PKI - An Insider's View (Extended Abstract)

Geraint Price

Technical Report
RHUL–MA–2005–8
27 June 2005

Department of Mathematics
Royal Holloway, University of London
Egham, Surrey TW20 0EX, England
http://www.rhul.ac.uk/mathematics/techreports

Abstract to Technical Report

This technical report is an extended abstract of a report written for the members of the PKI Club at Royal Holloway [1]. This extended abstract makes public the *Executive Summary* along with the original *Introduction and Summary* chapter of that report. The original document was a confidential document based on the cumulative experience of the members of the PKI Club. The sections of the original report which we include here are free of the direct references and quotations which made the original report confidential in nature.

[1] `http://www.isg.rhul.ac.uk/research/projects/pkiclub/`

PKI - An Insider's View

The views and practical knowledge of PKI users and designers

Dr. Geraint Price

October 2004

Information Security Group,
Mathematics Department, Royal Holloway,
Egham, Surrey, TW20 0EX.

geraint.price@rhul.ac.uk

Abstract

This report represents an insight into the experiences of the members of the PKI Club at Royal Holloway. The PKI Club is a group of parties from industry, Government and academia with an interest in PKIs. The club, which was formed in January 2002, has hosted a regular series of informal seminar meetings. At each of the meetings guest speakers have provided expert leadership in discussing the issues of particular relevance to the deployment and running of a PKI. This report has provided an extension to the core activity of the club and is an aggregate representation of the views of the club members on the issues that have been discussed.

Executive Summary

This report represents an annotated summary of the views and experiences of industrial practitioners who have had numerous and varied experiences with Public Key Infrastructures (PKIs). It is not intended as tutorial on PKIs, but as a precis of the issues which many of those in industry consider important to the future of PKIs in practice.

The main body of the work was collected through a series of very frank and in-depth interviews. These interviews were used as a tool to allow members of the PKI Club within Royal Holloway to express any opinion which they considered to be of relevance to PKI. In addition, we include some of our own analysis of the opinions expressed in the interviews. We believe that our ability to compare transcripts in a confidential manner has led to our analysis providing useful additional insights on many of the topics raised.

We begin by asking two simple, but important questions:

1 *Do we need PKI?*

We answer the first question in two parts:

- **Is there a market need for PKI?** We firmly believe there is. With systems such as BACSTEL-IP and CHAPS now mandating its use, PKI has a foothold in the industry. Asymmetric cryptography also provides a small but important set of services that symmetric cryptography fails to provide. Thus, in certain scenarios, PKI provides a competitive security solution when compared to other security technologies.

- **Can we manage asymmetric cryptography without a PKI?** We believe not. Although the shape and scope of the infrastructures might change, management of the public/private keys is a non-trivial task and requires some form of support mechanism.

2 *Where do we need PKI?*

In answer to the second question, we believe that PKI should always be used in support of some other business requirement. Even though the technology can provide very useful services, it is important not to lose sight of the applications which are going to use those services. We believe that the security services need to be promoted more in terms of the types of application services which the security services best supports. However, as many people have noted before, PKI is not a panacea and understanding its strengths and weaknesses is crucially important when comparing it to other technology.

We now overview the main issues covered within this report. We group them together, discussing the points raised in each category. As well as summarising the discussion contained in the main document, we identify key action points that need to be targeted within the industry.

Legal and Regulatory Issues

In a high proportion of implementations, the PKI touches upon a large part of an organisation's infrastructure. Thus, in most cases, the legislation which impacts on a PKI deployment is not limited to that governing digital signatures. Other areas of law raised during our interview process were: Contract legislation; Employment legislation; Data Privacy legislation; Regulation of Investigatory Powers Act; Freedom of Information Act.

The continued uncertainty surrounding the implementation of the legislation to support digital signatures was highlighted as a specific problem area. Due to this uncertainty, many of those implementing PKIs are using contracts as a means of ensuring that digital signatures are honoured, rather than relying upon statute.

In the case of the EU Directive on Electronic Signatures, the issue of control over the private key was seen as a legal hurdle which was having a large impact on the design process.

Regulation was considered to have already provided a driving factor in the uptake of PKI enabled applications. [1]

> **Regulation often provides a focus for what can be a lengthy and complicated design process. Thus, we believe that more should be done to further the impact of regulation regarding particular implementations or vendor products.**

Liability was a concern for many of those implementing a PKI. Finding ways of limiting that liability can increase the design and process costs. Operating a PKI under a membership agreement or scheme helped in this instance. By signing up to a specific rule set, the parties involved can better calculate the risks and liabilities involved.

If there is a requirement for a legal opinion within a project, it was considered to be important to get the lawyers involved as early as possible. The learning curve for those interpreting the legislation was seen as quite steep. Failing to take this into account early on could impact the project later on.

Technical Issues

One of the main selling points of PKI is its ability to deliver multiple security services within the same security architecture. The breadth of these services, coupled with PKI's ability to provide non-repudiation, was seen as a clear advantage when compared with similar security products in some scenarios.

We believe that the most important step is for there to be an increased understanding of how the security services provided by PKI should be used. This is because the services provided by the technology are only of business use in support of an application. Of the three main security services discussed by those we interviewed, here are the issues highlighted for each:

Authentication : Authentication is considered by many to be an enabler for other forms of technology. For example, one implementation used authentication to allow clients to migrate from leased lines to IP-based access.

> **Authentication is usually implemented in support of authorisation. We believe there needs to be greater promotion and understanding of what types of authorisation process are best supported by asymmetric cryptography.**

Integrity : The ability of a digital signature supported by PKI to provide *opaque* forms of integrity checks, where the document can be signed by multiple parties, is used in some fielded systems.

[1] In discussing regulation, we view it in as broad a scope as possible. For example, we consider schemes for individual applications (BACSTEL-IP, IBDE-2), as well as regulation for specific business sectors, etc.

We believe that the ability for digital signatures to provide integrity mechanisms is often overlooked. A few implementations we studied used the strength of this opaque integrity mechanism. Where else might signatures be deployed in this way?

Non-repudiation : There was still interest in the provision of non-repudiation, but the cost and legal uncertainty were seen as prohibitive in some cases.

Unfortunately, the uncertainty surrounding the legal issues and the understanding needed to build sufficient technical support is hampering non-repudiation's development as a service. These issues need to be ironed out if non-repudiation is to be more widely used in practice.

When it came to deciding whether to use hardware or software for the storage of the client private key, hardware was the method of choice for most projects. Having said this, some people saw software key storage as a means of reducing costs either in transition to a hardware solution, or in a scenario where there was a lower risk profile.

Integration has caused problems in many PKI implementations. Here were some of the views held by those we interviewed:

- Integration with specific vendor applications was considered more important than compliance with individual standards.
- Integration with Privilege Management Infrastructures (PMIs) was seen as important. We believe that, because of PMI's potential business benefit in the form of cost reduction, this is likely to become a key driving factor for PKI deployment in the future.
- Integration with other infrastructure components, such as directories, has added complication in many implementations. Solving the problems raised in these instances sometimes relied heavily on previous experience.

It was considered that standards should primarily be relied upon as a template, as well as for providing educational input. Their breadth and general lack of specificity can result in them being unwieldy in practice. As a consequence, adhering to standards was often sacrificed in order to get a working product.

In many situations infrastructure re-use was considered desirable, because of the high cost of deployment, although this is not always possible. For example, some implementations were provided by a third party as a *black box*, whereas others fell foul of scheme or legal requirements. If re-use is required, it would appear to be of benefit to design for such re-use from the outset, as migrating policies and procedures from one application domain to another is a non-trivial task. Although this can increase the initial costs, many believed the alternatives to be more expensive and more difficult in the long run.

PKI-enabled applications that deliver a service to external clients and users are much more difficult to manage than ones that apply to closed systems or user groups. This is mainly due to the fact that the external influences are more difficult to police. This can have a knock-on effect on the quality of the security provided, which can reduce the breadth of the types of transactions which an infrastructure can support.

Infrastructure Management and Administration

Registration of the end users was seen by many to be a critical problem faced by those managing a PKI. The procedures used to identify an individual reliably are not always clearly defined, but can have a large impact on the resulting security provided by the PKI.

In relation to revocation, the main focus of attention fell on whether to use a Certificate Revocation List (CRL) or an On-Line Certificate Status Protocol (OCSP). Broadly speaking there were three separate views expressed:

- CRLs are unacceptable, given the cost of implementing a PKI in the first place, coupled with the fact that security is likely to be used to protect a high value transaction.
- The decision on whether to use CRLs or OCSP should be based on the level of risk associated with the transaction.
- The decision on whether to use CRLs or OCSP should be tied to the time criticality of the action being supported. PKIs are often used to provide an authentication service which supports a separate authorisation mechanism. In some instances it was easier to use the authorisation mechanism to revoke the right to carry out any critical action.

A key point made regarding revocation was the difficulty of identifying who is responsible, and who has the authority for issuing the revocation request.

The confusion surrounding the use of Certificate Policies (CP) and Certification Practice Statements (CPS) was often cited as one of the key problems faced. When deciding whether CPs and CPSs were useful, many people felt that they only came into their own in an infrastructure used between multiple parties. CPs and CPSs could then provide important input into the legal process when defining liability etc. The *de facto* reference for building a CP and CPS – Internet RFC 2527 – was deemed an inadequate guide for specific instances when it came to writing CPs and CPSs.

> **CPs and CPSs are just the PKI instantiation of policy and management aspects of a security infrastructure. These are notoriously difficult to get right. We believe that the best way forward is for better examples to be built around business areas or types of process models. Such examples should be discussed and developed openly, leading to the provision of sets of best practice principles. While RFC 2527 provides a useful guide for those looking to develop their CPs and CPSs, it is too broad to be of use in all but a handful of cases. Having more specific classes and rigorous templates to work from should ease the pain of what is an important task in building a PKI. Having clearer examples should subsequently provide more adequate support for those attempting to draw up their own CP and CPS.**

Commercial Issues

Here are the main reasons given when justifying the business decision on whether to implement a PKI: **Cost** – PKI is an enabler for other cost reducing measures; **Security** – PKI provides a "strong" security measure; **Regulatory** – PKI ensures compliance with new regulation, which is sometimes mandated. Because there has been a fair amount of criticism in the past of PKIs

being put in place without a clear business case, two means of ensuring a strong business case stood out during our discussions:

- Make sure that there is early and regular involvement with the relevant business manager on the project.
- Design the PKI with a particular application environment in mind.

  A common business case we encountered was the need to protect high value transactions. While this is a well understood goal, we believe that if the PKI industry is to maximise its potential, further understanding of the types of transactions which can be supported is required. What are the types of communicating party relationships? What are the process flow models? What type of contractual environments provide the best scope for addressing the process issues? If these and other similar questions can be clearly answered, the industry will have a better understanding of where PKIs can best be used in the future.

Three possibilities stood out when we discussed the reduction of the large cost of implementing a PKI:

- Building shared infrastructure between several organisations.
- Infrastructure re-use within an organisation.
- Use of a hosted service provider.

When considering whether to use a hosted service or to build in-house, here are some of the main benefits demonstrated for each:

Hosted : Generally easier conformance to regulatory standards; clear cost savings; accumulated expertise within the hosted service provider.

In-House : Retaining control over the whole infrastructure can be an issue for high value transactions and sensitive applications; PKIs are often used to support applications and PMIs and these secondary services are generally not as easily outsourced, hence complicating the client/hosted service interaction.

Many felt that hosted services can provide a leading role in expanding the use of PKIs within the industry.

Human Understanding of the Technology

There was an almost unanimous call for an increase in education to support the deployment of the technology. It is critical that people at all levels of the organisation understand the impact. This is important given the wide-ranging effect that a PKI can have on the remainder of the organisation's infrastructures.

Even the knowledge of the technical staff was considered insufficient by many of those who had interacted with others on various projects. Consequently, very few organisations are equipped with the skilled individuals required to implement a PKI.

When it came to discussing the ease-of-use aspect, it was considered prudent to try and remove any trace of the technology from the user interface and, where possible, get it to mimic existing technology of a similar type.

In Conclusion

In concluding this summary, we wish to highlight two points:

- It is important not to underestimate how much impact the PKI has within an organisation. The breadth and variety of the legal, technical and business issues are often overlooked. This has potential for significant impact on the cost and complexity of the design of the PKI.
- Although the issues faced by those supporting a PKI are many and varied, none are insurmountable. In the cases where PKI has been deemed the technology of choice, it has been possible to resolve all the issues. Ensuring that a sound business case is built and well understood should allow for a successful deployment.

Contents

1 Introduction & Summary

In this section, we introduce the scope of this document and provide a summary of the following sections, each of which focuses on a specific area of relevance to PKIs. We also include a personal view of what we believe to be the main issues facing the PKI industry, along with some general conclusions.

1.1 Project Aims and Report Structure

This document is the result of an interview process conducted over several months in 2003. Section 1.2 provides a background to the formation of the project itself.

Section 1.3 provides a brief summary of each of the main sections of the report. Table 1 below represents a high level introduction to each of those sections. The content largely reflects the opinions and comments of those interviewed. Occasionally, we have added our own interpretations to this mix as a means of pulling together diverse comments made by different individuals.

There are occasions within the report where similar issues are raised in different sections. This is due to the breadth of work covered, and the way in which one issue can impact on another. We have kept this to a minimum, using cross referencing where appropriate, but we also have allowed some degree of repetition to aid in the reading of the report sections as stand alone discussions.

Section	Title & Description
2	**Legal and Regulatory Considerations**
	A description of the legal and regulatory issues that have potential impact on PKI deployment
3	**Technical Considerations of Asymmetric Cryptography**
	Technical issues that are of specific relevance to the implementation of public key cryptography
4	**Technical Considerations of the Infrastructure Deployment**
	Technical considerations presented by deploying the infrastructure in support of the cryptography
5	**Management and Administration of the Infrastructure**
	Issues surrounding the management and administration – including the people and processes – of the infrastructure
6	**Commercial Considerations**
	Commercial considerations surrounding the use and deployment of the infrastructure
7	**The Human in the Equation**
	The issues surrounding the human understanding of the technology along with its usability

Table 1: Overview of Chapter Contents

In Section 1.4 we provide some personal reflection on what we have learnt from the interviews. This is included to draw out some often unstated or global issues we encountered. The opinions presented there are not meant to provide a precis of the remainder of the document although there is, by necessity, some feel of a summary in its content.

In Section 1.5 we draw our conclusions from the remainder of this work. These conclusions provide a list of what we believe to be the main concerns for those in the industry, along with a brief overview of how we think these issues should be tackled.

1.2 Setting the scene

The project set up to deliver this report was an extension to the work of the PKI Club at Royal Holloway [2].

At the time of formation of the club, the PKI industry was suffering as the result of people scaling back their expectations for the technology. Although PKI had been around for many years it was perceived to have failed to deliver on its early promise. There were islands of use within the marketplace, but there was a debate over whether its use would extend much further.

Against this backdrop, the club was set up and hosted by the Information Security Group at Royal Holloway, University of London. The aim of the club was to bring together a diverse set of interested parties from industry, Government and academia to discuss issues of relevance to the future of PKI.

The main aim of the club was to provide a series of seminars given by specialists in each field. These were used to share ideas and examine future trends in the area, such as the technical, business and legal areas surrounding PKI. In covering these issues, here are some of the questions which were being discussed throughout the duration of the work of the club:

- What is a PKI?
- Who needs PKI and who gains most from it?
- What applications does PKI really enable and what are the benefits?
- What are the implications of recent electronic signature legislation?
- What is the future of PKI?

During the first year in which the club was running, it was felt that there was also an opportunity to provide the accumulated feedback from the club members. Thus, this report represents, in the main, the views of the members of the PKI Club. The content of the report itself was built up through the use of a dialogue process, where club members were interviewed by academics from the Information Security Group. The transcripts were then analysed and the information presented in an aggregate and structured manner.

There were twelve interviews conducted which each, on average, lasted around two hours. Eighteen people were interviewed and, through their experiences, some seventy projects with PKI relevance were discussed.

Although the interview process was restricted to the club members, those we interviewed represent both users and implementers in diverse industrial sectors. Thus, while every relevant

[2] `http://www.isg.rhul.ac.uk/research/projects/pkiclub/`

party within industry might not be represented, we believe that the breadth of the views covered does not detract from the contribution of this document.

The use of a relatively broad and loosely structured discussion allowed those being interviewed to present what *they* saw as important. In this way, the issues covered were not tainted by the pre-conceptions of those conducting the interviews. Due to the confidential nature of the interviews, the interviewees were generally extremely frank and open in their communication. The transcripts were then subjected to filtering and anonymisation before selected attributable quotes were released for this report. However, the rigorous level of anonymisation and aggregation means that almost all points raised during the interview process have been preserved in this report. We believe that the nature of our interaction with the interviewees – which was constructed free of any commercial considerations and constraints – led to a very open atmosphere where heartfelt views were expressed which might have otherwise been suppressed. Also, the wide-ranging nature of the discussions coupled with the level of analysis carried out on the transcripts has resulted in this report providing an in-depth examination of the issues raised. As a result we believe this report and its contents represent a unique insight for those interested in the future of PKI.

1.3 Annotated Summary

In this section, we summarise the main body of this report. It provides an overview of what we believe to be the salient points touched upon in each of the subsequent sections. For reasons of brevity we have had to restrict the discussion in this section, and refer to reader to the relevant section for a more in-depth view on each of the issues raised here.

Section 2: Legal and Regulatory Considerations

Due to the fact that, in most cases, a PKI touches a large part of an organisation's infrastructure, the legislation that can affect it is not limited to that governing digital signatures. In this respect, the legal impact of PKI broadens out quite quickly. Areas which were discussed during our interview process were:

- Contract legislation
- Employment legislation
- Data Privacy legislation
- Regulation of Investigatory Powers Act
- Freedom of Information Act

A PKI is only of real use in support of some application or set of applications. As a result, the legislation which can affect specific applications can have a knock-on effect on how the PKI is designed and built. One banking application was implemented in a jurisdiction in which divulging customer names was a criminal offence. This had a clear impact on a technology that primarily uses *identity certificates* and the PKI thus had to be built using anonymous certificates. These secondary effects felt by such legislation should be considered carefully.

When we reviewed the legislation drawn up to support the use of digital signatures, many cited as a problem the continuing uncertainty surrounding the implementation of the legislation. Another important issue was the notion of control over the private key – as discussed in the EU Directive on Electronic Signatures – and how to ensure that an implementation conformed to this part of the legislation. Such uncertainty has forced many of those implementing PKIs to rely upon contract law rather than on statutory law when ensuring that the signatures will be honoured. Currently, asymmetric cryptography is the only technology capable of meeting the criteria set down in the EU Directive outlining *Advanced Electronic Signatures*. Thus, it strikes us that the problems generated by the legal uncertainties are specific to those within the PKI industry. Subsequently, it is in the vested interest of those in the PKI industry to see that these issues are resolved.

Compliance with regulation was seen by many as an important driving factor. In reviewing the regulation used within the PKI industry, we considered the broadest possible scope. Thus, for example, we considered regulation which was relevant to schemes (such as BACSTEL-IP and IBDE-2), sectors (such as UK Government regulation affecting the NHS) and different industries. In discussion, the use of regulation was seen as an aid to developing realistic policies and procedures for implementing PKIs. Identrus was isolated as a useful tool to follow for those building an implementation. However, a few people did feel that regulation cannot exclusively be relied upon to create a market for PKI to work within.

Liability was seen as a concern for those implementing a PKI. Uncertainty as to where the liability lies has been responsible for driving up the cost in many implementations. Finding ways of limiting the liability was seen by some as a large part of the design and process cost.

Accreditation – such as that provided by *tScheme* – was seen as important, but also too limited in its exposure at present. Although much of what is being done in this sector is tied to accreditation for regulation, it was felt that this could still be improved. The awareness and appreciation of brands that provide accreditation could also be improved.

In closing our review of legislation that impacts upon PKIs, we highlight two points which came through clearly during our research:

- Writing a CP and CPS is inevitably going to drive up the cost of the project. A few felt that their use in some projects had been prohibitively expensive.
- If there is a need to get lawyers involved, then it is important to do this as early as possible. The learning curve for those interpreting the legislation in light of the technology was seen as quite steep. Failing to take this into account early on could have serious impact on the project later on. Such impact could increase the cost or time to deployment and could result in a significant portion of the infrastructure needing to be redesigned during the project.

Section 3: Technical Considerations of Asymmetric Cryptography

When considering the technical aspects of PKIs, we have separated the discussion of the core technology – public key cryptography – from the means of deploying the cryptography – the surrounding infrastructure. In this section we review the issues relating to cryptography.

In discussing the use of asymmetric cryptography for providing security services – such as confidentiality, integrity and authentication – the fact that a PKI can deliver multiple security services was seen as a benefit. This breadth of services, in combination with PKI's ability to provide non-repudiation, is a clear advantage in some scenarios. However, there is often a discrepancy between what the technology was originally designed to provide and what businesses ended up using it for. A clear example of this was provided in one discussion where the interviewee noted that, although Identrus had been designed to implement non-repudiation, it was primarily being used to provide authentication. In fact, asymmetric cryptography's ability to provide a strong means of authenticating an individual – when compared to weaker technologies, such as passwords – is seen as its main benefit. The order of importance of the services provided by PKI emerges as: authentication, integrity via signature, non-repudiation and encryption. For the first three of these services listed, here are the main issues which were raised:

Authentication : Authentication was considered by many to be an enabler for other forms of technology. For example, one implementation has used authentication to allow clients to migrate from leased lines to IP-based access.

Integrity : Asymmetric signature's ability to provide *opaque* forms of integrity checks, where the document can be signed by multiple parties, is used in some fielded systems.

Non-repudiation : There is still interest in the provision of non-repudiation, but the cost and legal uncertainty are seen as prohibitive in many cases.

When we were discussing the management of private keys, there were two interesting points made when asymmetric cryptography was being compared to symmetric cryptography:

- A few people noted that in a wide-scale, low risk deployment, PKI provided a benefit of easier key management than in the symmetric case.
- PKI is more easily able to provide a *one key* view of an individual, but in some cases this could cause its own problems. One example of this was where using the same key and certificate across multiple layers with different risk profiles increased the complexity of the design.

In terms of key generation, the consideration of whether the system was going to be used to provide non-repudiation increased the effort required to ensure that the generation was carried out securely. The decision on whether to store the key in hardware or software is also a contributory factor to the design of the key generation process.

The means by which the key is distributed has seen improvement through better use of directory integration. Another means of easing the distribution is to use existing organisational channels, such as the procedures for handing out physical access cards.

Probably the most important issue in relation to key management is whether the private key should be stored in software or hardware. The use of hardware is the most popular choice, although in some circumstances, software storage is seen as an acceptable compromise. Regarding software storage, the secure administration of the machines on which the keys are stored becomes increasingly important. Even if hardware is the preferred choice, care still needs to be taken with its design. Inadequate procedures and an inability to securely recover from a lost token can leave the system vulnerable to an attacker. As well as accessibility concerns regarding the key, the following quote clearly indicates the benefit of having a physical storage device:

"Make sure you've got something that you know you've lost when you've lost it"

When a project required the use of a security technology, the view of most of those we interviewed was that PKI compared favourably with other strong security mechanisms such as symmetric cryptography. The cost, as well as the relative merits of each technology within the particular application scenario then drove the implementation decision.

Section 4: Technical Considerations of the Infrastructure

We now review the issues raised in our discussion of the infrastructure that supports the cryptography.

The fact that PKI has been largely developed in isolation from specific applications is seen as a key reason why integration has been a problem. The fact that early implementations were built by vendors with no specific application or sets of regulations in mind meant that it was difficult to get the vendor's products to inter-operate. However, although this has been an issue for some deployments, many felt that within *closed* application environments, it was less of a concern.

Integration with specific vendor applications is considered more important than compliance with individual standards. Although supporting legacy applications still causes some problems, the use of middle-ware tools has reduced the potential headache to a manageable level.

When deciding where to position the integration with the PKI, one interviewee noted that, if the PKI was being used to support multiple applications, then it should be mediated through some form of middle-ware component. Only if the PKI was being used specifically for one application did they consider it wise to interface directly to the PKI.

Integration with Privilege Management Infrastructures (PMIs) is considered to be another important issue. A few people noted that PMIs have the ability to provide a more tangible

business benefit in the form of long term savings through cost reduction. PKI is then viewed as one of the means of providing authentication for a PMI. In this case, integration with PMIs is likely to become an important issue for PKI vendors. A related issue is the fact that many people saw the business decisions needed to support a PMI difficult to outsource. This could potentially have an impact on those companies offering PKI as a hosted service, as integration across an organisational boundary is generally more difficult.

Integration with other infrastructure components such as directories is also considered important. In some cases it causes problems equal to the more specific forms of integration described above. In various projects it has caused issues with: migration between PKI vendor components; support of revocation protocols; use of different communication standards between the various infrastructure components. These problems were often difficult to identify without prior experience. Dealing with these problems is seen by some as one of the biggest technical challenges faced by those deploying a PKI.

Regarding the impact of interoperability between different vendors, it was clear that many believe that it is still not as good as it should be. Interestingly, there are occasional additional sets of problems where more than one vendor is being used within the same hierarchy.

Some people have used standards as a means of getting away from the problems of interoperability, although others feel that this is not easily achieved. It is a commonly held view that, while standards are a benefit, they do not provide everything and are often used as a template for education rather than as a strict set of guidelines. One of the reasons for this is that they are often compromised in order to get things to work in practice. In particular, X.509v3 was criticised for trying to achieve too much.

When choosing a PKI vendor, most people have their own criteria based on the price and package offered. What does stand out is that different companies wanted different levels of flexibility in the products offered by the vendors. For example Baltimore was selected where the technologists wanted to modify the infrastructure to suit their needs, whereas Entrust was chosen primarily by those wanting to deploy a simple, rigid hierarchy.

Due to PKI's high cost, re-use of a deployed infrastructure is something that is often desired. This has happened in different ways. Some identify this up front with their deployment, whereas others see it as a by-product of building a PKI for some initial application. One of the reasons for re-using and infrastructure is to centralise the control of various PKI projects in a large organisation. Interestingly, re-use is not always possible, with some implementations being built as a *black box*, whereas others fall foul of scheme or legal requirements.

When designing and building an infrastructure, consideration needs to be given to whether the application, or the client using the application, is internal or external. PKI is seen as being able to provide an easier security design for external applications or clients. However, it is important to acknowledge that having the application running externally, or the client situated externally, can impact key management, as well as the policies and procedures used. In general, applications which are run externally, or have external clients, rely on more formal policies and procedures. As a result of this, migrating a PKI from an internal to external application, without careful consideration in advance, is likely to be expensive and difficult.

Many felt that the issues faced in terms of the time and management of a PKI project are similar to those faced in other large technology deployments. The long lead-time from project inception to a working PKI has been an issue for many. The opportunity to reduce this lead time is seen as a major benefit when considering the use of a hosted service.

When a security service is being migrated either from an old PKI to a new PKI, or from some other security technology to a PKI, knock-on effects are felt. Care needs to be taken when

considering any modifications that were carried out to the original infrastructure. For example, the objective of one project was to upgrade the PKI vendor technology. The upgrade process had to take into account changes which had been made to the original product, where these changes were not supported in the conventional upgrade package. Also, there can be other unintended upgrades as the new requirements are put in place. A good example of this is the new BACSTEL-IP application which requires a minimum of a Windows 2000 installation at the client site.

When setting up a new PKI, some thought it could be of benefit to make use of existing systems (such as *securID*) or software based keys as part of a phased migration over to the full PKI implementation.

In general, complex hierarchies, such as bridge CAs and cross-certification between CAs is not something that has materialised, even if it has been discussed at great length.

The fact that a PKI can touch a large part of an organisation's infrastructure can add a great deal of complexity to its management. Due to this breadth of impact, many feel that centralising control, if possible, is one way of dealing with this potential problem.

Another important issue that was raised is the need to ensure that the technological design reflects the risks inherent in the application. For example, software stored keys will be unacceptable if the PKI is put in place to support high value transactions. Core to managing the technology is realising that the management of the keys and certificates is non-trivial and has a large impact on the design of all aspects of the PKI.

In conclusion, we note that technological stability and maturity is still a problem, but that it is likely to improve through the development of specific schemes and applications.

Section 5: Management and Administration of the Infrastructure

In this section we consider the management and administration of the infrastructure. These subjects are of relevance to the procedural element of running a PKI.

Registration is seen by many to be a key problem faced by those managing a PKI. Identifying an individual reliably is imperative if certificates are to be based on an identity. In order to achieve reliable identification, most PKIs use the departments within an organisation as the registration points. Registration of an individual not associated directly with an organisation is seen as more difficult. This leads us to believe that PKIs which are outside a tightly constrained contractual environment are unlikely to be a success. Many noted that following industry or scheme regulations has been an advantage when building a registration process.

Management of the key and certificate life-cycles is also seen as an important factor in the running of a PKI. The prevailing view is that building successful life-cycles is a balance between the cost, technology and risk associated with the use of the keys. A few people noted that it is important to remember that there was often reliance placed on non-IT departments within an organisation in the processes implementing the life-cycles. For example, the processes might rely on information from the Human Resources division relating to changes in personnel. It is also important to make sure that the re-issuance or migration of certificates does not create undue burden on the system. In one example, the initial certificate expiry dates were all synchronised. As a result, the process had to be modified to re-issue the certificates before the expiry date in order to avoid a bottleneck during re-issuance.

Due to its high profile, we have isolated our discussion of revocation in a separate subsection. In discussing whether to use Certificate Revocation Lists or On-Line Certificate Status Protocol, there are generally three separate views:

- Using a CRL is considered unacceptable. If an organisation is going to the effort of imple-

menting a PKI, the fact that a CRL may be provide out of date information means they might as well have picked a cheaper alternative.

- The decision on whether to use a CRL or OCSP must be tied to the risk associated with the transaction. Thus, use a CRL for a low value transaction and OCSP for a high value transaction.

- The decision on whether to use a CRL or OCSP must be tied to the time criticality of the action being supported. In some cases the PKI only provides the front end authentication and there are separate mechanisms to lock people out of the authorisation modules on individual applications. In such cases where there are other means of revoking the right to carry out any time critical actions, CRLs are deemed sufficient for the authentication process.

This discussion also covers the issue that certificate revocation is only part of a broader security infrastructure. Here are two examples of how this expresses itself in practice:

- In a few implementations, whenever a revocation request was issued, the certificate would initially be suspended, with the revocation happening later on. Two reasons were highlighted for this: the revocation request could be verified by a more stringent process; there could be valid but as yet incomplete processes that rely on the certificate which might have been started before the revocation took place.

- In the case where the private keys were held on physical devices and piggybacked on existing mechanisms, such as entry cards, the revocation of the certificate was sometimes handled by the revocation mechanism for the existing card usage.

Identifying who is responsible, and has the authority, for issuing the revocation request is an important part of building a reliable revocation mechanism.

If we consider the management of the system security from a commercial perspective, then it is important to realise that the PKI policies and procedures exist in the context of a larger corporate security policy. This can be an important consideration when analysing how the corporate security policy might influence the PKI security policy and vice versa. Also, there are often other technical means of delivering the service we want from the PKI.

Even when PKIs are put in place to replace existing physical procedures, it is important that the procedures surrounding the PKI are still carefully managed. The fact that PKIs are sometimes envisaged as a means of getting rid of security management is a fallacy that was highlighted by a few of those we interviewed.

The logical and physical management of the infrastructure can, when necessary, be separated to aid those with different and possibly competing interests. A good example of this is the Identrus infrastructure which can be split into three parts: certificate management; user registration; application. These can then be run by separate entities. When it comes to using a hosted service, having the controls split between service provider and client is deemed standard practice. To finesse some of the concerns of the management of the infrastructure in a hosted environment, the concept of *in-sourcing* was discussed. The notion here is that a vendor would build and support the technology, but it would be physically located at the client site, thus allowing the client to progress to fully operating the technology if the need arose.

Maintaining the security of the Root CA private key in any infrastructure is important. Managing this key has been attributed as being partly responsible for the rising cost of running a

PKI. In a case where a hosted service is used to provide the Root CA, the client would normally be present at the key ceremony to generate the Root key pair. The control of this key pair could then be shared using commonly available technology. When we asked what would happen should the Root CA private key be compromised, very few people could provide a definitive answer.

Certificate Policies (CP) and Certification Practice Statements (CPS) provided a large amount of discussion during the interview process. Clarifying the content of each provided problems for some of those we spoke to. Many saw the confusion surrounding them as one of the biggest problems they faced. For completeness, we provide one simple definition given to us:

CP : The CP outlines what we want to achieve.

CPS : The CPS describes how we achieve it.

Most people thought that CPs and CPSs only really added any value outside of a *closed* infrastructure, for example, where parts of the infrastructure would reside in different organisations, or when multiple organisations want separate infrastructures to inter-operate. The *de facto* reference for building a CP and CPS – Internet RFC 2527 – was deemed by many to be insufficient for their requirements when they come to writing their own. If a CP and CPS are required within a project, then it pushes up the cost considerably, but if the PKI is operating within a framework which provides some level of existing regulation, then this helps to balance out that cost. One interviewee noted that, in a *contractual* infrastructure, they feel that using a CP and CPS is vitally important for ensuring the legal responsibilities of each party are more easily identified.

In closing the section on management, we note that there are a few comments that allude to how the design of the infrastructure, in terms of the logical architecture, should reflect the relationships between the communicating parties. An example of this is provided by comparing the generic Identrus model to the BACSTEL-IP model. The most commonly associated model for Identrus is a 4-corner model which is not designed to support any specific application. The BACS model is a 3-corner model which is designed around a particular application. It is felt that, by concentrating on the relationships reflected within a particular application scenario, the 3-corner model used by BACS has provided a design that was easier to implement and manage.

Section 6: Commercial Considerations

In this section we consider issues such as the cost, business requirements, etc, which are of commercial importance in any PKI project.

Many different business requirements were cited by those we spoke to, but in the main they could be categorised according to one of the following three reasons:

Cost : In some projects PKI was seen as a means of enabling an application which could provide an overall cost reduction to the organisation.

Security : PKI was seen as a means of providing "strong" security measures.

Regulation : In a few cases, new regulation mandated the use of PKI.

A number of people identified different ways in which it was possible to ensure that the implementation of a PKI had a sound business driver. The two clearest ways identified were: ensuring the involvement of a business manager relevant to the target area early on in the project

and maintaining close contact with them through the design process; designing the PKI with a particular application environment in mind. In following the second of these strategies, we believe that it is important to remember the separation between the security services provided by the PKI from the application that uses that service. It appears to us that many people who discuss the use of PKIs talk as though the security service provided by the PKI is directly attributable to the application. While it is important to isolate the security service from the application, it is also important to understand the impact the security service has on the application.

In closing our discussion on business drivers, we believe that too many PKIs are implemented speculatively. This has resulted in a large proportion of the projects failing to provide a genuine business benefit for the organisation.

Next we discuss the cost of ownership of a PKI. The cost estimates we were provided with ranged from $1M to £25M. The main reasons put forward for such a large cost were:

- The cost of building a reliable CA has been pushed up because of the security of the CA being crucial to the security of the system as a whole.
- The use of CPs and CPSs were isolated by many as a reason why costs have increased in many projects.
- The technical expertise required to implement a PKI is not available in many organisations. The cost of acquiring that competency is then attributed to the project.
- The impact on the surrounding infrastructure means the costs broaden out, and hence accumulate, quite quickly.

Many people feel that, although generally quite high, the cost of deploying PKI could be justified when compared to other technologies which provide a similar level of security. However, in certain cases where the use of asymmetric cryptography is clearly justified, in terms of the services provided, the cost has prohibited the use of a PKI. As a related issue, the provision of non-repudiation using a PKI is one of the factors that pushed up the cost dramatically. This was mainly attributable to the heightened expectation from what PKI could deliver coupled with the legal uncertainty surrounding the deployment of digital signatures.

In aiming to reduce costs, three possibilities stood out: building a shared infrastructure between several organisations; infrastructure re-use within an organisation; use of an outsourced provider to host the PKI.

When we questioned people about the use of a return on investment calculation against the PKI, most note that it was rarely seen within PKI projects.

One of the main business decisions facing those considering implementing a PKI is whether to build the infrastructure in-house, or use a hosted service. Here were some of the issues that were considered a benefit for each:

Hosted : A hosted service can provide for improved scaling in terms of recovery and fault tolerance. Once a hosted service has achieved regulatory compliance for a given scheme (e.g. Identrus), it makes it easier to implement new infrastructures within the same scheme on subsequent occasions. For the client of a hosted service, there are clear cost benefits, with one organisation we talked to having seen an order of magnitude drop in cost after moving to a hosted service. The expertise needed to complete some of the specific tasks within a PKI is generally accumulated by those in a hosted service. There is generally an improved roll-out time for the project because of the expertise and scaling issues.

In-House : For organisations looking to implement a PKI to secure high value transactions, the issue of retaining control over the whole infrastructure could influence the decision. PKIs are often used to support applications and PMIs, but those secondary services are generally not as easily supported in a hosted service. This means that there are increased integration and business process issues to tackle if the PKI is off-site while the PMI is on-site.

When choosing a hosted service, many different criteria were put forward. Here are some of those discussed: the PKI vendors supported by the hosted service; possible migration away from the hosted service in the future; Service Level Agreements, especially in support of revocation; the scheme regulation the hosted service had already complied with; the flexibility and scalability of the service on offer, allowing the client to modify their PKI investment as the project grew.

The issue of vendor lock-in is not seen as a great concern for many.

When it came to discussing the competition between vendors in the PKI marketplace, the following quote highlights the fact that it is still quite a congested market:

"There are clearly a number of organisations trying to sell into this confined space."

Although PKIs are more likely to be of use within large organisations, their deployment within those types of organisation is likely to bring its own set of problems:

- Where multiple business units are going to benefit from its implementation, it becomes difficult to isolate who precisely is going to fund and manage the infrastructure.
- Educating staff in diverse parts of the organisation is not easy within a large organisation. Due to the way a PKI may touch different parts of an organisation, it is important that people are aware of its impact.
- In the context of an organisation formed from a diffuse group of companies, the sector responsible for the group IT function can sometimes lack the authority necessary to direct the project within the separate companies. This can lead to greater interoperability problems.

Section 7: The Human in the Equation

Most definitions of what a PKI is will acknowledge the role of the people and processes within the infrastructure. In this section, we isolated the issues of relevance to the impact of the human understanding of the technology.

One of the issues raised quite clearly within our research was the conceptual difficulty posed by the technology.

Many people noted that it is important that those managing the business aspects of the organisation understand the impact of the technology. To achieve this, there needed to be effective education. In doing so, it was seen as important that people at all levels of the organisation were included in this process because of the wide-ranging effect of the infrastructure. Some people saw that, in a few market areas, there was still too much hype surrounding non-repudiation compared to the reality of how difficult it is to implement. Criticism is also levelled at the expectations some people were placing on software key storage, which reduces the benefits gained from using asymmetric cryptography. In this case, there is a requirement for increased understanding of the drawbacks and benefits of the various technological choices.

Even with technical staff the general level of awareness and knowledge is considered by many to be insufficient. Asymmetric cryptography is quite a difficult technology to understand fully,

especially when we consider the nuances of its practical implementation. Consequently, very few organisations are equipped with the skilled individuals to implement a PKI. This is highlighted by the following quote from one of our interviewees:

> *"If you look at where the money gets spent [. . .] it's all the consultancy you have to put around it."*

Thus, an increase in awareness of the implementation issues among technical staff would appear to be a priority within the industry.

The end-user's ability to use the technology, along with their acceptance of it, is considered an important issue. In terms of ease-of-use, trying to get the user interface to mimic existing technology of a similar type is seen as one way of improving usability. It is considered important that the user is not hampered when using the security, otherwise they might find ways around it. As well as reducing security, unusable security can reduce usage and uptake of new applications. In some scenarios this has led to new applications being dropped after that had been implemented because the security was too unwieldy.

Many felt that early PKI products had suffered in their usability by the fact that the technologists designing the security had also designed the user interface. In this light it is important to understand that user interface design is a very different discipline to cryptography.

When it comes to explaining the importance of the security, there is a fine line to be drawn between not saying enough and saying too much.

A few people note that a part of the problem faced when trying to design the process from an end-user's perspective is that very few users are used to managing an identity securely, much less a digital identity. Tackling this problem is more of a social issue, but many felt that schemes such as *"chip and PIN"* would help promote these issues in the real world.

1.4 Personal Contribution

In this section, we introduce our own thoughts which have been formed over the course of this research project. Some of what we say here mirrors what is said in the previous section and some points are purely our own. Our aim in adding this section is to draw an overall picture of where we believe the main issues lie.

Do we need PKI?

In answering this question we consider it from two points of view:

- **Is there a market need for PKI?** We believe there is, and with systems such as BACSTEL-IP and CHAPS now mandating the use of PKI technology, it has a firm, if limited, foothold within the industry. Also, asymmetric cryptography offers a small but important set of services which symmetric cryptography cannot provide.
- **Can we manage public/private key-pairs without a PKI?** We believe that this is not possible in the general case. The way in which we manage the keys, both technologically and procedurally, might change, but to think we can remove the need to manage them at all would overturn all conventional security practice. It might be the case that in the future such infrastructures might look different from the ones we have today, but the infrastructure itself instantiates that management process and hence cannot be done without.

Where do we need PKI?

Once we accept that there is a need for PKI, the natural extension to the first question is: Where do we need to implement a PKI? As we discuss in many areas of this report, a PKI is seen by many as a supporting mechanism and rarely the end goal itself [3].

The main goals that we have seen supported are: enabling other business solutions; providing Privilege Management Infrastructures with strong authentication; meeting regulatory requirements. We believe that, in order to provide better support for these, and potentially other services, the how and why of the way in which PKI supports them needs greater attention.

Even though PKI, as a technology, can provide a broad range of services, we must remember that the applications which are going to use these services must have a requirement for them. When PKI entered the marketplace, it was sold on the basis of the technology. As it became apparent that this did not result in effective use of the technology, those in the industry focused on selling the services provided by the technology (e.g. authentication, signatures, etc.). While this move is to be lauded, we believe those selling the technology need to go one step further and promote those services in terms of the types of application each service best supports. For example, one clear use of authentication is in support of a PMI. There needs to be more thorough investigation into the classification of application level services which can be supported.

We believe that, in providing the infrastructure in support of the technology, there has been too much focus to date on the subscriber-CA and CA-CA relationships and not enough on the subscriber-relying party and relying party-CA relationships. Such issues are being brought to light with the implementation of more 3-party models such as BACSTEL-IP, but they still require further development.

[3] Although, during the "hype phase" in the late 1990s, the opposite was often true.

We believe that pilot applications provide little benefit. Many of the management and procedural issues do not emerge until we move from a small PKI used for exploring the application possibilities to a full scale deployment.

Although the issues faced by those supporting a PKI are many and varied, none are insurmountable. In the cases where PKI has been deemed the technology of choice, it has been possible to resolve the problems. These successful cases need to be analysed to provide guidance on how these process can be made easier and cheaper.

Fit-for-use considerations

If we accept that PKIs are mainly required for the support of other processes, then this forces us to tailor each implementation in support of these processes.

We believe that the notion that it is feasible to have a broad multi-use certificate is misleading. While many in the industry have moved away from the idea, we believe that this needs to be refuted for once and for all. A certificate is a representation of some state that is of relevance to the security of an application. We believe that such a state is not an easily transferable commodity and is by necessity of relevance to a limited range of applications or environments.

What we need is a better understanding of how certificates and the means of managing them, such as the CP and CPS, can be tailored for similar applications. We also need greater understanding of what effects the changes to these processes have upon the application. Currently, there appears to be little in the way of industry best practice on certificate management. This highlights the relative lack of maturity within the PKI industry.

Currently, the security engineering of certificate management is not well understood.

How PKI fits into the broader security infrastructure

We need to appreciate how the PKI fits into the broader security policies of the business process and surrounding technological infrastructure.

The policies that are in force within a PKI are generally quite restrictive. One interviewee noted that an internal client requested that those procedures were changed. The interviewee had to point out to the client that this was as a result of industry regulation which the PKI had to comply with.

Care must be taken when calculating how policies and procedures of the PKI interact with the policies and procedures internal to the PKI. In the case of PKIs which are used to support applications which are delivered to external clients and users, it is much more difficult to manage the policies and procedures. This has a knock-on effect upon how secure the resulting use of the private key can be considered to be. A few implementations that were discussed during the course of the interviews have services supported by externally managed keys limited to a more benign set of processes.

It is crucial to remember that non-repudiation, which is one of PKI's few trump cards, is about much more than the technology. The processes and procedures need to be well implemented and the legal issues are still uncertain.

Those proposing an implementation of PKI within an organisation need to provide multiple cases together to clearly demonstrate its impact. The technical, business, regulatory and procedural elements need to be stressed individually and collectively to demonstrate the potential benefits, along with a sound business case.

Hosted Services

We believe that hosted services have the ability to provide a leading role in expanding the use of PKIs. The two main benefits demonstrated in practice are the cost savings to clients and the provision of the required expertise. It is our belief that these are two of the greatest problems faced by any organisation wishing to implement a PKI. The cost is undeniably large and any significant saving can increase the likelihood of a project providing a financial business benefit. The importance of accumulated expertise cannot be overlooked when implementing a PKI. Many of the problems encountered when implementing PKIs were difficult to deal with when encountered on the first occasion, but were common to many implementations. Given that most organisation will only attempt to implement a PKI once, this is likely to increase the difficulty of implementing the PKI. If the organisation does not have experienced engineers in-house, procuring the expert knowledge externally will leave them with a large consultancy bill.

The one question we see becoming important for providers of PKI hosted services is: How can one support a business process, when that process itself is often difficult to host? A clear example of this is the provision of PMI which, because of the business related decisions that authorisation are based upon, are not easily managed by a third party. The security process and technological integration issues faced when a PKI is supporting a service across an organisational boundary are likely to become more important.

Standards, Schemes and Regulations

When discussing standards and regulation, we are using the terms in the broadest possible sense. For example, we include schemes such as Identrus which can provide the focus on implementation issues which standards and regulation should provide.

We believe that both standards and regulation can provide a useful lead in implementing PKIs. This primarily happens in one of two ways: they provide the educational input required to help ease deployment; they provide the means by which integration can be built around a more rigid set of constraints, aiding the application developers.

There are instances where standards and regulation provide a clear focus for a given business application, which leads to a concrete market opportunity.

Providing the application developers with something concrete to support aids in developing the procedures and processes. This then makes deployment easier. One comment made to us during our research was that standards are often compromised in practice to get a system to work. We believe that this is a result of standards being far too general and not being developed to support specific scenarios. As a result, we believe that standards need to be more targeted in order to provide more benefit.

Liability

Many of those we spoke to voiced concern over the issue of liability within PKIs. It would appear to us that well-defined schemes have the potential to allow for greater control over such liabilities. By signing up to a specific set of membership or scheme rules, all parties should be able to make use of a PKI with a clearer understanding of the risks they are undertaking. To us, this would appear to imply that justifying the cost and investment in an *open* PKI (i.e. a PKI with little or no contractual framework) is going to be difficult if the inherent risks of liability are to be factored in.

Cost

Although the cost for deploying a PKI is large, most people we spoke to conceded that, in a lot of cases, it compares favourably with other offerings, such as symmetric cryptography, while providing a similar level of security.

The question then becomes: How much of a requirement is the additional security? It would appear to us that the ability to tie the cost to both business benefit and associated risk is still in its infancy. We believe this to be true of the risk analysis aspect, having attended several presentations where conventional risk management was deemed inadequate for IT security.

Leverage from Infrastructure re-use

Infrastructure re-use is desirable for many of those who implement a PKI. This is primarily down to cost, but can also be driven by the nature of a PKI and the way in which it interacts with the remainder of the organisational infrastructure.

Unfortunately, the re-use of an infrastructure does not come cheaply and by and large needs to be considered from the start of the design phase. We saw that PKIs which had initially been built to support an internal applications rarely translated well to support external applications. Specifically, the policies and procedures implemented are difficult to migrate from one instance and the other. A clear example of how this can happen is demonstrated by the fact that CPs and CPSs are regularly dropped for an internal application, whereas they provide greatest value when supporting an external application.

Designing this additional capability from the start can be desirable, but it does, by necessity, push up the cost. It also means that the design process needs to take into consideration a separate set of requirement and constraints. This can greatly complicate the security policies and procedures.

Too accessible and over-generalised

We believe that one of the reasons why PKI has suffered a tarnished image is that the technology can be too easily discussed, with many people learning a very small amount about it.

We believe that it is deceptively simple to talk about *keys* and *signatures*. The technological discussion can then be overburdened by the expectation and understanding of their physical counterparts. Also, with there being a broad range of technical issues to consider, the combined consequences of which means it can be extremely difficult to see the whole picture. This can result in a strong reaction against the technology when people realise that it is difficult to implement or does not provide what they originally thought it could.

Taking note of how this misunderstanding can cause things to go wrong is an important issue. Designing a unilateral security service (e.g. where only one party in a communication is strongly authenticated) through the support of asymmetric cryptography is one of the strengths of using a PKI. However, it does open up a variety of ways in which the security service can fail. Due to the inherently one-sided nature of asymmetric security services, detecting this failure can be very difficult.

We believe that another reason why PKI has suffered is because of the breadth of services it can provide. Although this is seen as one of its strengths, it is important to remember that being able to provide such breadth comes at a price. Not all PKIs should, or can, provide all possible services. Moreover, the requirements for each service, along with the impact on the running of the PKI, are going to vary dramatically.

CP and CPS

We have heard many people complain about the difficulty of composing a CP and CPS. Although, in principle, we agree with much of what has been said, we believe that, in blaming the CP and CPS, they are missing the point slightly. CPs and CPSs are just the PKI instantiation of policy and management aspects of a security infrastructure. These are notoriously difficult things to get right. What needs to happen is for there to be an improved understanding of how CPs and CPSs fit into the broader challenge of managing information security. This should be coupled with the evolution of more specific classes and examples of CPs and CPSs to aid those using them in their system design.

There is an analogy here with the process of designing secure protocols. A colleague once told us of an example where he was asked: *"Is my protocol secure?"*, to which his response was: *"Secure for what?"*. This might sound trivial, but until there is a better understanding of specific CPs and CPSs and how they are used in varying real world examples, we will struggle to provide adequate support for those attempting to draw up new ones.

Business Victim of its Technological Success

In closing our personal reflections on this project, we draw one overall conclusion that is purely our view.

Public key cryptography offers a set of security services which were unavailable prior to its development. From the outset, this has made it an attractive target for commercial development. This is true, in part, because its use offers the biggest potential benefits in environments where traditional symmetric key technology is difficult to deploy. However, irrespective of the choice of technology used to support security services, the design and maintenance of a secure system is an extremely difficult task.

Managing a CA – and hence a PKI – is a technologically demanding task. As we have seen throughout the report, the benefits offered by public key cryptography (most notably non-repudiation) has, directly or indirectly, raised these standards even further. In addition, the environments in which public key cryptography can offer the biggest advantage are those where strong pre-existing commercial relationships are unlikely to be in place. We believe that a combination of these two facts make it difficult to justify the additional expense in the risk-averse commercial world.

It is our opinion that this has resulted in a technology which has the ability to provide a broad range of security services being hamstrung by the technical and engineering cost of doing so.

In isolation, the technology itself compares favourably with the competition. However, the engineering and business considerations experienced in the environments in which it has an edge over the competition has been its drawback. This has resulted in it being costly and difficult to implement when compared to its potential benefits. We belive that this difficulty was not adequately taken into account in the early years of PKI deployment.

1.5 General Conclusions

In this section, we aim to draw together what we believe are the key points raised in this report.

The Need To Target Business Processes

During our work, when we asked where PKIs have been used, one of the most common responses was that it was used to protect high value transactions. We agree that the additional security achieved and cost incurred from deploying a PKI is best placed for supporting such transactions. It is also our view that considering the value of the transaction alone does not provide us with sufficient information when considering what types of applications PKI might best support in the future. We believe that there needs to be a more complete understanding of the nature of the transactions that PKI supports best: What are the types of communicating party relationships? What are the process flow models? What type of contractual environments provide the best scope for addressing the process issues? Answering these, and other similar questions, will provide the industry with a clearer understanding of where PKIs can best be used in the future.

CP and CPS

The difficulty of building CPs and CPSs were mentioned by many. We believe that the best way forward is for better examples to be built around business areas or types of process models. While RFC 2527 gives a starting guide for those looking to develop their CPs and CPSs, it is too broad to be useful for those taking on this specific task for the first time. With more rigorous templates to work from, it should ease the pain of what is an important part of building a PKI.

Lightweight Infrastructures

Although most PKIs, for the foreseeable future, will continue to be large and costly implementations, we also are aware of some scenarios where it is used for small internal applications with the minimum of fuss. Understanding what types of scenarios a lighter form of infrastructure is applicable to is an important step for improving the opportunities for PKI vendors.

Security Service support for Applications Services

We believe that the most important step forward is to enhance our knowledge of the how the security services provided by a PKI are best used. In the list below we cover the three main services used and some initial ways in which their use can be promoted. Expanding on these must be a priority for those looking to expand the use of PKIs:

Authentication : The ability to provide strong authentication is seen as the main benefit to be gained from using a PKI. Supporting other authorisation mechanisms is the natural place to use such authentication. There needs to be greater promotion and understanding of what types of authorisation processes are best suited to asymmetric cryptography.

Integrity : The ability for digital signatures to provide an integrity mechanism is often overlooked. A few of the implementations that we studied did use the strength of this opaque integrity mechanism (where many signatures can be concurrently added to the same document). Where else might they be deployed?

Non-repudiation : The potential to provide non-repudiation is still a desirable facet of the technology. Unfortunately, the uncertainty surrounding the legal issues and the understanding required to build sufficient technical support is hampering its development as a service. These issues need to be ironed out if PKI is to be widely used in practice.