

PKI Technology Survey and Blueprint

Peter Gutmann
University of Auckland

Abstract

This paper presents and examines the results of a series of interviews in which a cross-section of experienced programmers, system administrators, and technical project managers with many years of practical, real-world experience were asked which technologies they would use to solve some of the major problems which occur in PKI implementation. The results of the interviews and various significant issues identified by them are presented and discussed. Finally, a PKI technology blueprint based on recommendations made by respondents is presented. The resulting design is noteworthy in that it is almost completely unlike the one proposed in X.509 and related standards, which would indicate that at least some of the deployment difficulties being encountered with X.509-style PKIs are due to their sub-optimal choice of technology.

1. Introduction

PKI users have for many years been trying to build PKIs using the methods and mechanisms described in the X.509 series of standards, and more recently profiled in an ongoing series of IETF RFCs. During this time the feasibility and practicality of these technologies have been called into question as the result of negative implementation experiences [1][2][3][4]. In an attempt to determine which technologies would be the most suitable for implementing a real-world PKI, this paper looks at the following question:

If you asked experienced programmers, sysadmins, and technical project managers how they would implement certificate management, what would the technology framework for your PKI look like?

The intent of this work is to take a cross-section of technical computer users with many years of practical, real-world experience, and see which technologies they would select if asked to implement a PKI, or more specifically how they would solve the major problems which occur in PKI implementation. To this end the paper asks a series of “How” questions (such as “How would you store certificates”) rather than “What” questions (such as “What policy would you employ for private-key handling?”). The intent is solely to determine the most practical means of solving common PKI-related technology problems without trying to address policy and legal issues, which are best left to upper management, lawyers, and lawmakers. Some of the technology issues which need to be addressed are relatively obvious and sufficiently well-known that they have their own names, examples being the “Which directory?” problem and the “Which John Smith?” problem.

The remainder of this paper is organised as follows. Section 2 covers the method used to interview respondents and the questions they were asked. Section 3 presents the raw results obtained from the interviews, and section 4 analyses various common issues which showed up in the information provided by respondents, as well as presenting additional (unsolicited) feedback provided on areas which respondents saw as being a source of significant implementation or deployment difficulties. Finally, section 5 presents a PKI implementation blueprint based on the recommendations made by respondents.

2. Survey Method

A series of interviews was conducted in which a number of programmers, system administrators, and technical project managers were asked various questions about their choice of technology for implementing certificate management. In order to avoid problems with the type of self-selecting survey often posted to Usenet newsgroups and web pages, the respondents were explicitly selected from among the employees of various companies and organisations rather than by soliciting responses from volunteers.

The respondents were instructed to choose the technology they would use if they were responsible for implementing, deploying, and supporting the PKI. The intent of this was to emphasise the fact that they were being asked to design a practical, real-world system rather than one based on bleeding-edge or experimental technology. In addition they were told that there were no right or wrong answers to questions, and in particular that “I don’t know” was a perfectly valid answer since it indicated that the question corresponded to a particularly tough problem. Despite the fact that most of the respondents were male primates, several of them did admit to not knowing the answers to various questions.

2.1 Bias Removal

The author has for some time been exposed to user feedback on preferred PKI implementation technology as a side-effect of maintaining an open-source crypto toolkit which contains fairly extensive PKI functionality. This led to concerns that the questions might (inadvertently) be phrased in a manner which influenced respondents into replying in a manner which matched the author's existing experience with users.

In order to ensure that the results weren't biased because of the way the questions were phrased, they were sent to representatives of every major and some minor PKI technologies for comment. The intent of this solomonic bias-removal process was to ensure that none of the technologies could later claim that they had been unfairly excluded from consideration because of the way a particular question was phrased. For example instead of asking about revocation checking (which would bias the results towards a CRL-style solution), the question was phrased in terms of freshness/validity checking, which allowed for a variety of answers, including CRLs. Feedback from the PKI representatives was applied to the initial questions (although almost no changes were deemed necessary), leading to the final questions given in the next section.

2.2 Questions

Before being asked the questions themselves, the respondents were asked about any existing exposure to PKI which might bias the response. The actual questions which followed were broken down into five groups covering enrolment, identification of certificates, storing and obtaining certificates, checking certificate validity, and a miscellaneous section. Although these questions don't represent an exhaustive enumeration of all possible PKI technology issues (it's unlikely that any finite question list can), they cover all of the major areas and, in their answers, provide a good solid technology framework on which to build a practical PKI.

The identification, storing and obtaining certificates, and validity checking question groups were presented in that order because the results from one group tended to affect the following ones. For example a choice of domain names as a certificate identifier would lead logically to the use of a DNS-based certificate distribution mechanism in the following question group, and a DNS-based validity checking mechanism in the group which followed that.

Questions within each group were ordered logically and in some cases anticipated answers to earlier questions. For example experience with users indicated that email addresses were popularly used to identify certificate owners, so one of the questions which followed the initial identification question was how someone (or something) which didn't have an email address would be identified. Finally, two questions about the cost and/or complexity of the solutions given in previous questions were added in order to discourage impractical and extravagant schemes.

Enrolment

1. How do people sign up?
2. Can this be bypassed/made less labour-intensive?

Identification

2. How would you identify certificate owners?
3. What if there's more than one John Smith?
4. What if they don't have an email address?
5. How would you get an ID which is globally unique?

Obtaining certificates

6. How would you store a collection of certificates?
7. How would you access a collection of certificates?
8. How would you locate a collection of certificates?

Freshness/validity checking

9. How would you check the validity/freshness of a certificate?
10. How would you handle the cost of doing this?

Miscellaneous

11. How would you check that an operation was valid at a given time in the past?

Respondents who were particularly quick to leap in with an answer were asked to justify their choice, mostly by being reminded that they would be responsible for implementing and supporting their choice of technology. As mentioned earlier, the intent behind this was to weed out technology they had only heard of through trade magazines in favour of technology they were familiar with and believed was practical to deploy and maintain in the field.

3. Results

When it came to exposure to existing technology, it proved almost impossible to find anyone who hadn't been exposed to PGP and (to a lesser extent) ssh. In an attempt to locate someone who hadn't been subject to these potential sources of bias, the net was cast wider and wider, eventually landing non-technical managers who, unfortunately, weren't able to provide answers to most of the technical questions. However, since both PGP and ssh have little in the way of infrastructure, the results were probably not affected by preconceived notions of how a PKI was supposed to be implemented.

Only one respondent had had any significant exposure to X.509, and his responses to the questions differed markedly from the other responses (details are given further on). None of the respondents were aware of other PKI systems such as SPKI/SDSI or AADS.

The remainder of this section presents the responses to the questions provided by users. The section which follows this one contains an analysis of the results.

3.1 Enrolment

This group of questions was seen by most as being policy rather than technical questions. As a result, the non-technical managers were able to answer them while most of the remainder saw it as being an issue which was best handled by others. This position wasn't taken because they were trying to avoid work or responsibility, but because they were used to such decisions being made by management, the clients, or other external agents. In other words, they were unfairly being asked "What" rather than "How" questions.

The responses which were provided tended to be domain-specific. One respondent who worked for an organisation with a large number of clients suggested a paper-mail-based enrolment system in which the (known to the organisation) clients were enrolled using traditional paper documents, with all certificate details being filled out from the organisation's records. Another respondent with a healthcare background suggested creating certificates based on existing health records. The responses tended to leverage existing mechanisms for use in the enrolment process as much as possible, both for ease of deployment and because they represented established business practice and were therefore likely to be looked on favourably if a dispute over enrolment details arose. The main design goals of these schemes appeared to be a combination of ease-of-use and risk avoidance.

3.2 Identification

Almost all users immediately suggested the use of an email address as the primary identifier for a certificate. One or two users suggested domain-specific identifiers, for example in healthcare the patient ID or medical registration number might be used as the identifier, and an organisation with known clients would use the clients' account number.

In cases where the certificate owner didn't have an email address, various (obvious) solutions such as the DNS name or IP address were suggested (an alternative way of phrasing question 4 was "What if the certificate owner is a printer?" when the respondent had immediately suggested using an email address as the answer to question 3). Other responses included MAC and IPv6 addresses (the latter because they included more information than IPv4 ones). One respondent provided a nice generalisation to "the name you saw when you first encountered the device", so that if the printer mentioned earlier appeared on the local network as "Wallet Buster 300" (the name used in one department for a photographic-quality printer with a particularly high per-page printing charge) then it would also be identified in the certificate as the "Wallet Buster 300". Although names such as these were meaningful only in the local context, the fact that the identified item was only visible locally made this issue irrelevant.

Some users had problems coming up with an identifier. One user suggested using a personal name or company name and asking the user to select a certificate if several matching ones were found, but wasn't able to provide a solution which would be amenable to automated processing. It's probable that they misunderstood the nature of the question, however the author was reluctant to provide further prompting for fear of influencing the results.

Most users immediately suggested the use of a GUID (Globally Unique ID) as a unique value to identify certificates. Two users weren't aware of GUIDs but described (in some detail) an identifier built up in much the same way as a GUID.

3.3 Obtaining Certificates

As with the email addresses, almost all users immediately suggested using a database as the certificate storage mechanism, seasoned to taste (“Anything but Oracle”, “ODBC, because it's on every Windows machine”, “Whatever the company's using at the moment”, “Oracle, DBAs are a dime a dozen”, and so on). One respondent suggested LDAP “because that's what you store certificates with” but was unable to provide further information, and had no actual LDAP experience. Another respondent (the one with X.509 experience) also suggested using LDAP “because that's what you use”, but immediately followed it up with the comment that it wouldn't work in his organisation because users typically occupied multiple roles (leading to a multitude of possible entries in the directory, which could change several times a year), and the only way they had found to resolve the problem was to use the directory as a flat database. Another respondent initially suggested “Whatever I can find on SourceForge”, but eventually settled on a database like most of the others because of the availability of open-source solutions such as Berkeley DB and MySQL.

The unanimous consensus for the access mechanism was HTTP (“That's reality”). One user commented that they'd really prefer XML and SOAP if it were a bit more widespread, and another user suggested ODBC as another possibility, while acknowledging that there would be some problems due to it being a mostly Windows-only solution and having some problems with Internet traversal. Several users expanded their basic answer to address reliability and scalability issues (“We've had zero downtime for our web pages in the last year (except for link outages) even though individual servers have occasionally gone down”). They provided sketches for web architectures to handle almost any eventuality, based on their existing experience with web technology.

Most respondents suggested fetching the certificate from what can be generalised to “the most obvious place”. For example if the certificate belonged to someone at a given organisation, they would query the organisation's web server. If they needed a certificate for someone at their own organisation, they would query their main corporate file or web server. If they had an email address, they would query the corresponding web server, for example `www.hotmail.com` for a Hotmail email address.

Some of the respondents who worked for organisations with known clients indicated that their (custom) client software would be configured when it was built or deployed so that it would talk to their own servers. One respondent also suggested using the DNS as a certificate storage mechanism, but quickly decided that it wouldn't work for much of the standard reasons that DNS-based certificate storage has been regarded as impractical. Two users suggested the possibility of a certificate search engine which worked like existing web crawlers and indexers, extracting certificate information and providing a single portal from which multiple disparate certificate stores could be accessed.

Finally, several respondents commented that if all else failed, users would have to manually set preferred server URLs in the same way that they set preferred home pages in web browsers.

3.4 Freshness/Validity Checking

As with the previous responses, the almost unanimous response to the question of validity checking was to use the certificate store in the manner of a trusted directory. Freshness and validity checking would be performed through a simple fetch operation, with the result being either a known-good certificate or some form of error indication. The user with X.509 experience suggested using CRLs, but immediately followed it with the observation that they didn't really work, and something better would have to be found.

Since this question required a bit more information than the basic “Use HTTP”-type response to question 7, users provided a fair bit of detail on the operations involved. For example one respondent suggested communicating a checksum (meaning a cryptographic hash) rather than the full certificate to save bandwidth, and several used the GUID (the unique ID from question 5) to fetch the certificate. Other respondents added use-by dates to certificates to indicate the interval after which it should be re-fetched from the certificate store, or suggested the use of HTTP-type cache control mechanisms which served a similar purpose.

As with questions 1 and 2, most of the respondents regarded question 10 as being a policy issue and therefore someone else's problem (“That's beancounter material”). The main motivation for adding this question, as with question 2, was to discourage excessively extravagant solutions in the answers to the previous question.

Those who did answer this question suggested a variety of approaches such as a multi-tiered charging structure similar to the pricing schemes used by ISPs and for web hosting in which different levels of service and usage were billed in different ways. Other suggestions were to use per-query charging, to specifically charge the relying party rather than the certificate owner, or (in recognition of the fact that charging for queries would discourage use) the use of cost-sharing schemes to avoid one party carrying the cost while another party obtained all the benefits.

3.5 Miscellaneous

Again as with earlier responses, the unanimous response to the historical-query question was that the certificate store should maintain audit logs of certificate histories and use those to resolve historical queries. Since auditing was built into most databases and the certificate store was the ideal place to maintain this information, the consensus was that this was a job for whatever or whoever was managing the certificate store.

As with question 9, some respondents provided a fair amount of detail on the operations involved. For example one user suggested charging for the length of storage of historical information, and another user came up with the novel idea of storing historical information for previous certificates in the current certificate, so that anyone obtaining an end entity's current certificate (via the mechanism from question 9) could also use it to answer historical queries. While this is in theory impractical due to bandwidth considerations, in practice having a 5Kb vs. 1Kb certificate would make no difference with a PC-based application, affecting only highly constrained devices such as smart cards.

3.6 Summary of Results

A summary of the results of the survey, which contrasts the approach provided by the X.509 standards (which represent the most commonly-used PKI blueprint) with the approach suggested by programmers, system administrators, and technical project managers, is shown in Table 1.

	X.509	Survey response
Identifier	DN	email address/DNS name/IP address
Unique ID	DN	GUID
Storage	X.500 directory	Database
Access	LDAP	HTTP
Validity check	CRL	Repository presence check
Historical query	(Timestamping)	Authority records

Table 1: Survey results summary

4. Discussion

This section analyses the responses from users given in the preceding section. The major trends which were apparent in the responses are presented in their own sub-sections, with miscellaneous comments gathered at the end.

4.1 Consistency of Results

The most remarkable thing about the results presented in the previous section is the fact that almost all of the respondents agreed on one particular solution to the problem presented by each question. So consistent were the answers (somewhat akin to finding a straight line on a double-log graph) that the author felt it necessary to locate and question further respondents, leading to the eventual extension of the survey to non-technical managers as described in section 3.

The fact that the respondents had been specifically instructed to select the technology they felt was the most practical and feasible probably helped produce this consistent result. A few of the respondents were later informally asked what they would have suggested if they had been allowed to choose any technology (no matter how impractical), and came up with very different answers such as CORBA (although this was suggested as a joke by someone whose employer had a customer with a particular obsession with CORBA) or LDAP "provided that *the name of their main competitor* has to support it afterwards".

Almost all users suggested using a GUID or GUID-equivalent as a unique identifier for a certificate. This came as something of a surprise to the author, since the conventional approach (at least in PGP, SPKI/SDSI, and recently-issued X.509v3 certificates) is to use a value derived from the certificate's public key. When asked why they hadn't used the public key, the users responded that they hadn't considered it, but that that would also work. It appears that the widespread acceptance and use of GUIDs as general-purpose unique identifiers lead to this being the immediate choice for unique certificate identifiers as well.

4.2 Universality of WWW Technology

The penetration of the web into all aspects of computer use was very obvious in the responses. All respondents regarded HTTP as the universal glue to tie the PKI together. The use of web technology went far beyond the basic transport mechanism. Users suggested the use of HTTP cache-control mechanisms to handle certificate re-validation, web search engine technology to make locating certificates when their exact location wasn't known easy, and the use of various standard reliability and scalability-enhancing techniques such as round-robin DNS to address availability concerns.

The ability of the respondents to design sophisticated web-based solutions without too much effort reflects the extensive practical expertise available in this area, backed up by a large number of tools (both commercial and open source, to suit all tastes) and background technical information. For example the level of scalability planning extended beyond the basic bullet-point-on-a-PowerPoint-slide level to "we'll use these servers and this software because we've done it before and we know it works", a good indication that the resulting design would be practical under real-world conditions.

4.3 Key Management Issues

Several users expressed concern about the complexity involved in the key and certificate setup process. One user proposed a certificate-vending-machine type mechanism for which the only user interface task consisted of entering some form of authenticator and clicking a button labelled "Click here to generate a key and obtain a certificate". This was to be implemented using an HTTPS interface to the CA, submitting the public key and reading the resulting certificate back from the certificate store. Another user suggested "look at how Netscape does it and then do the exact opposite", a reference to the complexity of the browser-based enrolment process used to obtain Verisign certificates.

Yet another user, from a healthcare background, commented that many of their users would require per-site (rather than per-user) keys, since doctors expected many of the operations requiring the keys to be performed by nurses or administrative assistants, and keys were expected to be associated with roles such as "Duty doctor" (covering several GPs and assistants) rather than a particular individual. The inevitable result of this inability of per-user certification to match existing practice was that "they'll take whatever doctor turns up first in the morning's key and use that for the rest of the day", an observation arising from many years of experience with equivalent (non-public-key-based) solutions.

Other users also expressed concern about the enrolment/setup process. One user, working for a large organisation with known users, commented that a one-click enrolment process ("assuming Amazon hasn't patented that too") would be an absolute requirement, with an automated phone call-back being used to confirm that the user had indeed required the certificate ("cumbersome but functional").

This is clearly an area which needs further study to determine how low-impact the enrolment process can be made while still satisfying various legal concerns. Without any rigorous (and workable) framework for this area, users are coming up with solutions such as the alarming practice of having the CA generate the end user's private key and then sending it to them via email, either in plaintext form or with the password attached [5].

4.4 Miscellaneous

The fact that some respondents worked in a particular area influenced their replies to policy (rather than purely technical) questions. For example people working for organisations with clients or members tended to think of end entities in that role, with the organisation managing certificate issuance by taking advantage of its existing knowledge of users.

Several respondents spontaneously evolved SPKI/SDSI-type concepts such as local and global naming and timed re-validation of certificates, even though they had no previous exposure to PKI design. This mirrors experience with

psychological studies of complete non-programmers who spontaneously evolved programming-language-like constructs such as control statements when they were asked to create descriptions of algorithm-like tasks [6].

The innate tendency of sysadmins and technical managers to build in disaster-planning has already been pointed out elsewhere [7]. This was also apparent in many of the architectures laid out by respondents, with multiple development paths being possible in order to arrive at the final goal (one user summed it up with “Postgres if possible, Perl, Apache, and MySQL if they need it by Monday”). This was further reinforced by the fact that a number of respondents planned in future extensibility to handle scalability and reliability issues. The initial condition that users would be required to lie in the bed they had made for themselves appears to have been a powerful influence in both the choice of technology and the overall architecture.

Another fact which became apparent from the replies to the questions (although it’s not directly relevant to this paper) was that the respondents’ job position often matched their ability to provide answers to the questions. For example respondents who were working as senior programmers often had difficulty in architecting solutions to some of the more complex problems like identification or billing, while respondents with a similar amount of work experience who had been migrated into technical project management had little difficulty in this area. Although this has little effect on the results presented here (the respondents were chosen from a general cross-section of technical users without concentrating on one particular area), it is interesting to note that people seemed to have drifted into the job role they were most suited to, at least as determined by their ability to answer the survey questions.

5. PKI Implementation Blueprint

Using the results presented in the preceding sections, we can now look at how a PKI might be implemented with a particular goal of using the most practical real-world technology in order to increase the chances of successful deployment. As was already mentioned earlier, this implementation blueprint covers only the “How” aspect and leaves issues such as policy and legal concerns to the appropriate entities.

The basic certificate-management system is built on top of the database of choice, and uses an HTTP (or HTTPS) interface for communication. Certificates are generally identified by user name (CommonName in X.500 terminology) and email address, with alternatives such as an account number, IP address, or device name being used where this isn’t feasible.

Certificate issue is handled via a minimal one-click interface, which can be accomplished on most systems in a reasonably automated manner by reading the user name and email address from the user account information (for example the GCOS field under Unix or the Windows user information), and using it to populate the certificate request. The generated certificate is obtained by fetching it from the certificate store.

The process of obtaining a certificate is also the mechanism used for freshness/validity checking, with the certificate store returning only known-good certificates. Historical queries and similar issues are handled through the standard auditing and accounting mechanisms built into the database, which are used to track certificate additions and deletions and similar operations.

The basic mechanisms presented here can (obviously) be garnished to taste. For example some CAs may require a private-key proof-of-possession operation before issuing a certificate, which may require a two-stage process to be used when requesting a certificate. Potential implementers should however bear in mind that the goal of this work was to determine how to build a practical PKI. A workable (but not quite theoretically perfect) practical PKI is still better than theoretically perfect vapourware.

6. Conclusion

This paper has presented the results obtained from asking a number of technically skilled users with extensive practical experience how they would implement a certificate-management system. The resulting design is noteworthy in that it is almost completely unlike the one proposed in X.509 and related standards. This would indicate that at least some of the deployment difficulties being encountered with X.509-style PKIs are due to the sub-optimal choice of technology. To address this problem, the paper proposes a new certificate management technology blueprint based on information in the responses from users. This blueprint makes use of widely-utilised, mature technology and the extensive experience which users have working with it.

7. Acknowledgements

The author would like to thank various PKI theology representatives for feedback on the questions asked and Paul de Bazin, Tony Bryant, Tom Bowden, Nick Brooker, Russell Fulton, Paul Kendall, Suad Musovich, Edwin Ng, Steven Robb, Raymond Sellars, Chris Stephens, Russell Street, and Clifford Wilson for putting up with the questioning.

8. References

- [1] “Advances and Remaining Challenges to Adoption of Public Key Infrastructure Technology”, United States General Accounting Office report GAO-01-277, February 2001.
- [2] “Solution and Problems: (Why) It’s a long Way to Interoperability”, Jürgen Schwemmer, *Datenschutz und Datensicherheit*, **No.9, 2001** (September 2001).
- [3] “Prime-Time Player?”, Leo Pluswich and Darren Hartman, *Information Security Magazine*, March 2001.
- [4] “PKI: An Insider View”, Ben Rothke, *Information Security Magazine*, October 2001.
- [5] “Lessons Learned in Implementing and Deploying Crypto Software”, Peter Gutmann, to appear.
- [6] “What non-programmers know about programming: Natural language procedure specification”, Kathleen Galotti and William Ganong III, *International Journal of Man-Machine Studies*, **Vol.22, No.1** (January 1985), p.1.
- [7] “Apropos: Re-Routed Packets”, Tina Darmohray, ;*login*, **Vol.26, No.8** (December 2001), p.3.