

# Towards a Framework for Evaluating Certificate Status Information Mechanisms<sup>¶</sup>

\* **John ILIADIS<sup>1</sup>, Stefanos GRITZALIS<sup>1</sup>**

<sup>1</sup>*Department of Information & Communication Systems Engineering*

*University of the Aegean*

*Research Unit, 30 Voulgaroktonou St., Athens GR-11472, Greece*

*Tel: +30-10-6492.112, Fax: +30-10-6492.299, Email: {jiliad,sgritz}@aegean.gr*

**Diomidis SPINELLIS<sup>2</sup>**

<sup>2</sup>*Department of Management Science and Technology*

*Athens University of Economics & Business*

*76 Patission St., Athens GR-10434, Greece*

*Email: dds@aueb.gr*

**Danny DE COCK<sup>3</sup>, Bart PRENEEL<sup>3</sup>**

<sup>3</sup>*Dept. of Electrical Engineering ESAT/COSIC, Katholieke Universiteit Leuven*

*Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium*

*e-mail: {danny.decock,bart.preneel}@esat.kuleuven.ac.be*

**Dimitris GRITZALIS<sup>4</sup>**

<sup>4</sup>*Department of Informatics*

*Athens University of Economics & Business*

*76 Patission St., Athens GR-10434, Greece*

*Email: dgrit@aueb.gr*

## Abstract

A wide spectrum of certificate revocation mechanisms is currently in use. A number of them have been proposed by standardisation bodies, while some others have originated from academic or private institutions. What is still missing is a systematic and robust framework for the sound evaluation of these mechanisms. We present a mechanism-neutral framework for the evaluation of certificate status information (CSI) mechanisms. These mechanisms collect, process and distribute CSI. A detailed demonstration of its exploitation is also provided. The demonstration is mainly based on the evaluation of Certificate Revocation Lists, as well as of the Online Certificate Status Protocol. Other well-known CSI mechanisms are also mentioned for completeness.

---

<sup>¶</sup> *Computer Communications*, 26(16):1839–1850, October 2003.

This is a machine-readable rendering of a working paper draft that led to a publication. The publication should always be cited in preference to this draft using the reference in the previous footnote. This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

\* Corresponding author

## Keywords

Certificate, Certificate Revocation, Certificate Status, Certificate Revocation List, Certificate Revocation Status, Certificate Revocation Tree, Evaluation Framework.

## 1 INTRODUCTION

The deployment of Public Key Infrastructures, as an e-commerce and e-business enabling technology has been extensively studied. PKI applications could now be used on a large scale to enable electronic services, such as e-government, B2B transactions and e-commerce. It has rapidly become clear that these services would become successful only if PKI users can trust digital certificates and the respective digital signatures. In this context, “trusting a digital signature” means that the user who validates a digital signature has obtained enough evidence to believe that the signature was computed by the intended producer of the digital signature as stated in the certificate that corresponds to that digital signature, and that the complete certificate chain is still valid. Consequently, a digital signature is not to be trusted if one of those criteria is not met.

In this paper, we present an evaluation framework for mechanisms that collect, process and distribute information pertinent to the validity of digital signatures and certificates. This framework deals with revocation of the issued certificates, and the way revocation information can be made available to certificate users, using Certificate Revocation Lists (CRLs), CRLs variants, OCSP, etc. An early version of this paper has been published at the ISSE2000 conference in Barcelona, Spain [1].

The evaluation methods described in this paper strive to avoid unnecessary biases. The framework comprises of a set of qualitative and quantitative evaluation criteria, which can be applied to any mechanism that updates information on the status of certificates (Certificate Status Information — CSI). We use this model as a tool to identify potential problems in the mechanisms in a methodical way. Our evaluation framework splits the evaluation process into three main domains, namely management, performance and security.

Management of revocation mechanisms includes the way these mechanisms operate, the way information processing is being performed, the participating entities, and the respective timeframe.

Performance of revocation mechanisms refers to the efficiency characteristics of those mechanisms. These characteristics include the timeliness of the mechanism, the freshness of information it delivers, the scalability and adjustability of the mechanism, and the capability to immediately generate information on the status of a certificate (emergency certificate status information).

The security aspect of revocation mechanisms covers issues related to protecting the operation of the mechanisms themselves, as well as of the information they produce. Certificate status information has to be protected while generated, communicated, and stored.

While designing the evaluation framework, we took into consideration the requirements and restrictions imposed on the use of these mechanisms by the modus operandi required by the European Directive on a Community Framework for Electronic Signatures [2], the EESSI Expert Team Report [3] and the NIST PKI study [4], [5]. The framework can be used to evaluate mechanisms that are operated by a Certification Service Provider (CSP) that issues “qualified certificates” [2], [3]. Certificates are considered to be “qualified” [2] if they meet the requirements set forth in Annex I of the Directive and are provided by a Certification Service Provider meeting the requirements laid out in Annex II of the Directive. We have also considered the draft or final requirements and recommendations contained in [6], [7], [8], [9], [10], [11], [12]. Most European countries have rewritten their legislation so that digital signatures produced with a secure signature

creation device must be considered as a handwritten signature if that digital signature comes with a qualified certificate.

We demonstrate the application of our framework by evaluating in detail the two most frequently used mechanisms: Certificate Revocation Lists [6], [9], and the Online Certificate Status Protocol [11]. Other techniques that are less frequently used, but are also of interest, include the Certificate Revocation Status mechanism [13], Suicide Notes [14], Revocation Authority [15], Authenticated Dictionaries [16] and Certificate Revocation Trees [17]. These techniques are discussed in less detail.

This paper is organised as follows: In section 2, we present the aforementioned CSI mechanisms, while section 3 deals with the evaluation framework for CSI mechanisms. In section 4, we describe two representative CSI mechanisms and evaluate them based on the framework presented in section 3. In section 5, we provide a discussion, based on the evaluation presented in section 4 and the evaluation framework presented in section 3, while section 6 contains our concluding remarks.

## 2 CLASSIFICATION OF CSI-RETRIEVAL MECHANISMS

In this section, we provide an overview of CSI mechanisms. There are three distinct mechanism categories: mechanisms that provide negative CSI, mechanisms that provide positive CSI, and those that provide complete certificate status information. Each of these categories is discussed in the following subsections.

### 2.1 Mechanisms supporting negative CSI

All the methods described in this section specify negative CSI, i.e. CSI that concerns the certificates that have been revoked.

#### 2.1.1 *Certificate Revocation Lists*

The most common method of revocation is the publication of a Certificate Revocation List (CRL). This list consists of time stamped pointers to revoked certificates, which have not yet expired. The certificate identifier used in this list is the unique serial number assigned to the certificate by the Certification Service Provider that issued the certificate. A certificate that appears in a CRL is no longer considered valid. As the list is digitally signed by the CSP that issued the CRL, it can be retrieved from an untrusted CSI repository such as an LDAP server or a Web server. Each certificate contains a reference to the CSI repository, typically in the form of a URI.

#### 2.1.2 *Distribution Point CRLs*

A CRL can quickly become too large to be handled by dependent entities. Using CRLs requires a large storage capacity in the CSI repository, large bandwidth for the dependent entity that retrieves the CRL, and a considerable amount of time for that entity to process the CRL, each time a certificate has to be verified. With Distribution Point CRLs it is possible to split a full CRL into several smaller CRL partitions using a variety of partitioning criteria (e.g., certificate serial numbers, certificate issuing dates, revocation reasons, etc.). The collection of all these CRLs includes all the entries the full CRL contains.

Each certificate contains a “DistributionPointName” extension. That extension defines one or more distribution points. Each of these points provides the URI of the CSI repository, the reasons for the revocation of certificates contained in that CRL, and the issuer of the CRL.

Each Distribution Point CRL contains an “Issuing Distribution Point” extension containing flags indicating whether the CRL contains revoked CA or end-user certificates, a list of revocation

reasons for the certificates contained in the CRL and an indication of whether or not the CRL is an indirect CRL.

### *2.1.3 Delta-CRLs*

Delta-CRLs [6], [9] are the second major variation on full CRLs. A Delta-CRL is used to partition a CRL according to time criteria. It is issued with respect to a particular base CRL, and contains all updates since that base CRL was issued. CSI found in the base CRL is combined with the incremental information found in the most recently issued Delta-CRL to determine which certificates have been revoked. The extensions used for a Delta-CRL consist of the “deltaCRLIndicator” and the “baseCRLNumber”. Delta-CRLs reduce the necessary resources for the communication of CSI to dependent entities and for the processing of CSI by these entities.

### *2.1.4 Freshest CRL*

If near-real time CSI is required, a URI pointing at a “Freshest Revocation Information Pointer” [18] can be included in a certificate. The latter points to a CRL or Delta-CRL that has as a base the partitioned CRL pointed at by the `crlDistributionPoint`. In this model, Delta-CRLs are issued very frequently, at time intervals specified a priori, and each time an additional certificate has to be revoked. The dependent entity can retrieve the latest Delta-CRL while retaining the base-CRL, which is issued much less frequently.

### *2.1.5 Redirect CRL*

Adams and Zuccherato [18] introduced the concept of a Redirect CRL. This concept relies on Distribution Point CRLs resulting from dynamic repartitioning of a full CRL. It is particularly useful for devices with limited resources, as a CSP can “load balance” the CRLs by moving entries between them to keep the size of the individual CRLs about the same. This is done by using the status referral extension to refer to the CRLs containing moved entries. A dependent entity could then check the original CRL, and subsequently locate any CRLs the CSI has been moved to.

### *2.1.6 Indirect CRLs*

An Indirect CRL combines CRL entries from multiple CSPs into a single CRL. These CRLs are implemented using three extensions: the issuing distribution point extension for a CRL to identify itself as an Indirect CRL, the corresponding certificate extension that indicates which CA issued the original CRL (if different from the CA that issues the Indirect CRL), and a certificate issuer extension contained in each CRL, pointing to the CSP that issued the corresponding revoked certificate.

## **2.2 Mechanisms supporting positive CSI**

All the methods described in this section specify positive CSI, i.e. the information contained in CSI specify which certificates are valid.

### *2.2.1 Suicide Notes*

A certificate holder can revoke his certificate by signing a suicide note and sending it to a network of collaborating Suicide Bureaus [14]. These Bureaus issue signed statements called “certificates of health”, regarding a specific certificate. A certificate of health states that at that point in time, no evidence had been received by the Suicide Bureaus indicating the revocation of that certificate.

### *2.2.2 Revocation Authority*

A Revocation Authority typically issues, as opposed to a Certification Authority, short-lived time stamped certificates. These certificates assert on the validity of long-term certificates that are issued by a CSP. A long-term certificate contains information concerning freshness constraints on the CSI

a dependent entity must require, together with a pointer to the Revocation Authority that issues the short-lived certificates.

## **2.3 Mechanisms supporting complete CSI**

CSI mechanisms presented in the following sections give explicit information on the validity of certificates in question, i.e. CSI that specifies explicitly whether a certificate has been revoked or if it is still valid.

### *2.3.1 Online Certificate Status Protocol*

The Online Certificate Status Protocol (OCSP) was specified by the IETF PKIX Working Group in RFC 2560 and [11] to provide real-time status of certificates using an entity called OCSP responder. That responder queries for the status of certificates (possibly by direct interaction with the Certification Authority that issued the certificate in question, or by collecting CRLs from that CA), and returns digitally signed (but possibly outdated) status responses to dependent parties in a more timely fashion than CRLs. Depending on the allocation of certificates and CAs to responders, multiple queries may be necessary to validate a certificate chain. OCSP can also be used in conjunction with CRLs, as it provides an extension that can be used as a pointer to a CRL, in case timelier CSI is unavailable at a certain point of time.

The OCSP responder produces a digitally signed response on one of these states: good (the certificate has not been revoked), revoked, or unknown (the responder cannot determine whether the certificate has been revoked or not). The response has time-limited validity.

### *2.3.2 Certificate Revocation Status*

Micali introduced [13] the Certificate Revocation Status (CRS) CSI mechanism. Using this mechanism, a Certification Service Provider publishes to a public repository, on a daily basis, a digitally signed message for each certificate it has issued, which states that the certificate has been revoked or not. Although this method reduces the communication burden for the dependent entity, which no longer needs to download a (possibly) large CRL (a single query for each certificate in question suffices), the CA needs to produce and disseminate a digitally signed statement to the repository, for each certificate it has generated.

### *2.3.3 Certificate Revocation Trees*

The Certificate Revocation Tree (CRT) [17] consists of a binary hash tree, where each leaf corresponds to a set of statements about individual certificates. Each leaf specifies a lower and an upper boundary for certificates. A certificate has been revoked if its serial number appears as a lower-boundary in one of the tree's leaves. It is clear that it is not necessary to browse the complete tree in order to determine whether a certificate has been revoked. Note however, that it is computationally expensive to add a newly revoked certificate to the tree, as it requires the full restructuring of the binary hash tree.

### *2.3.4 Authenticated Dictionaries*

An Authenticated Dictionary [16] is similar to a Certificate Revocation Tree. It consists of a search tree that counters the CRT's update complexity. If a certificate has to be inserted, and the leaf that needs to be updated consists of two branches, a third branch is simply added. If the leaf consists of three branches, a new sub-tree is created, each branch of which consists of two leaves.

## 3 EVALUATION FRAMEWORK

The Electronic Signature Directive (Annex II) requires “the operation of a prompt and secure directory and a secure and immediate revocation service”. Furthermore, the Directive requires that the authenticity and validity of the certificate required at the time of signature verification are reliably verified, and that the verification result and the signatory's identity are correctly displayed. The EU Directive on a Community Framework for Electronic Signatures also requires that the date and time when a certificate is issued or revoked must be determined precisely. Finally, the last major requirement of the EU Directive that relates to CSI mechanisms is that the use of these mechanisms must be in accordance with data protection legislation, and respect the privacy of individuals.

The evaluation framework we propose covers the aforementioned, legislative requirements; our framework also comprises of technical requirements the EU Directive and the related regulatory framework does not deal with. The requirements our framework identifies fall in three main categories: *management*, *performance* and *security*. These requirements are, in fact, criteria that can be used both for evaluating existing CSI mechanisms, as well as for designing new ones. In the following paragraphs, these requirements are referred to as CSI Evaluation Criteria, or criteria for short.

### 3.1 Management

#### 3.1.1 Feedback

Feedback on the CSI that has been retrieved is mostly a matter of user interface of PKI-aware applications or devices that handle CSI on behalf of the dependent entity, and not a matter of CSI mechanisms themselves. Dependent entities must receive information regarding the intermediate operation results of a CSI mechanism and the final output it produces (i.e., whether the certificate is still valid), in accordance with Annex II of the EU Directive on a Community Framework for Electronic Signatures. This feedback must indicate at least the following information:

1. Location information on the CSI that relates to the certificate the dependent entity attempts to validate. This information could be a URI [19],
2. if the CSI location has been successfully contacted,
3. if CSI has been retrieved,
4. if the validity (integrity and authenticity) of the retrieved CSI can be verified. Also, if this verification can be based (directly or indirectly) on information the dependent entity has declared as trusted (e.g., CA certificates that are stored locally), or if other, possibly untrusted, information is also needed (e.g., more CA certificates which cannot be validated based on the locally stored set of trusted CA certificates),
5. if the CSI that was retrieved corresponds to the certificate the dependent entity wishes to validate,
6. the status (revoked, suspended, not revoked) of the certificate the dependent entity wishes to validate.

If the CSI mechanism can provide the dependent entity with the information above, either at the beginning or at the end of its execution, or in fragments while the mechanism is operating, then the feedback criterion is met.

### 3.1.2 Transparency

Users of information systems nowadays are no more computer experts. On the contrary, they are inexperienced computer users that are bound to lack information security awareness and training. Therefore, locating the CSI repository and verifying that the CSI contained in that repository is the one that corresponds to the certificate to be verified must be an automated procedure that requires no human intervention. The intervention of the user should be restricted, if possible, to requesting a validity check on a specified certificate. Going even further as far as abstraction is concerned, the dependent entity should not be required to request a validity check on a certificate but on a digital signature the dependent entity sees on a document it received.

If CRLs or Delta-CRLs are used, the added communication burden for disseminating certificates to dependent entities consists of the number of bytes needed for the *cRLDistributionPoints* X.509v3 certificate extension, which points to a valid URI where the specific CRL can be downloaded from. For Distribution Point CRLs based on reason for revocation, an extra URI should be included for every revocation reason. For the other types of CSI mechanisms, a single URI suffices.

If OCSP is used, a certificate should include the *AuthorityInfoAccess* extension, which points to the location of the authority that provides OCSP services for the specific certificate. The communication burden for this case equals to the number of bytes needed to include the *AuthorityInfoAccess* extension.

### 3.1.3 Delegation of revocation

The dependent entity could trust an authority, other than the CA, for generating CSI. “CSI generation” in this context means digitally signing CSI in order to protect its integrity and authenticity.

It could be either another CA or another authority that operates only as a Revocation Authority (RevA) and not as a Certification Authority. Moreover, it could be an authority, using a separate, distinct CA-issued key for signing CSI or it could be a distinct authority, local to the dependent entity, which is trusted by this entity.

In any case, the dependent entity must be able to verify the status of the keys used by that authority, based on certificates or other information it already considers to be trusted, before trusting and using CSI from that authority.

### 3.1.4 Delegation of CSI dissemination

The CA may delegate CSI dissemination to another authority. This second authority may need to be trusted by the dependent entities or not, depending on the mechanics of CSI dissemination.

The dependent entity has to be able to verify the authenticity and integrity of CSI it retrieves from that authority: if CSI delivered to the dependent entity is contained in a CA-signed field, then the repository (e.g., an LDAP Directory) used by the CSI dissemination authority need not be trusted. However, if CSI delivered to the dependent entity is not already integrity-protected, then the authority disseminating CSI must digitally sign CSI before delivering it to the dependent entity and the dependent entity must be able to validate the respective certificates. However, even if CSI is contained in a CA-signed field, the authority that disseminates CSI may behave maliciously, withholding fresh CSI from the dependent entity. A CSI mechanism must use a distinct method to let the dependent entity verify that the CSI it received is the freshest one (e.g., bounded revocation in CRLs, white-lists in CRS [13]).

The communication cost  $s(t)$  (in time units) of distributing CSI using a CRL to one authority grows linearly with the size of the CSI at time  $t$ . When  $n$  authorities/repositories must receive this

CSI, the communication cost equals  $n$  times the cost to distribute the CSI to a single authority, i.e., the communication cost for  $n$  authorities/repositories equals  $n \cdot s(t)$  time units.

When a CRL is distributed over  $p$  Distribution Point CRLs, and  $s_{pi}(t)$  denotes the time needed to deliver CRL partition  $pi$  based on the CRL produced at time  $t$ , then the total time needed to transport all CRL partitions equals to the sum of all  $s_{pi}(t)$ 's, making the communication cost equal to  $s(t)$  time units.

The time needed to ship the Delta-CRL equals  $s(t+l)-s(t)$ , being the difference in the time needed to send the new CRL and the time needed to distribute the BaseCRL. When  $n$  authorities/repositories must receive the new Delta-CRL, the communication cost equals  $n \cdot (s(t+l)-s(t))$  time units. However, if the authorities/repositories do not already have the BaseCRL, the communication cost becomes  $n \cdot s(t+l)$  time units.

In a system with  $n$  OCSP service providers, the communication cost of CSI dissemination equals that of the first case:  $n \cdot s(t)$  time units.

### 3.1.5 *Delegation of certificate path validation*

We will present the mechanics of delegation of certificate path validation from the dependent entity to another entity. The dependent entity should be provided with the means to review and verify the results of the validation process (Annex II of the EU Directive on a Community Framework for Electronic Signatures). In addition to that, the entity that performs the certificate path validation should be trusted by the dependent entity.

The validation of a certificate path takes as input the certificate to be validated, a number of other certificates the dependent entity trusts, and CSI regarding these certificates. The output is status information regarding the certificate to be validated.

In some CSI mechanisms one proceeds as follows to assert on the validity of a certificate: the dependent entity retrieves CSI regarding the certificate to be validated and uses as input the other certificates, the retrieved CSI and the certificate in question. Other CSI mechanisms can perform the certificate path validation on behalf of the dependent entity. In this case, the dependent entity uploads the certificate to be validated and the other certificates (or appropriate values that cryptographically identify in a unique way these certificates), to the authority that has access to the relevant CSI; this authority will then validate the certificate in question on behalf of the dependent entity. That authority then communicates the validity check result to the dependent entity. The dependent entity has to be able to verify the integrity and authenticity of the result returned by the CSI mechanism, using appropriate cryptographic mechanisms.

### 3.1.6 *Referral capability*

The CSI location function may lead the dependent entity to a CSI location that does not contain the requested CSI (the CSI may be less fresh than requested, or it may not contain information regarding the specific certificate to be validated). In that case, the CSI mechanism could refer the dependent entity to another CSI location in order to retrieve the requested CSI.

### 3.1.7 *Revocation Reasons*

When validating the path, the certificate path validation function could consider the reasons for the revocation of a specific certificate contained in a certificate path. The validation function may output different results depending on the revocation reason. The semantics of revocation reasons and the logic the validation function uses in order to include this information, while validating a certificate path, is a matter of policy.

If the certificate path validation function occurs locally to the dependent entity, then the dependent entity must be able to set the policy for using revocation reasons in the validation function. If the certificate path validation function is delegated to another authority then the inclusion of revocation reasons in the logic of certificate path validation is a matter of the policy used by that authority. The dependent entity must be aware of that policy. Alternatively, the dependent entity could communicate its own policy regarding revocation reasons to the authority that performs the validation.

### 3.1.8 *Notification of revocation or suspension*

A subscriber, whose certificate is being revoked or suspended, should be notified; it might be necessary to inform other entities, such as a Revocation Authority or a Suicide Note collecting bureau, as well (mentioned as a possible policy requirement in [7]).

Notification should not be integrated in the CSI mechanisms themselves, because failure to locate appropriate contact information for the certificate owner, the dependent entity or any other entities could lead to disruption of the CSI mechanism. This procedure must be implemented outside the CSI mechanism itself, however linked to the CSI generation or CSI storage function of the mechanism.

## 3.2 Performance

### 3.2.1 *Timeliness of CSI*

Dependent entities should be able to locate and receive CSI in a timely fashion, to allow them to use such information in authenticating entities or verifying the signatures of entities. This feature, along with the Emergency CSI capability criterion, satisfies the requirement for an “immediate revocation service” contained in Annex II of the EU Directive on a Community Framework for Electronic Signatures.

Timeliness concerns the amount of time between the generation of CSI from the appropriate authority until this CSI becomes available (this does not include the actual dissemination of CSI to dependent entities, but only its availability) to the dependent entities. Timeliness of CSI increases when this amount of time decreases.

### 3.2.2 *Freshness of CSI*

This criterion concerns the maximum period of time between the most recent CSI generation, regarding a specific certificate, and a request for CSI regarding that certificate. In the following paragraphs, we present quantitative metrics that can be used to estimate the timeliness and freshness of CSI.

The metric that estimates the cost of having a CRL, issued at time  $t$ , and posted to  $n$  repositories, equals  $n \cdot s(t)$ , where  $s(t)$  equals the time needed to distribute the CSI to one repository. Assume that it takes  $d \cdot s(t)$  time units to forward the same CSI to the dependent entities. The freshness requirement of  $k$  time units must be larger than  $n \cdot s(t) + d \cdot s(t) = (n+d) \cdot s(t)$ , and the CRL must be updated at least every  $k - (n+d) \cdot s(t)$  time units.

A similar metric can be given for the cost (in time) to distribute a Distribution Point CRL coming from the appropriate authority to the dependent entities; the freshness requirement of  $k$  time units must now be larger than  $(d+1) \cdot s(t)$ , and the Distribution Point CRL must be issued at least every  $k - (d+1) \cdot s(t)$  time units.

For a Delta-CRL that is issued at time  $t+l$ , the time needed to have the Delta-CRLs distributed among  $n$  repositories is given by  $n \cdot (s(t+l) - s(t))$ , and it takes  $d \cdot (s(t+l) - s(t))$  time units to send the CSI

to the dependent entities. Therefore, the timeliness metric equals  $k - ((n+d) \cdot (s(t+l) - s(t)))$ , which means that a Delta-CRL must be issued at least every  $k - ((n+d) \cdot (s(t+l) - s(t)))$  time units.

For OCSP, the timeliness metric is given by  $k - n \cdot s(t)$ : it takes  $n \cdot s(t)$  time units to update the CSI of the OCSP service providers, and the dependent entities retrieve their CSI on-line from the OCSP service providers. Note however, that an OCSP responder does not need to retrieve the freshest CSI available; it has to produce a fresh digital signature, but possibly on old CSI.

### 3.2.3 *Bounded revocation*

There should be an upper time limit for new, fresh CSI to be produced and made available (this does not include the actual dissemination of CSI to dependent entities, but only its availability). Dependent entities should reject CSI they receive, if the “date of fresher CSI generation” is in the past and not in the future.

There could be another requirement concerning the time of issuance of CSI, which restricts the placement of CSI issuance in time even more. We call this *time-complete* revocation. In time-complete revocation, CSI is generated in specific moments in time, neither sooner nor later.

### 3.2.4 *Emergency CSI capability*

This requirement concerns the ability of the CSI authority to generate CSI and make it available, immediately after receiving a valid revocation request. However, this does not include the immediate dissemination of this CSI to dependent entities, but only the immediate generation of CSI. This feature, along with the timeliness criterion, satisfies the requirement for an “immediate revocation service” in Annex II of the EU Directive on a Community Framework for Electronic Signatures.

### 3.2.5 *Scalability*

If the number of the authorities and users (e.g., CAs, RevAs, dependent entities, certificate holders) increases, new obstacles in the operation of the CSI-retrieval mechanism should not emerge.

### 3.2.6 *Adjustability*

The dependent entities (or the CA and the CSI authorities) should be able to adjust the location or validation function operation, in order to create a balance between performance and protection, depending on the requirements and the risk assessment in each case. Ideally, the dependent entity should be able to adjust the location or validation function, since it is the dependent entity that takes the risk by accepting this balance between performance and protection.

## 3.3 **Security**

The CSI mechanism features presented in this section meet the requirement (Annex II of the EU Directive on a Community Framework for Electronic Signatures) for a “secure directory” and a “secure revocation service”.

### 3.3.1 *CSI disseminator authentication*

The dependent entities must verify the origin of the CSI they receive. If authentication is not used, a malicious entity pretending to be a trusted CSI dissemination entity could disseminate false CSI to dependent entities, which appears to be valid.

### 3.3.2 *CSI integrity*

The integrity of the CSI must be protected, when it is stored in the CSI repository, while it is transferred to the dependent entities and when it is stored in the dependent entities’ local repository.

The integrity protection mechanisms must ensure that a malicious entity cannot modify either the stored CSI or the CSI while in transit. If this happens, the dependent entity should be able to know that the received CSI is old, partial, or invalid in any way. Note however, that this requirement may introduce another type of attack, i.e. the Denial of Service. The information that is transferred to the dependent entity must be integrity-protected. If a system such as OCSP is used, this means that the OCSP responder must produce a (mostly expensive) fresh digital signature on each response. On the contrary, when using systems such as Delta-CRLs, a single digital signature is computed on the Delta-CRL, which is subsequently forwarded and stored without additional (integrity) protection.

### 3.3.3 CA compromise

There should be a mechanism (e.g., an Authority Revocation List (ARL) which is similar to a CRL but enumerates revoked CA certificates) for the dependent entities to know whether a CA has been compromised. There should also be a mechanism to allow a CA to recover from compromise. The effects of a CA key being compromised should be minimised.

### 3.3.4 Revocation Authority (RevA) compromise

There should be a mechanism (e.g., similar to an Authority Revocation List as explained above) for the dependent entities to know whether the authority that revokes certificates (RevA) has been compromised. This mechanism must not be the same as the one used by the dependent entities in order to receive CSI on certificates that belong to entities other than the RevA.

### 3.3.5 Contained functionality

If RevA is compromised, it should not be possible for the entities that gained control of the RevA to issue new certificates.

### 3.3.6 Availability

The CSI dissemination mechanism has to be resilient against unreliable networks, Denial Of Service (DoS) attacks, etc.

## 4 EVALUATION OF CSI-RETRIEVAL MECHANISMS

In this section, we apply our evaluation framework to Certificate Revocation Lists [6], [9] and the Online Certificate Status Protocol [11].

Certain evaluation criteria are accompanied by quantitative metrics that can, in turn, be used in order to investigate possible enhancements to the CSI mechanisms we examine. The notation that will be used for providing quantitative metrics is explained in Table 1.

<b>Table 1: Notation of quantitative metrics</b>	
$t$	linearly increasing event timescale
$s(t)$	cost of distributing CSI
$s_{pi}(t)$	cost of distributing CRL partition $pi$ to the appropriate authority or repository
$n$	number of authorities or repositories that must receive the generated CSI
$d$	number of dependent entities seeking CSI
$k$	freshness requirement (in time units)

## 4.1 Management

### 4.1.1 Feedback

We believe that it is necessary to provide the dependent entity with feedback information. An assertion regarding the status of a certificate will be interpreted differently by the dependent entity, depending on the information that led to this assertion and the local policy. Existing standardisation efforts, which are related to the CSI mechanisms we examine, do not include suggestions for the user interface of such applications or devices.

It is also necessary to define a simple way to convey this information to the dependent entity, because of the complexity of the information as well as of the possible lack of understanding of nontrivial security and cryptography issues by the dependent entity. Standardisation efforts related to the mechanisms we examine should include high-level requirements for a user interface that provides such information to the dependent entities. This would result in an even higher level of user awareness on CSI mechanisms. Finally, this information could be provided to the dependent entity at no substantial operational cost for the CSI mechanism.

### 4.1.2 Transparency

Although this is an important feature of CSI mechanisms and does not introduce an important additional communication cost, it is not, as of today, a matter of common practice among commercial or non-commercial CAs. The fact that standardisation efforts, regarding CSI location and retrieval information in the certificates, are ongoing could be a reason for that.

### 4.1.3 Delegation of revocation

OCSP does not support delegation of revocation, because OCSP does not include a specific mechanism for generating CSI. It uses the CSI generated by the CA, retrieving it possibly from a database indicated by the CA, in order to construct the CSI to be sent to the dependent entities.

Plain CRLs, Distribution Point CRLs and Delta-CRLs and the other revocation variants support delegation of revocation. The CA can designate a distinct authority (or a unit of the CA itself, distinct to the certificate-issuing unit) to issue the various CRL types. The CA has to issue a certificate for that authority, which must contain the `cRLSign` key usage attribute.

### 4.1.4 Delegation of the CSI dissemination

Delegation of CRL dissemination can be performed since CSI, in the CRLs case, is contained in a field signed by the CA. However, since the authority that disseminates CSI and the respective repository may not be trusted, there has to be a specific policy for the issuance of these CRLs. Such a policy must ensure that the CRL provided by the CSI disseminating authority is always the one that corresponds to the needs of the dependent entities at the specific moment in time when the CSI query is performed. Furthermore, the aforementioned policy must ensure that dependent entities have the ability to understand whether the CSI disseminating authority is withholding a specific (fresher) CRL variant from them.

OCSP supports two levels of CSI dissemination delegation:

1. At the first level, the authority (CA Designated Responder [11]) that disseminates CSI (OCSP signed Responses) must have a certificate issued by the CA for that purpose.
2. At the second level, the authority that disseminates CSI can be a third, distinct authority (Trusted Responder [11]) whose public key is trusted by the dependent entity.

Partitioning the CSI space and assigning the responsibility of disseminating the CSI partitions to a number of authorities could decrease communication costs, incurred by CSI dissemination.

The use of Distribution Point CRLs reduces the communication cost to distribute the CRL information from the CA to the repositories. A combined solution wherein Delta-CRLs are used to update Distribution Point CRLs reduces the communication cost even further.

When the number of revoked certificates in a system does not grow significantly over time, Delta-CRLs are the best choice.

The ideal situation involves a hybrid system, in which OCSP is used whenever it is possible to post online CSI queries, and where Delta-CRLs are used in any other case. Note however, that OCSP allows a responder to pre-compute responses and attach a validity period to each response. Using this capability, the online queries lose their real-time nature, as each response is then equivalent to an one-entry CRL.

#### *4.1.5 Delegation of the certificate path validation*

CRL, Distribution Point CRL, and Delta-CRL do not support delegation of the certificate path validation function.

OCSP partially supports delegation of the certificate path validation function. Extensions to the supported delegation have been proposed [20]. Delegation of the certificate path validation function requires the availability of a large set of trust-related information (e.g., CA certificates, CRL or CSI in other formats) to the OCSP service provider. The aggregation of this kind of information could incur communication, maintenance or other costs.

#### *4.1.6 Referral capability*

The referral capability relates to transparency. OCSP partially supports this capability, through the *serviceLocator* request extension. If the OCSP service provider does not have CSI concerning the certificate the dependent entity enquires, it may refer the dependent entity to another OCSP service. Location information for that OCSP service is contained in the OCSP Response extension *serviceLocator*.

CRL, Distribution Point CRL and Delta-CRL do not inherently support the referral capability. If an LDAPv3 [21] Directory is used as the CRL repository, the CSI provider can refer the dependent entity to other CRL repositories by using LDAP referrals [21].

#### *4.1.7 Revocation Reasons*

Most of the mechanisms mentioned in this document can disseminate information to the dependent entity, regarding the reasons for the revocation of a certificate (*CRLReason* extension [9], [11]). However, the use of those reasons as input information in a certificate path validation function has to be studied further [22], [23]. These reasons should be used in the process of certificate path validation only if they can provide results that are complete, repeatable, and compliant with the policy of the dependent entity and the policy of the authority that executes the certificate path validation function. These requirements are currently not met by the mechanisms we examine.

#### *4.1.8 Notification of revocation or suspension*

The CSI mechanisms we examine do not provide this kind of functionality. This functionality must not be integrated to the CSI mechanism; it should only be bound to the CSI generation or storage functions.

## 4.2 Performance

### 4.2.1 Timeliness of CSI

The metrics presented in the respective part of section 3 can be used in order to evaluate the use of CSI mechanisms we examine, either on their own or their joint use, in specific communities of certificates users (certificate holders and dependent entities). The initial number of certificate holders, dependent entities, CAs, CSI providers and CSI requests has to be estimated. The rate of growth of those has to be estimated as well, over time or over specific time periods. Using the aforementioned data, we could evaluate the use of the mechanisms we examine for a specific community of certificates users.

### 4.2.2 Freshness of CSI

The freshness property of CSI for CAs in Europe depends on legal requirements ([2] and national laws). A freshness value that meets these requirements should be agreed upon, using certain mechanisms (see section 4.2.1 on timeliness). Less fresh CSI should also be available, offering other advantages.

### 4.2.3 Bounded revocation

CRL, Distribution Point CRL, Delta-CRL and OCSP support bounded revocation with the use of the *nextUpdate* CRL and Response extensions respectively. These CSI mechanisms could also support time-complete revocation. This is a matter of policy.

### 4.2.4 Emergency CSI capability

All the mechanisms we examine support this capability. Immediate generation of CSI, though, does not include necessarily immediate CSI dissemination.

A CSI authority that uses one (or more) of the mechanisms we present can generate CSI immediately after receiving a validated revocation request. However, this CSI will not be made available to the dependent entities immediately, with the exception of OCSP. If the OCSP service locator retrieves CSI directly from the repository, where the CSI authority stores it, then CSI will also be made available to the dependent entities immediately.

### 4.2.5 Scalability

For CRLs and OCSP, the communication cost in time units varies linearly with the size of CSI. This is also the case when the number of authorities or CSI repositories varies.

For Distribution Point CRLs, the communication cost in time units varies linearly with the size of CSI, and this is independent of the number of partitions.

The communication cost (in time units) of Delta-CRLs grows linearly with the number of authorities or CSI repositories. This cost grows also linearly with the differences between the BaseCRL and a Delta-CRL.

The timeliness metrics depend highly on the characteristics of the communication medium, the bandwidth, and the propagation delay of the communication channels between the CA and the repositories, and of the communication channels between the repositories and the dependent entities. For the CSI mechanisms we examine, the timeliness metric decreases when the CSI size increases, and when the number of repositories grows.

Note once more that an OCSP responder may be vulnerable to DoS attacks because it has to sign every OCSP response, which is time consuming.

#### 4.2.6 *Adjustability*

None of the mechanisms we examine supports adjustability of the CSI protection level. The freshness and timeliness of the available CSI is not adjustable. However, the CA policy could provide more than one level of protection, by requiring more than one, separate running instances of the CSI mechanism. Every instance could provide CSI with different characteristics, as far as timeliness and freshness is concerned.

### 4.3 **Security**

Features discussed in this section are the ones that would have to be met to comply with CSI security requirements set forth in Annex II of the EU Directive on a Community Framework for Electronic Signatures. These features mainly address the requirements contained in the European Directive concerning the “secure directory” and the “secure revocation service”.

#### 4.3.1 *CSI disseminator authentication*

All the CSI mechanisms we examine meet this criterion. CSI disseminator authentication in CRL, Distribution Point CRL, and Delta-CRL is achieved through the verification of the digital signature in the respective revocation lists. OCSP responses are also signed, thus OCSP also meets this criterion.

#### 4.3.2 *CSI integrity*

The integrity of CSI, while in transit or while stored at the dependent entities' local storage, is protected through the digital signatures of the CSI authority. This applies to all the mechanisms we examine.

#### 4.3.3 *CA compromise*

Except for Suicide Notes, none of the mechanisms we examine meets this criterion. If the CA is compromised, the entity who gained control of the respective CA keys is able to revoke certificates or issue new ones at will, until the CA compromise information reaches the dependent entities through an Authority Revocation List (ARL) or another, possibly out-of-band mechanism (e.g., in case the CA private key is no longer available to the CA personnel and it is only available to the entity who illegally gained control of it).

#### 4.3.4 *RevA compromise*

OCSP does not provide for delegation of certificate revocation. The authority that revokes certificates is the CA itself. In this case, dependent entities will be informed of a possible CA compromise through an ARL. The same applies for the other CSI-retrieval mechanisms we examine.

#### 4.3.5 *Contained functionality*

If the authority that disseminates CSI is a Trusted Responder, a CA Designated Responder [11], or a CA (CSI is distributed in CRL, Distribution Point CRL, or Delta-CRL and the CA uses a separate CA key to authenticate CSI), then the compromise of the keys used by these authorities do not enable the entities who gained control of these keys to issue new certificates. However, previously revoked or suspended certificates can be made valid again.

If the authority that disseminates CSI is a CA, and it uses the CA certificate signing key in order to sign CRLs, Distribution Point CRLs or Delta-CRLs as well, then a compromise of that authority (that is, the CA) will enable the entity that gained control of the respective keys to revoke certificates and issue new ones at will.

#### 4.3.6 Availability

There are no mechanisms in CRLs [6], [9] to protect the availability of CSI. If LDAP Directories are used as CSI repositories, LDAPv3 [21] replication and referral mechanisms could be used in order to increase the availability of CSI.

The OCSF *serviceLocator* extension and the mirroring of the CSI repositories used by OCSF in order to generate CSI from could increase the availability of CSI provided by OCSF.

## 5 DISCUSSION

We evaluated CRL variants and OCSF, using our evaluation framework, and we defined qualitative and quantitative metrics for estimating the timeliness, freshness and scalability of these mechanisms. The criterion of adjustability is not met by any of the mechanisms; therefore, freshness and other metrics are not adjustable by the entities that take the risk of trusting CSI.

Furthermore, none of the mechanisms we presented meets the feedback criterion, which we believe is crucial for the efficient operation of CSI mechanisms. Although the mechanisms we examined support the inclusion of revocation reason information in CSI, they do not process this information within their certificate path validation functions. There are still issues to resolve [22], before this is feasible. Moreover, revocation notification is not inherently supported by any mechanism, though external notification procedures could very well be synchronised with CSI mechanisms. The same applies for the referral capability of CSI mechanisms.

The consequences of having the CSI authority key compromised are in some cases contained while the consequences of having the CSP key are not; if the CSP signing key is compromised, certificates can be issued and revoked at will by the entity who gained control of this key. Note that it is generally considered bad practice, i.e. to use a single key pair for different purposes. It should therefore be avoided to reuse the CSP signing key for CSI-signing purposes.

All the mechanisms we present support delegation of CSI dissemination, bounded revocation, CSI disseminator authentication and CSI integrity protection. These are well-supported features of the mechanisms. Some of the mechanisms we presented could meet the transparency criterion; however, the respective transparency features of the mechanisms have not been widely used by the industry.

Delegation of certificate path validation is partially supported by OCSF, while delegation of revocation is only supported by CRL variants. Finally, Emergency CSI generation is supported by all mechanisms. However, rendering available the CSI that was generated in an emergency would cause problems in the policy of a CSI authority, should this authority wish to support time-complete revocation.

## 6 CONCLUSION

The evaluation framework we presented can be used by the research community for further research on CSI mechanisms, either for improving the existing ones or for developing new ones. The industry can also make use of this framework; until now, high-level PKI requirements were used on specific CSI mechanisms, based on empirical methods or ad hoc research. Our framework can be used by PKI implementers and policy makers to select the CSI mechanism or mechanisms that will be used in a PKI, depending on the underlying, high-level requirements.

Future work regarding our evaluation framework includes adding even more quantitative metrics and drawing more conclusions based on the comparative evaluation of CSI mechanisms. Quantitative metrics and corresponding graphs concerning the scalability of a mechanism can be

used as a quick reference in the selection of a CSI mechanism to use in a specific PKI environment, depending on the underlying requirements.

The mechanisms we have been evaluating, as a case study for our evaluation framework, are among the most widely used ones. These mechanisms do not meet the adjustability, feedback and transparency criteria, among others. We consider these criteria crucial for a CSI mechanism. These criteria, if they were met by a CSI mechanism, would help dependent entities comprehend the CSI mechanism and take advantage of its full potential. Further research on the most widely used CSI mechanisms is necessary, for them to meet these criteria. Alternatively, a meta-mechanism (i.e. a mechanism that operates as a functional wrapper of the aforementioned CSI mechanisms), meeting the adjustability, feedback and transparency criteria could be researched and implemented in order to deliver the CSI service in a more effective manner.

We believe our evaluation framework will become a useful tool for PKI researchers and developers. The latter can use our framework in order to select a CSI mechanism to implement in PKI environments, depending on the requirements of the specific PKI. Our framework can also be used to further develop or customise existing CSI mechanisms, depending on the requirements of a specific PKI implementation.

## ACKNOWLEDGEMENTS

This work was partially funded by the European Commission (Directorate General III, contract #ETD/99/502536: "Study on the Scalability of Certificate Revocation and Certificate Suspension and Proposals for Enhancements on the Respective Mechanisms"). This work was also partially supported by the Concerted Research Action (GOA) Mefisto-666-2000/06 of the Flemish Government.

## REFERENCES

- [1] Ioannis S. Iliadis, Diomidis Spinellis, Sokratis Katsikas and Bart Preneel, A Taxonomy of Certificate Status Information Mechanisms, Proceedings of Information Security Solutions Europe ISSE 2000, Barcelona, Spain, September 2000. European Forum for Electronic Business.
- [2] Directive 1999/93/EC of the European Parliament and of the Council on a Community framework for electronic signatures, 13 December 1999.
- [3] Nilsson H, Van Eecke P., Medina M., Pinkas D., Pope N., European Electronic Signature Standardization Initiative, ICTSB, Final Report of the EESSI Expert Team, 20 July 1999.
- [4] Berkovits S., Chokhani S., Furlong J. A., Geiter J. A., and Guild J. C., Public Key Infrastructure Study: Final Report. Produced by the MITRE Corporation for NIST, Apr. 1994.
- [5] US National Institute of Standards and Technology, 'A Public Key Infrastructure for US Government unclassified but sensitive applications', September 1995.
- [6] ISO/IEC 9594-8 (1994), Open Systems Interconnection — The Directory: Authentication Framework.
- [7] Chokhani S., Ford W., Internet X.509 Public Key Infrastructure, Certificate Policy and Certification Practices Framework, (Status: INFORMATIONAL) Request for Comments 2527, March 1999 (available at <http://www.ietf.org/rfc/rfc2527.txt>).

- [8] Adams C., Farrell S., Internet X.509 Public Key Infrastructure Certificate Management Protocols, Request for Comments 2510, 1999 (available at <http://www.ietf.org/rfc/rfc2510.txt>).
- [9] Housley R., Ford W., Polk W., Solo D., Internet X.509 Public Key Infrastructure, Certificate and CRL Profile, IETF PKIX Working Group, Request for Comments 2459 (Category: Standards Track), January 1999 (available at <http://www.ietf.org/rfc/rfc2459.txt>).
- [10] Santesson S., Polk W., Barzin P., Nystrom M., IETF PKIX Working Group, Internet X.509 Public Key Infrastructure, Qualified Certificates Profile, Internet Draft, February 2000 (available at <http://www.ietf.org/internet-drafts/draft-ietf-pkix-qc-03.txt>).
- [11] Myers M., Ankney R., Malpani A., Galperin S., Adams C., Internet X.509 Public Key Infrastructure, Online Certificate Status Protocol, IETF PKIX Working Group, Request for Comments 2560 (Category: Standards Track), January 1999 (available at <http://www.ietf.org/rfc/rfc2560.txt>).
- [12] Housley R., Ford W., Polk W., Solo D., IETF PKIX Working Group, Internet X.509 Public Key Infrastructure, Certificate and CRL Profile, Internet Draft, October 1999 (available at <http://www.ietf.org/internet-drafts/draft-ietf-pkix-new-part1-00.txt>).
- [13] Micali S., Efficient Certificate Revocation, Technical Memo 542b, Laboratory for Computer Science, Massachusetts Institute of Technology, March 1996.
- [14] Ronald L. Rivest. *Can we eliminate certificate revocation lists?*, Financial Cryptography FC'98, volume 1465, pg 178–183, Anguilla, British West Indies, February 1998, Springer.
- [15] Stubblebine S. G., Recent-Secure Authentication: Enforcing Revocation in Distributed Systems, In Proceedings IEEE Symposium on Research in Security and Privacy, pages 224–234, May 1995, Oakland
- [16] M. Naor, K. Nassim. Certificate revocation and certificate update, Proceedings of the 6th USENIX Security symposium, pg 217–228, January 1998.
- [17] P. Kocher. On certificate revocation and validation, Financial Cryptography FC'98, volume 1465, pg 172–177, Anguilla, British West Indies, February 1998, Springer.
- [18] Adams C., Zuccherato R., A General, Flexible Approach to Certificate Revocation, White Paper, Entrust Technologies
- [19] Berners-Lee T., Fielding R., Masinter L., Uniform Resource Identifiers (URI): Generic Syntax, (Draft Standard), August 1998 (available at <http://www.ietf.org/rfc/rfc2396.txt>).
- [20] Hallam-Baker P., OCSP Extensions, IETF Internet Draft, September 1999.
- [21] Chadwick D. W., Internet X.509 Public Key Infrastructure, Operational Protocols: LDAPv3 (Category: Standards Track), August 1999.
- [22] Fox B., LaMacchia B., Certificate Revocation: Mechanics and Meaning, In Proceedings of Financial Cryptography 98, LNCS 1465, New York, Springer – Verlag.
- [23] Iliadis I., On the Dissemination of Certificate Status Information, MSc thesis, Department of Mathematics, Royal Holloway and Bedford New College, University of London, UK, 1999.