector, once tagged objects are owned by persons, it is possible to be related to them. The ability of tracking objects might become an ability to track individuals. Using RFID-Technology retailers might track customers within their shops in order to create profiles of movement which can be used to improve marketing strategies. One should mention that this is possible only by connecting the information obtained by the tagged object that individuals carry with them and their customer or credit cards that they submit at the purchase point. Only in that matter the data stored on the EPC tag relates to the person carrying it. In shopping malls several shops might interlink tracks and analyse the popularity of different parts of the centres by analysing the favourite shopping routes of customers that have already been identified by one of the shops in the

mall. The advantage of it is a better management and promotion policy to increase consumption.

3.1.2 Unnoticed remote reading without line-of-sight

RFID tags can be read without line-of-sight and without overt evidence that they are being read. In addition their small size and their ability of working without any energy supply make them appropriate to be installed hidden. The problem is that radio waves allow data to be processed over a given distance without any need for a direct line-of-sight link with the chip and without the data subject having to take an active part in the process. In other words, data processing can take place without the knowledge of the data subject. Any data on RFID transponders that have not been destroyed or deleted can be read by visible or even invisible readers. The unnoticed remote reading may indeed be used for various purposes without the knowledge of the person in question, for instance for unnoticed surveillance of workers, unnoticed profiling of one's consuming preferences etc.

3.1.3 Use of RFID technology for law enforcement purposes

The state might have an interest on making use of personal data obtained through RFID applications for law enforcement purposes. Here all the applications mentioned above can be used by the Law Enforcement Authorities, under the conditions that every national legislation allow this, for the puropses of prevention, investigation and prosecution of criminal offences. We could imagine the interest of these authorities for the exact identification of the owner of a consumer good related to a criminal offence, or the lists of the movement of cars passing through the toll-controls, the tracking of people carrying RFID enabled IDs or passports, or even RFID implanted tags. Even the use of RFID tags in banknotes can be highly problematic in this perspective. Through RFIDs it will be possible to determine which banknotes were withdrawn by whom from which automatic teller machine, or where those banknotes were then used to buy certain products or services.

3.2 Infringement of the right to personality

RFID technology will contribute to the realisation of the Ubiquitous Computing: in a world of ubiquitous services the interaction of humans with computers should step behind and help us enter a digital world without realising it. The citizens must be fully aware of the innovation and of the data-processing procedures that enable this phenomenon but at the same time concerns them instantaneous[25]. Within a densely populated world of smart and intelligent but invisible communication and computation devices, no single part of our lives will per default be able to seclude itself from digitalisation[26]. Nevertheless one should always be able to retrace the data-processing procedures and have the right to switch onto an "of-line" world. If there is no possibility to do so, this will affect the free expression of the personality of a human being.

---

[25] See *Langheinrich,* Die Privatsphäre im Ubiquitous Computing - Datenschutzaspekte der RFID-Technologie, available at: http://www.vs.inf.ethz.ch/publ/papers/langhein2004rfid.pdf

[26] *Langheinrich,* Privacy by Design-Principles of Privacy-Aware Ubiquitous Computing, p. 7 available at: http://www.vs.inf.ethz.ch/publ/papers/privacy-principles.pdf

Identity Management
Biometrics
PKI-Solutions
Network Security

Identifying Patterns of Federation Adoption

Heather Hinton · Mark Vandenwauver

IBM Software Group
{hhinton | mvanden}@us.ibm.com

Abstract

"I don't see that [federation identity] happening this year, I don't see it happening next year or the year after that--that leaves 2009, and I'll leave that one open … There are a lot of issues, but basically it boils down to trust and antitrust." [Penn06]. Despite this pessimism, federation technology and models are being deployed and are in production now, only not in those areas that they were originally expected. In this paper, we describe several adoption patterns that we have observed and the characteristics that have driven these deployments. Existing business relationships between companies are often strong enough to support federated relationships and are being used as the foundation of present-day federated identity deployments.

1 Introduction

Federation has typically been cast as the "next best thing" to happen to a user's online experience. For example, the much-hyped travel agency example has been used to illustrate the benefits of a federation relationship – when booking travel tickets with an online travel agency, John can seamlessly access the airline (to book seats and identify special requirements), the car rental agency and the hotel. Another equally well discussed example is based on the Internet Service Provider as *uber*-Identity Provider, allowing Jane to authenticate to her ISP and then engage in online shopping. These examples, while excellent for explaining some of the values of a federated single sign-on environment, do not yet have widespread adoption, leading many to question if, not even when, federation technology will be adopted.

What is interesting with these early examples of (hypothetical) federation adoption is that they all describe scenarios where federated technology is used to bootstrap a business relationship between companies. It helped, of course, that the average Internet user could immediately see the value of this type of environment. What hurt, however, is that it was not immediately obvious to the federation partners what value they would realize with this environment, if they acted in any role other than an Identity Provider. And given that they all act as Identity Providers right now, why would they willingly give that up? This leads to the conclusion that federation technology can not (and should not) be driving federation adoption in and of itself.

Our customer experience shows us that federation technology and models are being adopted right now though. Federated technology is being adopted in many environments between companies with **existing** business relationships. Adoption is taking place where there is a need for tighter integration achieved through the loose coupling offered by federation. The adoption/deployment of federated identity solutions that we have been involved with (to date), have all been based on scenarios where:

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2006), 151-160

- Business agreements are already in place upon which a federation-driven trust relationship can be based, and,
- One of the participants is a clear "owner" of the user identities and is responsible for the identity lifecycle management

Federation solutions are being implemented today by companies that already have business agreements in place. It is not the case that trust must be established from the ground up as part of building these federations and hence trust is not an inhibitor for these deployments. Federation solutions are also being put in place in scenarios where one entity clearly owns a user's identity lifecycle. While other entities may need to own/manage user information specific to that user-entity relationship; these other entities have no business model that requires that they factually own the user's lifecycle management.

Over the past two years, we have been involved in many customer engagements regarding these types of federation adoption and deployment. Based on this, we are able to identify four patterns driving federation adoption. These patterns are characterized by business relationships and user lifecycle management characteristics. They can be classified as:

- Employer based federations
- Single Sign-On for companies/subsidiaries
- Specialized content provider federations
- Rich client based single sign-on,

In the remainder of this paper we will describe these patterns, the use cases driving them and the user lifecycle characteristics that drive the need for a federated identity solution. We will finish with our observations and conclusions.

2 Federation Adoption Patterns

In this section we provide an overview of the four patterns that we commonly see driving federation adoption. Each of these patterns has distinct characteristics that help define the pattern. The characteristics of each pattern that stand out as driving the type of federation adoption can be divided into pre-existing requirements (which can be viewed as pre-conditions to federation adoption) and business requirements that drive the need for a federated solution. These requirements are:

- Pre-conditions
  - Online presence
    Federation relationships are typically driven by the need to simplify a user's online experience, starting with the reduction in the number of authentication (or sign-on) steps that a user must undertake.
  - Existing Business Relationships
    Establishing the business and trust relationships over which federated functionality is to be leveraged is not simple. Starting from scratch with the establishment of these relationships to support a federation reduces the appeal of federations considerably. Leveraging existing relationships however makes a federation relationship easier to establish even if only from a process point of view.
- Business requirements
  - Identity Lifecycle Requirements
    Federation relationships allow one partner to off-load the majority of the user's life-

cycle management (account creation, management and deletion) to an authoritative source. Note that some business models dictate that a service provider does not wish to off-load this functionality; in these cases there is a much smaller incentive for federation.

Within these patterns, there are two types of federated single sign-on that may be adopted: *push based* (F-SSO request is pushed from Identity Provider to Service Provider) and *pull based* (F-SSO request is initiated from Service Provider to Identity Provider and then pushed from Identity Provider to Service Provider). Within either type of single sign-on, the user's accounts may be already *linked* (meaning that both parties already understand how to refer to the user) or *not linked* (meaning that each party has an independent means of referring to the user and no way to cross-reference or link these two).

In the remainder of this section we will briefly introduce each pattern example and then discuss the characteristics of this pattern that make it suitable for immediate adoption.

2.1 Employer Based Federations

Employer based federations are those where a company enters into a federation relationship as part of the provision of services to its employees. These services are typically employee-driven services, such as benefits (medical, dental, etc) and personnel (work-life balance resources, fitness plans, etc). This is in contrast to federations driven by the need to expand the company's business, such as supply chain management and customer relationship management.

Interestingly, we have found that employer based federations can be classified as *internal federations* and *external federations*. Internal federations are used to ensure single sign-on within an Enterprise. External federations are those federation relationships that are put in place between an Employer and its (external) third-party service providers.

2.1.1 Internal Federation

The first step in entering into an employer-based federation is for the "Enterprise-as-Employer" to put in place internal federation solutions to provide "Enterprise Single Sign-On". An internal federation solution is one that allows users (employees) to sign on to existing resources within the enterprise, typically from their employer portal. This approach to federation is most easily applied to Web based applications and access. It allows an enterprise to focus on getting their "internal house" in order before moving to an external federation, or employer-based Federated Single Sign-On (F-SSO), solution.

Benefits to an employer for implementing this type of Enterprise-based single sign-on solution include:

- Reduced in-house management costs
- Better user productivity
- Fewer lost passwords
- Easier access to resources required to complete daily tasks
- Preparation for Employer-based F-SSO relationships

Typically internal federation relies on push-based F-SSO and deals with linked users. That is, within an Employer's environment, even though a user may have multiple (seemingly) independent usernames, the Employer is able to establish a mapping or linking of all of these usernames, based on the Employer's ownership and control of the corporate directory.

2.1.2 External Federation

Once the internal house is in order, the enterprise is ready to move to external federation scenarios, exemplified by the "Enterprise-as-Employer" based federation. In this scenario, an employer leverages third-party services providers to provide benefits (e.g., medical, dental, retirement) and services (e.g., work-life balance, HR-related). Most enterprises already have these types of relationships in place; most medium-large employers already "out-source" the providing of these benefits to third-party providers; larger employers deal with tens or even hundreds of third-party providers.

In this federation adoption pattern, the enterprise enters into Federated Single Sign-On (F-SSO) arrangements with trusted third party partners. F-SSO to these third parties is typically triggered from the employer portal, supporting a push-based F-SSO model, where single sign-on information is pushed from the employer to the federation partner.

What makes these scenarios ripe for federated identity solutions?

- Pre-conditions:
  - Online presence:
    Both Employer and Third-party provider are making their services available through on-line techniques such as Web portals, meaning that these parties include password management and online account management services as part of their user lifecycle management functionality. As the employer adds *links* to third-party providers from the corporate portal, the completely independent nature of the two entities is exaggerated.
  - Existing business relationships:
    The Employer and Third-party provider have already had to establish the business and trust relationships required to exchange and manage user information, including such sensitive/private information as Social Insurance Numbers, health care statistics, and so on.
- Business requirements
  - Eliminate duplicate identity lifecycle requirements:
    While the Employer is the authoritative source of an employer's lifecycle management (account creation, account deletion, and all activities in between), the third-party providers must also provide their own, in-house lifecycle management for these users so that the user can access third-party resources only when appropriate. Because users have access to the third-party provider's resources on the basis of their status as an employee, there is no advantage to the third-party provider to manage the user's lifecycle – the third-party provider cannot typically make money by offering additional services to the user.

Typically external, or Employer-driven, federation relies on push-based F-SSO of linked users. The push-based F-SSO takes the employee from the Employer's site to the third-party provider through a link on the Employer's corporate portal. The Employer is able to assert a known, unique identifier to the service provider. In this scenario, it is possible that this value is the user's Social Security Number, or some other value, such as an email address or an employee serial number. Because the users are participating in this federation on the basis of their status as an employee, the employer has the ability to assert this type of personally identifiable information to the service provider.

In moving to a federated identity solution, the employer and the third-party service provider achieve multiple benefits, to both their user experience and to their "bottom line." For example, the employer can now provide seamless access to third-party providers based on federated identity driven single sign-on. This removes the need for the third-party provider to authenticate the employer's user, which in turn will reduce the help desk and user management costs for both companies. The third-party provider can also improve its integration and retention of customers (the employers whose users are now engaged in a federated single sign-on relationship with the provider). An important observation about this scenario is that the third-party provider does not have a business model that involves control of a user repository

Adoption of these federation scenarios is common across all sectors, with health and financial benefit providers being the initial targets for third-party benefit provider single sign-on.

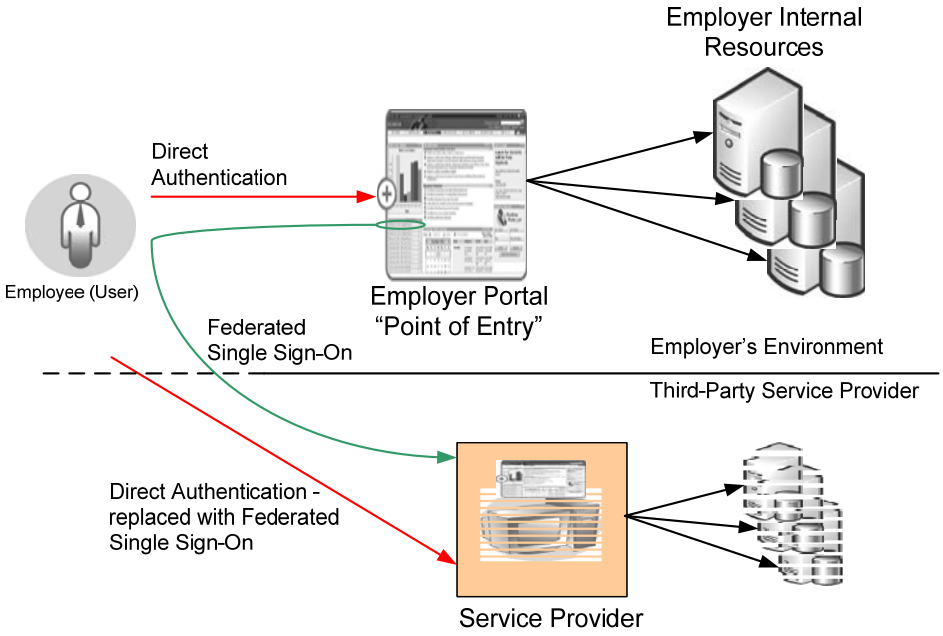

Figure 1: Employer Based Authentication & Single Sign-On

Despite this pattern being referred to as the "Employer Based Federation" pattern, we have found that to date, these employer-based, third-party federations have been driven by the third party providers, or service providers *and not the employers*. Service providers are driving this for several reasons, including reduced user management costs for users that are authoritatively managed by the employer and increased integration (lock-in) with these employers. The employer also stands to benefit from these relationships, namely in the form of reduced benefit provider costs. That is, as provider's costs go down, savings are likely to be passed to the employer.

2.2 Parent Company/Subsidiary; Mergers and Acquisitions

In this scenario, often referred to as a Mergers & Acquisitions (M&A) scenario, a company acquires a (new) company, with its own set of users. Federation provides a quick solution to allow these merged company's users to authenticate to the new parent company and continue to access resources of both the subsidiary and the new parent company. This M&A characterization of this scenario is simply one way of articulating a larger scenario, that of a large/multi-national company with multiple subsidiaries. Large companies often have a re-

quirement for separate "islands" of identity management, perhaps driven by geographic requirements of subsidiaries (employees of a company in Europe often have different identity lifecycle management requirements, if only from a legal point of view, than those in North America), resulting in two separate identity management systems, including authentication requirements for users in those geographies. These subsidiaries are often independent, autonomous entities driven by different regulations and legislation based on their line-of-business and country of operation.

In this overall scenario, federated single sign-on solutions are used to emulate enterprise single sign-on in an environment made up of mutually independent domains. These mutually independent domains may be based on the newly merged/acquired companies or independent subsidiaries. In either case, the desire of the parent company is to provide all of its employees, regardless of location, access to resources, again, regardless of the where these resources are hosted. This example, while common across industry sectors, is especially prevalent in the retail, manufacturing and financial industries.

What makes these scenarios ripe for federated identity solutions?

- Pre-conditions:
  - Online presence:
    Employees and users of both the parent and subsidiary/merged companies need to have quick and easy access to the consolidated company's resources and applications. The consolidated company's company portal should provide easy access to all of these resources and applications, regardless whether the nature of the two entities is exaggerated.
  - Existing trust relationships:
    The parent and merged/subsidiary companies manage independent identity management systems, including authentication systems, within a single company; trust relationships are thus implicit in the nature of the overall parent company.
- Business requirements
  - Eliminate duplicate identity lifecycle requirements:
    While the parent company is the authoritative source of its employee's lifecycle management (account creation, account deletion, and all activities in between), this does not represent the complete set of *all* of the parent companies employees. The subsidiary/merged company is the authoritative source of its employee's lifecycle management. While there may be no long-term benefit for a merged company to continue to manage its employees, subsidiaries often require this separate of users from the parent company's user registry.
    Regardless of where the employees call *home*, however, they should be able to seamlessly access resources across the entire company, be these resources hosted by the parent, the subsidiary/merged company, or yet another subsidiary/merged company of the same parent company.

Typically parent company/subsidiary type federations must rely on both push-based and pull-based F-SSO of linked users. This allows a user to authenticate to their authoritative source (the parent or the subsidiary) regardless of which resource they request (a parent or subsidiary hosted resource), thus triggering the authentication requirement. The parent/subsidiary companies are able to assert a known value referring to the user (thus indicating a form of account linking) because the users are participating in this federation on the basis of their status as an employee of either the parent or the subsidiary.

In moving to a federated identity solution the users and employees of the parent/merged/subsidiary companies will be able to have the same online experience, including seamless access to required resources, even if they are hosted by a different entity within the overall parent company. Duplicate lifecycle management can be removed, as users are migrated from a merged company's system to the parent company. Independent lifecycle management can be tolerated for those subsidiary companies that must maintain a level of separation from the parent company.

The scenario is a specialized form of the employer-based federation, as described before. This pattern allows merged/subsidiary company employee to access both the parent company's resources and resources from other subsidiaries, as required. This pattern adds a layer of single sign-on, allowing two employer-based environments to single-sign-on to each other. Federated identity management and federated single sign-on provide employees of these entities access to world-wide resources, regardless of the *location* and *ownership* of these resources.

2.3 Specialized Content Providers

In the specialized content provider pattern, an identity provider leverages third-party service providers (or specialized content providers) to provide specific content. This allows companies to extend their reach by providing their customers with more, targeted services, which in turn increases customer loyalty. This scenario is largely seen within the financial and telecommunications industry, where customer retention is a key business driver. In this scenario, the driver of the federation is the specialized content provider as an identity provider (where in the previous patterns, the driver for the federation adoption was actually the service provider).

In this pattern, the specialized content provider (SCP) has a relationship with the user, typically as what the user would view as a content or service provider. However, this SCP is also able to act as an Identity Provider because it currently is able to authenticate the user (the first step in acting as an IdP) and because it typically is able to provide additional services to the user, where these services in turn may be provided by an independent service provider (SP).

Unlike the employer-based federation pattern, the user's relationship with these SPs is driven by the user, not the identity provider. That is, a user is able to select which SPs they wish to have a federation (and thus single sign-on) relationship with, within the broader context of their relationship with the SCP. The SCP allows the user to select and customize the set of federation partners from an SCP approved list.

- Pre-conditions
  - Online Presence
    More and more users/customers of a service provider are leveraging an online based approach to management, including such features as bill payment, service management and so on. Both the service provider and the specialized service provider typically provide these types of profile management services for their customers.
  - Business Relationships
    A service provider (possibly Internet based services but not required to be) provides additional third-party, specialized services to its users and thus has the ability to act as an identity provider. These specialized services may or may not be branded for the service provider.
- Business Requirements

- Eliminate duplicate identity management requirements
  We have found that duplicate identity management costs are not a driver for this type of federation. In fact, this type of federation is not guaranteed to lower identity management costs. It is believed that this type of federation will increase customer retention due to a better overall user experience based on "one-stop-shopping" and seamless access to specialized content.
- Increase market share/market uptake for services
  In many cases, a specialized content provider wishes to increase its market share for its services. It is able to do this in part by providing specialized services to a partner's customers.

Typically specialized content provider type federations must rely on both push-based and pull-based F-SSO of unlinked users. This allows a user to authenticate to their authoritative source (their "main" service provider acting as identity provider) regardless of which resource they request (a service provider resource or a specialized content provider resource) triggering the authentication requirement. The service provider and specialized content providers are not able to initially (first time) assert a pre-determined value representing the users, therefore requiring the users to go through a form of account linking to cause this value to be established. As part of establishing this account linking, both the service/identity provider and the specialized content provider may need to collect a "consent" to participate in the federated relationship from the user.

Both the service providers and the specialized content providers tend to be large entities with large customer databases. The specialized content providers that are adopting this type of federation scenario are not *boutique* type providers; as an example, these specialized content providers may include all of the subsidiaries of a media company, where the service provider that provides the entry portal to the federation is a service provider such as a broadband service provider, and the specialized content providers may include VOIP services, search portal services, and cable television services offered by a subsidiary of the broadband service provider.

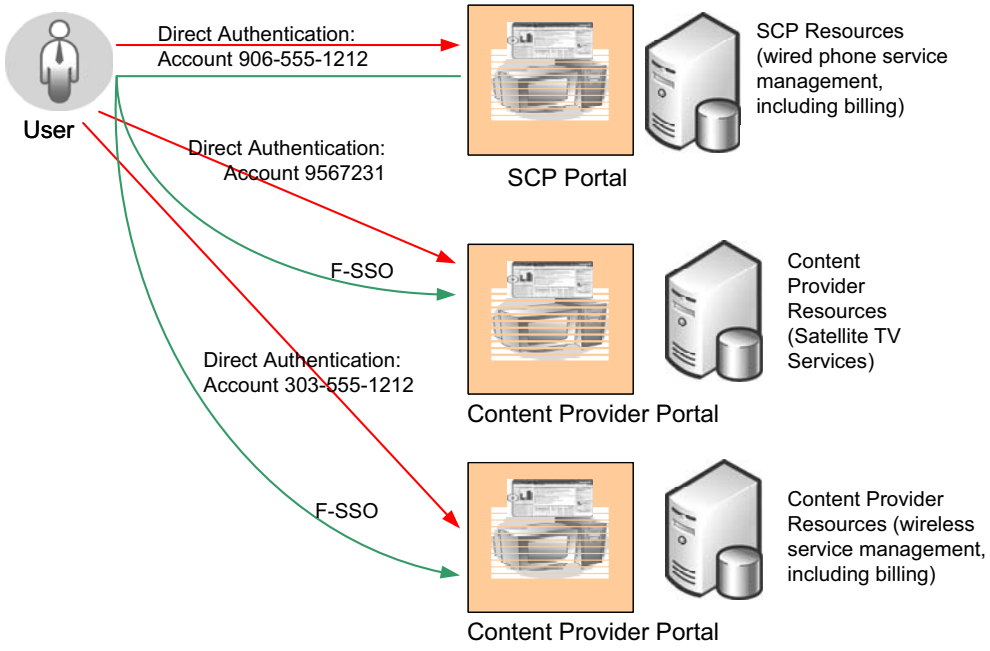

Figure 2: Specialized Content Providers: Federated SSO Replacing Direct Authentication

Security Management Applications

Centrally Administered COIs Using Cross-Organizational Trust

Kevin Foltz · Coimbatore Chandersekaran

Institute for Defense Analyses
Alexandria, VA, USA
{kfoltz | cchander}@ida.org

Abstract

We study collaboration among different organizations using the idea of a community of interest (COI). We consider not just sharing of data, but collaboration requiring tight integration of all elements of the collaborating organizations. Current models for COIs have shortcomings that limit their applicability to real-world situations. Our previous work addresses some of these shortcomings, but its restrictive architecture limits its applicability. All resources were required to collaborate within a single domain forest. In this work, we extend the prior model with new capabilities that allow collaboration across forests.

1 Introduction

Please note the following important formatting guidelines:

We study collaboration among different organizations using the idea of a community of interest (COI). We consider not just sharing of data, but collaboration requiring tight integration of all elements of the collaborating organizations. These collaborations may be long-term mergers, short-term goal-oriented tasks, or ongoing projects between different organizations. Rapid assembly of collaborating partners is required. However, they must still provide the proper functionality and security measures to ensure safe interaction and sharing of resources.

Ideas for COIs have been developed [Cent95][SJYS00][Khur02], but they all have shortcomings. We developed the Centrally Administered Community of Interest (CACOI)[FoCh05] to address these shortcomings. The centralized nature of the CACOI solved many problems, but its restrictive architecture limits its applicability.

Specifically, the requirement of importing all elements of the COI into a single domain or forest limits the way resources are shared and the way users can access the COI. Databases, for example, can be very large and cumbersome to import and export. Also, an organization unwilling to cede complete control of the resources it contributes would be unwilling to join such a CACOI. The CACOI structure simply does not provide flexibility in resource sharing. The resource is either in the CACOI or it is out, but there is nothing in between. This is a benefit for tightness of interaction and establishing a clear security boundary, but this is simply too restrictive for many types of collaboration. In addition, in a CACOI users must log in locally on a different account and different network. These issues limit the CACOI's applicability to real-world problems.

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2006), 311-317

2 Cross-Forest Collaboration Goals

In this work we propose a new model, the Cross-Forest COI (CFCOI), which captures much of the transparency and distributed nature of the P2P approaches, improves resource management, and maintains the functional richness and administrative simplicity of the CACOI model. The main idea of the CFCOI is to allow collaborations to extend across multiple forests instead of being confined like the CACOI to a single domain or forest. To achieve this while maintaining functionality and security the CFCOI implementation has three main goals.

The first goal is authentication. In a CACOI authentication is done within the domain or forest. This is a built-in function of domains, and is easy to implement. Once implemented, the identity is secure, and authentication and authorization are done in a natural way within the domain. When creating COIs across forests, the notion of identity must change, since the identities within the domains will not carry across domains, leading to problems with authentication. There must be a new way to authenticate across forests to allow consistent treatment of users within the COI.

The second goal for the CFCOI we will call “credential shaping.” In a CACOI authorization is based on the identity of the user within that domain or forest. When extending identities across domains a new issue arises. One user may be a member of multiple COIs with overlapping resource utilization in overlapping forests. A problem can arise when a user is a member of two COIs, one of which has access to a resource, and one of which is denied access. The problem is one of authorization. There is no clear-cut way to determine access to a user who is a member of both COIs. Credential shaping is a way to allow the user to specify which credentials to use for access.

The third goal is confinement. When users are distributed across different forests there is an increase in the risk of data being released without authorization from the COI. When files and users share physical hardware, the boundary between COI and non-COI is weaker than the CACOI, where the hardware and user identities are all contained within a single COI domain or forest. As a result, confinement must be managed at the level of individual resources, and not by the domain or forest.

3 Cross-Forest Collaboration Solutions

To address the three goals, we first establish the notion of identity across forests. This is done by credential providers (CPs). These CPs are located in each domain or forest, and they sign credentials for principles in their local domain or forest. The credentials include information such as Name, Group, Role, Domain, and COI. The principle requests the CP’s signature for a certain combination of the indicated fields. If that combination is valid (i.e. if that principle has that name, is a member of that group, acts in the indicated role, and is a member of the indicated COI), then the CP signs such a credential. CPs in different forests establish agreements about COI membership rights and privileges. When these credentials are sent across forests, the local CP in the destination forest can verify the signature on the credential, based on the agreement between the forests that is established in advance, and verify the indicated principle’s membership in the indicated COI. Thus, the identity is preserved across forests.

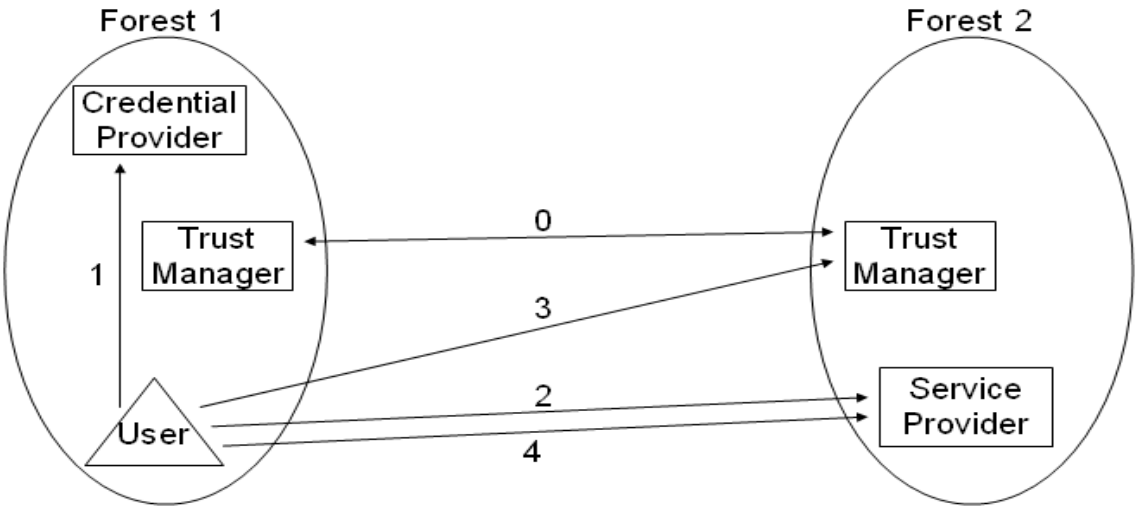

Figure 1: Cross-forest service access using WS-Trust. Step 0: TMs in the two forests establish rules and policies for cross-forest interactions. Step 1: A user first obtains a credential, indicating COI membership. Step 2: The user requests a service from Forest 2. Step 3: The user is referred to the TM in Forest 2 to verify the credential issued by Forest 1. Step 4: The TM in Forest 2 checks its policies, verifies the credential's validity, and refers the user back to the service provider with local authorization.

To address the first goal, cross-forest authentication, we make use of the Web Services Trust (WS-Trust) specification. This provides a standardized way to allow fine-granularity sharing of resources across different organizations. The specific implementation used for the CFCOI model is the Active Directory Federation Services (ADFS) add-on to Microsoft Server 2003. The organizations, modeled as domain forests, establish Trust Managers (TMs) that enforce sharing policies across the forests. When a principle from a CFCOI requests access to a CFCOI resource in another forest, the resource refers the principle to the TM in the resource's domain. This TM then checks its policies to determine local privileges and access rights to the resource. The trust relationships established in the TMs' policies allow a specified type of access to a specified resource to be granted to a specified principle in the CFCOI. If the principle is allowed access under the TMs' policies, the TM adds its own credential and refers the principle back to the resource, which then provides access to the principle. When modeling resources as web services, this transaction is very clean. The resource simply uses the http protocol to refer the principle to the TM, and the TM uses the same technique to refer the principle back to the resource, but with extra information that allows access. The net result is that resources in one forest can authorize principles in another forest with the help of the CPs and TMs. The basic idea is illustrated in Figure 1.

This approach is based on the industry practice of identity federation. In identity federation, different organizations agree to authenticate each others' users based on pre-negotiated trust policies. For the user, the net result is that logging into one domain allows access to the other domain's services transparently. Instead of requiring the user to authenticate in the second domain, the authentication is done automatically by previously established rules between the two domains. The Liberty Alliance and WS-Federation are two standards that address identity federation. Our implementation uses the WS-Federation approach, which fits cleanly with the CACOI domain concept.

To address the second goal, credential shaping, we use the CPs to provide the appropriate COI information for a user at a given time. The COI choice can be done at log-in, or it can be

chosen dynamically, to allow switching between two COIs. In either case, the principle is in control of which COI is on the credential, thus allowing the appropriate access to resources.

For the third goal, confinement, we use digital rights management. Documents and other resources are encrypted. To access them, members of the appropriate COI must prove membership in the COI. Rights Managers (RMs) are present in each forest involved in the COI, and principles consult the local RM for access to COI documents. Documents created within the COI are encrypted by the RMs to prevent unauthorized use. The document itself contains information about what credential is required for the requested access, and indicates a RM or COI where this credential can be obtained. If a principle is acting within a given COI, this credential can be automatically granted based on the CP-signed credential. This allows transparency of rights management to the appropriate principles.

4 Using the CFCOI

In comparing the CFCOI to the CACOI, the main functional difference is that the CFCOI allows a COI to be established across multiple forests, whereas the CACOI requires the entire COI to exist within a single forest. The benefit of the CACOI design is that authentication, authorization, resource sharing, and confinement all come at relatively no cost. The domain and forest structure includes built-in mechanisms to implement all of these features. However, the cost of using a single forest is often too high for the parties involved, and an opportunity for collaboration is lost. The CFCOI provides more flexibility while preserving the benefits associated with the CACOI, thus allowing collaboration to take place where it otherwise would not be possible. Figures 2 and 3 show the CACOI and CFCOI model architectures.

A simple example showing the benefit of the CFCOI is the sharing of a large, complex, and proprietary database. A large company has a database of customers and account information, and they would like another small company to use its proprietary analysis techniques on parts of the data to determine if a long-term agreement would be worthwhile. The large company wishes to provide limited access to the database, and the small company wishes to interact directly with the database in a tight way. A COI is a good way to do this, since it is established quickly and involves a short-term, goal-oriented task with tight integration of resources. A CACOI would be good, except there are two problems. First, the database must be exported to a new domain. This would be very difficult and time- and resource-consuming. Second, employees of the small company would have full access to the database in the COI domain, since the database would be under the control of the COI, not the large company. This is not acceptable to the large company, since the database is too sensitive.

The CFCOI is a good solution to the situation. Both companies can agree on a level of access for the small company, and keep the database in the large company. TMs can be used to allow cross-forest access. The CFCOI avoids the costly import/export of the large database, and allows the large company to maintain control of the database while allowing limited access, as specified by the TM policies. The functionality and ease of use are preserved, since the database authentication and authorization are transparent to the users.

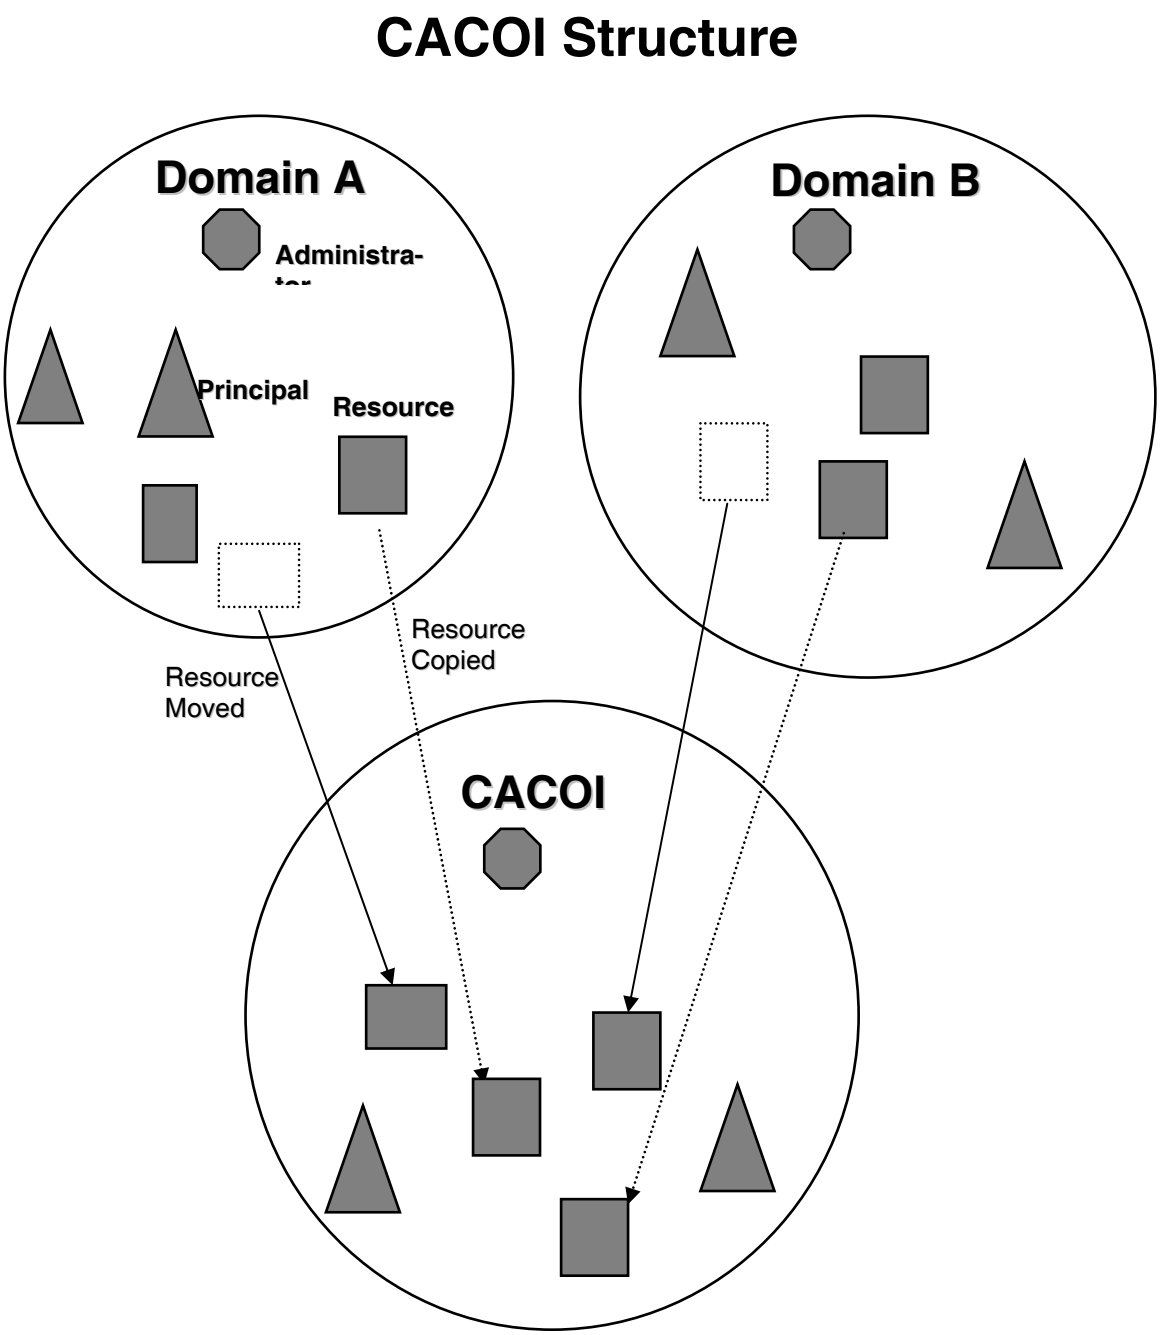

Figure 2: Structure of a CACOI. Domains A and B form a CACOI. All CACOI resources are either moved or copied into the community domain, and new principals are created in the CACOI. Once the CACOI domain is established, no outside communication is permitted. Users must log into the CACOI user account to function in the CACOI.

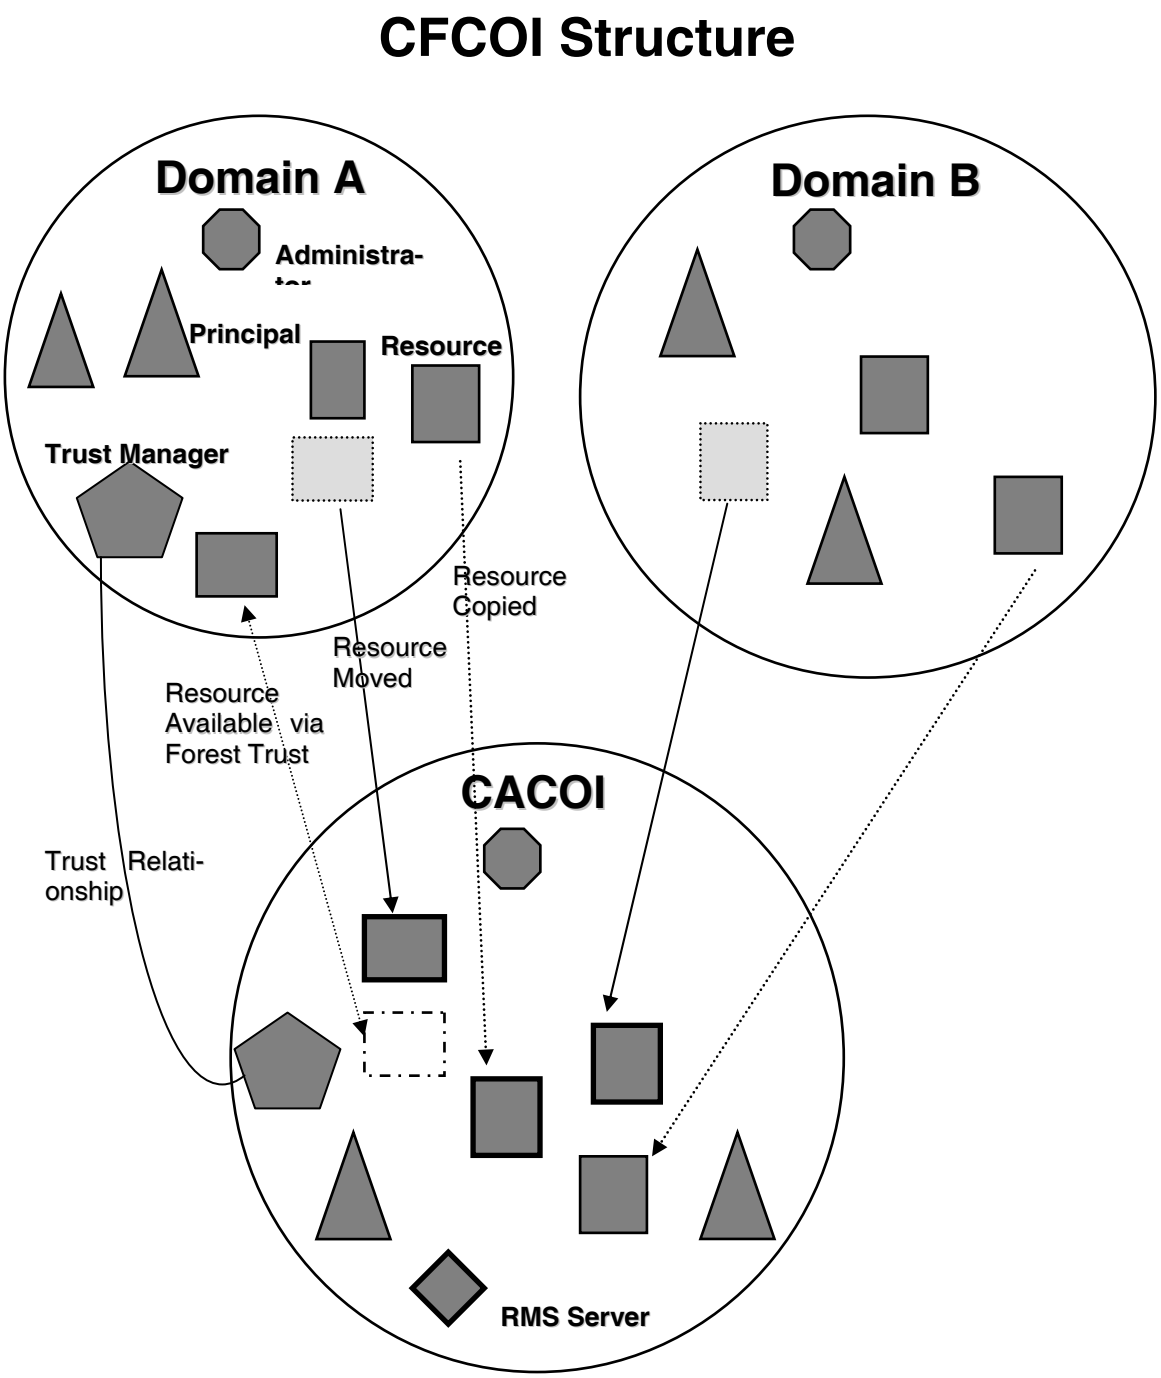

Figure 3: Structure of a CFCOI. Unlike the CACOI, some resources are made accessible to the community but remain in their original domains. Trust Relationships between Trust Managers allow secure sharing of resources. Resources with dark blocks are protected by Rights Management.

Another example, highlighting credential shaping, is work involving proprietary or otherwise restricted data. An individual may be assembling data to write a report on a topic, but the report must have only publicly available sources. The user in question might have access to many different COIs in different domains, allowing access not available to the public. In assembling the resources, the user specifies membership in no CFCOI when logging in. This prevents unintentional access to restricted data, and ensures that all work is releasable to the public. This could also be used when collaborating with an individual with no access privi-

leges. Logging in without COI access ensures that all interaction with that individual is safe, in the sense that sensitive information will not be disclosed accidentally. This ability comes for free in the CACOI, since users in the CACOI forest are isolated. However, the point here is that even with the CFCOI, which has the benefit of cross-forest interactions, the "isolation" of the CACOI is still achievable when desired.

An example showing the benefits of data confinement is the case of someone working on the road, where only public Internet access is available. The individual copies working files to the local machine, works with them, and sends updated versions back to the COI. When the individual leaves the terminal, files are left behind (on the desktop, in the "Recycle Bin," or in hardware locations accessible with basic or more advanced extraction tools). In a CACOI this type of access is not permitted, since all principles, resources, and services must remain within the CACOI domain or forest, but if it was possible, there is the potential for data leakage. In a CFCOI, this type of access can be permitted, but after the individual logs out of the COI any remaining files are protected by encryption, and the RM in the CFCOI, which is needed for decryption, is not accessible to the public. Hence, the files left on the computer are no longer accessible, even though the full text of the encrypted files may be present.

5 Conclusion

The CFCOI builds on the CACOI model, which itself addressed many of the shortcomings of previous models of collaboration. Drawbacks of the CACOI's single domain/forest model are identified, and the requirement of keeping all resources in one forest is relaxed. This allows a more flexible COI model, capable of supporting a more diverse array of collaboration types. This flexibility presents new challenges in authentication, authorization, and confinement. To address these concerns, the idea of a cross-forest COI-based identity, enabled by local CPs and TMs, is established. This enables cross-forest trust relations, cross-forest authentication, and digital rights management. This is easy to implement using existing hardware, software, and protocols. The resulting CFCOI model preserves many of the features of the CACOI while removing the strict single-domain/forest structure. This makes the CFCOI model more suitable to real-world collaboration problems.

References

[Cent95] Department of Defense (DoD) Goal Security Architecture (DGSA). Center for Information Systems Security (CISS), Defense Information System Security Program (DISSP), Version 3.0, 30 September 1995.

[SYJS00] Shands, D.; Yee, R.; Jacobs, J.; Sebes, E. J.: Secure Virtual Enclaves: Supporting Coalition Use of Distributed Application Technologies. Proceedings of the Network and Distributed Systems Security Symposium, San Diego, February 2000.

[Khur02] Khurana, H.: Negotiation and Management of Coalition Resources. Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Maryland, College Park, MD, 2002.

[FoCh05] Foltz, Kevin; Chandersekaran, Coimbatore: Sharing Resources through Dynamic Communities. Proceedings of ISSE 2005.

Improving Assurance of Information Security RoI

Michael D. Barwise

Integrated InfoSec
6 Maple Green, Hemel Hempstead, Hertfordshire HP1 3PY, UK
mbarwise@bcs.org.uk

Abstract

Changing business expectations of information infrastructures have imposed new demands on security architectures. Established technocentric perimeter-oriented security architectures are yielding ground to business-driven deperimeterised architectures that assume extensive information and resource sharing and global virtualisation. These changes provide the opportunity to take a new approach to security architecture specification, based at its highest level not on costing of reactive countermeasures to current technical threats, but on prioritising the allocation of resources to robustly and proactively protect business information assets against business-oriented exposures. This permits tighter specification of both requirements and budgeting with a concomitant improvement in RoI, but depends on a new approach to management described here.

1 Changing Security Architectures

Until quite recently, information security budgeting was a relatively simple affair. Perimeter-oriented security technologies and personnel awareness campaigns were the essence of the solution for the majority of businesses. The components, and thus the component costs, of such a security implementation were generally well-defined and their selection was driven almost exclusively by techno-centric architectural decisions. However, the resulting architectures (essentially an assemblage of secured cells each consisting of a hard shell surrounding a soft centre, coupled by hardened data corridors to similar cells) have never been optimum. They generally bear no close relationship to business structure, essentially echoing instead the geographical distribution of business premises, data centres and facilities, they suffer from single points of failure, and their greatest weakness has always been at the user interface endpoints.

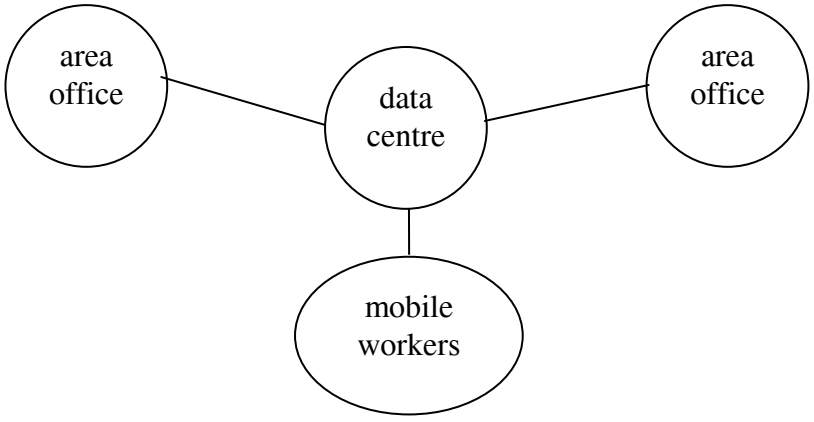

Figure 1: Traditional Perimeter-oriented Architecture

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2006), 318-325

Such architectures also tend to reinforce a reactive approach to security, with the emphasis on countering technical threats from outside the perimeter as they emerge. The reactive stance makes the relationship between implemented security technologies and levels of protection achieved hard to demonstrate, and there is seldom any strong correlation between the value of the protected assets and the cost of protecting them. Uncertainty both in predicting the nature of emerging threats and in ensuring that reactions are appropriate and proportional both contribute to this. Typical of the problems faced are "zero-day exploits" (software vulnerabilities that are exploited by an attacker immediately they are publicised, prior to any remedy being developed). Such threats are extremely difficult to predict, and in the absence of proactive management for robustness can cause losses that are not only potentially large but also very difficult to quantify. Clearly, although the budgeting process is relatively simple to execute for traditional perimeter-oriented architectures, assurance of security and cost-effectiveness are both low, so poor RoI is to be expected.

1.1 New Demands

The expansion of technological facilities driving twenty-first century business has created new departures in infrastructure, not only in terms of scale but also in nature. The increasing ubiquity of web services, virtual networks and grids, distributed and shared data, mobile working and a host of other business demands that tend to destructure corporate network boundaries result in the requirement for tight integration of security solutions with business processes, rather than primarily with network physical infrastructures as in the past.

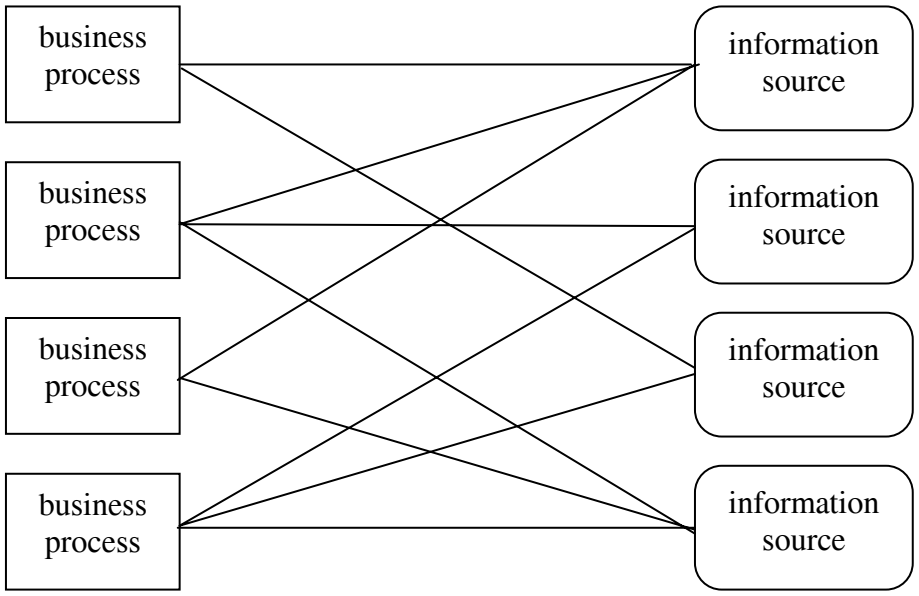

Figure 2: Business-oriented Architecture

Information security must become an intrinsic component of enterprise architecture. This need is driven by ever more complex mappings between data sources and users. Users may reside on partner networks over which data providers have little or no security authority, and business information is likely to be assembled from numerous data sources not necessarily under common control. These developments force us to take a more proactive approach to security architecture decision-making that is based on business structure and processes. Nevertheless, the ultimate security implementation remains a technical issue: the creation of a network of secure channels between segregated and secured assets and segregated and secured

users that has no soft centre and has a much less well-defined corporate perimeter or, ultimately, even no perimeter. The change of emphasis provides an associated opportunity to improve the quality of the budgeting process and consequently both the cost-effectiveness and the assurance of security, but it requires a new management approach that encompasses both business expectations and technical solutions, and a rigorous methodology allowing business and technical personnel to collaborate effectively towards the common goal of security without loss of focus.

2 A New Approach

For the technical security architect, this blurring of boundaries has made the primary requirement one of understanding the assemblage of business processes and their associated information flows. For security management and budgeting personnel, the aim has becomes to prioritise the allocation of a security budget with reference to the business value of the various information assets to be protected, the first task being to place values on those assets in the context of the business processes they serve and the hazards to which they are potentially exposed. For both parties this must be in the first instance a business-centric matter, not a techno-centric one, so we must speak of information, not of data, must consider detriments in business rather than in solely technical terms, and must define hazards within a business, not a purely technical, frame of reference.

Whereas in the past techno-centric security architectures were generally derived from a single-stage technical requirements analysis, the need is now for a multi-stage analysis that bridges the business-technical divide on the lines of a Zachman framework [Zach99], but, and critically, with an explicit mechanism for maintaining conceptual integrity between layers.

The components of this analysis process are

1. recognition of business structure
2. discovery of business processes and their boundary communications
3. discovery of the information assets used by each business process
4. for the identified information assets, discovery of the business detriment caused by specific hazards
5. costing of identified business detriments
6. assigning values to the information assets for each attribute by mapping back from the discovered detriments
7. mapping the information assets to the physical infrastructure
8. aggregating the information asset values by attribute against the infrastructure
9. specifying proportional budgets for securing infrastructure components against breach of the attributes
10. delivery of technical solutions with reference to the required attributes and technical threats, and within the proportional budgets

Clearly this is a more complex specification process than has generally been deemed necessary in the past, but it must be recognised that it has to fulfil a much more sophisticated brief, not least in bridging the business/technical cultural divide to ensure that the necessarily complex business security vision is accurately reflected in the final technical implementation. It is also imperative to recognise that the way the process is performed will have an overriding influence on the quality of the results. Most extant security management and architecture development methodologies gloss over this distinction, but experience shows that it is absolutely

critical to recognise it if consistency is desired. Uncertainty, and consequent variability in judgement quality, must be minimised as far as possible by some formalised process of debiasing [KaST99]. The key contributors to uncertainty in the current context are subjectivity and linguistic imprecision, which have been widely discussed in other spheres of risk judgement [MoHe90]. They will primarily affect components 3, and 4, 5 in cases where finite data are not available (for example when assessing regulatory liabilities), and possibly 2, depending on the complexity of the business process set. But these two causes of uncertainty can be minimised by making use of well-defined parameter sets, standard semantics and rigorously defined investigative methods.

2.1 Defining the Parameters

The precise definitions of business structure, processes and information assets will obviously be dependent on the nature of the enterprise. It is probable that existing business process management systems and tools can provide support here, remembering that for a complete security implementation any manual processes and information sources must also be included. However, the distinction between data and information is important. An information asset is defined jointly by the information used and the specific business process that uses it. The same actual information used by a different process constitutes a different information asset. This allows for the possibility of differing impact of a given breach depending on the nature of the business process that is affected.

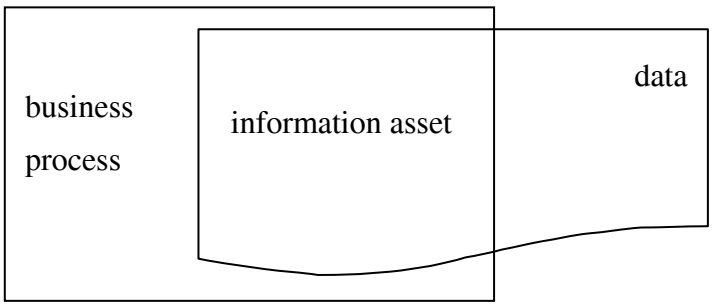

Figure 3: Venn Diagram of an Information Asset

2.2 Criteria for the Semantic Structure

To consistently assess information asset values, it is necessary to work within a well-defined semantic structure that expresses exposures and detriments unequivocally. Linguistic imprecision is one of the major contributors to inconsistency in risk decision-making, and much care must be exercised to minimise it. Any such structure must be transportable across the business-technical divide without ambiguity. Consequently it must be small in scale, and must exclusively make use of terminology that is simply expressed in both spheres. Considering first the exposure vocabulary, concepts such as “loss of reputation” from the business perspective, or “viruses” from the technical, are not useful as they are vague and their implications are subjective.

2.2.1 Information Attributes

The established information security trilogy of information attributes: “confidentiality”, “integrity”, “availability”; are a good starting point for an exposure structure. Each can be explained simply in terms of both business expectations and technical requirements. But this

Awareness Raising
Compliance
Data Protection
Cyberspace Regulation

Internet Early Warning System: The Global View

Norbert Pohlmann · Marcus Proest

Institute for Internet Security
University of Applied Sciences Gelsenkirchen
Neidenburger Str. 43, D-45877 Gelsenkirchen
{norbert.pohlmann | marcus.proest}@internet-sicherheit.de

Abstract

The constantly growing importance of the Internet for our knowledge and information society makes it necessary to analyze and be acquainted with its status beyond the limits of the individual network operators. Only precise knowledge of the normal status makes it possible to detect anomalies which influence the functionality of the Internet.
With the help of the probe-based Internet Analysis System, which is currently being implemented as a research and development project of the Institute for Internet Security at the University of Applied Sciences in Gelsenkirchen in collaboration with the German Federal Office for Information Security (BSI), it is intended to create and analyze local and above all global perspectives in order to make the generation of early warnings possible.
Particular focal points of the project are the collection of information in compliance with data protection regulations and optimization of the amount of information, so as to be able to store information in the long term and therefore allow the analysis of trends and developments over long periods.

1 Introduction

We have all experienced the situation: you are sitting in a traffic jam and all you can see is a long line of cars in front of and behind you. In this situation, without any assistance, you do not have an overview of the problem. There is no direct information concerning why the traffic jam has come about, how long it is, at what point of the traffic jam you are located or - the most important information - when the traffic jam will be over. As this is a problem faced on a day-to-day basis by thousands of motorists, solutions have been developed to overcome the lack of information. In Germany there is a close network of traffic counter loops which record the traffic volume and situation on the motorways/freeways. Important information about traffic jams is provided by means of radio announcements, SMS, telephone and the Internet, while modern navigation systems process the information directly when planning the route to be taken. Through the use of these resources, motorists are "liberated" from their constricted local view of the situation and can take decisions in good time on the basis of the global information available, e.g. leaving by the next exit and using an alternative route.

This situation can also be applied to the perspective that the network operators have today of the Internet. As a rule they have only a local perspective, i.e. an overview of their own network segments and the communication data that are transferred. If problems occur here and are detected, they can be rectified quickly and systematically. However, if it becomes apparent that a problem has occurred that is not within their own domain of action, or if the re-

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2006), 377-386

quired perspective is lacking, the situation is more difficult. It is often not clear where the problem comes from, and for the correction of the problem we are reliant on third parties.

The global view required in order to detect the problem and select the correct solutions is missing. Such a global perception is difficult to achieve on the Internet as people like to play their cards close to their chest. The precise internal network structure, communication connections and topologies are often treated confidentially by the network operators.

Furthermore, in order to achieve a global perspective, there are a few challenges that have to be solved: communication data are relevant in principle to data protection, the quantities of data are enormous, the data rates are sometimes so large that they cannot always be analyzed "live", while long-term storage of the communication data in order to observe long-term developments appears to be impossible. Moreover, the question also arises of who feels responsible for creating a global perspective?

Nevertheless, the Internet has developed into an omnipresent medium over the past few years, without which very large areas of the economy, research and private life would be unimaginable today. For this reason the analysis and knowledge of the medium known as the Internet in its totality is of particular significance in order to be able to assess its development and guarantee the future functioning of all the services it provides.

2 Aims and Task of the Internet Analysis System

The task of the Internet Analysis System on the one hand is to analyze local communication data in defined subnetworks of the Internet, and on the other to create a global perspective of the Internet by bringing together the large number of local perspectives.

The functions of the Internet Analysis System can be divided up into the four subsegments of pattern formation, description of the actual status, alarm signaling and forecasting.

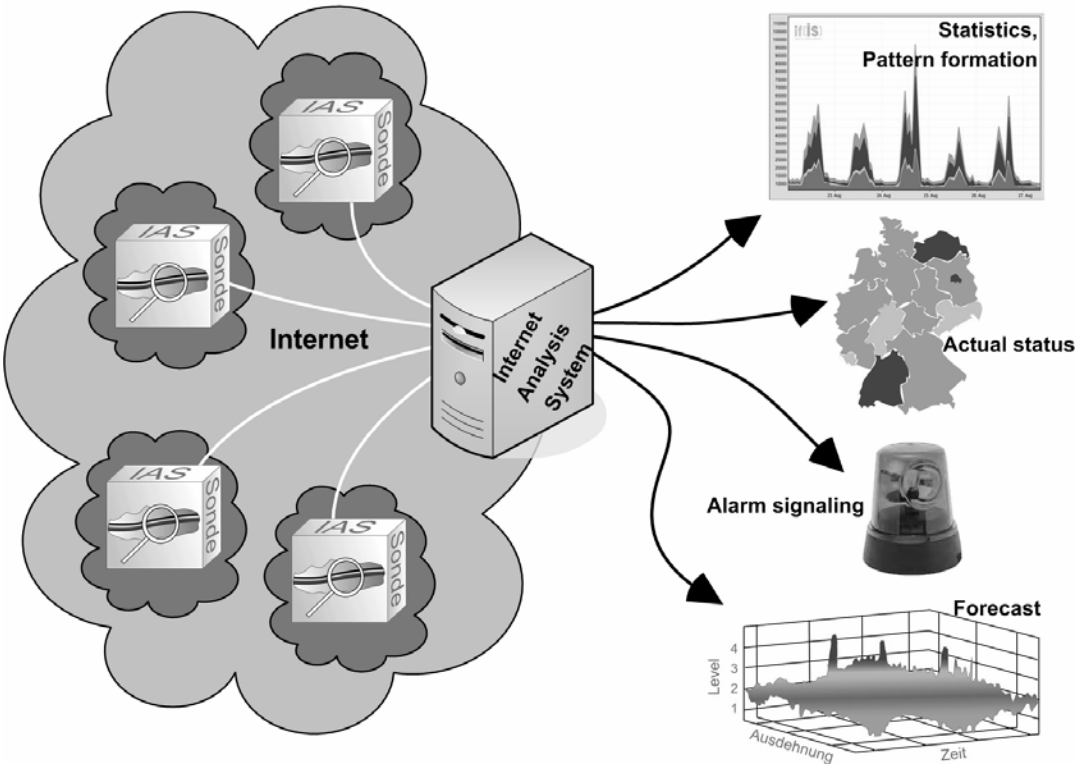

Figure 1: Tasks of the Internet Analysis System

The main task of pattern formation is a comprehensive analysis and interpretation of the communication parameters of Internet traffic, with the aim of detecting technology trends, interrelationships and patterns which represent the various statuses and perspectives of the Internet. On the basis of this knowledge a search is carried out for anomalies among the current measured values and the causes of status changes analyzed and interpreted. Here it is important to find out whether the status anomalies have a natural origin, for example as a result of a technological change, or whether they are attributable to a wanton attack. .

With knowledge of the current status of a communication line and the use of historical - i.e. previously collected - information (knowledge base) it is possible in the case of significant changes to traffic volumes or communication data to concentrate on analyzing these anomalies on the basis of which measures can be initiated to protect and maintain the correct functioning of the Internet.

A further important function is the visual depiction of the Internet status similar to a weather or traffic jam map. Here intuitive depictions are being developed with which the most important parameters are discernible at first glance.

Through the examination and analysis of the extrapolated profiles, technology trends, interrelationships and patterns it will be possible by means of an evolutionary process of the acquired results to make forecasts of Internet status changes. In this manner it is possible to detect indications of attacks and important changes at an early stage and forecast the effects of the damage [Pohlmann2005].

3 Mode of Operation of the Internet Analysis System

The Internet Analysis System consists of probes which passively access the network traffic of the communication lines of various networks and count communication parameters at various communication levels. In an evaluation system the communication parameters are evaluated from various aspects and displayed in a clear manner. Illustration 2 shows the interrelationships between the components involved in the Internet Analysis System.

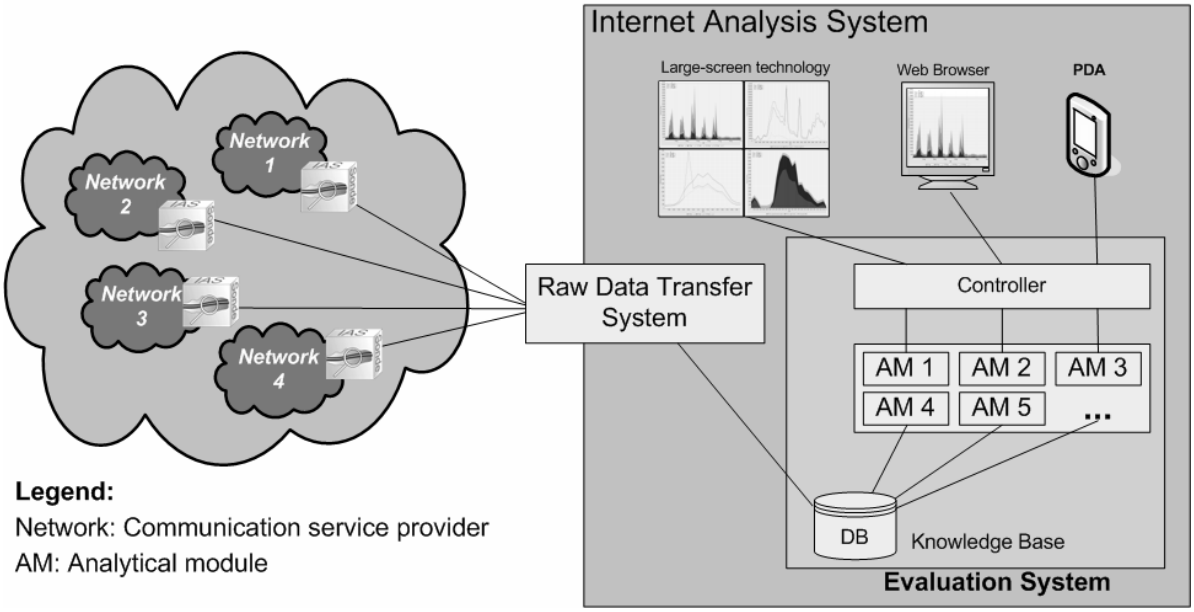

Figure 2: Components of the Internet Analysis System

In order for the Internet Analysis System to be able to supply meaningful results it requires as great a quantity of raw data as possible, i.e. very many counters of various communication parameters at all communication levels over time. All the analyses performed by the evaluation system are based on these raw data. They consist of aggregated counter readings with which the system is provided by various networks. The left-hand side of illustration 2 depicts the Internet, which consists of a combination of numerous networks. In each of these networks telecommunication service providers (ISPs, companies, universities …) assume the tasks, for example, of providing Internet access to end-users, content or other services. The Internet Analysis System acquires its raw data from probes implemented and operated by the network operators. The raw data are transmitted via the Secure Raw Data Transfer Protocol (RDTPs) specified for this purpose.

The probes can send the raw data to one or more evaluation systems. Each network is able to perform its own analyses with its evaluation system. In order to achieve a global and representative perspective of the Internet, probes have to be operated in various types of networks, such as the Global Tier One Provider, Transit Provider, Eyeball Internet Service Provider, Content Provider and Business Networks, as well as various regions (see also in Internet Germany [Dierichs2005]).

4 Tasks and Mode of Operation of the Probes

It is the task of the probes to extract information from a communication data stream which provides details of the status and use of the communication line and the network or networks behind them. Here all information should be retained which is required to detect misuse, a misconfiguration, trend developments or an attack situation. At the same time, however, the quantity of information should be restricted to the minimum required so that the information can also be considered and analyzed retrospectively over long periods. A further important point for the operation of the probes is that the information extracted by the probes, the so-called raw data, does not contain any information relevant to German data protection law.

Technically the network connection is accessed passively and the communication parameters of the various protocols at the communication levels counted. The results of the counts are transferred at defined intervals to the raw data transfer system.

5 Principle of Raw Data Collection

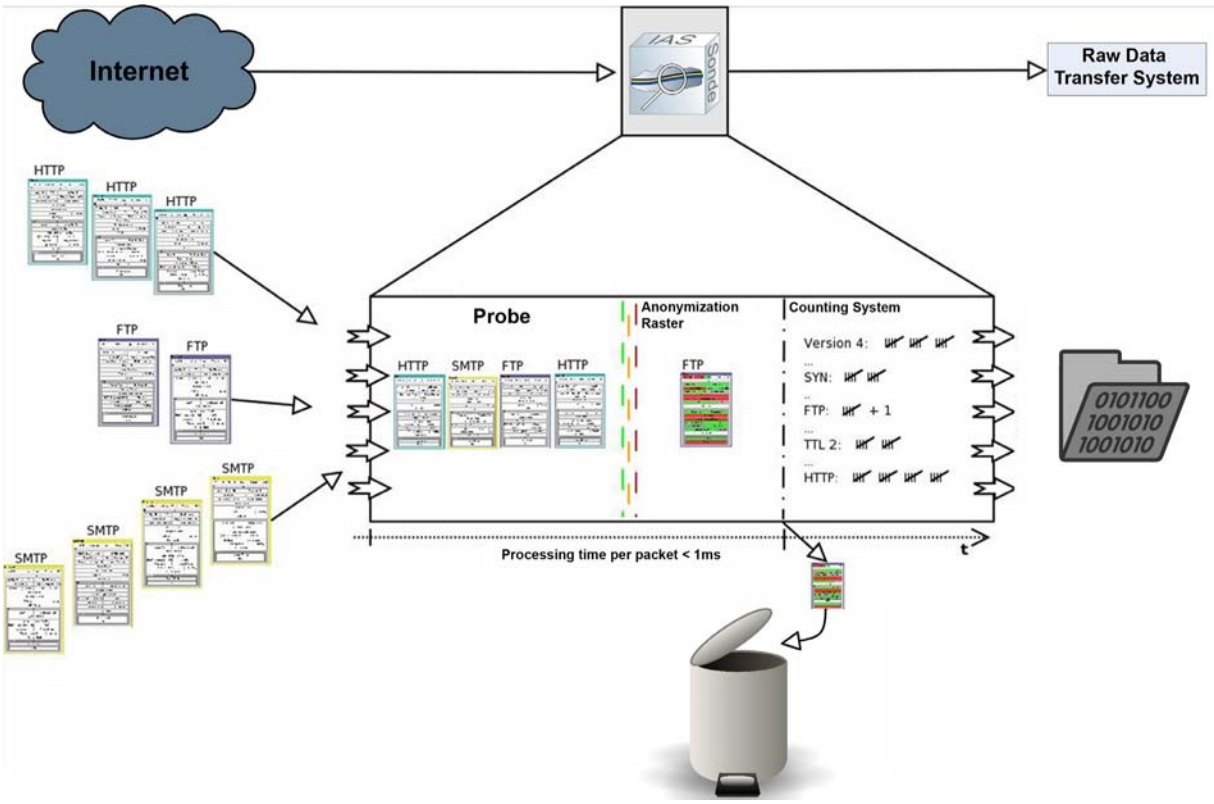

Figure 3: Principle of raw data collection

Illustration 3 shows the principle of raw data collection by the probes. This is divided up into three sections. The Internet is represented on the left. Packets of three different application sessions are shown: related HTTP packets, an FTP session and an SMTP session. The probe is located in the middle of illustration 3. The packets of the three applications are accessed passively by the probe one after the other in their random order and evaluated. The packet that is accessed is channeled through several analysis categories, each of which is responsible for a certain protocol. These evaluate strictly defined communication parameters in the protocol header at the various communication levels which are not relevant to data protection law. The counters allocated in the counting system are incremented according to how the header information of the packet is filled out. The frequency of certain header information is recorded in the same way as on a tally sheet. For example, in illustration 3 the accessing of the FTP packet is recorded by incrementing the FTP counter by 1. The raw data are therefore aggregates of counters, i.e. counters of communication parameters that have appeared at the various communication levels over a defined period. The packet - in illustration 3 an FTP packet - is immediately deleted physically, i.e. irreversibly and without trace, by the probe [Proest2005].

ID	*Description*	*Count*
131134	IP (Protocol Number 6)	: 18.854.151
131145	IP (Protocol Number 17)	: 1.123.149
327708	TCP (Flags: SYN)	: 334.435
327723	TCP (Flags: FIN/ACK)	: 480.697
327724	TCP (Flags: SYN/ACK)	: 275.779
545857	HTTP (Request Method POST)	: 2.026
545861	HTTP (Request Method GET)	: 293.616
545863	HTTP (Request Method HEAD)	: 18.992

Figure 4: Counting system in the probe

Reconstitution of the context of a packet or only a communication parameter is not possible or necessary. At definable intervals the counter readings (raw data) of the probes can be transmitted to the raw data transfer system. All of this information is completely anonymous, as shown in Illustration 4. On the right after the colon are the counter readings for the header information specified on the left. Each line stands for a counter. On the left-hand side of the colon is the count-if function (appearance of the corresponding communication parameters) and on the right the number of packets which contained the communication parameter during the defined measurement period. For example, line 2 of the raw data shown indicates that 1,123,149 packets with the IP protocol number 17 (UDP) appeared in the prescribed time. The count-if functions and their codings are specified in a versioned XML file.

The raw data transfer system functions as a server to which the probes can connect in order to transmit their raw data for a defined period. This is a unidirectional connection, meaning that a connection can only be established from the probe side. A probe can transmit the raw data to one or more raw data transfer systems. An example of a typical configuration is that every 5 minutes the raw data, for example 20 kilobytes in size, are sent to its own and a central raw data transfer system.

As the raw data are only a statistical formulation of the actual communication data, it would also be sufficient if not every packet were considered, but for example only every 10^{th} packet. This aspect can be a pragmatic solution in the case of very high communication data rates, without producing a different result from a statistical point of view.

6 Evaluation of the Collective Raw Data

The actual evaluation and processing of the collected information takes place in various analytical modules (AM) of the evaluation system. In Illustration 2 these are designated "AM1" - "AM5". The modules procure the information exclusively from the knowledge base (raw data and evaluation results). The aim of the various modules is the compilation of profiles, statistics and interrelationships, as well as the detection of where threshold values are exceeded and the graphic processing of the raw data and evaluation results.

As the raw data consist of completely anonymous information, they could also be exchanged between various network operators or collected at a central point in order to be used as the basis for a global perspective and analyses of an IT early warning system.

7 User Interface of the Internet Analysis System

There are many conceivable methods of displaying the results of the Internet Analysis System. Illustration 2 contains examples of the following: large-screen technology, a Web Client and a PDA. The large screen technology serves the purpose of continually updating the display of certain statistics, profiles and current statuses. By means of an intelligent client, more extensive analyses can be carried out with the evaluation system and the results recorded. Additionally, warning messages from the system can be received in mobile form, for example by a PDA, so that an initial overview of the hazard situation can be obtained. Currently the Internet Analysis System uses a stand-alone client as a front-end. Through this client all counter readings can be displayed for freely selectable periods. It is possible to include further functions by means of a plug-in system.

8 Results of the Internet Analysis Systems

For the purposes of illustration some results are presented in this section in order to provide an idea of the possibilities of the current status of the Internet Analysis System. At present there are approximately 300,000 different counters of communication parameters incorporated for the various communication levels. This large number clearly shows how complex the results can be.

8.1 Transport Protocol Distribution

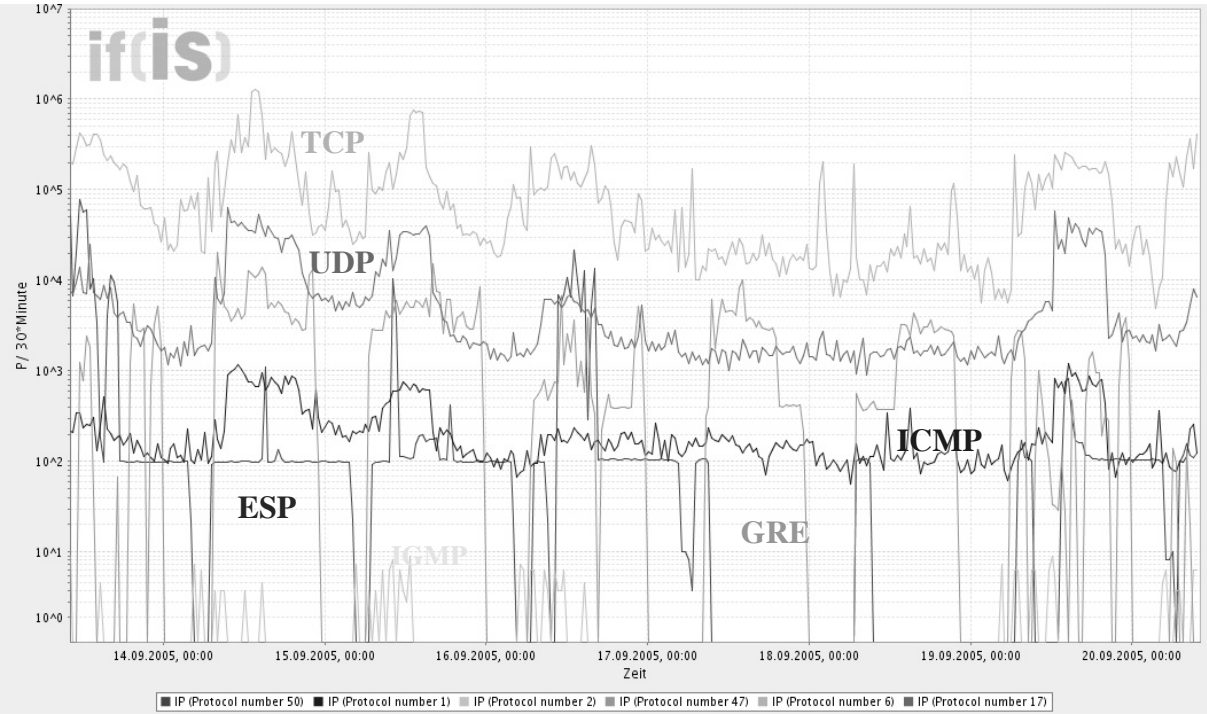

Figure 5: Protocols of the transport layer

Illustration 5 shows the distribution of the protocols of the transport layer used over a period of several days for a specific communication line. From the past the Internet Analysis System knows the profile, the standard deviation and from this can display an indication of untypical behavior. Additionally, the use of certain protocols can be determined, enabling capacity planning for the use of Virtual Private Networks (ESP protocol), for example. Protocol dependencies can also be detected: UDP appears to be proportional to TCP, which can be attributed to the dependencies of HTTP and DNS.

8.2 Browser Distribution (Technology Trend)

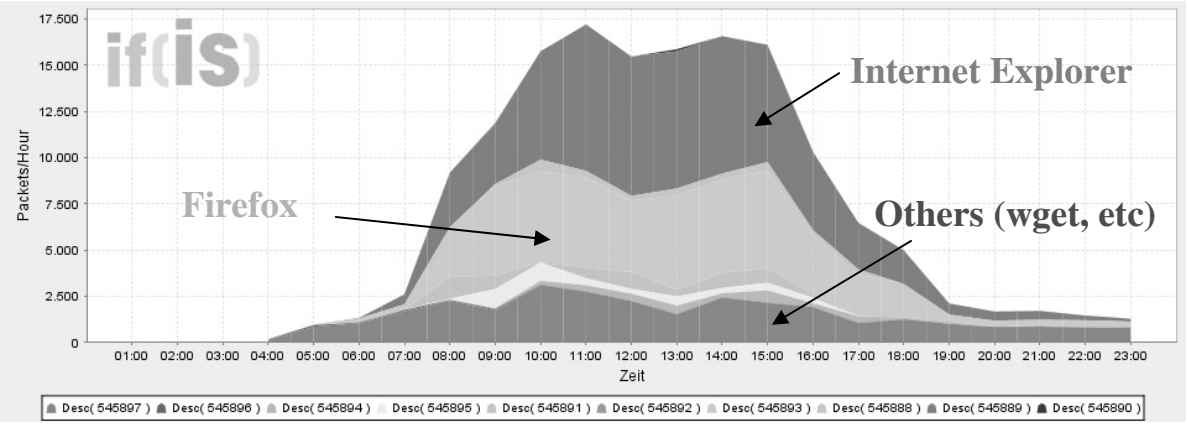

Figure 6: Browser distribution over time

Illustration 6 shows the distribution of various browsers over a period of one day for a specific communication line.

Here we can see the daily profile of the various browsers. The difference between manual use (e.g. Internet Explorer and Firefox) and automatic use (e.g. wget) over the course of the day can be clearly seen.

It is noticeable that these statistics in no way refer to a Web server, as is usual today. In contrast, these statistics refer to a communication connection.

8.3 Types of E-mail Messages

Illustration 7 shows the ability of the system to record the statistics of the headers of the e-mails sent by SMTP. The distribution can provide information on general communication behavior, as well as deviations from it. Illustration 7 shows an example of normal behavior in which the total number of messages without attachments represents 60% of all messages. These e-mails include messages with the text/plain, text/html and multipart/alternative content types. As a rule, e-mails with attachments are provided with the multipart/mixed content type. A mixed form is e-mails with the multipart/related content type. Here, for example, images are integrated directly into the text. If these e-mails are included in the e-mails with an attachment, approximately 36% of all e-mails are sent with an attachment. The remaining 4% essentially consist of confirmations of reading with the multipart/report content type. An abrupt change to these values in particular may indicate a wave of spam affecting a company from the outside, or indicate that a computer is sending spam from within the company.