

Side-Effects of Cross-Certification

James L. Fisher, Ph.D.
Mitretek Systems, Falls Church VA
jlf@mitretek.org

While many organizations lean towards cross-certification with bridge certification authorities (BCAs) for wider PKI interoperability, there are many hidden details that can affect operating capabilities as well as legal standings. The purpose of this paper is to share lessons learned to help the reader to understand implications of choosing a cross-certificate based trust model. Topics covered include: X.500/LDAP certificate directory implementation and interoperability requirements; transitive and asymmetrical trust; choosing the proper trust anchor; cross-certificate policy mappings; and Certification & Accreditation requirements.

This paper does not examine the trust path discovery process—it focuses on the necessary, enabling configuration components that enable path discovery and path validation to be automated.

There are technical solutions to the issues presented herein. However, due to their inherent complexity, a discussion of alternative solutions is beyond the scope of this paper.

Benefits of Certificate Bridges

One of the primary advertised benefits of a certificate bridge is that it allows the relying party to enjoy the benefits of a larger trust domain while not being required to be an integral part of the certificate hierarchy of that other trust domain, all while trusting only one trust anchor—a public key which is within its own trust domain.

The other benefit is that the relying party can also be relatively sure of the certificate holder's identity, based on the trust placed in others to validate an organization's identity. It's not that cross-certification

automatically grants that peace of mind, but rather that it is standard practice for each party considering cross-certification to scrutinize the other party's pre-issuance identity vetting policies, private-key protection policies, CA and directory infrastructure operational policies, etc. Thus, an issued cross-certificate represents a thorough background check with acceptable findings.

Issued cross-certificates can be used to dynamically assemble a chain of certificates called a trust path which spans the gap between the certificate issuer and the relying party. The complete set of certificates comprising the certificate trust path (including supporting time-specific validity statements) forms a tangible record of trust that can be stored for future evidence of due diligence. This might be necessary for institutional archival purposes or to satisfy National Archive and Records Administration (NARA) requirements.

Directory Interoperability

The operating authority of a BCA typically provides a publicly available X.500 or LDAP directory for publishing issued CA certificates, cross-certificates, and often certificate revocation lists (CRLs). Equally important, these X.500/LDAP directories are configured to facilitate trust path discovery and validation by providing chaining to, or referrals to, all other CA directories to which cross-certificates have been issued. The intent is to provide a one-stop-shop virtual directory from which all relevant certificates and CRLs can be retrieved during the path discovery and validation processes.

In practice, each cross-certifying organization typically has its own

X.500/LDAP directory to which its own users (certificate issuers and/or employees) point the LDAP clients in their local validation engines, and if the requested certificate or CRL does not fall under that local directory's base distinguished name (DN), a superior reference chains (transparently to the user) to the master BCA for retrieval. No matter where the certificate or CRL resides, it is returned to the LDAP client. Without such chaining, there is currently no way for a desktop LDAP client to discover what directory to connect to for retrieval of a certificate or CRL. This necessary processing has interesting implementation implications.

First, the BCA directory must be able to resolve any base DN that could *ever* form a trust path through that BCA. An illustration will help in understanding the details.

In the examples we use in this paper, assume the existence of the following fictitious PKI participants:

- A Government (Certification) Bridge Authority/Architecture (GBA)
- A Government Institution (GI)
- A Neighboring Nation Bridge Authority/Architecture (NNBA)
- A University Bridge Authority/Architecture (UBA)
- An older, legacy PKI system at the Enormous State University (ESU1)
- A newer PKI system at the same university (ESU2)
- A World Wide Council (WWC)
- A random transoceanic nation (RTN)

The following list describes the cross-certifications that have occurred between the above fictitious entities, and the resulting trust mesh appears in Figure 1:

- The Government Bridge—GBA— (with base DN `ou=GBA, o=Upper Government, c=US`) has cross-certified with only:
 - GI (with base DN `ou=GI, o=Upper Government, c=US`)
 - UBA (with base DN `o=edu, c=US`)
 - Neighboring Nation (with base DN `c=NN`)
- The University Bridge (UBA) has also cross-certified with:
 - The original Enormous State University infrastructure, ESU1 (with base DN `o=ESU Provosts, c=us`)
 - The new Enormous State University infrastructure, ESU2 (with base DN `o=ESU Provosts, c=us, dc=esu, dc=edu`)
- The Neighboring Nation has also cross-certified with:
 - The World Wide Council (with base DN `dc=int`)
- The World Wide Council cross-certifies with:
 - The Random Transoceanic Nation (with base DN `c=RTN`)




Figure 1: Cross-Certification Between Trusting Partners

Figure 1 allows us to see more clearly the many and varied trust paths that can be formed. For instance, through certificate bridges and cross-certificates, the Government Institution should be able to trust digitally signed messages from the Enormous State University, regardless of whether the signer's certificate was issued from ESU's old or new PKI infrastructures. Additionally, cross-certificates would also allow the Government Institution to trust digitally signed messages from the Random Transoceanic Nation—this trust path may or may not be intentional or desired, as we will discuss later. Let us now look at the certificate/CRL directory configurations required to enable a relying party to easily discover and validate a trust path.

Before discussing the supporting certificate directories, it will be helpful to review the definition of key X.500 directory knowledge references and concepts:

- Directory Server Agent (DSA): the software providing the X.500 directory service for a particular hierarchal directory information base (DIB)

- Name context: a subtree of a directory, and is identified by the DN of the topmost entry (the "base DN"); in many commercial databases, a DSA's database can contain more than one name context
- Cross-reference: specifies a name context (usually within a remote DSA) that is not a child of this DSA's name context; a cross-reference typically points directly to the entry in the name context hierarchy
- Subordinate reference: specifies a name context (usually within a remote DSA) that is a child of this DSA's name context
- Superior reference: specifies the parent DSA that holds entries outside this DSA; only one superior reference per DSA is permitted

For complete directory chaining, the (fictitious) US-based certificate directories are configured as follows, with Figure 2 representing the same information pictorially:

- The GI directory DSA is rooted at `ou=GI, O=Upper Government, c=US`, and has one superior reference to the GBA directory

- The ESU directory has:
 - One DSA rooted at `o=ESU Provosts, c=us`
 - A second DSA (or a second naming context under the first DSA) rooted at `o=ESU Provosts, c=us, dc=esu, dc=edu`
 - One superior reference to the UBA directory
- The UBA directory has:
 - One DSA rooted at `o=edu, c=US`
 - A cross-reference from the DN `o=ESU Provosts, c=us` to the original ESU directory
 - A second DSA (or a second naming context under the first DSA) rooted at `dc=edu`
 - A subordinate reference under the second DSA from the DN `o=ESU Provosts, c=us, dc=esu, dc=edu` to the new ESU directory
 - A superior reference to the GBA directory
- The GBA directory has:
 - One DSA rooted at `c=US`
 - A subordinate reference from the DN `ou=GI, o=Upper Government, c=US` to the GI
 - A subordinate reference from the DN `ou=UBA, o=edu, c=US` to the UBA directory
- A subordinate reference from the DN `o=ESU Provosts, c=us` to the original ESU directory
- A cross-reference from the DN `c=NN` to Neighboring Nation's border directory
- A cross-reference from the DN `dc=edu` to UBA's second DSA
- A cross-reference from the DN `dc=int` to the WWC's border directory
- A cross-reference from the DN `c=RTN` to the RTN's border directory
- (no superior references)

(We leave it as an exercise to the reader to consider the cross-references needed at the WWC and RTN border directories.)

Notice this very important fact: in order to be able to retrieve (via chaining) certificates issued by RTN's CA, the GBA directory must include a knowledge reference for the `c=RTN` DN (pointing to the RTN directory) *even though the GBA did not directly cross-certify with the RTN*. This requirement has potentially serious implications on how easily a BCA can dynamically expand to accommodate indirect trust agreements.

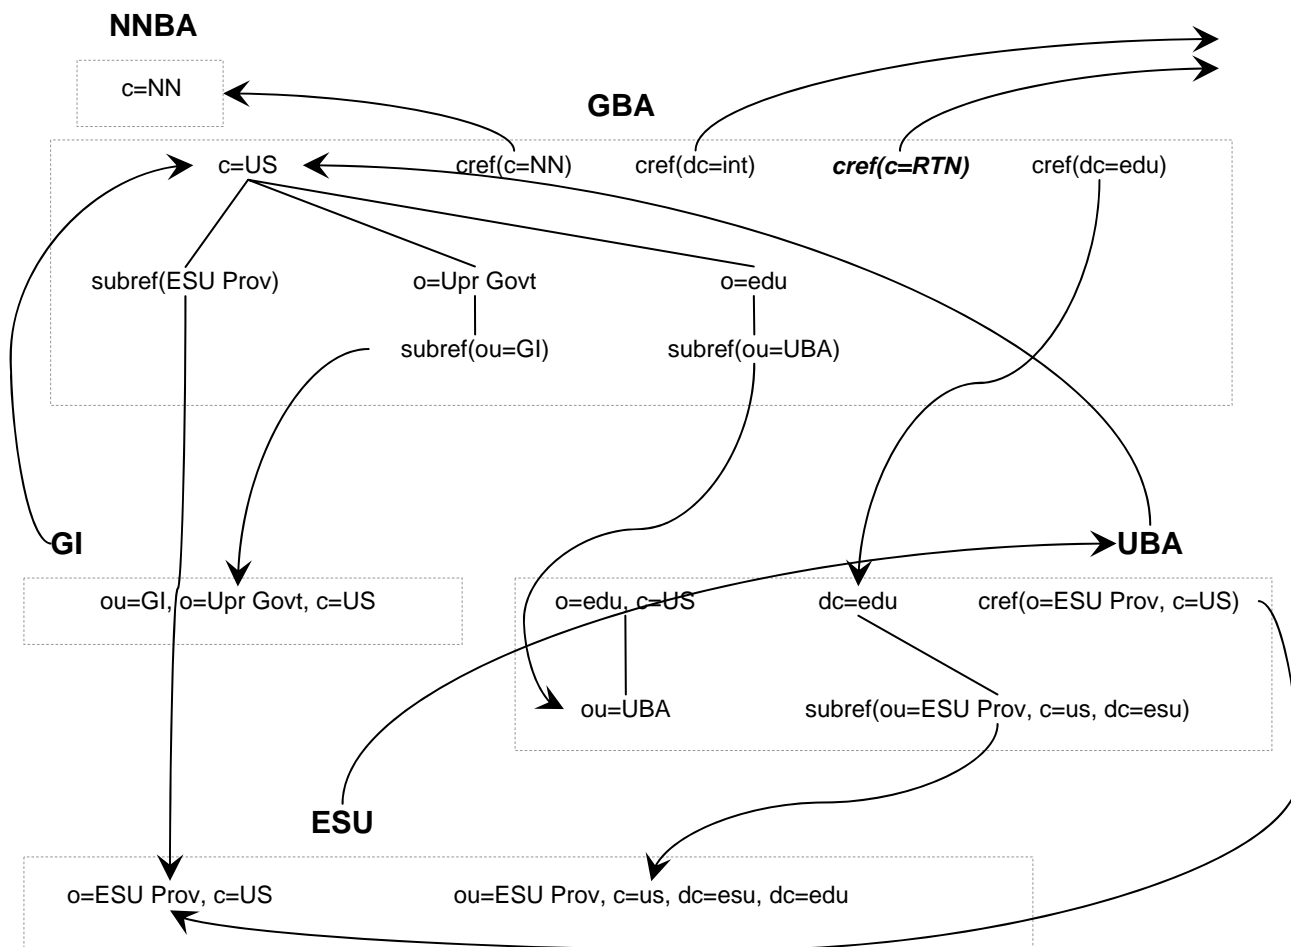


Figure 2: Directory Chaining Agreements, and Subordinate and Superior References

Notice that because of DN naming conventions chosen, and how each DSA is rooted, both GBA and UBA directories need a cross-reference from the DN `o=ESU Provosts, c=us` to the original ESU directory.

One can easily see how quickly border directory configuration becomes complicated when multiple BCAs cross-certify. Each BCA must be configured with knowledge of all possible directories traversed along a trust path, even if the BCA did not cross-certify directly with the corresponding CA.

As seen, each border directory needs the capability of supporting:

- multiple root DSAs or multiple name contexts within a single DSA

- cross-references
- subordinate references
- superior references

Fortunately, these capabilities are found in a number of commercial X.500 directories, such as: i500 by PeerLogic; eTrust by Computer Associates; M-Vault by ISODE. However, Microsoft Active Directory does not support superior references, non-local cross-references, or non-local subordinate references. Therefore, organizations wishing to provide a publicly accessible certificate directory (often called a border directory) *to their relying parties* suitable for use in automated path discovery should export all necessary certificates and CRLs from their Microsoft Active Directory and import into a directory supporting the aforementioned types of references.

Alternatives:

The BCA directory was required to provide a multitude of cross-references and subordinate references because directory chaining was desired. One design alternative is to have the BCA directory simply return *referrals* to only those directories with which it is directly cross-certified. But this presents two problems in today's environments:

- Many LDAP clients still can not process directory referrals
- Trust paths traversing more than one bridge would not be discovered if the pertinent certificate fields contained only DN-formatted information and not URI-formatted information

Directory chaining would not be necessary if all certificates had properly formatted AIA fields, and the local validation client could understand all AIA formats. However, one missing or improperly formatted AIA field would destroy the ability to discover a trust path.

Transitive Trust

When CAs cross-certify, the phenomenon of transitive trust comes into play. Such indirect trust may or may not be intended or welcomed. An internationally oriented illustration will make the side-effects clearer.

Let us extend the example of the previous section to include one additional cross-certificate pair. Assume that the very recently formed country of Forbiddenstan has also cross-certified with the World Wide Council for the purpose of discussing commercial trade. Consequently, there is now a new trust path from the GI, through the GBA, through a Neighboring Nation, through the World Wide Council, to Forbiddenstan. Let us also assume that "the United States maintains a broad embargo against trading with Forbiddenstan, and most commercial imports from Forbiddenstan are prohibited by law." And finally, to ensure compliance with federal

regulations, let us assume that GI directors would prefer that there be no valid certificate trust path from GI to Forbiddenstan.

To prevent the formation of such a trust path, cross-certificates must be re-issued to specify new name constraints or path length constraints. There are two logical places in the trust chain where such a filter could be positioned. Since this is a federal law, the GBA's cross-certificate issued to the NNBA could contain a name constraint to filter out Forbiddenstan's $c=FB$ DN. However, since other domestic humanitarian-oriented government agencies might have legitimate needs to trust Forbiddenstan-signed documents, a nationwide filter might not be appropriate. Therefore, a GI itself might need to insert name constraints into the cross-cert it issues to the GBA. Additionally, all other relying parties would similarly need to be aware of this political situation and place appropriate name constraints in their cross-certificates.

While this method works, it is very difficult to envision all necessary path constraint needs—expressed either as path permitting or path inhibiting rules—before cross-certificate issuance. And, re-issuance of a cross-certificate is not without its labor costs. Revoking previously issued cross-certificates will also be necessary to force validation engines that pre-cache the validation paths to refresh themselves. More importantly, the need to restrict certain trust paths is typically not realized until after an "inappropriate" trust path is formed.

Choosing the Proper Trust Anchor

Another challenge of a multiple cross-certificate environment is choosing the correct trust anchor and certificate policies on which to filter. Complicating factors include:

- Asymmetrical policy mappings
- The possibility of multiple trust paths

- Unidirectional cross-certification (e.g., a GBA may issue a cross-certificate to a military to facilitate GBA-centric trust path discovery, but the military BCA may choose to not issue a cross-certificate to that GBA)

Policy mappings are often placed in cross-certificates for the purpose of declaring how the levels of assurances (LOAs) in the subject's domain translate to the LOAs in the issuer's domain. In order to make use of these policy mappings, RFC3280 indicates that path discovery and validation algorithms must specify a trust anchor and a set of acceptable certificate policies (in the trust anchor's domain). Additionally, if the initial set of acceptable certificate policies is a subset of those mapped, then the effect is to filter out any trust paths involving certificates with LOAs that do not map to that initial set.

Consider the following example. The (fictitious) State of Algonk has one CA (with one private key) issuing end-entity certificates with one of four levels of assurance (LOAs): Level A, Level B, Level C, and Level D. Independently, GSA has three LOAs: Small, Medium, and Large. But when the State of Algonk cross-certified with the GBA, Algonk submitted documentation describing only their Level C LOA, therefore the cross-certificate issued by GBA to Algonk contains only one policy mapping: Algonk Level C maps to GBA Medium.

Furthermore, the cross-certificate pair between the GBA and the Algonk contains asymmetrical policy mappings. Why? Because each party's Policy Authority (PA) could independently evaluate the other party's Certificate Policy/Certification Practice Statement (CP/CPS), and there is no guarantee the parties will view each other's policies equally. Consequently, according to the policy mappings found in the cross-certificate issued by the GBA to Algonk, the GBA views Algonk Level C LOA as mapping to the GBA Medium LOA. Conversely, according to the policy

mappings found in the cross-certificate issued by Algonk to the GBA, the State of Algonk views GBA Medium as mapping to Algonk Level B.

Typically, the trust anchor of choice is a public key within one's own issuing hierarchy. However, let us consider the case where a centralized, organization-independent validation service (employed by a GBA) is being established to validate only certificates that are GBA Medium or equivalent. Let us examine two trust anchor options and their implications.

If the trust anchor is a GBA root public key, then the certificate policy OIDs on which to filter should be GBA Medium. In this case policy-filtered trust paths can be found from the GBA to (a) Algonk Level C certificates, but not to other Algonk LOAs, and (b) any other issuers' certificates that GBA maps to GBA Medium LOA.

Alternatively, the validation service operator could reason that since Algonk certificates are validated so often, the dynamically discovered trust path for Algonk certificates should be made as short as possible for reasons of efficiency. To shorten the trust path, the trust anchor is chosen to be the State of Algonk's root public key. Since the GBA views only Algonk Level C certificates as equivalent to GBA Medium LOA, the initial set of acceptable certificate policy OIDs is just the policy OID representing Algonk Level C. Obviously, trust paths will be found to Algonk Level C certificates. However, the policy mapping in the Algonk-issued cross-certificate (i.e., the cross-certificate going in the opposite direction) states that GBA Medium maps to Algonk Level B. Since Algonk's Level B certificate policy OID is not in the initial set of acceptable certificate policies, no other issuer's certificates mapping to GBA Medium (as determined GBA's point of view) will be accepted (i.e., no valid trust path will be discovered for those certificates). And the configuration decision complications increase as more BCAs become involved. It is therefore

recommended to explicitly state the validation service's acceptable certificate policy set—and not include the phrase “or equivalent”—and use only the corresponding trust anchor.

In practice, such asymmetrical mappings can be easily avoided. Both parties should explicitly state and agree on how each of the applicant's LOAs relates to the issuing party's LOAs. Continuing with our example, when applying for cross-certification, the State of Algonk should explicitly request that the GBA PA map Algonk's Level C LOA to GBA's Medium LOA.

Post-Issuance CA Subordination

One technique for reducing the number of certificate bridges is to establish one root under which all other certificates are issued. Recently, there have been discussions of establishing such a common root (CR) that would subordinate existing commercial CAs that meet government requirements.

A common root would also offer the advantage of needing to distribute only one self-signed public key in popular web

browsers.

For proper policy OID representation, one of following two items must occur:

- The new CR must assert all policy OIDs of all subordinated CAs, or
- The end-entity certificates of the subordinated CAs must be re-issued to include the new CR policy OID

Figure 3 depicts one phase of the proposed hierarchy. Assume the GBA has previously cross-certified with one of the commercial vendor's CAs (CVCA). The root CVCA has issued a proper subordinate CA A1, and A1 issues only special-audience end-entity certificates. These end-entity certificates, when processed through the certificate mappings in the GBA/CVCA cross-certificate, map to a GBA Medium LOA. Assume the common root, CR, was subsequently cross-certified with the GBA.

Then, in the final phase, the proposed technique for demonstrating that A1 qualifies as a Scrutinized Provider is for the CR to subordinate the CA A1. In practice, this would result in a new subordinate CA A1*. A1 and A1* have the same public key and subject DN (to validate the already-

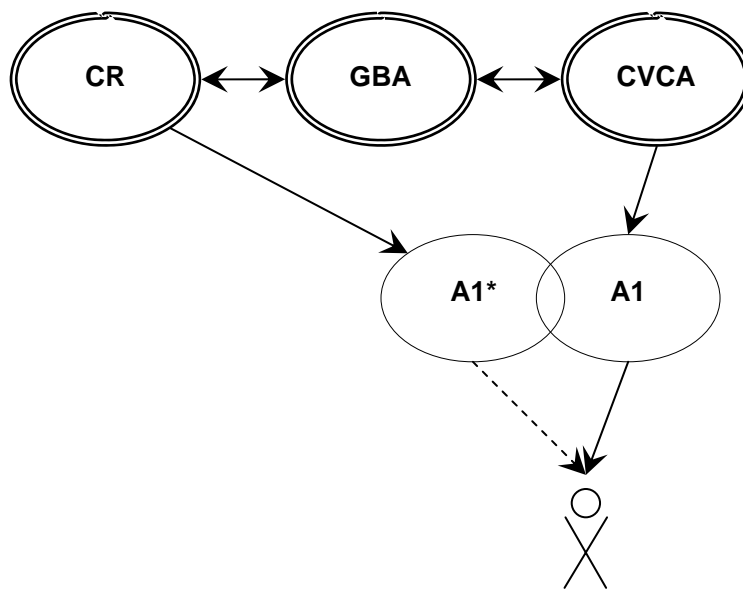


Figure 3: Subordinating Operational CA A1

existing end-entity certificates issued by A1), but different issuer DNs. This results in two possible trust paths from the CR root to the A1-issued certificate:

- One directly from the CR through the subordinated A1* CA certificate, and
- The second from the CR through the CR/GBA cross-certificate, through the GBA/CVCA cross-certificate, and finally through A1

When processed through the second trust path, the policy mappings in the cross-certificates will ensure the end-entity policies are properly mapped into the CR's certificate policy OID space. However, in direct trust hierarchies, such as from CR through A1* to the end-entity certificate, no policy mappings can take place. Therefore, A1-issued certificates must include two sets of certificate policy OIDs—one set for the CR policy name space, and the other set for the CVCA policy OID space—thus reflecting A1's two superiors.

Interestingly enough, if a relying party were given just the CR root, the subordinated A1* CA certificate issued by the CR root, and an end-entity certificate issued by CA A1, the relying party would find a perfectly valid hierarchical issuance path from the CR to the end-entity cert. However, if CVCA revoked A1 (say, due to a compromise of A1's private key), and the organization behind CVCA did not notify the GBA Program Management Office (PMO), then the above direct trust path through A1* would still appear valid when, in reality, it should be declared as "revoked."

Therefore, extreme caution should be used if such a "two master" topology is used.

Operations Policies

Typically, one focuses on technical and policy issues within one's own security domain. In this section we consider operations policies requirements such as Security Certification and Accreditation (C&A).

A popular misconception in writing system security plans (SSPs) is that when a BCA is cross-certified with the root CA of a certificate issuer hierarchy, the BCA PMO is required to see only the C&A report corresponding to the remote organization's root CA, and that organization can be trusted to silently perform C&As on their internal hierarchy. However, a careful review of NIST Special Publication (SP) 800-37, "Guide for the Security Certification and Accreditation of Federal Information Systems," reveals more stringent requirements.

SP 800-37 defines a certification agent as "an individual, group, or organization responsible for conducting a security certification, or comprehensive assessment of the management, operational, and technical security controls in an information system to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system." (p.15) Additionally, SP 800-37 states that since the certification agent "provides an independent assessment of the system security plan to ensure the plan provides a set of security controls for the information system that is adequate to meet all applicable security requirements," the certification agent needs to be independent from:

- "persons directly responsible for the development of the information system and the day-to-day operation of the system"
- "individuals responsible for correcting security deficiencies identified during the security certification"

Given a certification agent's independence from the managerial and operational chains of a CA, the resulting report cannot remain solely within the organizations chain of command—the report must be delivered externally. The organization most interested in and most impacted by such a

report is the BCA PMO, and therefore should be the recipient of the certification report.

Therefore, the PMO of each BCA should regularly see the C&As of every issuing CA to which the BCA can form a trust chain.

Conclusions

As we have discussed, "bridging" for PKI interoperability is not the panacea that many thought it to be. It is extremely complex and requires careful attention to detail. It is easy to structure unintended and difficult-to-detect consequences. Such complexity often results in significant opportunities for undetected errors that the security community often points out as exploitable vulnerabilities. We also implore each organization to consider these inherent issues when performing security C&A and Certificate Policy/Certification Practice Statement (CP/CPS) Compliance Audits.

References

NIST Special Publication (SP) 800-37, "Guide for the Security Certification and Accreditation of Federal Information Systems"

[<http://csrc.nist.gov/publications/nistpubs/800-37/SP800-37-final.pdf>]

RFC3280: "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile"

[<http://www.ietf.org/rfc/rfc3280.txt>]