DEVELOPING THREAT NETWORKS FOR RISK ANALYSIS OF INFORMATION SYSTEMS

Mark Branagan, Dennis Longley

Information Security Institute QUT
Information Security Institute QUT

m.branagan@isrc.qut.edu.au
Information Security Institute, Queensland University of Technology
GPO Box 2434, Brisbane Q 4001
Phone: (+61 7) 3864 9561
Fax: (+61 7) 3221 2384
d.longley@qut.edu.au
Faculty of Information Technology, Queensland University of Technology,
GPO Box 2434, Brisbane Q 4001
Phone: (+61 7) 3864 1931
Fax: (+61 7) 3864 1801

ABSTRACT

Risk analysis in complex systems must cope with sophisticated threats arising from attackers exploiting security loopholes, or a chain of accidental events triggered by some environmental incident. A security model for complex IT systems, termed the ISM (Information Security Model) stores system and risk information in a database. Associated software uses the database information to develop threat networks and threat countermeasure diagrams as a support tool for the risk analyst. A software tool is under development for this model and this paper explores future developments for the model. In particular, the role of the threat network and its possible extensions are discussed in detail.

KEY WORDS

Risk analysis, security models, countermeasures, threat networks.

DEVELOPING THREAT NETWORKS FOR RISK ANALYSIS OF INFORMATION SYSTEMS

1 INTRODUCTION

Risk analysis originated in a world where the threat and asset were geographically close, e.g. fire and a building. In such a world, an estimate of the risk could be obtained from the information on the probability of the fire and the value of the building. In the early days of mainframe computing risk analysis, threats to the system originated locally and could be treated in similar ways.

In recent years organisations have developed and become highly dependent upon, complex IT systems networked with cooperating organisations and vulnerable to a wide variety of local and remote attacks. The concurrent development of a universal threat highway, termed the Internet, extended the range of potential attackers geographically beyond the range of effective local law enforcement. Recent political events have extended the attacker's motivation from greed to ideological hostility.

Risk analysis in this modern world includes threats that may arise from a remote geographical location, or from an internal employee accidentally or maliciously exploiting some hidden loophole in a very complex system. The attack scenario in such complex systems will commonly exploit system dependencies in tightly coupled environments. Remote environmental incidents arising from severe weather, geological disturbances etc may also set in chain a series of threat events, within this tightly coupled environment, resulting in major disruptions to organisational IT systems.

The objective of risk investigations is to determine the potential business or societal impact, in terms of likelihood and impact magnitude, resulting from external events lying beyond the control of the organisational management. It was considered that a security model of the system, stored as a database with supporting software, may facilitate interactive risk analysis. Hence, the Information Security Model (ISM) [1] was developed for this purpose. This paper represents an extension of that early work and re-examines some of the assumptions made in that model.

The ISM employs an X.500 directory structure to allow for the recording of information regarding security related entities and relationships between these entities, in order to provide a security model. The stored information essentially comprises a local and a general component. The local component describes the local IT system and its defences, from a security viewpoint. The general component includes risk and security information to facilitate a study of the overall system risk scenario and the effectiveness of its security measures.

An essential part of the model is software to create and display interactive Threat Networks and Threat Countermeasure Diagrams. This software applies the recorded information in the general component to the local component data, representing the local IT system, to produce Threat Networks displaying potentially undesirable impacts. Development of Threat Networks may start with specified external threats (effect networks) or specified impacts (causal networks).

The ISM provides for ad-hoc addition of system or security/ risk data. Intrinsic to the ISM is the concept that there should be some abstraction of detail at various stages in order to allow for an overriding systems view. Much of the model concentrates on building causal or effect chains (Threat Networks) to allow the display of the entire chain, between initiating event and final impact.

A detailed description of the model is given in a previous paper [op cit] but an overview of the model is given below to provide a background for the concepts discussed in this paper.

2 OVERVIEW OF INFORMATION SECURITY MODEL (ISM) AND THREAT NETWORKS

2.1 Threat Networks

The ISM is intended for the progressive development of a comprehensive risk/security model for a complex IT system. It is designed to allow a top down interactive approach to risk studies and the importation of external security and risk expertise. Hence, the local IT system view, combined with security knowledge from a variety of external sources, is used to generate causal networks displaying potential threat paths from initiating events to their outcomes. Associated Threat Countermeasure Diagrams [2] provide information about the defence systems designed to minimise the probability of specified threat propagations in the Threat Networks.

Threat Networks consist of a series of nodes starting from some initiating events causing Threat Propagations to consequential events, as determined by the information describing the local IT system and the risk information contained in the model, until eventual impacts are attained. The nodes in a Threat Network are events analogous to the event definition in AS/NZS 4360:2004 [3].

A Threat Network (TN) thus comprises Threat Event (TE) nodes linked by Threat Propagation (TP) branches. Threat Propagation definitions comprise an Incident and a Target Threat Event. The TN development software, contained in the model, selects the initial Incident TE, and seeks those TP definitions containing this Incident TE. When such a TP definition is found, its Target TE is entered as the next node in the Threat Network, and the TP is recorded against the tree branch; the search for further TPs is then continued. The newly entered nodes become Incident TE nodes for the next phase of the search. The search ends with a set of impact nodes that provide no branches for consequential nodes (See Fig 1).

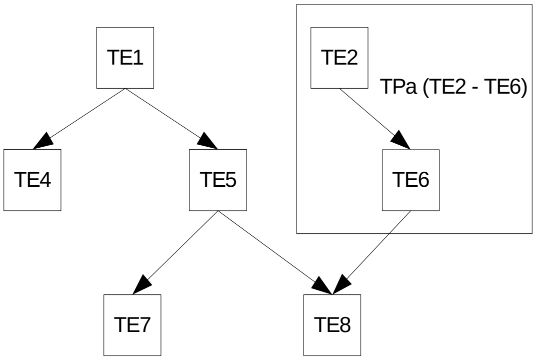

Figure 1 Threat Network

Threat Events (TE) represent a Threat acting on some entity in the model, e.g. Fire in Building, Virus in Application Software etc. These events are atomic in the present model (they occur as a single step and are not compound). TEs may cause consequential TEs, e.g. Fire in Building Causes Damage to Hardware as described by Threat Propagation (TP) definitions.

The TPs recorded against Threat Network branches relate to specific entities, e.g. *Fire in Room 2 Causes Damage to Accounting File Server*. However, the TP definitions stored in the ISM generally take the form of generic TPs providing relationships between model entity types e.g. *Fire in Building Causes Damage to Hardware* and the TN development software matches the specific entities in the TE nodes to the generic entities in the stored TP. The tree structured directory model of the ISM facilitates the matching of generic TPs to specific TPs, since entity types, used in the generic TP definitions, are recorded as parents of specific entities in the model.

To summarise, a Threat Propagation (TP) definition encompasses the causal linkage between Threat Events (TE). A TP definition contains two TEs, an Incident TE which is the initiating event, and a Target TE which is the consequence of the initiating event. TP definitions thus comprise a pair of Incident and Target TEs and may have both a generic and a specific form. The generic TP contains two TEs, with entity types, providing a template for generation of specific TPs.

Threats are propagated from one entity to another because the two entities have some form of relationship. Hence the Fire in Room 2 causes damage to the Accounting File Server because the Server is located in that room. Generic TPs must at least therefore store the relationship, e.g. Located between the Incident and Target TE entities, in addition to the generic Incident and Target TEs. This paper explores the nature of generic TPs in some detail.

2.2 Countermeasures

A countermeasure is represented, in the ISM, as some entity countering a TP branch in a Threat Network. The countermeasure is implemented to reduce the probability of one or more specific Threat Propagations. The countermeasures may themselves be subject to threats and commonly supplementary countermeasures are employed to counter such threats, e.g. Physical Access control for Firewalls. Threat Countermeasure Diagrams (TCD) [op cit] are used in the model to display the rationale of countermeasures and associated supplementary countermeasures structures. The Threat Networks and Threat Countermeasure Diagrams, provided by the model and its software, thus provide an interactive graphical view of the IT system risk scenarios and its defences.

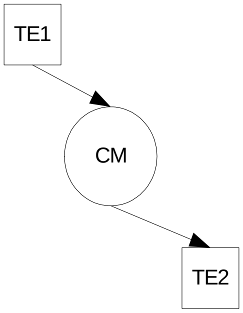

Figure 2 Role of a Countermeasure

3 ADVANCED ASPECTS OF THREAT NETWORKS

3.1 Overview

The development of an ISM software tool has provided a greater insight into the development of risk scenarios, indicating some potential areas for ISM extension and refinement. The following sections consider the mechanisms of threat propagation in more detail and discuss the development of threat networks for scenarios in which the network itself changes with time or events.

3.2 Detailed Consideration of Threat Propagations

3.2.1 Conditions for Threat Propagation

Threat Propagations are the cornerstones of Threat Network development, since Threat Events often have outcomes representing some consequential Threat Events. In some instances, the threat propagation may require a number of concurrent Incident Threat Events to cause a single outcome, as discussed below (See 3.2.3).

As a starting point, the discussion is limited to a single Incident Threat Event resulting in a Target Threat Event. A general definition of a simple Threat Propagation is *An Incident Threat Event (Incident Threat acting upon Incident Entity) will cause a Target Threat Event with some probability "p", if certain conditions of the Incident and Target Event entities are satisfied.* Some Threat Propagations may be seen to be self-evident; the development of others may require significant security expertise. The development of a comprehensive library of Threat Propagation definitions will be a major task and the ISM effectively codifies such security expertise so that such centrally developed TP definition libraries may be downloaded into local ISMs.

At the most elementary level, the condition for a TP may take the form of a specified relationship between the incident and target entities. For example, *Fire in Room 2 Causes Damage to Accounting File Server* implicitly requires that the item of hardware should be located in the room. This in turn implies that the local ISM model uses a *Located* relationship to specify the location of hardware items. The generic TP definition must therefore contain both the generic Threat Events *Fire in Location* and *Damage to Hardware*, together with the condition: the Incident Entity *Location* and the Target Entity *Hardware* share a *Located* relationship.

The current version of ISM software accommodates only such simple conditions for a Threat Propagation. The condition is thus constrained to relationships between the entities contained in the Threat Entities. The incident or target entities may themselves be relationships, but with the current version the actual entities in the relationships are not considered in the threat propagation.

To demonstrate some additional aspects required in a Threat Propagation definition, consider transmission of malicious code transmitted via a network to individual computers connected to that network. If we imagine that a security loophole, rendering the Operating System vulnerable to the virus, exists in a theoretical operating system "OSZ", and that a patch is available for the flaw, these facts need to be captured for the TP definition. Two Threat Events would be involved, *Malicious Code Transmitted on Network* and *Malicious Code Installed on Computer*. The simplest TP definition would be *Malicious Code Transmitted on Network* causes *Malicious Code Installed on Computer* where *Computer is Connected to Network.* With this simple definition, the TN indicates that malicious code on the network is transferred to every machine connected to the network (See Fig 3).

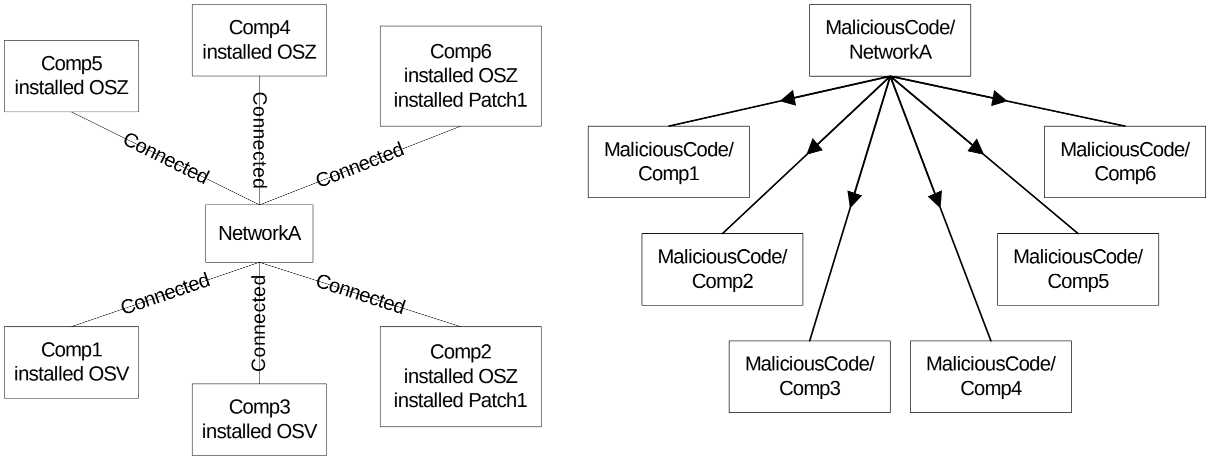

Figure 3 Network Setup and Threat Network

There are obvious deficiencies in this Threat Network. Since the security loophole only exists in OSZ, only those machines with this operating system installed should be included in the Threat Network. Thus, the example Threat Propagation definitions should be expanded to *Malicious Code Transmitted on Network* causes *Malicious Code Installed on Computer* where *Computer connected to Network* AND *Computer has installed OSZ.* The Threat Propagation definition thus includes the specific vulnerability exploited by the malicious code. Moreover, assuming the existence of a patch for the operating system which removes the software flaw, then computers with the patch installed are no longer subject to the malicious code. Including this fact in the threat propagation would require a final form such as; *Malicious Code transmitted on Network* causes *Malicious Code Installed on Computer* where *Computer connected to Network* AND *Computer has installed OSZ* AND NOT *Computer has installed Patch1*.

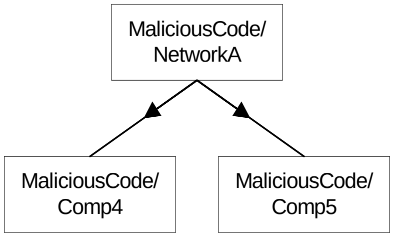

Figure 4 Final Threat Network after Improvements to TP Definitions

Figure 4 demonstrates the result of the improved Threat Propagation definitions. This example demonstrates the first proposed extension to the condition model for Threat Propagations, i.e. the ability to combine conditions using Boolean operators.

Conditions may also be applied to the Incident Entity e.g. suppose the malicious code requires a certain communication protocol in order to transfer itself across the network. The threat propagation from above would then become, *Malicious Code Transmitted on Network* causes *Malicious Code Installed Computer* where *the Network Protocol is TCP* AND *Computer connected to Network* AND *Computer installed OSZ* AND NOT *Computer installed Patch1*.

The model permits attributes to be supplied for entities in a general <TAG> <VALUE> format, and this permits a simple but somewhat rigid implementation of the conditions. Hence, in this case the TP definition would include the conditions:

- <PROTOCOL> <TCP> attribute in the network entity;
- <OPERATING SYSTEM> <OSZ> attribute in the target computers.
- <PATCH> <PATCH1> not in the attributes of the target computers.

Such an extension to Threat Propagation conditions would provide for a more realistic Threat Network minimising the number of false positive branches and invalid subtrees.

3.2.2 Relationships as Event Entities

Threat Propagations are not necessarily restricted to atomic entities. As mentioned above a Threat Propagation definition may also include relationships for the incident entity, target entity or both. Hence, a further extension to TP definitions would include conditions on the actual entities contained in these incident or target relationships. This extension would provide for a less rigid implementation of the conditions described above (See 3.2.1). For example, it would be then possible to define the target entity as the relationship *OSZ Installed on COMPUTER*, rather than use opaque attributes.

3.2.3 Multi-vector Attacks

The Threat Networks discussed so far represent OR conditions for the Incident Threat Events inasmuch as a node with two Incident TEs represents a Threat Event that will occur if either of the Incident TEs eventuate. In some cases, however two or more TEs must occur simultaneously, or at least within a specified time frame, for the Target TE to eventuate (See Fig 5).

For example In Fig 5 (a), the hardware in a building may be damaged by either a fire or flood. However, a file server and its backup server would both have to fail for loss of the file service (See Fig 5 (b)).

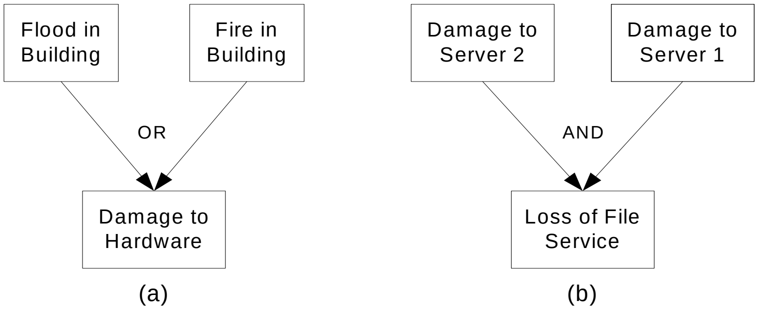

Figure 5 OR and AND Conditions for the Threat Propagation

Figure 6 describes an example where a web site is provided by three web servers with some mechanism of sharing requests between them. If a single server fails then the web site as a whole will not fail. There may be however some impact on the availability of the service. Perhaps there is a drop in the number of requests that the web site can simultaneously serve or some increase in response time from the web site. This situation raises the need for modelling two general concepts. Firstly, the failure of any single server does not indicate the failure of the service as a whole. Secondly, that the progressive failure of the servers leads to some degradation of the service as a whole, resulting in, some drop in availability of the service.

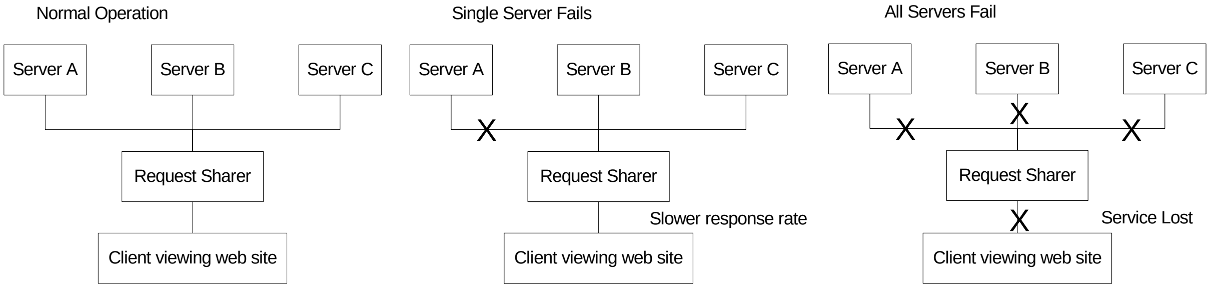

Figure 6 Servers with Single Failure and Total Failure

It can be said that the web site fails when Server A, Server B and Server C can no longer communicate with the request sharer as in the rightmost diagram in Figure 6. This could happen in multiple ways such as all three servers failing, power being lost to all the servers or the network cards in all three servers failing for example. Therefore, generation of Threat Networks needs to be able to account for such situations. This will require two separate functionalities:

- Determination of the set of Incident Threat Events that must eventuate before the Target Threat Event will be caused.
- The trigger action to produce the Target Threat Event upon the occurrence of the total set of Incident Threat Events as determined above.

It would appear that the first functionality implies the definition of multiple linked Threat Propagations, and these definitions themselves fall into two categories:

- Multi-vector attacks where the total set of Incident Threat Events is known at the time of construction of the TP definition.
- Availability type attacks where a multiple of similar Incident Threat Events is required but the number of such events is not known at the time of construction of the TP definition.

The first case is relatively straightforward since the definition of the TPs merely requires some linkage between the members of the set.

Loss of availability is a good example of the second set of TP definitions. Here the loss of availability of a service will depend upon the number of redundant entities providing the IT system for that service.

One approach to this problem may be to use causal Threat Network searches. So far, the description of Threat Networks has concentrated on the use of TPs to determine the effect of an initiating Threat Event. However, a minor enhancement to the search algorithm can use TPs to determine the set of Incident Threat Events that could cause a specified Target Threat Event.

Hence, it is possible to develop Threat Networks commencing with some undesirable impact and leading to the total set of Threat Events that could cause that impact. Any Loss of Availability Threat Event entered as an impact, or arising as an intermediate Threat Event then would automatically cause the search for the complete set of corresponding Incident Threat Events (e.g. loss of availability of redundant servers).

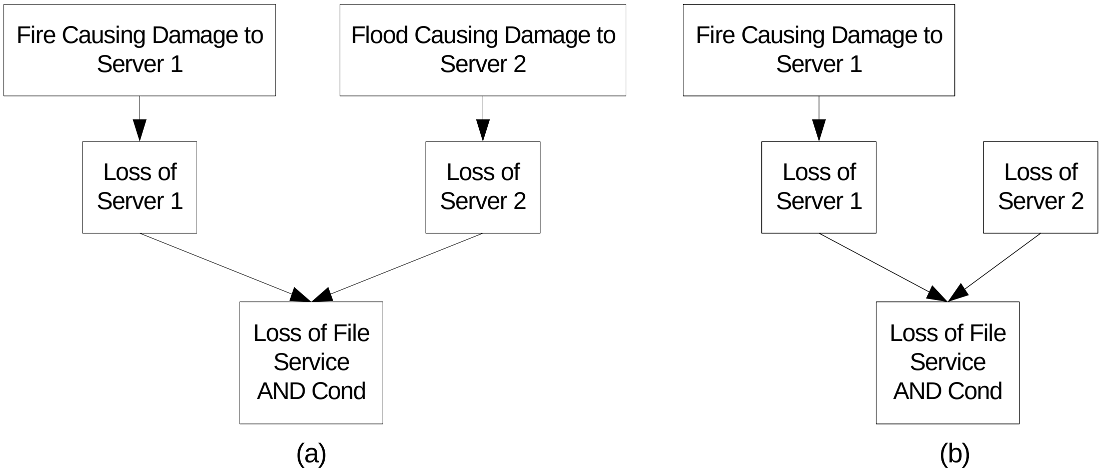
Fire Causing Damage to Server 1
Flood Causing Damage to Server 2
Loss of Server 1
Loss of Server 2
Loss of File Service AND Cond
(a)
Fire Causing Damage to Server 1
Loss of Server 1
Loss of Server 2
Loss of File Service AND Cond
(b)

Figure 7 Threat Networks developed bottom up, i.e. the eventual impact is specified and a search is conducted for causal TEs

Consider the two causal networks used to determine the set of Threat Events leading to the loss of availability of a File Service (See Fig 7). It is clear that a simultaneous fire and flood could cause a loss of File Service (See Fig 7(a)).

Suppose, however, in another IT system Server 2 is not prone to flood damage; this would produce the threat network of Fig 7 (b) where is there no indication of a threat causing Loss of Server 2. Given the AND condition in the Loss of File Service node it would appear that no threats leading to the Loss of File Service have been found for this particular IT system.

To summarise it would appear that causal Threat Networks provide an approach to multi-vector threats and loss of availability attacks, i.e. working up from some specified undesirable impact. The definition of the TPs requires some indication of the AND condition for the multiple incident TPs (moreover this AND condition is displayed in the Target TE node).

For multi-vector attacks, the complete set of Incident TPs is specified in the TP definition, and each of these TPs would be automatically selected in the causal search. In the case of loss of availability attacks the TP definition indicates the AND condition and the causal search automatically seeks the redundant servers (say).

It should be noted that additional complications may arise when a mixture of AND and OR conditions occur. For example, the loss of availability of the File Service may arise from Loss of All File Servers OR loss of network availability.

3.2.4 Probabilities

Risk assessment is necessarily probabilistic; unfortunately in IT system risk assessment there is a dearth of historical data required assign the component probabilities used to compute the final outcome. At best, the risk assessor can say - *If you assign these component probabilities then the impact probability can be computed.*

The ISM approach at one level provides a Threat Network indicating the path of threat events that could lead to some specified undesirable outcome. Long paths including some highly unlikely threat propagations may be eliminated so that other more probable paths receive due attention.

Probabilities have been included in the current ISM Threat Networks at a relatively unsophisticated level:

- Probabilities of initiating threat events.
- Probabilities of threat propagations.
- Probability of intermediate threat events.

An initiating threat event probability must be related to some timescale. If a threat network has a single initiating threat event and on average, this occurs once in p years, then entering the initiating threat event probability as 1/p suggests that the threat network may be assigned a one year time scale, i.e. all the computed probabilities are given in terms of expected occurrences per year.

If there are two or more initiating threat events and all the threat propagations are the OR type (See 3.2.3) then the model can compute the probabilities of intermediate TN nodes, on the assumption that these events are independent.

The situation becomes somewhat more complex if AND type threat propagations are present. Suppose an initiating Fire event causes downtime of File Server 1 (See Fig 7 (a)) and the expected recovery time is one day. A loss of File Service will occur if a second initiating Flood event causes a loss of Server 2 within that downtime period. In this case, the threat network should be assigned a one day timescale and the initiating probabilities correspondingly downscaled.

If it is assumed that threat propagations are virtually instantaneous (See 3.2.5) then we may include probabilities in each specific TP definition. However, the question arises on the assignment of probabilities to generic TPs, i.e. when a generic TP is instantiated to a specific TP should the same probability be assigned. Given the enormous problem of assigning probabilities in IT systems one may be tempted to take the simplest solution unless there is some compelling reason to be more precise. It may, for example, be possible to assign local probabilities to specific TPs based upon some attributes of the Target and Incident entities, e.g. a wooden workstation may be more likely to suffer damage from a fire than an asbestos one.

The calculation of the probabilities for Target Threat events depends not only upon the probabilities of the Incident Threat Event and that of the Threat Propagation, but also upon the OR/AND nature of the Threat Propagations (See 3.2.3).

The ISM has not produced the Alchemist Stone for probabilistic risk assessment. However, the interactive graphical nature of Threat Networks introduces more transparency into the computation of impact probabilities. Given some good graphical representation of Threat Event and Threat Propagation probabilities, the model should at least allow the risk assessor to identify critical events, and propagations, and to explore the sensitivity of the final outcome to those identified critical parameters.

3.2.5 Time and Threat Networks

At present, the ISM Threat Network is completely static. All events are considered to take place instantaneously. Countermeasures are assumed to fail or succeed instantaneously, threat propagations take place with no pause and threats act immediately. This is obviously not true in real life. In the real world, events are set in a temporal sequence; the following scenario demonstrates the limitations of the current ISM.

Consider a UPS system where the battery time available to run a particular computer is 20 minutes. The UPS will deal with any power failure less than 20 minutes, however after 20 minutes the system will have to shut down. At this point, any consequences which have been held at bay by the UPS will eventuate. This scenario is actually more complicated. The UPS will have a recharge time. After an extended operation, the UPS will be capable of less than 20 minutes of running time, until its recharging is complete.

This type of time dependent relationship is one that is fairly common in security incidents. Some other examples include the time taken to illicitly decrypt a protected message, to break into a safe, for a denial of service attack to subside or to restore a file from backup. Impacts may also have time implications, e.g. the acceptable period for a system downtime.

The variations of a threat network in time, however, are not likely to be structural but would appear to be limited to probability variations. An operational UPS will result in a reduction in the probability of a loss of service event; subsequent failure of the UPS will simply increase that probability. Hence, there is a case for the explicit display of probabilities in threat networks. For example, high probability nodes may be highlighted, and high probability branches (i.e. high probability threat propagations) may be shortened.

Hence, one may contemplate an interactive session when a particular threat event node is selected, e.g. loss of power causing activation of the UPS, and the dynamic changes of the network node probabilities are displayed for the subsequent time period.

4 CONCLUSIONS

The ISM was based upon some simple structures and concepts. Subsequent experience with the associated software development and trial projects has indicated its potential as a risk simulator. Such projects include its use in standards conformance testing, and extension from IT systems to the more general area of Critical Infrastructures. This paper has described the limitations of the current implementation but also indicates that the model can be extended into more complex risk scenarios.

5 ACKNOWLEDGEMENTS

The authors would like to thank Mr Robert Dawson in particular for his assistance in developing the ISM.

6 REFERENCES

1. Kwok, L.F. and D. Longley. *Security Modelling for Risk Analysis*. in *19th IFIP International Conference on Information Security*. 2004. Toulouse: Kluwer Academic Publishers Group.

2. Caelli, W., D. Longley, and A. Tickle. *A Methodology for Describing Information and Security Architectures*. in *IT Security the Need for International Cooperation, Proceedings of the IFIP TC11 8th International Conference on International Security '92*. 1992. Singapore: Elsevier Science Publishers.

3. Standards Australia, *AS/NZS 4360:2004 Risk Management*. 2004, Standards Australia, Standards New Zealand: Sydney, Wellington.