Evolving a Secure Internet

William J. Caelli¹, Lam-For Kwok², and Dennis Longley³

¹Queensland University of Technology, Brisbane

²Department of Computer Science, City University of Hong Kong

³International Information Security Consultants Pty Ltd.

{w.caelli,d.longley}@iisec.com.au, cs1fkwok@cityu.edu.hk

Abstract. Internet insecurity is inevitable if a high proportion of Internet users are insufficiently aware of the inherent risks involved, whilst those cognizant of those risks are denied the facilities to manage and control them. This paper highlights the first issue and discusses a potential approach to the second.

Keywords: Secure Internet, trust relationship model, Internet trust relationships.

1 Introduction

Forty years ago Scott Graham and Peter Denning [1] wrote:

On the basis of the foregoing argument, we conclude that the protection system is correct and will operate exactly as intended among trustworthy subjects. Untrustworthy subjects cannot be dealt with completely by mechanisms of the protection system. External regulation, together with a system for detecting and reporting violations, is required.

Forty years later this advice is relevant because it emphasises the key role of pre-existing trust relationships within information security. The focus of information security has evolved from government mainframes, to corporate information processing systems, small business and home computing and now the Internet and mobile devices. These developments involved major corresponding changes in the host environment trust relationships.

When information security developed as a topic in its own right, and as a profession, there was a natural tendency to concentrate on common themes of information security systems, perhaps at the expense of the fundamental role of pre-existing trust relationships in the host environment.

In the 1980's it was commonly assumed that the global communication offered by the Internet would facilitate social cohesion; but now it would appear that the Internet can produce some damaging impacts upon society. If the current impacts have been caused by petty criminals, one could well fear the potential future impacts from organised crime or a rogue government; hence the current widespread concerns for cyber insecurity, coupled with calls for greater regulation, sophisticated defence systems and even hints of assured mutual disruption.

The host environment of the Internet comprises that of the 2 billion Internet users. Whilst the Internet itself, may be viewed from a network security viewpoint, Internet security must be directed at the trust network of the host global user society. When governments are faced with some of the social evils arising from the Internet, e.g. cyber bullying, network security techniques alone offer no solution.

Taking a broad brush approach to Internet security we could recognise that the pre-Internet society had, over millennia, evolved a remarkably successful social trust network, bonding a set of highly complex social systems. The Internet, as a computer network, certainly offers an opportunity of increasing social cohesion. Taking the lead from Graham and Denning it seems clear that the role of Internet security should be defined in terms of enhancing the effectiveness of that trust network.

This paper seeks to explore the approach based upon the concept of the Internet as communication system designed to enhance a highly evolved social trust relationship infrastructure. It suggests that a better understanding of the role and functions of existing social trust relationships could lead to the evolution of a secure sub-Internet, expanding and gradually replacing the current anarchic Internet

The paper first describes an informal personal trust relationship model, and then reviews the current state of Internet trust relationships from a user - supplier viewpoint. Finally there is a discussion on the evolution of a secure Internet service complementing the trust relationships in traditional personal activities: banking, information retrieval, social networking, entertainment etc.

2 Personal Trust

2.1 Overview

The role of trust in society has been widely discussed [2]. The term trust has such emotive affiliations that for the purposes of this paper a pragmatic definition will be adopted:

An entity trusts another entity if it is confident that it can predict the behaviour of that entity in a specified context.

Example: a householder predicts that the plumber will successfully repair a dripping tap.

Trust relationships are formed in the cot; trust relationship training and experiential learning continues throughout a lifetime. Society successfully evolved from small hunting parties, through tribal societies, agricultural communities, the industrial revolution, to the pre-Internet trust relationships, because increasingly complex societies were rendered viable by a vast trust relationship infrastructure.

One of the consequences of this infrastructure was that individuals exploited trust relationships instinctively and became less conscious of their existence. Hence a late nineteenth worker in a small village was probably more acutely aware of the importance of such personal trust relationships than a teenager in the Internet era.

This apparent current lack of awareness had significant consequences in the explosive growth of the Internet.

Example: compare the attitude of a nineteenth century village child to a stranger, as compared with that of a modern teenager in a bulletin board session.

A trust relationship model is defined below and is used in a discussion on the role of such relationships on the Internet.

2.2 Trust Relationship Model

Overview. Fig 1 illustrates the Trust Relationship Model. The model parameters are:

- Context – the total set of potential actions requested from the Activator;
- Initiator – the party that establishes the trust relationship and subsequently sends transaction requests to the Activator;
- Activator – the party that performs the requested transactions;
- Security Attribute – details of the Activators ability and previous performance in performing the transactions;
- Trust Level – the Initiators' estimate of the probability that requested transactions will be successfully completed by the Activator.

In general the Activator predicts that the Initiator will provide some recompense for the completion of individual transactions. Hence most trust relationships are bilateral: party A predicts that party B will satisfactorily perform the transaction, whilst party B predicts that party A will pay the bill. Hence each party plays the Initiator role, in one of the constituent relationships, and the Activator role in the other.

Fig. 1. Trust Relationship Model: The Initiator predicts that within the context of the trust relationship the Activator will behave as predicted with a probability equal to the Trust Level value; the Trust Level is a function of the Activators Security Attribute, which in turn is based upon the Activator's qualifications and reputation in the relationship context

Forming a Trust Relationship. The Initiator evaluates the Activator via the security attribute (see Fig. 1) and estimates the trust level associated that security attribute. The Activator will also set a risk profile and conduct a risk analysis; estimating the impact value of a failed transaction and estimating the risk from the impact value and trust level. If this risk is within the boundaries of the risk profile, the Initiator will propose the formation of the trust relationship with the Activator, who will usually undertake a similar process for the second half of a bilateral relationship, and accept or reject the offer.

Zero Trust Relationship. A multitude of transactions are conducted outside the trust relationship model described above, every day such one-off zero trust relationships are employed when the potential impact is low, or the risk of not requesting the transaction is high. One common example of the zero trust relationship arises when one requests directions from a stranger. Such one-off zero trust relationships are based upon the assumptions:

- The net benefit from a large number of such relationships justifies their usage;
- There would appear to be no benefit to a malicious activator;
- The initiator has some limited basis, e.g. activator's appearance, demeanour to assume a beneficial outcome.

Surfing the Web provides an example of the common usage of zero trust relationships.

Transitive Trust. The difficulties associated with the task of identifying and evaluating activators is commonly bypassed with transitive trust relationships, e.g. recommendations from a trusted friend.

In a unilateral transitive trust relationship A trusts C and C trusts B, which can lead to a transitive trust relationship: A trusts B (See Fig 2). Corresponding bilateral transitive trust relationship Transitive trust can take one of three forms:

- Introduction: C merely passes B's identifier to A.
- Recommendation: A's trust level in B is influenced by C.
- Delegation: B is merely a component in the performance of the tasks; the trust level assigned by A in the A trusts B relationship, is equal to that of A's trust level with C, e.g. B is an employee in Bank C.

Fig. 2. Transitive Trust: A trusts C, C trusts B, A has a transitive trust relationship with B

A recommendation transitive trust is commonly employed when C is in a better position to evaluate B's security attribute than A, and in general A should have a high trust level in C; such transitive trust levels are highly conditional, e.g.

- The context of A's trust in B must be a subset of the context C's trust in B;
- If there are a large number of C's trusting a unique B then an average of these trust levels may be appropriate in the estimation of the transitive trust relationship;
- There should be no significant time gap between the last transaction of the C trusts B relationship and the first transaction of the A trusts B relationship.
- C's identifier for B must be unique in the union of population sets known to A and C.

The trust level between A and B can be no higher than the minimum of the A-C and C-B trust levels. If transitive chains are extended then trust levels between the initiators and activators and the ends of the chain could well fall below acceptable levels for even moderate impact transactions.

Transitive trust chains may become quite long and it is essential that they commence with a non-transitive trust relationship. This condition is significant for Internet users where transitive trust is commonly employed, because, as discussed below, there are extreme difficulties in forming non-transitive trust relationships on the Internet.

Appeal Systems. In bilateral relationships both parties experience a degree of risk with each individual transaction. An appeal system, trusted by both parties, can reassure them that any disputes may be resolved impartially. In effect the existence of the appeal system may be a component of each party's security attributes, effectively increasing each party's trust level.

Appeal systems may play an important role in transitive chains, counteracting the inevitable decline of trust levels along long chains; provided, of course, that the transitive chain does not extend outside the realm of an appeal system common to both parties.

3 Internet Trust Relationships

3.1 Overview

At the end of the 20th Century the Internet was viewed as a large computer network amenable to conventional network security. Today the Internet has evolved into a novel worldwide distributed social entity with a strong antipathy to security and regulation.

The past two decades have seen a worldwide migration of traditional activities to the Internet. Many of these traditional activities relied upon trust relationships evolved over centuries, and this paper discusses the issues surrounding the establishment of corresponding trust relationships on the Internet.

The establishment of trust relationships in a small community was facilitated by propinquity and multichannel communication, i.e. a few people in close proximity who using one or more their five senses in their interactions. The Internet user population, on the other hand, represents of the order of 30% of the world population, and a major part of its traffic is restricted to the exchange of text and graphics. This environment seriously complicates the establishment of trust relationships as discussed below.

3.2 Establishing an Internet Trust Relationship

Activator Identification and Evaluation. Initially the initiator must locate and identify a potential activator; the initiator must then select a set of activator attributes rendering that individual unique amongst the population known to the initiator, so that the initiator will be able to identify the activator when requesting future transactions. If the initiator and activator live in a small community then the activator's physical appearance will suffice. However, as activators are selected from increasingly large populations the number of personal attributes required for unique identification grows rapidly.

The unique identification of the activator becomes even more problematic as information is collected to determine the activator's qualifications relevant to the context of the proposed trust relationship. Again such evaluation in respect of a local plumber in a small community is straightforward, but becomes increasingly problematic when such information must be collected from a variety remote sources and the initiator must verify that the activator identifying information, provided by each source, corresponds only the selected activator.

The initiator must aim at the collection of sufficient information, about the activator's potential performance in the relationship context, to estimate a trust level for the proposed relationship. Then a decision is made on whether or not that trust level is consistent with the initiator's risk profile.

Having collected the security attribute information, either directly from the activator or one or more other sources, a chicken-egg problem becomes apparent. How does the initiator verify the integrity of the information collected over the Internet? If the initiator has a trusted source of security attribute information, how was the trust relationship formed with that source? A similar problem arises if the initiator bypasses the identification and evaluation process with a transitive trust relationship. Hence it is apparent that trust relationships can only be formed on the Internet after some root trust relationship has been formed outside the Internet.

When the initiator has selected a specific activator then the initiator's problem of ensuring that future transactions are initiated with the correct activator must be addressed. Since the communication is via the Internet, the initiator and activator need to agree upon some effective authentication process for future communications. If the relationship is non-transitive root then this exchange of authentication details, e.g. public key certificates, should also be undertaken securely outside the Internet.

Authentication processes employed in communication channels assume that the communicating parties do not deliberately provide masquerading entities with the

authentication parameters. Hence such authentication processes should be backed by some regulatory authority such that an individual would be held legally responsible for actions undertaken in any such masquerades.

Once the root trust relationship has been formed the Internet is, from an information security viewpoint, merely a computer network and the security facilities offered by public key cryptography may be deployed.

It would appear from the above discussion that the combination of trust relationships, formed in the traditional manner outside the Internet, may be used to provide secure roots for Internet transitive trust relationships and even trust relationship chains, thus combining the best of both worlds. This approach requires both a reconsideration of current approaches to Internet security, and a secure deployment of public key cryptography.

4 A Secure Internet Service

4.1 Overview

The period preceding the Internet experienced amazing technological advances in electronics and communications; these advances resulted in qualitative and quantitative changes; mass production of microelectronic devices provided mass access to computing and communication services. In the early years of computing developments Governments largely drove the agenda, but the Internet advances mainly resulted from market forces exploiting low cost consumer electronics. Such market forces tended to view security as a costly obstacle and the consumer was provided with a choice in many aspects of the Internet, except the level of personal security.

In conventional trust relationships the user weighs the advantages of undertaking transactions on the balance of potential gain and risk; the risk itself is measured in terms of impact probability and value. Over the past two decades it has become clear that a significant proportion of Internet users lack the detailed knowledge of Internet technology to evaluate their risks, e.g. they are unaware of potential unfavourable outcomes and associated impacts. For example: malicious code downloads, monitoring of users' Web usage, penalties of intellectual properties transgressions, lack of privacy etc.

Moreover a high proportion of Internet usage, i.e. surfing the Web, comprises apparently one off zero trust transactions (See Zero Trust Relationship in 2.2). Asking a stranger for directions is considered relatively safe because for any such single transaction the wrong information has limited impact, and the stranger is unlikely to benefit from deliberately malicious behaviour. Web surfing, however, may not fit this pattern, e.g.

- a malicious posted set of false information will impact upon multiple Internet users;
- monitoring of user Web actions may in some circumstances have significant long term impacts on specific users.

In recent years many users have been virtually compelled to employ Internet services because the off-line alternatives are not locally available or are too expensive. At the same time the average Internet user is provided with few opportunities to protect themselves apart from subscribing to anti-virus services and avoiding obvious pitfalls with passwords. It is proposed here that Internet users should be provided with the option of a security policy providing similar levels of security to that offered by traditional off line trust relationships, e.g. manual banking compared with current Internet banking.

The proposed secure Internet service is intended to complement, rather than replace, current Internet services, and could therefore commence with a few applications allowing user demand for security to determine its success.

Internet applications may be broadly listed in two categories: text/graphics and audio-visual. The text-graphics type Internet applications consist of financial services, commercial services, information retrieval, social networking and email, and interactive education; whereas audio-visual type Internet applications include music, movies and education presentation. The following sub-sections discuss the trust relationships of these Internet applications.

4.2 Text – Graphics Trust Relationships

Financial Services. These applications normally involve a pre-existing off-line trust relationship between the financial institution and the Internet user, and could be offered with a high level of user security. The client and financial institution are in a position to exchange mutual authentication data off-line and the Internet user security then depends upon the technology employed for the authentication process (See 4.4 PKI and Secure Interface Devices) , and the security of the institutional computers.

Commercial Services. Most on-line shoppers have no pre-existing trust relationship with the Internet supplier, and few Internet users would be in a position, or prepared, to establish such relationships with individual Internet suppliers. In this case transitive trust relationships are the only option for the secure Internet service, which implies some organisation, termed here root organization, is prepared to provide a root trust relationship for Internet users (See Transitive Trust in 2.2). Consumer Protection Authorities, for example, would be well placed for such a root organization role. Such agencies not only have some role with suppliers in specific geographic areas, they often have regulatory powers and access to some appeals system.

In the proposed scheme the suppliers would register with the root organization, and supply authentication data, off line. The Internet user would also register with this root organization offline and collect its authentication data. Having located a supplier on-line, and checked that it is registered with the root organisation, the user collects the supplier authentication data established with the root organization, and initiates the transaction with that supplier. The user supplier trust relationships could also be in a bilateral form, in which case the user would supply authentication data at time of registration with the root organization.

The obvious objection to this proposed scheme lies in the limited range of suppliers associated with a particular root organization; implying that Internet users would be restricted to a few suppliers, according to user's ability to register off line with various root organizations. If, however, diverse root organizations are prepared to offer reciprocal regulatory protection then they may merge off-line and provide their registered users with high trust level relationships over a much larger range of suppliers. One of the major advantages of the proposed scheme lies in the regulatory powers of the root organisation increasing the trust level of the user - supplier relationship.

The success of such a proposal is dependent upon market forces but it has the advantage that such a scheme could evolve from a small base.

Information Retrieval. Information retrieval probably represents the most important single application on the Internet. Search engines provide access to a host of relevant information. Web surfing is commonly regarded as a low risk activity, although the associated privacy risk to users is underplayed. However, users are on occasions concerned about the authenticity of accessed information; if one considers print based information retrieval it is apparent that the users on such occasions are strongly influenced by the provenance of written text ranging from newspapers, legal documents to respected text or reference books.

The risk associated with Internet information retrieval was apparent in Australia when an environmental activist, armed only a mobile phone and laptop, produced a fake Internet press report causing a sharp fall in the stock market value of a mining company. Medical and health warning information is now commonly accessed over the Internet by medical practitioners and the general public; one hesitates to list the potential dangers of this situation.

In traditional information retrieval the user can easily distinguish between a leaflet, a respected newspaper or reference book in the library, whereas URLs provide only limited guarantees of provenance and even these are commonly ignored. Even if users access a reputable site, they can be seriously misled by malicious hacker alterations in the text.

In the proposed scheme a reputable publisher acts as a root organisation (see Commercial Services in 4.2) for various Web publishers, although such a root organisation would normally have limited, if any, regulatory powers. The user obtains the Web publisher authentication data from the root organisation and checks a digital signature included in the Web page, thus providing assurance on the source and integrity of the displayed text.

Social Networking and Email. Social Networking did not feature in the early security vulnerable areas of the Internet, but the current level of cyber bullying is now a topic of government concern. Faced with the Pandora's Box opened by youthful entrepreneurs one can only compare the dangers faced by teenage parents in the pre-Internet days with their woe-begotten current counterparts.

Most parents traditionally adopted individual and group strategies in relationship to their off springs' companions. At an individual level the identity, security attribute

and trust level of the candidate activator were routinely assessed, because this information was available. At the group level the culture surrounding clubs for various age groups and interests highlights the interesting trust relationships between club members and the organising committee. The members trusted the local committee to vet potential members and effectively used that trust in the formation of transitive trust relationships between current and new members. Is it conceivable that a similar approach could be used to form high trust transitive social trust networks on the Internet?

Email security is particularly interesting in the context of this paper because PGP (pretty good privacy) [3] addressed email security concerns and used both public key cryptography, and transitive trust for the development of a certificate chain. Email addresses can be masqueraded and from time to time one receives emails with a colleagues email address but with associated text clearly derived from another source. There are good arguments for more use of PKI in emails, particularly in large organisations, where sender certificates could include attributes informing the recipient of the role and authority of the sender within the company.

The fundamental risks associated with emails arise, however, from the curse of immediacy. In the pre-Internet era many people re-read their outgoing correspondence before sealing the envelope, reflecting on the contents and potential reactions of the recipients.

Interactive Education. If the Internet is to have an increasing role within education, in particular higher education, the risks associated with online tutorials and assessment deserve detailed consideration. The potential pitfalls, and associated litigation, arising from tutor – student interactions and assessment decisions suggest that a serious re-consideration of academic trust relationships, and their implementation in a global Internet based educational network, would be advisable.

4.3 Audi-Visual Trust Relationships

Overview. Originally Internet security was based upon conventional network security dealing with primarily with textual data. One of the perhaps more surprising outcomes of enhancing a single sense channel (sight) with a second channel (sound) is that two teenagers with Web cams can now securely mutually authenticate over the Internet without the aid of cryptography.

Audio visual applications have at least two significant security implications: intellectual property and privacy. Governments have been concerned both with protection with corporate profits and the impacts of cyber bullying, responding with strict legislation for the one, and serious hand wringing for the other.

Intellectual Property. The trust relationships associated with theatres and cinemas normally took the form that the client predicted the quality of entertainment provided would be compatible with the price of the entrance ticket, the supplier predicted that the audience would not express significant disapproval during or after the

performance. Such trust relationships are significantly different when the client plays back some digitised music or video. The supplier now takes the risk of loss of income from piracy of the digitised information, whilst playing back such digitised information make involve the risk of penalties for intellectual property legislation transgressions.

In a previous generation the music industry employed technology that made illegal copying expensive; the publisher bore the cost of producing vinyl disks and the consumer purchased a hi-fi system capable of playing, but not reproducing, that disk. Digitisation revolutionised this industry, the supplier was no longer burdened with the cost and distribution of the disks, and the consumer purchased equipment capable of both play and reproduction. The downside from the supplier's viewpoint was the potential theft of their intellectual property. Their solution is to pass the responsibility of the protection of their intellectual property to the user, with major financial penalties for transgression; at least some governments have actively supported this initiative. An alternative solution to the intellectual property dilemma involves a reversion to the previous situation in which the supplier product was supplied in a form that could be played but not cheaply reproduced.

Cryptography combined with a special purpose secure playback devices could provide such a solution. In effect, the encrypted digitised data supplied could only be decrypted and played in the secure device holding a private key technologically protected against illegal access.

Privacy. There are no current technological solutions to the age old problem of the presumed friend who maliciously passes on intimate secrets. Social networking has unfortunately provided an international broadcast audio visual system for such indiscretions. As such it has exemplified the problems of the Internet world lacking effective social trust relationships. In previous generations teenagers were at least inculcated into the imperfect world of personal trust relationships (see Social Networking and Email in 4.2).

4.4 PKI and Secure Interface Devices

PKI. An Internet-wide PKI would provide a parent-child hierarchy of Certificate Authorities and presumably unique identification for each user. However such a system could pose a significantly enhanced threat of identity theft since it would rely upon the cryptographic strength of a particular public key algorithm, and the integrity of a vast host of employees charged with issuing certificates as well as the underlying computer systems used to create, store and distribute the base certificates themselves.

The proposed Internet trust relationship networks with root organizations uses a local PKI, and extends it with sibling certificates issued by the root organizations to their trusted individual Internet suppliers and, if appropriate, to their registered users. The user, and supplier, offline registration process with the root organization would thus involve face to face authentication and exchange of certificates.

Secure Interface Devices. The theme of this paper is that an Internet user should have the opportunity to benefit from trust relationships similar to that enjoyed in the pre-Internet era. Exploiting the proposed Internet transitive trust relationship network requires:

- secure end to end authentication;
- security of communication channels;
- some means of estimating trust levels over transitive chains, which may extend beyond the aegis of local appeal systems.

PKI Certificates exchanged between activators and initiators can facilitate unilateral or bilateral mutual authentication, and the exchange of cryptographically secured messages. The advantage of this proposed system is that it can mimic the conventional trust relationships practised outside the Internet where users make value judgements on transactions based upon impact and trust levels.

The user private key in this arrangement is the keystone to user security; its value and processing must be protected from the malicious code inevitably residing in the user's computer. Current technology has produced a plethora of handheld smart devices and hence a cost effective secure interface device capable of protecting cryptographic private keys and public key certificates can be reasonably postulated.

In keeping with the principle that the Internet trust relationships are either established outside the Internet or via transitive trust chains, certificates for non-transitive trust relationships, e.g. Internet banking, and root transitive trust servers would be loaded directly into the secure device.

The secure device has the task of extracting and checking certificates and supplier public keys derived from certificates, including sibling certificate chains, aided by attributes of the various certificates, and performing the corresponding cryptographic operations.

The sibling certificates may also be employed to facilitate end to end trust levels in long transitive chains. Attributes of these certificates may contain details of link trust levels and existence/ non-existence of end to end appeal facilities. Given some monetary impact value for the transaction the secure device could even provide warnings of risky transactions.

5 Conclusions

The user's security role is perhaps the most significant issue arising from this paper. In conventional information security environments, similar to those addressed by Scott Graham and Denning [1], the host organization information security system was designed to strengthen the pre-existing organisational trust framework. To this extent the user had a somewhat passive security role, e.g. protection of passwords etc. With Internet security, however, there is no host organization and the users have a major security role, including responsibility for their own risk analysis and management. Unfortunately the average user is neither equipped to fulfil this role, nor in a position to establish a requisite level of security. It is therefore of some concern

when suggestions are made that users should be held liable for security breaches, e.g. penalties for harbouring botnets.

This paper emphasises the key role of users in Internet security and highlights two major factors of that role: user Internet security education, and facilities for the deployment of trust relationships with trust levels consistent with user risk. Such a user education requirement is not particularly novel; vehicle drivers are not legally permitted to use public highways with skills limited to manipulation of automobile controls, and total ignorance of road traffic interactions. The current Internet user security awareness situation may perhaps be traced to an ill-informed replacement of traditional IT education with minimalist mouse icon click training, and should be redressed as a matter of urgency.

The, hopefully increasing, proportion of Internet users with sufficient knowledge and skills to protect themselves will be the key drivers, and only hope, of a future adequately secure Internet. This paper discusses the harnessing of traditional trust relationship skills, and the facilities required to implement secure Internet trust relationships. The fundamental problem of establishing secure Internet trust relationships is addressed with a proposal for transitive trust relationships, supported with secure authentication, rooted on traditional trust relationships formed off line

References

1. Graham, G.S., Denning, P.J.: Protection Principles and Practice. In: Proc. AFIPS Spring Joint Computer Conference, pp. 417–429 (1972)
2. Fukuyama, F.: Trust: The Social Virtues and the Creation of Prosperity. Penguin Books (1996)
3. Zimmermann, P.: The Official PGP User's Manual. MIT Press (1995)