

How to Incorporate Revocation Status Information into the Trust Metrics for Public-Key Certification *

Kemal Bicakci
Vrije Universiteit
Amsterdam, The Netherlands
kemal@few.vu.nl

Bruno Crispo
Vrije Universiteit
Amsterdam, The Netherlands
crispo@few.vu.nl

Andrew S. Tanenbaum
Vrije Universiteit
Amsterdam, The Netherlands
ast@few.vu.nl

ABSTRACT

In a traditional PKI, the trust associated with a public key is expressed in binary either by 0 or 1. Alternatively, several authors have proposed trust metrics to evaluate the confidence afforded by a public key. However their work has a static point of view and does not take into account the issue of public key revocation. In this paper, we make the first attempt to incorporate the revocation status information into the trust metrics for public key certification. To achieve our goal, we use a tailored form of a vector of trust model recently proposed. This would allow us to reason formally about when there is a need to check revocation status and how reliable the revocation mechanism should be in a given security application.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General-Security and Protection

Keywords

Revocation, Trust Metrics, Public Key Certificates, PKI

1. INTRODUCTION

Based on public key cryptography, digital signatures are valuable tools to enable secure transactions over open networks. This basic cryptographic mechanism provides authentication, integrity and nonrepudiation services on the premise that binding of public keys to signing entities was properly performed. In other words, in order to validate a signature, first of all, one should trust that there is an authentic link between the corresponding public key and its owner.

*This research is supported by SecureE-Justice project funded by EU within the Sixth Framework under the contract IST-2002-507188.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'05, March 13-17, 2005, Santa Fe, New Mexico, USA.
Copyright 2005 ACM 1-58113-964-0/05/0003 ...\$5.00.

Such a link is generally provided by a public key certificate which is simply a signed statement (or a chain of signed statements). Now we are in a situation where in order to trust the public key of a user, we need to trust the public key of the authority issuing the certificate(s) as well as the security of its certificate issuance policy.

In a traditional PKI model, the trust relationship between any two entities is expressed in binary simply by either 0 or 1 i.e. 0 means not trusted and 1 means trusted. In this basic model if somebody trusts a certificate authority which issues a certificate for a public key, then this public key is trusted.

However, several authors [8] [13] claimed that trust can vary in a range and it should be possible to assign trust (confidence) values between 0 and 1 thereby it becomes possible to model for instance the increase in the trust if multiple paths certifying the same public key is utilized. How much trust improvement can be achieved using multiple paths is a question they have tried to answer by proposing different kinds of trust metrics.

On the other hand, none of these approaches consider the revocation problem. In order to establish the authenticity of the public key, we should also pay attention to the issue of public key revocation.

Revocation means invalidating the public key before its expiration date. Revocation is required for instance when there is a suspect that the private key has been stolen. How can the verifier assure that the public key was not revoked at the time of signing (or verifying)? The previous work implicitly assume that there would be no revocation at all or in a fully-trusted way there is always a fresh check verifying that the certificate is unrevoked. But we know that revocation happens and not rarely. Moreover fresh revocation check is expensive and not practical most of the time.

Our primary goal in this paper is to incorporate revocation status information into the trust metrics for public key certification¹. By doing this it becomes possible to answer questions like:

- If we check the revocation status and see that the public key is not revoked, what would happen to the trust metric?
- What if public key is revoked?

¹In fact, this research problem was considered as an interesting but non-trivial open problem in both [8] and [12].

- How would the freshness of revocation status information affect the trust metric?
- What do we lose if the revocation is not checked at all?

The problem of revocation is well-studied and various technical solutions have been proposed. However in all of the previous studies, they start from an assumption. They assume that revocation status would be always checked. In contrary, we observe that revocation is rarely checked in practice. In this paper we also would like to formally discuss the reasons of this real-life behaviour and give some directions regarding when somebody really needs to check the revocation status and how recent this check should be.

For clarity, we would like to emphasize that proposing a new trust metric for public key certificates is not the goal of this paper. Instead, we use prior work on trust metrics e.g. [8, 2, 12, 6] and propose a method to incorporate revocation status information into it.

The rest of the paper is organized as follows. Section 2 gives some background information. Section 3 explains our extended PKI model which includes the revocation. Section 4 shows how one can assign the weights and confidence values in our PKI model. Section 5 mentions some open problems and concludes.

2. BACKGROUND

In this section, we give some preliminaries on PKI models, trust metrics, vector model of trust and certificate revocation.

2.1 PKI Models

Maurer has proposed two models for a public key infrastructure [8]. He called the first one as the deterministic model which is very similar to what we have presented as the traditional model in the introduction section. However this model takes into account not only certificates but also recommendations. Recommendations are again signed statements like certificates but they have a more general meaning. One can specify the trustworthiness of another entity in a particular context and assign different recommendation levels. For instance a recommendation of level 1 means that the recommended entity is trusted to certify others and a recommendation of level 2 means that the entity is trustworthy in recommending others for certification. As Maurer suggested, certificates can be regarded as recommendations of level 0. Although the topic of recommendation has a great importance, we do not discuss it here further and concentrate on the simpler model involving only the certificates.

Let us now see the fundamental rule of deterministic model when we exclude recommendations:

$$\forall X, Y : Aut_{A,X}, Trust_{A,X}, Cert_{X,Y} \vdash Aut_{A,Y} \quad (1)$$

What this equation says is basically if A trusts somebody X and A can authenticate² X and when there is a certificate issued by X to Y then this implies that A can authenticate

²It is also possible to think that trusting somebody means authenticating him in the first place. How would you trust somebody whom you can not authenticate? But in this model these two things are differentiated.

Y (A trusts the certificate of Y when the three statements on the left side are satisfied).

It is obvious that this simplified form does not say something new. Nevertheless it is important in the sense that the probabilistic model can be best explained using it. We use the deterministic model to explain also a probabilistic model.

The deterministic model would not be satisfactory to model a PKI from a user's point of view. As first presented by Zimmermann [13], trust is not something binary and can range from marginal to ultimate³. As a matter of fact, intermediate degrees of trust is all possible.

Maurer's second model is a probabilistic model where the confidence value of a certificate is calculated as the product of confidence parameter of authentication of entity issuing the certificate (denoted by $p(Aut_{A,X})$), confidence parameter of trust on the entity issuing the certificate (denoted by $p(Trust_{A,X})$), and confidence parameter of the certificate issued (denoted by $p(Cert_{X,Y})$)⁴. Confidence value and parameters are in the range of 0 to 1 in the following example.

Example 1: Let us suppose that

$$p(Aut_{A,X}) = 0.8, p(Trust_{A,X}) = 0.9, p(Cert_{X,Y}) = 0.7.$$

Then we have

$$conf(Aut_{A,Y}) = p(Aut_{A,X}) * p(Trust_{A,X}) * p(Cert_{X,Y}) = 0.8 * 0.9 * 0.7 = 0.504.$$

If this confidence value would not be enough to perform the transaction, it is possible to have multiple paths certifying the same public key and increase the confidence value since the confidence level in case of multiple paths adds up. We do not discuss here how the confidence value can be calculated if multiple paths are used. For interested readers, we recommend the original paper [8]. See also [3] for an easy-to-understand explanation of why multiple paths are useful and desired.

In the example above, the first two confidence parameters are assigned by the user and the third one is assigned by the entity issuing the certificate. Notice the difference between the confidence value of the certificate and confidence parameter of the certificate assigned by its issuer. The latter can be thought as another way of specifying certificate classes. As discussed in [7], CA companies like Verisign issue certificates in different classes and these classes are assigned depending on the reliability of the method used to identify the certificate owner. If the certificate issuer is fully trusted, the confidence value of a certificate would have the same value as its confidence parameter assigned by the issuer.

2.2 Trust Metrics

The term "trust metric" can be defined as the measure of amount of trust attached to something. This "something" can simply be anything: a person, a file, a URL, a country etc. Public key certificates are just one other application for trust metrics.

In the Maurer's probabilistic model of PKI, probabilities as parameters of subjective belief (denoted as confidence levels) are used as the trust metric for public key certificates. Other than this one, there are all sorts of different ways to represent trust values. In the following, we will mention only

³However, PGP is flaky because of the way these trust levels are defined and used as discussed in [12] and [6].

⁴In Maurer's terminology, the term "parameter" is used when the confidence level is assigned and the term "value" is used when the confidence level is derived.

one other and leave the survey of the rest to [10].

Assigning confidence parameters between 0 and 1 to various trust statements has raised questions among other researchers. They argue that there is an ambiguity in the semantics and there is no easy way to determine the value of these parameters in real world. For instance, Reiter and Stubblebine [12] suggested a different approach to assign such values and to use numeric labels as a trust metric representing the amount of money the certificate issuer is liable if the authentication of the corresponding public key fails.

It might be true that this insurance-type of trust metrics fits better to real world and most of us understands better if money is the concern. However we stick ourselves to use Maurer’s metric since the way Maurer’s confidence levels are used can include different approaches because

- It is always possible to convert a value in one trust metric to a value in another metric when the metrics have a finite range. For example, money can always be translated in confidence values by normalizing the value to range [0-1].
- These confidence parameters can be based on past experience. For instance it can reflect the risk associated or the reputation of the certificate issuer.
- For some applications, the parameters can be automatically assigned by the system based on some published security policy (this policy can be updated dynamically). If this can be done, the end-user does not need to bother assigning them.

On the other hand if we look at the trust problem from the reverse angle we see that for a wide range of security applications available today, it would be very convenient to set a minimum confidence value required (as a part of the aforementioned security policy) so that the decision whether the trust established would be enough to perform the transaction is not left to the end user.

2.3 Vector Model of Trust

Trust is something difficult to measure, compare and combine. For a better reasoning of trust, Ray and Chakraborty proposed a vector model to formulate trust in a very recent work [11]. In their model, trust is defined as a vector of several parameters each of which contributes to the overall trust but not in same amount.

As this trust model suggested, one important characteristic of trust is the ”propensity of trust” which means that two trusters may assign two different trust values even when all the factors influencing the trust has exactly the same value. The main reason of this phenomenon is that the truster may assign different weights to different factors during the evaluation of the trust value. The authors have introduced the concept of *trust policy vector* to capture this characteristic. In section 3.2, we will see one example of this policy vector which is simply a vector of weight values.

2.4 Certificate Revocation

As we said the status of a public key certificate can change during its lifetime due to unexpected events (i.e. loss of private key, etc.) thus a solution to reflect this possible change should be found.

A straightforward solution is to use instant or short-lived certificates which have a very short lifetime therefore eliminates the requirement of revocation check. Since this method requires the users to frequently communicate with the CA to get new certificates, it can not be always used.

For long-lived certificates (with a validity period of months even years), a separate mechanism for revocation check needs to be implemented. By checking revocation information, the verifier will know if the certificate(s) he is going to rely upon is still valid or it has been revoked.

The simplest and most widely used mechanism to implement revocation are *certificate revocation lists: CRLs*. With CRL, the revocation authority issues periodically a signed and timestamped list of the serial numbers of all the certificates that have been revoked at the time of issuing. To reduce the required bandwidth and most important to reduce the need for on-line connectivity, CRLs should not be issued too frequently. However, increasing the interval between two subsequent CRLs increase also the probability that a certificate that has been reported compromised by its owner has not yet been published and distributed in the list.

Delta CRLs and *Certificate Revocation Trees (CRTs)* [5] are two different improvements to basic CRLs. They succeed in decreasing the bandwidth of revocation information.

Since most of the time, the verifier is not interested in knowing which certificates have been revoked but rather checking if the certificate he just received has been revoked or not, a new mechanism, *On-line Certificate Status Protocol (OCSP)* [9] has been introduced.

With this mechanism, the verifier can query a designated server about the status of a specific certificate. The server always online will reply by signing a statement which asserts the fresh status (valid or revoked) of the certificate in question. The obvious disadvantage of this method is the requirement of on-line connectivity.

3. OUR PKI MODEL

To incorporate revocation status information into the trust metrics for public key certification, we extend Maurer’s PKI model. Like the original model we have two cases:

3.1 Deterministic Model

As expected, incorporating revocation into the deterministic model is the easier one. We simply modify the rule given in equation 1 as follows:

$$\forall X, Y : Aut_{A,X_1}, Trust_{A,X_1}, Cert_{X_1,Y} \vdash CondAut_{A,Y} \quad (2)$$

$$\forall X, Y : Aut_{A,X_2}, Trust_{A,X_2}, NotRev_{X_2,Y} \vdash UnRev_{A,Y} \quad (3)$$

$$\forall X, Y : CondAut_{A,Y}, UnRev_{A,Y} \vdash Aut_{A,Y} \quad (4)$$

Again these equations state the common sense. Conditional authentication denoted by ”CondAut” means that authentication will be completed only if revocation status is checked as equation 3 says and public key is confirmed not to be revoked.

We observe a couple of inherent weaknesses of this deterministic model:

- This model does not distinguish between revoked public key and the public key without any check of revocation.
- This model does not say anything about the freshness of revocation check and reliability of revocation information.
- The authority used to authenticate the public key might be different than the authority for getting the revocation information. In fact, these two authorities should not be the same otherwise there might be some security vulnerabilities [4]. Again this model does not take into account the different trust values for these two different authorities.

3.2 Probabilistic Model

The model introduced in this paper takes into account only the confidence value of the certificate, the result of Maurer’s formula denoted by $conf(Aut_{A,Y})$. The trust metric and the method used to calculate this value really do not matter. The method can be either Maurer’s method or any other one. Our model is flexible in the sense that it allows us to embed other models and other metrics.

Now we would like to introduce our extended probabilistic PKI model which contains the revocation status information. In our model the confidence value for the public key certificate is assigned explicitly without revocation. We would like to take into account also the revocation therefore in a similar manner we have a confidence value for the revocation information. While doing this, the confidence value for revocation is overloaded also by freshness constraints in our model. For instance a confidence value of 1 would mean that the revocation authority is fully trusted and the revocation information is gathered securely exactly at the time when it is needed. No later, not before.

Now we have two confidence values, one for the public key certificate and the other is for the revocation information of the public key. How can we combine these two values in a reasonable manner? This is where vector of trust model introduced in section 2.2 comes into the scene.

Definition 1: The trust vector $(A \rightarrow Y)$ for our model has two dimensions. One of them is $Conf_A$ (confidence value of public key certificate) and the other is $Conf_R$ (confidence value for revocation information).

$$(A \rightarrow Y) = (Conf_A, Conf_R) \quad (5)$$

Definition 2: The trust policy vector W for our model has two dimensions. One of them is W_A (weight for checking authenticity of certificate) and the other is W_R (weight for checking revocation).

$$W = (W_A, W_R) \quad (6)$$

Then the normalized trust relationship between A and Y can be expressed as

$$(A \rightarrow Y)^N = W \odot (A \rightarrow Y) \quad (7)$$

The symbol \odot denotes vector multiplication.

Definition 3: The value of normalized confidence value for the public key certificate is in the range $[0, 1]$ and is defined as

$$v(A \rightarrow Y)^N = W_A * Conf_A + W_R * Conf_R \quad (8)$$

Up to this point, our usage of vector of trust model is very similar to what have been proposed by Ray and Chakraborty [11]. We only tailor it in two aspects:

- The trust lies within $[0, 1]$ not within $[-1, 1]$ in our model. We believe this difference is not so important.
- The sum of all elements in the trust policy vector should always be 1 in [11]. However this is not always the case in our model. This property is captured in the following definition.

Definition 4: The trust policy vector has a length of 1 (the sum of W_A and W_R is equal to 1) when the public key is unrevoked and 0 (the sum of W_A and W_R is equal to 0) when the public key is revoked. Because if there is revocation this negates the trust for public key certificate.

4. DISCUSSION

After introducing our extended PKI model, we would like to discuss how the trust policy vector is set up and how one can determine the confidence values for revocation. We also say a few words on how our model can help to decide when to check revocation status.

4.1 Determining the Trust Policy Vector

Since our trust policy vector with a specified length has only two dimensions, it would be sufficient to present a methodology for the value of one of the weights.

In our model as equation 9 says, the weight for checking revocation information (W_R) is fixed to be equal to the ratio of revoked certificates to the total number of certificates issued so far. Here, our rationale is that this ratio gives us a good estimate for the probability that a given certificate is revoked.

$$W_R = \frac{\text{Number of revoked certificates}}{\text{Number of all certificates}} \quad (9)$$

We would like to state that this ratio can change from system to system and from application to application. There are number of parameters that can affect this ratio. For instance, it would be different ratios in applications where public keys support credit card transactions and in applications where public keys are used for IPsec connections.

Therefore while calculating the ratio above, it is crucial to consider only the certificates that can be used in the specific domain. For instance if the model is to be used in an application where credit card transactions are supported by public key certificates, only the certificates supporting such a transaction should be taken into account.

We have searched the previous work on revocation to see the possible values this ratio can take in different applications, but unfortunately we could not find much⁵. We believe this piece of information is essential and in our PKI model, we require CA or RA (Revocation Authority) to publish with the revocation list (or the other revocation mechanisms) not only the list of revoked certificates but also the

⁵We have only found in an outdated document that approximately 10% of public key certificates are revoked in general [1].

ratio of revoked certificates or at least the total number of current valid issued certificates so user can compute the ratio and assign weights accordingly.

Example 2: Suppose instant certificates as discussed in subsection 2.4 are used. That means there is no revocation and the weight of checking revocation information is simply zero.

Example 3: Assume that the ratio of revoked certificates is 0.1 in a given security application, we assign value of 0.9 to W_A and 0.1 to W_R when there is no revocation. As we have stated in Definition 4, if the certificate has been revoked, this has to negate the effect of certification and therefore W_R should be assigned to (-0.9) if this is the case.

4.2 Assigning Confidence Values for Revocation Information

Next, we would like to discuss another issue in our model, assigning confidence values for revocation information. As we have mentioned in subsection 3.2, the meaning of confidence value for revocation is overloaded in our model because it does not involve only the trust relationship between entities but also the reliability of revocation mechanism in use. Suppose for a moment that the revocation authority is fully trusted but the latest CRL it has issued is not a fresh one, do you put full trust on the revocation information?

It is a rule of thumb in our model that the confidence value if CRL-type methods are in use is lower than the value if OCSP is used. It is also possible to differentiate between fresh CRL and stale CRL by assigning different confidence values even when the authority issuing the lists is the same. Due to space limitations we do not go into further details.

4.3 When to Check Revocation Status

We believe our model explains why revocation is rarely checked in real-world. Users believe intuitively that for their security applications it is sufficient to check only the authenticity of the public key certificate since the confidence value of the public key certificate has a much greater weight (the probability of revocation is low). In fact, it is true that some of the applications really do not have stringent trust requirements. But we expect that when certificates start supporting higher-value transactions, using our model people can realize that even when public key certificate is fully trusted, the overall confidence value is not adequate and therefore there is a need to check the revocation status.

We need revocation check definitely when the minimum confidence value required is above W_A . In case it is lower than W_A , revocation check should again be performed if the confidence value of certificate is lower than 1 and we would like to reach the minimum confidence value required.

5. CONCLUSION AND FUTURE WORK

In this paper, we have considered the problem of certificate revocation in the context of trust metrics for public key certification. We have tackled this problem by extending Maurer's PKI model with a tailored form of a vector of trust model. Our model was not abstract, we were able to justify our model by showing how we can assign values properly to various parameters of the model i.e. the weights and the confidence value for the revocation information.

Since this is only the first attempt, there are so many open problems left. We have a crowded future work list:

- Trust metrics having a finite range can be embedded into our model by normalizing them to the range $[0, 1]$. But we need to consider also the trust metrics with an infinite range.
- There is definitely a need to search for the ratio of revoked certificates and its distribution over time in different security applications.
- More work is required to establish the minimum confidence value required for the public key certificate. This policy work seems to be the most challenging one.
- Our model is a simple one and there are some interesting cases not considered. For instance what if there are multiple paths certifying the same public key and one of them is revoked. This is again a question not easy to answer.

6. REFERENCES

- [1] S. Berkovits, S. Chokhani, J. Furlong, J. Geiter, and J. Guild. Public key infrastructure study: Final report. *Produced by the MITRE Corporation for NIST*, 1994.
- [2] T. Beth, M. Bocherding, and B. Klein. Valuation of trust in open networks. In *Proc. of ESORICS'94, LNCS 875*. Springer-Verlag, 1994.
- [3] M. Burmester and Y. Desmedt. Is hierarchical public-key certification the next target for hackers? *Commun. ACM*, 47(8), August 2004.
- [4] B. Crispo and M. Lomas. A certification scheme for electronic commerce. In *Proc. 1996 Security Protocols Workshop, LNCS 1189*. Springer-Verlag, April 1996.
- [5] P. Kocher. A quick introduction to certificate revocation trees (crt). <http://www.valicert.com/company/crt.html>.
- [6] R. Kohlas and U. Maurer. Confidence valuation in a public-key infrastructure based on uncertain evidence. In *Proc. of PKC'00, LNCS 1751*. Springer-Verlag, 2000.
- [7] A. Levi and C. K. Koc. Risks in email security. *Commun. ACM*, 44(8), August 2001.
- [8] U. Maurer. Modelling a public-key infrastructure. In *Proc. of ESORICS'96, LNCS 1146*. Springer-Verlag, September 1996.
- [9] M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 internet public key infrastructure - online certificate status protocol - ocsp. *IETF, RFC 2560*, June 1999.
- [10] A. A. Rahman. Survey of trust models for computer networks. <http://www.cs.ucl.ac.uk/staff/F.AbdulRahman/thesis/csreview.pdf>.
- [11] I. Ray and S. Chakraborty. A vector model of trust for developing trustworthy systems. In *Proc. of ESORICS'04, LNCS 3193*. Springer-Verlag, September 2004.
- [12] M. Reiter and S. Stubblebine. Authentication metric analysis and design. *ACM Transactions on Information and System Security*, 2(2), May 1999.
- [13] P. Zimmermann. *PGP User's Guide, vol. I and II, version 2.6*. 1994.