# CSIRO Submission 09/353

## RESPONSE TO HOUSE OF REPRESENTATIVES STANDING COMMITTEE ON COMMUNICATIONS INQUIRY INTO CYBER CRIME

July 2009

**Enquiries should be addressed to:**

Mr Peter King

Executive Manager

CSIRO Government and International

PO BOX 225

Dickson ACT 2602

Australia

Ph: (02) 6276 6361

Email: PeterG.King@csiro.au

**Main Submission Author**:

Dr John Zic

Principal Research Scientist

ICT Centre

PO Box 76

Epping NSW 1710

Australia

Ph: (02) 9372 4122

Email: John.Zic@csiro.au

# Contents

# 1. Executive Summary

CSIRO welcomes the opportunity to provide input to the House of Representatives inquiry into cyber crime and its impact on Australian consumers. In addressing the Terms of Reference CSIRO has responded to those areas where we have appropriate knowledge and expertise.

CSIRO has developed a world leading research and development capability in trusted systems. The group has a proven track record with its ability to address requirements and develop solutions in co-operation with key players in the finance and government sectors, as evidenced by the successful patenting and commercialisation of the Trust Extension Device, as well as the acceptance of seventeen internationally refereed academic papers this year.

This report to the House of Representatives Standing Committee on Communications addresses Terms of Reference point (f) “Emerging technologies to combat [these] risks” to consumers by criminal activity on the Internet. Two technologies are presented, both of which raise the level of trust in the Internet.

Our submission also addresses two important caveats: first, any technological solution must be able to deal with the basic design principle behind the Internet as a minimally constrained network system. As such, it offers only a best effort service on the delivery of data, and that at a fundamental level, offers no assurances about security, privacy, delivery or a minimum quality of service level. Second, any technological solution must not be divorced from a critical examination of the context of its use: from the social, to legal, to business and personal. Any technology deployed without consideration of these two is bound to fail in its mission of providing any consumer protection.

# 2. Introduction

This report will present an overview of technologies that are designed to combat the risks to consumers incurred when they use the Internet.

However, this point is addressed in the context of the following two important caveats.

First, the Internet *as an entity* is not going to fundamentally change in its nature because of the current investments placed into its infrastructure. It is a best effort service that delivers data from one network connected machine to another: no more, no less. Academic wishes for a "clean slate" approach to the architectures and protocols that underlie the Internet to make them more reliable, safer, assure quality of service and so on, while worthwhile in terms of research into possible future internets[1] (sic) are unlikely to impact on the current infrastructure for many years unless there is a globally compelling economic argument for change. As a consequence, any attempt to circumvent criminal activity by consumers using the Internet must be able to deal with its inherent untrustworthiness and no assurances of data being delivered correctly (or at all). This fundamental principle guided the development of the Internet from its inception. It is entirely the responsibility of the implemented client applications and their underlying protocols that then build on top of these basic Internet protocols to assure reliable and timely delivery, security, privacy, and authentication of messages being sent as data on the Internet.

Second, people, their interactions with each other and organisations must be considered in any design of secure or trusted system that protects them from criminal threats and actions, whether this be Internet based or not. It must be emphasised that technology on its own cannot protect against malicious criminal activity. This point is supported by leaders in the trust and security communities, such as Ross Andersen from Cambridge University and Bruce Schneier, Chief Security Technology Officer of BT. Any attempt at minimising threats from criminals, and raising security of individuals, must carefully consider the intersection between:

1. the needs and rights of individuals,
2. the needs of businesses interacting with their customers,
3. legal and regulatory requirements,
4. societal needs and expectations and finally
5. the technologies best suited to fit in this intersection.

[1] Note that there is a differentiation between the Internet (with a capital "I") and an internet (small "i"). An *internet* is defined to be a network of networks. The *Internet* is one instance of an internet; other instances are possible.

A question that should be asked for each of these categories is: "*What are the measures and costs involved in the mitigation of malicious or criminal behaviour?*"

Schneier[2] states that despite the seductiveness of technological solutions, it is the inability or unwillingness of decision makers, engineers and security specialists to elaborate these complexities to the public. These complexities are the heart of the apparent failures of security solutions that are supposed to protect the consumer against criminal activity. In fact, it is easier to propose a technological solution as "the solution" to all of the consumers' security concerns. However, Schneier states that:

> "Technology is generally an enabler, allowing people to do things. Security is the opposite: it tries to prevent something from happening, or prevent people from doing something, in the face of someone actively trying to defeat it. That's why technology doesn't work in security the way it does elsewhere, and why an overreliance on technology often leads to bad security, or even to the opposite of security."

At another level, the recent US National Academies of Science report (2007) on Security in Cyberspace[3] makes the observation that addressing cybersecurity through the uniform adoption and use of technologies has not been successful, despite the best efforts of the security communities over the past fifteen years. The reasons for this are related to the inability of decision makers, whether in government or enterprises, to co-ordinate with each other and recognise the complexities and costs of securing the many systems against identified threats. To quote (Ibid, p228):

> The various cybersecurity reports issued to date have not provided the sufficiently compelling information needed to make the case for dramatic and urgent action. If so, a sufficiently ominous threat cloud will inspire decision makers to take action. But it is well known that detailed and specific information is usually more convincing than information couched in very general terms—unfortunately, detailed and specific information in the open literature about the scope and nature of the cyberthreat is lacking.

---

2 "Beyond Fear: Thinking Sensibly About Security in an Uncertain World", Bruce Schneier, p13 Corr. 2nd printing, 2006, VIII, Springer Publishing, ISBN: 978-0-387-02620-6

3 "Toward a Safer and More Secure Cyberspace", Seymour E. Goodman and Herbert S. Lin, Editors, Committee on Improving Cybersecurity Research in the United States, National Research Council, 2007. ISBN: 0-309-66741-0

Even with the relevant information in hand, decision makers discount future possibilities so much that they do not see the need for present-day action. In this view, nothing short of a highly visible and perhaps ongoing cyber-disaster will motivate actions. Decision makers weigh the immediate costs of putting into place adequate cybersecurity measures, both technical and procedural, against the potential future benefits (actually, avoided costs) of preventing cyber-disaster in the future—and systematically discount the latter as uncertain and vague.

The costs of inaction are not borne by the relevant decision makers. The bulk of the nation's critical infrastructure is owned and operated by private-sector companies. To the extent that these companies respond to security issues, they generally do so as one of the risks of doing business. But they do much less to respond to the threat of low-probability, high impact (i.e., catastrophic) threats, even though all of society at large has a significant stake in their actions.

With respect to consumers being protected against criminal activity, there are localised (businesses and government departments) integrated security solutions that work well. However, with the proliferation of Internet enabled and connected businesses and systems, localised solutions are, by definition, fragmented and uncoordinated. Consequently, a criminal attack may be no longer on individual consumers, but on the very infrastructure that they have come to rely upon, or even upon the enterprise to enterprise levels. Coordinated security solutions are going to be difficult to justify until there is a large enough cost incurred across the Internet connected systems that are likely to be affected.


Cartoon by Peter Steiner. The New Yorker, July 5, 1993, Vol.69 (LXIX) no. 20, P61

With these important caveats in mind, we review and present two promising technologies that could to raise consumer trust in their day-to-day dealings on the Internet.

# 3. Emerging technologies

The cartoon on the previous page highlights only one of the key issues that are faced on the Internet: nobody knows who you are, or (just as importantly) how you will behave. The caption on the cartoon should be modified to read: “On the Internet, nobody knows that you are a dog, and no one can prove that you will behave exactly as a dog – yet”. Great lengths have been taken to provide authentication mechanisms and technologies to help provide proof to a claim of identity. However, identity is only one part of a complete solution. On its own, an identity does not give any indication or assurances of how you will behave. All it provides is a constant, unique and abstract identifier. The only assurances that are made are the same identity is being claimed, and your own prior experiences will strengthen your belief that the identifier is actually the correct and expected individual or organisation. It is simple enough (and widely done) to claim to be a leading bank and provide proof of identity on the Internet to ordinary users. When this is coupled with a typical user's naivety, or lack of care, this lack of validation of behaviour allows fraud to be easily and successfully deployed by criminals.

Consider the following banking example and the ease with which fraud (from the point of view of the criminal) may be perpetrated on the consumer.

It is difficult, uneconomical, and presents an unacceptably high risk to the criminal of being caught if they were to set up a fully functional bank branch in a shopping area, including provision of teller staff and the physical site and fit out of a real bank. Setting up a fake (physical) ATM is easier than setting up a branch, but is largely uneconomical and represents an unacceptable high risk to the criminal of getting caught.

It is much easier, and potentially more lucrative, to use an existing, installed ATM and use non-obvious attacks such as the installation of malicious software into the ATM to allow the criminal to harvest card and account numbers, PINs,etc. The recently reported[4] malware infection in Eastern European ATMS was one such case, offering a rich and not easily detected method, of allowing criminals total control over the infected ATM (including dispensing all of its

---

[4] “Automated Teller Machine (ATM) Malware Analysis Briefing”, TrustWave SpiderLabs, May 28, 2009

https://www.trustwave.com/downloads/alerts/Trustwave-Security-Alert-ATM-Malware-Analysis-Briefing.pdf

cash). Here the issue is that despite the ATM's claimed identity (verified in this case by its physical appearance and “normal” or expected behaviour), the actual behaviour of the system has been modified by the malicious software.

Central to the tenant of this report is that any technological solution to the problem of consumer fraud relies on the ability to provide incontrovertible evidence (proof) of correct, expected behaviour for all identified parties involved in transactions on the Internet. This means that all parties will be able to trust each other, since they know all their identities and their behaviours are as expected.

It is important to note that there is a distinction between trust and security. According to the Internet Engineering Task Force (the standards body that defines the protocols and systems behind the Internet) in the RFC4949:

```
$ trust

   (I) /information system/ A feeling of certainty
   (sometimes based on inconclusive evidence) either (a)
   that the system will not fail or (b) that the system
   meets its specifications (i.e., the system does what
   it claims to do and does not perform unwanted
   functions). (See: trust level, trusted system,
   trustworthy system. Compare: assurance.)

$ security

   1a. (I) A system condition that results from the
   establishment and maintenance of measures to protect
   the system.

   1b. (I) A system condition in which system resources
   are free from unauthorized access and from
   unauthorized or accidental change, destruction, or
   loss. (Compare: safety.)
```

The distinction comes about since trust and security are, like privacy, statements about how information is treated within a system.

As a way of clarifying the distinction between trust and security, consider the following example. A consumer's PC may establish secure, encrypted connections with another machine claiming to be a bank, but there is no assurance given to the consumer that the machine they are interacting with can be trusted to really behave as the real bank and not steal the account information and contents (just as in the ATM case above). Conversely, if the consumer is indeed provided the incontrovertible proof from the bank that it is a real bank, and will behave in exactly the expected manner, transferring confidential information across the Internet in clear text (rather than in secure, encrypted manner) will also be a problem (in that the medium itself is fundamentally untrusted) since that information can be easily intercepted.

Summarising this paper's position, we believe that consumer fraud is intimately tied to misplaced or simplistic trust of both the identities and behaviours of the

parties involved in transactions, coupled with the assumptions that the underlying technologies used to exchange information are secure.

None of these should be assumed to be assured. Providing current, general technologies on an Internet scale is going to be practically unmanageable due to the size, complexity and heterogeneity of the current Internet.

A better approach is to develop a set of specific technologies with a particular use in mind. These are tailored in such a way as to allow all the participants in transactions to establish a trusted, secure, private collaborative system. The transactions and information are contained entirely to the collaborative system. Each participant first proves their identity, and provides to the other participants a mathematically provable statement about their behaviour. The participants formulate a formal agreement between themselves, and then once they are all in agreement, carry out their critical transactions. Once the collaboration is complete, either by agreement between participants, or over a specific time period, the system used to interconnect the participants together is torn down.

Establishing trust and proof of behaviour may come from either previous experiences, or from using reputation (as does Amazon and eBay), or by using special hardware-based trust solutions promoted by the Trusted Computing Group[5]. Of course, there is always a cost involved in establishing and maintaining trust and security levels. Therefore, a careful examination of these costs, as well as a thorough risk assessment and quantification must be carried out on a regular basis.

Critical application of trust technologies in their intended scenarios of use, along with security and authentication mechanisms, offer a set of methods that will help protect the consumer from incidence of cybercrime and fraud.

Again, establishing this system requires a set of inter-related technologies that assume that the consumer's machine is in an unknown and unpredictable state. Also, it must be a design assumption that the network that connects the participants may be intercepted and listened to, and that participants themselves are unknown (despite their claims).

We now proceed to describe two new technologies that could improve the trust and security of Internet based services.

5 http://www.trustedcomputinggroup.org/about_tcg

## The Trusted Computing Group's *Trusted Platform Module*

The Trusted Platform Module, or TPM, is a widely deployed[6], but minimally used cryptographic microcontroller chip. This chip contains certificates, private cryptographic keys, secure storage and a set of specialised cryptographic functions. When combined with an appropriate infrastructure (network connections, application and certifying authority services), the TPM allows the validation of a user's identity, the identity of the machine, and through the use of an attestation protocol, prove the correct functioning (no unexpected programs or behaviours) of the machine. This correct functioning of the machine is done through a set of measurements on all components of the machine: the hardware layer, including the TPM chip, CPU, memory, disks, etc., through to the operating system and finally the applications on the machine. Measurements are cryptographically encoded for exchange between the participants in a transaction (e.g. the consumer's home computer and a bank's server). Each participant decodes this measurement, and if it is as expected, proceeds with participating in the transaction. In short, the use of this TPM assures the identity of a machine and that the machine is in a recognised, agreed upon configuration.

Should any variation occur from the expected configuration, such as the addition of a new version of application software, or a change in the hardware, or the inadvertent introduction of a piece of malicious code into the machine (e.g. as was the case in the ATM attack described above), the TPM would generate a measurement that was unrecognised by the remaining participants and so the transaction would not proceed.

The TPM chip has not been widely used because of several inhibiting factors. The most important technological inhibitor is the maintenance complexity of a system that attempts to use the full capability of the TPM. Any changes in the machine's configuration – hardware, operating system or application software – will cause any TPM-based transactions to deliberately fail. It is possible to address this, but at substantial cost and effort. Large scale deployment, given the variety of machines, operating systems, makes this task largely infeasible. This management problem has been recognised by the Trusted Computing Group, and there is some research now being done on addressing this through the use of a "minimal" unique feature set that can be used to prove the correct, expected functioning of the machine.

The TPM chip and associated system also has other inhibitors besides this technical management problem. Other inhibitors fall into legal and social privacy concerns, as the genesis of the TPM has been in the Digital Rights Management systems proposed by media content providers and owners.

---

6 "About 300,000,000 PCs have shipped with a chip called the Trusted Platform Module (TPM), with capabilities beyond traditional tokens or smart cards." – from the TCG web site.

The technological inhibitor largely disappears in the TPM-based system if it can meet two conditions:

1. The hardware/software system that uses the TPM chip is well-controlled and possibly proprietary, and unalterable by anyone, including the customer.
2. The cost of providing the system goes down by at least an order of magnitude below current computing solutions.

In the case of the malware affected ATMs referred to earlier, the introduction of TPM-based system would have prevented the malware fraud from succeeding. Had TPM been installed into each ATM, along with the appropriate supporting software services, each ATM could be authenticated against its claimed identity, as well as authorised to proceed with any transactions. Again, each ATM would be able to prove its integrity to a controlling server (certification authority), application server and ultimately, its correct configuration to the consumer. The introduction of the malware software would have been detected during the measurement of the TPM enabled ATM environment.

Another approach that can be used to meet the above conditions is by simplifying the software and hardware system associated with a TPM, as well as limiting its capacity and ability to be altered. In this specific case, proving the integrity of a component within a TPM based system becomes easier once again.

CSIRO has developed, patented[7] and is commercialising such a device, that is small, cheap and portable, referred to as the Trust Extension Device (or TED). This device is in the form factor of a USB memory stick, and is intended to be issued to the consumer (and revoked if lost or stolen) by a trusted authority such as a bank, enterprise or government agency. It has its own “locked down” applications and operating system, running on a dedicated and isolated CPU with a TPM chip used to attest the integrity of the device. The device plugs into any host PC with a USB port and provides an isolated environment from the host PC. As it is isolated from the host machine, the consumer can be confident that there will be no viruses, trojans or other malware inserted onto it, and if there were any such changes, the transactions would not proceed as these changes would be detected. The TED itself uses the host machine to simply establish a set of secure network connections to the application server (e.g. in a bank), and then all transactions and processing is sent through these secure channels, with no information being held on the host machine.

From the consumer's point of view, the TED has been issued by a trusted party. The consumer can be confident that the TED can also be trusted with their information as it will fail if there have been any changes to the TED software or hardware. The system itself makes the usual consumer problems of “phishing”,

---

[7] International Patent TW7941/WO, “A portable device for use in establishing trust”, John Zic and Surya Nepal, 11 September 2006.

memory attacks through the use of trojans, or man-in-the-middle attacks difficult to mount, not because of radically new security technologies, but rather because of the way that these technologies are used together to prove trustworthiness of the participants' (machines) in a transaction.

## Collaboration Services for Trusted, Secure Information Exchange

The use of Darknets in distributing anonymous content between machines is well known within the networking community, with systems such as Freenet[8] and WASTE[9] being freely available and used to prevent the interception of content (outside of the Darknet) through the use of anonymising and encryption technologies.

Although sometimes associated with illegal activities (hence the term Darknet), we can take some inspiration from the Darknet model – we would like to establish a private collaborative system that like a Darknet, hides content exchanged within the system from outsiders. Unlike Darknets, the content is legal, the participants are known to each other, and their behaviour is completely (and provably) specified to each other.

Facilitating this sort of collaboration requires that each participant identifies and states their own information sharing policies. Typically, these participant policies, and associated collaboration policies may be captured and shared in an electronic document or contract. The participants must first reach agreement between their policies before the collaboration is started. The CSIRO has previously developed and demonstrated such a system as one of the final milestones[10] of the Centre for Networking Technologies for the Information Economy (CeNTIE), funded through the Department of Communications, Information Technology and the Arts (DCITA).

For example, consider the following medical scenario. A Patient consults with their local GP, who then needs to refer the Patient to a specialist S1. In this collaboration example, the Patient's Medical records need to be exchanged

---

[8] http://freenetproject.org/ – "Freenet is free software which lets you anonymously share files, browse and publish "freesites" (web sites accessible only through Freenet) and chat on forums, without fear of censorship. Freenet is decentralised to make it less vulnerable to attack, and if used in "darknet" mode, where users only connect to their friends, is very difficult to detect."

[9] http://waste.sourceforge.net/ – "WASTE is an anonymous, secure, and encrypted collaboration tool which allows users to both share ideas through the chat interface and share data through the download system. WASTE is RSA secured, and has been heralded as the most secure P2P connection protocol currently in development." Interestingly, the source code for WASTE has been recently removed from sourceforge – http://en.wikipedia.org/wiki/WASTE

[10] CSIRO CeNTIE Milestone Report 52 to DCITA, "Implementation of storage networking and trust elements of a service operator architecture in appropriate sections of the foundation network", Sheldon Dealy et al, May 2007.

between the participants. GP knows and trusts S1, and the Patient knows and trusts the GP. However, the Patient does not trust the Specialist S1. Clearly, this conflict needs to be resolved before any of the Patient's medical records are exchanged. This can be done through the usual negotiation, with the GP and Patient agreeing to a second Specialist S2, whom both know and trust.

In this multiparty medical example, this preliminary discussion and subsequent negotiation between the Patient, GP, and referral to one of two specialists could have been done through the sharing of a contract document that captured these policies.

Once the multi-party agreement has been completed, the private content can be shared through the configuration and deployment of secure network connections augmented with some privacy preserving techniques. This combination would assure the content remains private to that collaboration. The same contract document can be used to enforce access control information, security level classification, encryption requirements and so on within any stored information that is shared between the collaborators.

It is possible to deploy such a system at the Internet infrastructure level, and this was the basis for the prior CSIRO CeNTIE research project. In such a system, the Internet routers themselves are programmed and configured to behave in this manner. Alternatively, by adopting a higher-level, *service oriented* view of the world, the configuration etc is now the responsibility of service providers (both network and storage). These providers can be used to rapidly deploy a collaborative system without needing to be concerned about Internet router manipulation. Of course, an agreed upon trust and security service interface for each of the providers is required in order for them to communicate with each other in a dynamic, reconfigurable manner. Again, there has been some progress in developing some initial prototypes and this is reported in recent publications[11,12]

The configuration, maintenance, and termination of the collaboration infrastructure may be automated through a dedicated management service. This management service is responsible for the interpretation of the service contract between the participants and is also responsible for contacting and establishing the appropriate service providers required to construct the collaboration service.

A final component within the proposed trusted, secure and private collaboration system is a specialised accountability service. This service interprets the

---

11 "A Service-Oriented Architecture to enable virtual storage services: a dynamic collaboration context"; Shiping Chen et al; International Journal of Ad Hoc and Ubiquitous Computing 2009 - Vol. 4, No.2 pp. 95 - 107; DOI: 10.1504/IJAHUC.2009.023900

12 "Monitoring Contract Enforcement within Virtual Organizations"; Anna Squicciarini, Federica Paci, in Proceedings of CollaborateCom 2008 (Orlando,America)

agreed upon contracts, and once the collaboration starts, observes the participant's behaviours, intelligently generates and logs events that are in breach of the contract (either deliberately or accidentally). This service can, if required, be used to provide incontrovertible evidence of misbehaving entities within a system, and as such, be used as a basis for raising the trust levels for participants within the collaboration.

# 4. Conclusions

Consumer fraud is intimately tied to misplaced or simplistic trust of both the identities and behaviours of the parties involved in transactions, coupled with the assumptions that the underlying technologies used to exchange information are trusted and secure.

This report has brought forward two technologies that could be used to increase the trust (and security) levels for consumers interacting on the Internet. However, as emphasised in this report, due care must be taken so as not to carelessly proclaim, or worse still, deploy a particular technology and claim that it is "the solution" to preventing cybercrime.

The Internet was developed on a "minimal assurance" basis – it offers, at its lowest level, no assurances of correct, secure or even reliable delivery of data (and information). Further, its fundamental architecture and deployment is not going to change for a long period of time because of the scale, heterogeneity and huge financial investments put into its infrastructure. As such, all trust, privacy and security technologies must have this as a design factor.

Second, the complexities of the types of interactions that occur between people, organisations and governments on the Internet must also be carefully considered when a technological solution is seen to be the answer to protecting the consumer from cybercrime.

## APPENDIX A – TERMS OF REFERENCE

The House of Representatives Standing Committee on Communications shall inquire into and report on the incidence of cyber crime on consumers:

a) Nature and prevalence of e-security risks including financial fraud and theft of personal information, including the impact of malicious software such as viruses and Trojans;

b) The implications of these risks on the wider economy, including the growing economic and security impact of botnets;

c) Level of understanding and awareness of e-security risks within the Australian community;

d) Measures currently deployed to mitigate e-security risks faced by Australian consumers

   i. Education initiatives

   ii. Legislative and regulatory initiatives

   iii. Cross-portfolio and inter-jurisdictional coordination

   iv. International co-operation;

e) Future initiatives that will further mitigate the e-security risks to Australian internet users;

f) Emerging technologies to combat these risks.