Sachar Paulus
Norbert Pohlmann
Helmut Reimer

Securing Electronic Business Processes

Highlights of the Information Security Solutions Europe 2005 Conference

Contents

Preface

ENISA is proud to be working with eema, TeleTrusT, the Hungarian Ministry of Informatics and Communications and the German Federal Ministry of Technology and Economics for this year's 7th annual Information Security Solutions Europe Conference.

The aim of ISSE has always been to support the development of a European information security culture and especially a cross-border framework for trustworthy IT applications for citizens, industry and administration. ENISA is committed to these goals, in our work to assist and advise the European Commission, Member States and business community on network, information security and legislative requirements.

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 2005, 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 Buellesbach, *DaimlerChrysler (Germany)*
- Lucas Cardholm, *Ernst&Young (Sweden)*
- Roger Dean, *eema (United Kingdom)*
- Marijke De Soete, *Security4Biz (Belgium)*
- Jos Dumortier, *KU Leuven (Belgium)*
- Leonard J.N. Franken, *ABN AMRO (The Netherlands)*
- Boaz Gelbord, *ENISA (Greece)*
- John Hermans, *KPMG (The Netherlands)*
- Frank Jorissen, *Control Break International (Belgium)*
- Jeremy Hilton, *eema (United Kingdom)*
- Matt Landrock, *Cryptomathic (Denmark)*
- Manel Medina, *UPC (Spain)*
- Karel Neuwirt, *The Office for Personal Data Protection (Czech Republic)*
- Sachar Paulus, *SAP (Germany)*
- Norbert Pohlmann, *University of Applied Sciences Gelsenkirchen, Chairman of the Programme Committee (Germany)*

- Bart Preneel, *KU Leuven (Belgium)*
- Helmut Reimer, *TeleTrusT (Germany)*
- Paolo Rossini, *TELSY, Telecom Italia Group (Italy)*
- Ulrich Sandl, *BMWA (Germany)*
- Szigeti Szabolcs, *Budapest University of Technology and Economics (Hungary)*
- Wolfgang Schneider, *Fraunhofer Institute SIT (Germany)*
- Robert Temple, *BT (United Kingdom)*
- Jurgen Truyen, *L-SEC (Belgium)*

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):

For 16 years, EEMA has been Europe's leading independent, non-profit e-Business association, working with its European members, governmental bodies, standards organisations and e-Business initiatives throughout Europe to further e-Business technology and legislation.

EEMA's remit is to educate and inform around 200 Member organisations on the latest developments and technologies, at the same time enabling Members of the association to compare views and ideas. The work produced by the association with its Members (projects, papers, seminars, tutorials and re-ports etc) is funded by both membership subscriptions and revenue generated through fee-paying events. All of the information generated by EEMA and its Members is available to other members free of charge.

Examples of papers produced in recent months are:- Role Based Access Control – a User's Guide, Wireless Deployment Guidelines, Secure e-Mail within the Organisation, The impact of XML on existing Business Processes, PKI Usage within User Organisations. EEMA Members, based on a re-quirement from the rest of the Membership, contributed all of these papers. Some are the result of many months' work, and form part of a larger project on the subject.

TeleTrusT (www.teletrust.de):

TeleTrusT was founded in 1989 to promote the security of information and communication technology in an open systems environment.

The non-profit organization was constituted with the aim of:

- achieving acceptance of the digital signature as an instrument conferring legal validity on electronic transactions;
- supporting research into methods of safeguarding electronic data interchange (EDI), application of its results, and development of standards in this field;
- collaborating with institutes and organizations in other countries with the aim of harmonizing objectives and standards within the European Union.

TeleTrusT supports the incorporation of trusted services in planned or existing IT applications of public administration, organisations and industry. Special attention is being paid to secure services and their management for trustworthy electronic communication.

Microsoft and its European Initiatives

Europe plays a key role for Microsoft with regard to the dynamic development of information and communication technologies and applications. The increasing potential of the combined member states of the European Union offers a globally operating company significant opportunities for profitable cooperation. Furthermore, the attractiveness of the European education, research and economic region is supported by the political drive towards a European information society with maximum innovative ability and competitiveness.

With "i2010: European Information Society 2010" the European Commission proposed a new strategic framework with a broad political orientation. i2010 encourages the development of an open and competitive digital economy and emphasizes ICT as a driving force for inclusion and quality of life. As a key element of the renewed Lisbon partnership for growth and employment, i2010 will work towards an integrated approach in the areas of audio-visual media policies in the EU and a modern information society.

Microsoft considers these goals, also identified as priorities by their own recent research, as an opportunity for the advancement of its own strategy and for the development of European partnerships.

The highest priority is the further implementation of the Trustworthy Computing Initiative. Trustworthy, secure and dependable ICT-devices, system and application software and infrastructure services are vital requirements for a broadly applied implementation of digital information processes. Microsoft intends to actively support a strategy for the secure information society proposed in the i2010 initiative.

Digital convergence requires interoperable devices, platforms and services. Microsoft has also included this aspect in the list of strategic tasks.

> *Bill Gates: "Interoperability is more pragmatic than other approaches, such as attempting to make all systems compatible at the code level, focusing solely on adding new layers of middleware that try to make all systems look and act the same, or seeking to make different systems interchangeable. With a common understanding of basic protocols, different software can interact smoothly with little or no specific knowledge of each other. The Internet is perhaps the most obvious example of this kind of interoperability, where any piece of software can connect and exchange data as long as it adheres to the key protocols."*

Since 1999 the Information Security Solutions Europe (ISSE) conference has played an important role in implementing a European strategy for a secure information society. It is an annual meeting place for European experts with international partners. Microsoft is pleased to be able to support the publication of the most important presentations of this years ISSE in Budapest.

Technology

Federated Identity: A Progress Report

Gerry Gebel

Midvale, Utah
Burton Group
ggebel@burtongroup.com

Abstract

Over the last two years, the concept of federated identity has emerged as a pragmatic and credible solution. Burton Group defines federated identity as *the agreements, standards, and technologies that make identity and entitlements portable across autonomous domains*. The Security Assertion Markup Language (SAML), Liberty Alliance, and WS-Security are already in the early adopter phase of implementation and deployment, across multiple industries. The Shibboleth model for limited attribute sharing is gaining adoption in the higher education community in North America and Europe. Furthermore, the WS-Federation specification is supported in several commercial products and will be included in Microsoft Windows 2003 Server R2, expected in the Fall 2005.
Federated identity is the right architecture for Internet authentication and is also applicable across business units within the enterprise. Applications developed with Web services and federated identity leverage loosely coupled interfaces, service oriented architecture (SOA), and eXtensible Markup Language (XML). Each of these characteristics favors component reusability, vendor independence, platform independence, and location independence – federation in the broadest sense.

1 Basic Federation Use Cases

Fundamentally, enterprises that deploy federation technology are either producing or consuming identity assertions. Figure 1 illustrates the basic case where an enterprise is on the consuming side of the transaction and creating identity assertions on behalf of its employees that will be consumed by partners or service providers.

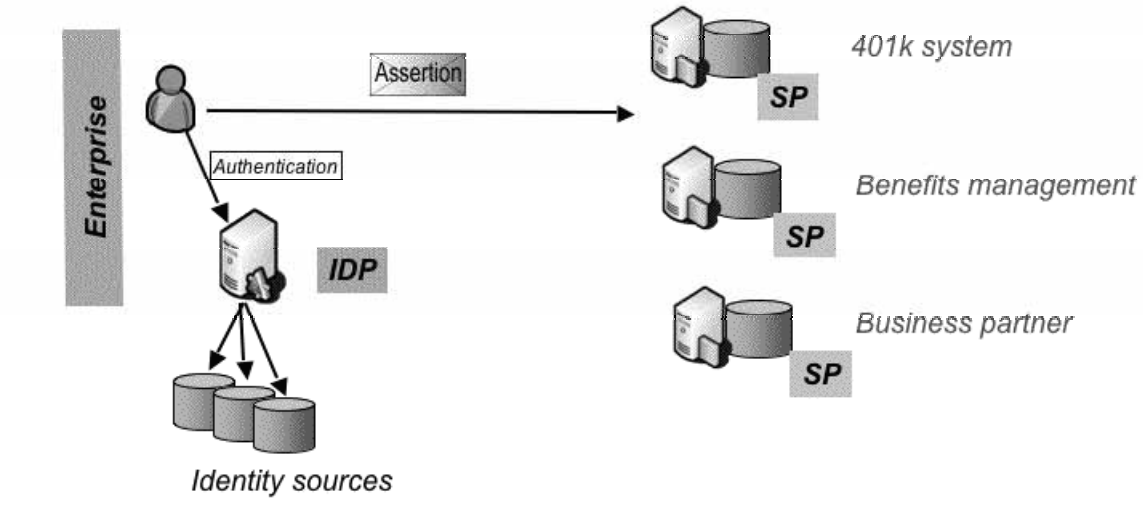

Fig. 1: Federation Producer Scenario

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2005), 3-9

The correlative use case is illustrated in Figure 2, in which an enterprise accepts identity assertions from external parties, whether they are customers or business partners. There are several business and technical drivers for implementing federation, including improved user experience, reduced administrative costs, better control of personal identity information, appropriate distribution of liability, ability to quickly meet new business opportunities, and others. To satisfy these requirements, enterprises can start with basic federation scenarios and subsequently implement more complex functions, incorporate third parties, and develop federation communities.

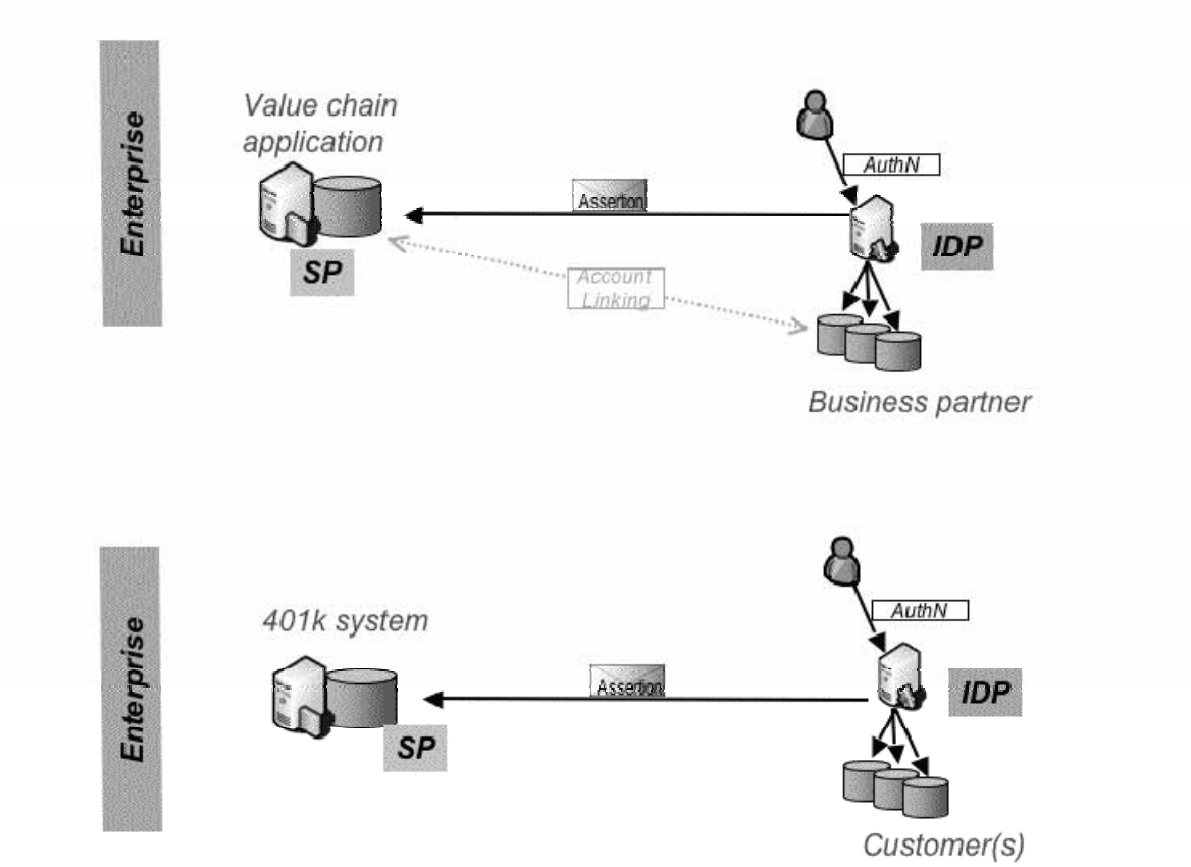

Fig. 2: Federation Consumer Scenario

Most early adopters start with basic use cases and many are planning to expand their deployments to more partners and larger user populations. As the use of federation grows, products must keep up the pace and provide the features, tools, and capacity to handle large communities.

1.1 Convergence Point for Browser-based Federation

In the area of federation standards, SAML version 2.0 offers a significant convergence point for the industry as it combines the previous work of SAML 1.x, Liberty Alliance Identity Federation Framework (ID-FF), and Shibboleth. Figure 3 shows federation protocols and indicates where there are relationships between the different standards and specification development efforts. But it will take 2-3 years for enterprises to migrate to, or implement the latest version of SAML. This transition period will present a number of technical challenges that must be addressed, along with business issues such as liability, responsibility and risk appor-

tionment. During this transition period, enterprises must also consider WS-Federation, which will be offered by several vendors in the coming months.

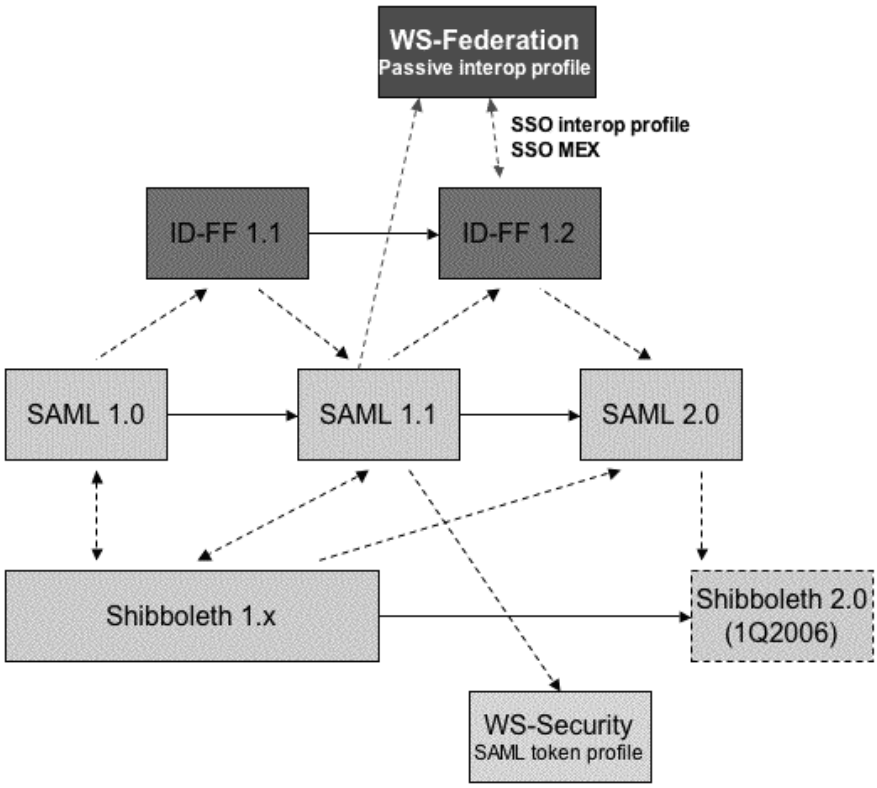

Fig. 3: Federation Family Tree

1.2 Federation Product Landscape

Many federation products support multiple versions of each protocol, as shown in Table 1. Currently, SAML enjoys the greatest vendor support, with an increasing number of vendors supporting or planning to implement the published WS-Federation Passive Requestor Interoperability Profile. All this coverage provides enterprise customers with lots of options to meet their needs for many federation scenarios and enables products to co-exist with partners that may be at different levels. But enterprises should seek to converge on the latest protocol versions to reduce operational complexity. Burton Group expects that SAML 2.0, WS-Security, and WS-Trust will be the key federation standards that customers and vendors rally around.

Table 1: Federation Protocol and Version Support

Vendor	SAML 1.0	SAML 1.1	SAML 2.0 *	ID-FF 1.1	ID-FF 1.2	Shibboleth	WS-Federation + SAML token
BMC	X	X	X	X	X	X	X
CA	X	X	X	P	P		F
Entrust	X	X	X				
Evidian		X					
HP	X	X	X	X	X		
IBM	X	X		X	X		X
Internet2	X	X	X			X	F
Microsoft							X
Novell	X	X	X	X	X		
Oracle **	X	X	X	X	X		X
Ping Id **		X	X				X
RSA	X	X	P				F
Sun	X	X	X	X	X		X
Symlabs			X	X	X	P	F
Trustgenix	X	X	X	X	X		F

Key: *=partial or under development; X=supported; F=future; P=professional services; ** toolkits support additional protocols

The vendors in this market can be segmented according to their strategy for packaging federation capabilities. Three basic approaches have emerged: federation embedded in WAM products, standalone federation products, or a dual strategy that offers both product types. Figure 4 illustrates where each vendor is positioned at the current time, but product strategies could change in the future and customer requirements cause adjustments.

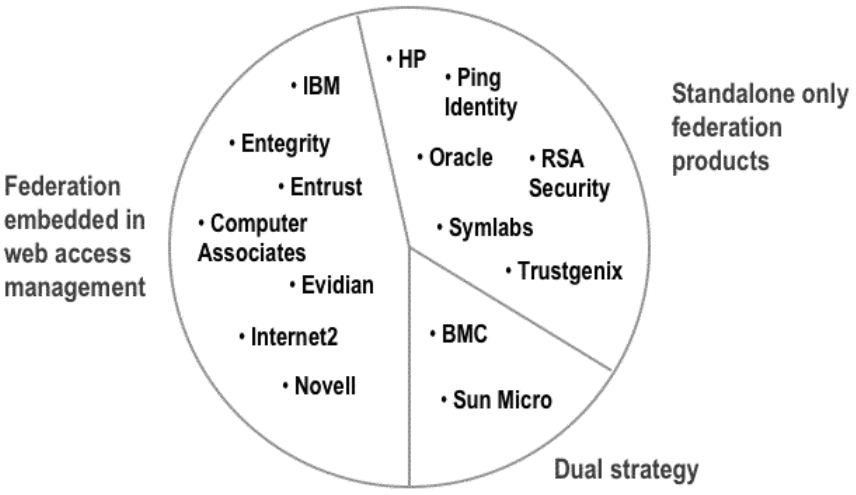

Fig. 4: Federation Vendor Product Strategies

Some vendors believe that federation capability should be integrated directly with the WAM platform to ensure a secure operation, less complexity, and easier administration. Limited resources can also influence a vendor to add federation to an existing WAM product instead of

creating a new, standalone system. Oracle and RSA Security decided to separate the federation functions from the WAM product, which created a product focused entirely on federation, freed it from the release schedule of other products, and permitted integration with competing WAM products. HP also chose this route after reaching an OEM agreement with Trustgenix and will no longer add federation improvements to its Select Access WAM product. Ping Identity, Symlabs, and Trustgenix are boutique vendors that didn't have WAM offerings, but entered the market as federation specialists intent on moving faster than their larger competitors.

BMC and Sun are the only two at this time with a dual product strategy. Such an approach permits these vendors to supply the needs of the large federation hub operator or the smaller spoke company seeking to join a federation community.

Microsoft has announced plans to release Active Directory Federation Services (ADFS) as part of the Windows Server 2003 R2 release in the Fall 2005. ADFS, available in beta today, supports the WS-Federation Passive Requestor Interoperability Profile, which is also supported by at least six other vendors. With the release of ADFS, Microsoft brings federation directly to the many Windows Server customers and presents an opportunity to raise the awareness and adoption of federation at a faster pace.

1.2.1 Emergence of Standalone Products

The federation market began as heavyweight WAM systems with embedded federation functionality. This configuration is appropriate for the large enterprise, service provider, federation community operator; but it's usually more than a small partner can handle. For example, it's unreasonable to expect a 100-500 person company to install a complete WAM system just to leverage the federation functionality. Early proprietary deployments installed specialized modules, such as CA (then Netegrity) affiliate agents. Several pioneers working with initial federation standards developed their own agents were installed at partner sites and enabled SSO. Recognizing this market opportunity, vendors responded by offering several standalone or spoke federation products.

One objective of these lighter weight solutions is to facilitate rapid and simplified installation and operation for smaller or less IT savvy partners. The other goal is to provide a full function, scalable federation server that supplements existing WAM and other identity infrastructure.

1.2.2 Growing Federation Ecosystem

The WAM vendor's reign as the exclusive purveyor of federation technology was a short lived phenomenon. First boutique vendors such as Trustgenix and Ping Identity entered the federation market to provide specialized attention on this relatively new technology. However, over the last 12-18 months a broad assortment of product types has introduced functionality to support federation, resulting in a growing federation ecosystem. Other types of products supporting federation include:

- **Toolkits:** Open source toolkits from Ping Identity and Internet2, plus a commercial offering from Oracle allow developers to integrate federation functions into applications or build federation systems.
- **SSL VPN:** Juniper Networks and PortWise support SAML assertion creation to extend SSO for users to WAM protected resources or other identity infrastructure that is federation enabled.

Applications

Privacy Policy Enforcement in Enterprises: Addressing Regulatory Compliance and Governance Needs

Marco Casassa Mont · Robert Thyne
Pete Bramhall · Kwok-Nga Chan

Hewlett-Packard Laboratories, Trusted System Lab
Filton Road, Stoke Gifford, Bristol, UK
{marco.casassa-mont | robert.thyne | pete.bramhall}@hp.com

Abstract

This paper describes issues and requirements related to privacy management as an aspect of improved governance in enterprises. It focuses on the privacy enforcement aspect, in particular related to privacy-aware access control and enforcement of privacy obligations: this is still a green field and, at the same time, is a key aspect to be taken into account to ensure compliance both to regulations and an enterprise's IT governance objectives. We introduce our HP Labs work in these areas: core concepts are described along with our policy enforcement models and related technologies. Two prototypes have been built as a proof of concept to: (1) enforce privacy policies on personal data by extending HP Select Access; (2) manage and enforce privacy obligations on personal data, integrated with HP Select Identity. We describe their technical capabilities and our next steps.

1 Introduction

Privacy management is important for enterprises that handle identities and personal data of customers, employees and business partners: it has implications on their compliance with regulations, their reputation and brand [CaTB05, Casa04a]. Enterprises have been heavily investing in identity management solutions for the last few years and want to leverage them also for privacy management, including: authoring, managing and enforcing privacy policies when provisioning and handling identity information and personal data; auditing and monitoring these policies for compliance. We focus on the specific problems of enforcing privacy policies and privacy obligations on personal data within enterprises: these areas are still a green field. Section 2 and 3 describe core privacy management concepts, addressed problems and core requirements. Section 4 describes related work. Section 5 introduces our work i.e. our privacy policy enforcement models and technologies. As a demonstration of the feasibility of our work we describe how we leveraged and extended two HP OpenView Identity Management solutions - HP Select Access and HP Select Identity - to respectively enforce privacy policies and privacy obligations. Current results and next steps are illustrated in section 6.

2 Privacy Management

Dealing with privacy is an important aspect of enterprises' regulatory compliance efforts and it is required by law [Laur03,Onli04,Oecd80]. Large enterprises that are geographically distributed across different nations might need to comply with different privacy laws.

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2005), 137-148

Privacy policies can be used to represent and describe privacy laws, guidelines and privacy statements. They express rights, permissions and obligations, usually in natural language that needs to be interpreted and understood by people. They need to be enforced and audited.

Most of the technical work currently done in this space focuses on the provision of auditing and reporting solutions to analyse logged events and check them against privacy policies. The enforcement of privacy policies is very important for regulatory compliance: often privacy policies are hardcoded into applications and services or managed with very vertical, ad-hoc solutions, in specific contexts. This approach is not adaptive to changes and does not scale.

The enforcement of privacy rights, permissions and obligations on confidential and personal data requires the mapping of these concepts into rules, constraints and access control, the meaning of which must be unambiguous so that it can be deployed and enforced by software solutions. This still requires following best practices and good behaviours: however, automating aspects of the enforcement of privacy policies and reducing the involved costs is important for enterprises. The (technological) enforcement of privacy permissions and rights requires extended access control and authorization mechanisms on stored personal data that check these privacy permissions against data requestors' rights and intents, data subjects' consent and the stated data purposes [CaTB05]. This applies, for example, to enterprise web services or applications that need to access and manipulate personal data for business reasons.

Even more complex is the case of dealing with the enforcement of privacy obligations. Privacy obligations dictate criteria for a privacy-aware lifecycle management of data. They might require the deletion or transformation of confidential data after a predefined (potentially very long) period of time, periodic notifications and requests for authorization to data owners or data subjects, fulfilment of opt-in/opt-out choices made by data owners, ongoing compliance with laws' obligations and internal guidelines. The events that trigger the fulfilment of privacy obligations can be completely orthogonal to the ones relevant to privacy permissions. Privacy obligations can have ongoing aspects that need to be monitored and satisfied over a long period of time. All these tasks are challenging for enterprises because of the need for specific IT infrastructures and processes able to manipulate confidential data as dictated by privacy obligations.

3 Addressed Problems, Issues and Requirements

This paper focuses on two core enterprise privacy problems: (1) Privacy policy enforcement on personal data; (2) Privacy obligation management and enforcement.

We address these aspects by analysing and developing a privacy enforcement framework that can be deployed within current enterprise identity management solutions to leverage current enterprises' investments in this area. In this context, we want to enable privacy management scenarios where data subjects can specify their privacy preferences, give explicit consent, limit the usage of their data and get degrees of control on their personal data. Enterprises must be able to explicitly author, deploy and enforce privacy policies and obligations during the access, manipulation and transmission of personal data. They need tools and solutions to achieve this.

3.1 Privacy Policy Enforcement on Personal Data

The enforcement of core privacy principles [Laur03,Onli04,Oecd80] on personal data has implications in terms of access control: enterprises must state the purposes for collecting data and data must be accessed only for that reasons. The consent given by data subjects impose

limitations on how these data are accessed. Similarly, the limitations on data usage, disclosure and retention dictate conditions and constraints that need to be satisfied before accessing personal data.

Traditional access control systems are necessary but not sufficient to enforce privacy policies on personal data. They are mainly based on "access control lists" and enforcement mechanisms that keep into account the identities of data requestors, their rights and permissions and the types of actions that are allowed/disallowed on the involved resources (data resources). These systems do not keep into account additional aspects relevant to privacy enforcement: the stated data purposes and data subjects' consent - i.e. properties usually associated to collected data - the intent of data requestors and any additional enterprise or customized data subjects' constraints.

It is necessary to build "privacy extensions" of traditional access control systems that can author and enforce privacy policies. To address the above issues and move towards privacy-aware access control systems, it is important to satisfy the following core requirements: (1) explicit modeling of personal data stored by enterprises; (2) explicit definition, authoring and lifecycle management of related privacy policies; (3) explicit deployment and enforcement of privacy policies; (4) integration with traditional access control and identity management systems; (5) simplicity of usage of all the involved system; (6) support for auditing. A more comprehensive analysis and discussion of these aspects can be found in [CaTB05].

3.2 Privacy Obligation Management and Enforcement

Privacy obligations on personal data can be defined by data subjects, by laws and by enterprises. Privacy obligations dictate responsibilities on how data must be handled and processed, given specific contexts, for example with respect to disclosure of personal information. Obligations can be expressed in terms of notice requirements, opt-out options, limits on reuse of information and information sharing for marketing purposes. Privacy obligations can dictate very specific requirements. For example, privacy obligations can require that personal data must be deleted after many years, e.g. 30 years, (long-term commitment) or in a few days if user's consent is not granted (short-term commitment) or their account is closed. Privacy obligations can have "ongoing" and long-term commitments for enterprises or might apply only for a short period of time and be transient. The enforcement of privacy obligations can be independent from access control (e.g. the deletion of personal data after 7 years has to happen independently from data accesses). It is important that privacy obligation management solutions address the following core requirements: (1) explicit modeling and representation of privacy obligations; (2) association of obligations to data; (3) timely enforcement of privacy obligations; (4) mapping obligations into enforceable actions; (5) compliance of refined obligations to high-level policies; (6) tracking the evolutions of obligation policies; (7) dealing with long-term obligation aspects; (8) accountability management; (9) monitoring of enforced obligations for compliance; (10) user involvement; (11) complexity and cost of instrumenting applications and services. A comprehensive analysis and discussion of these aspects can be found in [Casa04a, Casa04b].

4 Related Work

A common approach to enforce privacy policies on personal data consists of hardcoding them within applications and services or building ad hoc solutions. This approach is suitable for very simple and static environments: it shows all its limitations and maintenance costs in case of complex and dynamic organizations that need to adapt to changes. As described in the re-

quirements section, to explicitly address the problem, a model of the relevant personal data is required. Privacy policies dictating how these data must be accessed need to be authored, deployed, enforced and audited. This requires the definition of a comprehensive privacy-aware access control model and systems that implement it.

Relevant work in this direction, for privacy management and enforcement in enterprises is described in [KaSh02, KaSW02a, KaSW02b]. An Enterprise Privacy Architecture is introduced and described in [KaSW02b]. This approach is further refined and described in the Enterprise Privacy Authorization Language (EPAL) specification [Ibm04a]. However, these papers mainly provide general guidelines.

Important related work on actual privacy enforcement on personal data has been done by IBM with their research on Hippocratic databases [AKSX02]. The drawback of this approach is that it mainly focuses at the database level, specifically on RDBMS data repository architectures and related data schemas whilst the enforcement of privacy policies might need to deal with LDAP directories, meta and virtual directories, file systems and legacy systems. It might need to incorporate higher-level views and perspectives than just the database-level perspective.

In terms of commercially available solutions, IBM Tivoli Privacy Manager [Ibm05] provides mechanisms for defining fine-grained privacy policies and associating them to data. This approach dictates strong constraints on how applications need to be developed, the required operational and software environment and how personal data has to be stored and administered: it might require some duplications of administrative and enforcement frameworks.

Other products, such as HP Select Federation [Hp05a], focus on single-sign-on and address related privacy aspects: they enforce privacy rules on personal data in federated environment when these data are disclosed by an organization (or an identity provider) to other parties.

Our work specifically addresses the problem of enforcing privacy policies on personal data stored in a broad variety of data repositories and used *within* enterprises. Our work aims at not being invasive for applications and services: privacy policies are managed in an explicit way, in conjunction with traditional access control policies and not hardcoded in applications and services. We want to avoid duplications of efforts by providing a single, integrated framework for authoring, administering and enforcing both traditional access control and privacy policies. To demonstrate the feasibility of this, as a significant example, we leveraged and extended HP Select Access [Hp05b] to enforce privacy policies on personal data.

In terms of managing and enforcing privacy obligations, relevant work is described in [KaSh02, KaSW02a, KaSW02b, Ibm04a]: in particular, the EPAL approach to privacy obligations is driven by an access control perspective. However, privacy obligations cannot be managed at their best only from an authorization-based perspective as they can include aspects that are not driven by data accesses, for example the deletion of data at a predefined time.

We believe that modularity and separation of concerns are important aspects. In our approach obligation policies are first-class citizens with their explicit management, as a self-standing component of a more comprehensive policy management framework. Our architecture has high-level commonalities with the architecture described in [KaSh02, KaSW02a, KaSW02b] but in our work we further refine the concept of obligations and their enforcement. We split the enforcement mechanisms in two parts by including a scheduling mechanisms and an enforcement mechanism allowing for workflow automation and human intervention.

Approaches to deal with (privacy) obligations have already been implemented in products, in particular for data retention [Ibm04b] and in a variety of document management systems. Nevertheless, these approaches are very specific, focused on particular domains and handle simple obligation policies on files and documents, not really on personal data. Our work aims at pushing the barrier even further to create an obligation management framework that can be leveraged in multiple contexts, for different purposes.

Relevant work on mechanisms to associate policies to data is described in [KaSh02, KaSW02a, KaSW02b, CaPB03]. We can leverage aspects of this work, in particular [CaPB03] to provide a stronger association of obligation policies to confidential data.

5 Our Work

This section provides technical details of HP Labs work to enforce privacy policies and privacy obligation on personal data stored within enterprises. To demonstrate the feasibility and deployability of our work in real world identity management solutions, as a significant example we deployed our prototypes within HP Identity Management solutions: we extended HP Select Access to deal with privacy policy enforcement on personal data, integrated with the enforcement of "traditional" access control policies. We have also implemented a prototype of an obligation management system, integrated with HP Select Identity, to represent, schedule, enforce and monitor privacy obligations.

5.1 Privacy Policy Enforcement and Integration with HP Select Access

Our approach to enforce privacy policies is based on a privacy-aware access control model that extends traditional access control models (based on users/groups, users' credentials and rights, access control lists and related policies) by explicitly dealing with the stated purposes for which data is collected, checking - at the access request time - the intent of requestors against these purposes, dealing with data subjects' consent and enforcing additional access conditions and constraints on personal data defined by data subjects and/or enterprise administrators [Laur03,Onli04,Oecd80] – see Figure 1. The main aspects of this model are:

- A mechanism for the explicit modelling of personal data that are subject to privacy policies, including the type of the data repository (database, LDAP directory, etc.), its location, the schema of these data, types of attributes, etc.;
- An integrated mechanism for authoring privacy policies along with traditional access control policies: it is a Policy Authoring Point (PAP) to allow privacy administrators to describe and author privacy policy constraints and conditions along with more traditional access control policies based on security criteria (such as who can access which resource, given their rights and permissions);
- An integrated authorization framework - Policy Decision Point (PDP) for deploying both access control and privacy-based policies and making related access decisions;
- A "Data Enforcer" - Policy Enforcement Point (PEP) - for intercepting run-time attempts to access personal data and enforcing decisions based on privacy policies and contextual information. This mechanism is in charge (among other things) of dealing with the transformation of queries to access personal data (e.g. SQL queries) and filtering part of the requested data, if their access is not authorised for privacy reasons.

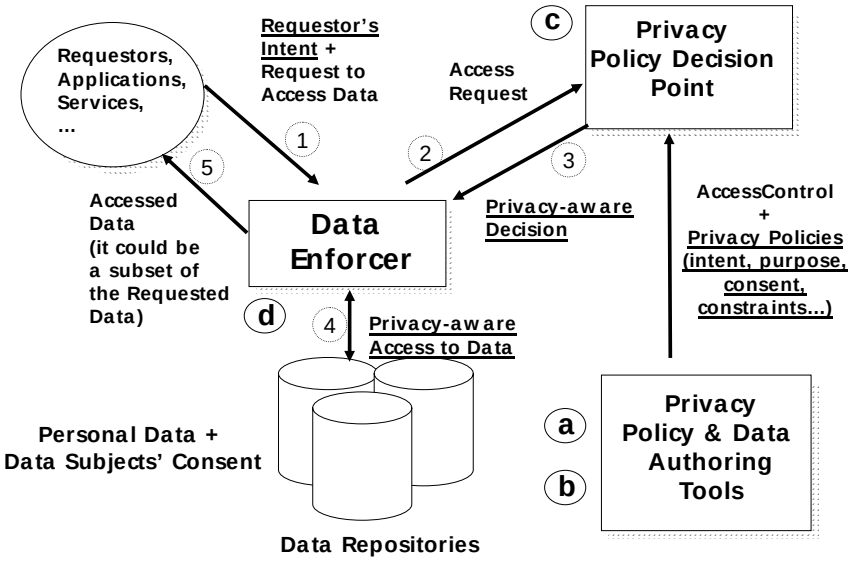

Fig. 1: Model of our Privacy-aware Access Control System

The data enforcer plays a key role to enforce privacy policies on personal data. At "run-time", attempts to access personal data are intercepted and managed in the following way - Figure 1:

1. A request from a data requestor to access personal data is intercepted by the data enforcer;
2. The data enforcer interacts with the privacy policy decision point by passing information about the request (including the requestor's intent) and the requestor's identity;
3. The privacy policy decision point makes a decision, based on available privacy policies and the context, and sends it to the data enforcer. It can be any of the following types:
   - No: access to data is denied;
   - No & conditions: conditions usually require stronger authentication;
   - Yes: access to data is granted;
   - Yes & conditions: access to (part of the) data is allowed, under the satisfaction of the attached conditions (i.e. data filtering, transformations and manipulations).
4. The data enforcer enforces this decision. In particular, if the decision is "Yes & conditions" the data enforcer might have to manipulate the query (query pre-processing) and/or transform the requested personal data (result prost-processing), before returning the result to the data requestor;
5. Data (or alternatively no data) is returned to the data requestor, based on the enforced decision.

To implement the above model we leveraged and extended HP Select Access. HP Select Access [Hp05b] is a leading-edge access control product. It provides policy authoring, policy decision and policy enforcement capabilities respectively via the following components: Policy Builder; Validator; Web Enforcer plug-in.

The current commercial version of HP Select Access does not handle data as managed resources: it only deals with traditional access control policies on web resources. New functionalities have been added to HP Select Access in our prototype, to explicitly deal with privacy-

aware access control on personal data, as shown in Figure 2. The following extensions of HP Select Access have been implemented in our prototype:

- **Estension of the HP SA Policy Builder to represent data resources** (databases, LDAP directories, virtual-directories, their schemas, etc.) in addition to traditional IT resources (such as web resources);
- **Extension of the HP SA Policy Builder to graphically author privacy policies on data resources**: a set of additional plug-ins has been implemented, including the ones that check (at the enforcement time) the requestor's intent against the stated data storage purposes, take into account data subjects' consent & data retention policies and describe how the accessed personal data must be filtered, obfuscated or manipulated, etc.;
- **Extension of the HP SA Validator to make privacy-aware decisions**. Plug-ins, correspondent to the ones used in the Policy Builder, have been implemented. This enhanced-version of the Validator can now make "Yes & constraints" decisions as described in our model;
- **A Data Enforcer has been built and added to the framework**: this is a new functionality added to HP Select Access. It enforces privacy decisions made by the Validator. It intercepts incoming calls to data resources, interacts with the Validator, performs fine grained manipulation of data resources and deals with the interpretation and enforcement of additional constraints as defined by the privacy policies. The data enforcer sits nearby managed data repositories: we envisage that a family of data enforcers (sharing a common logic but differentiated by add-ons dealing with different types of data resources) need to be built, because of the different semantic of different data repositories. The data enforcer currently implemented is a JDBC proxy for RDBMS databases.

The above functionalities address and satisfy the core requirements described in section 3 for privacy enforcement on personal data. More details can be found in [CaTB05].

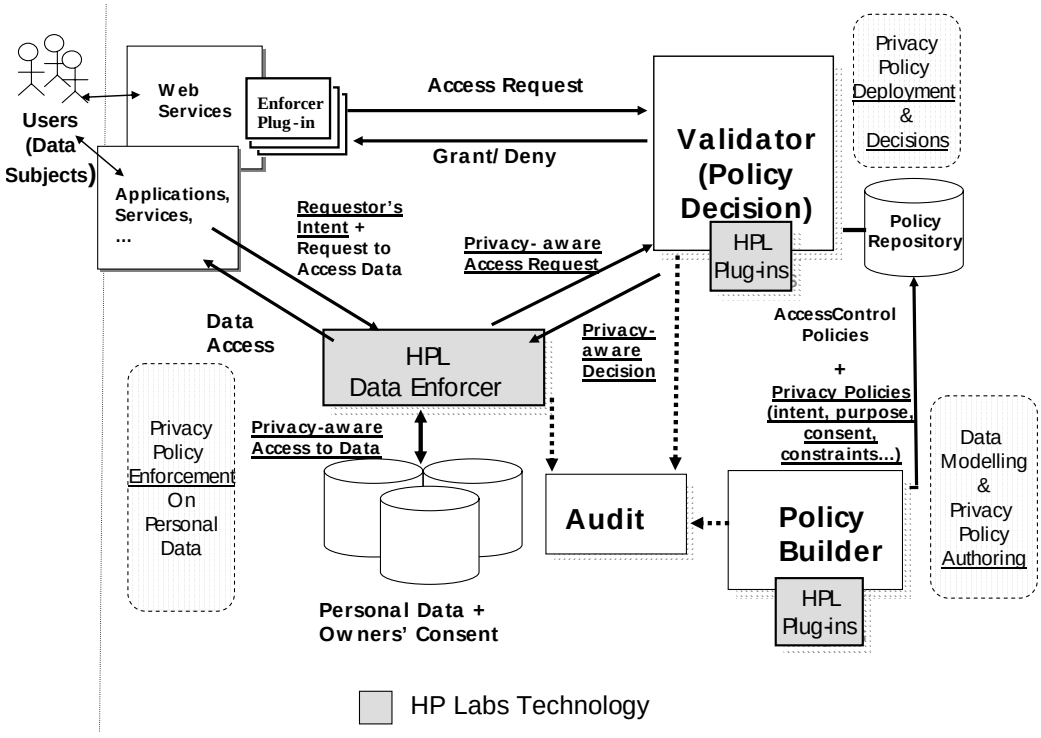

Fig. 2: Extended HP Select Access to deal with Privacy Policy Enforcement

5.2 Privacy Obligation Management and Integration with HP Select Identity

Our technical work focuses on the explicit management and enforcement of privacy obligations for personal data stored by enterprises. In our model, privacy obligations are "first class" entities, i.e. they are explicit entities that are modeled and managed to provide a privacy-aware lifecycle management of personal data: this includes data deletion, data transformation, dealing with notifications, etc. A related obligation management framework is introduced to manage these privacy obligations.

From a technical perspective, a privacy obligation is an "object" (currently expressed in XML) that includes (at least) the following aspects:

- **Obligation Identifier:** unique identifier of the privacy obligation;
- **Targeted Personal Data**: specifies links to the managed data;
- **Triggering Events**: one or more relevant events that can trigger the fulfillment of the obligation;
- **Actions**: set of actions that need to be executed when the obligations has to be enforced;

Different categories of privacy obligation need to be managed and enforced:

- **Transactional obligations**: privacy obligations to be immediately enforced, when transactions and interactions involve personal and confidential data. For example, they might require to notify the data subject or create audit logs every time personal data is accessed;
- **Data retention and handling obligations**: these privacy obligations describe criteria for the management and deletion of personal data, usually driven by time-based events. For example, they might require the deletion of data after a predefined period of time (ranging from days to years) or at a specific time agreed with the data subject;
- **Other types of event-driven obligations**: these privacy obligations are triggered by events that relate to contextual and application-relevant information, based on usage of personal data, trust information about the systems dealing with personal data, etc.

A complementary classification of our managed privacy obligations is based on their activation timeframe and period of validity:

- **Short-term obligations**: these obligations have a short period of validity;
- **Long-term obligations**: these obligations might have long term implications in terms of resources needed for their fulfillment (months or years);
- **Ongoing obligations**: these obligations might be short or long termed and imply an ongoing, periodic, fulfillment of activities related to the management of personal data.

Figure 3 shows the conceptual model underpinning our Obligation Management Framework.

Data subjects can explicitly define privacy obligations on their personal data at the disclosure time (e.g. during a self-registration process) or at any subsequent time. Enterprise privacy administrators can further associate other privacy obligations, for example dictated by laws or internal guidelines. Our obligation management framework handles these obligations by providing the following core functionalities:

- **Scheduling the enforcement of privacy obligations**: the system schedules which obligations need to be fulfilled and under which circumstances (events);

- **Enforcing privacy obligations**: the system enforce privacy obligations once they are triggered. The enforcement ranges from the execution of simple actions to complex workflow involving human interventions;
- **Monitoring the fulfilment of privacy obligations**: the system monitors and audits the enforced obligations, at least for a predefined period of time, to ensure that the desired status of data is not changed and to report anomalies.

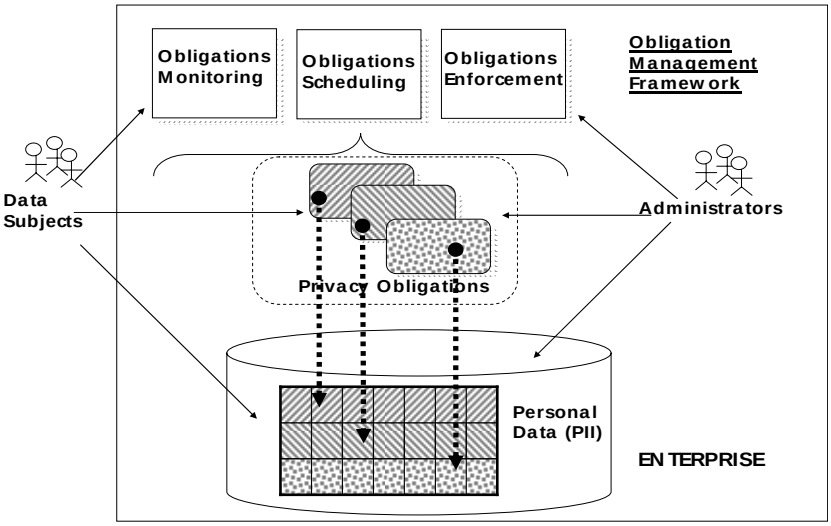

Fig. 3: High-level model of our Obligation Management Framework

More details can be found in [Casa04a, Casa04b]. These functionalities can be accessed by enterprise privacy administrators and potentially also by data subjects, for example to monitor their personal data and check for privacy compliance. Figure 4 shows the high-level architecture of our obligation management system, based on the model shown in Figure 3. Our obligation management system consists of the following modules:

- **Obligation Server**: it deals with the authoring, management and storage of obligations. It explicitly manages the association of privacy obligations to confidential data. It pushes active obligations (i.e. obligations to be fulfilled) to the "obligation scheduler";
- **Obligation Store and Versioning**: it stores obligations and their mapping to confidential data. Multiple versions of obligations are also stored in this system;
- **Obligation Scheduler**: it is the component that knows which obligations are active, ongoing obligation deadlines, relevant events and their association to obligations. When events/conditions trigger the fulfilment of one or more obligations, this component pushes them to the "obligation enforcer";
- **Obligation Enforcer**: it is a workflow system containing workflow processes describing how to enforce one or more obligations. The enforcement can be automatic and/or could require human intervention, depending on the nature of the obligation;
- **Events Handler**: it is the component in charge of monitoring and detecting relevant events for privacy obligations and sending them to the obligation scheduler;
- **Obligation Monitoring Service**: it is orthogonal to the scheduling and enforcement components and monitors enforced obligations and the expected status of data;
- **Information tracker**: it is a component that focuses on intercepting events generated by various system components and providing this information to the event handler.;

- **Audit Server**.

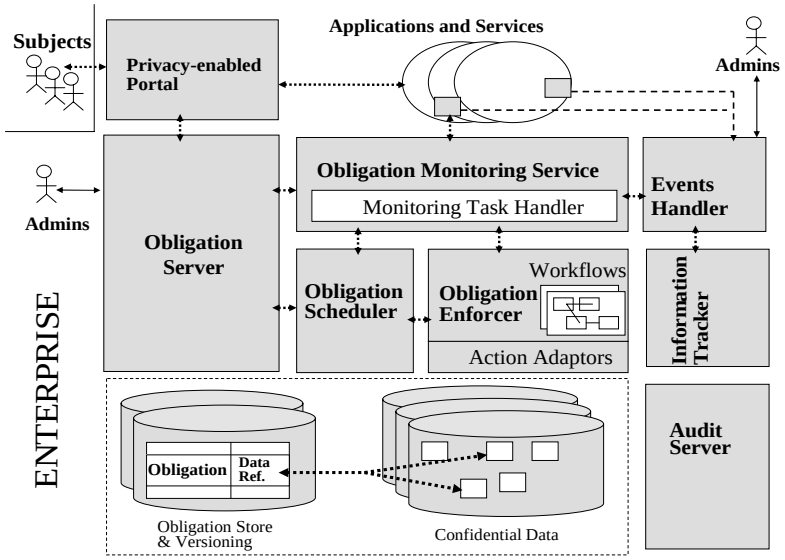

Fig. 4: High-level Architecture of our Obligation Management System

A working prototype has been implemented in the context of the EU PRIME project [Prim04], as a proof of concept, providing the core functionalities: scheduling, enforcement and monitoring of privacy obligations. At the moment the managed obligations are restricted to handling time-based and access based events. The supported actions include deletion of data and notifications. Short-term, long-term and ongoing obligations are supported. Our work addresses the core issues and requirements described in section 3. More details can be found in [Casa04a, Casa04b].

We believe that an obligation management system should be considered as an additional component of current enterprises' identity management solutions. These solutions already provide identity management functionalities for identity federation management, user provisioning and account management, access control and privacy management that can be leveraged. Our obligation management system can be integrated with the self-registration, customization and account management capabilities of identity provisioning systems to allow users and administrators to describe and handle privacy preferences and turn them into privacy obligations for the enterprise. In this context our system allows for the explicitly representation and management of privacy obligations, along with the coordination of their overall enforcement and monitoring.

To demonstrate how this can be achieved for real, we integrated our Obligation Management System with HP Select Identity [Hp05c]. HP Select Identity is a state-of-the-art solution to manage digital identities within and between large enterprises. The Select Identity solution automates the process of provisioning, managing and terminating user accounts and access privileges across platforms, applications, and corporate boundaries. Specifically, the key features of the Select Identity system include: Centralized User Management; User Provisioning; Administrative Delegation; User Self Service Registration; Password & Profile Management; Audit and Reporting.

In our integrated prototype we use HP Select Identity self-registration and user provisioning capabilities to specify and capture privacy constraints and preferences on how to handle personal data. These preferences are then passed to (via a connector) and processed by our obligation management system that transforms them into privacy obligations. Privacy obligations

are then scheduled, enforced and monitored by our system. We leverage the workflow and user/identity management capabilities of HP Select Identity to enforce aspects of privacy obligations. Our system retains control of the supervision of obligations and their monitoring. HP Select Identity is leveraged to enforce obligations constraints, such as deletion of identities, data transformation, etc. At the moment the deletion of personal data (as the effect of enforcing obligations) is achieved by triggering HP Select Identity workflows (via its Web Services API), whilst the obligation management system handles the notifications to users.

6 Discussion and Next Steps

Our prototypes are proof of concepts. They show the feasibility of our work in addressing core issues and requirements in real contexts and are ready for potential future commercial exploitation. A working demonstrator [CaTB05], based on a healthcare scenario and using both prototypes, has been implemented to show the integration of privacy policy enforcement and obligation enforcement in an identity management context.

At the moment the enforcement of privacy policies in HP Select Access mainly enforces data subjects' consent, constraints on data purposes and data expirations via data filtering. This has been achieved by intercepting and transforming incoming SQL queries by our data enforcer (query pre-processing). Current performance tests and analysis are promising. No noticeable loss of performance has been registered so far, on common SQL queries. More tests and experiments are in progress on different varieties of SQL queries. We are also planning to: (1) explore the implications of post-processing queries (post-processing of query results) to extend the current set of managed privacy constraints; (2) explore the enforcement of privacy policies on LDAP repositories and virtual directories.

In terms of privacy obligation enforcement, we are currently refining the integration of our obligation management system with HP Select Identity, specifically to leverage as much as possible the provisioning and workflow capabilities of HP Select Identity for obligation enforcement. Additional work and research in the space of privacy obligations is going to be done in the context of the EU PRIME project [Prim04]: in particular we plan to work in the area of stickiness of privacy obligations to personal data, management of complex obligation actions, end-to-end graphical management of privacy obligations, compliance feedback and longevity/survivability of the obligation management system.

7 Conclusions

Privacy management is important for enterprises to ensure their compliance to regulation and governance objectives and address customers' needs and rights. This paper focuses on privacy policy and obligations enforcement for personal data stored and accessed by enterprises: these aspects are still a green field. We discussed a privacy-aware access control model to enforce privacy constraints (including handling the purpose of data, checking data requestors' intent against data purposes and enforcement of data subjects' consent). We also analysed aspects and concepts related to privacy obligations, considered as "first-class" entities (including data deletion, data transformation, notifications, etc.) and introduced our obligation management framework to schedule, enforce and monitor them. Working prototypes have been implemented and integrated with state-of-the art HP identity mangement solutions. HP Select Access and HP Select Identity, to show their deployability in real world identity management solutions. These technologies are ready for commercial exploitation. Research and development work continues to refine our technolgies and implement adiditional functionalities.

Security Management

Using ISO 17799, COBIT & ITIL for solving Compliance Issue

Yves Le Roux

Computer Associates Int.
25 Quai Paul Doumer
F-92408 COURBEVOIE CEDEX
Yves.leroux@ca.com

Abstract

Businesses of many different shapes and sizes have compliance projects to manage, whether in conformance with specific vertical regulatory issues and horizontal legislation.
Sarbanes-Oxley Act, the French "loi sur la sécurité Financière", the UK Turnbull Report, National laws implementing the EU Data Protection Directive 95/46 are typical examples of horizontal legislation/regulations.
Basel II, Healthcare Privacy laws, Products Tracking are typical examples of vertical regulatory issues.
Furthermore, to be successful, organizations must implement effective compliance and ethics programs that go beyond mandating compliance with minimum legal requirements.
Management must disclose any material weakness and is unable to conclude that the company's internal control over financial reporting is effective if there are one or more material weaknesses in such control. Furthermore, the framework on which management's evaluation is based must be a suitable, recognised control framework that is established by a body or group that has followed due process procedures, including the broad distribution of the framework for public comment.
Addressing the Compliance issue, Computer Associates Int. (CA) has decided to work in two different directions

1. To study the roles of Standards in Compliance. CA selected two standards quoted by the European Commission COBIT and ISO Standard 17799 plus the IT Infrastructure Library (ITIL)
2. To develop a Compliance Management Framework

In this paper, we will present the results of those two approaches

1 Introduction

While governmental regulations cover a wide range of target areas, regulations that impact IT generally fall into one of three major categories:

- **Governance.** These regulations deal with issues related to the transparency and accuracy of financial records, the retention of records within the corporation, and requirements of disaster recovery and business continuity. Most notably with the Sarbanes-Oxley Act, this type of regulation was heavily driven by corporate scandals and financial fraud cases. In short, these regulations are intended to ensure that proper controls exist to guarantee that corporate reporting is accurate, timely, and complete.
- **Privacy.** These regulations are often specific to a single vertical market, and dictate how a user's personal information must be handled by the corporation. There are regulations that specify what type of personal information may be kept, how that informa-

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2005), 313-323

tion may be handled (including who, if anyone, it may be given to), and what actions are required in the event of a breach of established privacy restrictions.

- **Security.** The role of security regulations is to protect a corporation's critical infrastructure. These regulations specify how users will be identified, how their access to sensitive resources must be controlled, and how that access may be tracked and audited.

While there are a large number and wide variety of regulations, each has unique requirements for compliance, many of which cannot be solved merely through technology and/or procedural changes.

However, one element common to all regulations is the need for strong and effective controls over various enterprise business processes. A control is set of procedures or steps that can be used to ensure the successful operation of a business practice or transaction. Internal controls can be weak, strong, or anywhere in between. It is the job of compliance auditors to ensure and attest that these controls are effective enough to meet the requirements of the regulation.

To be considered conformant, the following internal control elements must be in place:

- **Accountability.** It must be clear which person performed a given operation, when it took place, the results of that operation, and whether they were authorized to do so. The generalization of this concept is that it must be easy for an auditor to determine the access rights and privileges of any user (or group of users), and that the effective infrastructure exists so that those access rights are enforced securely.
- **Transparency.** All business processes and internal controls must be able to be analyzed fully, so that their functions, as well as side effects, can be understood and measured. Any process that is opaque cannot, by definition, be in compliance since its operation cannot be understood fully. The concept of transparency is a familiar one in financial reporting, since the recent corporate scandals have given rise to a strong demand for full accountability (transparency) in the reporting of all information relating to a company's financial state of health. Compliance extends this notion to internal controls, so that the effectiveness of these controls can be assessed completely.
- **Measurability.** Processes that cannot be measured cannot be managed successfully. Compliance generally includes measuring these internal processes and quantifying their success or failure for the auditors. The ability to measure the current operation of a set of controls (through logging, auditing, event correlation, visualization and the like) is a cornerstone of any compliance effort. Manual, paper-based controls are often difficult to measure, and are therefore relatively ineffective when used to establish compliance.

Computer Associates Int. (CA) has decided to work in two different directions around Compliance Management:

1. To study the roles of Standards in Compliance. CA selected two standards quoted by the European Commission COBIT and ISO Standard 17799 plus the IT Infrastructure Library (ITIL)
2. To develop a Compliance Management Framework

2 The roles of Standards in Compliance

Generally, a governmental regulation does not specify what technology is required in order to meet its requirements. In fact, many regulations do not even specify any details of an effective internal control. Therefore, administrators and compliance officers are left to determine what methods they will use to meet the often vague requirements within each regulation.

The European Commission took a different approach. 0n 22 March 2005, the European Commission adopted the Commission Regulation (EC) no 465/2005 aimed at tightening information systems security across the European Union's 25 member states.
(europa.eu.int/eur-lex/lex/LexUriServ/site/en/oj/2005/l_077/l_07720050323en00060008.pdf)

Paying agencies associated with the European Agricultural Guidance and Guarantee Fund (EAGGF) are now required to select either COBIT, ISO Standard 17799 or the Bundesamt für Sicherheit in der Informationstechnik: IT-Grundschutzhandbuch/IT Baseline Protection Manual (BSI) as the basis for their information systems security.

The regulation directs that one of the three standards must be used retroactively from 16 October 2004. From financial year 2008, starting 16 October 2007, auditors must provide a statement on the security measures in place based on the chosen standard.

During the period 2004-2007, the annual auditors' reports are required to include a score for each domain of the chosen standard based on a maturity model developed directly from COBIT's Generic Process Maturity Model. Even if a member state chooses one of the other two standards, the auditor still needs to use the COBIT-based maturity model as part of the reporting mechanism.

2.1 COBIT

COBIT stands for Control Objectives for Information and related Technology and is an open standard for control over information technology, developed and promoted by the IT Governance Institute.

The main objective of the COBIT project is the development of clear policies and good practices for security and control in IT, for endorsement by commercial, governmental and professional organisations, world-at-large. It is the goal of the project to develop these control objectives primarily from the business objectives and needs perspective. (This is compliant with the COSO (Committee of Sponsoring Organisations of the Treadway Commission—Internal Control-Integrated Framework, 1992) perspective, which is first and foremost a management framework for internal controls). Subsequently, control objectives were developed from the audit objectives (certification of financial information, certification of internal control measures, efficiency and effectiveness, etc.) perspective.

COBIT is aimed at addressing business objectives. The *Control Objectives* make a clear and distinct link to business objectives in order to support significant use outside the audit community. *Control Objectives* are defined in a process-oriented manner following the principle of business re-engineering. At identified domains and processes, a high-level control objective is identified and rationale provided to document the link to the business objectives. In addition, considerations and guidelines are provided to define and implement the IT Control Objective.

The classification of domains where high-level control objectives apply (domains and processes), an indication of the business requirements for information in that domain, as well as the IT resources primarily impacted by the control objective, together form the COBIT *Framework*. The *Framework* is based on the research activities that have identified 34 high-level control objectives and 318 detailed control objectives.

The COBIT *Framework* consists of high-level Control Objectives and an overall structure for their classification. The underlying theory for the classification is that there are, in essence, three levels of IT efforts when considering the management of IT resources. Starting at the bottom, there are the activities and tasks needed to achieve a measurable result. Activities

have a lifecycle concept while tasks are more discrete. The life-cycle concept has typical control requirements different from discrete activities. Processes are then defined one layer up as a series of joined activities or tasks with natural (control) breaks. At the highest level, processes are naturally grouped together into domains. Their natural grouping is often confirmed as responsibility domains in an organisational structure and is in line with the management cycle or life-cycle applicable to IT processes.

Thus, the conceptual framework can be approached from three vantage points: (1) Information Criteria, (2) IT Resources and (3) IT Processes. For example, managers may want to look with a Quality, Fiduciary or Security interest (included in the *Framework* as seven specific information criteria). An IT manager, on the other hand, may want to consider IT resources for which he/she is accountable. Process owners, IT specialists and users may have a specific interest in particular processes or activities/tasks. Auditors may wish to approach the *Framework* from a control coverage point of view.

These three vantage points are depicted in the COBIT Cube.

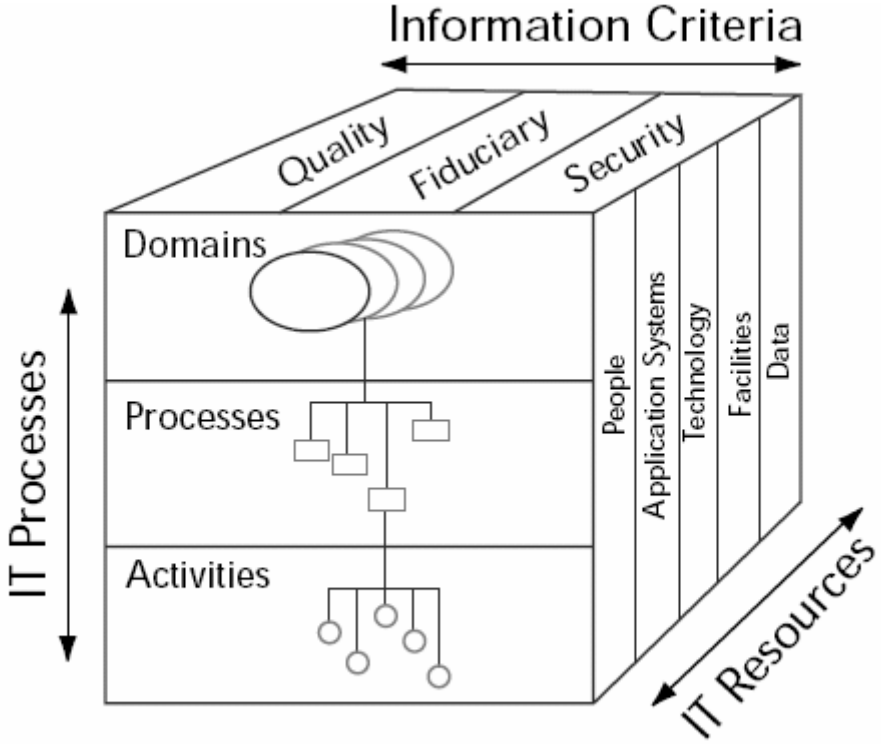

2.2 The Family of BS 7799 and ISO/IEC relatded Standards

ISO/IEC 17799 is a code of practice for information security management. It is designed to serve as a single reference point for identifying the range of controls needed for most situations where information systems are used. This code of practice started life as the UK standard BS 7799-1 in 1995 and was then released as an ISO/IEC standard in 2000. On 15 June 2005, a new version of this standard ISO/IEC 17799:2005 has been approved

BS 7799-2 provides a specification for an information security management system (ISMS). This includes a number of processes for designing, implementing, maintaining and updating an ISMS. BS 7799-2 can be used for ISMS certification according to the European standard EN 45012 and the accreditation guidelines EA 7/03.

To be published in November 2005, BS ISO/IEC 27001 is the new complementary standard to BS ISO/IEC 17799:2005 (BS 7799-1). The standard will provide a specification for ISMS and the foundation for third party audit and certification. It is harmonized to work with other management system standards such as ISO 9001 and ISO 14001 and will assist in the integration and operation of an organization's overall management system. The new standard, when published, will replace BS 7799-2:2002. BS ISO/IEC 27001 will also ensure effective information security management is established and maintained through a continual improvement process, and will implement the OECD principles governing the security of information systems and network

The standard ISO/IEC 17799:2005

The new standard now contains eleven 'core' chapters, as opposed to ten previously. The existing chapters have also been renamed and re-organized. The new chapter structure is as given below:

- Security Policy
- Organization of Information Security
- Asset Management
- Human Resources Security
- Physical and Environmental Security
- Communications and Operations Management
- Access Control
- Information Systems Acquisition, Development and Maintenance
- Information Security Incident Management
- Business Continuity Management
- Compliance.

The new version of the standard also introduces a range of new controls (seventeen in total) to address a number of emerging issues not previously covered. These include topics such as provision of outsourcing, external service delivery, and patch management. Equally, other areas have been substantially extended or re-shaped, such as employment termination, and mobile/distributed communication. Several old controls have been retired, or merged into others. These modifications give a new total of 134 Controls, a net increase of 7.

In addition to the content itself, steps have also been taken to enhance the "user friendliness" of the standard. The standard has also been normalized to position itself to sit more comfortably alongside related security standards in the future.

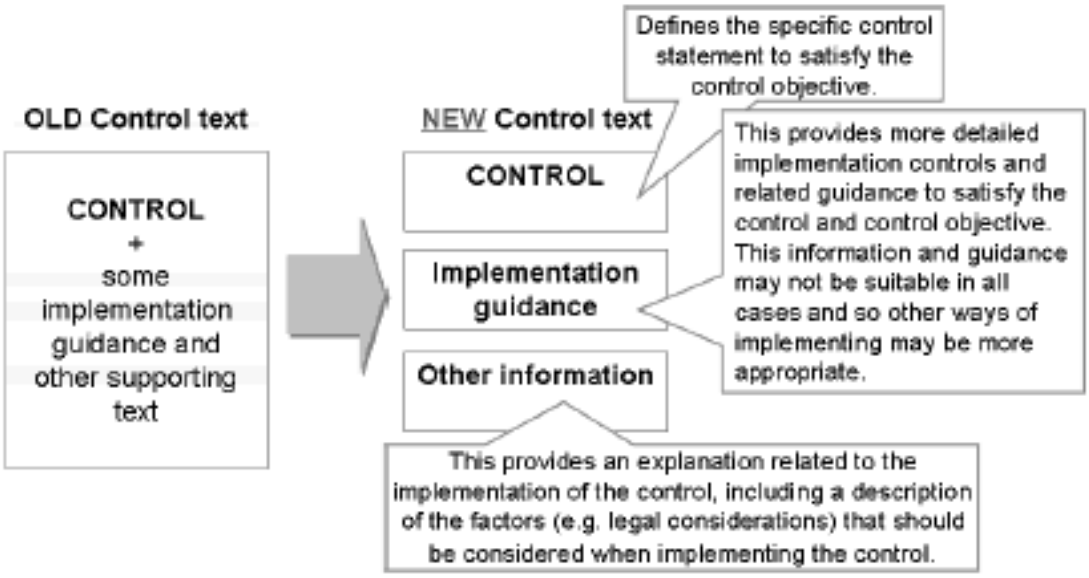

2.3 IT Infrastructure Library (ITIL)

ITIL is a set of best practices for IT Service Management that has been evolving since 1989. It began as a set of processes for use by the UK government to improve IT service management and has been adopted by the industry as a basis for successful IT service management.

It is gaining worldwide acceptance as the standard for IT service management. ITIL describes the processes that are executed as part of the management of IT. These processes are combined in nine sets. *Security Management* is one of the ITIL processes. The *Security Management process* has important relationships with other processes

ITIL is concerned with the *best practice* in the management and exploitation of the IT infrastructure. ITIL has arisen from practical experience and this makes ITIL a recognisable and practice-based approach.

Practice has also shown that using ITIL increases the quality of the IT service. ITIL focuses on managing an existing working environment, and managing changes in that environment.

ITIL is not specifically concerned with system development. Nor is ITIL concerned with the strategic and tactical processes required for developing the IT architecture and infrastructure. Also, ITIL does not specifically focus on corporate policy.

The following model describes the relationship between Security and the ITIL Core set. Every aspect of IT Service Management has Security Management considerations. There is a specific relationship with Availability Management – one of the prime aspects of security is Availability – and through this Business Continuity, but this should not be allowed to detract from its importance throughout the Service Management scenario.

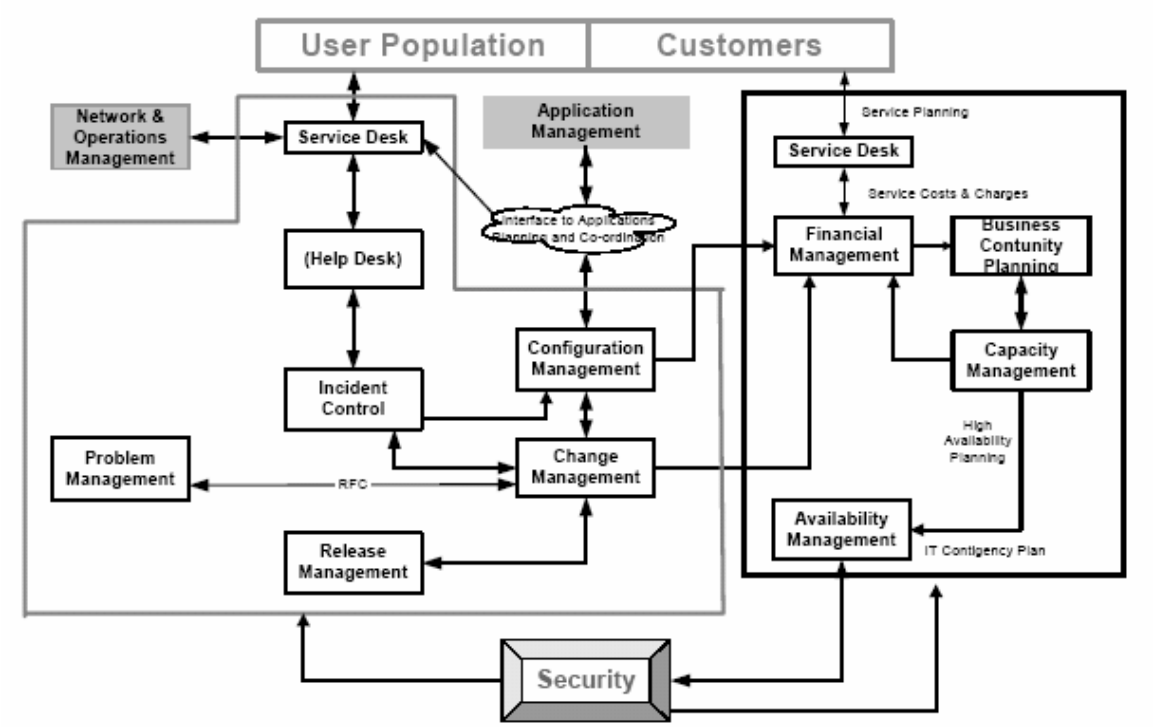

2.4 Integration

Firstly, we mapped the COBIT Control Point vs. ISO 17999 control point and, clearly, COBIT encompasses a larger spectrum where ISO 17799 is stronger in security controls. In one case for example, we found 7 ISO 17799 control points mapping on 1 Cobit control point. During this exercise, we don't identify contradictions between the two standards

Both COBIT and ISO 17799 does not address IT processes in details where ITIL does.

All of those standards are technology-independent.

Our recommendation will be

- To use CobiT and ISO 17799 for determining current status and identifying weaknesses in processes and controls.
- To use ITIL for improving your process
- To use the COBIT-based maturity model as part of the reporting mechanism.

3 The Compliance Management Framework

A Compliance Management Framework is the blueprint for protecting your business and its service offerings.

Such a framework, consisting of people, processes and technology, provides a concise yet high-level and comprehensive strategy to shape your tactical architectural requirements in relation to compliance objectives, To help you formulate customized compliance strategies, CA provides a framework that serves as a foundation for designing and constructing capabilities to set policy, monitor activities relating to compliance and respond to any potential non-compliant situations.

A Compliance Management Framework represents organization-wide priorities. It clearly defines the value of information assets and the underlying business requirements and assump-

tions that drive compliance activities. Development of a Compliance Management Framework begins with an understanding of what legislation-mandated controls must be in place. Next, your organization should assess whether such controls currently exist and, if so, are effective. By going through this process, your organization can make the hard decisions on the Compliance Management Framework up front, making implementation of the rest of the program much easier. Furthermore, an organization that regularly reviews and assesses its current controls associated with people, process and technology can identify key missing and ineffective elements within its compliance program. The implementation of a compliance management framework is not a one-time event. Successful implementation of compliance management requires that it be incorporated into the daily business functions of the organization. This will provide management, and those who have a need to know, with the strategic and operational pulse of the risk and compliance process.

3.1 Components of the Compliance Management Framework

The implementation of the Compliance Management Framework enables awareness of elements affecting the vulnerability, risks, policies and procedures relating to financial transactions, applications, systems, people and the security controlling them. The compliance management framework is comprised of three primary layers:

- **Presentation Layer** which shows a compliance dashboard, providing a visualization and at-a-glance view of the compliance process, vulnerabilities and risks.
- **Process Layer** which establishes the linking control of objects, activities, schedules and events to corporate compliance policy and procedures. Intelligent compliance monitoring and business rules help ensure that vulnerability risks are mitigated and corrective actions are initiated
- **Technology Layer** which provides a wealth of knowledge about networks, security, hardware, software and IT processes.

The framework provides integration among each layer to enable a holistic method of managing compliance components associated with compliance policies and procedures. Organizations traditionally manage compliance requirements by functional departments that have little knowledge of the state of other objects that affect the level of compliance.

The lack of a compliance framework places a heavier burden on manual processes, compliance tasks and communications across departments and functional units.

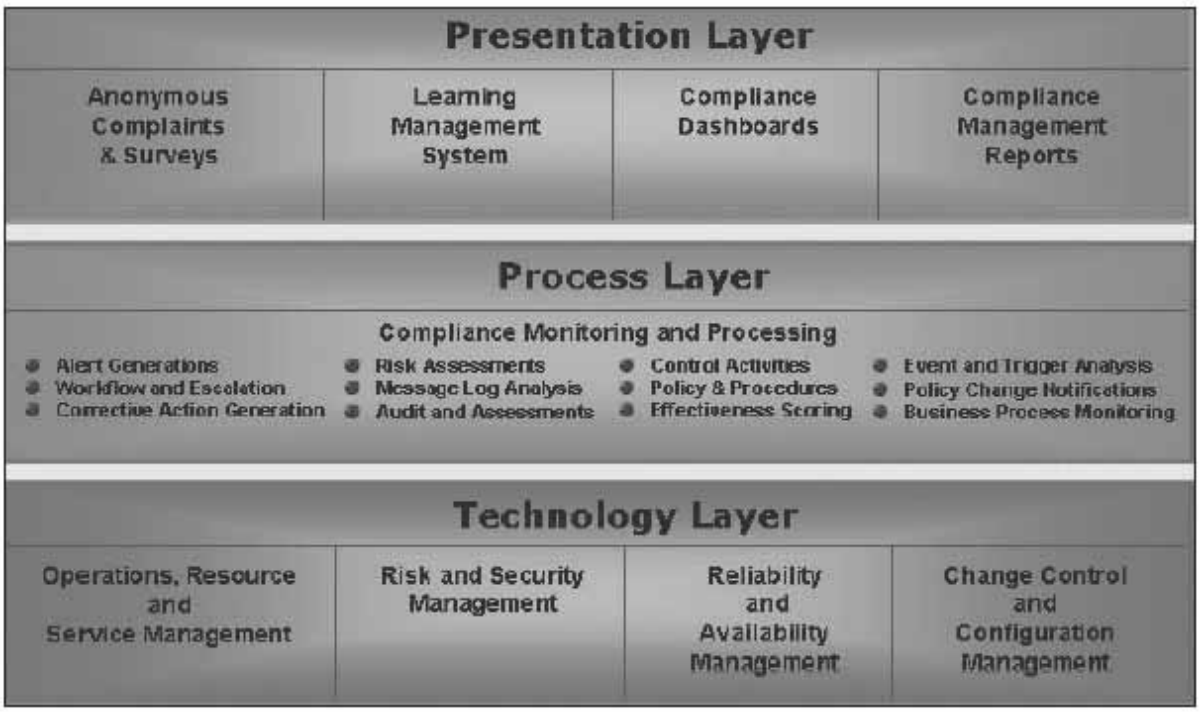
Presentation Layer
Anonymous Complaints & Surveys
Learning Management System
Compliance Dashboards
Compliance Management Reports
Process Layer
Compliance Monitoring and Processing
Alert Generations
Workflow and Escalation
Corrective Action Generation
Risk Assessments
Message Log Analysis
Audit and Assessments
Control Activities
Policy & Procedures
Effectiveness Scoring
Event and Trigger Analysis
Policy Change Notifications
Business Process Monitoring
Technology Layer
Operations, Resource and Service Management
Risk and Security Management
Reliability and Availability Management
Change Control and Configuration Management

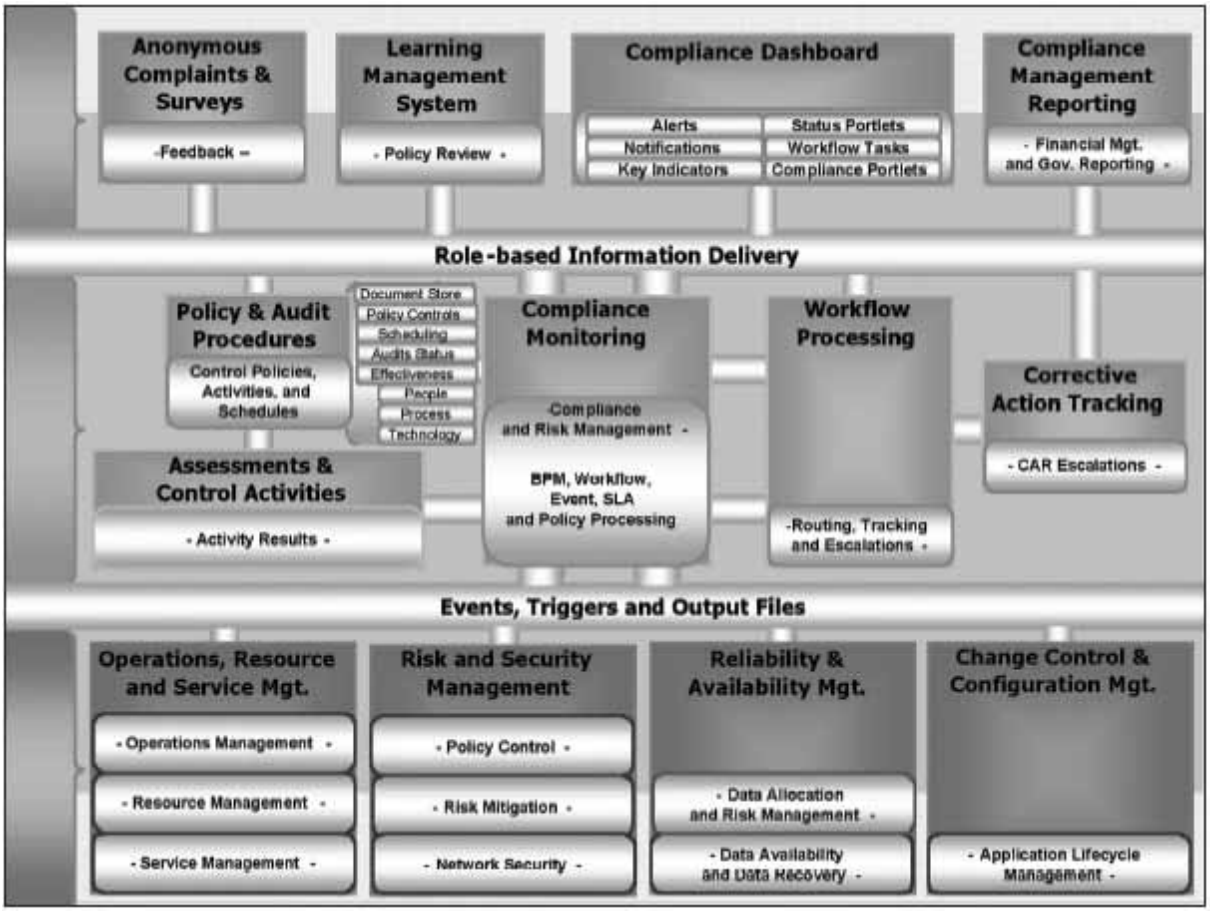
Anonymous Complaints & Surveys
-Feedback –
Learning Management System
- Policy Review -
Compliance Dashboard
Alerts
Status Portlets
Notifications
Workflow Tasks
Key Indicators
Compliance Portlets
Compliance Management Reporting
- Financial Mgt. and Gov. Reporting -
Role-based Information Delivery
Policy & Audit Procedures
Control Policies, Activities, and Schedules
Document Store
Policy Controls
Scheduling
Audits Status
Effectiveness
People
Process
Technology
Compliance Monitoring
Compliance and Risk Management -
BPM, Workflow, Event, SLA and Policy Processing
Workflow Processing
-Routing, Tracking and Escalations -
Corrective Action Tracking
- CAR Escalations -
Assessments & Control Activities
- Activity Results -
Events, Triggers and Output Files
Operations, Resource and Service Mgt.
- Operations Management -
- Resource Management -
- Service Management -
Risk and Security Management
- Policy Control -
- Risk Mitigation -
- Network Security -
Reliability & Availability Mgt.
- Data Allocation and Risk Management -
- Data Availability and Data Recovery -
Change Control & Configuration Mgt.
- Application Lifecycle Management -

3.1.1 Presentation Layer

The Presentation Layer of the framework provides the user interface into the compliance world. Organizations could provide:

- Anonymous Complaint System—Report questionable activities
- Online Surveys and Collaboration—Compliance knowledge sharing
- Learning Management System—Internal policy and procedure training
- Role-Based Compliance Dashboard—Relevant compliance information and alerts
- Compliance Reporting—Financial and governmental reporting management

The compliance dashboard typically shows key indicators on the health or risk relating to compliance.

Executives and managers normally will have a view of the overall compliance process, indicating whether they are on target with the implementation of their policies and procedures, control activities, training, assessments, audits and reporting. User-specific Alerts and Notification provide at-a-glance warnings of key issues that are generated from the Compliance Monitoring and Intelligent Processing component of the Process Layer.

3.1.2 Process Layer

The Process Layer of the framework allows for the intelligence monitoring of compliance activities with a very robust business rules engine. The business rules can initiate activities, including setting alerts and notifications, workflow transactions, corrective action issue tracking or updating compliance portlets' key indicators and assessing risk or compliance scores.

The complexity of today's business environment, IT infrastructure, industry regulations, company policies and regulatory compliance requires a framework that ties all elements together and monitors the risk and compliance process.

The Process Layer leverages existing IT technologies by associating business rules with objects from the IT layers. Compliance objects vary, from actual hardware devices to software (including operating systems, applications, databases) to configurations and settings (such as open ports, patches, security settings) to processes like transactions, events, triggers and log files.

IT applications will often have policies to control behavior. These policies may be used directly to monitor the status or state of a policy, yet often a compliance policy may require additional controls or the combining of several controls together. The Process Layer facilitates the augmentation of existing policies and provides rules from multiple disciplines.

Compliance policies can also extend into posting of transactions using business rules to help ensure that the proper debit and credit accounts are normal, or that a sales order has not already been posted in another sales accounting application.

3.1.3 Technology Layer

The Technology Layer includes every piece of hardware and software that is associated with the integrity of the financials. This includes file servers, application servers, personal computers, routers, hubs, connectivity, networks, access points, mainframes, financial applications, supporting applications, security and authentication, as well as policies and procedures governing these devices.

The compliance policies and procedures define key financial processes and the guidelines and controls required. Every control point becomes a cascading web of interconnective components, each with their own vulnerabilities and risks. Objects are defined once in the system, with multiple associations and states. These objects are defined in the process layer which will associate and inherit attributes based on business rules. For example, each of the Sales Order Entry Systems will have a unique object defined and will be associated with other objects, such as Sales Order Entry One, Release, User Security, Operating System, OS Patch, OS Security Settings, OS Trusted Connections, Machine Type, Database and so forth.

Every key attribute can be defined and used with rules. A rule may state for all servers that contain financial information that they shall not have any open ports or connectivity to a trusted server with open ports. Rules in this example would be written once and every server that was associated with financials would be propagated with that rule.

Output logs or events may be defined as an object. This enables the leveraging of existing monitoring and policy-based applications. For example, the reliability of a financial database requires that reliable backups are created and that databases have significant free space to grow during processing cycles.

Compliance rules would be written to monitor the backup software to help ensure that risk is minimized on the data integrity and reliability. If the risk is too high, alerts and corrective actions would be generated and shown in the dashboard of those that are affected. That could include the CIO, CFO, database administrator and even the compliance officer. The business rules would define who gets notified, when and by what method. If the risk or severity was low, perhaps only the database administrator would be notified. If the risk or severity was high, the CCO, CFO or CIO may find an alert on his or her dashboard. If a corrective action was not addressed in a specified period of time, the issue would automatically be escalated to another role, such as the compliance officer.

4 Conclusion

Auditing control objectives like COBIT or ISO 17799 approach IT governance from the top down, describing the results expected from good governance. ITIL approaches governance from the bottom up; it describes practices that result in good governance. There is no compulsion to institute ITIL practices to pass audits and comply with laws. However, ITIL practices are one route to IT governance. Often, an ITIL practice is the easiest and surest way to attain a control objective. Control objectives pair with control activities. A control activity is something done to attain a control objective.

We hope our Compliance Management Framework will help you to define a suitable internal-control framework

Using GIS Tools to assess the Vulnerability of the Internet

Neil E. Robinson

RAND Europe (Cambridge),
Unit 1, Westbrook Centre,
Cambridge, CB4 1YG
neilr@rand.org

Abstract

This paper outlines the results of preliminary research undertaken to explore how Geographic Information System (GIS) tools may be useful in the assessment of the vulnerability of the Internet. High-speed optical fibre networks, owned by global telecommunications companies, form an important element of the Internet. They are most dense in major cities. In these areas, Metropolitan Area Networks (MAN) form the important link between long haul national and international networks, and networks constrained within buildings. Geographical visualisation and spatial statistical approaches, based on commercially available data, are put forward as a useful means of describing the location of physical vulnerabilities within a constrained geographical area. Once identified, these physical vulnerabilities could then be matched against the logical architecture of telecommunications networks to inform a broader assessment about the vulnerability of the Internet.

1 Introduction

Government bodies have recognised the growing importance of telecommunications networks as a vital component of the Critical National Infrastructure (CNI). Fibre optic cables form an increasing part of the physical element of these networks. However, the events of September 11th, 2001 and other large scale physical events, including a fire in a tunnel in Baltimore in the United States in 2001 and two incidents in Manchester, UK in 2002 and 2004, have highlighted the vulnerability of these networks to physical disruption. In the United Kingdom this infrastructure is at its most dense in London, where several providers have built large networks, known as fibre-optic 'rings', and based on high-speed optical network technology capable of carrying tens of Gigabits per second (Gbp/s) of data. Together, these networks collectively constitute the London Metropolitan Area Network (MAN) infrastructure. The MAN is an intermediary type of telecommunications network, servicing the access requirements of local customers (who need lots of connections close to their premises and smaller bandwidth) and the transit needs of intercity and intercontinental 'long-haul' providers (whose requirements are high bandwidth and minimal local access requirements). Within the MAN, Internet Exchanges (IX) form a key part, allowing providers to exchange or 'peer' traffic with one another, facilitating the efficient transmission of data and cost effective use of their networks. In London, the London Internet Exchange (LINX) is the pre-eminent IX and has facilities at several locations across the capital.

A spate of recent incidents in the last few years have highlighted the importance of vulnerability of telecom networks to policy makers, scientists, those in the security world and the general public. These include the recent failure of British Telecommunications 'Colossus'

S. Paulus, N. Pohlmann, H. Reimer (Editors): Securing Electronic Business Processes, Vieweg (2005), 324-334

Internet Protocol (IP) network [King01] outages of various Internet Service Provider (ISP) networks and more recently, two fires that occurred at network facilities in Manchester, UK. Further abroad, the aftermath of the terrorist attacks of September 11th 2001 in New York prompted a review of the resiliency of the Internet and its ability to withstand such damage. These incidents are detailed more fully later on, but suffice to say that they continue to inform efforts by governments to address this issue. This investigation aims to assist in this process, by providing an adjunct to the 'traditional' method of assessing vulnerability. Clearly, the assessment of the vulnerability of such a complex system as the Internet is no trivial task, but it is hoped that this investigation, which highlights how vulnerabilities in the physical MAN infrastructure might affect the logical network and hence the wider Internet, may be a step towards a more holistic view.

Two aspects are evident in defining the Internet infrastructure. Firstly, the physical infrastructure – e.g. those cables (whether they be fibre optics or copper) and ducting which make up the networks. Secondly, the infrastructure can be described logically – e.g. the routing of data over these wires, how communications traffic flows around this infrastructure, and the role of specific devices. Understanding resiliency in the framework of this model allows for ease of assessment in the complex environment of interconnection at the backbone level, and also places any investigation in the sound engineering frame of commonly understood telecommunications principles. In order to place this understanding into a commonly accepted model, the Open Systems Interconnect was used as the frame of reference for this study.

The Internet infrastructure is a commonly accepted element of the Critical National Infrastructure (CNI). The UK government defines the CNI as:

> *"Those parts of the United Kingdom's infrastructure for which continuity is so important to national life that loss, significant interruption, or degradation of service would have life-threatening serious economic or other grave social consequences for the community, or any substantial portion of the community, or would otherwise be of immediate concern to the Government."* [NISC01]

Protecting this infrastructure has, in the last few years, become a key concern for government. The announcement of the existence of the National Infrastructure Security Co-ordination Centre (NISCC) in 2001 marked the first public acknowledgement by government that the resiliency of the Internet infrastructure had become important.

2 Aim of this study

The aim of the investigation was to assess whether a GIS based methodology is useful in the context of evaluating the vulnerability of the Internet infrastructure. A GIS based approach was chosen for the following reasons:

- To provide a realistic physical spatial based adjunct to existing logical assessments of vulnerability that comprise the majority of CNI assessments.
- To enable a more thorough and robust appreciation of vulnerability, by combining what is already known about logical vulnerabilities, with that of a clearer understanding of physical vulnerabilities.

This approach was adopted due to limitations in current methodologies, some of which are outlined in the literature review. Such assessments are generally based upon overly simplistic asset lists of the location of nodes, links and devices, rather than potentially more important issues such as the location of these devices, the proximity of these elements to others of a

similar type and their proximity to other locations that could potentially be attractive as targets or vulnerable in some other way (for example, transportation links).

The use of GIS in assessing risk or vulnerability has a clear precedent in the insurance industry, where it has been a standard part of the process of the calculation of insurance premiums for many years. Risks to certain types of hazards or perils, specifically ones of an environmental nature have been assessed using GIS techniques and are now common practice. In the UK, these perils are classified as wind, flood, fire, freeze, theft and subsidence. Companies such as Intermediary Systems Limited provide services to insurance companies by developing models showing risk of different geographic areas to these hazards [Insu04]. Such models often use area data as the main building blocks.

The identification of risk relies heavily on geographical information (often, with such environmental hazards, based on area or raster interpretations of the data) because the calculations are often based on spatial units that enable delineation of different at risk areas. However, care must be taken that the appropriate resolution is used when undertaking these calculations. Too low a resolution, and the description of the 'at risk area' covers a space too great that is not commensurate with the type of hazard. Too high a resolution and the space defined as being at risk from a particular hazard may exclude those who may fall into the category.

In a similar way that accurate information on environmental risks is useful to the insurance industry, accurate information on the spatial distribution of optical networks is valuable to those charged with protecting the Internet infrastructure. It is also useful to customers of providers. Consider the following example:

A company with a high degree of reliance upon the Internet, e.g. for e-commerce activities, decides to purchase connectivity off two different service providers, so if one network fails, a second is immediately available, thus helping to provide for redundancy and seperacy. This is a common approach by those aggressively pursuing Business Continuity Planning (BCP). The customer is separately assured by both providers that each infrastructure is redundant and thus places a contract with both, meeting its requirements for the assurance of a redundant and separate Internet connection. However, due to the complexities of the market both providers utilise larger upstream providers that run their optical networks along either the same RoW, or highly spatially co-related paths. The customer is operating in a false world due to his assumption that both links are diverse and separate – which they may be at the level of the logical topography, but are certainly not at the level of the geography of the network.

The results of this investigation, which may reveal the nature of this spatial correlation, could assist companies in identifying where they have erroneously placed faith in two providers to meet seperacy requirements but who actually have bought a service that is neither spatially diverse nor separate. Similarly, identification of the spatial correlation helps those who must assist in the provision of resiliency measures for the Internet infrastructure more generally.

3 Other relevant literature

A small but growing body of literature is evolving concerning geographical investigations of the vulnerability of the Internet infrastructure. However, some of this research does not take into account the technical qualities of the infrastructure, nor how logical and physical network properties interact to provide the scale free architecture of the Internet. For example research conducted in 1999 by the University of Dartmouth highlighted issues regarding the network of fibre optic cables across the United States [Calo02] and their vulnerability to single points of failure. Several logical infrastructural maps were presented: firstly, as evidence of the ease

in which such information could be obtained, and secondly as an aid to visualising the clear vulnerabilities inherent in the spatial layout of the architecture. The authors commented that although the maps were of low resolution, single points of failure were clearly evident. However, this research described the *logical* links between nodes rather than the specific *physical* path down which provider traffic travels. It is crucially important to understand the difference between these two. From the perspective of a logical link, for example between London and Manchester, such a direct link may indeed exist and be the primary route, but it does not show that it is vulnerable or does not exhibit adequate resiliency. Indications of vulnerability taken from superficially obvious evidence such as logical network maps generally represent unrealistic appreciations of the real situation. Unfortunately this is the approach many security researchers have taken in the study of this topic.

A somewhat similar approach was put forward in Grubesic et al, where a matrix was established to describe different provider arcs between major US metropolitan areas [GrOM03] This model recognised Barabasi's work, rightly identifying the Internet as a hub and spoke network, with a large number of logically poorly connected nodes (such as home computers and end user sites with a single Internet connection) and a small number of logically highly connected nodes (such as telecom hotels and Internet Exchanges). Grubesic's approach considered the multitude of providers supplying connectivity to each city, but no investigation was conducted of the logical issues surrounding loss of connectivity for each node. For example, the use of peering in each of the major metropolitan cities outlined in the paper is so extensive that it could reasonably be expected that in the event of one intercity link failing, traffic would be rerouted along infrastructure owned by unaffected peers.

The method of identifying the single point of failure at a national or regional level was approached in another study [LeFV02] Researchers at the Forschungszentrum Julich developed a model based on their Internet Security Simulation (INESS) that aimed to simulate the impact of a targeted strategic disturbance and whether there are specific points that, if disturbed, would result in consequences for the network. Failures in two points were modelled – the transatlantic backbone between Europe and the United States – called 'Euro1' and all Internet Exchanges (IX) in London, collectively called 'London'. The conclusions of this model were interesting, in that firstly the effects of the loss of the backbone and an IX were worse than the loss of just an IX and secondly, regarding the issue of capacity. The report concluded that the picture of the capacity of different providers was not positive from a perspective of resiliency. There are many providers in different metropolitan areas, but only a limited number with extensive bandwidth. These providers dominate the network topology. All the providers operate their networks at extremely small margins of capacity, due to the competitive nature of the telecommunications market and need to maximise efficiency. This manifests itself in very limited spare capacity in nodes and arcs to deal with more traffic than is provisioned for in normal day- to-day operation. Although nodes such as routing and switching hardware, network devices and the like can be easily reconfigured to handle more demand (if the operating parameters of the provider allows it) this is not necessary the case of links between physical facilities. This model also showed that the London – New York route was especially critical in the provision of Internet connectivity between Europe and the rest of the world.

Finally, Gorman's classified thesis outlined an investigation that was, according to the Washington Post, a comprehensive and detailed map of what companies provided connectivity to businesses, government departments and regional areas, across the globe [Blum03]. This paper was classified at the behest of the authorities and so was unavailable for assessment but according to the Washington Post the project highlighted which companies provided what

connectivity to a number of government and commercial organisations to an unprecedented degree.

4 Summary of results

This section discusses the role of the MAN in the Internet Infrastructure and describes the methods used in this investigation, as well as substantive conclusions from two of the approaches.

MANs are defined as telecommunications networks that interconnect users with computer resources in an area greater than that covered by a Local Area Network (LAN) (which typically covers a building or floor of a building) but smaller than a Wide Area Network (WAN), which is 'geographically dispersed' and may cover several countries over a single continent [Tech03]. It is also a term applied to the interconnection of networks in a city into a single larger network, allowing for more efficient connection into the WAN. Finally, the term describes the connection of LANs to backbone infrastructure.

From the perspective of the architecture of the Internet, MANs occupy a very important position. They act as an intermediary layer between the long haul carriers who provide intercity and intercontinental connectivity, and the provision of 'last mile' connectivity to customer premises at appropriate bandwidth. MANs are thus an intermediary layer that connects 'transport' or 'transit' traffic and local or access traffic destined for end users [Tele02].

4.1 Methods

This section covers the methodology and elaborates on three approaches adopted to describe the issue of physical vulnerability. It also presents summary highlights of the results of the modeling using two of these approaches that represent the most interesting results.

These approaches were:

- Visualisation of vector data
- Visualisation of raster images of varying cell sizes (accomplished by summing the instances of provider optical networks in squares in grids of varying sizes) –the first descriptive statistical approach known as the raster based approach.
- Assessment of distances from selected nodes to provider networks, (using a vector method, accomplished by measuring the distance to the nearest point on a provider network from a selected node) – the second descriptive statistical approach known as the vector based approach.

Concerns remain over the accuracy of the data used for this exercise. However these constraints were accepted with a view to demonstrating the utility of the use of GIS tools to assess vulnerabilities in the physical geography of the Internet infrastructure, rather than obtaining any highly accurate conclusions. It was also recognised that the data are largely from one single source, a significant restriction on the scope.

4.1.1 Results from the visualisation method

Data was extracted from electronic files provided by commercial market research company Telegeography and imported into the GIS as separate layers containing different coloured vector information. This dataset was then overlaid with other layers representing road infrastructure and land / water boundaries and point information detailing the location of major PoPs, co-location facilities and telecom hotels, purely to provide context. This layering proc-

ess represents the first method, which was simple visualisation of physical and human geographic elements to provide context. An example of the detail around the area of the Docklands / Isle of Dogs is given at Fig. 1. It was recognised that although this data looked visually accurate, the projection was not real world (being in Non Earth metres rather than Latitude & Longitude). Concerns also remained over its accuracy. However these constraints were accepted with a view to demonstrating the utility of the use of GIS tools to assess vulnerabilities in the physical geography of the Internet infrastructure, rather than obtaining any highly accurate conclusions. It was also recognised that the data are largely from one single source, a significant restriction on the scope.

Fig. 1:Visualisation of networks in the Eastern part of the bounding area