

White paper

Audit based public key infrastructure

DRAFT

Stephen Wilson

April 2000

Abstract

This paper outlines a new model for constructing open PKI where the relationships between CAs at different levels are based primarily on compliance audit. A strong trend across most PKIs at present is for CAs to seek independent audit, to assure users and relying parties that certificates are issued in accordance with agreed standards and procedures. Until now, the results of such audits have only been made available in the form of paper compliance reports.

To improve the ability for relying parties to decide online whether or not to accept a certificate, the results of a successful audit could be asserted in a digital certificate, issued to the CA by (or on behalf of) the auditor. Such digital audit certificates can conveniently be conventional X.509 'identity' certificates, signing the public key of the subject CA. This enables the audit certificate to be automatically processed by any ordinary X.509 parser, during the course of verifying the end user certificate. A regular X.509 certificate can represent a complete and precise summary of the audit. The Policy OID of this certificate would point to documentation of the conditions of the audit and any applicable standards. In turn, the Policy OID of the end user certificates issued by the CA would have to indicate the actual certification policies and procedures covered by the audit.

A complete PKI can be built upon the principles of contestable open audit and existing mechanisms for accrediting auditors. Approved auditors would likewise have their status confirmed by a digital certificate issued by a recognised accreditation body, authoritative over the audit standards in force.

The audit based PKI model is most applicable to "open-but-bounded" e-business environments. It 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 also provide the most logical and most practical pathway to transnational "root CAs", under the auspices of existing international accreditation associations and mutual recognition arrangements for same.

What's wrong with cross-certification as a model?

Most thinking about PKI interoperability has been explicitly dominated (or implicitly informed) by cross-certification: the mapping by CAs of one another's Certificate Policies in order to establish the equivalence of the certificates they issue. Cross-certification is widely acknowledged now to be impractical yet it seems to persist as some kind of ideal. While many policy makers are exploring "cross-recognition" as a more practical model, they often cling onto cross-certification for certain applications. For instance, the influential APEC Electronic Authentication Task Group recently stated that "Cross-recognition is generally not suitable for cross-border transactions that require a high level of trust as it is procedurally less rigorous¹ than cross-certification".²

But the fundamental problem with cross-certification is the underlying aim of establishing equivalence of certificates. Very often, such equivalence is not merely impractical, it is entirely moot to begin with. If for example a doctor and a government official were transacting, any equivalence of their credentials would never even be considered. The same goes for their respective digital certificates.

Often certificate equivalence shouldn't even be considered

The trend towards domain specific certificates

Increasingly, distinct Certification Authorities³ are being established by communities of interests or e-business domains. For enhanced trust, the membership rules of the community or the business rules associated with the domain are built into the certificate registration process. A unique Certificate Policy Object Identifier (OID) is assigned to such domain-specific certificates, and relying party applications need only check for the appropriate OID in an incoming certificate in order to verify the transaction.

Examples of e-business domains that might make use of their own specific certificates include:

- healthcare (doctors, medical specialists, nurses, pharmacists etc. all licensed by their respective professional bodies)

¹ Until detailed auditable evaluation procedures are laid down for cross-certification and cross-recognition, it is really impossible to describe either as more rigorous than the other. After all, it is easy to imagine high assurance certificate cross-recognition processes and low assurance cross-certification processes.

² See *Achieving PKI Interoperability* under the APEC Electronic Authentication Task Group home page <http://apii.or.kr/telwg/eaTG/eaTG-cont.html>.

³ In this paper, "Certification Authority" is used to describe the overall function of identifying and registering subscribers, and generating and distributing certificates, as defined by a particular Certificate Policy. Any Registration Authority and/or outsourced Certificate Manufacturing Authority involved in the process are rolled into the "Certification Authority". For the purposes of determining compliance with Policy and standards, it is necessary that the CA's boundaries include both the front end and the backend.

- the law and the judiciary
- accounting
- stockbroking (where a stock exchange certifies licensed brokers)
- government-to-business (for purposes of corporate taxation, securities regulation etc.)⁴
- government-to-citizen (for purposes of personal taxation, rates paying, social security services, voting etc.).

The move towards domain specific certificates has several drivers:

- high-value e-business (like online taxation and medical practice management) is conducted using dedicated applications or highly customised web forms, to support unique features and rich user functionality
- domain authorities are better trusted for specific applications than are general purpose CAs and can register certificate holders more economically
- certificate issuers minimise their risk by constraining the purposes for which their certificates may be used
- relying parties minimise their risk by using certificates from a recognisable issuer, preferably one which is authoritative over the business domain.

Domain specificity helps the privity problem

Note that the infamous problem of contractual privity between CAs and relying parties is not such an issue in this environment. Relying parties will tend not to be faced with incoming certificates from unexpected sources. Rather, CAs will typically be bound to specific domains and applications, bringing signatories and relying parties into close legal proximity, even in the absence of explicit contracts.

The trend towards independent audit of CAs

Another strong trend is towards independent assurance of the compliance of CAs with best practice benchmarks or standards. Some jurisdictions impose explicit requirements on CAs for certification or audit. In free markets, risk management makes it prudent for CAs to seek independent attestation of the quality of their operations, and for users to seek the services of CAs carrying such assurances. In time, it will become impossible or prohibitive for CAs to find insurance without showing successful independent audit under recognised standards. Currently popular audit methodologies include SAS 70 and the ISO 9000 series. A new WebTrust for CAs is under development by the American and Canadian accounting institutes.

⁴ A good current example is the Australian Business Number Digital Signature Certificate, to be issued by certified commercial CAs on behalf of government, and intended to support all transactions a business may have with government; see www.noie.gov.au/projects/govt/ABNDSC.htm.

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
- 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 its audit status available online to relying parties. WebTrust for CAs proposes doing this by way of a seal on the CA website, but this requires human intervention by the relying party, at least on occasion.

An improvement would be for the relying party application to be able to see the audit status automatically. This can be achieved if the compliance certificate resulting from a successful audit is issued in digital form instead of on paper.

It is proposed that a conventional X.509 ('identity') certificate, issued by (or on behalf of) the auditor be used to assert the result of a successful CA audit. The audit certificate would sign the public key of the subject CA and so would be capable of being parsed by conventional X.509 software library routines. The Policy OID of this certificate would point to documentation of the conditions of the audit and any applicable standards. In turn, the Policy OID of the end user certificates issued by the CA would have to indicate the actual certification policies and procedures covered by the audit.

Thus, a valid, current certificate chain extending from an end user back to a recognised auditor could 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.

Note that the Policy OID of each certificate in the chain would convey distinct meanings. For end user certificates, the OID would indicate the intended purpose of the certificate. For audit certificates, the OID would indicate 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. The Policies at different levels of the PKI are in fact orthogonal, in so far as a single reasonably generic audit Policy could cover a wide range of user CA Policies, since the audit methodology does not depend strongly on the operational details of the subject.

Accreditation (or Who audits the auditors?)

In any contestable audit market, there tends to be pressure for independent assurance of the proper conduct of the audits. In many cases, there are generic standards that govern the conduct of audits, and auditors may be evaluated against those standards in a process commonly known as *accreditation*. Such standards include ISO/IEC Guides 25 (for laboratories and test facilities), 62 (for quality and environmental management systems) and 65 (for product certification). The prime concerns of accreditation are the impartiality of auditors and their competence to work in a given domain.

Accreditation bodies, responsible for accrediting auditors and auditing bodies, are sometimes designated as ‘fourth parties’ to reinforce their independence from the third parties. Accreditation bodies can have a broad scope across a whole family of standards and disciplines, or they can confine themselves to particular verticals or domains. Examples include:

- Joint Accreditation System of Australia and New Zealand (JAS-ANZ; see www.jas-anz.com.au) has authority in those two countries for accrediting auditors for many different ISO standards and others
- US National Institute for Standards and Technology (NIST) covers such areas as the National Voluntary Laboratory Accreditation Program (NVLAP)
- US Registrar Accreditation Board governs quality certification bodies (known in the US as ‘registrars’) under the ISO 9000 series
- The Information Systems Audit and Control Association (ISACA; see www.isaca.org) ensures the proficiency of information systems auditors through its CISA programme⁵
- Various institutes of chartered accountants⁶ accredit or formally license financial auditors and, increasingly, security-related professionals, such as WebTrust auditors.

The diagram below illustrates the relationship between suppliers, auditors and accreditation bodies in a traditional accreditation/certification scheme (such as applies under the information systems security standard AS/NZS 4444). Also shown is how the same relationships can apply in an audit based PKI.

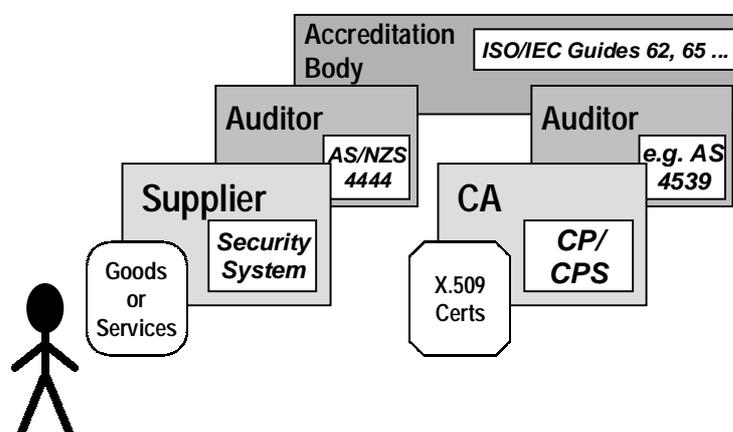
⁵ Because ISACA asserts the competence of auditors to perform independent audits, rather than actually performing any audits itself, the qualification *Certified* Information Systems Auditor (CISA) is another example of the inconsistent use of the terms.

⁶ Also known as ‘certified public accountants’ in some jurisdictions – yet another example of inconsistent use of ‘certify’.

The main reference for the PKI audit is the Certificate Policy and Certificate Practice Statement of the CA (just as the main reference for an IS security audit is the documented security systems procedures). There will be a number of relevant standards covering the operation of a CA; for brevity just one is shown in the diagram – the new Australian Standard AS 4539 (Public Key Authentication Framework).

Existing accreditation standards and bodies can govern PKI

Crucially, existing accreditation principles, standards (and indeed bodies) can be applied to govern CA audits, with little or no modification.⁷ And of course there is also an existing international pool of security auditors with the basic skills needed to evaluate CAs.



Accreditation body as root CA

The audit based PKI model would have an accreditation body acting as a type of root CA. The top level accreditation certificates would be signed by a key pair under the control of the accreditation body (the actual root CA system might be outsourced).

Rethinking the meaning of the root CA

The idea of a root CA carries some baggage and so it is important to carefully delineate the meaning of an accreditation body root CA. In the audit based PKI model, the root CA issues certificates exclusively to CA auditors who have been formally accredited against recognised standards of impartiality and competence. The significance of such certificates appearing at the end of a given certificate chain is simply that the end user CA has been evaluated by an *accredited* auditor.

It is not the case that every end user certificate issued under such an accreditation body root CA is automatically equivalent. As discussed

⁷ Initial discussions suggest that either ISO/IEC Guide 62 or Guide 65 is a good fit for managing PKI audit but a firm decision would need to be made through the due process of establishing a formal accreditation scheme (JAS-ANZ, personal communication, 1999).

above, it is increasingly common for relying parties to have to verify a range of different certificates intended for distinct purposes. Establishing equivalence is not an important objective; rather, relying parties need independent evidence of fitness for purpose, as provided by the audit certificate chain.

Of course, nothing stops an unaccredited auditor evaluating end user CAs and issuing digital audit certificates. Such certificates might chain back to some other type of root CA; alternatively the auditor CA might be self-signed. But market forces can be expected to favor accredited CAs, and it is reasonable for relying parties to discriminate between different types of root CAs – those that are associated with accreditation bodies, conferring special extra significance to the certificate chain, and those that are not.

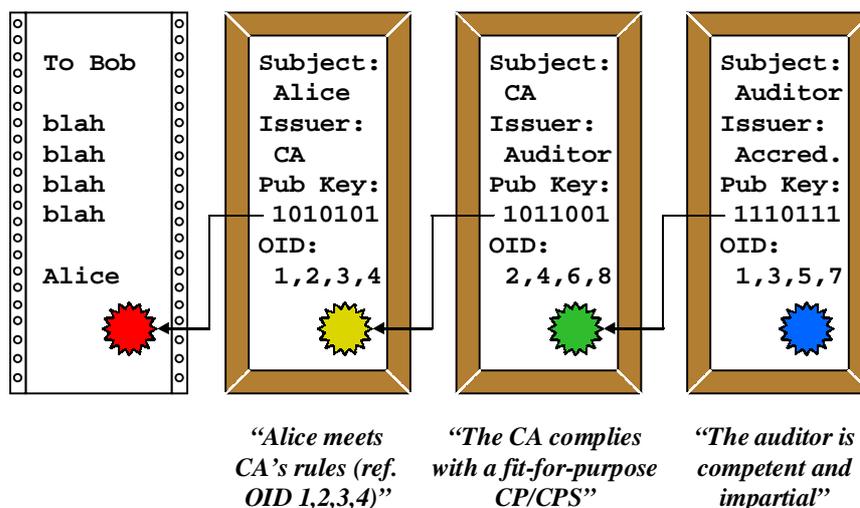
Resilience of the root CA

These new views of the root CA can lead to a more resilient PKI. In fact the root need not even be mission critical. The accreditation of auditors will by its nature be painstaking and gradual. It follows that certificates (and CRLs) will be issued infrequently by the root – and never suddenly. Additions to and subtractions from the set of accredited auditors will be significant events and in practice would be pre-empted by out-of-band announcements of non-conformities, providing the opportunity for extra security checks.⁸ In the event of root key compromise, CAs can continue operation by virtue of the intrinsic stability of the intermediate auditor CAs.

Automatically verifying fitness for purpose

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

⁸ It is probably safe to say that in case of a transaction of such high value that revocation of the auditor merits consideration, an out-of-band check of the auditor's current status (and indeed of other conditions) would be reasonable. Thus the mission criticality of the root CA is not much affected by the value of the transaction!



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,⁹ and (2) the accreditation body’s root public key.

Advantages of the audit based model

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.

And the scheme requires no special legislation. Yet it will still confer legal benefits by introducing transparent, independent assurance of the compliance of a given certificate with published practice statements, policies and standards.

Opt-in and 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

⁹ As discussed above, high value e-business usually involves special purpose applications and the appropriate OID will be configured in the application.

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.

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.¹⁰ 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.

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.

Normalises 'trust'

Finally the audit based model helps to put the 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

¹⁰ Indeed, ISO/IEC Guide 65 *requires* auditors to carry insurance before they may be accredited.

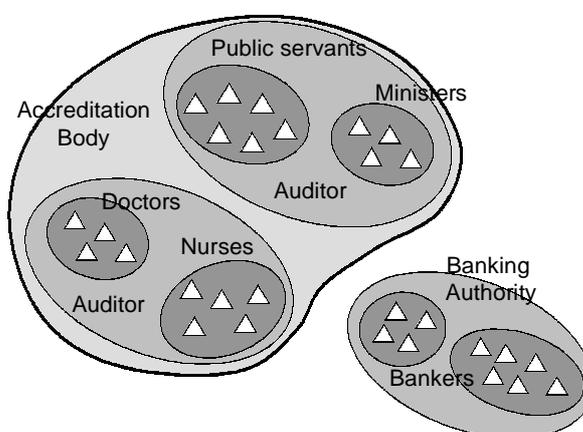
holder has met the specific rules of the CA, so that relying parties can make informed decisions as to whether or not to accept the certificate in support of a given transaction.

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 authoritative over a closed community that 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.¹¹ 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.

A picture is worth a thousand words!

The ‘visual language’ of 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

¹¹ Other common misconceptions arising from the tree chart include the notion that all CAs are ‘online’ at the same time, that they all participate in the issuance of each end user certificate, that all registration data is transmitted up the tree to the root CA, and that network connections to the root CA must have enormous bandwidth and availability.

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.