Modeling and Evaluation of Certification Path Discovery in the Emerging Global PKI

Meiyuan Zhao[1] and Sean W. Smith[2]

[1] Communications Technology Lab
Intel Corporation
Hillsboro, OR 97124
meiyuan.zhao@intel.com
[2] Department of Computer Science
Dartmouth College
Hanover, NH 03755
sws@cs.dartmouth.edu

Abstract. Establishing trust on certificates across multiple domains requires an efficient certification path discovery algorithm. Previously, small exmaples are used to analyze the performance of certification path discovery. In this work, we propose and implement a simulation framework and a probability search tree model for systematic performance evaluation. Built from measurement data collected from current PKI systems in development and deployment over more than 10 countries, our model is (to the best of our knowledge) the largest simulated PKI architecture to-date.

1 Introduction

Public key infrastructure (PKI) is a powerful tool for protecting information. Current development and deployment of PKI systems shows a trend toward an emerging global PKI, where individual PKI domains by governments, institutions, and enterprise establish trust relationships via cross-certification technology. However, as a PKI becomes more complicated, so does the work required for validating an individual certificate. The first step is *certification path discovery*: constructing a "chain of certificates" that connects the certificate in question to a trust anchor. It is challenging to locate appropriate resources to establish a candidate path and to maximize its chance of being valid.

The global PKI spans many countries and consists of many domains, CAs, repositories, and users. PKI protocols need to be robust in such a complex network environment. By establishing trust relationships between domains, cross-certification confronts us with a complex "certificate topology". Moreover, users in different PKI domains may display completely different behaviors that may impact the effectiveness of PKI protocols.

Previous analyses of certification path discovery focused mostly on using small examples to understand algorithm options. In this study, we evaluate its performance in the context of the emerging global PKI. The power of *simulation* allows us to model such complex certificate topologies and to simulate realistic situations. It also enables us to explore a wide range of algorithm options and different network environments, and to examine the effect of user activities as well. We make the following contributions:

A.S. Atzeni and A. Lioy (Eds.): EuroPKI 2006, LNCS 4043, pp. 16–30, 2006.

- We design and implement a PKI simulation framework for general-purpose PKI performance study. This framework implements classical X.509 PKI services and is flexible to allow new types of models and performance studies.
- We design and implement a *PathBuilder* module for this framework. This module uses novel probability search tree models to simulate a variety of algorithm behaviors for certification path discovery.
- We model a global PKI architecture using measurement data collected from current PKI system deployment over more than 10 countries. To the best of our knowledge, this is the largest simulated PKI architecture to-date.
- Using these tools, we evaluate performance of certification path discovery using a range algorithm options. We show that the performance is sensitive to algorithm options, PKI architectures, and user activities.

We hope to make our tools publicly available, as open source.

In the rest of this paper, Sect, 2 discusses the background of PKI system and certification path discovery. Sect, 3 presents previous research. Sect, 4 discusses our simulation framework for general purpose PKI systems. Sect, 5 discusses details of our work on modeling certification path discovery and performance analysis. Finally, we conclude this work with discussions in Sect, 6 and 7.

2 PKI and Certification Path Discovery

PKI was first proposed [14] for securely distributing public keys. It has now evolved to architectures providing comprehensive services for public key *certificates*; these services include storing and retrieving certificates, maintaining and updating certificate status, and validating certificates. In a traditional X.509 [10] PKI system, the certificate storage service is provided by a repository that supports protocols for users to store and retrieve directory information; the protocol used most commonly here is the *Lightweight Directory Access Protocol (LDAP)* [23]. The *certificate status information (CSI)* service communicates the validity status of certificates. A certificate is typically considered as "valid", "revoked", or "unknown". Classical approaches to CSI includes periodically updated data structures such as a *certificate revocation list (CRL)* [10], and online protocols such as *online certificate status protocol (OCSP)* [17].

2.1 Certification Path Discovery

The user who tries to validate a certificate is referred to as *relying party*. A certificate validation service handles *certification paths*, sequences of certificates representing a trust path to the certificate of interest. In such a sequence, the issuer of the first certificate is called a *trust anchor*; a trust anchor is an entity the relying party trusts by default. The last certificate in the sequence is called the *target*; the target certificate is the one that the relying party is trying to validate. In a path, consecutive certificates are linked together by having the *subject* of the previous certificate match the *issuer* of the next certificate.

A certificate validation service is composed of two stages: certification path *discovery* and certification path *validation.* The latter stage is well-established. RFC3280