Chapter 12
PKI Trust Models

Audun Jøsang
University of Oslo, Norway

ABSTRACT

A PKI can be described as a set of technologies, procedures, and policies for propagating trust from where it initially exists to where it is needed for authentication in online environments. How the trust propagation takes place under a specific PKI depends on the PKI's syntactic trust structure, which is commonly known as a trust model. However, trust is primarily a semantic concept that cannot be expressed in syntactic terms alone. In order to define meaningful trust models for PKIs it is also necessary to consider the semantic assumptions and human cognition of trust relationships, as explicitly or implicitly expressed by certification policies, legal contractual agreements between participants in a PKI, and by how identity information is displayed and represented. Of the many different PKI trust models proposed in the literature, some have been implemented and are currently used in practical settings, from small personal networks to large-scale private and public networks such as the Internet. This chapter takes a closer look at the most prominent and widely used PKI trust models, and discusses related semantic issues.

INTRODUCTION

Trust is a directional relationship between two parties that can be called the relying party and the trusted party. One must assume the relying party to be a 'reasoning entity' in some form (Jøsang, 1996), meaning that it has the ability to make evaluations and decisions about trust based on received information and past experience. The trusted party can be anything from a person, organization or physical entity, to abstract notions such as information or a cryptographic key.

A trust relationship has a scope, meaning that it applies to a specific purpose or domain of action, such as "to be authentic" for a cryptographic key, or "to provide quality service and repair" for car

DOI: 10.4018/978-1-4666-4030-6.ch012

mechanics (Jøsang et al., 2005). The literature uses the term trust with a variety of meanings (McKnight and Chervany, 1996), so it is not always clear what authors mean by it. In order to avoid misunderstanding it is always useful to be specific and define the meaning of trust when using the term in a particular context.

A distinction should be made between interpreting trust as an evaluation or as a decision. When interpreting trust as a subjective evaluation of the reliability or correctness of something or somebody, it is called evaluation trust. When interpreting trust as a decision to enter into a situation of dependence on something or somebody, it is called decision trust. This distinction can appear subtle but is in fact quite fundamental. For example, having high evaluation trust in an entity is not necessarily sufficient to make a decision to enter into a situation of dependence on that entity if the risk is perceived as being too high. Evaluation trust reflects the reliability of the trusted party and is application and context independent, whereas decision trust depends on the particular application and on the context in which it is embedded. It can be shown that decision trust is a function of evaluations trust and risk (Jøsang and Lo Presti, 2004).

Both evaluation trust and decision trust reflect a positive belief about something on which the relying party potentially or actually depends for his welfare. Evaluation trust is most naturally measured as a discrete or continuous degree of reliability or belief, whereas decision trust is most naturally measured in terms of a binary decision. Several authors have proposed to let certificates express levels of trust on a discrete or continuous scale, e.g. (Kohlas et al., 2008). However, this would only be meaningful in case CAs are uncertain about the correctness of what they certify, and expressing levels of trust in the certificate seems to be incompatible with CA business models. It would be rather strange if a CA states in a certificate that the certified public key e.g. is authentic with probability 0.9, as no user would want to buy such certificates. Certificates are issued according to a certification policy. In practice this policy is often published as two separate documents called the Certificate Policy and the Certificate Practice Statement where the former specifies high level requirements and the latter how these requirements are fulfilled in detail. We will here refer to both with the term "certification policy." The relying party can judge the adequacy of the policy for the intended certificate usage. The relying party must also consider whether the certification policy is properly adhered to by the CA. Evaluation trust in a validated certificate can be defined as "the quality of the certification policy combined with the belief in the CAs adherence to that policy." However, relying parties often do not have the expertise to judge the certification policy, and it would be practically difficult for relying parties to audit the CA's adherence to the certification policy.

A validated certificate never provides 100% assurance that the public key actually is authentic. It could for example be possible for an attacker to trick the CA to issue a public-key certificate with the wrong name, thereby enabling the attacker to spoof the corresponding identity (Microsoft, 2001). It is up to the relying party how the assurance provided by a particular public-key certificate is to be used and depended upon in a real situation. For example, a relying party may take a validated public-key certificate as evidence of authenticity but still only be 90% convinced that the certified public key is authentic, which would be equivalent to 90% evaluation trust. The same relying party can nevertheless decide to accept and use the public-key certificate despite not being totally convinced about the correctness of the identity, and this would be a case of binary decision trust. Decision trust in a validated certificate can be defined as "the acceptance of the certificate based on evaluation trust and other contextual factors."

A certificate only partially provides the trust needed for a particular transaction. Relying parties should interpret a validated public-key certificate as evidence of its authenticity, but not as evidence