

# A Response to “Can We Eliminate Certificate Revocation Lists?”

Patrick McDaniel<sup>1\*</sup> and Aviel Rubin<sup>2</sup>

<sup>1</sup> EECS Department, University of Michigan, Ann Arbor [pdmcdan@eecs.umich.edu](mailto:pdmcdan@eecs.umich.edu)

<sup>2</sup> AT&T Labs – Research Florham Park, NJ [rubin@research.att.com](mailto:rubin@research.att.com)

**Abstract.** The massive growth of electronic commerce on the Internet heightens concerns over the lack of meaningful certificate management. One issue limiting the availability of such services is the absence of scalable certificate revocation. The use of certificate revocation lists (CRLs) to convey revocation state in public key infrastructures has long been the subject of debate. Centrally, opponents of the technology attribute a range of semantic and technical limitations to CRLs. In this paper, we consider arguments advising against the use of CRLs made principally by Rivest in his paper “Can we eliminate certificate revocation lists?” [1]. Specifically, the assumptions and environments on which these arguments are based are separated from those features inherent to CRLs. We analyze the requirements and potential solutions for three distinct PKI environments. The fundamental tradeoffs between revocation technologies are identified. From the case study analysis we show how, in some environments, CRLs are the most efficient vehicle for distributing revocation state. The lessons learned from our case studies are applied to a realistic PKI environment. The result, *revocation on demand*, is a CRL based mechanism providing timely revocation information.

---

\* This work was completed at AT&T Labs in Florham Park, NJ as part of the AT&T summer internship program.

## 1 Introduction

The value of the commercial, educational, and personal services Public Key Infrastructures (PKIs) are likely to enable cannot be understated. However, identifying PKI architectures that meet the requirements of even existing services has proven to be difficult. One particularly contentious aspect of PKI design is the mechanism used for distributing certificate revocation information. Public key *certificates* are the vehicle used by an authority to state identity or authorization. The ability of an authority to later UNDO these statements allows longer certificate lifetimes and less exposure to incorrect or compromised certificates. However, revocation is inherently difficult. No solution has been found that meets the timeliness and performance requirements of all applications and environments.

*Certificate revocation* is the act of invalidating the association between the public key and attributes embodied in a certificate. Generally, it is difficult to find revocation solutions that address both the timeliness and performance (resource usage) requirements of all parties. One mechanism, the certificate revocation list (CRL)<sup>1</sup> has received a particular amount of attention. A certificate revocation list is a digitally signed and time-stamped enumeration of all certificates within a domain that have been revoked, but not expired. Therefore, the revocation state of any certificate within the domain can be obtained from a suitably recent CRL.

It has been argued [1–3] at length that CRLs are both semantically and technically inferior to other approaches. This paper is in particular a response to [1], which identifies a majority of arguments present in the literature. We illustrate the positive and negative aspects of CRLs by applying them to three PKI environments. Through these case studies, we show that while CRLs may be sub-optimal in some environments, they adequately address the needs of other (non-trivial) environments.

Some confusion arises from the different terminologies used in PKI literature. Throughout this document, we will refer to certificate issuers as *CAs*, the subject of a certificate as the *principal*, and the party accepting certificates as the *verifier*.

We use the taxonomy presented by Myers in [4] to describe the current revocation design space. Myers identifies four classes of revocation mechanisms; *CRLs*, *trusted dictionaries*, *online*, and *short lifetime certificates*.

In systems supporting CRLs [5–8], the revocation state for all certificates within a domain is announced in a singular periodic statement. Thus, once a verifier has determined the revocation state of a certificate, she knows *a priori* the revocation state of all other certificates within the same domain. There are a number of mechanisms that allow the costs of traditional CRLs to be mitigated [6, 9–12].

---

<sup>1</sup> Throughout, we use the term CRL to represent any scheme in which revocation information is distributed through periodically generated statements encompassing all certificates within a domain.

Trusted dictionaries [10, 13, 14] provide pre-generated proofs of revocation state. Verifiers obtain the state for each certificate independently, subject to the periodicity of proof generation, application requirements, and verifier policy.

In *On-line* approaches, [15] proofs of a certificate's (non) revoked state are generated and distributed in real-time. Thus, each re-assertion of a certificate's validity is handled individually and potentially independently of others. Other approaches [16–18] provide an on-line protocol for initial retrieval, specifying a *time-to-live* during which the retrieved certificate may be used without more recent validity information.

Typically, architectures not supporting revocation issue certificates whose lifetimes are short. Because exposure is small, there is less of a need for revocation. Short term certificates are semantically identical to short term symmetric key associations (e.g. [19] Kerberos tickets).

For flexibility, a number of systems provide multiple mechanisms for distributing revocation state. [20, 21].

The service provided by these revocation mechanisms is similar. Within a known timeliness bound, the verifier is able to obtain a proof of certificate's revocation state. Presumably, this information will help determine the appropriateness of a certificate for some use. Note that we specifically do not address the meaning of a certificate revocation [22]. Revocation reason is a central determinant in the processing of revoked certificates, and is typically left to application/verifier policy. This paper addresses the *mechanism* used to distribute the revocation state.

A central policy issue is the allowable length of time between a statement of validity and the use of the certificate. This policy defines the amount of exposure to a revoked certificate the verifier is willing to tolerate. Any number of factors may contribute to this policy; the type of transaction the certificate is to be used for, the process in which the certificate was acquired, or simply as a function of the trust held in the certificate owner or issuing CA.

We assert a central tradeoff of these approaches is between performance and timeliness. Clearly, obtaining revocation state for a single certificate using CRLs is more costly than other approaches. However, as the reference locality rises (certificates from a single authority are used), so do the advantages of CRLs.

An often stated objection to CRL based mechanisms is that they do not provide near real-time revocation state. This statement assumes PKI users are not willing to accept any exposure to revoked certificates. Secondly, it assumes it is impossible to achieve or it does not make sense to have real-time CRLs. We believe these assumptions are based on pre-conceptions about the uses and environment in which PKI systems are to be deployed.


In the remainder of this document, we analyze the classes of revocation mechanisms in an attempt to uncover the salient features of CRLs. We demonstrate how, in some environments, CRLs are the most efficient vehicle for distributing revocation state.

## 2 Certificate Revocation Lists

Recently, a number of arguments advising against the use of CRLs have been advanced [1–4]. While these arguments are compelling, further investigation of their assumptions and foundations is warranted. We distill the majority of these arguments in the following propositions:

1. As the verifier is the party assuming risk, he should have control over the recency guarantees [1]. CRLs require the verifier to accept a guarantee bounded by the rate at which CRLs are generated. Thus, in CRLs, the recency guarantees are always under the control of the CA (or party generating CRLs).
2. For efficiency, the principal should supply all relevant validity evidence [1]. Thus, principals must acquire or generate all the appropriate proofs of revocation state for each transaction.
3. The demand for “high-value” transactions necessitates the availability of online revocation mechanisms [4, 3]. While this assertion does not directly argue against the use of CRLs, it implies other mechanisms (with better timeliness guarantees) must also be supported. This argument is based on two assumptions; a) there are inherent latencies in any solution using CRLs, and b) “high-value” transactions are commonplace. As defined in [3], a transaction is deemed “high-value” if the relying parties’ policy requires real-time revocation state.
4. The cost of CRL management and distribution is too high [10, 2]. Because of the potential size of CRLs, scaling to large communities can be difficult. This is a commonly cited argument.
5. CRLs are inappropriate for transactions that require real-time revocation state [3]. That is, the inherent costs of CRL generation and distribution prohibit online CRL generation.
6. CRLs do not provide a positive response [4]. Because CRLs only identify revoked certificates, the existence of a (non-revoked) certificate cannot be determined solely from validity information.
7. New certificates are the best evidence of recency [1]. If a (new) certificate with a guaranteed validity period is available, then the acceptance process may be reduced to the validation of a single certificate signature. As the revocation state is implied by the existence of the certificate, CRLs are unnecessary.
8. Certificates in traditional CRL based schemes do not have any inherent recency information other than the certificate lifetime [1]. Thus, each time a certificate is accessed, the verifier is *required* to obtain and validate a suitably recent CRL. Combined with proposition 7, this makes a strong argument for the use of online revocation mechanisms [15].

In general, these propositions state that CRLs are limited by mechanism and performance. More precisely, they state that “CRLs cannot provide the required service” and “the service CRLs provide is too costly”. Note that the service is defined by application and environmental requirements. Without an understanding of the range of possible requirements, it is difficult to make general statements about the applicability of CRLs.



**Fig. 1.** electronic commerce PKI - Customers authenticate vendor sites. Revocation state is retrieved from a limited number of authorities.

In the following sections, we investigate the correctness of these propositions when applied to several PKI environments. Furthermore, the ways in which CRLs may be adapted to address performance and security requirements are investigated. Finally, we identify heuristics for the use of revocation mechanisms.

### 3 PKI Analysis

In this section, we analyze CRLs by looking at the requirements of distinct PKI environments. The selected environments represent three important classes of applications commonly associated with certificate based authentication; electronic commerce, intranet services, and Internet mail. These environments are representative of the types of workloads that future PKIs are likely to encounter.

#### 3.1 Electronic Commerce

PKIs supporting electronic commerce enable transactions between vendors and customers. Vendors act as servers and customers as clients. The client, acting as the verifier, initiates transactions by authenticating the server (typically) via a challenge-response protocol [23]. However, the client is not typically authenticated. In the normal case, the payment channel (credit card) provides sufficient authentication for the vendor. As is true for most CA based PKIs, CAs state the validity of certificates through digital signature and distributed revocation state. This architecture is described in figure 1.

<sup>2</sup> The PKI described in this section enables typical web-site vendor to end user customer transactions. We note the existence of other PKIs present on the Internet

During the transaction, the client is depending on the validity of the server certificate to protect her payment channel. The client risk is directly determined by her liability to the exposure of that channel. In most cases, clients have a maximum liability for the loss of credit cards. Conversely, the server risks its reputation. Customers are unlikely to purchase goods from vendors who have historically unsafe operation. A single publicized compromise of the private key can irreparably damage an electronic business. Although the risk is less tangible for servers than for the clients, it may be significantly higher.

Recall that proposition 1 states,

As the verifier is the party assuming risk, he should have control over the recency guarantees. [1]

Risk in our model of electronic commerce is not clearly greater for verifiers. Thus, for this environment, the proposition does not hold.

We note in today's Internet there are few widely used authenticating bodies (CAs). For example, the Netscape Communicator [24] version 4.51 ships with the certificates of 42 CAs. The number of servers is large, but is significantly smaller than the number of clients. Potentially, revocation state for millions of certificates needs to be distributed (either directly or indirectly) by a few authorities to tens or hundreds of millions of users. CAs are clearly heavily loaded in this environment.

A central reason revocation is not currently supported in commercial transactions on the Internet is performance; scaling existing mechanisms to the Internet is prohibitively expensive. Recall that proposition 2 states,

For efficiency, the principal should supply all relevant validity evidence [1].

The vendors (principals) are likely to be heavily loaded. Requiring vendors to obtain and distribute revocation state only exacerbates the existing performance problems. Thus, for electronic commerce applications, this proposition does not hold. In this case, verifiers are more likely to have the available resources for obtaining revocation state.

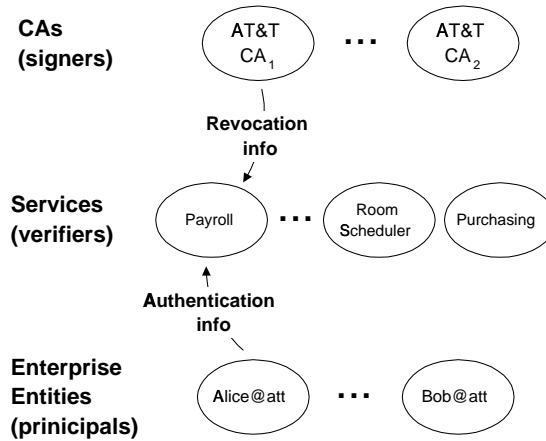
There are a number of known techniques that reduce or distribute the cost of supporting certificate revocation. Most frequently, an authority delegates the revocation duties to other services. Thus, the private key used to sign certificates need not be used for revocation. This has the advantage that the compromise of revocation service does not compromise the CA.

Online approaches require the CA to generate a digital signature for each request. In environments where even modest loads can be observed, the CA quickly can become compute bound. Thus, replication of the CA or delegation of the revocation responsibilities becomes necessary.

CRL based mechanisms avoid much of the costs associated with signature generation in the critical path of the transaction. However, because the size of the

---

(e.g. used for business to business electronic commerce). However, the risk, usage patterns, and trust relationships of these environments can differ significantly from our model of electronic commerce. Hence, the analysis presented in this section may not directly apply to these environments.



**Fig. 2.** intranet service PKI - Services authenticate enterprise entities. Revocation state is retrieved from enterprise local authorities.

CRL is potentially large, the cost of retrieval can consume significant bandwidth and introduce long latencies. This demonstrates a chief performance tradeoff between online and traditional CRL mechanisms; CPU cost vs. bandwidth.

Trusted dictionary approaches [10, 13, 14] can be used to meet the requirements of electronic commerce applications. These approaches avoid both the signature generation costs of online revocation and the distribution costs of CRLs. The advantages of the performance and timeliness compromise found in trusted dictionaries has led to the adaptation of certificate revocation trees [10] in several commercial applications.

An interesting question is, “Is real-time timeliness a requirement of commercial transactions on the Internet?”. Based on risk, is it reasonable to assume the participants are willing to accept five minute latency? An hour? More? Clearly, the lack of a revocation mechanism in today’s electronic commerce infrastructure has not significantly limited its acceptance. Citing current infrastructure use as evidence, it can be inferred that most electronic commerce transactions do not fit the definition of “high value” presented in [3]. Recall that Proposition 3 states;

The demand for “high-value” transactions necessitates the availability of online revocation mechanisms [4, 3].

We assert that this proposition does not apply to vast majority of electronic commerce transactions. Customers and vendors are willing to accept the (short) timeliness guarantees provided by other, less costly, revocation mechanisms.

An important aspect of all revocation mechanisms approaches is availability; relying on a server to distribute revocation state introduces a single point of failure. Where justifiable, providing multiple, independent sources of revocation state seems prudent.

### 3.2 Intranet Service

Certificates can be used as the mechanism for client authentication in intranet information services. In the model presented in this section, a *service* provides useful content to clients within enterprise internal networks. The clients, typically employees, authenticate themselves to the service before being allowed access to the service content. Unlike in the electronic commerce example, the clients act as principals and the servers act as verifiers. An intranet service architecture is depicted in figure 2.

Since all of the principals exist within a single administrative domain, the certificates may be serviced by a small number of CAs. However, we cannot assume the CA workloads are manageable by singular hosts. For example, AT&T has over 126,000 employees, any one of which can be the principal in a number of certificates. Recall Proposition 4 states

The cost of CRL management and distribution is too high [10, 2].

This proposition does not hold in this environment. There are a small number of verifiers and fewer CAs. Thus, the acquisition or subsequent validation of CRLs should not present a significant burden on the enterprise network infrastructure. We investigate how this particular feature can also be used to reduce latency below. Myers identifies CRLs as a potential solution for similar, albeit smaller, environments in [4].

Certificate usage in these services exhibits the one characteristic that makes CRLs attractive; reference locality. Because the certificates are issued from a small number of CAs, we can obtain recent revocation state for many certificates simultaneously. Moreover, the obtained revocation state is likely to be useful over many transactions.

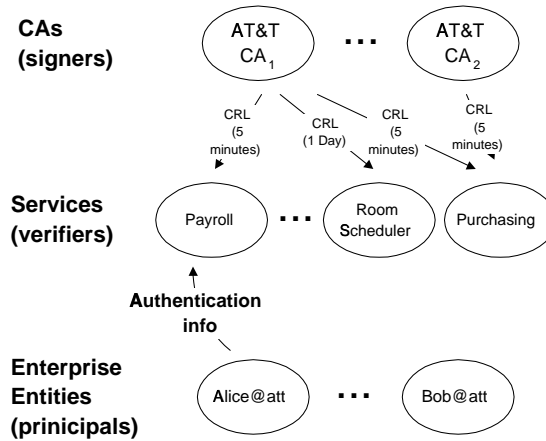
The value of the service content directly determines risk for both the clients and the service. If the service allows access to the direct-deposit or salary information, then it is important that the validation process be strong. If however, the service provides an interface to conference room scheduling, less diligence is necessary. This is another example of a fundamental axiom of security; the protection need only be as strong as the value of what it protects [25].

In this model, the services (verifiers) are likely to be the most heavily loaded entities. Each server must perform certificate validation, user authentication, and service itself. Furthermore, because services may be visited frequently, there is economic motivation for reducing the latency of the certificate validation process.

As transaction value is the determinant of risk, it also should determine the timeliness requirements. In some enterprises, vast sums of money or stocks are transferred using local services. Clearly, these transactions should meet the “high-value” definition. Proposition 1 holds for these services; the verifier is in the best position to assess risk, and as such should have control over the recency guarantees. Risk among services is not uniform; some services have stronger timeliness requirements.

We now introduce a CRL-based solution addressing the requirements of the intranet service. *Revocation on demand* (ROD) uses a publish/subscribe [26]





**Fig. 3.** Revocation On Demand (ROD) - Verifiers subscribe to a CRL delivery service from each CA in which they are interested. The CAs generate and deliver CRLs per a schedule commensurate with verifier subscription requests. Thus, the subscribers always have a suitably recent CRL. This removes the certificate validation process from the critical path of transactions.

mechanism for CRL delivery. In this approach, verifiers subscribe to a CRL service associated with each CA in which they are interested. Verifiers state the rate at which they wish to receive CRLs during the subscription process. Afterward, CAs generate and deliver CRLs in accordance with the subscription requests. We describe this approach in figure 3.

The rate at which the CRLs are delivered is limited by the speed at which CRLs can be generated and delivered. Due to the near constant cost of signature generation, the speed of today's networks, and the limited number of verifiers, ROD can provide timeliness guarantees that are essentially equal to those provided by online protocols. Recall that proposition 5 states,

CRLs are inappropriate for transactions requiring real-time revocation state [3].

Because of the characteristics of the environment in which it is deployed, ROD can deliver near real-time revocation state. Thus, the proposition does not hold for this environment and mechanism.

Note the subscriber approach removes all revocation related operations from the critical path of any transaction. Verifiers never need to wait for the retrieval of revocation state. If the verifier uses sufficient certificate caching, many transactions may be completed without the direct involvement of a CA.

In using CRLs, we take advantage of the environment's natural reference locality. Because the certificates are issued by a small number of CAs, the probability that retrieved revocation state will be useful in later verification is high.

CRL delivery has some constant bandwidth cost. As the number of verifiers grows, so does the the CA's bandwidth consumption. We present several ap-

proaches that may be used to mitigate these costs. First, we could reduce the size of the delivered CRLs. In delta CRLs [6], infrequently generated “base” CRLs are generated. The more frequently generated delta CRLs indicate only those certificates that have been revoked since the last base CRL. In a related approach, one can limit the period over which revocation is reported (e.g. Windowed Revocation [11]).

A second approach is to deliver CRLs via multicast. Using a tiered quality of service approach similar to [27], one could provide channels delivering CRLs at several rates. This is closely related to freshest CRLs [9]. However, because of the unreliable nature of multicast, some additional protocol engineering may be necessary.

CAs may wish to avoid placing their private key on hosts connected to the Internet. However, because of timeliness requirements, the key used to generate the CRLs must reside on a highly available host. Thus, it may be advantageous to separate a CA’s certificate issuance and CRL generation duties into separate services using different keys.

The revocation on demand architecture supports the X.509v3 distribution point extension [6]. Distribution points are used by CAs to delegate CRL generation duties. CAs in our approach may delegate CRL generation to one or more distribution points. This may lead to a more efficient design; each verifier may receive all pertinent revocation state through a single CRL. Also, the overhead associated with the reception and processing of multiple CRLs may be avoided through CRL aggregation.

It has been claimed that the information embodied in a CRL is limited. Recall that proposition 6 states,

CRLs do not provide a positive response [4].

As Myers suggests in [4], the existence of a certificate can not be determined from the serial number and CRL alone. We believe an existence proof service is fundamentally different from current definitions of revocation. Thus, precluding the use of CRLs based on this argument does not seem warranted. If such a service were required, altering CRL specifications to include valid identifier ranges (instead of serial numbers) is trivial. We state that CRLs in ROD supports both explicit serial numbers and identifier ranges.

### 3.3 Electronic Mail

Email has become a primary medium over which parties on the Internet communicate. Thus, a PKI supporting electronic mail should be able to establish authentication between arbitrary endpoints. Verifiers may or may not know anything about the principals or their authenticating bodies. Given this definition, providing certificate services within a global environment seems intractable.

Early attempts to project a global authentication framework on the Internet have failed. This is due to the intransitivity of trust, the difficulty in finding a set of entities in which all users trust, and a myriad of other technical, political, and social issues.

In response to the failure of global approaches, various groups have introduced infrastructures constructed within independent communities. Some approaches [28, 18, 20] construct interconnections mirroring trust derived from personal relationships. This approach generally leaves the certificate acceptance process to the user. Other proposals adapt the hierarchical approaches to enterprises [7].

To be widely accepted, authentication frameworks should model the social environments in which they operate. History has shown that while global PKIs (e.g. PEM [5]) are not readily accepted, approaches whose trust model is derived from the supported community (e.g. webs of trust [28]) are more successful. The success of ICE-TEL [7] system further demonstrates the connection between underlying social structures and PKI acceptance.

The ICE-TEL [7] system was designed to support loose interconnections of highly structured local domains. The separate local domains were, at the administrative level, aware and trusted each other. Thus, the interconnections were a physical manifestation of trust that already existed. Within each domain the certificate services mirrored the trust embodied in the enterprise structure; users (employees, students, ...) trusted a hierarchy of local authorities.

ICE-TEL is comprised of previously existing, but not widely accepted, technologies. It can be inferred from the ICE-TEL experience that the success of a PKI is not completely defined by its underlying mechanisms, but also from its connection to the population that it supports. Many PKIs that adequately address their environmental requirements have not been accepted because of a failure to model real world trust.

Because of a lack of real global trust, it is unlikely that a global PKI will ever be successful. Any architecture projecting a structure on the Internet would embody a trust that simply does not exist.

Because of the differences of the communities using electronic mail, it is unlikely that any one PKI (or revocation mechanism) will be used in all environments. We expect the independent communities will continue to deploy a range of PKI architectures. By necessity, the independent communities will interconnect through well known, but not necessarily trusted, gateways. Existing PGP public key servers currently provide a gateway service.

Similarly, deployed revocation mechanisms will be tailored to the PKIs in which they operate. CRLs will be used in environments in which they are suited, and other techniques where they are not.

## 4 Short Term Certificates

In [1], Rivest asserts that frequent certificate re-issuance provides the best evidence of recency. A recently issued certificate is efficient; it provides enough information to determine both authenticity and validity. Re-issued certificates reduce the possibility of error by avoiding misinterpretation or falsification of mappings between certificate serial numbers and revocation state. However, certificate re-issuance also has inherent costs.

Re-issuance is a CPU intensive operation. Where CAs are heavily loaded, the cost of re-issuance may be prohibitive. Recall that propositions 7 and 8 state, New certificates are the best evidence of recency [1].

**and**

Certificates in traditional CRL based schemes do not have any inherent recency information other than the certificate lifetime [1].

While propositions 7 and 8 may be true, providing short term certificates in some environments is infeasible. One must weigh the advantages of short term certificates against performance issues.

An interesting feature of the short term certificates defined by Rivest is the *guaranteed* period. A guaranteed period is a CA defined period during which the certificate is necessarily valid. The guaranteed period represents a contract between the CA and verifiers. The contract states, for the guaranteed period, the CA will not revoke the certificate for any reason. This has the unique advantage that the CA need not be contacted until the certificate expires. Because the CA does not have control over certificate compromise, additional infrastructure is required. The proposed approach defines a *suicide bureau* that distributes (online) positive statements of certificates' non-compromised status.

In conjunction with short term certificates, the guaranteed period can be used to greatly reduce the cost of revocation. Because compromise is the only reason these certificates are revoked, we eliminate the costs associated with administrative revocation. As short term certificates are used, the time over which a compromise needs to be reported is limited. Windowed revocation [11] uses a similar mechanism to reduce the period during which revocation is announced.

## 5 Conclusions

Throughout, we have investigated the applicability of recent arguments against the use of CRLs in a range of PKI environments. We note that while these arguments are true for certain classes of applications, CRLs provide a useful and efficient service for others.

We assert that the need for real-time revocation state is not present in a large portion of current Internet transactions. Certificate based electronic commerce has grown immensely in the absence of widely used revocation mechanisms. The requirements of timeliness can be met with short, achievable, periods using any number of revocation techniques.

CRLs are most suited to tightly coupled environments where reference locality can be observed. This is best demonstrated in service oriented environments, where the services must authenticate many users from a limited number of CAs. However, other mechanisms may be more efficient in environments with many CAs.

It is possible to achieve near real-time revocation state using CRLs. Using the publish/subscribe *revocation on demand* mechanism, CRLs can be generated

and delivered to a limited number of verifiers with minimal latency. Moreover, the timeliness can be tailored to meet the differing requirements of many verifiers simultaneously.

Because of the lack of global trust, we believe finding a general purpose, fully automated, global authentication framework is intractable. Thus, in the future, we expect the certificate and revocation services will mirror the social structure of the communities which the service, leading to loosely connected islands of independent PKIs.

The answer to Rivest question, “Can we eliminate certificate revocation lists?”, is both yes and no. CRLs are clearly the wrong mechanism for a large class of PKI environments. Addressing PKI requirements in large, loosely coupled environments using CRLs is difficult. However, in other environments, CRLs are a useful tool for limiting the costs associated with revocation.

Ultimately, the the design of a revocation mechanism must be driven by the applications it supports. Much of the arguments for and against particular revocation technologies, while correct, are derived from assumptions made about the target environments. Thus, while these arguments provide good design heuristics, they do not apply to all environments.

## 6 Acknowledgements

We would like to thank Rebecca Wright and Sugih Jamin for their many helpful comments. We would also like to thank Carl Ellison for his advice and perspective on Public Key Infrastructure technologies and environments.

## References

1. Ronald L. Rivest. Can We Eliminate Certificate Revocation Lists? In Rafael Hirschfeld, editor, *Financial Cryptography FC '98*, volume 1465, pages 178–183, Anguilla, British West Indies, February 1998. Springer.
2. J. Millen and R. Wright. Certificate Revocation the Responsible Way. In *Post-Proceedings of Computer Security, Dependability, and Assurance: From Needs to Solutions*, volume ix, pages 196–203. IEEE, 1999.
3. B. Fox and B. LaMacchia. Online Certificate Status Checking in Financial Transactions: The Case for Re-issuance. In Rafael Hirschfeld, editor, *Financial Cryptography FC '99*, volume 1648, pages 104–117, Anguilla, British West Indies, February 1999. Springer.
4. M. Myers. Revocation: Options and Challenges. In Rafael Hirschfeld, editor, *Financial Cryptography FC '98*, volume 1465, pages 165–171, Anguilla, British West Indies, February 1998. Springer.
5. S. Kent. Internet Privacy Enhanced Mail. *Communications of the ACM*, 36(8):48–60, August 1993.
6. R. Housley, W. Ford, W. Polk, and D. Solo. RFC 2459, Internet X.509 Public Key Infrastructure Certificate and CRL Profile. *Internet Engineering Task Force*, January 1999.
7. D. Chadwick and A. Young. Merging and Extending the PGP and PEM Trust Models - The ICE-TEL Trust Model. *IEEE Network*, 11(3):16–24, May/June 1997.

8. P. McDaniel and S. Jamin. A Scalable Key Distribution Hierarchy. Technical Report CSE-TR-366-98, Electrical Engineering and Computer Science, University of Michigan, July 1998.
9. C. Adams and R. Zuccherato. A General, Flexible Approach to Certificate Revocation, June 1998. <http://www.entrust.com/securityzone/whitepapers.htm>.
10. P. Kocher. On Certificate Revocation and Validation. In Rafael Hirschfeld, editor, *Financial Cryptography FC '98*, volume 1465, pages 172–177, Anguilla, British West Indies, February 1998. Springer.
11. P. McDaniel and S. Jamin. Windowed Certificate Revocation. In *Proceedings of IEEE Infocom 2000*. IEEE, March 2000. Tel Aviv, Israel. (to appear).
12. P. Hallam-Baker and W. Ford. Internet X.509 Public Key Infrastructure - ENHANCED CRL DISTRIBUTION OPTIONS. *Internet Engineering Task Force*, August 1998. (draft, expired) [draft-ietf-pkix-ocdp-01.txt](#).
13. S. Micali. Efficient Certificate Revocation. Technical Report Technical Memo MIT/LCS/TM-542b, Massachusetts Institute of Technology, 1996.
14. M. Noar and K. Nassim. Certificate Revocation and Certificate Update. In *Proceedings of the 7th USENIX Security Symposium*, pages 217–228, January 1998.
15. M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. RFC 2560, X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP. *Internet Engineering Task Force*, June 1999.
16. J. Galvin. Public Key Distribution with Secure DNS. In *Proceedings of the 6th USENIX Security Symposium*, pages 161–170, July 1996.
17. D. Eastlake and C. Kaufman. RFC 2065, Domain Name System Security Extensions. *Internet Engineering Task Force*, January 1997.
18. R. Rivest and B. Lampson. SDSI A Simple Distributed Security Infrastructure, October 1996. <http://theory.lcs.mit.edu/~rivest/sdsi11.html>.
19. B. C. Neuman and T. Ts'o. Kerberos: An Authentication Service for Computer Networks. *IEEE Communications*, 32(9):33–38, September 1994.
20. C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. RFC 2693, SPKI Certificate Theory. *Internet Engineering Task Force*, September 1999.
21. C. Adams and S. Farrell. RFC 2510, X.509 Internet Public Key Infrastructure Certificate Management Protocols. *Internet Engineering Task Force*, March 1999.
22. B. Fox and B. LaMacchia. Certificate Revocation: Mechanics and Meaning. In Rafael Hirschfeld, editor, *Financial Cryptography FC '98*, volume 1465, pages 158–164, Anguilla, British West Indies, February 1998. Springer.
23. T. Dierks and C. Allen. RFC 2246, The TLS Protocol Version 1.0. *Internet Engineering Task Force*, January 1999.
24. Netscape Corporation. Netscape Communicator, 1999. <http://www.netscape.com/>.
25. D. Kahn. *The Codebreakers*. Macmillan Publishing Co., 1967.
26. R. Wilhelm. Publish and Subscribe with User Specified Action. In *Patterns Workshop, OOPSLA '93*, 1993.
27. S. McCanne, V. Jacobson, and M. Vetterli. Receiver Driven Layered Multicast. In *Proceedings of ACM SIGCOMM '96*, pages 117–130. Association of Computing Machinery, September 1996.
28. P. Zimmermann. PGP User's Guide. Distributed by the Massachusetts Institute of Technology, May 1994.