Electronic Information Security Documentation

Peggy Fung[1] , Lam-for Kwok[1], Dennis Longley[2]
[1] Department of Computer Science
City University Of Hong Kong
Kowloon, Hong Kong
[2] Information Security Research Centre
Queensland University of Technology
Brisbane, Australia

peggy@cs.cityu.edu.hk cslfkwok@cityu.edu.hk d.longley@qut.edu.au

Abstract

Effective security management depends upon good risk management, which is itself based upon a reliable risk assessment, involving data collection of all the facets influencing system risk. Such data collection is often an extremely onerous task, particularly if a substantial proportion of the required information is not adequately documented. Hence comprehensive, updated information security documentation is a keystone of good information security management. Whilst the recently emerging information security management standards provide some implicit guidance on the development of documentation; there is relatively little support available for security officers attempting to develop and maintain such documentation.

Traditionally textual security documents are not necessarily the most appropriate format for describing the security of large complex, networked systems, subject to frequent updates. It has been suggested [1], [2] that a security officer's workstation, with a database and GUIs, may present a more effective form of security documentation. However, such a tool requires a well-developed model of the information system and, as discussed in this paper, a standardised means of representing security entities.

This paper proposes an information security model to facilitate the development of electronic security documentation. A proposed security entity classification scheme is first described. Such a classification scheme and the use of object identifiers to identify security entities greatly facilitates the development of a security officer's workstation. The potential of the model for risk assessment and security design is described.

A prototype model was developed in Visual Basic to test the concepts proposed, and a Java based model is currently under development at the City University of Hong Kong.

Keywords: Information Security Management, Risk Analysis, Information Security Standards, Information Security Documentation.

1 Introduction

In the past three decades there has been a sharp increase in the awareness of the potential deleterious impacts, arising from inadequate information security. Unfortunately the scale of the problem has escalated more rapidly than the commitment to combat it. Moreover, in many cases, the media emphasis on hackers and viruses has distorted the debate and tended to divert senior management awareness from the more fundamental aspects of information security.

In particular there may be a sharper focus on technical solutions, to well advertised attacks, than on the fundamental necessity to view information security as an organisation wide business/ management / technology issue.

Organisational security officers are charged with ensuring the security of information assets and systems. As such, they are perilously located between management and technology. They are required to ensure that the technological systems are implemented and operated in such a manner, that the business risk to organisational information assets and systems is contained within acceptable boundaries. In effect they are required to assess the level of business risk from an information security viewpoint, and to recommend operational or technical changes designed to bring that risk down to some acceptable, but often unspecified, level.

The first step in such a risk assessment involves a major data collection and evaluation process. This process is often extremely time consuming, disruptive and expensive. Hence, there is a temptation to work with over-simplified models of the information system, and to request highly subjective estimates of risk-related data from I.T. staff.

Subjective risk assessments bode ill for a security officer in a highly complex, networked environment, particularly when information security failures may have significant impacts on the financial well being or the regulatory or contractual obligations of that organisation. In the aftermath of a serious information security failure, security officers may well be called upon to supply convincing, documented, evidence that their risk assessment recommendations, to senior management, were well founded.

Hence one can easily demonstrate the importance of comprehensive, timely, risk and security documentation, to organisational security officers. Unfortunately, there appears to be minimal support systems available to

security officers tasked with the development, maintenance and interpretation of such documentation.

Information security management standards such as German IT Baseline Protection Manual Standard Security Safeguards [3], BS7799 [4], and ISO17799 [5] do provide an infrastructure of information security management and hence some guidance on the structure of security documentation. Nevertheless it is interesting to compare the emphasis on bookkeeping in the training of financial auditors, with the average educational/training courses for security personnel. In general there is a significant lack of guidance, let alone tools to aid the security officer in the documentation task.

In this paper, we discuss the importance and role of information security documentation. In particular it is suggested that a commonly agreed information security model, and a common method of security entity classification, would facilitate the development of software tools for the production and utilisation of such documentation.

2 Role of Information Security Documentation

2.1 Support for Risk Assessment

The information security industry has made significant advances to meet the perceived threats to organisational information security. Originally, outside the military sector, the major threat identified by the finance and banking industry was the security of electronic transactions, and security manufacturers supplied hardware cryptographic systems to this market. The advent of viruses in the late 1980s spurred a new industry in anti-viral software. Later the development of the Internet as a common communication channel for organisations, expanded the hacker community and the production of firewalls to thwart them. PKI companies provided cryptographic software the emerging E-commerce market, and many organisations now invest in various access tokens such as smart and magnetic stripe cards. The biotechnology industry is also continually gearing itself up for its promised future.

Nevertheless security officers can face a difficult task, in convincing management that these vendor products represent only a part of the solution. Individual countermeasures must be embedded within a coherent information security infrastructure, if the organisational operations are to be adequately protected.

The development of such an infrastructure must itself be guided by effective risk assessment projects. The importance of effective risk analysis was recognised in the early 1970s[8], and there was a strong move by some governments to facilitate the adoption of such methodologies in sensitive computing systems.

Risk analysis includes the identification of assets, threats, vulnerabilities, countermeasures and the evaluation of loss expectancy. An information security risk analysis study defines the IT environment under consideration and recommends corrective actions.

Risk analysis projects were relatively expensive, even in the mainframe computing era, because they involve the collection and evaluation of a significant volume of data including: – the intrinsic threats, the IT system, its physical and operating environment, the assets to be protected and the business functions dependent on those assets.

Such risk studies were either conducted by in-house staff or external consultants. In general the in-house staff often lacked extensive experience of the subjective aspects of risk evaluation, and consultants had no previous knowledge or experience of the organisational system under study. Generally the existing documentation was inadequate, in terms of its content, detail and currency, for risk assessment. Hence the initial familiarisation process was normally accompanied with a major task of data collection.

The magnitude of this initial familiarisation task escalated rapidly, as systems evolved from batch processing mainframes to current complex, multi site networked, client server scenarios. Moreover, the batch processing mainframe environment was stable for long periods, usually between purchases of the mainframe equipment. Hence risk assessment recommendations had a long half-life, significantly reducing the average annual cost of such studies.

In the current climate the complexity and volatility of information systems is such that:

- The risk assessors, must at the outset, have significant knowledge of the organisational system, its environment and the business functions that it supports.
- The system documentation must be sufficiently versatile, comprehensive and timely to reduce the data collection task to achievable levels.
- The cost of risk assessment updates must be minimised.

There appear to be two conclusions from the above:

- IT systems must be fully documented, from a security viewpoint, and such documentation must be regularly updated.
- The abovementioned security documentation must be in a format that significantly reduces the cost and effort of risk assessment exercises.

2.2 Due Diligence

The evolution of IT systems, described above, clearly escalated the magnitude and complexity of the organisational security officer's task. This development in IT systems was moreover accompanied by increasing integration of the IT systems into the organisational business functions, to the extent that the health of the business functions were inextricably linked to that of the supporting computing and computing systems. Computing downtimes, causing merely minor irritation in the erstwhile mainframe era, would be life threatening to most modern corporations.

Hence the security officer is not only faced with a major task of risk assessment in a complex environment, the potential penalties associated with inadequacies, in the subsequent recommendations, have also escalated. Unfortunately given the probabilistic nature of risk assessment, there can be never be a guarantee of incident free operation for the IT system over a long period of time.

In a post security incident environment the security officer must demonstrate that the security systems implemented were reasonably compatible with the true level and nature of the system risk. Moreover, current I.T system failures may have serious consequences for the financial well being of the organisation, and for its compliance with regulatory and contractual obligations. In the current climate management may well be formally required to demonstrate due diligence in the protection of information assets and systems.

Macro risk assessments, based upon apocryphal, subjective assessments, are likely to be unconvincing in the witness stand. Today's security officers would be well advised to equip themselves with comprehensive security documentation, and associated risk assessment strategies, as evidence that they had acted with a high level of professional competence.

2.3 Security Documentation Requirements

It is much easier to make a case for the development of comprehensive security documentation, that to actually produce the documentation itself. In many cases advice takes the form " I would not start from here".

The information security management standards do provide an infrastructure for information security management, which at least suggests a structure for the documentation. A recent paper by the authors [6] suggested the type of current organisational documentation and data that should be collected and packaged to form an initial set of information security documentation.

In this paper the necessary facets of security documentation are described and some insight into recent work on an Information Security Model is discussed.

At the outset the question arises – what is being described by the security documentation? Most system documentation is designed to assist operators and developers in the performance of their tasks. Security documentation is not however aimed normal system operation, but rather at the circumstances in which the system fails, in some sense. Hence security documentation should provide a detailed description of an agreed security model for the system. In other words an organisation's security documentation should contain the local parameters of a generally accepted information security model.

The proposed model need not be described in conventional textual format. Given the complexity, magnitude and volatility of modern information systems, some form of database representation is more appropriate. Moreover such a database should be supported with software tools and GUIs to facilitate the development, updating, investigation, risk analysis and security reporting.

If a common model were employed by organisations then third party vendors would be encouraged develop support software. Moreover, given a common format of security documentation one could envisage situations in which external security advice and expertise were readily absorbed by an organisation. Hence it is possible to envisage a system in which CERT Advisories are automatically downloaded and added to the security database. The security software could then generate a report on the implications of the reported attack for the organisation.

3 A Proposed Model

3.1 Overview

The Risk Data Repository [1], [2] is a risk analysis model, developed some years ago, which aimed to integrate all available organisational data related to security. The model had the ability to evolve over time as it incorporated newly acquired data. The RDR described entities in term of their roles from a security viewpoint, and demonstrated the inter-relationships of security data. The RDR essentially comprised three domains: Environment, Platforms and Assets. The environment domain included elements that effectively hosted or supported the operation of the information processing system: equipment, building, staff. The platform domain was the logical description of the information processing system and its defences. The assets domain described the data and processes, to be protected, because misuse of these assets would have a deleterious effect on the organisational business operations.

The RDR comprised a database and graphical facilities to trace the inter-relationship of security entities. Hence it was possible to trace the effect of a threat of fire in a building to the potential business impact. Experience with the RDR demonstrated three significant aspects of such security modelling:

- the difficulty of describing the wide range of security entities concerned with risk assessment and security modelling;
- problems of importing data from other RDRs; and
- problems arising from the hard coding of security expertise in the model.

It was clear that a major problem in the development of such an organisational risk database lay with the classification of the various entities. There appears to be no common directory to describe such items as: Threats, Computing Hardware, Buildings, Services, Users, Information Assets, Access Control Policies, etc.

In the development of an Information Security Model, to replace the RDR, the concept of Environment, Platform and Assets was extended to five categories:

- ***Systems***: includes hardware, software, platforms, networks, applications, users and information assets.

- ***Environment***: includes locations (sites, buildings, floors and rooms) and services (power, cabling, air conditioning, water and communications).

- ***Security***: includes threats, countermeasures, Threat Trees and Threat Countermeasure Diagrams.

- ***Procedure***: includes external procedures, such as government legislation and international standards, and internal procedures: organization policies, guidelines etc.

- ***Relationships***: security depends critically upon the context of entities and this context is described by relationships. For example, hardware is located in a building, networks are connected to other networks, and a security policy complies with a Standard's recommendation. Relationships among the various entities are defined here.

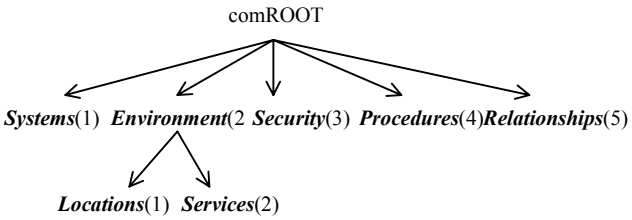

Fig 1. Directory Tree For Security Entities

Each of the above classes has a number of subclasses and the whole set of entities can be described as a directory tree. Borrowing the concepts of X.500[7] the various subclasses and subsequent entities can be classified with object identifiers, representing the set of nodes traversed from the root to that entity (See Fig 1).

The proposed classification system has a number of immediate advantages, from a risk assessment and security documentation viewpoint. Firstly each entity is uniquely and succinctly identified by its object identifier, indicating its position in the directory tree.

Secondly the classification provides a top down model with the major entities specified at an early stage of development. For example, a building, floor and room are each subclasses of the parent - site. It is well recognised in risk assessment that the preliminary investigation involves consideration of the major entities, followed by a subsequent refinement into more detailed areas, as the analysis identifies the risk priorities. Risk assessment models that require full system details to be entered at the outset hit major data collection problems.

The top down approach is also facilitated by the ***Platforms*** entities under ***Systems*** entities. ***Platforms*** are large IT systems comprising all the other ***Systems*** entities, ***Hardware***, ***Software***, ***Networks***, ***Users*** and ***Assets***. Defining ***Platforms*** at an early stage facilitates a large-scale organizational model, e.g. ***Platforms***, located on ***Sites***.

A further advantage of the classification scheme is that it facilitates the importation of data from another risk database, assuming both databases have followed the same classification model. Hence mergers within branches of an organization, with consequent integration of systems, can be readily handled, from a risk assessment / security documentation viewpoint.

The classification system described so far provides only an inventory of the security entities. Security relevant details of those entities, e.g. vulnerability to flooding for a site, communication protocol of a network, issue date of a security manual, are also stored in the database. Given the diverse nature of the entities such attribute information is stored as a <TAG, VALUE> tuple, e.g. <COM PROTOCOL, TCP/IP>.

Risk assessment and security documentation are, essentially concerned with the relationships between these entities, i.e. the Web Server is *Located* in the IT Building, and there be will a wide diversity of such linkages. Given the importance of these linkages, to the role of the model, they are themselves classified as security entities i.e. ***Relationships***. Hence the linkages, or relationships can be structured into classes and sub-classes, with each class and sub class given an object identifier. Such linkages can be stored as a simple tuple: < Linkage OI, Incident Entity OI, Target Entity OI>, represents a linkage between two entities, similarly linkages involving three or more entities can be unambiguously defined.

For example, the relationship

Server A is *located* in *Building B* can be represented by the tuple <5.1.1.2.1.3, 1.1.1.3.2, 2.1.3.2>. Where

- 5.1 **Relationships between two entities**
- 5.1.1 incident entity is a ***Systems*** (1)
- 5.1.1.2 target entity is an ***Environment*** (2)
- 5.1.1.2.1 relationship class is ***Environment/Locations*** (ID = 1)
- 5.1.1.2.1.3 particular Location Link (ID = 3).
- 1.1 ***Sytems/Hardware***
- 1.1.1 Computing Hardware (ID = 1)
- 1.1.1.3 Server Class (ID = 3)
- 1.1.1.3.2 Server A (ID = 2).
- 2.1***Environment/Location***
- 2.1.3 HQ Site (ID = 3)
- 2.1.3.2 Building B (ID = 2).

The model entities, attributes and relationships can provide an overview of the current systems, e.g. major platforms, the major components of such platforms: networks, computing systems, users, information assets, the sites where the platforms are located, the services they

depend upon etc. The model can also be refined with increasing level of detail, e.g. the sub-networks that form the major networks etc.

The relationships can be employed to facilitate cross-references between documentation. For example, the ***Procedures*** Class can refer both to internal and external documentation. Hence chapter and paragraphs of standards, and security manuals may be given object identifiers. A *Compliance* relationship, between paragraphs in internal security manuals and corresponding paragraphs in BS 7799, would facilitate internal audits.

3.2 Threat Trees

Risk Assessment is concerned with the ultimate effect of intrinsic threats, e.g. fire, loss of external services, international network failures, on business operations (See Fig 2). An important role of the security documentation, and hence the proposed model is to facilitate the tracing of such threat scenarios.

From the work on the model conducted so far, it would appear that the classification scheme, and in particular the classification of relationships, significantly facilitates such threat tracing.

The threat transmission illustrated in Fig 2 is in effect a series of statements along the following lines:

Incident Threat acting on Incident Entity causes Target Threat to act upon Target Entity (Fig 3). For example:

- *Fire* acting upon *Building* causing *Physical Damage* to *Equipment* (located in Building).

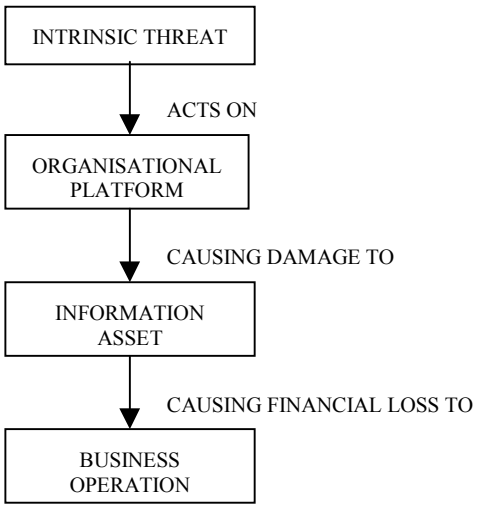

Fig 2.Effect Of Intrinsic Threat On Business Operation

Threats are security entities classified in the model and are classified within the ***Security*** class. The concept of a Threat acting upon an Entity is embodied in a Threat_Entity relationship i.e. the tuple <Threat Entity_OI, Threat_OI, Entity_OI>

The Risk Assessment diagram (Fig 2) may hence be represented as a Threat Tree (Fig 4) where each node represents a Threat_Entity relationship caused by the parent Threat_Entity. Relationship. The Threat Tree recognises that a Threat_Entity may spread to many target entities. At this stage it should also be stated that the

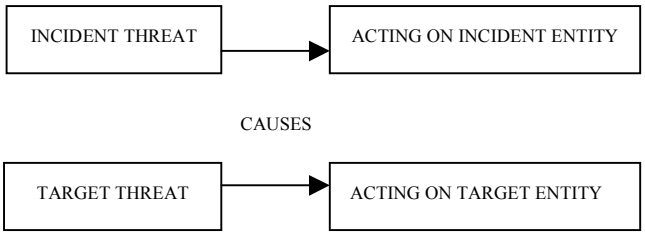

Fig 3. A Threat Entity Causes A Resultant Threat Entity

Threat_Entity transmission need not be restricted to a tree, since a Threat_Entity node can have more than one parent. The model can deal with these situations but for simplicity they are not discussed here.

The concept of threat trees is well known, but a major problem with such trees lies in the effort required for their development. One of the more interesting facets, of the proposed model, is that it opens up the possibility of an automatic construction of threat trees. Consider first manual development of threat trees in the context of the model.

The starting point is the root node, i.e. interest is focused upon the effect of a particular threat acting upon a particular entity, or more simply upon a particular Threat_Entity.

At this stage some security expertise is required to predict the effect of this Threat_Entity on other entities in the organizational database. For example, a security officer would predict that a fire in a room would damage equipment in that room. In effect a ***Relationship*** between Threat_Entities, which are themselves ***Relationships***, is developed. This ***Relationship*** between Threat_Entities is termed a TETE in the model. Hence:

- Incident Threat_Entity <TE_1_OI, Threat_Fire_OI, Room_OI>
- Target Threat_Entity <TE_2_OI, Threat_Physical Damage_OI, Equipment_OI>
- TETE defines the linking of the Threat_Entities < TETE_1_OI, TE_1_OI, TE_2_OI>

Given a database of all possible Threat_Entities and TETEs, developed by a security officer, then threat trees could be automatically produced for any root Threat_Entity as described below:

1. Starting with the root Threat_Entity, TE_1_OI, check all TETE entries <TETE_a_OI, TE_b_OI, TE_c_OI> for those where TE_b_OI = TE_1_OI.
2. Extract TE_c_OI from TETE_a_OI – this is a child node in the threat tree.
3. Repeat 1 – 2 until no more TETEs found.
4. Repeat 1- 3 for the next child node in the threat tree.

This procedure does provide for the automatic development of threat trees, but at a massive cost of manual development of possibly billions of TETEs. Some results of the model, however, suggests that multiple TETEs describing, for example, fires in every room in the organization, and the equipment stored in each individual room, can be replaced by a single TETE using object identifiers with wild cards.

As a simple example of this approach consider the observation that a fire in a building, with OI 2/1/1/1, is could affect all floors of that building, and such floors can be represented with wild card OIs 2/1/1/1/*. Hence we can replace individual TETEs representing the spread to each individual floor with a single TETE along the lines < TETE_a_OI, TE_b_OI, TE-c_OI > where

TE_b_OI is < TE_b_OI, Threat_Fire_OI, 2/1/1/1 >

TE_c_OI is < TE_c_OI, Threat_Fire_OI, 2/1/1/1/* >

Using a comprehensive wild card approach security expertise can be embodied in a minimal number of TETEs, which can then be used to develop automatic threat trees.

The work conducted so far has found that this approach is quite versatile, to mention a few of the findings:

- TETEs can be defined to incorporate the concept of required linking between incident and target entities. For example for a fire in a room to spread to equipment, such equipment must be L*ocated* in that room. This type of condition can be included as an attribute of the TETE
- The transfer of a Threat is not deterministic, it is required that some estimate of the probability of the threat transfer be included as an attribute of the TETEs.
- If wild card TETEs is defined then the probability of a particular threat transfer can be made dependent upon some attribute of the target entity.

TETEs effectively represent security expertise, and are therefore developed by the security officer. Suppose however a large organisation has adopted this model for its various branches, each with its own security database. Given the common means of classification it is clear that TETEs representing common security knowledge can be developed by head office (say) and imported into branch databases.

3.3 Security Design

Security documentation should also play a role in the design of security systems, following the identification of significant areas of risk.

The threat trees provide an insight into the path from an intrinsic threat to an undesirable business impact. Having identified such a path, as a priority security task to be addressed, the role of the security design is to reduce the probability associated with this path. Consider the threat tree illustrated in Fig 4, it can be considered that additional security is required to reduce the probability of the three transfer Threat /Entity 1 – Threat/ Entity 1.2 and / or Threat /Entity 1.2 – Threat/ Entity 1.2.1.

The security measures, physical or procedural, to be deployed clearly depend upon the nature of the TETE linking the nodes of the tree. In effect, the role of the countermeasure is to reduce the attribute of the TETE describing the probability of the threat transfer.

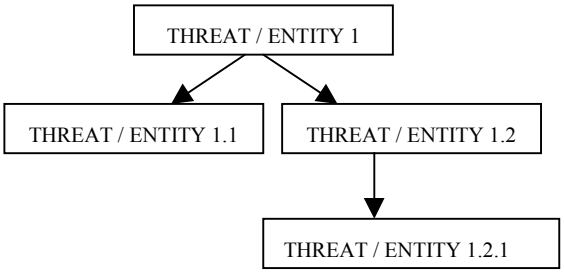

Fig 4. Threat Tree

The threat trees can thus play an important role in security design, inasmuch as they help to define the type and placement of the countermeasures.

The RDR included the concept of Threat Countermeasure Diagrams (TCD) to describe that aspect of security design concerned with the effectiveness of countermeasures, and such diagrams have been incorporated into this model.

The TCD is based upon the concept that countermeasures are themselves subject to threats that can either result in the countermeasure being bypassed or rendered ineffective. Threats to countermeasures are countered by additional countermeasures. For example, it is well known that firewalls are vulnerable to illicit reconfiguration, and must be protected by effective access control. Threat Countermeasure Diagrams are trees of countermeasures designed to ensure the security effectiveness of the root countermeasure.

TCDs like TETEs represent security expertise, since they demonstrate the effective deployment of countermeasures. Hence given acceptance of the classification scheme they can be imported into databases. Interestingly the use of object identifier wild cards seems to allow a TCD to be customized to its environment. Hence it would appear to be possible for a generic imported TCD to take account of local conditions.

4 Conclusion

The information security environment has undergone radical changes over the last decade. Organisations are now highly dependent upon the effective operation of their information systems, and these systems have become complex and highly vulnerable to external influences. Hence effective information security risk management is now a vital component of an organisation's viability.

Such risk management has also been impacted by the escalation of system complexity coupled with the increasing vulnerability and strategic importance of the information systems. Effective risk management, in turn relies upon reliable and timely risk assessments.

The cost of risk assessment exercises increases sharply with system complexity, and a major component of such costs lies in the collection of the wide range of security relevant data. Moreover in an security officers now must provide convincing evidence of the actions taken by the organization, to identify and address the threats to their information systems.

This paper has emphasized the importance of effective security documentation in the above scenario. It has also noted the lack of tools and support to assist security officers in the development of such documentation.

The paper suggests that conventional textual documentation may be replaced by an electronic database and supporting software. Such a database, and associated software tools, must developed around a common information security model and this paper describes such an approach.

It has been demonstrated that a standardised classification of security entities, using object identifiers, facilitates the development and implementation of such a model. The work conducted so far has indicated how the model may be deployed in risk assessment and security design. Moreover the model provides an opportunity for the importation of security expertise from vendors, advisory bodies, etc.

A prototype model based upon Visual Basic has been developed to test the concepts and a more comprehensive Java based software package is currently under development at the City University of Hong Kong.

5 References

[1] Kwok, L.F. (1997): A hypertext information security model for organizations, Information Management and Computer Security, Vol. 5 No.4, pp 138-48.

[2] Anderson AM, Longley D and Kwok LF (1994): Security Modeling for Organizations, Proc. 2nd ACM Conf on Computer and Communications Security, Fairfax VA, pp. 241-250.

[3] IT Baseline Protection Manual Standard Security Safeguards, URL:http://www.bsi.bund.de/english/index.htm

[4] British Standards Institute (1999), BS7799: 1999 Information security management, Part 1 Code of practice for information security management, Specification for information security management systems.

[5] ISO/IEC 17799 (April 2001): Code of practice for information security management URL: http://www.bsi-global.com

[6] Kwok, L.F, Fung, P.K., and Longley, D (2001): Security Documentation, information Security Management & Small Systems Security, IFIP TC11.1/WG11.2, 18th Annual Working Conf. On Information Security Management & Small Systems Security, Las Vegas, USA, pp127-140.

[7] The Directory. CCITT REC. X.500-X.521 ISO/IEC Standard 9594:1993

[8] Federal Information Processing Standards Publication 31. Guidelines for Automatic Data Processing Physical Security and Risk Management, Springfield: National Technical Information Service, June 1974.