

Designing for Trust

L. Jean Camp

Associate Professor of Public Policy

Kennedy School of Government

Harvard University

1. Abstract

Designing for trust requires identification of the sometimes subtle manner in which trust can be embedded in a system. Defining trust as the intersection of privacy, security and reliability can enable or simplify the identification of trust as embedded in a technical design. Yet while this definition simplifies, it also illuminates a sometimes overlooked problem. Because privacy is an element of trust, purely operational definitions of trust are inadequate for developing systems to enable humans to extend trust across the network. Privacy is both operational (in the sharing of data) and internal (based on user perception of privacy). Designing trust metrics for the next generation Internet, and indeed implementing designs that embed trust, requires an understanding of not only the technical nuances of security but also the human subtleties of trust perception.

2. Introduction

Trust is built into all systems, even those without security. Trust assumptions are built in when data collection is enabled, or coordination is made feasible. Trust is built in when resources are reserved (as shown by denial of service attacks). If trust is an element of all systems, what does it mean to design for trust?

Trust is a complex world with multiple dimensions. There has been much work and progress on trust since the first crystallization of this concept. Combining the three-dimensional trust perspective with studies of humans, I conclude that a new approach to understanding and designing mechanisms for peer to peer trust are critically needed. There is a need for greater understanding of how individuals interact with computers with respect to the extension of trust.

The first section of this work gives a quick overview of the alternative perspectives on trust: rational trust exhibited through behavior and internal trust which cannot be directly observed. The second section revisits the definition of trust offered in Camp 2001, by considering privacy, security, and reliability. At the end of that second section is an examination of trust might be built into whois. Thus at the beginning of the third section there is a clearly defined concept of trust as used in this work. Using that definition and building on the problems defined in the previous sections the third section argues for a trust system that allows users to aggregate trust, make transitive trust decisions, and manage their own electronic domains. This leads to the conclusion - that current trust management systems are hampered by designing for computers rather than humans. Trust systems for the next generation Internet must be built on solid conceptions of human trust drawn from the social sciences.

This work was supported by the National Science Foundation under Grant No. 9985433. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

3. Alternative Perspective on Trust

Multiple authors have offered framings of trust. In this section the three dimensional concept of trust is contrasted with other selected concepts of trust. Trust is a concept that crosses disciplines as well as domains, so the focus of the definition differs. There are two dominant definitions of trust: operational and internal.

Operational definitions of trust require a party to make a rational decision based on knowledge of possible rewards for trusting and not trusting. Trust enables higher gains while distrust avoids potential loss. Therefore risk aversion a critical parameter in defining trust.

In the case of trust on the Internet operational trust must include both evaluation of the users intention and the users' competence. Particularly in the case of intention, the information available in a physical interaction is absent. In addition, cultural clues are difficult to discern on the Internet as the face of most web pages are meant to be as generic as possible to avoid offense. In game theory competence is not at issue. A person is perfectly capable of implementing decisions made in a prisoner's dilemma without hiring a graduate of Carnegie Mellon.

In the three dimensional definition of trust privacy, reliability, and security are based neither entirely on intention or competence. Both good intention and technical commitment are required to ensure security. The result for the user (fraudulent use of data, usually to charge services) from a failure in either intention or competence are the same. Thus an operational approach arguably supports a focus on the types of harms resulting from trust betrayed¹.

One operation definition of trust is reliance. (Golberg, Hill & Shostack, 2001). In this case reliance is considered a result of belief in the integrity or authority of the party to be trusted. Reliance is based on the concept of mutual self-interest. In that way, reliance is built up the assumptions of human beings as homo economicus (Olson, 1965). Therefore the creation of trust requires structure to provide information about the trusted party to ensure that the self-interest of the trusted party is aligned with the interest of the trusting party. When reliance is refined, it requires that the trusted party be motivated to insure the security of the site and protect the privacy of the user. Under this conception trust is illustrated by a willingness to share personal information.

The definition of trust offered in (Camp, 2000) is operational. The concerns are with risk rather than risk perception.

Another definition of trust, popular among social psychologists assumes that trust is an internal state. (e.g., Tyler, 1990; Fukuyama, 1999) From this perspective, trust is a state of belief in the motivations of others. Based on this argument, social psychologists measure trust using structured interviews and surveys. The results of the interviews often illustrate that trust underlies exhibited behavior, finding high correlations between trust and a willingness to cooperate. Yet trust is not defined as but rather correlated with an exhibited willingness to cooperate.

The difference between these perspectives is a difference in conception of trust as being as basis rather than the behavior itself. To some degree this can be modeled operationally as the difference between perceived (e.g., internal sense of) versus measurable risk (statistical or deterministic). (e.g., Morgan et. al, 2002).

The definition of trust as a function of privacy, security and reliability is operational. It is based on risks rather than user perception of risk. This definition focuses on the existence or risk rather quantifying the risk. In that way it is neither deterministic nor stochastic but rather Boolean. From the operational perspective, privacy is a measure of willingness to share information. The assumption is that this willingness is based on the risk of secondary use of information rather than a psychological sensitivity to information exposure. Risks in the United States include loss of employment or medical insurance. Risks in the United Kingdom include loss of employment. In both nations medical issues are considered private. An internalized definition of trust would assume roughly equivalent sensitivity of information exposure in both nations assuming both had the same cultural sensitivity to medical privacy. An operational perspective would argue that medical privacy is more important in the US because the risks are greater.

¹ Betrayal is used in operational definitions in part because to choose not to cooperate is always a function of intent. The same ill intent or moral implications are not appropriate in failures of technical competence; however, the word is still useful for the results of trust ill-placed.

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlang (Berlin) 2003.

These definitions of trust will merge only when observed behavior can be explained by internal state. Yet without considering the internal state, and examining trust behaviors designs for enabling peer to peer trust over the digital network will be flawed.

3.1 The Three Dimensions of Privacy

Privacy is an overloaded word. No only political battles but philosophical debates rage about the nature and role of privacy. There are quite nearly as many taxonomies for the evaluation of privacy as there exist scholars interested in the subject.

In particular privacy is often distinguished based on data topic, jurisdiction, or spatial metaphors. Before moving forward into the conceptual taxonomy of privacy it is worthwhile to address why none of these approaches were selected.

Basing an analysis on a topical taxonomy would require the understanding of the political and cultural interaction for each topic.

The most notable topical distinctions are illustrated in the American mosaic of statutory privacy protections at the state and federal level. For example, the privacy of video rentals records has long been subject to national protection while national protection for medical information still lags. In the case of video rental records, the availability of the rental records for a Supreme Court nominee motivated action at the Federal level. In the case of compilations of medical records the largest compilation of medical records is organizationally base din Massachusetts, and thus effective regulation did not require federal intervention until the cost of medical data compilations decreased sufficiently to make other medical data compilations common. The operational theory requires presenting a rational apolitical unifying argument for this distinction, and thus would be left with an explanation of organizational theory and processes. The internal definition would allow for an explanation based on perceptions without addressing the larger organizational framework. Clearly neither is adequate.

A second common approach to the examination of privacy is based on jurisdiction. As travelers cross jurisdictional boundaries their privacy rights, indeed basic human rights, are altered. Any consideration of privacy in the case of the Internet must be sufficiently flexible in order to describe any legal regime of privacy, yet an exhaustive examination of privacy in the members of the United Nations would provide little guidance, as well as exceeding the patience of the reader.

A third concept of privacy is based on cultural concepts of space. Spatial privacy is of particular interest on the Internet because of the lack of cultural or social clues in virtual spaces. For example, disputes over domain names are exacerbated by the nature of the virtual spaces defined by those names. Virtual spaces differ from physical spaces with respect to simultaneity, permeability and exclusivity. (L. Jean Camp & Y.T. Chien, 2000). Permeability is the ability to move seamlessly between spaces. (Shapiro, 1998) Simultaneity is ability to move into one space without moving out of another - even when there is no overlap. For example, one my have multiple threads in discrete email lists, or view multiple new sources from a single framed browser. Exclusivity refers to the ability to create spaces that are not only private, but also invisible from the outside. (Nissebaum, 1999) Firewalls are aptly named, as the design goal of a firewall is prevent those outside from passing through to the inside. Clearly different privacy rules and expectations are appropriate for the marketplace, the avant-guard theater, and the home. Yet there is no single analysis that offers a single coherent theory about spatial privacy across the globe despite some progress on this track, and this paper is unlikely to move the frontier of the understanding of privacy. Beginning with an operational approach, I necessarily fall back on process and structure in order to define privacy. The American federalist legal system provides an effective parsing of privacy into those issues that are criminal and civil, and those which fall to the states as opposed to those which must be a matter for Federal consideration. In contrast the privacy regimes of Europe are designed to provide protection against violations of data protection. The data protection regimes can fit well within the taxonomy presented here if the distinct section only if the myriad acts are addressed under privacy as a human right and privacy as a property right.

Thus my operational framing and the carefully structured (if not particularly rational in outcome) American legal system offer a conception of personal data as property, and thus of privacy as a property right as well as a right of autonomy and seclusion. At the risk of self-plagiarism I consider the concepts of privacy as embedded in United States law.

Any design addressing privacy requires some definition of privacy that states clearly the perception of privacy built into the code. Of course the three dimensions I described here are no the only framing, but

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlang (Berlin) 2003.

any true framing would exclude some dimensions of privacy. Definitions of privacy such as those provided by iPrivacy in which transactions are said to be as private "as in the offline world" are meaningless. The offline world of political action, idle gossip or commercial transactions? As private as cash transactions or credit card transactions? By including the world in the definition, no concept of privacy is excluded. (See iPrivacy.com for definitions.)

3.1.1. Privacy As Autonomy - The Human Right

Privacy is the right to act without observations. People under constant surveillance are not free.

Arguments against privacy on the basis of autonomy often imply that the ability to act freely and without surveillance offers only the ability to commit those acts normally subject to social sanction. Privacy is sometimes presented as a moral good only an to the sinner. Yet privacy as an element of autonomy also enhances the public good. Few would argue that the right to privacy allowed by the Supreme Court "right of members to pursue their lawful private interests privately and to associate freely with others." in 1956 was a right to pursue justice. Yet at the time the members of the NAACP were seen by law enforcement as troublesome at best and subversive at worst.

Besides the historical arguments for privacy as autonomy for the greater good there are empirical arguments. Making this argument on the basis of empirical research requires three assumptions. The essence of these assumptions is contained in the second sentence of the first paragraph in the section. First, assume that the opposite of privacy is recorded surveillance. That is, not only is some act observed via surveillance but there is also a record of the act created. Second, assume that when privacy is violated the user is aware of that fact. (If this is true is the basis of some debate. Certainly some data compilations are obvious. and some technical mechanisms to obtain user information are devious.) Lastly assume that the existence of the record implies some ability to coerce either by rewarding good behavior or punishing bad behavior. (In this case good or bad can be defined by the party with surveillance capacities.)

Based on the three assumptions above, homo economicus would increase his or her good behavior. Yet the arguments that individuals respond in a strictly irrational way when faced with rewards (Kahan, 2001) or punishment (Lawler, 1988) are not reflected in empirical studies. When individuals are paid, required, or recorded in some "good" act the motivation to do that act decreases. A documented real world example of this is the drop in blood donations when individuals are paid (Titmuss, 1971).

Privacy as autonomy offers free people the right to act freely. It enhances not only the power to choose socially prohibited acts, but also the tendency to choose socially optimal acts. Surveillance alters action. The constraint on action created by observation is the basis of the autonomy right of privacy.

The Constitutional right to privacy is grounded in the First, Third, Fourth, Fifth, Ninth and Fourteenth Amendments (Compaine, 1988; Trublow, 1991).

The First Amendment states:

"Congress shall make no law respecting an establishment of religion, or prohibiting the free exercise thereof; or abridging the freedom of speech, or of the press; or the right of the people peaceably to assemble, and to petition the government for a redress of grievances."

The right to read is the right to read anonymously. (Cohen, 1996). The argument above suggest that people are not only less likely to go to assemblies and support organizations subject to official sanction. Additional recent in trust argues that in fact, people are less likely to offer their efforts to those socially sanctioned public actions. If every appearance at a social function is marked and credited, then the internal motivation is diminished. People are less free, less autonomous, and less active.

The Third Amendment states:

"No soldier shall, in time of peace be quartered in any house, without the consent of the owner, nor in time of war, but in a manner to be prescribed by law."

The Fourth Amendment asserts:

"The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no warrants shall issue, but upon probable cause, supported by oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized."

The Third Amendment illustrates better than any story or argument that privacy is no strictly modern, rather when privacy was rare it was arguably more valued. Yet when combined with the Fourth Amendment, the Third Amendment creates of space safe from government intervention or casual surveillance.

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlang (Berlin) 2003.

The element of the Fifth Amendment that is relevant to privacy states:
"No person shall be compelled in any criminal case to be a witness against himself, nor be deprived of life, liberty, or property, without due process of law; nor shall private property be taken for public use, without just compensation."

In terms of privacy the limits on forced testimony are of greatest interest. One cannot be held to testify against oneself. Neither can a person's movements nor his thoughts be constrained by governmental force.

The Ninth Amendment states that the set of Constitutional rights is neither exclusive nor exhaustive. The Ninth Amendment allows the right to privacy to exist in a Constitutional sense. Under the Fourteenth Amendment (which is primarily concerned with Civil War debt) states cannot remove those rights given the individual by the Federal government.

The question with respect to rights of autonomy on the Internet are with respect to economic and corporate power. The coercive power of the state was well-recognized by the eighteenth century. Yet the modern corporation did not yet exist. The ability to gather information and violate privacy was held by the state alone until the rise of the popular press in the nineteenth century. Because of the First Amendment, weak privacy rights were necessarily trumped by strong speech rights. Yet the debate on the Fourteenth Amendment asks if the state has a positive responsibility to guarantee those rights, or simple the responsibility not to violate them directly.

When building a system specific to digital government, an understanding of autonomy is required. Yet the legal understanding of autonomy in the commercial corporate world is yet inchoate. Therefore technical designs that address the issues of autonomy must meet a high standards in order to avoid failure. Systems designed to provide autonomy include Zero Knowledge's Freedom Suite and the Anonymizer. (Camp and Osorio, 2002).

3.1.2 Privacy and Seclusion: The Right to be Let Alone

"The right to be let alone." Warren & Brandies' alliteration of privacy has come to be a definitive work. A century and half later the work was either refined (Prosser, 1941) or destroyed (Bloustein, 1968.) by determine that violations of seclusion consists of four possible torts: intrusion upon seclusion, appropriation of name and likeness, false light, and public disclosure of private facts.

Each of these torts is framed by the technology of the printing press. Understanding their meaning in a networked digital world requires a reach across an economic and technological chasm. In fact, the work singled out the merging popular press for reprobation: "Gossip is no longer the resource of the idle and of the vicious, but has become a trade which is pursued with industry as well as effrontery." (Warren and Brandies, 1890). Now gossip is not only the vocation of journalist but also the avocation of many with a modem.

Appropriation of name and likeness may include names in meta-data in order to associate with a more successful site. It may include the use of domain names in order to obtain the attention of those seeking a related site, as when Colonial Williamsburg was used by the service employee unions. (Mueller, 2001) Or it may include the use of a person's name or publications in order to attract those interested in related materials. Yet such appropriations are treated very differently. Meta data is not generally actionable while domain names have been subject to action based on trademarks and personal names.

False light is so common on the web that making it actionable seems impossible. When everyone is a journalist, everyone has the rights to frame content. Private persons need show only falsehood, yet how can one be an active participant in networked conversations and remain a private participant? False light is entirely content based. It is unlikely that false light would ever be used in a design decision.

Public disclosure of private facts implies individual control over information. Registration systems that send information about the user (also known as spyware) arguable violate this concept of privacy. Spyware is used in browsers and peer-to-peer systems including Kazaa and Limewire. The sharing of this information and targeting of ads provides the financial incentive for the systems to continue to function. Arguably the networks would not exist without the spyware. Yet the design for trust perspective would allow such designs only if the systems were easy to delete, and adequate notice was part of the design. Adequate notice may be as simple as allowing a add-on to be disabled during use rather than asking for a one-time installation permission

3.1.3 Privacy as Property: The Right of Ownership

For those who believe that privacy is property, what is required is a fair trade for private data. Much of the legislative debate about privacy concerns the existence and intensity of concerns about privacy. Observations of diffusion of the Internet commerce are in contrast with surveys identifying privacy concerns.

The privacy as property argument is enhanced by considering private information as a form of intellectual property. In that case the transfer of data subject to data owner is fairly conceptually simple.

the concept of privacy as property can explain this conflict. Individuals are ready to provide information to Amazon. Amazon decided that legal risk prevented the personalization and affinity marketing provided by user data. Therefore Amazon issued a privacy policy removing all possible expectation of privacy from users. The Free Software Foundation and Computer Professionals for Social Responsibility issued a call for a boycott. Amazon was only marginally affected. Amazon used consumer information for consumer benefit.

Geocities in conflict used consumer information only for the benefit of Geocities. Geocities, like Amazon, depends entirely on customer relationships. After the Federal Trade Commission announced that Geocities had substantially violated the privacy of it's the total value of Geocities fell nearly \$1,000,000 for each minute that the stock market remained open. Geocities never recovered the value.

If privacy is property then programs that send personal information or trap personal information in this case are theft. In that case the most basic moral frameworks are all that is required in designing for privacy. Yet this hypothesis may be rejected in the face of the decades of effort put into security without a dependence on identity.

3.2. Trust as Security: Distinguishing Privacy and Security

Security is not privacy. Confidentiality is an element of both security and privacy, yet this commonality does not them the same. Confidentiality allows a person to communicate to another without eavesdroppers. In the general case, the control of information enabled by security does not imply privacy. Security enables the control of digital information, while social and organizational forces determine who exercises the power of that control. Privacy requires that a person be able to control information about his or her self.

Security provides to privacy the ability to generate privacy in a specific case (as with confidentiality of communication). Security also provides the capacity for cryptography. Cryptography is the art of hiding information. When the information that is hidden is identifying information then security can be said to provide anonymity. Anonymity is a technical guarantee of privacy.

Thus, unlike many social values, the concept of privacy has an excellent mapping into implementation because of anonymity. Yet the simplicity of removing individual names is misleading. Inclusion of date of birth, current residence and place of birth will uniquely identify most Americans.

3.3 Trust as Reliability

Trust implies more than secure endpoints -- it requires that such security not come at the expense of survivability. Two of the greatest strengths of the Internet Protocol are that it is distributed, and it exhibits graceful degradation. Graceful degradation means any person can connect to the a network without altering others' access, and the loss of one machine should not effect those not using its services. Even during the most effective assault to date on the Internet, the Morris worm incident, staying connected proved to be the best strategy for recovery. Obtaining defenses against the worm, and information regarding these defenses, required remaining connected. Those who disconnected were isolated, with only their own resources to develop defenses. The ability of any network -- the Internet or an intranet -- to degrade gracefully rather than suffering catastrophic failure a critical component in survivability.

Trust architectures have developed significantly in the past decade. Yet despite that innovation, trust and security systems come at the cost of reliability and survivability. Security systems (as well as a lack of security systems) both enable denial of service attacks. Security systems that are computationally intensive or intolerant of user input make a user experiences of a lack of reliability more likely.

An element of design for trust should be designing the survivability of distributed trust mechanisms. Proposals for trust include short-lived attribute-specific certificates (M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis, 1999); long-lived multipurpose certificates (Anderson, 2001;

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlang (Berlin) 2003.

Sirbu & Tygar, 1995); certificates signed by multiple parties (Visa, 1995); a Web of Trust (Garfinkle, 1994) and or a combinations of these into a Web of Hierarchies. Yet other than the Web of Trust, few of the distributed trust mechanisms have been evaluated with respect to their ability to recognize an attack, reduce the damage of any attack, and subsequently recover. To design for trust, it is necessary to determine if, and under what conditions trust mechanisms are brittle.

4. A Design for Trust Application: the case of whois

Were whois to function as designed there would be no privacy considerations. Recall that the design goal of whois is to provide technical information in the case of technical errors or malicious action. The goal of whois is to provide contact information for technical action.

Yet the Internet has changed, and the administrative structures of the Internet have changed as well. Therefore looking at whois enables us to examine a system at a point in time, economics, and politics vastly different than the one in which it was designed.

whois was designed for narrow technical contact information for a relatively small Internet where the users were predominantly technical. Additional fields were added to whois when the trust assumptions began to fail. Additional fields are administrative and billing contacts. Had the trust model implicit in whois been recognized, the lack of wisdom in adding the additional fields would have been obvious. A technical contact would be appropriately contacted if a server were taking part in a DDoS attack. Yet the webmaster or billing contact would be appropriately contacted if content in a web site were under dispute.

The additional fields in whois are useful primarily by enforcement authorities. A significant problem with the traditional approaches to obtaining law enforcement information is that web sites cross jurisdictions. There already exist treaties and cooperation in terms of obtaining subscriber information from telephone companies across the borders of jurisdictions. Such policies, worked out over more than a century, provide a basis for law enforcement to obtain information. However, such information and policies do not naturally apply to whois specifically and would be only confused by application to the DNS system in general. As the Internet and traditional network services converge, and the possible business and legal arrangements between a network service provider and content provider explode. Thus attempting to track those complexities in whois seems unlikely to be feasible as well as being unwise.

By limiting whois information to technical contact and the appropriate registrar, motivation for incorrect contact information will be significantly decreased. Default automated access to whois information could reasonably be limited to those with network responsibilities. Feasible limitation of automated access to whois, and thus the ability to increase the integrity of the information, may require technical coordination at a level the holders of whois information have yet to achieve. A necessary first step for cooperation is trust. The trust may be enabled by removing the spotlight brought to bear by the unwise expansion of whois, and thus decreasing the resulting focus of intellectual property and other enforcement authorities on whois.

In addition to the changes in community the domain name itself has changed. Originally simply a mnemonic the domain name is now commercial property, political speech, personal expression or artistic moniker. As a result very different models of privacy apply. It is these differences in privacy models that is the core cause of the trust models in whois. It is unlikely that IBM.com considers the contact information in the domain registration as constraining institutional autonomy in the political domain. etoys.org was notoriously noncommercial.

The trust failure is a function of the expansion of whois to include billing and administrative fields without reconsidering the core trust assumption: that all Internet users are created equally powerful. Billing and administrative contact became necessary as the use and users of the Internet, and thus the trust relationships on the Internet, were changing.

In this case the original design was narrow and suitable for the initial environment. Failing to expand the function and fields of whois beyond the minimal necessary technical requirements would both have served the whois system more effectively and allowed the trust assumptions to remain valid in the rapidly changing realm of the Internet. This is because the trust framing was for technical individuals empowered over some small section of the network. By limiting the fields to technical information, that trust model would have been more likely to remain consistent, and therefore the service was more likely to remain effective.

5. Design for Trust

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlag (Berlin) 2003.

At this point there is a clearly articulated concept of trust as consisting of privacy, reliability and security. Also there has been one small example, arguing that design for trust would have resulted in a more limited and possibly more reliable whois. IN this section that modest core is expanded to a broad call for trust systems that are multidimensional, transitive, and aggregate.

Trust in today's Internet is based on all-or-nothing trust relationships. A network resource request is not trusted before authentication, and after authentication it is granted the full credentials of the corresponding user. Executable content from within a protected network is completely trusted, but content from outside the firewall is strictly disallowed. A network connection, once established, has equal priority with all other network connections on the system. These all-or-nothing trust relationships fail to match the expectations of users and the needs of next generation network applications. This mismatch promotes security breaches among users, as users undermine simplified trust models to meet their own complex resource-sharing needs. As for the specific example of executable content, it is one of the keys to providing advanced functionality in network applications, but is typically disallowed by firewalls. The firewall model of trust is too simple to distinguish secure sources of executable content. When sophisticated users find this exclusion unacceptable and use methods like tunneling to work around it the security of the entire protected network can be compromised. There is a need for distributed trust modes that will allow distinctions to be made in the trustworthiness of network entities. In order to do this it is necessary to provide a better match both to peoples' intuitive notion of trust and match that understanding with the needs of next generation applications.

Security in today's Internet is focused on a centralized model where strong security requires a firewall. The firewall may be a formidable obstacle, but once it has been compromised the entire network that it protects is compromised, making the firewall a single point of failure. The tunneling example demonstrates how this centralized approach can allow a single breach of security to compromise the security of the entire protected network.

The Microsoft/Verisign approach to regulating executable content is to centralize trust. In this approach, a presumably trustworthy third party uses a digital signature to verify the identity of an executable module. Although there is some commonality in purpose, their security model is the antithesis of most human approaches. It assumes that the same level of trust appropriate for all and give a binary approval to some developers. Further, it requires users to manually examine the source of executable content to provide more subtle variations of trust.

Currently proposed cross-domain trust mechanisms seek to minimize computational costs and management overhead. For example, commerce systems minimize key generation by linking all attributes and rights to a single commerce-enabling certificate. These keys are validated by a single root. This creates a single point of failure for the entire system (the root) as well as a single point of failure for the consumer (the key). The only similar system in the United States is the currency system, where the failure of the US Treasury would yield complete collapse. In family systems, individual businesses, and even religions there are multiple levels and power points. In physical security, any key is a part of a key rings, so that the failure of the validity of one key does not destroy the strength of all electronic locks. .Net ("dot net") or Passport exacerbate this problem by allowing cross-domain failure from a single lost pass phrase.

SSH and SSL are used for securing Internet connections. SSH is commonly used to provide secure terminal connections, whereas SSL is commonly used to implement secure HTTP connections. The endpoints of these connections have to be ready to extend trust before the mechanism are called into play. They are extremely useful technologies for the prevention of snooping but are not useful for implementing organizational or individual trust across many dimensions (including time).

Yet in real life and in social networks the "security models" (including drivers licenses, check clearing, credit cards, etc.) distribute the resources that implement authentication and authorization. In network security there are still single roots, and control is often held in a centralized point of control.

A network service can be rendered unusable when the number of requests it receives exceeds the rate at which they can be served. This creates an important relationship between performance and security for Internet servers. Although users on today's Internet are accustomed to server failures due to overload, the NGI will require better service guaranties. Decentralization is necessary to provide stable peak performance, even when a site as a whole is experiencing overload, until network capacity becomes the limiting factor. Decentralization provides defense against a large class of denial of service attacks. In contrast, overload in conventional systems typically results in thrashing behavior such as paging, leading to significant performance degradations.

Decentralization requires utilizing processing power at the endpoints more effectively. Decentralized trust requires enabling users to be their own trust managers. There is a need for a peer-to-peer distributed trust mechanism in order to implement increasing scalability of the network. The network needs to scale not only to an increasing number of devices but also in terms of complexity of tasks. Yet as there are increasingly complex interactions and task on the network, simplicity is critical to user-managed resource-specific security.

In order to allow users to share information it is necessary both to communicate trust states and allow users to manipulate their own internal trust states. Trust must support the complexity of life, in that users function in multiple dimensions. For example, a spouse will have access to all shared family and personal information. Yet a spouse should not have access to all company and employer information. Trust in these two dimensions is managed off-line because of the reality of physical space.

In addition to having multiple dimensions, users should be able to aggregate within a dimension. With aggregate trust the initial extension of trust is based on the introduction, which is provided by any entity or security mechanism. Any additional extension of trust is then based on aggregating different mechanisms (e.g., attaching value to different attribute-based certificates and summing) and/or extending trust to a machine based on interactions over time. Such a mechanism would be modeled more on observed social networks than on the strengths of cryptography. Users who find multiple independent paths to another user would increase the trust to that person accordingly, in a more generous manner than proposed in (Beth, Borcharding, & Klein, 1994).

In short, trends in distributed system security computing are on a collision course with system survivability through the construction of brittle trust mechanisms. The lack of understanding of the human interface exacerbates this problem. If trust extensions are not effectively communicated to the very human users, those users cannot react effectively when and if the trust system fails.

An early example of a user-centered approach to distributed trust, the UNIX philosophy gives users responsibility for setting security controls on their own resources. For example, Unix systems allow users to set file protection level. Yet neither of these approaches are not adequate for two reasons. First, for those using UNIX based system the security mechanism is hampered by its lack of simple mechanisms for authentication and resource sharing across domains.

The UNIX security system requires subtle understanding of the operating system and the interface violates the rules of good human-computer interaction (HCI) design. Truncated commands (e.g., `chmod`), a text line interface, and obscure error codes make this model infeasible. In addition the function has too many parameters, and these parameters are not clearly specified. For these reasons, even if they were well implemented UNIX file protections fail to meet the needs of the modern Internet user.

Similarly peer to peer systems allow users to determine which files are be shared. Peer to peer systems are built to implement coordination and trust across administrative domains. Peer to peer systems allow for sharing trust across domains, yet are notoriously hampered by problems of accountability. (e.g., Oram, 2001) Peer to peer systems allow users control over their own files in a more trans[parent manner than UNIX controls, but the P2P code itself is often untrustworthy. (e.g., Borland, 2002)

Any optimal trust approach would benefit from experience with Pretty Good Privacy (PGP), which lets users increase trust in a transitive manner. Transitivity means that users select their own sources of validation; e.g. if A trusts B and B validates C, then A trusts C. There is no central server of tree-like hierarchy that validates users. PGP also lets users select their own sources of trust, and select a key length appropriate for the situation. PGP is specific to a single application, electronic mail. In PGP users select specific individuals to trust based on their ability to verify the identity/key carried in a PGP certificate. This research extends and enhances the distributed security model of PGP to the more generic problem of sharing resources.

PGP is weak in that there is a single dimension of trust. Regardless of the definition of trust, it is certain that there are different dimensions of trust. Social admonitions not to mix friendship and money illustrate this, as well as concepts of family trust versus trusting in a business transactions. Trusting one's sister and trusting IBM are very different matters indeed. Users should be able to express trust in more dimensions, more richly, than with PGP. Yet unlike whois, PGP has maintained its efficacy by refusing to expand beyond its design base of email.

The attempt to minimize system management is a fundamental error because it is doomed to fail in a world of increasingly complex trust arrangements. Oversimplified security paradigms which limit implementations will result in users subversion. Security management should be distributed and simplified by automation, rather than simplified at by the administrative assumption of a single trusted entity.

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlang (Berlin) 2003.

Humans are capable of managing quite complex tasks (consider in the abstract the task of driving an automobile) if enabled by an interface that provides adequate and useful feedback.

Rather than minimizing computational and management costs, future trust designs ideally will recognize the high value of distributed security and empower the resource owner to be a security manager. Security management must become more complex in the Next Generation Internet because peer-to-peer, international resource sharing is more complex than intra-network sharing. Peer to peer systems recognize the need to share resources, yet the trust problems in peer to peer systems have not been solved. In fact, in 2002 most trust systems require users trust a central software distributor or administrator. The trust problem has only begun to be solved.

In order to provide simple mechanisms to enable users to take responsibility for their own resources, the design must implement an understanding of trust based on an understanding of trust among human users and social networks. While such a design basis may appear initially too complex for implementation, such a model would inherently provide better scalability and better resistance to attacks than the current, popular, centralized model.

6. Conclusions on Design for Trust

Experts focus on the considerable technological challenges of securing networks, building trust mechanisms, and devising security policies. Although these efforts are essential, that trust and security would be even better served if designs more systematically addressed the (sometimes irrational) people and institutions served by networked information systems. In order to address human concepts of trust, privacy must be a consideration and not an enemy or afterthought of the implementation.

Efforts at securing systems should involve not only attention to machines, networks, protocols and policies, but also a systematic understanding of how social agents (individuals and institutions) participate in and contribute to trust. Security is not a separable element of trust. An interdisciplinary perspective will enable protocols for trust over the network to be optimized for human trust.

That the human is a critical element in security systems has been recognized both from a usability point of view (Whitten, 1999) and from the analysis of systematic failures of security (Anderson, 1994). However, little work integrates methods from the social sciences, philosophy, and computer science to evaluate mechanisms for trust on-line. Previous work on integrating privacy and security (Friedman, Howe & Felton, 2002) has been complicated by the lack of a definition that can be used across disciplines. Efforts have been made to find a single definition of trust that can be used effectively within philosophy, computer security, and those social scientists embracing an operational definition of trust, as shown in (Camp, McGrath & Nissenbaum, 2001).

Design for trust requires examining all assumptions about a system and the user of the system. Sometimes those assumptions are based on class (e.g., the user has a credit card). Sometimes those assumptions are based on the capacities of the human (e.g., the user must select a large number of context-free random passwords). Sometimes the assumptions are necessary to enable a functioning design.

Design for trust requires enumerating the social assumptions and examining how those assumptions can function to put some user of the system at risk. In order to understand and design trust systems, acknowledgement of the social and human elements are required.

Bibliography

- R. Anderson (2001) *Security Engineering*, Wiley (New York, New York).
- Thomas Beth, Malte Borchert, Birgit Klein 1994 Valuation of Trust in Open Networks, Proc. 3rd European Symposium on Research in Computer Security -- ESORICS '94
- M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis. (1999) "The role of trust management in distributed systems security" in *Secure Internet Programming*, vol. (1603) of *Lecture Notes in Computer Science*, pp. 185-210. Springer-Verlag Inc. (Berlin).
- Bloustein, A., 1968. Privacy as an aspect of human dignity: an answer to Dean Prosser. *New York University Law Review* 39:962-970.
- J. Borland (2002) "Stealth P2P network hides inside Kazaa", CNET Tech News, April, 2002. <http://news.com.com/2100-1023-873181.html>
- L. Jean Camp (2001) *Trust and Risk in Internet Commerce*, MIT Press (Cambridge, MA).

L. Jean Camp, "Design for Trust," *Trust, Reputation and Security: Theories and Practice*, ed. Rino Falcone, Springer-Verlag (Berlin) 2003.

L. Jean Camp & Y.T. Chien (2000) "The Internet as Public Space: Concepts, Issues and Implications in Public Policy", Readings in Cyberethics, eds. R. Spinello and H Tavani, Jones and Bartlett Publishers (Sudbury , MA) January 2001. Previously published in ACM Computers & Society, September 2000.

L. Jean Camp, Cathleen McGrath & Helen Nissenbaum, "Trust: A Collision of Paradigms", Proceedings of Financial Cryptography, Lecture Notes in Computer Science, Springer-Verlag (Berlin).

L. Jean Camp & Carlos Osorio, (2002) "Privacy Enhancing Technologies for Internet Commerce", Trust in the Network Economy, Springer-Verlag (Berlin) Winter 2002.

Cohen, J. 1996, "A Right to Right to Read Anonymously: A Closer Look at Copyright Management in Cyberspace", 28, Conn. L. Rev. 981.

Compaine B. J., 1988, Issues in New Information Technology, Ablex Publishing (Norwood, NJ).

Friedman, B., Howe, D. C., and Felten, E. (2002). Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. Proceedings of the Thirty-Fifth Annual Hawaii's International Conference on System Sciences. Abstract, p. 247; CD-ROM of full-paper, OSPE101. IEEE Computer Society: Los Alamitos, CA.

Fukuyama F. (1996), Trust: The Social Virtues and the Creation of Prosperity, Free Press, NY, NY.

Golberg, Hill & Shostack (2001) "Privacy Ethics and Trust", Boston University Law review, Vol. 81, N. 2 April. pp. 407 -422.

D. Kahan (2001) "Trust, Collective Action, and Law", Boston University Law review, Vol. 81, N. 2 April. 333-347.

Simson Garfinkle, (1994), Pretty Good Privacy, O'Reilly Publishing (Cambridge, MA).

E. J. Lawler, (1988) "Coercive Capability in Conflict: A Test of Bilateral versus Conflict Spiral Theory", Social Psychology Quarterly, Vol. 50, pp 93-96.

M. Granger Morgan (Editor), Ann Bostrom, Baruch Fischhoff, Cynthia J. Atman (2002) Risk Communication : A Mental Models Approach, Cambridge University Press, (Cambridge, UK).

Olson (1965) The Logic of Collective Action: Public Goods and the Theory of Groups, Harvard University Press (Cambridge, MA).

Oram, ed. (2001) Peer-to-Peer Harnessing the Power of Disruptive Technologies, O'Reilly and Associates (Cambridge, MA).

S. Shapiro (1998) " Places and Space: The Historical Interaction of Technology, Home, and Privacy", The Information Society, No. 14, Vol. 4, pp. 275-284.

Sirbu, M. and Tygar, J. D. 1995. "NetBill: an Internet commerce system optimized for network delivered services". IEEE ComCon. San Francisco, CA. 6 March, 1995.

Richard M Titmuss (1997) The Gift Relationship: From Human Blood to Social Policy, Expanded and revised edition, Ann Oakley and John Ashton (eds.) The New Press, (New York, New York).

Trubow, G., ed. (1991) Privacy law and practice. Times Mirror Books (New York, New York, US).

T. Tyler, (1990) Why People Obey the Law. Yale University Press, (New Haven, NH).

Visa. 1995. Secure transaction technology specifications. Version 1.1, Visa International (New York, New York).