A General, Flexible Approach to Certificate Revocation

Dr. Carlisle Adams & Dr. Robert Zuccherato
Entrust Technologies

1. Introduction

Since public key certificates have a relatively long lifetime, the information they contain can become invalidated during their lifetime with a significant probability. Therefore, checking certificate revocation information is necessary. The standard method of conveying this information is the Certificate Revocation List (CRL). X.509 v1 CRLs have some problems, however. In order to deal with these problems, variants and alternatives to v1 CRLs have been proposed. The problems and their proposed solutions include the following.

- Since an X.509 v1 CRL issued by an authority must include all valid certificates issued by that authority that have been subsequently revoked, it can become excessively large. Its size is directly proportional to the size of the subscriber community, the lifetime of the certificates, and the probability density of a revocation. The network resource consumed by the distribution of the CRL is directly proportional to the size of the relying party community, the size of the CRL, and its frequency of update. If the subscriber community and the relying party community are the same, then the consumption of network resource is proportional to the square of the size of this community; hence this approach is impractical for large communities. In order to correct this situation, the following solutions have been proposed:
  - CRL Distribution Points, as specified for X.509 v2 CRLs, fragment the full set of certificates issued by the authority into sub-sets, so that each fragment can have its own smaller CRL. Each X.509 v3 certificate has a pointer to the CRL fragment where its revocation status is indicated;
  - The Online Certificate Status Protocol (OCSP), which is currently the subject of an Internet Draft, provides users with revocation information for individual certificates. Thus, users do not receive information about certificates they have no immediate interest in.
- In order to contain the network resource consumed by the distribution of X.509 v1 CRLs, they are typically updated relatively infrequently. Thus, it is difficult to provide timely, fresh revocation information. The following solutions were developed to address this issue:
  - Delta CRLs use a base CRL or CRL distribution point which may be issued relatively infrequently. The Delta CRLs are then issued more frequently and only contain updates to the base CRL. Since they are small, they can provide timely information without unduly consuming network resources;

    - OCSP can be used to provide timely information as well. Since OCSP responses must be signed in order to provide the necessary evidence, the scalability of this protocol when providing timely information is doubtful.
- When there is only one CRL, it can be produced by only one entity. The number of certificates or the frequency of CRL issuance can become too large to be handled by one entity.
    - Indirect CRLs are CRL Distribution Points that are signed by other entities. This allows a more scalable solution.
- Once a strategy employing X.509 v1 CRLs (or Distribution Points, Delta or Indirect CRLs) has been implemented, it is impossible to change the partitioning strategy. Unexpected revocation behaviour then becomes difficult to deal with.
    - The Open CRL Distribution Point proposal attempts to solve this problem by providing unsigned pointers to CRLs that contain a "scope" statement, describing the range of certificates that the distribution point covers. The use of unsigned pointers, however, leaves this proposal open to denial-of-service attacks.

Another criticism of CRLs that has been raised is that they do not contain a positive response (because the absence of a certificate from a CRL indicates that it isn't revoked). This is simply a matter of efficient encoding, based upon the expectation that, at any given time, a higher proportion of valid certificates will not be revoked than are revoked. Certificates and CRLs contain the information necessary to identify the appropriate CRL fragment, and the existence of the appropriate CRL that does not list the certificate in question as being revoked is evidence that the certificate has not been revoked.

Unfortunately none of the alternatives listed above is fully satisfactory because none of them solves all of the problems associated with v1 CRLs. So the question remains: "Is there a single method of providing certificate revocation information that solves all of the perceived problems with X.509 v1 CRLs?" This document will attempt to answer this question.

2. Fresh Revocation Information

There is a cost associated with obtaining timely revocation information. This cost must be balanced with the risk of using less timely information. Some applications may be content with revocation information that is a day old. Others require revocation information that is more-or-less instantaneous. It seems reasonable to assume that, if a relying party requires timely revocation information, then it will be willing to pay for this "upgraded" service in the form of increased processing time and/or network costs. Relying parties that have a more relaxed freshness requirement should not be required to incur increased costs. Alternatively, the authority may insist that only its freshest revocation information should be accepted.

Since there are different levels of service that can be offered, it would seem natural that there should be different methods of providing this service to relying parties depending

on their needs. For this reason we propose the use of two different certificate extensions: the familiar CRL Distribution Point and a new Freshest Revocation Information Pointer. Authorities may populate certificates with none, one or both of these extensions.

Under this proposal, no changes need to be made to the CRL Distribution Point as defined for v2 X.509 CRLs. It points to a CRL fragment that contains revocation information for the certificate in question. These CRL fragments can be issued on a regular schedule. The schedule can be tailored to the needs of the relying party community and the limitations of the CA. The CRL fragments could be distributed by LDAP, HTTP, FTP, etc., so no additional trusted parties are required. The information in the certificate and the corresponding CRL Distribution Point ensure acceptable evidence of certificate revocation.

Relying parties that require fresher information can then use the Freshest Revocation Info Pointer extension. This pointer would point to the freshest source of revocation information that this particular CA is prepared to provide, in the form of a Freshest CRL. The Freshest CRL is just a CRL that is updated more frequently than the usual CRL Distribution Point. The issuance interval for a Freshest CRL need not be fixed, and can be changed when needed to provide more timely or relevant revocation information. In fact, the system can be engineered so that the issuance interval for Freshest CRLs is less than the administrative time required for processing a certificate revocation request.

In a typical implementation using both extensions, the Freshest CRL would be a Delta CRL having, as its base, the CRL fragment pointed to by the CRL Distribution Point. The full fragment is obtained by combining the Delta CRL with the base CRL. Thus, the Delta CRL would be very small and could provide very timely revocation information. We illustrate this situation below:

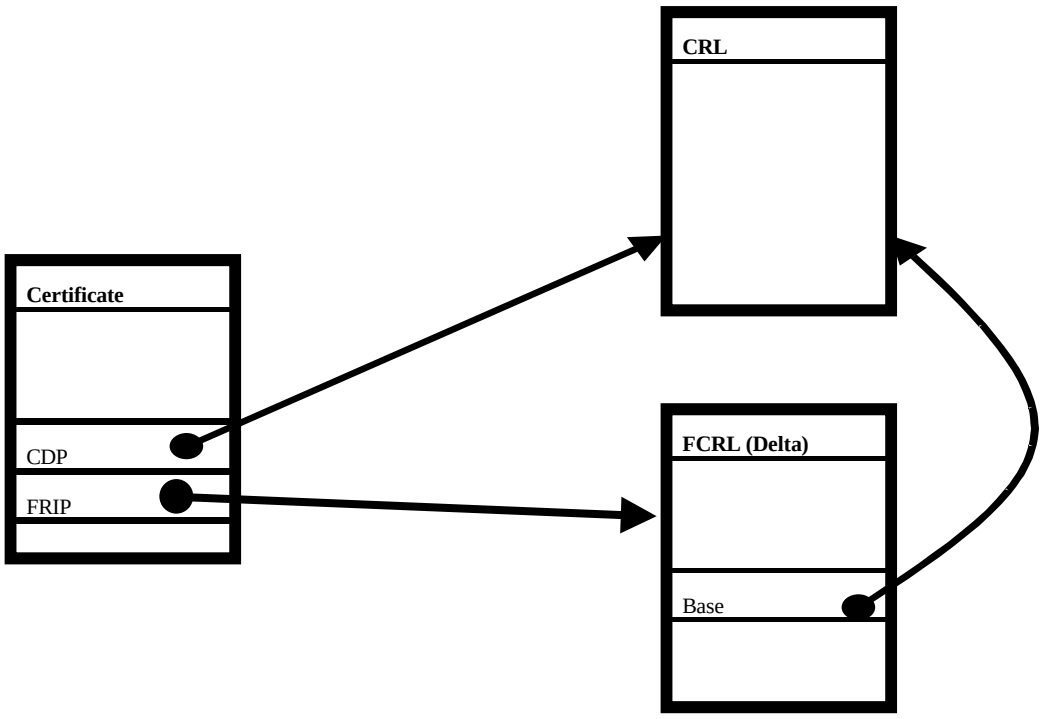

Thus, relying parties that do not require up-to-date revocation information can process the CRL Distribution Point, whereas applications that do require up-to-date revocation information can process the Freshest CRL. There is a cost associated with accessing the up-to-date information. Another CRL (the base) needs to be requested and processed and the Freshest (Delta) CRL is very unlikely to be cached. Since the Delta CRLs and base CRLs can be issued on a fixed schedule, however, the authority can control its cryptographic workload.

A key point of this proposal is that it uses existing standards, simply defining a new extension for the X.509 v3 Certificate. It also provides flexible control over the parameters that affect performance.

3. Controlling the Size of CRL Fragments

In the previous section a method was proposed which makes use of CRL Distribution Points. Distribution Points allow for partitioning of the user space into fragments that can be efficiently processed as a single data object. For most environments, this segmentation of the user space will result in CRLs that do not grow unmanageably large.

The Open CRL Distribution Points proposal has however introduced an interesting new concept, that of dynamically updating the space partitioning. In some environments the actual revocation behaviour may be drastically different from that which was expected. This may result in some CRL fragments being excessively large. Highly constrained devices which may not be able to handle CRL fragments larger than a certain limit may find this a particularly difficult problem.

For those circumstances in which this is a requirement, we propose a new CRL extension called a Redirect Pointer. This extension will contain a set of scope statements and their associated pointers. Each scope statement will indicate a range of certificates that are covered by the CRL fragment indicated by the pointer. For example, if a CRL Distribution Point covers serial numbers 1000 to 1500, one of the scope statements in a Redirection Pointer contained in that Distribution Point may indicate serial numbers 1100 to 1200. The pointer then points to another CRL fragment that contains the revocation information for all certificates within that range. Thus, after a PKI has been established, and certificates have been issued with a particular distribution point extension, it is still possible to re-partition the distribution points if, for whatever reason, the CRL for that distribution point grows to be unmanageably large.

Redirect Pointers would be used in Redirect CRLs. Redirect CRLs are simply CRLs that contain no entries but do contain the (critical) Redirect Pointer extension. The Redirect CRL defines new scope rules and the corresponding location for the new CRL Distribution Point or Freshest CRL. In most instances a Redirect CRL would not be required and so the penalty associated with processing them will not be incurred, but if circumstances dictate re-partitioning, then a solution is available.

We illustrate one possible implementation of this proposal in the following diagram:

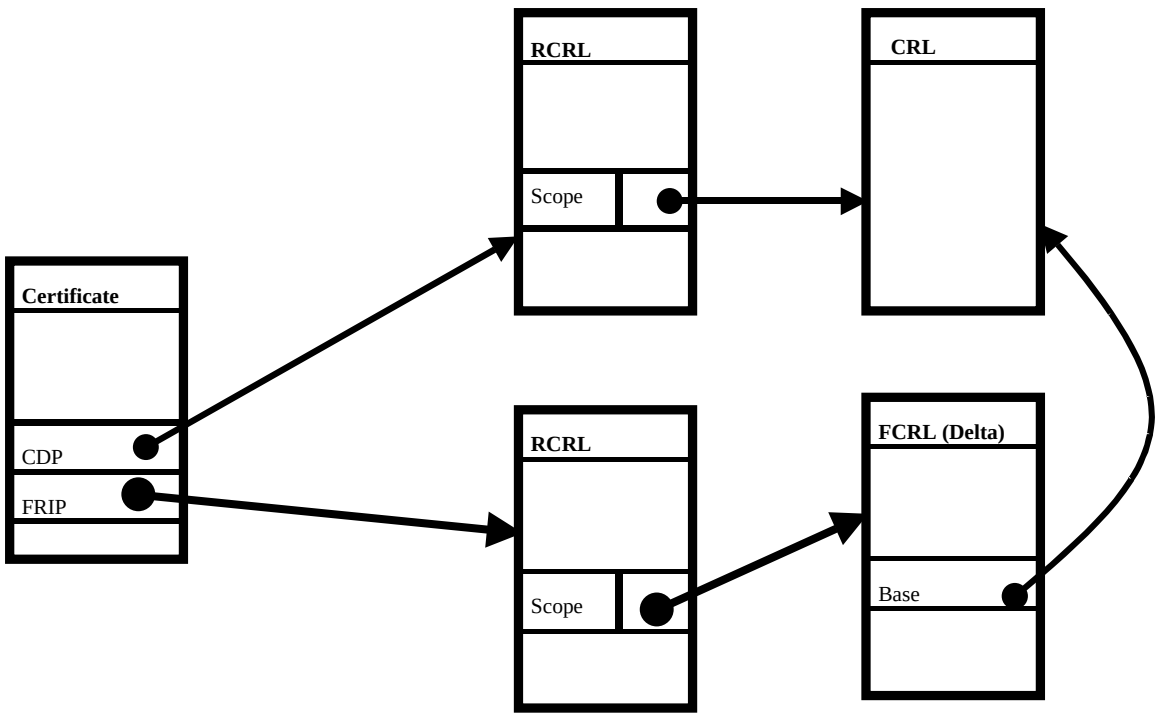

This proposal borrows heavily from the Open CRL Distribution Point proposal. It differs, however, in that the redirection information is signed and thus the user is less open to denial of service attacks.

Again, this proposal uses existing standards, simply defining a new extension for the X.509 v2 CRL. It also provides flexibility and the associated cost is incurred only when absolutely necessary.

4. Third-Party Service Providers

The X.509 v2 CRL Issuing Distribution Point extension may contain an Indirect CRL indicator that allows the CRL Distribution Point to be issued by a third party service provider. In most applications this facility will not be needed, but it may be useful in large PKIs in which revocation is administered at a point in the organization which is remote from the point at which enrollment is administered. Alternatively, as a convenience to a community of relying parties, a service provider could produce revocation information for a number of different CAs. If these CAs offer their revocation information in accordance with a variety of different protocols, then the third-party may present a single interface to relying parties, in order to simplify the processing required of them.

In order to take advantage of this facility, we propose that the Redirect Pointer introduced in the previous section also contain an Indirect CRL indicator. This would allow third parties to be enlisted to aid in processing revocation information.

Notice that Redirect Pointers now not only allow for dynamic repartitioning of the space requirements, but also of the scale requirements. Combined with the Freshest Revocation Information Pointer, which allows for dynamic repartitioning of the freshness requirements, what has been described is a general, flexible model for revocation processing which fits completely within the framework of existing standards.

5. Conclusion

We have presented a framework that allows CRLs to be small, to provide timely information when needed, to scale and to be flexible. It also has the advantage that it can be engineered at the outset to provide a reliable service meeting certain space, timeliness and scale requirements and can be modified later if these requirements change. This framework is based on established standards and requires only the definition of one additional certificate extension and one additional CRL extension. We believe that there is no need to define new structures, protocols, and trusted third parties to provide small, timely and scalable revocation information. These objectives can be met with minor additions to existing standards.

6. Patent Statement

Techniques described in this document may be covered by US patent 5,699,431. A license to use this patented technique is available on a royalty-free, non-onerous basis.