beTRUSTed™
An e-security business of
PRICEWATERHOUSECOOPERS

Demystifying international cross-recognition of PKI

We've been barking up the wrong tree

Information Security Solutions Europe Conference, London, 2001

Stephen Wilson
Director Policy & Strategy, PricewaterhouseCoopers beTRUSTed

Abstract

Cross-certification and cross-recognition continue to be stumbling blocks in PKI. Cross-certification has been a lofty goal for many years but has proven to be expensive and impractical. And when we look at it closely, we find that it wouldn't give users much benefit in any event. Cross-certification establishes the equivalence of certificates from different PKIs, yet two users on either end of a transaction are usually asserting different types of credentials which will never be equivalent. The fundamental issue for users is not equivalence; it is fitness for purpose.

We're accustomed to the role of independent audit reports helping us to decide if a CA can be relied upon, but the decision is traditionally made out-of-band. This paper will present a new way of making a CA's audit report machine-readable, as a standard X.509 certificate. The approach is based on existing international audit standards and mature accreditation systems. It thereby demystifies PKI, clarifies liability, cuts compliance costs, and preserves sovereignty in communities of interest and national schemes.

Introduction

This paper presents a new interoperable PKI model where each CA is certified by a higher level CA on the basis of its audited compliance with agreed policies and practices. The model leverages the current strong trend in most PKIs for independent audit, but makes the audit reports machine readable, as standard X.509 certificates, so that Relying Party applications can process them automatically.

Furthermore, the new model utilises existing auditor accreditation bodies, and thereby paves the way for robust national and international Root CAs, while avoiding the controversy and complexity historically associated with such roles.

Making CA audits machine readable

A regular X.509 certificate can be made to represent a complete and precise summary of a CA's audit. Such a certificate would be issued by or on behalf of the auditor. Now note that the Policy OID of an end user certificate customarily indicates the certification policies and practices (together, the CP/CPS) under which that certificate was issued. These are the subject of the CA's audit. In this new PKI model, the Policy OID of the CA's certificate indicates the provisions and standards of the audit methodology. In turn, the auditor may carry their own X.509 certificate, conveying their accreditation under conventional audit control standards (such as the ISO/IEC Guides 62 and 65).

In effect, auditor accreditation bodies can act as Root CAs in this PKI model. The significance of such a Root CA is that any unbroken chain of certificates beneath it can be assumed by a Relying Party to have the following meaning:

- the end user's certificate signifies that the user has satisfied the enrolment requirements laid down in the CP/CPS of the CA
- the CA's certificate signifies that the CA passed its most recent audit, and that it was found to be in compliance with all its practice disclosures as well as all relevant standards
- the auditor's certificate signifies that the auditor's accreditation is current.

Because the certificate chain is readable by any standard X.509 parser, any Relying Party software therefore has the in-built capability to determine the fitness-for-purpose of certificates issued under this PKI.

The new audit-based PKI model promises to save costs by leveraging existing bodies, audit standards and methodologies. Importantly, many auditor accreditation standards operate independently of the technical domain of the auditors themselves, and so may be applied at the top of the PKI with little or no modification. At the same time, by stressing conventional assurance and risk management principles, the model makes PKI more comprehensible and accessible to business users. The model may provide the most logical and most practical way towards national and international Root CAs, under the auspices of existing international accreditation associations and mutual recognition arrangements for same.

Overcoming the problems of cross certification

The model incorporates *cross-recognition* of end user CAs (as opposed to cross-certification) in that independent audits are used to signify that each CA is conforming to its appropriate practices and that its certificates are therefore fit for purpose. It does not matter that different CAs in this system might be representing distinct communities of interest, issuing certificates for entirely different purposes.

In contrast, cross-certification is explicitly concerned with the detailed direct mapping of CAs' respective Certificate Policies and CPSs, in order to establish the equivalence of the certificates they issue. It is widely recognised that CA-to-CA cross-certification cannot scale up to practical numbers of CAs, and there are few if any examples of full cross-certification having been achieved from scratch by independent CAs. Moreover, equivalence is often entirely moot. If for example a doctor and a health insurance claims agent are transacting, there is no question that their respective credentials and authorisations are quite distinct. By the same token, their digital certificates must be expected to be non-equivalent.

The autonomy of communities of interest

Increasingly, dedicated certification authorities are being established by (or on behalf of) distinct communities of interests. For enhanced trust, the membership or business rules of the community are built into the certificate registration process. A unique Policy OID can be assigned to such community-specific certificates, and Relying Party applications need only check for the appropriate OID in an incoming certificate in order to verify the capacity of the sender to act in transaction.

Examples of communities that are making use of their own specific certificates include healthcare professionals, the law and the judiciary, accountants, stockbrokers, and registered company officers.

For e-business to thrive within these groups, it is imperative that they be allowed to preserve their membership rules and their identities. In PKI terms, the implication is that their registration rules and Certificate Policies be self-determined. This in turn elevates the importance of independent audit, so that Relying Parties outside the communities have appropriate assurance of their proper conduct in registration and certification.

The trend towards independent audit of CAs

There is a strong trend towards independent audit of CAs, driven by the desire for CAs to have autonomy over their own rules, and by the need for conventional risk management (in particular, liability insurance is impossible to arrange without independent audit taking place). Some jurisdictions (such as Italy) impose explicit requirements on CAs for certification or audit, and some vertical market segments do the same (see for example, Identrus in banking and finance).

Currently popular audit methodologies include SAS 70, ISO 17799, and the WebTrust for CAs (which is rapidly gaining popularity after being endorsed by Microsoft and by Identrus).

The need for audit is naturally strongest in cases where:

- the transaction value or business risk is high
- relying parties are at arms length from the certificate issuer
- certificate management is not a core business function, or
- certificate management has been outsourced.

PKI as a chain of digital audit certificates

To make the most of an audit result, the CA should make it available online to relying parties. CA Trust does this by way of a seal on the CA's website, but this requires out-of-band examination by the Relying Party, at least on occasion. It is far preferable for the Relying Party application to be able to recognise the audit status automatically.

A conventional X.509 certificate, issued by (or on behalf of) the auditor may be used to assert the result of a successful CA audit. The audit certificate signs the public key of the end user CA and so is capable of being parsed by conventional X.509 software. Thus, a valid, current certificate chain extending from an end user back to a recognised auditor can be interpreted to mean that the user certificate is fit for purpose, and that it has been issued by a CA that was, at the time of the last audit, in compliance with its own policies and procedures as well as any other prescribed standards.

It is noteworthy that the Policy OID of each certificate in the chain conveys distinct meanings. For end user certificates, the OID points to the conditions under which the certificate was issued and therefore provides an indication of the intended purpose of the certificate. For CA certificates in this audit based PKI model, the OID indicates the terms and conditions of the audit. The Policy of an end user CA in this model clearly does not map onto the Policy of an audit, nor is it a sub-set as is the case in some more traditional PKI implementations.

Who audits the auditors? Re-inventing the role of Root CA

There are general purpose international standards that govern the quality and proper conduct of audits over a wide range of fields. Such standards include ISO/IEC Guide 25 for laboratories and test facilities, Guide 62 for quality and environmental management systems, and Guide 65 for product certification. Audit firms may be accredited against those standards by independent national accreditation bodies.

Examples include the Comite Francais D'Accreditation (COFRAC), the German Accreditation Council (DAR), the Dutch Accreditation Council (RvA), the United Kingdom Accreditation Service (UKAS), the US ANSI Registrar Accreditation Board National Accreditation Program (ANSI-RAB NAP) and the Joint Accreditation System of Australia and New Zealand (JAS-ANZ).

Crucially, existing accreditation standards and accreditation bodies can be applied to govern CA audits, with little or no modification. And there is an existing international pool of qualified information security auditors with the skills needed to evaluate CAs.

In effect, national accreditation bodies can act as Root CAs in the audit based PKI. Their role in PKI is the same as their role in traditional certification schemes – that is, they audit the auditors – and so their potential liability is well understood and tends to be well limited. Because of their maturity and long established authority, accreditation bodies solve the problem of infinite regress in PKI (that is, how far back do you go before you find a CA you can trust?). And most importantly, because there are existing protocols and agreements for national accreditation bodies to recognise one another, this approach to PKI provides the most natural and robust means for cross-recognition in PKI.

Automatically verifying fitness for purpose

To summarise, the diagram below illustrates the meaning of each certificate issued within the audit based PKI.

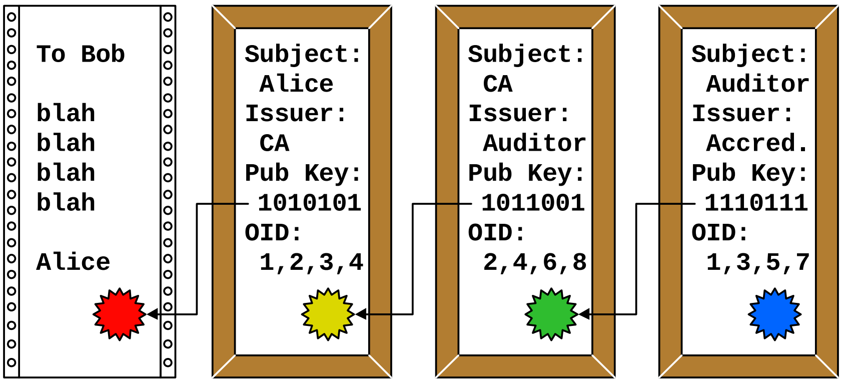

"Alice meets CA's rules (ref. OID 1,2,3,4)"

"The CA complies with a fit-for-purpose CP/CPS"

"The auditor is competent and impartial"

For Bob to be able to automatically process Alice's signature and certificate, he needs to be equipped ahead of time with just two pieces of information: (1) the expected Policy OID appropriate for the transaction at hand, and (2) the accreditation body's root public key. High value e-business usually involves special purpose applications and the appropriate OID will be readily configured in the application.

Advantages of the audit based model

It has a light touch

There need be little or no government involvement in running the audit based PKI scheme. It leverages existing accreditation bodies, an established contestable marketplace of information systems auditors, and existing accreditation standards. These existing structures support the ready creation of brand new accreditation schemes, including this PKI model, so long as complete technical standards are available for auditors to reference.

The scheme requires no special legislation, and so is compatible with both technology neutral and 'two tier' legislative regimes. Yet even in technology neutral regimes the model will still confer legal benefits by introducing transparent, independent assurance of the compliance of a given certificate with published practice statements, policies and standards.

It is opt-in and builds bottom-up

The model starts with the assumption that even in the absence of regulatory mandate, market forces will drive the audit of CAs. Auditors may be expected to compete on the basis of service level, industry specificity, reputation, price and so on. Depending on the value and risk of the transactions, and on the openness or other preferences of their community, CAs might start out without external audit, then bring in auditors as their certificate population grows or as their market demands it. Auditors might not necessarily be accredited but again market pressures will apply.

Of course, some communities or jurisdictions may mandate audit as well as particular standards. The model can accommodate different audit standards, which may be asserted in the auditor CAs' Policy OIDs.

It clarifies liability

Liability in PKI, especially for the higher level CAs has hitherto seemed an almost intractable problem. But the audit based model, even without legislation, will clarify liability in most jurisdictions.

Liability is no mystery in any regular standards accreditation scheme; risk can even be insured away under errors & omissions policies (indeed, ISO/IEC Guide 65 *requires* auditors to carry insurance before they may be accredited). In practice, liability actually diminishes as you go up the chain. By way of comparison, it is exceedingly rare for quality management or product certification auditors to be sued, and there appears to be no precedent at all for legal action against an accreditation body.

Allows for fitness for purpose

The model caters for different CAs implementing their own business rules and autonomous registration practices, fit for the purpose of particular application purpose. The two levels of audit and accreditation allow for complete flexibility of Certificate Policies and CPS at the end user CA level.

It normalises the language of PKI

Regulators, legislators, lawyers and insurers – to name just some of the non-IT specialists involved in formulating PKI – can now better understand the roles of higher level CAs and the root CA, because the relationships can be seen as conventional ones of audit and accreditation. By normalising the language, we will better engage all interested parties, and improve the decisions they make.

It normalises 'trust'

Finally the audit based model helps to put the often problematic concept of 'trust' into its proper perspective by emphasising the fitness for purpose of a certificate. PKI should not be overloaded with broad aims of conferring 'trustworthiness' to certificate holders or CAs. Rather, a certificate should only be seen as demonstration that the holder has met the specific rules of the CA, so that relying parties can make informed decisions as to whether or not to accept the certificate in support of a given transaction.

A note about sovereignty

'Sovereignty' has been one of the stumbling blocks in development of transnational PKI but the audit based model demystifies the issue. We should no longer see the role of root CA as being to push policy from the top down, nor even to approve policy for CAs. Instead, user CAs can have autonomous control over their policies and sovereignty over their communities. The root CA certainly does not hold the "keys to the kingdom" as some have asserted.

Certificates issued within different communities of interest remain completely distinguishable from one another, by virtue of their Policy OIDs. There is no imputation of automatic equivalence for certificates issued within this PKI. Rather, it is up to relying party applications to check the Policy OID before accepting any certificate. Note that this check is essential practice regardless of the type of PKI in operation, since it is impossible to prevent anyone trying to use a certificate outside the domain in which it was issued.

The following diagram shows how different communities of audited CAs can co-exist under the one accreditation body. We can expect some auditors to specialise in different verticals, just as they do for example in quality certification. Therefore, CAs in the government domain might be covered by different auditors from those in say health. Yet the same accreditation body can audit the auditors in whichever domain, and with no intrusion into the respective communities' ways of doing business. The diagram also shows the case where an industry body (banking) is sufficiently authorative over a closed community that it it can forego accreditation altogether.

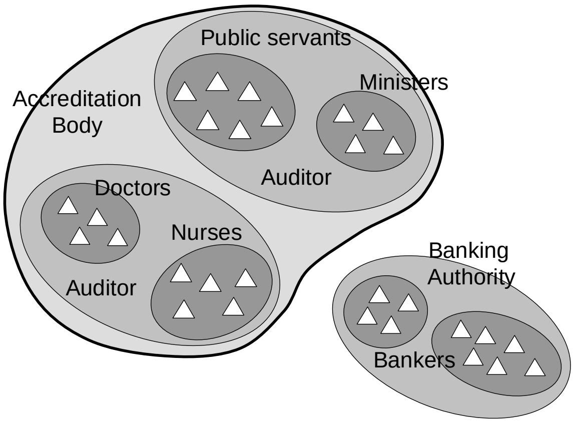

Communities vs. hierarchies

It is common for PKI to be deprecated purely on the basis of being hierarchical. Regardless of whether this is really reasonable or not, it is true that the common depiction of PKI as a tree carries some baggage. Tree charts carry inescapable authoritarian overtones. They confer some sort of position of superiority to the root and other high level CAs, and it is hard to shake off the impression that the root CA imposes rules on its 'subordinates'. Furthermore, the typical tree chart confuses the relationships between different levels of CA, by making all the links in the tree appear the same. Some commentators have tried drawing the tree upside down, to make the root seem less imposing, but this fails to convey the essential bottom-up growth of modern PKI, where connection to the root CA ought to be optional.

The 'visual language' of nested communities within communities in the diagram above seems less threatening than the stark tree charts, despite the fact that the depictions are topologically the same. Note too that the stand-alone banking community looks somehow less peculiar as an island than it would as a severed branch hanging off the side of a tree. The concept of certificates asserting membership (of groups satisfying certain rules) rather than absolute identity, is also conveyed by the diagram.

Frequently asked questions

Q. I find it hard to accept that a root CA could be set up and run by anyone other than a peak security establishment. Just what sort of security expertise do these typical accreditation bodies have?

A. The act of certification (of some entity's compliance) is separate from that of accreditation (of an auditor's competence and impartiality). Different standards apply to the conduct of accreditation compared to certification, and so the accreditation body does not need to be expert in the domain of the auditor. Accreditation bodies, as governed by standards like ISO/IEC Guide 65, have tried and proven processes for assembling advisory committees with the necessary domain expertise to conduct accreditation reviews.

For example, the Australian National Association of Testing Authorities (NATA) accredits independent test facilities for the Australasian Information Security Evaluation Programme (AISEP; see www.dsd.gov.au). NATA itself has no expertise in cryptographic security but it does have processes for assessing the competence of those who claim such expertise. These processes have to be generic so that accreditation bodies can 'boot strap' certification schemes for any new domain. So, technical security should not in fact be the prime concern of the root CA; it should be governance.

Q. A follow-up question. Great care is still going to be needed over the 'root key'. Does the typical accreditation body have the skills or resources?

A. The root key probably wouldn't be kept under the direct sole control of one organisation. Rather, it would be broken into components and held in separate hardware devices, stored in safety deposit boxes or the like. The root key components need only come together on the odd occasion that a new CA auditor is accredited, or an existing one renewed or revoked. The environment and systems for using the root key would of course be critical, but these functions could be outsourced to a high end CA operation.

Q. You say the model provides for fitness for purpose to be asserted in the certificate chain. But isn't it a conflict of interest for an independent auditor to make assertions about the appropriateness of a Certificate Policy?

A. Yes it would typically be beyond the scope of an audit for the auditor to make their own assessment of the fitness for purpose of the Certificate Policy. Nevertheless, the auditor can look for evidence that the CA has written (or otherwise adopted) the Policy with proper care and attention to the application domain. A parallel is the area of contract management under the ISO 9001 quality management standard. ISO 9001 auditors examine the contracts written between a manufacturer and its suppliers. The auditors do not directly judge the appropriateness of the contracts but they do seek documentary evidence of supplier consultation, contract review, dispute resolution and so on, as per the standard. In audit based PKI, we would expect similar processes for assuring the fitness for purpose of a Certificate Policy to be in place.

Q. I have never even heard of these accreditation bodies. How can they form the root of all trust in e-commerce?

A. One part of the answer is that despite their low profile, accreditation bodies are in fact ubiquitous in business today. Our dependence on the integrity of financial audits, product safety, occupational health and safety, environmental inspection, cryptographic systems, and more, all rest on systems of independent qualified auditors and accreditation bodies.

The other part of the answer is that maybe we shouldn't imagine 'trust' to be anchored at some all powerful location. Certainly, responsibility for Certificate Policy needs to be de-centralised, along with CAs' business rules, in order to preserve the autonomy of communities of interest. Trust as such actually needs to be created between CAs and users. The proper role of auditors and of PKI itself is to provide reliable assurance that correct procedures are being followed. The role of root CA should not be to push policy from the top down. And it certainly should not hold the "keys to the kingdom" as some have asserted.

Q. Who pays for all the overhead?

A. The certificate holder will usually pay, in the form of a premium price paid for certificates issued under the system. CAs and certificate holders are motivated to seek the services of accredited auditors, because it maximises acceptance of their certificates. The same economic rationale underpins all voluntary certification schemes, such as ISO 9001. For risk management, it is possible furthermore that insurance companies will only offer policies to CAs that are independently audited. The overall reduction in systemic risk and cost may lead large scale application hosts, such as tax departments and online healthcare operators, to underwrite some of the audit costs.