Rethinking PKI: What's Trust Got to Do with It?

Stephen Kent

Chief Scientist – Information Security, BBN Technologies, USA

Much of the literature related to public key infrastructure (PKI) uses terms such as “trust” extensively and assumes that certification authorities (CAs) are trusted third parties (TTPs). It is certainly true that the best known CAs today are commercial TTPs, and such CAs have played an important role in making the general public aware of PKIs. But, not all PKIs need adopt this sort of CA model, in which relying parties are required to make value judgments about the trustworthiness of the organizations that operate CAs. PKIs are not intrinsically valuable. They are infrastructures that, if successful, facilitate authentication and authorization services based on the use of public key cryptography. Thus it is appropriate to ask questions about these services:

- In what context are these services being employed?
- What forms of identifiers are meaningful for the context?
- Does the context relate to existing physical world, or does it exist only in cyberspace?
- Are the services offered to anyone, or are they intended for identifiable user populations?
- Are their existing organizational entities that are authoritative for the authentication or authorization information contained in the certificates issued by the CAs?

In many of the situations in which PKIs are being used today, or which have been proposed, the contexts have been transplanted from the physical world to cyberspace. Often, the users (subjects) of these PKIs already have been assigned identifiers in the physical world, identifiers managed by organizational entities that are considered authoritative for assigning these identifiers to these users. In such contexts it would seem natural for these organizational entities to act as CAs, identifying users in cyberspace in the same fashion as they have identified them in the physical world.

An implicit form of trust relationship exists here, between relying parties and these organizations, but this trust is often based on long established business or social relationships, contracts, or by statute. Trust in this context is not created out of thin air in cyberspace. It is a very different form of trust from what is usually described in papers on “trust management.”

This presentation argues that, in the best circumstances, CAs should not have to be trusted explicitly. Rather, CAs of the sort noted above merit an implied trust due to their position as authoritative entities responsible for name spaces, authorization information, etc. The presentation describes why this style of PKI has numerous advantages relative to the models commonly described

in the literature, and promoted by commercial TTPs. It examines examples of authoritative CAs in many aspects of everyday life, and analyzes the nature of the relationships that give rise to these CAs. It describes how standard X.509 constructs can be used to facilitate cross- certification among organizations in the context this model, how several forms of cyberspace identities could be certified consistent with this model, and how organizations can issue certificates easily to users with whom they have existing, client relationships.

A variety of photographs (shot by the speaker) will punctuate the presentation.