Global Information Assurance Certification Paper

Certificate Revocation in Public Key Infrastructures

Scott Fairbrother

Scope

This paper presents a description of the more popular certificate revocation options that have been proposed within the Public Key Infrastructure environment. Although discussions relate specifically to Public Key certificates as defined in X.509 Version 3, many of the principles and models may be applied to other certificate variants.

Background

The revocation of public key certificates is still very much an area of debate with various ideas and proposals surrounding every aspect including meaning, methodology, technology, and even the requirement for revocation.
This discussion will continue while public key infrastructures (PKIs) are still maturing and as application developers, service providers, enterprises and standards bodies begin to understand the implications of large-scale PKI deployments.
It's not that a PKI or digital certificate is a new concept. In-fact these ideas have been around since the late seventies, but it is the ability of industry to implement an online model that meets the requirements of current business processes and maintains the important trusted environment, which has fostered this examination.

The Meaning of Revocation

As a fundamental requirement of PKI is to maintain a path of trust, it is therefore essential to preserve the ability to verify the validity of a public key certificate. An essential element contributing to this is the ability to perform certificate revocation.
So, what is certificate revocation?
The meaning of "certificate revocation" in a literal sense can be defined as removing a certificate's validity prematurely. That is, the elimination of the trusted binding of a certificate to the issuing party prior to the certificate expiration date. This change in status must then be communicated via various distribution methods (to be described later), to certificate end users and systems.
Certificate revocation may be required for a number of reasons. Some of which include:

- Change of status of a subject
- Change of subject name
- Change of relationship between a subject and an organisation
- Severing or suspension of the relationship between an issuing party and an organisation
- An issuing authority ceases to operate
- Suspected private key compromise
- Detected private key compromise
- Media containing a private key is compromised or lost
- There has been improper of faulty issue of the certificate
- Certificate no longer required by the subject

The Revocation Process

A source of much discussion is the logical structure of certificate revocation and its implications to end-user applications. In order for a certificate to be valid, the relying party (the certificate user) must trust the issuing party (signer of the certificate) utilising any of the revocation status verification mechanisms as described in this paper.
In order to maintain this environment of trust it is vital that the revocation process is well defined, implemented, and enforced without any ambiguity existing as to the status of a certificate.

Authority to Request Revocation

A revocation request should only be initiated by:

- The certificate holder
- Authorised third parties such as parties with Power of Attorney from the certificate holder, or law enforcement agencies with appropriate jurisdiction
- The certificate issuing authority if they believe circumstances that warrant revocation exist

The issuing authority may receive a revocation request either from the certificate holder in person, via physical mail, or by facsimile. The proper authentication of the requestor by the issuing authority is crucial in ensuring that Denial of Service attacks are not caused by fraudulent requests.

Certificate Revocation

Once the revocation requestor has been authenticated the issuing authority processes the request. A process similar to the following would be implemented:

- Notice of the revoked certificate is published
- Confirmation of the key and certificate revocation including the date and time revocation was actioned is sent to the certificate holder (eg. via email)
- The revoked certificate remains published until the certificate's validity period elapses.

Methods of Revocation

Even though the concept of Public Key Infrastructures has been around for a number of years the business environment has been cautious in its implementation. This is due in part to the complexity of deploying this online architecture, with little data available on large-scale operations to aid infrastructure planning and dimensioning. As we are still in the "pioneering days" of PKIs there are a number of certificate revocation models that exist with different implementation and infrastructure requirements. Those models selected will depend on business requirements and issues to be considered including:

- Cost
- Infrastructure
- Availability
- Redundancy

- Timeliness
- Relying party risk and liability
- Volumes and patterns of transactions

An analysis of discussions on certificate revocation reveals four major concepts behind validation and the need for revocation of certificates. These consist of periodic revocation mechanisms such as Certificate Revocation Lists (CRLs), which have several forms of implementation; online query mechanisms such as the Online Certificate Status Protocol; re-issuance of certificates in response to validity checks; and a case for no revocation where certificate validity periods are set within the policy requirement to revoke them.
Within this paper we will concentrate on the fundamental revocation concepts of periodic and online revocation.

Certificate Revocation List (CRL)

Among the standards that exist today and those being developed, the Certificate Revocation List is generally the preferred model. There are many variations to this model but they are all based around the same basic structure.
The Certificate Revocation List is a periodically published data structure that contains a list of revoked certificate serial numbers. The CRL is time-stamped and digitally signed by typically the issuer of the certificates. However, other trusted third-party entities such as those providing "revocation services" may also publish and sign the CRL. Generally a CRL is published within an X.500 directory which also stores the certificates for the particular CA domain.
The publishing period should be determined by relying party business needs and therefore the associated Certificate Policy. Protocols used to extract revocation information need not employ signed transactions as the digital signature on the CRL maintains the integrity of the CRL itself.
CRLs are currently defined in the X.509 standard with two CRL versions being defined since 1988. As with many standards the detail of implementation is open to interpretation, and other works such as RFC2459 have been developed to address specific requirements.

Complete or Base CRL

This is the implementation of a CRL as described above which is limited to containing all revocation information of a single CA domain. Successful use of this model would only be effected provided the number of end-entities was relatively small.
Drawbacks with the use of complete CRLs include:

- Scalability issues due to the volume of posted data, which can escalate significantly given that revocation information in the CRL must remain available until expiry of the certificate. (If certificate validity periods are reduced this helps to alleviate this problem.)
- As with all CRL variants timeliness can be an issue so it is important to align the posting periodicity with business requirements and certificate policies. As the volumes of posted data increase this may enforce a minimum CRL refresh period due to unacceptable performance degradation beyond a certain threshold.
- The single data structure of a base CRL limits the ability to distribute the network load especially as the data size expands.

Authority Revocation List (ARL)

An ARL is a CRL that is used exclusively to publish revocation information for CAs. It therefore does not contain any revocation information pertaining to end-user certificates. As an ARL is used to revoke the certificates of CAs it is typically issued by either a superior CA (one which has the responsibility of revoking subordinate CAs) or the issuing CA is revoking a cross-certificate issued by that CA. The ARL is identified using the *Issuing Distribution Point* extension as implemented in X.509 for Version 2 CRL extensions.
When validating a certificate path, a valid ARL must be available for each CA that has signed certificates in the path. The exception being for a self-signed root CA performing the function of a "trust anchor" within the particular domain.
It can be expected that an ARL will remain empty or small, as revocation of CA certificates is likely to be rare. This will ensure that performance is significantly improved over Complete CRLs.

CRL Distribution Points

The CRL Distribution Points (also known as Partitioned CRLs) scheme allows a single CA domain to post revocation information on multiple CRLs. Certificates have knowledge of the CRL Distribution Point by utilising the *CRL Distribution Point* extension as specified in X.509 Version 3. This therefore ensures that the relying party does not need to have prior knowledge of where the revocation information for a particular certificate might be located. Another and more significant advantage is that revocation information is spread across a number of more manageable partitions to enhance scalability and improve performance by distributing both the maximum and average loads. The drawback with this is that each end entity is likely to require access to multiple partitions therefore increasing the average request rate. Scalability may also be restricted in that the CRL partitions are fixed or static, which leads to the next model.

Redirect CRL

A Redirect or Referral CRL is based on existing standards and protocols but expands the concept of standard CRL Distribution Points by allowing for a more flexible partitioning approach. This structure is basically an empty CRL in that it contains no certificate revocation entries, but importantly it does contain specific Redirect Pointers within CRL extensions that identify the location of CRL partitions. If re-partitioning is required then this provides a flexible means of re-defining pointers with the important feature of being an issuer signed structure.

Dynamic CRL Distribution Points

The concept of Dynamic CRL Distribution Points (previously referred to as Enhanced CRL Distribution Points) was developed to overcome the static partitioning of CRL Distribution Points. Once the associated certificate is issued, the CRL partition pointed to by the CRL Distribution Point extension is fixed for the life of that certificate. In addition to this, the issuing CA must have prior knowledge of the partitioning structure and that this structure cannot change over time. It is of course

desirable, especially with evolving PKIs that implementations are able to evolve with the needs of the PKI community.
The answer to this is to implement a flexible and dynamic partitioning capability, which is defined through scoping statements with associated pointers. A scope statement specifies a range of certificates that are covered by a particular CRL partition while the pointer defines the associated CRL Distribution Point.
These requirements have been met by proposing a CRL *Scope* field and a CRL *Status Field* to be incorporated in amendments to X.509 certificate extensions.
The original proposed concept known as Open CRL Distribution Points employed this strategy but with unsigned pointers open to denial-or-service attacks. The concept of Dynamic CRL Distribution Points addresses this issue by defining pointers in Redirect CRLs as previously discussed.

Delta CRL

A Delta CRL is a digitally signed list of revoked certificates that have occurred since the last posting of a Base CRL or CRL partition referred to by a CRL Distribution Point. Each successive periodic posting of a Delta CRL contains the latest certificate revocation update including all updates listed in previous Delta CRLs. The Delta CRL is therefore an increment to a single base CRL only and never to another Delta CRL. A full CRL is referred to as the *base posting* while the Delta CRL is considered an *incremental posting*. In order to obtain the most accurate revocation information, a relying party must retrieve the most recent base CRL posting and the most recent Delta CRL.
Delta CRLs allow the validity period of base CRLs to be extended, as the Delta CRL is relatively small and therefore can be issued on a much more frequent basis than the base CRL. This both improves latency and reduces the overall network load.
In accordance with X.509 Version 2 CRL the *Delta CRL Indicator* extension is used to differentiate between a Base CRL and Delta CRL.

Freshest CRL

There are often varying requirements by relying parties on the timeliness or "freshness" of available certificate revocation information. The Freshest CRL is a method of meeting these ranging needs in a cost and infrastructure efficient manner. For those relying party applications that require more timely information, then a CRL (typically the latest Delta CRL) with short latency postings is made available via a *Freshest CRL* extension which has been incorporated into the *Final Proposed Draft Amendment on Certificate Extensions, April 1999*. Other relying parties with less stringent latency requirements may opt to utilise a Base CRL or other alternative. This service differentiation could be based on a "user pays" approach ensuring service level segmentation is applied on a business needs basis thereby distributing network load.

Indirect CRL

A PKI domain that may utilise multiple CAs or a trusted third party service provider is able to publish revocation information in a single CRL structure using this mechanism. The relying party is therefore able to avoid the retrieval of revocation

information from multiple CRLs being issued by multiple CAs. This Indirect CRL can therefore be considered as an aggregate of a number of Base CRLs.
The implementation of an Indirect CRL follows the same construct as that of a Base CRL, except for an identifying attribute in the *Issuing Distribution Point* extension. When this extension is set for an Indirect CRL there is also a requirement for a CRL entry to specify the issuing CA of the certificate using the *Certificate Issuer* extension. Both of these CRL extensions are specified in the X.509 Version 2 CRL standard.

Online Certificate Status Protocol (OCSP)

Currently the preferred online revocation mechanism amongst standards and implementations is the Online Certificate Status Protocol. This is specified in the proposed standard *X.509 Internet Public Key Infrastructure – Online Certificate Status Protocol* [RFC2560], June 1999.
The basic process consists of a request/response protocol that obtains online revocation information from a trusted entity referred to as an OCSP Responder, with all responses being digitally signed. Although these responses are available in real time, the validity interval of the provided information is only as timely as the back-end mechanisms generating the data.
It is important to note that OCSP purely provides revocation status information and does not verify whether a certificate is within its validity period.

Other Certificate Revocation Mechanisms

This paper has been limited to the descriptions of certificate revocation mechanisms that are supported or likely to be supported in current or future standards. Although the following revocation methods are relatively obscure, they have been listed here for the sake of completeness and so that the reader may research these as desired.

- Over-Issued CRL
- Sliding Window Delta CRL
- Windowed Revocation
- Certificate Revocation Tree (CRT)
- Certificate Revocation 2-3 Tree (23CRT)
- Certificate Revocation System 1
- Certificate Revocation System 2
- Hierarchical Certificate Revocation Scheme
- Easy Fast Efficient Certification Technique (EFECT)
- Online Certificate Status Protocol – Extensions (OCSP-X)
- Simple Certificate Verification Protocol (SCVP)
- Data Certification Server (DCS)
- Dependers

Summary

As with all system designs it is important to analysis current business practice and to determine any future needs based on business growth. These should then be factored in when making decisions on suitable revocation schemes. Many deployed systems

will consist of a combination of revocation mechanisms depending on such things as business risk due to less timely revocation data, network capacity, and relying party mode of operation. It is also important to ensure that deployed models support the requirements of certificate policies.
In selecting a revocation system, the general principle of *protection criteria* should be applied and states that "the protection need only be as strong as the value of what it protects" [Kahn, David. The Codebreakers. Macmillan Publishing Co., 1967]

References

Gunter, Carl A. & Jim, Trevor. "Generalized Certificate Revocation".
URL: http://www.acm.org/pubs/articles/proceedings/plan/325694/p316-gunter/p316-gunter.pdf (5 March, 2001)

Arnes, Andre; Just, Mike; Knapskog, Svein J.; Lloyd, Steve; Meijer, Henk. "Selecting Revocation Solutions for PKI.
URL:
http://citeseer.nj.nec.com/cache/papers2/cs/17243/http:zSzzSzwww.pvv.ntnu.nozSz~andrearnzSzcertrevzSzcrpaper_final_fullpage.pdf/selecting-revocation-solutions-for.pdf (5 March, 2001)

Gassko, Irene; Gemmell, Peter S.; MacKenzie, Philip. "Efficient and Fresh Certification".
URL: http://citeseer.nj.nec.com/cache/papers2/cs/11673/http:zSzzSzwww.bell-labs.comzSzuserzSzphilmaczSz.zSz.zSzresearchzSzcert-pkcfinal.pdf/gassko00efficient.pdf (8 March, 2001)

Naor, Moni & Nissim, Kobbi. "Certificate Revocation and Certificate Update".
URL:
http://citeseer.nj.nec.com/cache/papers2/cs/14211/http:zSzzSzwww.wisdom.weizmann.ac.ilzSzhomezSzkobbizSzpublic_htmlzSzpaperszSzrevoke_JSAC.pdf/naor98certificate.pdf (8 March, 2001)

Millen, Jonathan K. & Wright, Rebecca N.; "Certificate Revocation the Responsible Way"
URL:
http://citeseer.nj.nec.com/cache/papers2/cs/14250/http:zSzzSzwww.csl.sri.comzSz~millenzSzpaperszSzneeds.pdf/millen99certificate.pdf (9 March, 2001)

Adams, Dr. Carlisle & Zuccherato, Dr. Robert. "A General, Flexible Approach to Certificate Revocation"
URL: http://www.entrust.com/resourcecenter/pdf/certrev.pdf (5 March, 2001)

McDaniel, Patrick & Jamin, Sugih. "Windowed Certificate Revocation". 1 November, 1999.
URL:
http://citeseer.nj.nec.com/cache/papers2/cs/12714/http:zSzzSzwww.eecs.umich.eduzSz~pdmcdanzSzdocszSzCSE-TR-413-99.pdf/mcdaniel99windowed.pdf (5 March, 2001)

Cooper, David A.; "A More Efficient Use of Delta-CRLs"
URL: http://csrc.nist.gov/pki/documents/sliding_window.pdf (5 March, 2001)

Fox, Barbara & LaMacchia, Brian. “Online Certificate Status Checking in Financial Transactions: The Case for Re-issuance”
URL: http://www.farcaster.com/papers/fc99/fc99.htm (5 March, 2001)

Adams, Carlisle & Lloyd, Steve. Understanding Public-Key Infrastructure: Concepts, Standards, and Deployment Considerations. Indianapolis:Macmillan Technical Publishing, 1999. 76 – 80, 109 – 131.