

Trust in the Internet: Observations on Being Highly Connected

John Zic
CSIRO ICT Centre
INCO TRUST, New York, May 2010





**“the soul invariably travels
at the speed of a camel”**



**“the soul invariably travels
at the speed of a camel”**

**Our desire to communicate
is greater**

So we built the Internet

Dec 1969

4 destinations

University of California Santa Barbara; UCLA;
SRI; University of Utah

- 8 IP addresses
- 4 IP links

So we built the Internet

Dec 1969

4 destinations

Jan 2000

154,104 destinations

- 220,533 IP addresses
- 374,013 IP links
- 5,107 Autonomous Systems

So we built the Internet

Dec 1969

4 destinations

Jan 2000


154,104 destinations

Jan 2008

~ 48,000,000 destinations

- 4,853,991 IP addresses
- 5,682,419 IP links
- 17,791 Autonomous Systems

With growth comes complexity




THE ARPANET NETWORK


DEC 1969

4 NODES

IPv4 & IPv6 INTERNET TOPOLOGY MAP JANUARY 2008

AS-level INTERNET GRAPH





Easy, rapid interchange and aquisition of vast amounts of data

Driver (1) - What are the major issues and problems faced by your organisations? **

Control of information across and within organisations (including departments) when collaborating!

Realising solutions requires establishment and maintenance of:

- Trust relationships
- Information privacy
- Security

between collaborators.

** From members of CeNTIE2 Project *Enterprise Systems Focus Group*, September 2005 (2 major Australian banks & 2 government departments)

Driver (2)

So: who/what is going to save us?

“It's every man for himself,” he says, grinning.

In the end, it seems every machine has to defend itself.

The Internet was designed that way.”

Vint Cerf, 1 Oct 2008

<http://www.guardian.co.uk/technology/2008/oct/02/interviews.internet>

Jack Schofield meets Vint Cerf, the 'father of the internet'

This article was published on guardian.co.uk at 00.01 BST on Wednesday 1 October 2008. A version appeared on p1 of the Technology Guardian section of the Guardian on Thursday 2 October 2008. It was last modified at 00.09 BST on Thursday 2 October 2008.

Working together requires *trust*




Hxl project, 2008

NICTA, CSIRO, DSTO

The Internet Dog


“On the Internet,
nobody knows
you're a dog”



The Internet Dog

“On the Internet,
nobody knows
you're a dog”

Let alone *breed* ...




The Internet Dog

“On the Internet,
nobody knows
you're a dog”

Let alone *breed* ...

or *temperament* ...



Refining the meaning of trust

Something can be ***trusted*** when

- It can be unambiguously **identified**
 - It operates **unhindered**
 - A user has access to either:
 - First hand experience of its **consistent good behaviour**
- or
- Someone who can vouch for its consistent good behaviour.

Graham Proudler, HP Labs Bristol;

See also IETF RFC4949 *Internet Security Glossary, Version 2*

Trust from the systems viewpoint

- A Trusted System behaves *exactly* as specified and no more, despite
 - Disruptions by environmental factors
 - Errors cause by human or automated interaction, or
 - Hostile attacks on the system.

Examining the elements of trust

- *Identity*
- Unhindered operation
- Proof of good behaviour

Examining the elements of trust

- *Identity*

- Unhindered operation

- Proof of good behaviour



Examining the elements of trust

- *Identity*


- Unhindered operation

- Proof of good behaviour




Operational

Identity



Identity



Complete identity

- Relation over a set of *all information* associated with an entity that can be used to describe the entity's distinguishing features in *all possible contexts* of use.

Relative identity

- Relation over a *subset of information* associated with an entity in a *particular context* that can be used to uniquely characterise that entity

Nyms

- Term coined by Ian Goldberg** to represent the types of identifiers used on the Internet and the amount of information that they reveal about your identity.
- Taxonomy
 - Verinym
 - Persistent Pseudonym
 - Anonym
 - Linkable
 - Unlinkable

** Ian Goldberg, "A Pseudonymous Communications Infrastructure for the Internet", PhD thesis, University of California at Berkeley, 2001

Taxonomy of identifiers and identities

Verinym

- “true name” - an identifier that links to a sufficiently large subgraph of (complete) identity.
- Linkable, Persistent
- Examples (for humans) may include:
 - Passport
 - Driver’s License
 - Credit card number

Persistent pseudonym

- No linkage to a *relative* identity graph
- New “node” in *complete* identity graph
- Persistent; can determine that the same thing or individual carried out prior activities

Nyms

Anonym (no links to identity graph)

- Linkable
 - Limited, non-persistent identifier;
 - Can however determine that same thing individual carried out prior activity or transactions e.g.
 - Prepaid SIM card using cash to purchase
 - Loyalty cards in some circumstances
- Unlinkable
 - No identifier;
 - No way of finding out if same individual carried out some prior activity or transaction
 - Cash payment for shopping

Managing Identity Information

- **Complex!**
 - Needs a set of processes associated with creation, use and verification of information that links to an identity graph.
- **Includes processes for:**
 - Verification of supplied identifying information.
 - Authentication of a claimed identity.
 - Access control mechanisms for the control of available information based on authenticated identity.
 - The configuration of applications or systems to provide the access required in response to the provision of authenticated identity information.
 - The provision and registration of authentication credentials
 - The secure storage of identity information ensuring the availability of this information where and when it is needed.
 - The safe storage of identity information, including the ability to recover any information that may be lost, damaged, or corrupted.

Review

Something can be trusted when

- It can be unambiguously identified
- **It operates unhindered**
- The user has either:
 - First hand experience of consistent good behaviour
 - or
 - Someone who can vouch for consistent good behaviour.

Unhindered operation

- *Unhindered operation is synonymous with fault tolerance*
 - System/component needs to operate as specified, despite faults caused by errors that are a result of human or automatic responses, or deliberate malicious acts, or environmental factors
 - Requirement: identification of faults to be tolerated
 - Requirement: Need to understand and distinguish between potential *adversaries*

Operates unhindered despite ...

“... there are known knowns. There are things we know that we know.

There are known unknowns. That is to say there are things that we now know we don't know.

But there are also unknown unknowns. There are things we do not know we don't know.”

http://en.wikiquote.org/wiki/Donald_Rumsfeld -- sourced 23/09/09

Unhindered operation...

- Relies on critical examination of expected faults and attacks on the system
- Take appropriate counter measures when designing the system to allow it to progress despite the presence of faults
- Requires understanding of expected, or known, attacks and other misbehaviours
- In short: good architecture, design and careful engineering required to assure that Quality of Service metrics are met

Review

Something can be trusted when

- It can be unambiguously identified
 - It operates unhindered
 - **The user has either:**
 - **First hand experience of consistent good behaviour**
- or
- **Someone who can vouch for consistent good behaviour.**

Consistent good behaviour

Consistent good behaviour

Example: *A* claims to *B* that it always behaves well.
How can *B* be assured that this is true?

Possible approaches:

- Previous experience of *A*
- Complete characterisation of *A* by *B* (and v.v.)
- Current agreement between *A* and *B* as to terms and conditions of exactly constitutes “good behaviour”
 - Define and agree on “good behaviours”
 - Manage “bad behaviours” or exceptions
- Combination of the above
- Any of these may be done
 - by *B* directly or
 - another third party who *B* already knows and trusts.

Consistent good behaviour

Example: *A* claims to *B* that it always behaves well.
How can *B* be assured that this is true?

Possible approaches:

- Previous experience of *A*
- Complete characterisation of *A* by *B* (and v.v.)
- Current agreement between *A* and *B* as to terms and conditions of exactly constitutes “good behaviour”
 - Define and agree on “good behaviours”
 - Manage “bad behaviours” or exceptions
- Combination of the above
- Any of these may be done
 - by *B* directly or
 - another third party who *B* already knows and trusts.

Past

Consistent good behaviour

Example: *A* claims to *B* that it always behaves well.
How can *B* be assured that this is true?

Possible approaches:

- Previous experience of *A* Past
- Complete characterisation of *A* by *B* (and v.v.) Present
- Current agreement between *A* and *B* as to terms and conditions of exactly constitutes “good behaviour”
 - Define and agree on “good behaviours”
 - Manage “bad behaviours” or exceptions
- Combination of the above
- Any of these may be done
 - by *B* directly or
 - another third party who *B* already knows and trusts.

Consistent good behaviour

Example: *A* claims to *B* that it always behaves well.
How can *B* be assured that this is true?

Possible approaches:

- Previous experience of *A* Past
- Complete characterisation of *A* by *B* (and v.v.) Present
- Current agreement between *A* and *B* as to terms and conditions of exactly constitutes “good behaviour”
 - Define and agree on “good behaviours”
 - Manage “bad behaviours” or exceptionsFuture
- Combination of the above
- Any of these may be done
 - by *B* directly or
 - another third party who *B* already knows and trusts.

Previous experience...

- A record of a set of actions that *A* performs is kept for *B*
 - What actions are recorded, and what actions can be ignored?
 - Need for an agreement to be reached as to what is “good behaviour”
 - Where are the actions recorded?
 - Who records the actions?
- From this record, and agreed upon measure of good behaviour, *B* then assigns some confidence in the truth the claim that *A* always behaves well.
- The record can be used to develop a measure of *reputation*
 - Reputation based systems for vendors on eBay or Amazon
- The record can also be used to find, and act upon, *misbehaviours*
 - Accountability system
 - Provide undeniable evidence for resolving disputes

Example

- A Contract-based Accountability Service [ICWS 2009]
 - A schema for specifying accountable contracts between service providers.
 - A contract defined based on the schema is auditable in an automatic way
 - Protocols to support this model and define mechanisms for evidence management.
- What is in a contract?
 - The set of identified participants
 - A set of **states** that can be observed by all parties
 - A set of **activity logging events**
 - A set of **specifications** that relate activities to state transitions

Example misbehaviours

Misbehaviour	Defense mechanism
1. The service provider ignores or fails to commit a valid data state change request.	<p>A_c & A_p remain in the old data state after the corresponding “put_started” activity log expires. The faulty party is not distinguishable. It can be caused by either a service consumer does not submit the request in time, or the service provider ignores the request.</p>
2. The service consumer fails to submit a proposed operation.	
3. The service provider reverts or tampers data created by a service consumer.	<p>A_c & A_p can detect that the data is not consistent with the data state they observe when the data is requested via an accountable-get operation. The faulty party can be identified as c is provable not initiate the operation.</p>
4. The service consumer denies that it has submitted an operation.	<p>As a service request is signed by the service consumer, the fact is undeniable by its signature.</p>
5. The service consumer denies that it has submitted an put operation to the service provider at data state T .	<p>As the service provider needs a token from A_c & A_p to submit an accountable-put. The fact is undeniable through the logs in A_c & A_p.</p>
6. The service provider uses old tokens to revert data states (replay attack).	<p>A token carries a data state and a timestamp. The contract specifies the timeout value for a party to acknowledge a token. They can be used to defend replay attack.</p>
7. The service provider allows unauthorized get access from other parties to the data.	<p>Not detectable in our model.</p>

Complete characterisation

- Before any transaction between two parties *A* and *B*, an exchange of information is made characterising all their possible behaviours to each other.

- Requires secure transfer of information
- Requires authentication

If *measure is as expected*

- transaction proceeds

Else

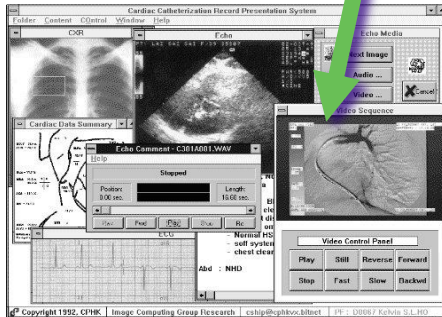

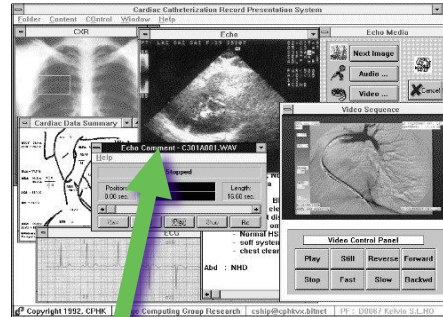
- transaction is cancelled.

- Fundamental basis for the widely deployed (but not utilised) *Trusted Platform Module*


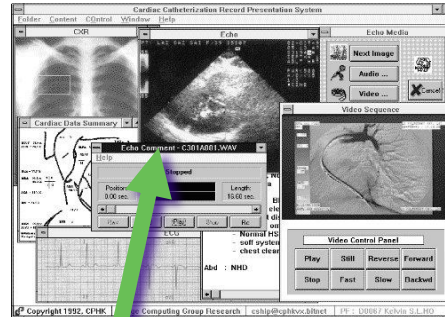
- Cryptographic microcontroller chip in most PCs
- Needs supporting infrastructure and protocols between components of the infrastructure to be effective

Example use of complete characterisation

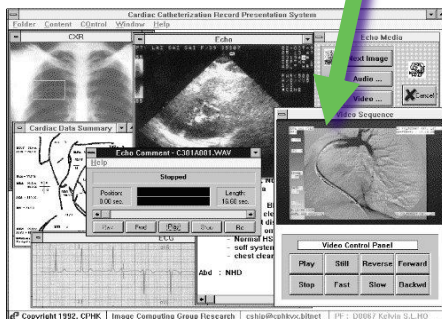
Example use of complete characterisation



Example use of complete characterisation



Secure, trusted transfer of medical records using Trusted Platform Module




Complete characterisation -- issues

- Any change to the operational environment requires updating all characterisation information
 - Applications -- new or versions
 - Operating System -- updates and patches
 - Device drivers -- new or revised
 - Hardware
- OK for extremely well managed systems
- *Not good* on general purpose machines
- **Alternative** – small, dedicated and controlled environment known by each participant
 - e.g. Trust Extension Device (TED); sensor network node (SecFleck)

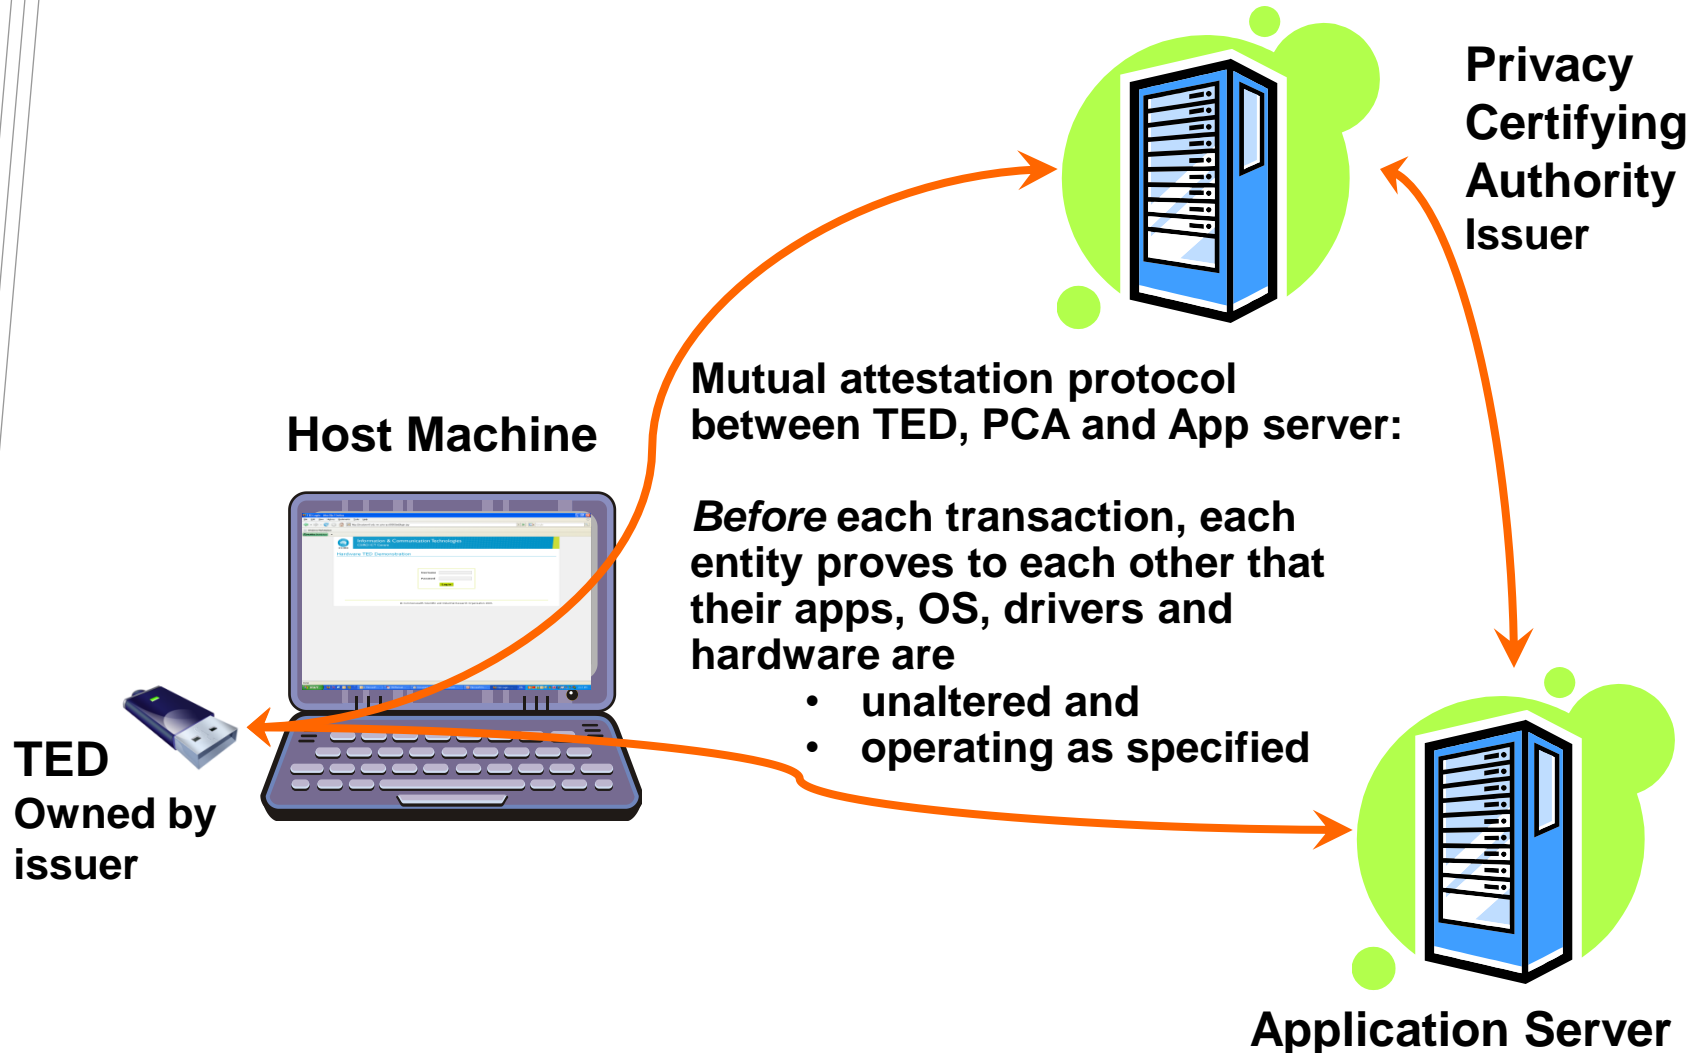
Trust Extension Device (TED)

- One solution to end-point trust
- Two instantiations: software TED + hardware TED
- Both identity and operational integrity (behaviour) are proven before any transaction or information is exchanged

Trust Extension Device (TED)



TED Enterprise Architecture



Future behaviour agreement

- Again requires a *Contract* between participants, managed by
 - Either the participants themselves,
or
 - a trusted third party
- Recall that a contract contains:
 - The set of identified participants
 - A set of **states** that can be observed by all parties
 - A set of **activity logging events**
 - A set of **specifications** that relate activities to state transitions
- Requires a complementary *monitoring and accountability* system
 - Contractual obligations are always satisfied
 - Exceptions (misbehaviours) are captured in an irrefutable manner

Examples

- **A Contract-based Accountability Service** [ICWS 2009, SCC 2010]
 - Set up between multi-party collaborations between enterprises
 - Identities known
 - Contract expresses agreed upon:
 - Policies for sharing
 - Workflows for collaboration
 - Work flows agreed
- **Accountable Internet Protocol** [SIGCOMM'08]
 - Self-certifying IP address
 - Routers are trusted parties
- **Accountable Network Storage (CATS)** [TOS 2007]
 - Use an external publishing medium to store evidence
- **Ensuring correctness over untrusted private database** [EDBT'07]
 - Periodically freezes the values of data items to generate proofs
 - The proofs guarantee the authenticity of the frozen values.

Conclusions

- Trusted interactions rely on three components:
 - Identity information
 - Assuring reliable operation despite outside interference or environment
 - Proofs that the system will always behave in a known manner
- The amount of interconnectedness, and ease with which information is accessible, offers many research challenges how to do this on a *sufficient* scale.

Conclusions

- Trusted interactions rely on three components:
 - Identity information
 - Assuring reliable operation despite outside interference or environment
 - Proofs that the system will always behave in a known manner
- The amount of interconnectedness, and ease with which information is accessible, offers many research challenges how to do this on a *sufficient* scale.

“It's every man for himself,” he says, grinning. In the end, it seems every machine has to defend itself. The Internet was designed that way.”

www.csiro.au

