

PKI, The What, The Why, and The How

(GSEC V1.3 Practical Assignment)

© SANS Institute 2002, Author retains full rights.

Duncan_Wood_GSEC

Table of Contents

1	Summary	3
2	What is PKI?	3
2.1	The Basic Technology	3
2.2	Certificates.....	6
2.3	Attribute Certificates	7
3	Why are there constraints?	9
3.1	The common business drivers	9
3.2	Why is there legislation?.....	10
3.3	The Organisations Interests and Obligations.....	11
3.4	Protection of Personal Information	11
4	How to select an Architecture.....	12
4.1	What type of CA/RA facility is to be used	12
4.2	What information is to be included in the certificate versus providing it elsewhere.	13
4.3	Seed application versus the use of PKI within the Organisation.....	14
5	Conclusion	14
6	Glossary	15
7	References	18

PKI, The What, The Why, and The How

1 Summary

I want to title this paper as the 'PKI, The good, the bad and the ugly' even though it might attract the curious it was a little too off beat to qualify (even with my sense of humour) as a good title. The purpose of this paper is to: -

- Describe what Public Key Architecture (PKI) is, and how it can help organisations and individuals during the enactment of electronic transactions. (The GOOD);
- Detail why governments worldwide have, or are, introducing legislation and guidelines covering information privacy and the operation of PKI facilities. Due to the nationality of the author, Australian regulations are drawn upon throughout this document. (The BAD); and
- Finally, to show how with honest definition of business requirements, and careful architectural decisions, a way forward satisfying both the business requirements and the legalities. (The UGLY)

Note. If I am honest it is not that ugly, but then a title something like 'PKI, The good, the bad and the somewhere between stunning and ugly' doesn't have the same ring.

If this paper achieves its aim, it should leave the reader with the impression that for an organisation to set up its own complete infrastructure is no trivial matter. It is left as an exercise for the reader to determine that acquisition costs are significant, but hopefully it is shown that the life cycle costs, especially for facilities providing Identity Certificates can be huge. Those organisations that cannot justify the resources to supply their own PKI in house, can probably justify a level of outsourcing the infrastructure, or even relying on 3rd Party accredited commercial facilities. If a level of outsourcing still cannot be justified, then there is a very good chance that there are no business drivers for using PKI. For those organisations that can justify using PKI, the paper offers guidelines in determining an appropriate architecture.

2 What is PKI?

2.1 The Basic Technology

PKI provides cryptographic services through the use of special algorithms. These do not as traditionally done rely on symmetric keys for encrypting and decrypting. With symmetric algorithms the same key that is used to encrypt data is used to decrypt. Whereas PKI uses matched key pairs where one key is used to encrypt data and the other matching key is used to decrypt the data. By convention one key is referred to as the private key and the other is referred to as the public key. The latter key gives PKI its name. The basic concept is that one key (the public key) can be published (made public) while the other key is maintained in secrecy. A more detailed explanation of the algorithms can be obtained by reading Applied Cryptography [11].

With PKI these services become a lot simpler to manage, mainly due to key management issued as follows: -

- With symmetric encryption the number of cryptographic keys to be managed for 'n' users will be 'n²' keys. With symmetric encryption every user has to share an encryption key uniquely with each user they wish to communicate with (this means that for 10 users one hundred keys have to be securely generated, for 1,000 users one million keys

need to be generated). It is easy to see that this does not scale well. Also once the keys have been generated each key has to be securely delivered to the two users sharing that key.

- With PKI based encryption the number of cryptographic keys to be managed for 'n' users is from $2n$ keys. Additionally PKI can allow half of those keys to be published in a directory or a key server, and the other half of the keys only has to be distributed to one user. Also the technology itself only requires one user to know any private key (conforming to the adage that if 2 or more people know a secret, it is no longer a secret). As will be discussed later, there will be good reasons (yes there are bad reasons as well) for private keys used for decrypting data to have secure copies made, and very good reasons for keys used for signing data NEVER to be copied.

This significantly simplifies the cryptographic key management problem of key distribution, a problem that even the traditional encryption world (symmetric) is turning to PKI to solve.

There are downsides to PKI when compared to symmetric processing, the primary one being that numerical processing to perform PKI operations is very intensive especially when compared to symmetric processing, estimated at somewhere in the order of 2 to 3 orders of magnitude. For large volumes of data this can be VERY significant (If it takes a symmetric algorithm 3 seconds to process, it could PKI 300 to 3000 seconds). This is obviously unacceptable for everyday commercial type transactions, however (explained later in this section) there are significant efficiencies to be gained which, also provide additional functionality.

The remainder of the issues are usually based on mathematical comparisons of algorithms. This results in an assessment, or normally an opinion, on the 'secureness' or 'trust', they provide. These issues only become significant when extremely high trust models are required, where typically risk management becomes risk avoidance (cost is at best a secondary concern).

Although there various PKI algorithms, and significant efficiencies can be obtained by combining the use of PKI with symmetric algorithms, there are only four basic user operations, which can be performed with PKI: -

- **An individual (the originator) can encrypt with their private key** – This allows any other user (the recipient) with access to the originators public key to decrypt. This is not providing a very secure confidentiality service (anybody can decrypt the data), but it does allow the recipient to know that the data could only have come from the originator, provided the decrypted data is not 'garbage' ('Garbage' would imply that the data was either not encrypted with the originators key, or it has been modified/corrupted since encryption). This process underpins the concept of a applying a **digital signature**.
- **Individuals (the recipient or recipients) can decrypt with the originators public key** – By being able to decrypt the data (into data that has meaning (instead of 'garbage'), the recipient knows the only person who could have applied the encryption was the originator (the only person who should have access to the private key). This process underpins the concept of verifying a digital signature.

- **An individual (the originator) can encrypt with the recipients' public key** – This encryption can only be decrypted with the recipients' Private Key. This basic process effectively provides a **confidentiality** service similar to the service provided by the traditional symmetric cryptographic algorithms in that only the intended recipient has access to the data.
- **Individuals (the recipient or recipients) can decrypt with their private key** – This decryption can only be performed on data encrypted with the recipients public key otherwise "garbage" will be produced. If garbage is produced then it is likely that the data has been corrupted/modified subsequent to encryption. Although it is possible the data was encrypted with the wrong public key, this essentially fails safes, where the confidentiality is maintained although a self-inflicted Denial of Service occurs.

The operations above are matched pairs where the first two could loosely be considered the inverse of each other with one using the private key of a key pair and the other using the public key of a key pair. Similarly the second pair of operations use the other key pair. It should be noted that at all times it is only the key holder that has access to and use of the Private key.. By adjusting the way we use the above processes, it is possible to provide Integrity and Non-Repudiation services, and at the same achieve significant performance improvements.

For signing (and subsequent verification of a signature) a transaction, instead of encrypting the entire transaction with the originators private key to digitally sign a transaction, a hash (a complex checksum) of the transaction is calculated and only it is encrypted with the originators private key, providing an alternative mechanism of signing a transaction. The encrypted hash is now included with the otherwise untouched transaction and can be used as a separate piece of data namely the signature. This not only improves performance, but also provides an Integrity service, by including a cryptographically protected checksum with the message.

The problem with the original process is, if for some reason the contents of a transaction has changed (due to transmission errors or deliberate modification) there is no way for the recipient to detect it (at least electronically). Errors could be detected manually by inspecting the transaction and seeing 'garbage', however for complex data structure and binary data, errors could very easily be missed. In fact the original process is not very effective as a digital signature, and the revised process of signing a hash is much stronger, as there is now a checksum included with the transaction, which is signed by the originators private key. This means that the recipient can now verify the integrity of the transaction by regenerating a hash of the received transaction and comparing it to the received hash. If the generated hash of the received transaction matches the received hash, and the received hash was produced by decrypting the received encrypted hash using the originators public key, the following can be assumed: -

- The originator encrypted the hash; and
- The hash and transaction were not changed during transmission.

Note: The inference is that the transaction and hash from originator arrived unchanged.

If they didn't match, the following can be assumed: -

- Either the originator did not encrypt the hash of the transaction; and/or

- The transaction, and/or encrypted hash, was modified/corrupted during transmission.

Note: It is not possible to tell which or if both events occurred, only that at least one of them occurred.

An excellent improvement in performance is obtained as the processing required to generate a hash is comparable to the processing required to symmetrically encrypt, and the PKI encryption is only performed on up to a few hundred bits (dependant on actually hashing and signing algorithms used)

This can work with multiple recipients very efficiently by, symmetrically encrypting the data, and only PKI encrypting a copy of the Symmetric Encryption Key (SEK) for each recipient. The main data is only encrypted once and can be any size, and multiple encrypted copies of the SEK (very small with commercial algorithms only using up to 192 bits) should still have a combined size that is insignificant when compared to the main encrypted data. These encrypted SEKs are transmitted to the recipients over the same media as the encrypted data and thereby provide a Key Distribution mechanism, which the symmetric world is starting to use for their own key distribution.

2.2 Certificates

The basic technology as explained makes some assumptions for it to be useful, these are: -

- The public keys can be trusted to match the private key of the organisation/individual that business is to be conducted with;
- An organisation/individual will not repudiate a transaction; and
- That a third party does not need to be shown proof that a transaction has taken place.

Of course there are other more basic assumptions, such as the algorithms are sufficiently resistant to crypto-analysis, and that the private keys have in fact been kept private. Covering the first assumption is outside the scope of this paper, and requires a comprehensive understanding of number theory. The second assumption relies on ensuring the owner of a private key is responsible for its privacy (much like the PIN number of a credit card, indeed smart cards used for PKI often require a PIN number to be entered before they can be used).

When conducting business, especially when dealing with organisations for the first time, assumptions can be imply risk. A 'technology based' solution is of little use unless the risks associated with it are reduced to an acceptable level.

To address some of these risks, PKI is usually enhanced by the use of certificates issued by Certificate Authorities (CA). In simple terms, a CA digitally signs a collection of data, consisting of at least an identifier tag, a public key and a validity period. Certificates normally conform to the ITU X.509 recommendation [12], and there are compatible IETF Request For Comments equivalents. Additional data is required, to ensure the correct operation of certificates. representing an organisation or individual. Providing the CA¹ can be trusted to issue certificates only to identities

¹ It is fairly common practice to split the CA into entities called a CA and a Registration Authority (RA), where the RA performs identity checks on applications, and if successful submits a request to the CA for certificate creation. This separation can simplify the process of satisfying legislative requirements on privacy. It also allows different parts of the organisation, or even different organisations best suited to performing each function to do so.

they have undergone a process of proving their identity (usually through documentary evidence), then a trust hierarchy has been created. This is where the CA is at the top of the hierarchy, and is implicitly trusted by users of the certificates it signs. Nesting of CAs can be achieved to create a network of CAs that can trust each other.

Now when a transaction needs to be verified, a recursive process is undertaken which follows the initial checking of the transaction against the originators public key. The public key (in the form of the certificate) is validated against the certificate of the issuing CA. This process is repeated until a common trust point between the originator and recipient is reached. If no common trust point is reached, the transaction cannot proceed, as an acceptable trust model does not exist between the two parties. If a common trust point is reached, and all signatures verify (transaction and certificates), then the transaction can proceed according to each organisations business rules.

In more complex architectures, such as required for national and international architectures, mechanisms such as cross certification can be used to create a 'trust' path between CA that are not part of the same infrastructure.

CA facilities generate certificates, which also include a reference to a Certificate Policy (CP) that identifies what the certificate can be used for, what the responsibilities/liabilities of the three parties (the certificate holder, the relying party and the CA) involved in transactions using the certificate undertake if they accept the use of the certificate. CA facilities also publish a Certification Policy Statement (CPS), which describes how the CA operates and the types of certificates it can use. When cross certifying, these documents needs to be interpreted to assess acceptability for cross certification and end user applications need to be cognisant that different CPs could be associated with the certificates they are now dealing with.

A final item to cover concerning certificates is that as a certificate is issued with a validity period, it could become necessary to cancel that certificate, this cancellation process is known as revocation. When a recipient application is processing an identity certificate, it needs to check the status of the certificate. It can do this by examining a Certificate Revocation List (CRL) which contains a list of all the certificates issued by the CA that are revoked. This list is signed by the CA, and must therefore also undergo signature validation. The CRL is published on a frequent basis, to a location that is either known to the recipient application, or is identified in the identity certificate, thereby permitting an up to date list to be available. New protocols have emerged which can improve the performance of CRL processing by setting up services, which validate a certificate on behalf of the recipient application and return a Pass/Fail notification.

2.3 Attribute Certificates

PKI services can provide the building blocks necessary to develop electronic business applications. However it is important to ensure that the appropriate trust models are in place. An important part of the business 'decision making' process is ascertaining where responsibility, liability and consequences lie and then developing any systems to address the business interests in those areas. Before ANY technical solution or system should be considered, it is essential that there is a sound and effective business case justifying its implementation.

Without a business case there is a huge risk that the technology is being used for technologies sake, which rather than enhancing the organisations operations can in fact cause serious harm to the organisation. PKI is capable of significantly improving

the security and efficiency of an organisation when implemented appropriately, but incorrectly implemented PKI can open huge vulnerabilities while giving the organisation a false sense of security in the belief that they are now safer than before PKI was implemented.

In recent years, certificate technology has come to recognise that with the use of additional types of certificates known as Attribute Certificates, attributes about the certificate holder can be bound to that holder without having to include the information in the identity certificate. This is achieved by creating additional certificates, which instead of referencing the identity by name and including their public key, it instead includes a reference to the certificate (normally the identity of the issuing CA, and the certificate serial number of the identity certificate).

It is important to note that not all attributes need to be concerted to attribute certificates. In many cases a simple lookup table, database query, or directory attribute referenced by the relying application. Attribute certificates only need to be used for those attributes, which grant significant access or privilege for the identity.

This creates several advantages. Changes to attributes associated with an individual can be effected without reissuing an identity certificate to the certificate holder (To decrease risks of key exposure, it is 'best practice' when issuing a new identity certificate to associate it with a new key pair). Different validity periods can be associated with the attributes, even down to validity for a single transaction, but not longer than the validity period of the identity certificate.

Usually attribute certificates will identify access rights and/or privileges associated with an identity. They can then be used by identities to authenticate the validity of identities to perform specific transactions. This is normally accomplished in one of two ways: -

- The transaction originators application can include the appropriate attribute certificates with the signed transaction. The recipient application validates the signature and the validity of the attribute. This approach allows the originator to control who gets access to the attribute certificate.
- The transaction originators application only signs the transaction and does not include any attribute certificates. The recipient application validates the signature and then looks in a certificate repository for the appropriate attribute certificate, and the recipient application then validates it. This approach allows the certificate repository to control access to the attribute certificate.

Depending on the validity period of the attribute certificate, CRL processing may be required: -

- If the validity period is short enough, that the risk of not revoking a certificate is acceptable, or the validity period is comparable to an achievable CRL publication frequency, then a CRL does not need to be published. This is known as 'short lived certificates' and can simplify the PKI processes.
- If the validity period is not short enough, such that risk of not revoking a certificate is unacceptable, then a CRL does need to be published. This is known as 'long lived certificates', CRL processing needs to occur.

There is still debate occurring as to whether long lived certificates should be converted to short lived certificates, by using certificate servers which re issue certificates until a validity period expires. There is also debate as to whether certificate servers should be used as the certificate repository instead of directory (which could be replicated as required). Certificate servers can definitely get around the CRL processing requirements for attribute certificates, and they could definitely act as data repository. Certificate servers also introduce yet another high availability service requirement into the infrastructure, and as they don't currently support the replication capability of directories can be considered a single point of failure. Additionally, the tables in the certificate servers defining how long an attribute certificate should continue to be reissued and what value the attribute actually could require the PKI protection, this means that the certificate server may require to be a high performance machine to process the additional short lived certificates it needs to issue versus the less frequently issued long lived certificates and appropriate associated CRL frequency.

I believe that for attributes that have fairly static values the 'long lived' and CRL option is appropriate, and that for attributes that are rapidly changing, the 'short lived' certificate is appropriate. I also believe that on-line certificate servers should be reserved for extremely short lived or the use one transaction only certificates, and that attribute certificates should be published into a directory like service that can be replicated for availability and performance.

PKI can help significantly with providing authentication services (Identity certificates), authority checks (Attributes and attribute certificates) and transaction records (appropriate storage of the transactions, signatures, and attributes). In fact, one of the key drivers behind the growth of PKI is the ability of PKI to perform these functions with a high degree of trust.

3 Why are there constraints?

3.1 The common business drivers

Many organisations are under increasing pressure to provide on-line services just to maintain their market position. Competitor organisations are adopting on-line marketing and conducting business utilising electronic means. Customers are developing an appetite for dealing with organisations from the comfort of their own premises. All these pressures drive a need to follow the market trend of dealing with customers and other organisations on-line, quite often to the point where business is conducted where the end parties have never met, or even done business, before.

- Having said all this, the basic drivers that apply to conducting business online, are the basic drivers that organisations have when conducting business face to face and on paper, principally: -
- What confidence is there in the Identity of the organisation/individual we are dealing with?
- What type and value of transactions are acceptable with the identified party?
- What is the contractual/legal framework applicable to the transaction?
- At what point is the transaction binding on all involved parties?
- What are the residual risks, liabilities and responsibilities with the transaction after mitigation?

Although this section is not strictly covering the legal requirements, there are two documents that guide an organisation as to the minimum information required to record a transaction (The Financial Transactions Report Act [6]), and the Guidelines on Identification [7] issued subordinate to the FTR Act.

3.2 Why is there legislation?

There are legal constraints², or obligations, that should be included in considerations for the implementation of PKI. Probably the most important of these are those involving the protection of privacy information, the requirement to ensure the integrity of the organisations operations, and conformance to the appropriate government or corporate legislation, as well as protecting organisational IP. Public concern over the access and use of personal information is very high, due to the damaging, and sometimes irreversible, effect of inappropriate or accidental disclosure.

There are several legislative acts, bills and government guidelines (having the weight of legislation behind them, which have an impact on an organisations choice to use PKI to conduct some of its business. Therefore they have to be taken into consideration when selecting an architecture and developing policy and procedures.

Some of the legislation and guidelines provide the necessary legal framework to enable electronic business, and interoperability, to take place. Probably the most significant of these are: -

- The Electronic Transactions Bill 1999 [10], which stipulates: -
 - 1) A transaction is not invalid because it took place by electronic means; and
 - 2) The following requirements under Commonwealth law can be met in electronic form: -
 - Give information in writing;
 - To provide a signature;
 - Produce a document;
 - Record information; and retain a document.
- The Gatekeeper Strategy [8], which defines the minimum standards to be followed in order for electronic commerce to be undertaken with Government Agencies. An accreditation process is imposed upon organisations wishing to claim Gatekeeper compliance to ensure that, not only are the minimum standards complied with, but also that the organisations operations conform to their policy and procedures are suitable for the uses of PKI undertaken. This accreditation is subject to regular reviews. Gatekeeper in some ways is a 'top level' document in that it references other documents, which need to be complied with. In summary it requires CA facilities dealing with Government to comply with ACSI 33 (security guidelines) [1], requires critical components to be security evaluated against the Common Criteria [4], resulting in the listing on the Evaluated Products List [2]. One of the biggest risks to

² All discussions concerning legislation, and Government Agencies, are based upon the Australian environment. Even so, similar legislation is appearing in nations globally. In fact, legislation is becoming more restrictive as organisations are becoming more dependent on the Internet and On-line transactions. Therefore even though the specifics of the legislation may differ, there basic intent appears to be the same, this paper should retain significant validity in many countries.

personal information is at the registration process resulting in guidelines covering Registration Authority Accreditation [5]

The remainder of the legislation and guidelines are there to protect the individuals and third party organisation rights and personal information. The most significant of these is the Privacy Act of 1988 (recently amended and enacted December 2001). This act has resulted in the development of Information Privacy Principles [3], which government agencies and organisations (of significant size) are required to comply with (refer to section in this document entitled Protection of Personal Information). Subsequent to the release of the IPPs the Office of the Federal Privacy Commissioner has produced Privacy guidelines specific to PKI [9].

Implementation of PKI therefore needs to be carefully managed to ensure that the appropriate measures are in place, usually to the satisfaction of a Privacy Commissioner, to meet the legislation. Failure to meet legislation could obviously result in legal action where the implementation of PKI could come under extreme scrutiny by people who would rather be somewhere else in a case where the organisation is often perceived as big brother and the complainant is perceived as the defenceless individual.

3.3 The Organisations Interests and Obligations

Under company laws, with similar (and usually stricter) legislation for Government Agencies, the organisation is required to maintain appropriate records of its operations, and its officers are to protect the organisations interests (It is usual for an organisation to be considered a legal entity in its own right).

For transactions this is accomplished by identification of the identities involved in the transaction, risk assessment of those identities based on the transaction value and known attributes of the identities involved, and ensuring sufficient audit material recorded concerning the transaction. This not only complies with the applicable legislation but also provides legal recourse for organisations, and individuals, to seek compensation for unsatisfactory transactions.

3.4 Protection of Personal Information

In addition organisations are required to provide an appropriate amount of protection concerning any personal information it may be required to collect, store and process as part of its business.

The Australian Government has developed the following Information Privacy Principles [3] to support interpretation of the Privacy Act of 1988: -

- 1) The manner and purpose of collection of personal information.
- 2) Solicitation of personal information from the individual concerned.
- 3) Solicitation of personal information generally.
- 4) Storage and security of personal information.
- 5) Information relating to records kept by record keeper.
- 6) Access to records containing personal information.
- 7) Alteration of records containing personal information.
- 8) Record keeper to check accuracy etc of personal information before use.
- 9) Personal information to be used only for relevant purpose.
- 10) Limits on use of personal information. And
- 11) Limits on disclosure of personal information.

These IPPs appropriately inhibit an organisation from: -

- Collecting all personal data they like without justification;
- Not providing appropriate attention to the accuracy of personal data;
- Performing any processing of personal data without justification;
- Data aggregation or profiling without justification; and
- Disclosure of personal information without justification.

This impacts PKI because certificates are all about publishing or distributing personal information. Inappropriate controls over the quality of data collected, and how and where it is used, can very easily be shown to be non-compliant with the legislation.

4 How to select an Architecture

This section does not define the architecture for any organisation instead it attempts to identify key decision areas and some of the possible decisions.

For organisations that are serious about electronic business, the architecture needs to be developed that, meets the business requirements of the organisation and the legal constraints in a cost effective manner. Addressing the following Key Criteria should assist in identifying an appropriate architecture.

4.1 What type of CA/RA facility is to be used

Identifying the business requirements is paramount here. There are basically 6 options, which are as follows: -

- The organisations implements its own (independently accredited) CA;
- The organisation implements its own unaccredited CA;
- The organisation authorises an accredited agent to provide its CA capability (Certificates are branded as if the organisation issued them);
- The organisation employs the services of commercially available accredited CAs (The organisation accepts appropriate certificates issued to their clients, and the certificates are not branded as issued by the organisation);
- The organisation employs the services of a commercially available un-accredited CA (the certificates are not branded as the organisations, and the CA and RA facilities have not undergone independent accreditation).
- The organisation accepts self signed certificates, which means certificates which have not been issued by any CA and in fact are signed by the private key matching the public key in the certificate.

These are listed approximately in order of cost of acquisition set up, although life cycle costs need to be considered especially with outsourced certificates where licensing costs can vary considerably and result in recurring or escalating costs.

As the costs of setting up an accredited CA and RA facility can run into the millions, it is usually only organisations with a large client base (10s if not 100s of thousands) who are very well funded, and organisations with very specialised requirements that can both justify and fund their own CA.

To assist in determining the business needs the organisation should determine: -

- Determine the strength of identification required when dealing with clients. Technically this should be a risk assessment based on the value of the business transactions undertaken weighed against the

likelihood that the potential client could abuse or renege on the transaction. The result of this assessment should result in a requirement for 0 points, 50 points, 100 points or 150 points of identification required. (0 points implies there is no risk, 150 points implies there is significant risk due to value and likelihood. As a guideline 100 points is the normal requirement to open a bank account). Guidance from the FTR Act and Guidelines on Identification can assist here.

- Determine whether the organisation needs to identify clients themselves, or whether an agent could do it for them, or an independent accredited CA, or for a 0 point identification a self signed certificate or independent non accredited CA. It needs to be remembered that it takes a specialised skill set to perform EOI checks where very few organisations would have the skills or resources to equal or exceed the ability of an accredited CA.

An assessment needs to be made that if there are any additional costs to the organisation running its own facility versus an otherwise acceptable form of outsourcing that they can be justified by specialised requirements (i.e. If the number of clients does not provide a direct cost benefit for running their own CA, then is there a justification such as exceptional EOI requirements or particularly sensitive privacy information that can only be exposed within the organisation).

The number of facility options increases by considering the CA facility separate to the RA facility. This approach can offer easier compliance with the Privacy Act by separating the EOI collected at registration, from the facility issuing the certificates, but it does complicate the liability and trust model by including an additional link in the chain.

4.2 What information is to be included in the certificate versus providing it elsewhere.

As an Identity certificate is a certificate binding an individual or organisation to a public key, any additional information that could be usefully included in the certificate will usually be about that identity. Additionally it will usually be identifying a privilege, access right or values associated with that identity. This means that the more information stored in a certificate the greater the difficulty in meeting, or proving compliance, with the Privacy Legislation.

A balance needs to be struck between: -

- The costs of meeting the privacy requirements. As information about an identity is increasingly aggregated in one place the level of protection and restrictions on disclosure increase. Each Application should only have access to the privacy information absolutely necessary to perform its function and applying those access controls is increasingly difficult if all the information is stored in the identity certificate;
- The costs of satisfying the information requirements of the organisation, by means other than the identity certificate (eg databases containing the additional information to support the PKI application, or attribute certificates separate but bound to the identity certificate, which contain the information). There are potentially additional costs here in providing more intelligent PKI enabled applications that interrogate the

database or process an attribute certificate. However it is much easier to put appropriate access controls in place, and increase the likelihood of the Identity certificate having greater utility whilst still meeting the privacy requirements.

- Can the attributes (in certificate or non-certificate form) be managed or sourced from a third party. As the use of attributes separate to the identity certificate increases, scope for third party organisations to provide an attribute service for other organisations. The third parties may have their own internal requirement, and the provision of a service could be a cost offsetting exercise. Alternatively their primary business could be in providing an attribute service as they may be better positioned to validate those attributes. This option is where the Privacy Act, and IPPs will really start to come into play.

4.3 Seed application versus the use of PKI within the Organisation

It would be unusual for a non-PKI'd organisation to be able to sit down and correctly identify all the business processes and applications that they will ever want to become PKI enabled. More usually, there is a 'killer application', which by itself can justify the substantiation of a PKI.

The trick with the killer application approach is to ensure that the thought processes, used to determine an appropriate architecture, considers a wide enough scope. The type of potential future applications such as physical access, financial delegations, resource management, and electronic purse, and not necessarily actual business applications, should be taken into account to protect the selected architecture from becoming obsolete each time a new PKI application comes online.

It will be expensive enough incorporating new attributes as certificates (or as data in databases or directories), without having to rip out the existing infrastructure due to insufficient flexibility to accommodate new requirements. Remember PKI applications themselves are often uniquely customised to meet the PKI architecture, especially bespoke applications, and revoking all certificates, re-customising your applications and reissuing new certificates is certainly not cheap, not fun and not instantaneous (i.e. there could be considerable disruption to service during transition), although I suppose some individuals may see it as an excellent method of job preservation.

In short the architecture should be selected that allows the organisation to define a static set of data for inclusion into identity certificates (revoking and reissuing is minimised), common methods for accessing certificates should be predefined such that application customisation gets simpler with each new application (the benefit of experience when repeating a process), and hopefully generates a common toolset to allow easier customisation of applications. With this approach privacy issues can be addressed in a methodical manner, and the rules and approval process for incorporating new data can be pre-defined.

By doing this it should be easier to develop the appropriate documentation to support the expansion, and if necessary provide an easier path to re-accreditation.

5 Conclusion

Despite the impression that the business requirements can be at odds with the legislative requirements, it has been shown that a suitable architecture should be possible that meets both sets of requirements. Critical areas for consideration need

to be what data (particularly privacy information) is being used and how much of the PKI can be outsourced. It should be remembered that effective Systems Engineering and Acquisition processes should be followed in an undertaking as large as developing a PKI, even though this paper did not cover specifically cover the topic

6 Glossary

Acquisition Process	A business process, or method, to ensure that the acquisition of a system is done in accordance with company policy.
Attribute Certificate	An electronic document, signed by an Attribute Authority, which is bound to a (Identity) Certificate, and contains additional attributes about the key holder of the (Identity) Certificate.
Attribute Authority	A body that signs and issues digital certificates, which binds the issued certificate to an identity certificate.
Bit	A contraction of the words Binary digiT, representing the smallest piece of information possible in digital form a '1' or a '0'. The interpretation of a bit is dependent on context.
(Identity) Certificate	An electronic document, signed by a CA, which identifies a key holder, binds the key holder to a key pair (by specifying the public key), references a certificate profile and defines the validity period of the certificate. It may contain additional data.
Certification Authority (CA)	A body that signs and issues digital certificates, which binds clients to their keys.
Certificate Revocation List (CRL)	A special certificate issued by a CA, which is a signed list of certificates that have been revoked by the CA.
Confidentiality	A cryptographic service to restrict the viewing of a document to key holder(s).
Crypto-analysis	The analysis of cryptographic algorithms, and their implementation, to determine any weaknesses.
Cryptography	The development or use of special algorithms to apply to data, to improve its confidentiality, integrity or authenticity.
Digital Signature	An electronic signature created by using a Private Key, thereby binding the document to the key holder. Usually applied to the hash of a document to provide a document integrity check.

Evidence of Identification (EOI)	The evidence provided at registration, which demonstrates the identity of the applicant.
Endorsed Products List (EPL)	The list of products, which have successfully undergone an independent evaluation. The evaluation checks a products ability to provide specified security functions as detailed in a Security Target in accordance with Common Criteria.
Gatekeeper Accreditation	An Australian Government accreditation granted on the basis that the CA meets the criteria set out in the Gatekeeper Strategy report.
Hash	The result, of applying a special cryptographic algorithm to data, which is of fixed length and with an extremely high likelihood that the specific data is the only data capable of resulting in that particular hash. Another important feature of the hash is that even for an extremely small change in the data (even 1 bit), a very large change occurs in the hash.
Identity	The name (normally the well known name) associated with an individual or organisation. For the purposes of transactions, and PKI, an Identity is indivisible meaning that the existence of sub-identities cannot be determined from the name.
Integrity	Ensuring, or testing, that a document has not been changed. Usually applied in conjunction with a digital signature
Information Privacy Principles (IPP)	A set of guidelines to assist organisations in complying with the Privacy Act 1988 (Amended 2001)
Key	A data element used by a cryptographic algorithm used to encrypt or decrypt information. Can refer to the single key for symmetric algorithms, and one or both keys of a Private and Public key pair.
Key Pair	A pair of asymmetric cryptographic keys (i.e. one decrypts data encrypted with the other)
Non-repudiation	The ability to prove an entities role in a transaction. To be useful this proof needs to be obtainable independent of the entity.

Originator	Used to refer to refer to the initiator of a specific leg of a transaction.
Public Key Infrastructure (PKI)	The particular implementation of Public Key Technology described in (Accredited) documents, under which Keys and Certificates are issued
Private Key	One key, of a key pair, that must be kept secret by the Key holder. It can be used to decrypt data or sign data (rarely used to do both due to technical and privacy reasons)
Public Key	One key, of a key pair, which must be accessible to third parties intending to perform transactions with the key holder. It can be used for encrypting data or verify a signature already applied to data (rarely used for both due to constraints on the matching private key).
Recipient	Used to refer to the responder of a specific leg of a transaction.
Registration Authority (RA)	An entity, usually closely related to the Certificate Authority, which registers applicants for keys and certificates. The RA usually performs the specified EOI checks before submitting the application to the CA.
Transaction	Used to refer to an individual piece of business involving two or more identities. A transaction can be synchronous and real time requiring all identities to be concurrently involved. It can also be asynchronous where as few as one identity is involved at any particular stage through the transaction (e.g. e-Mail)

7 References

1. ACSI 33 http://www.dsd.gov.au/infosec/acsi33/acsi_index.html
2. Evaluated Products List www.dsd.gov.au/infosec/aisep/EPL.html
3. Information Privacy Principles www.privacy.gov.au/publications/ipps.html
4. Common Criteria www.commoncriteria.org/cc/cc.html
5. Registration Authority Accreditation www.govonline.gov.au/publications/RAAccreditationCriteriaV10.pdf
6. Financial Transaction Reports Act 1988 <http://scaleplus.law.gov.au/html/pasteact/0/59/top.htm>
7. Guidelines on Identification <http://www.austrac.gov.au/text/guidelines/guidelines/guid3.html>
8. Gatekeeper Strategy www.govonline.gov.au/publications/GatekeeperStrategy.pdf
9. Privacy and Public Key Infrastructure www.privacy.gov.au/publications/pki.doc
10. Electronics Transactions Bill 1999 www.aph.gov.au/parlinfo/billsnet/99131.pdf
11. Applied Cryptography Author B. Schneier – Publisher Wiley
12. X.509 Certificates ITU X.509 Recommendation

Note 1 With the exception of the references to 'Common Criteria', 'Applied Cryptography', and 'X.509 V3 Certificates', references have been sourced from Australian Government based web sites, and have been treated as authoritative (or representing authoritative) at least for Australian purposes. The remaining references are to industry recognised sources or documents.

Note 2 All URLs were valid as of 20th March 2002