Global Information Assurance Certification Paper

The Promise of PKI
Kerry Bryan
February 17, 2002

Abstract

This paper assumes an understanding of *public key cryptography* (PKC) and will only provide a high-level overview, focusing mainly on the issues surrounding PKI and why, after years of debate and research, it continues to be anticipated. For a more detailed discussion of the technology of PKC there are numerous papers, articles and resources available on the Web.

PKI is a concept that has been bandied about since the late 1990s but has not yet fully taken hold on a large scale. Public-Key Infrastructure (PKI) is purported to be the cornerstone to secure e-commerce across the Internet. The promise of PKI is that in building a third-party "trust" system into the arena of electronic transactions businesses, as well as individuals and government, can securely interact across the Internet. Smaller, contained PKI installations, such as those within private corporate, government or education networks, are beginning to slowly appear. But, the "public" key infrastructure and the expansion of the use of external third-party CAs and digital certificates has not yet evolved to the degree predicted, why is that?

Introduction

What is PKI?

The core technology of PKI is *public key cryptography* (PKC). *PKC* is an asymmetric key-pair cryptography scheme developed by Whitfield Diffie and Martin Hellman in 1976. With PKC each entity has two keys, a private key and then a public key, which is derived from the private key to form a *key-pair*. Using this method the public key can be published widely and used by anyone who needs to electronically communicate securely with the owner of the key-pair. Since the public key was derived from the private key, any message encrypted with the associated public key can be decrypted with the private key of the pair, and only by that key. PKC, implemented properly, provides for confidentiality, integrity, authentication and non-repudiation.

PKC can, and is, used independent of any formal infrastructure. However, issues arise when trying to utilize this technology on a large scale and/or in an environment where neither the sender nor the receiver knows each other. It becomes increasingly difficult to manage public keys and to verify the owners of keys, and therein the validity of the encrypted message, as the user population expands. What assurances do you have that the holder of the key-pair is indeed who or what it professes to be? Has the key-pair been compromised? Is the key-pair valid any longer?

For PKC to be truly useful what is needed is a means of managing public keys and many-to-many relationships between disparate unknown entities, while still providing assurance that electronic communications/transactions are secure. This is where PKI fits in. PKI is, as it's name states, an infrastructure. The fundamental objective of PKI is to introduce a means of forming a trust relationship between entities interacting in a purely electronic arena therein

establishing a foundation for the implementation of the security mechanisms provided by public key cryptography. The means for introducing this trust relationship is the use of third-party authentication processes and procedures provided by a *certificate authority* (CA). The CA's primary role is to provide for authentication and management of a *digital certificate*, much like a registrar or notary in the physical world.

A *digital certificate* is essentially a pre-defined set of identifying information about a person or entity that has been verified through the CA. Once verified the entity's public key is bound to the certificate and validated with the CA's own digital signature, thus providing assurance that the presenter of the digital certificate, and owner of the public key, is indeed who or what they claim to be. In essence the digital certificate is a digital identity, or electronic passport for use in electronic interaction.

PKI Basic Overview

Although all PKI is built around the ISO X.509 Standard, there is a lot of room for differentiation in the execution; not all PKI implementations are the same. The basic components of a PKI are a Certificate Authority, a directory service that provides for certificate storage and retrieval and client software, which may be proprietary or not. The Larger PKI installations may also include the added role of *Registration Authority* (RA). The RA's role is delegated from the CA, to actually verify the contents of a certificate for the CA, sometimes involving review of physical documents and/or presence of the actual person requesting the certificate – much like applying for a passport. A CA may have multiple trusted RAs.

The basic functions of PKI are:

- Certificate generation, which includes the creation of the digital certificates and *certificate revocation lists* (CRLs) within a defined format – based on ISO/IETF x.509v3 (detailed information about this standard can be found in the IETF RFC 2459).
- Certificate distribution, which involves the certificate repository, or directory. Here all public key certificates, CRLs and CA certificates would be stored. Because of the high availability requirement and normally distributed nature, the repository would most likely be LDAP compatible.
- Certificate management, which includes the management of keys and certificates over time. More robust implementations of PKI, particularly those of external providers, would also include services such as key backup and recovery, support for non-repudiation, automatic update of key-pairs and certificates, historical key management and support for CA cross-certification.

Two key documents involved in the implementation of any PKI are the Certificate Policy (CP) and the Certificate Practice Statement (CPS). The CP is the high-level guide describing a set of rules for the qualification of a certificate to a class of applications with similar security requirements. For instance, guidelines regarding authentication of certificates that are to be used in the exchange of goods for monetary value. The CP would include statements describing the purpose of the PKI and specific business requirements the PKI addresses.

The CPS is a statement of practices and procedures. Essentially the CPS is the work instructions for the administration of digital certificates. It is in this document that you will find the

processes describing all facets of the lifecycle of a digital certificate, including certificate generation, authentication, storage, distribution and revocation procedures, among others. Here is the heart of PKI, where the trust is instilled that encourages users to participate in the model.

Obstacles

So, if PKI is truly the answer to secure electronic interaction, why then is it not more fully realized? There are many obstacles surrounding the widespread acceptance of the PKI model that must be addressed before PKI can be fully actualized; such as interoperability problems, process issues (authentication and revocation), privacy and legal concerns to name a few. Additionally, there are external factors that have an impact; such as economic/market conditions, consumer education and business management's understanding of the importance of information security.

Lets start with the foundation of the PKI model, trust, and more specifically trust in the CA. What does that mean? In order for the CA to gain and maintain this trust the user/subscriber has to believe that the CA is first, acting with due diligence and care, and second willing to backup their claims should something go awry. These concerns should be covered in the CP and CPS. These documents should explicitly spell out, among other things, procedures for adequately managing certificates end-to-end, roles and responsibilities as well as warranties and liabilities. Unfortunately these documents are extremely long, overly wordy, use obscure phraseology, are mired in legalese and all too often minimize the CA's liability. This mechanism for instilling trust has not to-date been overly embraced. Trust in the physical world is built around complex relationships over time and not easily or quickly mimicked in the electronic world. Incidents such as Verisign's issuance of two Microsoft code-signing certificates in error only proves that the system is not infallible and makes trust that much more difficult to attain. The question is do you trust that the CA is following secure and thorough practices and that they will maintain these practices over time?

What is actually contained in the digital certificate? And further, what is authenticated? Digital certificates are not necessarily readable, so the user may not even know what they are accepting. What the certificate does assure is that a message was generated by some entity that had access to a given private key at the time said message was created. Further, that the CA that issued the certificate had reason to believe that a particular entity was associated with this private key at the time the certificate was issued. A digital certificate does not provide assurance that the entity named by the certificate in fact originated the message. For PKI to be effective the following must be true:

- Assurance that the private key is actually in the possession of the appropriate party, the valid sender/originator, and no one else.
- There has been sufficient security measures undertaken to assure that no other party can ever obtain access to or use this private key.
- The private key used is the correct / valid one and not a fraud
- Processes exist to address the compromise / invalidation of a private key, ensuring revocation of the associated public key, all within a meaningful timeframe and with minimal impediments to enact.

Unfortunately, the existing PKI model provides no real mechanisms for the assuring the above requirements.

It is a given that mistakes and breaches will occur in the process, as has already been proven quite publicly by Verisign. Therefore, there needs to be a process to revoke compromised or invalid certificates. Given breadth of the Internet, there exists a real issue with certificate revocation. The current wisdom is to use Certificate Revocation Lists (CRL), however the processes for *timely* notification of a revoked certificate, removal of applied trust of the revoked certificate and the issue of retro-active revocation all pose significant hurdles to widespread implementation of PKI. Do all clients / applications that utilize certificates have built-in mechanisms for validating against a CRL? The reality is no, even current browsers do not automatically check CRLs. Is it realistic to assume the user will know to periodically check CRLs?

In PKI there is one piece that is vital to the entire process, protection of the private key. Weather you are just an average computer user, a user within a company, the company or the CA itself, securing the private key is absolutely critical. Inherent in the PKI model is the belief that entities will be able to appropriately secure the associated private key. Corporate environments are notoriously poor at implementing appropriate measures to ensure the level of security necessary to safeguard private keys. Getting management to fully comprehend the need and fund the necessary measures has proved to be a daunting task. Assuming that the need arises that individual consumers use PKC the issue is even more perilous. Your average computer user has little or no understanding of computer security. Additionally, retail computer products currently provide poor accommodations for the level of security that would be necessary to protect something as sensitive as your private key. Getting the user, weather he/she is an individual or acting as an agent of a company, educated to the level necessary to assure appropriate safeguards is a difficult and long-term endeavor.

The standards and specifications that PKI is being built around tend to be technical in nature, complex and only provide a framework from which to work within. Given this, there is a good deal or room for interpretation by PKI implementers. It is inherent in business for vendors within a product group to differentiate their product, to add or "enhance" features, which by design make them unique, creating implementation and interoperability issues. Added to this are legacy systems that do not necessarily interact well with PKI.

Non-repudiation presents a problem in that, as mentioned previously, in the proposed PKI model there is no inherent validation that the entity initiating the message is in fact the owner of the certificate. The verification process used to initially link a public key to an identity and create a certificate is vendor and time specific and heavily dependent on the practices and procedures of the issuing CA. There are many links in this chain where inaccuracies, mistakes and fraud can enter the process. Consumers have lived with certain safeguards within the physical world that have become customary and accepted and would be assumed within the Internet arena. To repudiate a signature in the physical world is always within the right of the alleged signatory. Further, that once repudiated the burden of proof that the signature is valid has fallen on the recipient, not the person alleged to have signed the document. With digital signature this concept is being reversed, putting the burden of proof of a repudiated signature on the certificate

holder. It is inconceivable that an individual would have the knowledge, access to the necessary information or the resources for such an undertaking.

Adding to the inherent issues with PKI, external factors are also hindering its growth. The Dot Com fallout of 2000, coupled with the following economic downturn has played havoc with corporate budgets. Managers are reluctant to embark on complex and expensive projects that do not have a clear and quantifiable cost-benefit objective. Information security, at least from a traditional management view, is not the easiest concept to sell since there's no immediate and tangible benefit to the bottom line. Until some event occurs that can be directly tied to a negative financial impact upper management does not want to hear about the importance of information security. In an Information Security 2001 industry survey one-third of the respondents indicated that their security budgets had been cut or frozen. Additionally, only 24 percent of those budgets were eventually restored to the original funding levels.

Perhaps the factor presenting the most difficult hurdle is ease of use and understanding. PKI is at best complex. The concepts of third-party trust, digital certificates, digital signatures, public key cryptography, transitive-trust etc, are not inherently simple to grasp. A significant portion of the people using computers today are by no means secure in their understanding of electronic communication and computer security. The most your average person hears about computer security is from the media when it's been breached in the form of a virus or hacking attack. It is the nature of human beings to be reluctant regarding concepts that are complex and not well understood. Particularly when it comes to legal or financial impact, people want to be able to understand the basic process and their role in it. PKI is a matter of trust and it is people who must accept this idea and trust in it in order for it to flourish.

Summary

There remain significant issues that must be addressed for PKI to survive and flourish. Digital certificates, digital signatures and CAs are a long way from being the be-all-end-all of computer security. Each aspect and step in the process is fraught with misinterpretation and ambiguity, not to mention people. The technological issues of interoperability, real-time certificate revocation; Social and legal issues surrounding trust, privacy and non-repudiation; And, process issues regarding identity and authentication must all be sufficiently worked-out in order to provide a level of comfort to allow for widespread buy-in to the model. Given the far-reaching implications of PKI, and level to which key aspects remain unresolved, it will continue to be some time before it is widely accepted and used.

References

1. Roger Clarke (2000). *Conventional Public Key Infrastructure: An Artefact Ill-Fitted to the Needs of the Information Society,* Version of 13 November 2000. Retrieved February 2002, from http://www.anu.edu.au/people/Roger.Clarke/II/PKIMisFit.html

2. Andy Briney (October 2001). *2001 industry survey.* Information Security. Retrieved February 2002, from http://www.infosecuritymag.com/articles/october01/images/survey.pdf

3. C. Bradford Biddle (1996). *COMMENT: Misplaced Priorities: The Utah Digital Signature Act and Liability Allocation in a Public Key Infrastructure*, 33 San Diego L. Rev. 1143. Retrieved February 2002, from http://www.acusd.edu/~biddle/mp.html

4. Bob Walder (August 1999). *Public Key Infrastructure*, An NSS Group White Paper. Retrieved February 2002, from http://www.nss.co.uk/pki/introduction_ed4.htm

5. Carl Ellison and Bruce Schneier (2000). *Ten Risks of PKI: What You're not Being Told about Public Key Infrastructure*. Computer Security Journal, v 16, n 1, pp. 1-7. Retrieved February 2002, from http://www.counterpane.com/pki-risks.html

6. A. Perez . *My Response to 10 Risks of PKI*. Retrieved February 2002, from http://home.pacbell.net/aram/ResponseTenRisks.html

7. Entrust, Inc. (August 2000). *Trusted Public Key Infrastructures*, version 1.2, An Entrust White Paper. Retrieved February 2002, from http://www.entrust.com/resources/pdf/pki.pdf

8. Elaine Blair and Alison Aiton (February 2001). *Issues in the Implementation of PKI in UK HE/FE*. Retrieved February 2002, from http://iii.gla.ac.uk/scotmid/gendocs/imppki-smp.html

9. Ed Gerck, Ph.d. (May 1998). *Why is Certification Harder Than It Looks?*. Retrieved February 2002, from http://www.mcg.org.br/whycert.htm

10. Ed Gerck, Ph.d. (July 2000). *Overview of Certification Systems:X.509, PKIX, CA, PGP & SKIP*. Retrieved February 2002, from http://www.mcg.org.br/certover.pdf