Trust as Qualified Reliance on Information

Ed Gerck
egerck@nma.com

Summary

"If the world were really random, chemistry, cooking, and credit would not be possible, so our models cannot be figments of our imagination."
P. Cheeseman, "Finding the Most Probable Model," p. 91, 1990.

Editor's Introduction

Trust is a word that is commonly applied to many situations and consequently has many shades of meaning. The following essay by Ed Gerck focuses on one precise set of coherent meanings: the concept of trust in the context of communication. More specifically, in the context of the engineering problem of Internet communications. At the same time he demonstrates why trust is needed in this context. Trust is considered something essentially communicable, but with specific rules for its communication. Gerck's exposition also discusses the induction (communication)of trust in heterogeneous environments, from human to machine, machine to machine, and machine to human.

By allowing trust to bridge the many gaps between human and machine, people will be able to tailor their own human to human communication needs via the Internet. What this means is that communities of interest, as networks of people, can build their own networks within the Internet according to their needs, without any limitation imposed on them by their Internet connectivity.

No one is at the edges of the Internetwork, while everyone is at the center of their own network. In this sense the flat, edge controlled Internet that we wrote about in the lead article in the December 2001 *COOK Report* is really just a local vision of a multi-dimensional network of networks made up of many different user groups and their networks, who actually act as control centers of such local networks.

Einar Stefferud observes:

I have known Ed Gerck since 1997. I have discussed with him and read many of his previous papers on trust; so this essay serves to bring together many different threads that we have discussed on- and off-line since 1997. So I now see that all my previous talk about the Net being edge controlled needs to be revised in some new framework. In short, the Internet does not really have a center or edges. It only has connection points, each of which can be connected to any other such connection point for the purpose of packet exchange. One reason that the Internet does not have an edge (as I just realized) is that at any termination connector, it is possible to extend the Net beyond that point by relaying packets, or by relaying messages, via dial-up modem, FAX or printer, or word of mouth, for that matter. So we suddenly discover that we cannot define any edge of the Internet.

Gerck's communication concept of trust may be just the beginning of a broad-based understanding of a new view of the Internet where security is a core part of the design. This new view of the Internet is that of a large collection of local network centers and edges, of potentially overlapping subsets of the total Internet. It is built around local common interests and purposes (communities of interest), but with a global communication pattern that closely resembles how we humans communicate across such boundaries and how our commerce works; and it looks like an assembled collection of networks, each of which has its own local centers and edges when we observe them closely, but the global collection of these local networks into the whole Internet doesn't have a single global center or any edges.

Now, since we have so many available connections, Gerck is saying, let's use sets of connections to enable us to transfer trust using distinctly separate multiple channels. Except that, in the essay that follows, he leaves for a next article the discussion about how one can use those multiple channels to induce trust, and how many channels to use. First, one needs to establish the need to use multiple channels, before explaining how to use them.

The problem is that if the Internet is this thing with users and servers attached to its interconnection spigots with nothing but connection pipes between them, and where any attached user or system has protocol-based communication access to all others so mounted, then we must ask what controls the whole thing? And where might we mount a controller for the whole Internet?

The essay explains that the answer depends on how you use the communication concept of trust. You may choose to be at an edge of some local network by joining a mailing list or participating in a message board on some website, or subscribing to some information services, trusting that which has been authorized for you. Or you may choose to be at the center of your own network where you control the nature of all the connections. Or you may choose to be at both, edges and center, and from this position you will be able to realize the full potential of the Internet. However, quite independently of your choices, the Internet is still just a Network of Networks—which prevents anyone, from anywhere, from hosting a single control center for all elements of all networks in the Internet. We do not even know what local networks exist inside the Internet.

Trust as Qualified Reliance on Information by Ed Gerck

When I say that the key to solve the fundamental problem of Internet communication is trust, I usually get two reactions. The first is "what is the fundamental problem of Internet communication?" The second is "what is trust?"

Let's answer these questions.

In 1948, Claude E. Shannon created information theory and stated that the fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one selected from a set of possible messages. The system must be designed to operate for each possible selection. [quoted from Shannon, C. "A Mathematical Theory of Communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379-423, July 1948. Available at http://cm.bell-labs.com/cm/ms/what/shannonday/paper.html]

Fast forward to 2001. The fundamental problem of Internet communications is that of reproducing at one point exactly the message selected at another point. No one can control at the same time both ends of a connection to another party, neither sending nor receiving. The route followed by the messages cannot be controlled or positively verified by any party. All messages have meaning relevant to the engineering problem of transferring the bits, but with different meanings, some of them intertwined, at different protocol layers; that is the messages contain packets, headers and fields that need to be formatted, interpreted, or verified, in correlation with references, by suitable protocols at each end and en route.

Reading the two paragraphs above, we realize that Shannon's information theory fails to model Internet communication. Internet protocols are much closer to human communication than to Shannon's idealized communication systems. We, speakers of the same natural language, communicate with one another by trading contents, not by exchanging uninterpreted strings of symbols (bits). Each bit of information sent by a human to another must either contribute to the content or be discarded. Content may have different meanings, some of them intertwined, at different layers of our understanding. And, in the same way that content must be conveyed in human-to-human (H2H) communication, we find that content must also be conveyed at different protocol layers in Internet machine-to-machine (IM2M) communication – not just bits. In internetworking, machines are not just trading un-interpreted strings of symbols, or bits. They are trading bits and meaning, machine to machine. They are talking.

Before anyone thinks that I intend to turn Internet engineering into some sort of socio-philosophical-techno babble, let me comment that the objective here is to discuss a technical solution to the engineering problem. We are still happy with Shannon's definition of information as a measure of the decrease of uncertainty at a receiver. In other words, information is what you do not expect. However, the problem has now an added dimension. We must be able to convey meaning in IM2M communications. But this meaning is not the same meaning conveyed in the H2H communications using those same machines, and not the same meaning at every protocol layer either. Meaning must be conveyed in heterogeneous environments, from human to machine, machine to machine, and machine to human.

Introducing meaning into information theory, so that the communication of meaning can be described, has been an open problem since 1948. I assert that the way to communicate meaning is to first communicate trust and bits, and then use them to define the meaning. This may sound like jumping from a frying pan into the fire, because we must still communicate trust. However, my assertion is based on the observation that trust is essential for H2H communication and needs to come first before we can rely on the contents being communicated. Thus, since we can readily observe that communication processes in general are in many ways very close to H2H communication, as exemplified by the IM2M communication discussed above, we should expect that trust may also be essential for communication in general and also needs to come first. In other words, rather than introducing meaning into information theory, we introduce trust. Meaning will be introduced and conveyed implicitly.

But what would trust be in the context of IM2M communication? Or in the context of communication processes in general? It would need to represent the same abstract idea of trust in the context of H2H communication. What is this abstract idea?

The answer to these questions must be useful for a wide range of communication systems, such as H2H, IM2M and others that need to interoperate.

The only answer that turns out to be viable is that trust in communication systems must have nothing to do with feelings, emotions or other psychological and multiple-variable concepts. Trust is to be understood as something potentially communicable.

Further, trust must bridge different instances and observers, otherwise communication would be isolated in domains with islands of user interoperability that could not be bridged over time. This means that different subjective, objective and intersubjective (see Glossary) realizations of trust must depend on some common, basic and abstract expression. This expression is simply:

"Trust is that which is essential to a communication channel but cannot be transferred from a source to a destination using that channel."

Additional reason to use this expression

as the definition of trust in communication systems and why it is useful for Internet security protocols is given in the sidebar, quoted from the book *Digital Certificates: Applied Internet Security* by Jalal Feghhi, Jalil Feghhi and Peter Williams.

It is important to note that there are also "poetic" or "everyday" uses of the word trust that permeate some security work and Internet communication protocols. This may explain trust's "bad name" as a difficult concept and as an overloaded terminology. The problem is not however in the concept of trust by itself, but in using trust quantifiers that are either artificial or limited in the context of communication systems. A common limitation is to use trust as a synonym for authorization. However, this is valid only in a network, not in an internet. For example, in a network a trusted user is a user authorized by the network management to access some resources. But where is the "Internet management" in the Internet? It does not exist.

Let's use the definition of trust just given and move toward an understanding of it by using simple examples and problems that can later on be translated to more complex H2H and IM2M communication.

Examples and Problems

To make progress in understanding all this, we probably need to begin with simplified (oversimplified?) models and ignore the critics' tirade that the real world is more complex. The real world is always more complex, which has the advantage that we shan't run out of work. [quoted from Ball, J. "Memes as Replicators," *Ethology and Sociobiology*, vol. 5, p.159, 1984]

Let us suppose that a lion sees a lamb and tells the lamb "I'm not hungry." What should the lamb do? Usually, the lamb would run away – but not necessarily. There would be no danger if the lamb were able to know with a high level of reliance, acceptable to the lamb, that the lion is not hungry. Can the lamb

Trust Points

Commencing with a quotation from [SHAN48], Egardo Gerck leads an Internet discussion http://www.mcg.org.br/trustdef.txt with the assertion, used with permission, that:

"In Information Theory, information has nothing to do with knowledge or meaning. In the context of Information Theory, information is simply that which is transferred form a source to a destination, using a communication channel. If, before transmission, the information is available at the destination, then the transfer is zero. Information received by a party is that what the party does not expect—as measured by the uncertainty of the party as to what the message will be.

Shannon's contribution here goes far beyond the definition (and derived mathematical consequences) that "information is what you do not expect." His zeroth-contribution (so to say, in my counting) was to actually recognize that unless he would arrive at a real-world model of information as used in the electronic world, no logically useful information model could be set forth!

Now, in the Internet world, we have come to a stand off: either we develop a real-world model of trust or we cannot continue to deal with limited and fault-ridden trust models, as the Internet expands from a parochial to a planetary network for e-commerce, EDI, communication, etc.

And, what would be this "real-world model of trust" for the Internet world? Here, akin to Information Theory, trust has nothing to do with friendship, acquaintances, employee-employer relationships, loyalty, clearance, betrayal and other hard to define concepts.

In the concept of Generalized Certification Theory (see http://www.mcg.org.br/cie.htm), trust is simply "that which is essential to a communication channel but which cannot be transferred from a source to a destination using that channel."

Dr. Gerck's underlying observation that the integrity of certificate-based security systems (such as the X.509 Authentication Framework) hinges upon the very notion of trust is very valuable. Although trust is defined by the X.509 standard in terms of integrity—a CA is expected to reliably perform its user authentication and certificate registration duties—this does little to establish any conceptual properties of trust itself as a basis for building secure systems.

As anticipated by the ISO process, local environments are expected to profile and tailor the X.509 Authentication Framework. In this way, they can address the integrity requirements of national, application, community, or personal needs. ISO data communications standards are generally constructed assuming that islands of user interoperability will form, and the economic or social benefits of networking will inevitably cause systems to link together over time. X.509 does not impose a particular economic or social model of integrity, however. Such telecommunications standards generally limit their scope to stating technical matters. Models of integrity and trust in a particular space are best left to the communities of interest, governments, and industry forums, which are most familiar with these groups'specific needs.

Excerpted from *Digital Certificates: Applied Internet Security* by Jalil Feghhi, Jalal Feghhi and Peter Williams, Addison-Wesley, ISBN 0-20-130980-7, p. 194-195, 1998. Copyright work.

trust that the lion is not hungry?

Let us follow some idealized steps of this communication protocol.

The lion does not need to receive any information from the lamb besides that which is communicated in the communication channel itself – the lamb is there. The lamb obviously needs to know whether the lion's assertion "I'm not hungry" is truthful. The truthfulness of this assertion is not information and cannot be transferred using that same channel.

Why not? The truthfulness of the lion's assertion cannot be information because in Information Theory information has nothing to do with knowledge or meaning, and it cannot be transferred using that channel because how could the lamb know that the lion was not lying?

The same situation would apply to two machines on the Net, or to humans communicating. Information is surprise, and not always nice. Therefore, we see that "trust me" is an empty affirmation; a self-affirmation cannot communicate trust.

In other words, a decision to trust someone, the source of a communication, the name on a certificate, or a record must be based on factors outside the assertion of trustworthiness that the entity makes for itself.

Loosely speaking, we can say that "information is what you do not expect" and "trust is what you know." In the lamb example, the lamb needs to trust whether the lion is hungry. This could not have been information, because information is what the lamb does not expect – information is surprise. To be sure, the lamb does not want surprises in regard to the lion's appetite.

Linking both concepts of information and trust, we can say that "trust is qualified reliance on received information." This definition, derived from the first definition, is just one out of dozens of other trust definitions that can be derived, all coherent with the first definition.

This example also shows the interplay between trust and power. A very large difference in power, of one agent over another, implies that the more powerful agent can offset and control the other agent to such a degree that the other agent's actions are immaterial, even if the actions are already occurring – hence, a vastly powerful agent does not need to trust the least powerful agent. On the other hand, the least powerful agent needs trust in the other agent's behavior, since it cannot offset or control the other agent's actions to any degree – it needs to know with high reliance what the other agent's actions can be and, in some cases, what they cannot be, before they happen.

Let's look at some problems. How can I trust whether the message received by a party is the message that I sent, if I cannot control both sides of the communication channel? This question has gained in importance lately. Over time, we are finding that everything we see on our screens just might be false, including e-mail that says it was mailed by friends, or even digitally signed by them. And we are never totally sure that the website pages we are looking at are really from where they say they are from, or that what they say was not tampered with. Isolated networks do not help. Why? Because we are also finding that authorized users moving data from exterior to interior networks can compromise purportedly secure networks even if the networks are fully isolated.

These problems have no solution today. And it is well known that digital certificates and cryptography cannot help.

With this in mind, let me ask the reader how you could trust that the above text originated from myself? Of course, "It is printed in the *COOK Report* and Gordon Cook trusts it to be yours!" would be the most probable answer. But should you not be concerned about how Gordon Cook might trust that the above text is my own?

This question is not just rhetorical. Ignoring questions like this is at the base of flaws in many Internet protocols, even recent ones. For example, the recently developed SAML (security assertion mark up language) protocol for expressing authorizations in access control and payment systems has no answer to this question. This results in lack of support for audit of assertion dependency between co-operating authorities. For example, suppose that Bob authenticates to the Widget Marketplace using assertion A and receives Assertion B from the Widget Marketplace, whereupon Bob purchases machinery from a parts provider hosted at the Widget Marketplace. The parts provider authorizes the transaction based on Assertion B. If there is a problem with Bob's purchases at the Widget Marketplace (Bob will not pay his bills) there is nothing in the SAML flow that ties Assertion B to Assertion A. In other words, even though Assertion B has been issued by the Widget Marketplace in response to assertion A, there is no way to represent this information within SAML.

Ignoring fundamental trust issues is not the only problem. One must also emphasize the extraordinary flexibility of the requirements that are served by the use of common names and related identifiers in the Internet. There is a danger here for the creators of protocols that they will (accidentally or deliberately) misuse ordinary, trusted words and familiar concepts in ways that have been artificially restricted by special definitions to fit them for their purpose in the context of formal structures. The result is that they will mislead users as to the actual capabilities of the protocols.

One example is the debate about the word non-repudiation in X.509 and PKIX digital certificates. Even though the only possible meaning of the word is to be something that cannot be repudiated, since X.509 and PKIX cannot provide that meaning but people nonetheless like the "business" sound of it, the word is now oftentimes used to mean merely evidence of authentication.

Misusing trust is also related to the scenario of a spoofing attack. In spoofing, a user trusts a fraudulent service or infor-

mation that pretends to be legitimate. Spoofing attacks cannot be prevented by using SSL, digital signatures or encryption. The U.S. National Science Foundation recently sponsored a conference dealing with Internet security issues for voting applications, where spoofing was declared an open and very serious problem without a solution today and long-term research on the subject was recommended.

The problems mentioned above are diverse and touch upon many different aspects of the reliability of Internet communications. We need to solve these problems. A large section of our economy and our lives are already riding on the Internet.

A common question is whether these problems could not be solved by more control. Trust is good but control may be better. However, what to control and where? Unless every user is watched 24x365, or a filtering program is massively used denying functionality to users (as Earthlink does, denying SMTP and NEWS connections to their users, who must then only use Earthlink's SMTP servers notwithstanding the privacy, security and delay considerations), users are pretty much free to do whatever they wish at their connection – including using a different port for a route-around SMTP connection. Savvy users can evade controls. Thus, control does not appear to be effective in an internet. The essential point is that the Internet is a network of networks that has no central control point to be controlled. Not even by means of the DNS.

Here, a different approach imposes itself. Since it is illogical to break communications in order to ensure reliable communications, we ask: can reliable communications depend on trust? And, if so, what is controlled, and how?

Trust vs. Control

In the discussion of trust versus control, it is instructive to view trust as an open-loop control process in control theory terminology, i.e., a control process that does not rely on a closed feedback loop in order to achieve its purposes. This comparison allows one to recall the advantages and disadvantages of open-loop control vis-a-vis closed-loop control and apply them respectively to trust and close surveillance (also called control).

In control theory, the basic parameter to measure performance is position-error, which translates here to the actual response as compared to its expected or estimated (i.e., trusted) response. In open loop-control, one method frequently used to decrease position-error is to introduce periodic checks of any convenient system variable, not necessarily the control variable. This is not a feedback loop because it is done after the actuation. This method is equivalent to the well-known dictum "trust but verify" – implying the need for a pre-defined policy of checks and balances that can periodically adjust the trust estimator as a function of observed behavior.

Thus, trust can also be explicitly defined as "trust is an open-loop control process of an entity's response on matters of X" or, less precisely but more concisely, as "trust is to rely upon actions at a distance."

Interesting qualities resulting from this approach to trust in communication systems vs. close surveillance can be exemplified by the just mentioned control theory analogy regarding the main advantages of open-loop control over closed-loop control. These advantages include: simpler systems (hence, better fault-tolerance); immediate response (i.e., nothing needs to be measured in order for it to actuate); easier design (e.g., avoiding probable but unknown pitfalls of complex designs); easier interfacing (i.e., suffers less influence from and also exerts less influence on the rest of the system); modular design (i.e., complete and interchangeable); and less cost.

Where Is The Center? Where Are The Edges?

Using trust tools to solve the Internet problems we see today is, thus, not only a natural answer in terms of enhancing the very IM2M communication that is failing, but also possibly easier, cheaper, quicker, simpler, more secure and more successful than trying to take control of the Internet.

There is an additional problem that flows from a strategy of control. Strengthening centralized control would make that single handle of control a single point of failure. Strong centralized control also becomes the one basket for all eggs, which everyone wants to possess. These arguments were presented by the author in the April 2000 issue of *The COOK Report* in terms of domain name issues and are also valid here.

An attempt to bring centralized control to the Internet would also need to deal with that vexing question –what to control and where? In the Internet, no single person can tell which networks are included, because no one is there to tell which networks to include and which not to. Any user can add any number of networks to the system. Centralized control is impossible in an internet made out of open-ended networks of networks.

But if the answer does not lie in centralized control, what is the answer to the problems mentioned above? Could Internet control be decentralized? How would this be effective and not generate even more confusion? How would this provide for planetary reliable interpretation of protocols and their messages by machines?

First, one must dispel the notion of "center" and "edges" existing in the Internet. Of course, a network has "edges" and "centers" but the Internet is not a network, it is a network of networks. Surely, one often hears about "edge control" and "center control" and "trust on the edges of the Net" but let's ask – where is the center of the Universe? Where are the edges? Some structures have no edges and no center. The Internet, we must realize, is one such structure. Talking about "edges of the Net" is like talking about the "last Web page of the Net" – where is it?

The questions above are not even de-

fined. In the Internet, clients and servers are connected with peer-to-peer internetworking capacity, which capacity may or may not be used by an entity. In other words, all clients and servers are able to route IPpackets because of their internetworking capacity. So any client or server can become a center of its own network, as well as an end of IP protocol connections. And there is no edge because the Net does not end there, past that edge. At every point, internetworking can be extended to a neighborhood of points all of which lie inside the Net. Further, there is no privileged center that might be called "the center." Thus, there is no meaning that can be assigned to the phrase "the edges of the Net" in the same way that there is no "center of the Net" either.

This also shows that the argument for "trust at the edges of the Net" is a fallacy. An edge that is able to route IP packets becomes a center. And there is no privileged trust location at an edge in the same way that there is no privileged trust location at a center.

Indeed, as we observe the real world and seek to model the trust mechanisms that allow business and human communication to function, what do we see? We do not see a hierarchical trust structure controlling business from a single center.We do not see "edge control" either, and where it was tried it resulted in anarchy. What we see are inter-entity (intersubjective, see Glossary) relationships heavily qualified in many ways. We see internets, networks of networks, a manifold of networks with multiple control boundaries and lacking a common single reference.

Likewise in the Internet, we see a set of edge-edge, center-center (yes, there are many centers), center-edge and edge-center relationships that we can use to induce (communicate) trust. Trust is always formed from relationships between entities because it is induced by, or results from, a communication event. Such relationships may be unilateral (e.g., I do not know you and I do not know that you trust me) but not singular. There can be no information transfer between a sender and itself, since all is known. The net transfer is zero (see sidebar "Trust Points"). The only possibility for an entity to transmit information to itself is to do so to the future, but then if and only if there has been some loss of memory.

Note that so far we are not yet talking about how we use multiple channels to induce trust, as first we need to establish the need to use multiple channels before explaining how to use them.

Second, we need to direct our attention to the last sentence of the sidebar "Trust Points," quoted from the book "Digital Certificates: Applied Internet Security" by Jalal. Feghhi, Jalil. Feghhi and Peter Williams, where the authors state:

"Models of integrity and trust in a particular space are best left to the communities of interest, governments, and industry forums, which are most familiar with these groups'specific needs."

This sentence reminds us that trust is always local to and is earned in communities of interest. This dispels the idea of a cookie-cutter approach to Internet control, since each community of interest (for example, your company) will have different goals, different control objectives. Increasingly, users want more freedom in controlling their own connections, bandwidth, services and rules of use. Maybe that's why NAT/IPv4 is broadly used and IPv6 is not widespread – even though IPv6 is better in many aspects.

In summary, the answer needed to solve the fundamental problem of Internet communications is trust. Not trust as blind faith, compliance, belief, or ignorance, but trust as qualified reliance on information through open-loop control. Trust is that which provides meaning to information. Trust is something essentially communicable between machines and humans, something that can flow in our existing TCP/IP, dial-up and other networks.

But there are rules to this communication. Self-assertions cannot induce trust. Client-server communication is not enough to induce trust. We must move to a network model where not two, but four entities need to be in communication in order for trust to be induced.

Why Four?

Trust, as qualified reliance on information, needs multiple, independent channels to be communicated. If we have two entities (e.g., a client and server) talking to one another, we have only one channel of communication. Clearly, we need more than two entities. It seems unreasonable to require a hundred entities.

Editor's Note: This discussion will continue - most likely in the next *COOK Report*.

Glossary

Intersubjective – also called inter-entity; pertaining to more than one entity. For example, making a medical diagnosis is intersubjective because physicians of a same class (i.e., equivalent as observers) diagnosing patients with the same illness (i.e., equivalent as observables for medical purposes) may arrive at different results. The results depend on a patient-physician interaction. A medical diagnosis is thus not objective (i.e., the diagnosis is not the same for equivalent patients and equivalent physicians) and also not subjective (i.e., the diagnosis does not depend only on the physician). The same happens in other cases, most notably in risk assessment.