AutoPKI: A PKI Resources Discovery System*

Massimiliano Pala and Sean W. Smith

Dartmouth College, Computer Science Department,
6211 Sudikoff, Hanover, NH 03755, US
{pala,sws}@cs.dartmouth.edu
http://www.cs.dartmouth.edu

Abstract. The central goal of *Public Key Infrastructure (PKI)* is to enable trust judgments between distributed users. Although *certificates* play a central role in making such judgments, a PKI's users need more than just knowledge of certificates. Minimally, a relying party must able to locate critical parameters such the certificate repositories and certificate validation servers relevant to the trust path under consideration. Users in other scenarios may require other resources and services.

Surprisingly, locating these resources and services remains a largely unsolved problem in real-world X.509 PKI deployment. In this paper, we present the design and prototype of a new and flexible solution for automatic discovery of the services and data repositories are available from a *Certificate Service Provider* (CSP). This contribution will take real-world PKI one step closer to achieving its goal.

Keywords: PKI, Service Discovery, Certification Authority, Digital Certificates.

1 Introduction

The central goal of *Public Key Infrastructure (PKI)* is to enable trust judgments between distributed users. At is core, PKI depends on certificates: signed bindings of public keys to keyholder properties. Effective use of PKI requires use of these certificates; however, effective use of certificates requires many additional services, such as OCSP servers, CRL repositories, timestamping services, etc. As a consequence, client-side PKI tools need to be able to discover and use these services; server-side PKI tools need to be able to provide these services and enable client tools to discover them.

Unfortunately configuring these tools to carry out these tasks is painful for both server administrators and end users, thanks to badly written User Interfaces

* The authors would like to thank Stephen Kent, Frank Pooth, Ashad Noor, Sravan and all the PKIX WG for several discussions and comments. This work was supported in part by the NSF (under grant CNS-0448499), the U.S. Department of Homeland Security (under Grant Award Number 2006-CS-001-000001), and Sun. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of any of the sponsors.

J. Lopez, P. Samarati, and J.L. Ferrer (Eds.): EuroPKI 2007, LNCS 4582, pp. 154–169, 2007.

(UI) and overly detailed configurations. Certification Authorities barely publish access details on their official websites; even data as basic as the URLs for provided services and repositories are usually omitted. As a result, if a CA provides a new service (e.g. OCSP [1]) or a new data repository (e.g. LDAP [2]), users and administrators have difficulty learning of these changes. Furthermore, certificates already issued could not carry any sign of the new services. It is unlikely that users (and applications) will be easily aware of the new services if not directly contacted. This problem impacts even more on users from enterprises other than the issuing organization, as they have very limited knowledge about CA's practices and service locations.

In this paper, we present a new approach to provide a flexible way to automatically discover which services and data repositories are available from a CA. This flexibility would also facilitate interoperability across different infrastructures. Section 2 presents the core aspect of our solution: the design and the implementation of a new (and simple) *PKI Resource Query Protocol* (PRQP) easing PKI management both for administrators and final users. Section 3 presents our prototypes. Section 4 evaluates the performance of our prototypes and the effectiveness of our solutions. Section 5 reviews other approaches to solving the problem. Section 6 concludes with some directions for future work.

2 The PKI Resource Query Protocol

To solve this problem, we define the *PKI Resource Query Protocol* (PRQP) for finding any available PKI resource from a particular CA. In PQRP, the client and a *Resource Query Authority* (RQA) exchange a single round of messages:

1. the client requests a resource token by sending a request to the server;
2. the server replies back by sending a response to the requesting entity.

The client embeds zero or more resource identifiers (OIDs)—when specifying exactly the data the client is interested into—in the request token, in order to specify which subset of CA resources she wants. If the client does not specify any services by providing an empty list of OIDs in the request, all of the available data for a particular CA should be returned by the server in the response. The resources might be items that are (occasionally) embedded in certificates today—such as URLs for CRLs or OCSP or SCVP—as well as items such as addresses of the CA homepage address, the subscription service, or the revocation request.

Fig. 1 shows an example of this protocol: an SSL web server needs to retrieve the revocation status of a user's certificate. (Here, the Web server is the PQRP requesting client.) At first (step 1), the web server receives the user's certificate from the browser. The web server looks at the issuer identifier in the certificate and builds up a PRQP request asking the RQA for the location of the OCSP server of the issuing CA (step 2). The RQA provides (step 3) the web server with the URL of the requested service, as configured on the RQA. In this particular example only the OCSP URL is requested, and therefore only the locator for such service is put in the response. The web server, then, continues with the