Stepping Beyond the PKI Pilot

Steve Purser

Steve Purser is Director ICSD Cross-Border Security Design and Administration at Clearstream Services, Luxembourg. Steve is also a founder Member of the "Club de Sécurité des Systèmes d'Information au Luxembourg (CLUSSIL)" and author of "A Practical Guide to Managing Information Security" (Artech House, 2004).

1. Introduction

Despite much of the negative publicity that it has suffered recently, Public Key Infrastructure (PKI) continues to play a critical role in underpinning security on the Internet and other potentially hostile network environments. In such environments, it is common to deploy secure protocols to deliver one or more standard network security services (authentication, confidentiality protection, integrity protection and sometimes non-repudiation) so that individuals and/or organisations can communicate securely. Whilst all this can be achieved without using PKI, setting up and maintaining the trust relationships that support these protocols turns out to be a major challenge when there are many parties involved. This is the challenge that PKI solves – it provides a framework, which allows a Trusted Third Party (TTP) to extend trust to other individuals or organisations (known as *subscribers* in this context), thereby providing them with a context to support some kind of exchange or transaction. In addition to providing a trust framework, PKI is also a useful way of standardising the implementation of the security services as it is based on the use of (usually X.509) certificates and standard, public key algorithms.

This all sounds quite neat and tidy. However, if we stop to consider what we actually mean by trust, it quickly becomes apparent that this is not a simple concept [1]. In fact, the extent to which PKI succeeds in implementing a trust framework depends on many factors, including:

- The degree to which participating parties agree on what trust means.
- The criteria used to decide whether a potential subscriber to the system is trustworthy or not.
- The checks that are made to ensure that subscribers meet these criteria.
- The way in which use of the trust is defined and limited.
- The length of time it takes to revoke a trust if it has been abused or is no longer appropriate.
- How disputes are handled when things go wrong.

Providers of PKI respond to these issues by defining a 'set of provisions', or operating model, which is usually documented in the form of a Certification Practice Statement (CPS) [2]. In addition, the use of particular certificates may be governed by specific Certificate Policies (CP) [3] and all these documents are likely to be supported by appropriate legal documentation [4].

Even a description as short as this makes it quite evident that implementing a PKI is no small task. It is therefore surprising to note that many organisations have gone to the trouble of designing and implementing PKI to support an initial requirement (the 'pilot project'), but have been unable to build on this initial step in order to realise the many advantages that a full deployment brings. This is unfortunate because, as its name suggests, PKI is only infrastructure and businesses require secure applications.

The purpose of this article is to help managers realise the benefits of an extended deployment by showing how to define and execute a strategic plan for PKI.

2. Key steps in implementing PKI

One of the factors that have hampered the widespread adoption of PKI is the associated cost. For all but the smallest infrastructures both the investment and the day-to-day running costs are likely to be important. For this reason, achieving a reasonable Return on Investment (ROI) in the area of PKI requires a well thought out strategy, supported by a solid business plan.

In particular, organisations considering implementing a PKI need to be sure that they can progress pass the pilot phase and achieve a high-level of integration with the applications they use or plan to use. A high level of integration will enable the organisation to standardise administration procedures and thereby achieve economies of scale. Similarly, integration often has the effect of standardising the underlying technologies, which reduces complexity and therefore increases control. However, not all organisations will be in a position to realise these benefits and in many cases outsourcing might prove to be a better business decision.

This document proposes a stepwise approach to PKI implementation, beginning with the initial assessment of the requirement and covering all subsequent implementation steps. The key activities comprising this approach are as follows:

- Assessing the requirements.
- Choosing between an outsourced solution and an in-house approach.
- Building the business case.
- Defining the strategic roadmap.
- Planning the implementation.
- Resolving typical issues and problems.

Each of these steps is discussed in the sections below.

3. Assessing requirements and alternatives

Developing a thorough understanding of the requirement and eliminating alternative approaches is the first (and most important) step in implementing PKI.

Given that PKI usually requires a substantial investment, it is important to understand the business drivers that support a potential deployment. Where in-house deployments are being considered, these drivers are more likely to represent long-term goals than short-term concerns for many reasons:

- Deploying PKI is time consuming – outsourcing may be the only realistic option when deadlines are tight.
- Whereas it may be possible to achieve an acceptable ROI on the basis of a short-term business opportunity, the associated revenue stream would need to be very strong.
- Setting up and maintaining an operating environment for PKI requires a lot of effort and expertise. The business case for PKI is likely to be more convincing if it can be shown that the infrastructure can add value to other applications.

Here, it is important to understand that a lot of the effort that goes into designing and implementing a PKI is likely to be in non-technical areas, such as process design, liability control and ensuring legal and regulatory compliance. Indeed, designing and documenting the operating model is likely to be the hardest part of a PKI implementation project.

It can be useful to derive the requirements in a stepwise manner:

- Identify the initial business requirements.
- Identify other applications that could benefit from PKI integration in the future.
- Identify alternative approaches for satisfying both short-term and projected long-term requirements.
- Define criteria for comparing the different alternatives – the relative importance of different criteria is likely to vary from organisation to organisation.
- Ensure that PKI provides the best alternative.
- Define detailed requirements.

By analysing alternative approaches according to their ability to meet both short-term and projected long-term requirements, we show that the proposed solution is the most appropriate approach for the enterprise (which should help bolster the business case). Only once this analysis is complete should work start on the detailed requirements.

4. Choosing between in-house and outsourcing

Once it has been established that there is a requirement for PKI, the first major decision that has to be taken is whether to implement an in-house system or to adopt an outsourced solution. The advantages and disadvantages of each approach are compared in the following table:

In-House PKI	Outsourced PKI
Direct risk control for the enterprise The enterprise can completely define the operating model, thereby tailoring it to its own requirements for mitigating risk. Issues relating to scalability and flexibility can be planned in advance.	Indirect risk control The operating model is largely determined by the organisation providing the service. Customisation can be achieved through contracts.
Direct control of costs The enterprise has direct control of the majority of the costs.	Indirect control of costs Cost must be periodically negotiated with the service provider. This can be difficult when the enterprise is 'locked-in' to the solution.
Requires considerable investment Implementing an in-house PKI requires extensive infrastructure and considerable deployment effort. The payback period may be of the order of several years.	Initial investment likely to be lower Outsourced solutions usually require less investment in the short-term.
Requires maintaining specific skill-sets In-house models require the enterprise to maintain a wide variety of unusual skill-sets covering design, implementation and administration of cryptographic systems.	Skill-sets are outsourced Skill-sets are maintained by the supplier in an outsourced model.
Ability to meet regulatory requirements In general, regulatory bodies require organisations to demonstrate control over their own services and may mandate	May be disallowed by regulatory bodies This is not an option if not allowed by the appropriate regulatory bodies.

limited access to data for third parties. This will probably be easier for in-house implementations.	
Liability is with the enterprise For in-house implementations, the enterprise is usually liable in the event of an incident.	Service provider may offer limited liability cover Service providers may offer limited liability cover.

In general, in-house PKI is likely to be a more attractive proposal for large enterprises that are answerable to regulatory authorities and have to handle a significant level of risk. This is due to the following factors:

- It is easier to customise an in-house PKI to the particular needs of the enterprise. In particular, this solution gives the enterprise more direct control over the legal framework and liability control mechanisms.
- Large enterprises are often in a position to reduce the overall costs by leveraging existing infrastructure and processes. For instance, most large enterprises could reasonably be expected to have in place primary and backup data centres with sufficient space. Similarly, organisations operating in an international environment will probably have already set-up multi-lingual customer support facilities.
- As long as the enterprise can demonstrate sufficient control over its PKI infrastructure and procedures, this solution leaves the enterprise in complete control of the process, which is usually an advantage where regulatory compliance is an issue.
- Large enterprises are more likely to achieve a positive Return on Investment (ROI) in a reasonable timeframe by exploiting economies of scale. In addition, large enterprises will be able to cover the significant, initial investment.
- It may be difficult for such enterprises to find a service provider capable of offering the required service levels to the customer base – this is particularly true for companies with a strong international presence.

Smaller enterprises are less likely to opt for an in-house solution for similar reasons. Smaller enterprises that are subject to strong regulatory constraints should consider alternatives to PKI before embarking on an in-house implementation, due to the high costs and difficulty in maintaining the required skill-sets.

5. Building the business case

Thinking in terms of applications is the key to defining a successful strategy for PKI deployment and a realistic plan for integrating future applications into the PKI infrastructure can form the basis of a convincing business case. A suggested method for realising the full benefit of PKI and constructing the associated business case is as follows:

- Ensure that the pilot implementation is for a high-profile application and that it will be possible to demonstrate results in the near future (ideally 1 year or less).
- Estimate initial (fixed) implementation costs and (ongoing) running costs.
- Identify benefits and classify them as those that can reasonably be quantified and those that cannot.
- Create the financial analysis and supplement this with supporting arguments based on the non-quantifiable benefits.

Agreeing on the right application for the pilot project is critical, as management (and other colleagues) are likely to judge the future potential of PKI based on their initial experiences. A successful implementation will set a convincing precedent for further rollout and it is therefore a

good idea to choose a highly visible application for which the implementation is not too complex. Finally, it should be possible to demonstrate clear results within a reasonable timeframe – this will help convince managers to include the necessary expenditure into the following year's budget to progress with the rollout. If there are no discernable results within a year, it may prove difficult to obtain budget to take the rollout further, which can lead to a loss of momentum and ultimate lack of support.

Estimating costs and benefits is standard practice for most investments, but it is important to ensure that ALL costs have been identified and correctly estimated. Example fixed costs include initial consultancy, hardware and software acquisition, rollout costs and all staff costs associated with the initial deployment. Here, it is important to correctly assess the need for external consultancy and this is especially true where there are stringent requirements on the documentation (CPS, CP and operating procedures). Example ongoing costs include insurance and liability cover, costs associated with administering the system and maintenance costs. When estimating the ongoing administration costs, special attention should be paid to the division of responsibilities (such as the four-eyes principle), as this can have a major impact on the final figure.

Where benefits are concerned, it is worth noting that the whole point of PKI is to reduce risk, and for most organisations the biggest benefit associated with deploying PKI will therefore be a reduction in risk. Whilst it is true that several organisations have succeeded in generating a revenue stream by selling PKI-related services (such as certificate issuance and management), this is becoming increasingly difficult and corporate clients are more likely to be interested in investing in secure applications than infrastructure. Other benefits potentially include increased scalability of operations, standardisation and cost savings through economies of scale. Of these benefits, scalability is a direct result of using public key cryptography techniques and is not therefore strictly related to PKI – but we have already noted that public key cryptography needs an underlying trust network to support the distribution and verification of keys. Standardisation normally reduces complexity and therefore makes it easier to analyse and deal with problems. Finally, by designing applications to use the PKI infrastructure, it is possible to implement a common administration model across multiple applications.

Once we know the costs and benefits associated with the deployment, it would be useful to calculate the return On Investment (ROI). Better still, we could include an estimate for risk reduction by calculating the Total Return on Investment (TROI) [5]:

$$\text{Total Return on Investment} = \frac{\text{Generated Revenue + Generated Cost savings - Value of change in risk}}{\text{Investment}}$$

The difficult part of this exercise is likely to be quantification – consider for example the problems associated with estimating the financial value of a loss of reputation. More generally, it is difficult to quantify changes in risk and this is where the major benefit of deploying PKI is likely to lie. Whilst methods such as Annual Loss Expectancy (ALE) [6,7] and Cost-Benefit Analysis (CBA) [8] can be useful in this respect, most organisations will probably not have the empirical data to estimate potential losses in this area.

Where quantification proves to be difficult, a useful compromise is to use 'order of magnitude' calculations and to support the proposal using strong qualitative arguments. The goal is to

convince executive management that the initiative is worthwhile and this is not simply an exercise with numbers – it is a question of putting forward convincing arguments.

6. Defining the strategic roadmap

The key point to bear in mind when defining a strategic roadmap for PKI deployment is that it is essential to involve the right people. The strategic roadmap needs to reflect the goals of the enterprise and, as such, should be closely aligned with the business strategy. Hence, the first step in deriving the roadmap is to identify the stakeholders. Normally, the important stakeholders will be the different business lines within the enterprise, but it is also worthwhile considering some of the 'support' services as stakeholders. It is likely that risk management, legal, IT production and IT development teams will all play a role in any significant deployment strategy, so it is logical to consider them as stakeholders. Whilst most audit departments would probably not like to be considered as stakeholders, they should at least be comfortable with the approach.

Once the stakeholders are on board, it should be possible either to obtain a copy of the business strategy or at least to identify the key business initiatives for the next two to three years. For more mature organisations there might already be a defined product roadmap, which can then be used directly for planning the PKI rollout. Typically, this product roadmap will involve both bespoke development and product acquisition. If this doesn't exist, it will be necessary to piece together likely future requirements through discussions with business managers. Whatever the situation, it will be easier to integrate PKI with new applications than to retrofit legacy applications.

Although the need for new applications or functionality is likely to be the major input to the strategic roadmap, there are other factors to take into consideration. Development teams for instance play a key role in deciding whether a particular approach will succeed or not and introducing new skills or methods should be thought out in advance. A good example is provided by applications that are likely to be developed using security toolkits. Here it is important that development teams have the time and support required to learn how the tools work and how they fit in with current development methods and tools. Wherever possible, it is useful to introduce such tools at an early stage, because once developers are familiar with them this experience can be used for future projects. Where product acquisition is concerned, it is important to ensure that requirements relating to the PKI are included in Request For Proposal (RFP) documentation or that they are taken into account as part of any comparative studies. Finally, it is useful to look upon PKI as part of a broader set of technologies and procedures that collectively constitute the security architecture – an architectural approach helps ensure that the way in which PKI services are offered is consistent with other security services.

In most cases, the strategic roadmap will probably reflect a phased approach to deployment and will identify important milestones, so that significant progress can be reported back to the management team. Once the initial learning curve is over, it usually makes sense to prioritise important applications, but to allow enough buffer to allow for the implementation of 'quick wins'.

7. Implementation & rollout

Due to the specialised skill-sets involved and the possibility that unusual hardware will be used, there are challenges which are quite particular to implementing and rolling out PKI infrastructure. Some of the most common implementation and rollout issues are:

- Requirements for parallel development and test environments.
- Need for specialised training

- Use of special equipment, such as HSM and smart cards.
- Help desk, technical support and multi-lingual customer support.
- Estimating the administration workload.
- Ensuring business continuity in the case of an incident.
- Integration with current procedures.

Where developments involving PKI are concerned, the requirement for parallel development and test environments can be costly (especially where specialised equipment such as HSMs are being used). Supporting such environments can also be quite onerous, as these environments should not use production keys for obvious reasons. Where test environments are regularly re-built, administering cryptographic modules and performing the associated key management is likely to be time-consuming. Organisations should therefore understand the impact of testing in terms of cost and resource availability and ensure that this is compatible with planned schedules.

Staff training starts in the implementation phase and will continue throughout the rollout process. Initial training will probably be aimed at development and testing staff and aim to ensure that toolkits and specialised hardware are understood and handled correctly. In many cases, it will be necessary to provide an introduction to the important concepts before this. During the rollout phase however, the emphasis changes and the primary goal will be to provide adequate training to support staff and end-users. Correct training for end-users is particularly important because users need to appreciate how cryptographic keys are used and their responsibilities when using them – indeed, for some applications, users may be required to sign some sort of user agreement that defines how these keys can be used together with any restrictions. End-users should also have a basic appreciation of the security mechanisms work. For instance, passwords that are used to unlock a local key store have a different function to those used to authenticate to a remote system, which results in a different set of risks for these two different types of applications (e.g. it might be possible to copy the key store and to try to break the password off-line).

Organisations choosing to deploy specialised hardware, such as HSMs or smart cards will need to ensure that they maintain the necessary skill-sets to support all environments and will also have to budget for upgrades and replacements. Even when local skill-sets are well developed, it is important to ensure that the appropriate vendor maintenance agreements are there as a safeguard.

All the required support procedures should be defined before or during the implementation phase. Typically, this will involve different types of support ranging from internal technical support to formal first and second-line support agreements to support external customers. Where around-the-clock cover is required, this may have an impact on staffing levels. Similarly, there may be a requirement to support staff or customers in several different languages, although organisations requiring such services will probably already have established teams in this area. Finally, maintenance contracts should obviously reflect these support commitments (which might involve 7/7 and 24h/24h vendor support). The testing phase provides an ideal opportunity to re-estimate the amount of administration work involved in supporting the target system.

For most organisations, disaster recovery or business continuity policy will require the establishment and maintenance of a backup site, together with procedures for ensuring that data remains current on the latter and procedures for switching between the two in the event of an incident. Note of course that the backup site has to be configured with live production keys and, after a switchover, revocation information needs to be current also.

Finally, it is worthwhile to take the time to smoothly integrate the PKI administration procedures within the existing procedural framework. This can both avoid unnecessary duplication of work and help ensure that the control framework as a whole is consistent.

8. Resolving typical issues and problems

This section briefly looks at some of the issues commonly encountered when designing and implementing PKI. This set of examples is by no means exhaustive, but hopefully it does provide a good idea of where and how things can go wrong.
Examples of issues associated with the design phase include:

- Selecting a coherent set of standards.
- Deciding on the most appropriate revocation model.
- Choice of the trust root.

Standards are important in the area of PKI because they help ensure that different user communities have a common understanding of what PKI is and how it is used. Standards also facilitate interoperability, which is extremely important when unforeseen circumstances require different PKI implementations to work together in some way. Standards in the area of PKI do not all originate from the same source (e.g. the X.509 standard is a part of the ITU (formerly CCITT) standards, whereas the idea of a CPS originated with the American Bar Association) and they can sometimes be in conflict with each other. Furthermore, when integrating PKI with applications, organisations will be confronted with a whole host of cryptographic standards. This includes standardisation of algorithms (e.g. Federal Information Processing Standard (FIPS) 197, Advanced Encryption Standard [9]), general cryptography standards (e.g. PKCS - Public Key Cryptography Standards [10]), protocols (e.g. RFC documentation) and specialised material (the most notable being the FIPS 140 standard [11]). A good strategy is to decide on a coherent set of standards and to stick to it as tightly as possible.

The revocation model is a critical design decision because it largely determines the *window of risk* between the initial revocation action and the time at which this information is available to relying parties. The classical approach is to publish a Certificate Revocation List (CRL) at periodic intervals, but there are many alternatives to this simple approach [12]. It is worth briefly mentioning the Online Certificate Status Protocol (OCSP) [13], which if correctly implemented can support a close to real-time revocation model.

The choice of the trust root is equally important, as changing an existing trust root is likely to be both difficult and disruptive to users. The trust root tells relying parties who they are trusting and documents such as the CPS and other operational procedures are closely associated with this root. Organisations that change their name, merge with other organisations or are acquired might want to change the name of the trust root without actually changing the root itself (although in the latter two cases it is possible to argue that the trust has changed because the organisation setting up the trusts has itself changed). Alternatively, certain third-party software may require a root certificate from one of the major providers (i.e. present by default in the browser certificate store).

Examples of issues associated with the implementation phase include:

- Interpretation of certificates.
- Inter-operability of commercial products.
- Support for Tamper Resistant Equipment (TRE).

When developing in-house applications to integrate with a PKI the question of how certificates are to be interpreted is under the control of the organisation and it is therefore unlikely that unexpected behaviour will occur. This is not the case however where third-party products are concerned and it is important to check what such products actually do with the certificates they

use. In particular, organisations should check how products detect an handle revoked certificates, which certificate extensions (see [14]) they support and how they process these extensions.

Interoperability problems occur for a variety of reasons and are particularly common when dealing with server-side cryptographic equipment. Hardware Security Modules (HSM) for instance may offer limited possibilities for integration with other software components. Web server software might support a PKCS#11 interface for certain versions of HSM software, but not for others. When support is provided, it may only be partial – continuing with the example of the PKCS#11 interface, the basic cryptographic operations may be supported, but there might be limited support for some of the ancillary functions. Similarly, there may be limitations at the network level, such as support for the SCSI interface but not for an Ethernet interface – this is an important consideration if there is a requirement for an HSM to support several systems.

Examples of issues associated with the rollout phase include:

- Gaining user acceptance
- PKI expertise is expensive and difficult to find.

Ironically enough, although PKI can simplify things for end users in the long-term by standardising the authentication process, the first step may involve additional passwords and new practices and may be perceived negatively. In some cases, users may be required to sign usage agreements (e.g. if they are provided with a signing key) and there may be a non-trivial local installation. The way to deal with these and similar issues is to start preparing users for the impact early in the process – a proactive approach to training should result in users that are capable of weighing short-term inconvenience against long-term benefits.

Although the question of expertise has been included here, it could have been put almost anywhere in this section. Most organisations are unlikely to have the necessary experience in-house at the start of a PKI deployment initiative and this will translate into a requirement for external service providers. Organisations should strive to build up the necessary skills in-house however and this will only happen if enough emphasis is put on transferring skill-sets. The danger here is that companies will want to maximise the added value of highly paid service providers by affecting them to the critical tasks whilst permanent staff continue supporting daily operations. Apart from the obvious problems associated with this approach, it is likely to be demoralising for staff.

Conclusions

Public Key Infrastructure (PKI) continues to play a critical role in underpinning network security services on the Internet and other, un-trusted networks. Simply put, PKI allows Trusted Third Parties (TTP) to extend trust to subscribers, thereby enabling them to participate in some form of secure data exchange or transaction. However, closer examination of this model shows that the extent to which PKI is really providing an acceptable trust framework is dependent upon many factors. These factors include the degree to which users understand the trust model, the effectiveness of verification procedures, limitations on how the trust can be used, revocation procedures and how disputes are handled. PKI providers summarise how they resolve these issues by defining and implementing an operating model, which is usually documented in a Certification Practice Statement (CPS). The acceptable use of particular certificates is often governed by Certificate Policies (CP) and this operational documentation is supported by an accompanying set of contracts. Given the significant effort associated with deploying PKI, it is surprising to note that many organisations have not developed their initial PKI implementation further than the pilot project. This is unfortunate because such organisations are not realising the

potential benefits that an in-house PKI can offer, such as a scalable and flexible way of controlling certain risks and the advantages associated with economies of scale.

Given that PKI usually requires a substantial investment, it is important to understand both the short-term and long-term business drivers that support a potential deployment. It is useful to derive these requirements in a stepwise manner. Before identifying detailed PKI-related requirements, different alternatives should be evaluated against all known requirements to ensure that PKI really provides the best alternative.

Once the requirement has been clearly understood, a key decision to be taken is to decide whether to satisfy this requirement by deploying an in-house system or whether to choose an outsourced solution. The major advantages associated with the in-house model include the increased level of control and the possibility of customising the implementation to meet enterprise-specific requirements. Disadvantages include the substantial initial investment and running costs, the need to maintain unusual skill-sets and the fact that liability in the event of error will be with the enterprise. The outsourced model allows for less direct control, but is likely to be cheaper. In some industries, regulators may not allow companies to outsource PKI. Due to these reasons, in-house PKI is likely to be more attractive to large multinationals with a significant risk profile.

A successful business case will be built on a convincing strategy and it is a good idea to base this on a realistic approach to integrating future applications with the PKI. By selecting a high-profile application that can be integrated with reasonable effort for the pilot project, security managers should be able to demonstrate convincing results within a reasonable timeframe, which will help secure the necessary budget to continue with further phases of the rollout. The resulting business case should include all costs, both fixed and ongoing and compare these with the perceived benefits. One way to achieve this is to calculate the Total Return on Investment (TROI) using ‘order of magnitude figures’ where necessary and to supplement this with strong qualitative arguments.

In order to derive a meaningful strategic roadmap for PKI deployment it will be necessary to involve the right people from the start. Whilst the most important stakeholders will be the managers of the major business lines, other functions such as risk management, legal and audit should be involved in the process. The PKI rollout strategy should ideally be closely aligned with the business strategy of the organisation – if there is no formal business strategy, it is often possible to put together a vision of the future by questioning business managers. Although the plan for acquiring or developing applications is likely to be the major driver for the roadmap, other factors affecting success should also be taken into account. Examples of such factors include the complexity of the work required to integrate particular applications and the time required to train project staff.

Implementation and rollout of PKI and PKI-enabled applications is complicated by a number of issues. Firstly, there is usually a need for separate development, (possibly parallel) test environments and a stable production environment, which is both costly and requires a lot of support. Supporting such environments will probably require flexible key management procedures and moving HSMs between environments might require considerable extra administration. Much of the required training should also be carried out in parallel with implementation activities to ensure that teams are up to speed for the rollout period. Support procedures need to be developed to cover a range of activities and any external support should be aligned with the requirements of the organisation. In particular, vendor support for cryptographic hardware and software should be consistent with response times agreed with end-users. The final implementation will need to be compliant with business continuity policy.

Finally, by way of example, this paper has discussed a number of issues that can occur at different phases of the deployment project and has proposed ways of dealing with these issues.

References

[1] S. Purser, "A Simple Graphical Tool For Modelling Trust", Computers & Security, Vol. 20, No 6, pp. 479-484.

[2] S. Chokhani, and W. Ford, "Internet X.509 Public Key Infrastructure: Certificate Policy and Certificate Practices Framework (RFC 2527).

[3] ITU-T Recommendation X.509, "Information Technology – Open Systems Interconnection – The Directory: Authentication Framework" (1997).

[4] C. Adams & S. Lloyd, "Understanding Public-Key Infrastructure: Concepts, Standards & Deployment Considerations, Second Edition", Macmillan Technical Publishing (2002), pp. 183-194.

[5] S. Purser, "Improving the ROI of the Security Management Process", Computers & Security, Vol. 23, No. 7, pp. 542-546.

[6] S. Berinato, "Calculated Risk", CSO Online, December 2002.

[7] "Guideline for Automatic Data Processing Risk Analysis", FIPS publication 65, 1979.

[8] R. T. Mercuri, "Analysing Security Costs", Communications of the ACM, Vol. 46, No. 6, 2003.

[9] Federal Information Processing Standards Publication 197, "Announcing the ADVANCED ENCRYPTION STANDARD (AES)" (2001).

[10] Public Key Cryptography Standards, RSA Laboratories, http://www.rsasecurity.com/rsalabs/node.asp?id=2124

[11] Federal Information Processing Standards Publication 140-2, "Security Requirements For Cryptographic Modules" (2002).

[12] C. Adams & S. Lloyd, "Understanding Public-Key Infrastructure: Concepts, Standards & Deployment Considerations, Second Edition", Macmillan Technical Publishing (2002), pp 105-129.

[13] M. Meyers et al, "Internet X.509 Public Key Infrastructure: Online Certificate Status Protocol (OCSP)" (RFC 2560).

[14] C. Adams & S. Lloyd, "Understanding Public-Key Infrastructure: Concepts, Standards & Deployment Considerations, Second Edition", Macmillan Technical Publishing (2002), pp 72-74.