Developing Electronic Trust Policies Using a Risk Management Model

Dean Povey

Security Unit, Cooperative Research Centre for Enterprise Distributed Systems**,
Level 12, S-Block, Queensland University of Technology,
Brisbane Qld 4001, Australia,
povey@dstc.edu.au

Abstract. Trust management systems provide mechanisms which can enforce a trust policy for authorisation and web content. However, little work has been done on identifying a process by which such a policy can be developed. This paper describes a mechanism for developing trust policies using a risk management model, and relates this to a conceptual framework of trust. The process uses an extended risk management model that takes into consideration beliefs about the principals being trusted and the impersonal structures and systems involved.
The paper also applies the extended risk management model to a hypothetical case study in which an individual is making investments using an electronic trading service.

1 Introduction

Regardless of the strength or robustness of a given security mechanism, its effectiveness is limited without the existence of trust. Security protocols, cryptographic devices and digital signatures rely on the ability to trust either one or more parties, mechanisms or equipment to be sure that the assets they protect remain safe.

In the physical world we derive much of our notions of trust from the tangible nature of things. For example, we perceive the information in a book to be worth reading because we know that it costs a lot of money to print a book, because the logo on the side shows that it has been reviewed by a publisher of repute, and often because a library has thought it worthwhile enough to stick it on their shelf. Similarly, we are convinced by the stability and trustworthiness of a bank, because the difficulty of licensing a fraudulent organisation and the cost of setting up branches, ATM networks and marketing etc, would make it prohibitively expensive.

However, the shift toward e-commerce means that we can no longer infer trust from physical, tangible things. We need to rethink our approach to trust

** The work reported in this paper has been funded in part by the Co-operative Research Centre Program through the Department of Industry, Science & Tourism of the Commonwealth Government of Australia

so that we can rely on the information and actions of people in a virtual world, with the same degree of confidence that we do in the real world.

Trust management systems such as PolicyMaker[1], KeyNote[2], and REFEREE[3] provide mechanisms that can enforce a trust policy for authorisation and web content. However, little work has been done on identifying a process by which a trust policy for such systems can be developed.

This paper describes a mechanism for developing trust policies using a risk management model, and outlines a hypothetical case study to illustrate the usefulness of such a scheme.

2 Risk Management

Risk management is the total process of identifying, controlling, and minimising the impact of uncertain events [4]. The Common Criteria [5] outlines a model for relating different elements of the risk management process, which is given in figure 1. In general, risk management for information security involves the following process:

1. Identify the assets to be protected, the threats to these assets, and the expected impact if those assets are compromised.
2. Identify the vulnerabilities or weaknesses which can lead to these threats arising.
3. Analyse the risk (i.e. the likelihood and consequences) of the vulnerabilities leading to these threats being exploited.
4. Determine whether to accept or treat the risk.

Risk is treated using countermeasures which seek to reduce either the likelihood or consequence of a risk, or defer the risk to some third-party (e.g. insurance). Implementing a countermeasure has a cost associated with it, which must be balanced against the expected utility of implementing the measure. Countermeasures may also expose additional risks, or retain residual risk which must be considered in the risk management process.

Risk management is well understood, and numerous standards and methodologies exist to describe the process (e.g. [6][7][8]). Integrating risk management into the trust management process is therefore useful, as it will enable us to leverage off this existing body of work.

3 Trust

To integrate trust with risk management, it is necessary to provide a framework by which different aspects of trust can be described and related. One of the more comprehensive frameworks for trust was developed by McKnight, Cummings and Chervany, and results from a survey of sixty papers across a wide range of disciplines[9][10]. McKnight et al's model provides a classification system for different aspects of trust, as well as a system for showing how trust can influence behaviour and defines the following constructs:

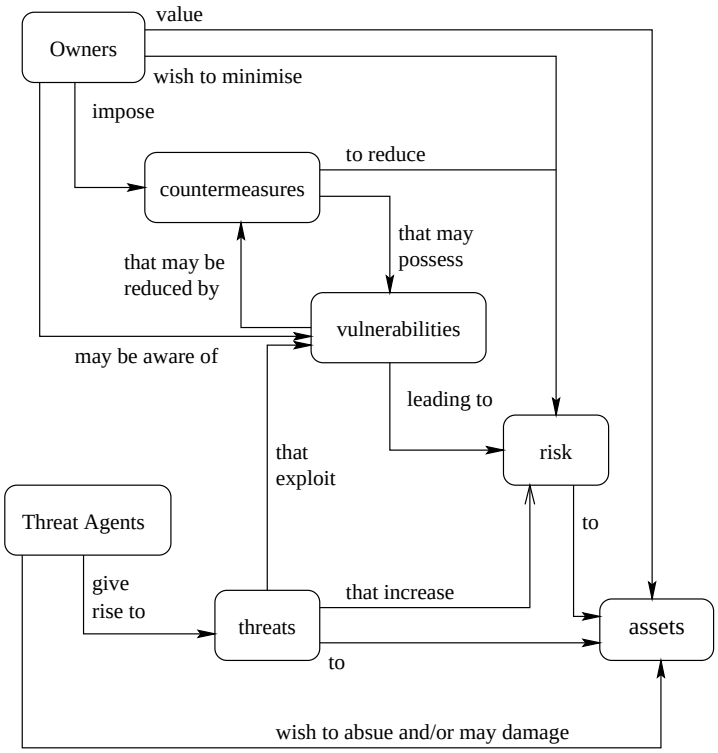

Fig. 1. Security concepts and relationships from the Common Criteria

Trusting behaviour the extent to which one person voluntarily depends on another person in a specific situation with a feeling of relative security, even though negative consequences are possible. This construct is in effect describing the "act" of trusting, and implies acceptance of risk (negative consequences) by the trusting party.

Trusting intention the extent to which one party is willing to depend on the other party in a given situation with a feeling of relative security, even though negative consequences are possible. A trusting intention usually leads to trusting behaviour. Trusting intentions relate directly to the security policy which determines how entities in the system are trusted. A trusting intention essentially specifies a willingness to trust a given individual in a given context, and implies that the trusting entity has made decisions about the various risks and benefits of allowing this trust.

Trusting beliefs the extent to which one believes and feels confident in believing that the other person is willing and able to act in the trusting party's best interests. A trusting intention will be largely based on the trusting par-

ties cognitive beliefs about the other person. McKnight et al describe four categories of trust belief:

1. Benevolence - the belief that a person cares about the welfare of the other person;
2. Honesty - the belief that a person makes agreements in good faith;
3. Competence - the belief that a person has the ability to perform a particular task; and
4. Predictability - the belief that a person's actions are consistent enough to forecast what they will do in a given situation.

Trusting beliefs characterise the information by which we make our trusting decision about a given individual. They may be based on evidence, recommendations from third parties (which themselves must be trusted), and often by simple intuition. We can think of trusting beliefs as being the measures by which we will determine whether a given entity should be trusted given a specific risk profile. It should be noted that not all beliefs need to be strong in order to trust an individual in a given context. In business transactions, the issue of benevolence is rarely important (although the presence of malevolence may be) when compared to the issues of honesty, predictability, and most importantly competence. Also, some beliefs are easier to be confident about than others. It is usually simpler to obtain a measure of an organisations competence (by accreditation and recommendations), and predictability (by past dealings); than it is to obtain a measure of their benevolence and honesty.

Like trusting intentions beliefs may also be specific to a context (e.g. belief in the competence of a lawyer to write contracts, does not extend to their competence to perform neurosurgery).

As we shall see it is trusting beliefs which are the most important to ascertain, as they will determine the confidence by which we establish our trusting intentions.

System trust the extent to which one believes that proper impersonal or institutional structures are in place to enable one to anticipate a successful future endeavour. An important difference between system trust and trusting beliefs, is that while trusting beliefs relate to the attributes of another person whom is being trusted, system trust relates to the actual system/infrastructure under which the trusted action is taking place.

System trust is important, as it provides stability to our interactions with people and organisations. Legal and regulatory systems provide punitive mechanisms to discourage malicious behaviour, and accreditation and certification schemes provide systems which allow us to evaluate an organisations competence. Like trusting beliefs, system trust is a critical component of determining a trusting intention.

Dispositional trust the extent to which one has a consistent tendency to trust across a broad spectrum of situations and persons. A person may have dispositional trust because they either believe in the general good nature of people, or they believe that they will achieve better outcomes by tending to trust people.

Situational trust the extent to which one intends to depend on a non-specific party in a given situation. Situational trust is related to dispositional trust in that it is a general intention. However, it is differentiated by the fact that where dispositional trust refers to a broad spectrum of situations and persons, situational trust is related only to a specific situation.

Belief formation processes The process by which new beliefs are developed and integrated into our schema about the world.

These constructs do not exist in isolation, but have well-defined relationships between them. We can clearly see that a trusting behaviour relies on the existence of a trusting intention, which in turn is created through the existence of one or more of trusting beliefs, system, dispositional or situational trust. Figure 2 shows the various constructs and their dependencies.

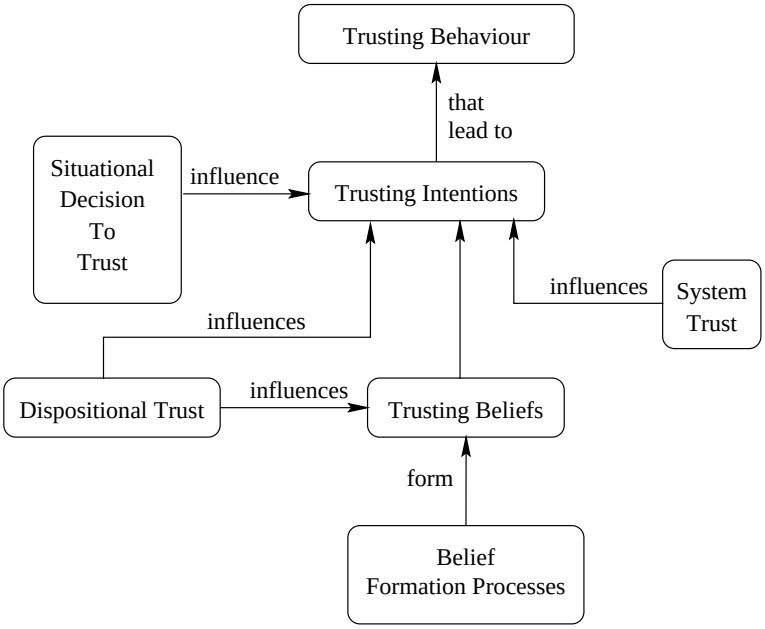

Fig. 2. Related Trust Constructs

McKnight et al's conceptualisation of trust as multi-dimensional is both powerful and compelling. It also goes some way to explaining the difficulty that researchers in many disciplines have encountered in the formulation of a single broad definition of what trust is. In addition, their wide consultation of literature from many disciplines including management, communication, sociology, social psychology and economics, positions their model within a context that it is sufficiently broad to categorise most definitions of trust.

4 An Extended Risk Model

To extend the risk model to encompass trust, it is important to see the goal we are trying to achieve in developing a trust policy. A risk management process seeks to identify risks, and to determine whether those risks should be treated or accepted. Trust management on the other hand, seeks to identify the circumstances under which we are prepared to accept risks that may be exposed by relying on certain entities. The key to merging these two concepts is to focus on risk as the common element. We can see that the definition for trust management can be related to the decision about risk acceptance/treatment. In effect, trust becomes a risk treatment option, i.e. you are prepared to accept risks if you trust the entities that can expose them.

This fact is intuitively obvious to most people. The more someone is trusted, the more we feel we can rely on them, and consequently the more risk we expose ourselves to. When we talk about levels of trust, we are really discussing the level of risk that we are prepared to accept for relying on a trusted entity.

4.1 Relating Trust Policies to McKnight et al's Model

The constructs described in section 3 provide a vocabulary for describing how trust is formed. By combining this process with the risk management process we can show how trust policies can be captured from the environment using a structured process.

In McKnight et al's model, the trusting intentions form the trust policy, which is essentially a statement of the conditions under which we are prepared to trust a given entity. As noted in section 3, these intentions are formed from a number of sources: our dispositional trust, our beliefs about the entity we are trusting, how we trust the systems which we look to to support and protect us, and our tendency to trust in the given situation. As described, it is important to consider the risks of the behaviour of entities that we intend to trust. However, it is also important to consider the utility or value of trusting this entity, as this can considerably alter the decision to accept or treat a risk, or to not allow the behaviour to occur.

On further analysis, we see that there is also important interactions between components of the trust framework that we must consider. One of the most important elements of forming the trusting intention is the existence of trusting beliefs about an entity. These are important, as they are the only input into the trusting intention decision which is specific to a given entity. McKnight et al's model identifies a *belief formation process*, which is an iterative mechanism that uses information and experience gathered from the environment to form one or more trusting belief about an individual. In this extended risk management model, the information that is input into this process is called a *trust metric*. Trust metrics contribute to our understanding about the four trusting beliefs (competence, predictability, honesty and benevolence); and include:

- information based on previous experience;

- recommendations from third parties;
- certifications or qualifications;
- memberships of professional organisations;
- certified histories (criminal records, credit reports etc.); and
- brand.

As we can see from this list, the trust metrics themselves can be subject to trust decisions about their accuracy. Thus, the belief formation process is recursive.

Another important observation, is that metrics may have a cost associated with them (e.g. obtaining a credit report may cost money). In developing a trust policy, we must be careful to ensure that the costs of gathering metrics do not outweigh the utility gained from trusting, and that we maximise the value of our metrics, such that the cost reflects the contribution to our understanding of the trusting beliefs.

Figure 3 shows how these constructs relate to form an extended risk model.

4.2 Using the Extended Risk Model for Trust Management

By combining the concepts from risk management with the extended risk model, we can establish the following process for establishing a trust policy:

1. Identify the entities and situations you want to determine a trust policy for. This allows the establishment of *trust contexts*, which encapsulate the security context within which trust decisions will be made. Note that such a context should include both all probable trusted entities and threat agents.
2. Identify the assets to be protected within this trust management context, the threats to these assets, and the expected impact if those assets are compromised.
3. Calculate the expected utility of trusting entities in the given situations.
4. Identify the vulnerabilities or weaknesses which can lead to these threats arising.
5. Analyse the risk (i.e. the likelihood and consequences) of the vulnerabilities leading to these threats being exploited.
6. Determine the adequacy of existing countermeasures which may mitigate these risks.
7. Determine the required beliefs and confidences in these beliefs required to trust (or distrust) entities which may expose the given risks.
8. Identify the various impersonal structures or systems which have an impact on the given trust context. Common systems will include legal or regulatory frameworks. Analyse our confidence in these systems to mitigate risks.
9. Identify metrics which will can help make decisions about the required trusting beliefs, and determine the confidence we have in the accuracy of these metrics (in itself a mini trust-management decision).
10. Evaluate the costs of gathering these metrics, and relate this to the expected utility, and their contribution to confidence in the trusting beliefs. Use this evaluation to select the subset of metrics which can be used to establish the trusting beliefs.

11. Using the metrics, establish the beliefs identified in step 7 and determine whether they meet the required confidence levels.
12. Based on this evidence and the levels of system trust, either unconditionally accept a trusting intention for the evaluated entity in the given situation; reject the trusting intention; or treat the risk and reevaluate.

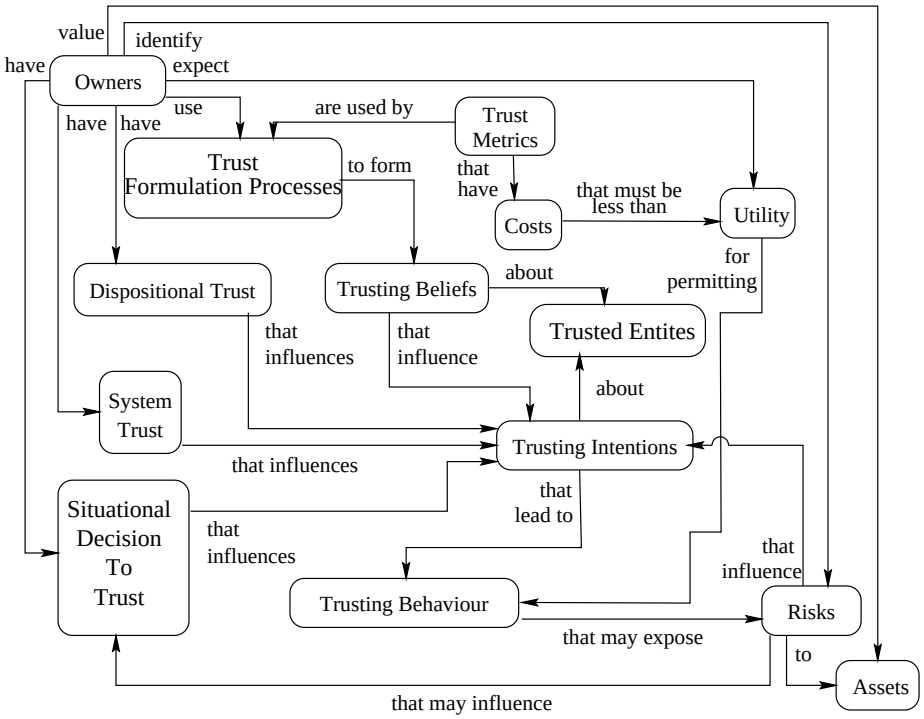

Fig. 3. Extended Risk Model

Trusted entities and threat agents may be either known or unknown. In the case that they are known, then the policy should include the actual measurements for this entity obtained using the trust metrics. In the case that they are unknown, the policy should contain the list of metrics which are required to determine whether an entity should or should not be trusted.

If a trusting intention is rejected, then risk may be treated by a number of mechanisms:

1. Add countermeasures which decrease the risk
2. Defer risk to a third party (e.g. insurance)
3. Increase the required belief trusting confidences by obtaining more or better metrics.

We can see that this process extends the risk management by integrating components of the trust model.

5 Writing Trust Policies

The outcome of the trust management process should be a policy which documents the decisions made. The policy should include:

Trust metrics A list of the metrics used in a trust policy, the trusting beliefs they measure, and their appropriateness for given trust contexts.
Confidence levels A description of the list of qualitative or quanti tative labels that indicate our level of confidence in a given trusting belief.
Trust context policies An articulation of the policy for making trusting decisions for each of the identified trust contexts.

These items are described below.

5.1 Trust Metrics

One important component of the extended risk model is the use of trust metrics. These are mechanisms which can be used to enhance our confidence about certain beliefs. An important thing to note is that the trust metrics themselves need to be trusted, and we will have a confidence level associated with their precision. A trust policy should begin by evaluating the trust metrics which it will use, and providing confidence levels which we have in their measurements in a given context.

When specifying the metrics used in the policy, the policy writer should state:

- the contexts in which that metric is trusted;
- the belief(s) that they measure;
- the confidence in that metric for those contexts in which it is trusted, and how this is measured (NB: metrics can be evaluated using other metrics); and
- the cost of evaluating that metric.

In general, metrics should require close scrutiny as we are exposed to more systemic risk by trusting them.

5.2 Confidence Labels

The policy writer should include confidence labels which may be attached to particular beliefs in the trusting decision. Confidence labels can be either qualitative or quantitative, and are similar to the likelihood measurements which are commonly used in risk management. The confidence label represents the likelihood that the belief it is attached to is correct, i.e high confidence means a high probability of correctness. Figure 4 gives an example of qualitative and

Label	Quantitative	Qualitative
Very Low	A belief with this label has very low confidence, it should only be relied on if the risk is negligible.	$p <= 0.5$
Low	A belief with this label has low confidence, it may be relied on only if the risk is low.	$p > 0.5$
Medium	A belief with this label has medium confidence, it may be relied on for contexts at medium risk	$p > 0.95$
High	A belief with this label has high confidence, it may be relied on in high risk situations.	$p > 0.995$
Very High	A belief with this label has very high confidence, it may be relied on in all situations.	$p > 0.999$

Fig. 4. Example confidence labels

quantitative labels which might be used in a trust policy. In the quantitative descriptions, p is the probability that a belief with that given label is correct.

Confidence levels may be combined to obtain new confidence measures. This is useful when for example a number of metrics are being used to determine the level of a trusting belief. Quantitative metrics can be combined, simply by summing the probabilities (i.e. only one of the metrics has to be correct, for the belief to be correct). Qualitative metrics can be combined either on an ad hoc basis or by using rules to combine levels (e.g. HIGH = 3 × MEDIUM).

5.3 Trust Contexts

These are all the situations and environments that are under consideration for the trust policy. For each trust context, the trust policy should detail:

- a description of the context;
- the risks inherent in trusting entities for a given context;
- the expected utility of trusting entities for the given context;
- a list of the possible trusted and threat agents;
- a list of the beliefs and confidences required to trust/distrust entities in the trust context; and
- a list of the required/available metrics appropriate to establish these beliefs.

Contexts may be included within other contexts, for example a context which covers user access to a web site, may also include a sub-context for privileged access to files. This allows a simple hierarchical organisation of trust policies.

If specific entities are to be trusted for a given context, these entities should be listed along with the rationale for trusting them. Where the policy is specifying criteria for trusting unknown entities, it is sometimes useful to separate out the requirements in terms of the type of entity which is to be trusted. For example, entities could be divided into customers, employees and contractors. The policy writer may wish to express differing levels of required beliefs and confidences in each of these, as there are varying levels of utility for differing classes to exploit threats, and as such varying likelihood of threats occurring.

6 Hypothetical Case Study

In this section, the extended risk management model is applied to a hypothetical case study in which an individual is making investments using an electronic trading service. The case study serves to illustrate the complexity involved in evaluating a given trust decision, as it shows how making one trust decision relies on many other trust decisions. It should be noted that in the following example, some of the steps described in section 4.2 have been consolidated together. The aim is to give a general feel to how a trust policy can be developed using the mechanism, and not to explicitly show how the policy should be expressed.

6.1 Scenario

Bob is a naïve investor, with a small amount of cash to spend. He is contemplating some direct share investments, and so asks his friend Alice who is wise in the ways of sharemarket for her advice. Alice suggests he makes a number of investments, but recommends in particular, a recently listed small Internet company – ComDot.com. She says that she has heard on the grapevine, that this company is likely to do spectacularly well, once it releases the next version of its new Website construction software. Alice also suggests that rather than fork out for brokerage fees, Bob purchase the shares directly from E-Shares, an online brokering firm which allows small purchases using a credit card. Bob contemplates whether to take Alice's advice.

6.2 Trust Management Process

Based on the information from Alice, Bob has to make a number of decisions about whether to invest in ComDot.com shares. Doing this requires a number of trusting decisions, which may also involve gathering information, and determining whether that should be trusted. Following the trust management process outlined in section 4.2, Bob sets about determining his trust policy.

Establishing Trust Management Context The scenario above constitutes Bob's trust management context, i.e. he is making decisions about trust within the context of making a specific decision about buying a certain set of shares using an electronic trading service. There are a number of trusting intentions which Bob must have before he can make this decision:

1. Bob must trust Alice to give good advice about the shares;
2. Bob must trust ComDot.com to conduct their business competently; and
3. Bob must trust E-Shares to respect his privacy, and keep his credit card details secure.

In addition, Bob must also consider threats from the following sources:

– hackers, who wish to steal Bob's credit card details and make fraudulent purchases;

- ComDot.com's competitors who may wish to spread misinformation in order to gain market advantage; and
- marketeers, who may wish to use knowledge of Bob's share purchase as fuel for direct marketing campaigns.

Calculate Expected Utility By purchasing the shares, Bob aims to make at least an 8% per annum return on his investment. By using the online trading scheme he hopes to save up to 20% in brokerage fees.

Identify Assets to be Protected In this scenario, Bob determines that he has three main assets under threat:

1. The cash he is investing (could be lost due to poor investment)
2. His credit card number (there is a threat of disclosure leading to fraudulent transactions on his credit card).
3. His privacy (Bob doesn't want people knowing how he spends his money).

Vulnerability Analysis By analysing his assets and the possible threats, Bob determines the set of vulnerabilities which may lead to those threats being realised.

1. Information Bob uses to make decisions could be inaccurate.
2. Companies which Bob invests in might go out of business
3. E-shares might disclose private information
4. E-shares might disclose Bob's credit card details
5. Hackers might intercept Bob's credit card details over the Internet.

Risk Analysis For each of the above vulnerabilities, Bob identifies the likelihood and consequences of these vulnerabilities causing threats to be realised. Likelihood is measured qualitatively (RARE, UNLIKELY, MODERATE, LIKELY, CERTAIN), and the label UNKNOWN is used where making this judgement is not possible in this first analysis (usually due to lack of knowledge about trust levels). Consequences are also indicated qualitatively with the labels: INSIGNIFICANT, LOW, MODERATE, SIGNIFICANT, CATASTROPHIC). This analysis is summarised in figure 5.

Identify Required Beliefs and Confidences Bob now needs to determine the level of required beliefs in order to accept the risks he has identified. We shall briefly outline these decisions for two of the identified vulnerabilities:

Information Bob uses to make decisions could be inaccurate Given the risk identified, Bob determines that he has to trust the information he receives about given shares with a HIGH degree of confidence (see figure 4). In order to trust the information he receives, Bob determines he has to know that the sources of the information are competent, honest and predictable; and that his confidence in these beliefs must either be HIGH, or the information must be confirmed from other sources, such that the total confidence for each of these beliefs is HIGH.

Item #	Likelihood	Consequences	Comments
1.	UNKNOWN	SIGNIFICANT	Likelihood depends on how much we trust the source of information
2.	MODERATE	SIGNIFICANT	–
3.	MODERATE	SIGNIFICANT	–
4.	MODERATE	LOW	Low consequences, as vendor bares liability for all but $50 of fraudulent transactions
5.	UNLIKELY	LOW	As above, but SSL encrypted link which makes it less likely.

Fig. 5. Risk analysis Summary

E-shares might disclose Bob's credit card details Given the identified level of risk, Bob decides he needs to only have MODERATE confidence in E-shares competence to protect his credit card details.

Identify and Evaluate Metrics When relying on information or actions, Bob determines the following metrics to be used to determine the confidence he has in certain beliefs about that entity.

- previous experience with the entity (MEDIUM-HIGH confidence);
- recommendations from other trusted sources (MEDIUM confidence);
- established brands (MEDIUM confidence);
- contractual obligations (HIGH confidence); and
- regulatory controls (MEDIUM confidence).

In addition, he determines the following additional metrics to be used where specific software countermeasures (e.g. the SSL enabled browser he uses) are used to combat risk:

- ITSEC or Common Criteria evaluation (HIGH);
- open source software which has been heavily scrutinised (MEDIUM);
- well known product or vendor (MEDIUM); and
- recommendations from other trusted sources (LOW-MEDIUM).

Lastly, Bob determines the following metrics which are used where he is relying on a third party security system.

- disclosure of security practices and procedures (LOW);
- third party audit by a trusted auditor (MEDIUM-HIGH); and
- certified quality system (e.g. ISO9000) (MEDIUM-HIGH).

Belief Analysis

Information Bob uses to make decisions could be inaccurate Bob has already determined the following beliefs about two entities he will rely on for information:

- Alice: Competence (MEDIUM), Honesty (HIGH), and Predictability (HIGH). Alice can be trusted for information, providing the information is confirmed from at least one other mediumly trusted source. These beliefs were determined solely from a long history of past experience with Alice.
- Reuters News-Wire service: Competence (MEDIUM-HIGH), Honesty (HIGH), and Predictability (HIGH). Reuters can be trusted to report information, providing it can be confirmed by at least one other LOW-MEDIUM trusted source. These beliefs are determined by Reuter's good brand, recommendations from Alice and other friends, and previous experience.

E-shares might disclose Bob's credit card details Bob determines a HIGH level of confidence about E-shares' competence to keep his credit card details secure. This belief is determined from the existence of a certified ISO9000 quality system and a third party audit from KPMG which E-shares describe on their web sites.

Trusting Decisions Figure 6 summarises Bob's trusting decisions for each of the identified vulnerabilities.

Item #	Trust decision	Comments
1.	Accept Risk	Trust Alice's information (confirmed by a Reuter's article), and a policy is described for trusting subsequent information
2.	Accept Risk	Sufficient information is available to trust ComDot-.com's competence to do well. A policy is described for obtaining the required trust in other companies whose shares Bob wants to purchase.
3.	Accept Risk	Bob determines E-shares' privacy policy is sufficient, and trusts them to enforce it.
4.	Accept Risk	Bob is convinced by third party evidence that E-shares' is competent at keeping its site secure enough to mitigate this risk.
5.	Accept Risk	Bob trusts the SSL mechanism used to secure communications with E-shares, and trusts his browser and E-shares web server to implement this mechanism correctly.

Fig. 6. Trusting decisions summary

6.3 Summary

This hypothetical case study outlines the application of the trust management process based on the extended risk model. It should be noted that only a sub-

section of the full analysis is presented. Nevertheless it serves to illustrate the plausibility of such a technique in a real world situation.

7 Related Work

Khare and Rifkin [11] describe how trust management philosophies can be applied to the World Wide Web, and describe how trust policies can be designed. However, Khare and Rifkin's work is very much focused on the expression of trusting intentions, i.e. they describe how to express a trust policy, but do not provide a methodology for how to derive it.

In [12] Jøsang describes general criteria for modelling trust in information security and critiques some other existing formal schemes. Further work by Jøsang [13] develops these ideas into a formal model based on a concept called *subjective logic*. Subjective logic allows us to reason about beliefs or opinion using an algebraic notation, and would be useful in the context of working with trusting beliefs in the extended risk model.

As indicated, there have been several attempts to build trust management systems [1][2][3]. Of these REFEREE[3] is probably the most notable, as it provides way to integrate with third party recommender systems like the PICS [14] labelling scheme. The REFEREE architecture is also extensible, making it simple to integrate new components into the system. Future work on automating the trust management process could benefit highly by utilising REFEREE as a platform for gathering and evaluating information.

8 Future Work

Decision support systems is a catch all for a wide variety of systems which provide computer support for decision making [4]. There is a significant body of work on using decision support systems for risk management [15][8], which could be leveraged to develop similar systems for trust management based on the extended risk model.

Another direction for this work might be the development of trust metrics which could be used to automatically establish beliefs about pages on the World Wide Web. Examples of such metrics might include:

- number of pages linking to a given web page;
- trusted pages linking to given web pages;
- third party recommendations (e.g. PICS labels); and
- number of hits on a given web page.

Search engines might be useful sources for such information. In particular the Google search engine [16] already uses link counts in order to rate matched pages.

Lastly, the importance of considering dynamically changing policies needs to be investigated. Beliefs and trust are not static, but change as new information is received. It would be useful to investigate how policies could be defined which cope with dynamic changes.

9 Conclusions

This paper has presented a scheme for developing trust policies based on an extended risk management model. The scheme was applied to a hypothetical case study, which shows the utility of the process to real world applications. The paper has also discussed related work and given some firm directions for future research in this area.

References

1. M. Blaze, J. Feigenbaum, and J. Lacey. Decentralized trust managment. In *Proceedings of the 1996 Symposium on Security and Privacy*, pages 164–173, 1996.
2. Matt Blaze, Joan Feigenbaum, and Angelos D. Keromytis. Keynote: Trust managment for public-key infrastructures. In *Cambridge 1998 Security Protocols International Workshop*, England, 1998.
3. Yang-Hua Chu, Joan Feigenbaum, Brian LaMacchia, Paul Resnick, and Martin Strauss. Referee: Trust management for web applications. In *Proceedings of the 6th International WWW Conference*, 1997.
4. Dennis Longley, Michael Shain, and William Caelli. *Information Security: Dictionary of Concepts, Standards and Terms.* Macmillan, 1992.
5. Common Criteria for Information Technology Security Evaluation – Part 1: Introduction and general model, May 1998.
6. Standards Australia/Standards New Zealand. AS/NZS 4360:1999 Risk Management, 1999.
7. Communications Security Establishment (CSE) Government of Canada. A guide to Security Risk Managment for Information Technology Systems MG-2, 1992. URL: http://www.cse.dnd.ca/cse/english/Manuals/mg2int-e.htm.
8. Dennis Longley, Michael Shain, and William Caelli. *Information Security: Dictionary of Concepts, Standards and Terms*, pages 450–453. Macmillan, 1992.
9. D. Harrison McKnight, Larry L. Cummings, and Norman L. Chervany. Trust formation in new organizational relationships. In *Information and Decision Sciences Workshop*, October 1995. URL: http://www.misrc.umn.edu/wpaper/wp96-01.htm.
10. D. Harrison McKnight and Norman L. Chervany. The meanings of trust. Technical report, MISRC Working Papers Series, 1996. URL: http://www.misrc.umn.edu/-wpaper/wp96-04.htm.
11. Rohit Khare and Adam Rifkin. Weaving a web of trust. *World Wide Web Journal*, 2(3), 1997.
12. Audun Jøsang. Prospectives for modelling trust in information security. In Vijay Varadharajan, editor, *Proceedings of the 1997 Australasian Conference on Information Security and Privacy.* Springer-Verlag, 1997.
13. Audun Jøsang. A model for trust in security systems. In *Proceedings of the Second Nordic Workshop on Secure Computer Systems*, 1997.
14. W3C. Platform for Internet Content Selection (PICS) technical specification. URL: http://www.w3.org/PICS/.
15. Giampiero E.G. Beroggi and William A. Wallace, editors. *Computer supported risk management.* Kluwer Academic Publishers, 1995.
16. Google Inc. Why use Google?, 1999. URL: http://www.google.com/why_use.html.