

# Empirical Analysis of Certificate Revocation Lists<sup>\*</sup>

Daryl Walleck, †Yingjiu Li, and Shouhuai Xu

University of Texas at San Antonio, 6900 North Loop 1604, West San Antonio, TX  
78249

†Singapore Management University, 80 Stamford Road, Singapore 178902  
dwalleck@cs.utsa.edu, yjli@smu.edu.sg, shxu@cs.utsa.edu

**Abstract.** Managing public key certificates revocation has long been a central issue in public key infrastructures. Though various certificate revocation mechanisms have been proposed to address this issue, little effort has been devoted to the empirical analysis of real-world certificate revocation data. In this paper, we conduct such an empirical analysis based on a large amount of data collected from VeriSign. Our study enables us to understand how long a revoked certificate lives and what the difference is in the lifetime of revoked certificates by certificate types, geographic locations, and organizations. Our study also provides a solid foundation for future research on optimal management of certificate revocation for different types of certificates requested from different organizations and located in different geographic locations.

**Keywords:** public key infrastructure, certificate revocation, empirical analysis

## 1 Introduction

With the rapid growth of the Internet over the last decade, new challenges appear daily. Of these challenges, perhaps none is more important than the need for protecting sensitive transactions. By means of digital certificates, public key infrastructures provide a degree of authentication to protect sensitive transactions. However, digital certificates must be revoked if the corresponding private keys have somehow become compromised, perhaps due to attacks launched by worms or viruses. Thus, managing certificate revocation efficiently has become a major issue in public key infrastructures [14].

Previous research on certificate revocation management has primarily focused on the tradeoffs that can be made among different revocation mechanisms [6, 15], including certificate revocation list (CRL) [5], certificate revocation system (CRS) [12], certificate revocation tree (CRT) [7], and on-line certificate status protocol (OCSP) [11]. Though various tradeoffs have been studied, little effort has been made toward understanding the distribution of certificate revocations, especially from real-world data. Understanding the distribution of certificate

---

<sup>\*</sup> The work of Shouhuai Xu was supported in part by ARO, NSF and UTSA CIAS.

revocations would enable certificate authorities to optimize their operations over time.

**Our contributions.** We collected five real-world certificate revocation files from VeriSign for different types of certificates, and conducted an in-depth empirical study to understand the distribution of certificate revocations from different perspectives. This paper reports the major findings of our empirical study, which can be summarized as follows.

- The types of certificate revocation files, which are used for different purposes, do not appear to be a fundamental factor regarding the behavior of certificate revocation distributions. This is so because all the five individual certificate revocation files exhibit exponential distribution patterns, so is the merged dataset. Nevertheless, different types of certificates can still be clustered into two groups based on their mean certificate lifetimes, where each certificate’s lifetime is defined to be the difference between its revocation date and its issue date. This may suggest that certain classes of certificate-enabled systems (e.g., code signing and financial applications) are better protected than others under the assumption that other factors that affect the certificate lifetimes remain similar in the comparison.
- Although certificate revocations in different geographic locations still exhibit exponential distributions, the distribution parameters vary significantly. This implies that different strategies should be used to disseminate certificate revocation information for different countries or continents. Moreover, the average certificate lifetimes may serve as a measure for the security levels of certificate-enabled systems in different geographic locations provided that no other factors that affect the certificate lifetimes are significantly different in comparison.
- The number of revoked certificates is bouncing on a daily basis. In particular, many certificate revocations occur during weekdays, whereas few occur during weekends. This indicates that an attacker who compromises a public key certificate during weekends may have a better opportunity to conduct unlawful activities before the compromised private key (i.e., the corresponding certificate) is effectively revoked. We also observed that the numbers of revoked certificates in January and February in both 2005 and 2006 are always significantly lower than their respective counterparts in other months of the same year. This is not because the certificate-enabled systems are better protected, but because fewer certificates are requested and issued due to seasonal reasons.
- Different organizations exhibit different characteristics in terms of their certificate lifetimes. Although the certificate lifetimes still follow exponential distributions, the average certificate lifetimes vary widely among different organizations, even within the same industry group such as financial institutions. This result may stimulate organizations to improve their security levels and security awareness in a competitive market.

**Limitations of this paper.** The nature of empirical study restricts us from extrapolating our results to the whole universe. In particular, the following limitations of the present study are identified for possible future improvements.

- The major findings of this paper are based on a number of CRL data sets collected from VeriSign only. Though considered to be representative for commercial use of public key certificates, VeriSign’s data may not demonstrate the same revocation patterns as other data sets. In addition, our findings cannot be extrapolated to OCSP responders.
- We do not have access to the certificates that are issued by VeriSign but never revoked. While it is meaningful to investigate the ratio of revoked certificates to the certificates that are never revoked, we experience difficulties in collecting such data from VeriSign or any other resources in the public domain (in most cases, only the information regarding the revoked certificates is available to the public).
- We do not investigate why the certificates are revoked. Understanding various revocation reasons will definitely help us understand the relationship between certificate revocation and the security levels of certificate-enabled systems. For example, one can suspect that no one cares too much about their SSL certificates if they lose the private keys, as they can get new certificates minted, maybe from someone cheaper than VeriSign. For another example, it is very important to revoke public key certificates if someone loses his company’s smart cards. Unfortunately, it is very difficult to obtain such information as the revocation reasons are often considered sensitive in commercial applications (to some extent, this is similar to the situation in which financial organizations are disinclined to publish any security breaches to the public).
- We do not consider many other factors in certificate revocations except the security factor on which we focus. There could be a host of other factors affecting certificate revocations: (i) the errors made in data entries, (ii) the purposes of the certificates being used, (iii) the reasons of the certificates being revoked, (iv) the administration policies for certificate revocation, and (v) the fraction of all issued certificates that get revoked. We assume that all these factors are similar when we make connections between certificate revocations and security levels in certificate-enabled systems.

**Related work.** The work most closely related to this work is the paper “On the Release of CRLs in Public Key Infrastructure” by Ma, Hu, and Li [9], which builds analytical models on how often a certificate authority should release CRLs in order to minimize its operational cost based on empirical analysis on real-world data. However, their analysis of the data gathered from certificates is not as in depth as what is proposed in this paper. In particular, they did not consider the impact of geographic location and organization to the distribution of certificate revocations. Another difference is that they proposed optimal CRL releasing strategies, while the main purpose of our study is to characterize certificate revocations based on real empirical data.

Except [9], most of previous researches are not based on any empirical analysis of real-world data; instead, they focus on theoretical aspects of certificate revocation including the meaning of revocation [3, 4], the model of revocation [2], communication cost of revocation [12], tradeoffs in certificate revocation schemes [16], and risk management in certificate revocation [8]. Rivest has once proposed to use short-lived certificates so as to eliminate certificate revocations [13]. However, his approach places a high burden on certificate servers which need to sign more certificates as compared with traditional certificate revocation solutions; it also creates the problem of key compromise which cannot be addressed without using a separate mechanism [10].

**Paper organization.** The rest of this paper is organized as follows. In Section 2, we discuss the methodology we used to collect and analyze real-world data from VeriSign. In Section 3, we analyze the VeriSign data from various perspectives including differences in certificate revocation between certificate classes, geographic factors in certificate revocation, trends in certificate revocation rates over time, and trends in certificate revocation rates by organizations. We also discuss how to derive optimal certificate revocation policies based on our empirical results. Finally, in Section 4, we summarize our research and point out possible future directions.

## 2 Methodology and Data Collection

To investigate certificate revocation, we used VeriSign’s Certificate Revocation Lists (CRLs) to find certificates that have been revoked over the last several years. After gathering a large sample of revoked certificates, VeriSign’s database was queried using its web interface to determine relevant information about each certificate such as when the certificate was issued, what organization requested the certificate, and its country of origin. However, not all of their certificate data is publicly accessible. Though VeriSign allows users to determine the status of some certificates through a web interface, we could not find information about certificates from all CRL files through it. Because of this, our analysis is limited to the data we could gather from the files mentioned later.

We also encountered similar problems when considering analyzing data from other certificate authorities such as Thawte and GeoTrust. Since the CRL file contains the revoked date for each certificate, we would require some way to determine the date a certificate was issued to determine its lifetime. Though both certificate authorities do publish CRLs, neither of them offers an interface to search their certificate database, making any analysis of their CRL files impossible.

We were also interested in discovering the number of active certificates (including those never revoked) so that it would be possible to compare the number of revoked certificates to the number of active certificates at a given time. Unfortunately, aside from searching the Internet to find live VeriSign certificates, there is no easy way to determine this. Though VeriSign’s web interface does allow users to search its database by organization name which does return some valid

certificates, organizations can also request that their certificates not be viewable through that interface. Because of this, even if we had attempted to build a list of valid certificates, there would be no guarantees of its completeness.

Using the CRLs available, we were able to analyze the data gathered to characterize certificate lifetimes by different sub-fields. In this paper, we will try to characterize the lifetime of revoked certificates by the following criteria:

- The lifetime of certificates over time
- The difference in the lifetime of certificates by type
- The difference in the lifetime of certificates by geographic location
- The characteristics of certificate lifetimes by organization

File Name	Issue Date	# Certificates	Dates Covered	Purpose
SVRIntl	3/26/2007	21192	2/15/2005-3/26/2007	Global Server certificates
RSA	3/13/2007	10100	12/18/2004-3/13/2007	Secure Server certificates; also a root CA
Secure	3/26/2007	11898	12/18/2004-3/26/2007	Secure Server certificates
Financial	3/26/2007	326	5/7/2002-3/26/2007	OFX certificates
Code signing	3/13/2007	1413	9/28/2004-3/13/2007	Code signing and object signing certificates for use with Netscape browsers, Microsoft Internet Explorer browsers, Microsoft Office, Sun Java Signing, Macromedia, and Marimba


**Table 1.** Breakdown of the Composite Data Set by CRL File

For the purpose of this paper, five CRL files were used to find revoked certificates which were used to create our data pool. The files chosen for this research are described in Table 1. Since VeriSign removes most certificates from its CRLs after they expire [1] (most certificates have a one to two years issued lifespan before their expiration), most of the certificates contained in the lists cover the past two years. Between these five CRL files, 44,929 certificates were gathered. Since each CRL file only includes the serial number and revocation date for each certificate, a Ruby script was used to search VeriSign’s database for each certificate’s issue date, country of origin, and the organization that requested the certificate.

### 3 Empirical Analysis

First, we would like to examine the trend of certificate lifetimes for revoked certificates from all of the CRL files. The lifetime of a certificate can be defined as follows:

$$\text{Actual Lifetime} = \text{Date Revoked} - \text{Date Issued}$$



**Fig. 1.** Number of Certificates Revoked vs. Certificate Lifetime **Fig. 2.** Number of Certificates Revoked vs. Percentage of Lifetime Achieved

We begin our analysis by plotting the lifetime of a certificate against the amount of certificates revoked for that lifetime. As can be seen in Figure 1, the lifetime of a revoked certificate is fairly short. In fact, the average lifetime for certificates in the composite data set is 28 days. However, this plot does not take into account the fact that certificates expire at different rates. While some certificates may only be valid for a year, the issued lifetime of other certificates may be two or three years. The issued lifetime of a certificate can be calculated as:

$$\text{Issued Lifetime} = \text{Expiration Date} - \text{Date Issued}$$

To see what kind of difference this might have, in Figure 2 we take this into account by plotting the percent of a certificate's normal lifetime against the number of certificates that were revoked after that percentage. As can be seen, the trends displayed in Figure 1 still hold. We discovered that the mean percent lifetime of any given certificate is 4.8%.

By using the `dfittool` and `expfit` functions of Matlab, it was determined that this data follows an exponential distribution. The common form of the probability density function (PDF) for the exponential distribution is as follows:

$$R(t) = ke^{-kt}$$

However, Matlab uses an alternate form of the exponential distribution. This form is:

$$f(x) = \frac{1}{\mu} e^{-\frac{x}{\mu}}$$

The composite data set was discovered to follow the exponential distribution with the parameter  $\mu$  being 27.56 at a 95% confidence interval. When testing the percent lifetime view, it was also determined to follow the exponential distribution with  $\mu = 0.0479$  at a 95% confidence interval. This is an interesting finding: most revoked certificates have lifetimes shorter than a month, or 4% of their issued lifetimes, even though they have one to two years issued lifetimes. As it is mentioned earlier, our research is restricted to the certificates that get revoked.

### 3.1 Differences Between Certificate Classes


Now that we have examined the characteristics of the data set as a whole, we were also interested in breaking down the data into the individual files and seeing how well the distribution holds. Table 2 shows the mean lifetime for the revoked certificates from each CRL file. While the International, RSA, and Secure server certificates have relatively similar mean lifetimes, the mean lifetime for Code Signing and Financial certificates is nine to ten days (or about 25%) longer than the others.

CRL File	$\mu$ (mean)
International	26.83
RSA	27.12
Secure	27.99
Code Signing	35.72
Financial	37.08

**Table 2.** Mean Certificate Actual Lifetime by CRL File

Figure 3 plots each of the five CRL files separately to see how well the distribution holds. Though there is some difference in scale, the data from each file still follows the exponential distribution with the parameter  $\mu$  shown in Table 2, all at a 95% confidence interval.

This result indicates that the type of certificates is not a fundamental factor regarding the distribution of certificate revocations. An exponential distribution is observed for each of the five certificate revocation files and for the merged dataset. The difference in the mean value of certificate lifetimes may suggest that certain classes of certificate-enabled systems (e.g., code signing and financial) are better protected than others under the assumption that all other factors affecting certificate lifetimes are similar in comparison. It should be noted that the protection levels are not the only factor affecting the certificate lifetimes. For instance, the differences in certificate lifetimes in different certificate-enabled systems could suggest that the administrators of certain systems such as financial servers (for which certificates get revoked slowly) work in environments in which



**Fig. 3.** Certificate Actual Lifetime by CRL File

it takes longer to get authorization for revocation. On the other hand, some certificates get revoked quickly because errors were made in data entries, or because the certificates were for tests or experimental systems. To get a comprehensive understanding about protection levels, one needs to know about all these factors that affect certificate lifetimes, which may include the purposes of the certificates being used, the reasons of the certificates being revoked, the administration policies for certificate revocation, and the fraction of all issued certificates that get revoked. The certificate lifetime in a certificate-enabled system can be considered as a multivariate function of many variables; in our discussions, we focus on the variable of protection levels while assuming that the other variables are fixed. A more comprehensive study on all such variables is an obvious topic for future work if sufficient data is available.

### 3.2 Geographic Factors


Now that a standard has been created to compare against, we would first like to discover if geographic location has any influence on the lifetime of a revoked digital certificate. In all, 136 countries were identified in the CRL files we used. To begin, we first investigated CRL usage of the country with the most total certificates revoked, the United States, and plotted the results in Figure 4.

Considering that certificates from the United States make up a large portion of the composite data, it is not surprising that Figure 4 is very similar to Figure 1. Before coming to any conclusions, we then plotted the results for four of the other leading certificate holders in Figure 5. The difference in the amount of certificates used by these countries is significantly smaller than that of the United States, so a smaller scale will be used to display these graphs.


We also examined the behavior of certificate lifetimes over the first 30 days more clearly, Figure 6 shows the same data from Figure 5 on an even smaller scale (i.e., Figure 6 is a “zoom in” of a portion of Figure 5).

Like the composite data set, when divided by country these data sets also follow the exponential distribution. Table 3 gives the parameter  $\mu$  for each of





**Fig. 4.** Number of Certificates Revoked vs. Lifetime (United States) **Fig. 5.** Number of Certificates Revoked vs. Lifetime (Remaining Countries)



**Fig. 6.** Number of Certificates Revoked vs. Lifetime (Remaining Countries - magnified)

the data sets (all at a 95% confidence interval). Interestingly enough, the trends shown in the initial results hold true when the data is broken down by geographic region. In all cases, a large number of certificates are revoked within the first month before falling off to a few revocations per day after that. By these results, it can be inferred that location plays only a minor role in certificate revocation rates. However, it is also of interest to note that the average lifetime of a revoked certificate in Japan is less than half that of any of the other countries shown here. Since the average certificate lifetimes vary significantly for different geographic locations, different strategies may be used to disseminate certificate revocation information for different countries or continents. Moreover, the average certificate lifetimes may serve as a metric for the security levels of certificate-enabled systems in different geographic locations provided that all other factors that affect the certificate revocation are the same.


### 3.3 Trends in Revocation Rates over Time

Another view of the data we were interested in was tracking certificate revocation rates over time. If surges in revoked certificates could be found, we would ex-


US	JP	GB	CA	AU	DE
29.48	13.15	25.18	37.26	30.62	29.36

**Table 3.** Mean Certificate Lifetime by Country (in days)

pect that these surges could be traced back to the occurrence of major security incidents such as widespread worms or viruses. Figure 7 displays the amount of certificates revoked per day from the composite data during the period of 1/1/2006 to 12/31/2006.




**Fig. 7.** Certificates Revoked Per Day (2006)



**Fig. 8.** Certificates Revoked Per Day (2005)


Though we did observe small peaks in the amount of certificates revoked per day, there were no extraordinarily large spikes in certificate revocations like we thought there might possibly be. What we did notice was the wave-like bouncing pattern that certificate revocations follow. Upon further investigation, we found that nearly all certification revocations happened between Monday and Friday, with only minimal revocations occurring on weekends. To make sure that this trend did not only occur in 2006, we also investigated the data from 2005 in Figure 8 and compared the trends between 2006 and 2005 in Figure 9.

Surely enough, the pattern still holds. From these two figures, we observed that the number of certificates revoked per day in January and February is significantly smaller than the number of certificates revoked per day in other months. Another observation is that the number of certificates revoked per day increases significantly from 2005 to 2006. These changes are primarily due to the changes in the total number of certificates being issued at different times. To make this clear, we also investigated the percentage of the valid certificates revoked each day. Since we did not have access to VeriSign's database to determine the true number of certificates active at a given time, we instead used the certificates from the CRL files to determine the number of not-yet-revoked certificates daily. Figure 10 below plots the percentage of certificates revoked daily



**Fig. 9.** Number of Certificates Revoked Per Day (2006 compared to 2005)


(over the total number of certificates that have not been revoked at the beginning of the day, which would vary on a daily basis) for the period of 1/3/2005 through 3/26/2007.



**Fig. 10.** Percentage of Certificates Revoked Daily


Because every certificate in a CRL file is eventually revoked, the end of the curve in Figure 10 is skewed because at the end of the time period, 100% of the certificates are revoked. However, this is artificial and does not affect the data before it. On average, 4% of the total revoked certificates were revoked daily. From this plot it can be seen that no matter how many total revoked certificates are in existence, the percentage of certificates revoked daily stays fairly constant with some small growth over time. Next, we were interested in seeing how the number of not-yet-revoked certificates plots over time. In Figure 11, we plot the number of not-yet-revoked certificates over the same period of time.

Since these CRL files contain only certificates that were eventually revoked, the number of certificates active at the end becomes zero. Other than the rise



**Fig. 11.** Number of Not-Yet-Revoked Certificates over Time **Fig. 12.** # Not-Yet-Revoked Certificates Compared to # Revoked Certificates

and decline at the starting and ending periods, there is only one sharp change in active certificates, as the number of active certificates double in number in March, due to expired certificates being removed from the CRLs. Even with this large increase of active certificates, the percent of certificates revoked daily only gradually rises. This implies that the number of revoked certificates changes in a similar trend as the number of active certificates does and this similarity in trend is illustrated in Figure 12 on a daily basis.



**Fig. 13.** Number of Certificates Revoked Per Day - Breakdown by Length of Lifetime

In Figure 13, we take a different approach to viewing the actual lifetime of revoked certificates. In this chart, we take the number of certificates revoked over a two week period and break down the certificates revoked each day by the lifetime of certificate before it was revoked. We observed that almost one third of certificates were revoked within one day after they were issued, and that only about one third of certificates enjoyed lifetime greater than 10 days. Since we do not have access to the reasons why these certificates are revoked, we cannot further interpret this result.

### 3.4 Trends by Organization

From the files collected, 15,341 organizations were identified. However, due to differences in how the company name was placed on the certificate, it is likely that there are fewer than that amount. To make our data as correct as possible, records that had similar names but only differed by punctuation (ex. Verisign Inc. and Verisign, Inc.) were modified and merged into one standard name. We began our analysis by visualizing the amount of certificates used by each organization. Since it would be impossible to list the number of certificates revoked for every company, we decided to instead focus on a smaller portion of the organizations. Figure 14 shows the top 30 organizations by certificates revoked.




Fig. 14. Number of Certificates Revoked by Organization (Top 30)


Another one of the ways we would have liked to examine the data is from a per certificate perspective, judging the distance between when a given certificate is revoked and when the next certificate for the organization is issued. Since each company uses multiple certificates at any given time, it is impossible to determine the average time between when a certificate is revoked and when its replacement is issued. Instead, we will have to use other methods to try to measure the security of an organization. First, we determine the mean certificate lifetime for each company in Table 4.

For these top five companies, the means vary widely. It is interesting to note that the mean lifetime for certificates issued to Bank of America and Citigroup, both financial institutions, differ by almost 45 days. While this may not for certain say that one company is more secure than the other, it does show that these organizations have either mishandled their certificates or possibly have had breaches in their security.

Since these numbers vary so widely, we next decided to fit the lifetime data for each of the above organizations to a probability distribution. We determined that when the data is divided by organization, it still follows the exponential

Organization	Mean Certificate Lifetime (in days)
Bank of America Corporation	15.92
Citigroup	60.32
Ford Motor Company	42.28
JPMorgan Chase	50.00
Hewlett-Packard	14.34

**Table 4.** Mean Certificate Lifetime by Organization



**Fig. 15.** Number of Certificates Revoked vs. Lifetime (By Top 5 Organizations with Revoked Certificates)

distribution, as shown in Figure 15. The parameter  $\mu$  for each organization is given by the mean listed in Table 4, all at a 95% confidence interval.

Clearly, different organizations exhibit different characterizations in terms of their certificate lifetimes. While the certificate lifetimes still follow exponential distributions, the average certificate lifetimes change from organization to organization, even within the same industry group such as financial institutions. If the average certificate lifetimes are treated as a reflection of the organizations' security level or security awareness, those organizations in a competitive market should investigate why their certificates are revoked more or less frequently than their competitors and how to improve their certificate lifetimes at organizational levels. It is imaginable that the publication of more empirical analysis on certificate lifetimes would stimulate organizations to increase their security levels or security awareness, especially in a competitive market.

### 3.5 Discussion on Optimal Management of Certificate Revocation

Our empirical analysis provides a solid foundation for optimal management of certificate revocation for different types of certificates requested from different organizations located in different geographic locations. The reason is that our study enables us to understand the distribution of certificates being newly revoked and the distribution of certificates being cumulatively revoked both on a

daily basis. Given these distributions, a certificate authority (CA) can minimize its operational cost for any type of certificates based on the analytical models proposed in [9], where the CA's operation cost consists of three parts: (i) the expected liability cost per certificate revocation if CA delays publishing the revocation for one day; (ii) the fixed cost for CA to publish a CRL regardless of its size; and (iii) the variable cost for CA to include each individual certificate into a CRL. The CA needs to balance between the liability cost of not releasing CRL on time and the fixed and variable costs of releasing CRL too often for optimal management of certificate revocation.

We should note that the distribution of certificates being newly revoked and the distribution of certificates being cumulatively revoked are not derived directly from empirical data in [9]; instead, they are deduced from the exponential distribution of certificate lifetimes. Consequently, these distributions become constant after the time reaches the issued lifetime. However, as it is shown in our paper, these distributions may fluctuate over time in reality. The analytical models proposed in [9] therefore need to be revised so as to capture this phenomenon.

## 4 Conclusion and Future Work

The certificate revocation is a very complicated issue and is affected by many factors. This paper analyzes the influence of these factors empirically from the Verisign's data. Our research represents the first step towards linking empirical observations to mathematical models in description of the complicated problem of certificate revocation. We have focused on the empirical part in this study. In the future, we plan to conduct extended research on optimal management of certificate revocation based on our empirical analysis. We also hope to conduct a more thorough examination of the per organization data from a larger and more continuous data pool.

## Acknowledgement

The authors would like to thank the anonymous referees for their helpful comments, including a suggestion to revise the title of this paper to be more appropriate.

## References

1. VeriSign certification practice statement, version 3.4. Internet proposed standard RFC 2560, April 2007. <http://www.verisign.com/repository/CPS/VeriSignCPSv3.4.pdf>.
2. David A. Cooper. A model of certificate revocation. In *ACSAC '99: Proceedings of the 15th Annual Computer Security Applications Conference*, page 256, Washington, DC, USA, 1999. IEEE Computer Society.

3. Barbara L. Fox and Brian A. LaMacchia. Certificate revocation: Mechanics and meaning. In *FC: International Conference on Financial Cryptography*. LNCS, Springer-Verlag, 1998.
4. Carl A. Gunter and Trevor Jim. Generalized certificate revocation. In *Symposium on Principles of Programming Languages*, pages 316–329, 2000.
5. R. Housley, W. Ford, W. Polk, and D. Solo. RFC 2459: Internet X.509 public key infrastructure certificate and CRL profile, January 1999. Status: PROPOSED STANDARD.
6. Gaurav Jain. Certificate revocation: A survey. <http://citeseer.ist.psu.edu/511985.html>.
7. Paul C. Kocher. On certificate revocation and validation. In *FC '98: Proceedings of the Second International Conference on Financial Cryptography*, pages 172–177, London, UK, 1998. Springer-Verlag.
8. Ninghui Li and Joan Feigenbaum. Nonmonotonicity, user interfaces, and risk assessment in certificate revocation (position paper). In *Proceedings of the 5th International Conference on Financial Cryptography (FC'01)*, number 2339 in LNCS, pages 166–177. Springer, 2002.
9. Chengyu Ma, Nan Hu, and Yingjiu Li. On the release of crls in public key infrastructure. In *Proceedings 15th USENIX Security Symposium*, pages 17–28, Vancouver, Canada, 2006.
10. Patrick McDaniel and Aviel Rubin. A response to “can we eliminate certificate revocation lists?”. *Lecture Notes in Computer Science*, 1962:245+, 2001.
11. Michael Myers, Rich Ankney, Ambarish Malpani, Slava Galperin, and Carlisle Adams. X.509 internet public-key infrastructure — online certificate status protocol (OCSP). Internet proposed standard RFC 2560, June 1999.
12. Moni Naor and Kobbi Nissim. Certificate revocation and certificate update. In *Proceedings 7th USENIX Security Symposium*, San Antonio, Texas, 1998.
13. Ronald L. Rivest. Can we eliminate certificate revocations lists? In *Financial Cryptography*, pages 178–183, 1998.
14. Stuart Stubblebine. Recent-secure authentication: Enforcing revocation in distributed systems. In *Proceedings 1995 IEEE Symposium on Research in Security and Privacy*, pages 224–234, May 1995.
15. Petra Wohlmacher. Digital certificates: a survey of revocation methods. In *MULTIMEDIA '00: Proceedings of the 2000 ACM workshops on Multimedia*, pages 111–114, New York, NY, USA, 2000. ACM Press.
16. Peifang Zheng. Tradeoffs in certificate revocation schemes. *Computer Communication Review*, 33(2):103–112, 2003.