PROCEEDINGS OF SCIENCE

A Study of Certification Authority Integration Model in a PKI Trust Federation on Distributed Infrastructures for Academic Research

Eisaku Sakane*, **Takeshi Nishimura, and Kento Aida**
National Institute of Informatics
E-mail: sakane@nii.ac.jp, takeshi@nii.ac.jp, aida@nii.ac.jp

Among certification authorities (CAs) in an academic PKI trust federation such as IGTF (Interoperable Global Trust Federation), most of the academic organizations that operate a CA install by themselves the CA equipment in their building. To keep the CA trustworthy, it is necessary to maintain specialized CA equipment and to employ specifically trained operators. The high cost thereby incurred for CA operation weighs heavily on the CA organization. For research institutes whose primary duties are not the CA operation, the burden of the high cost of CA operations is an earnest problem, and cost reduction by increasing the efficiency of the operation is an important issue.

Instead of focusing on any further operational optimization of a single individual CA, in this paper we will review cost reductions by way of integrating more than one CA in a PKI federation. This paper considers the issuing and registration authorities that constitute a CA, and proposes the following integration model: it integrates the issuing duties, and each organization carries out the registration duties as before. In the proposed model, integrating the issuing duties means that one issuing authority (IA) takes over the duty of the other IA. Since each registration authority (RA) performs the registration duty as usual, most of procedures such as the application process to obtain certificates remain unchanged, so that it does not confuse the users.

Based on this proposed model, we discuss how to connect the superseding IA with the RA(β) operated by the organization that closes its own IA(β). Among possible connections, we examine not only a direct connection between the superseding IA and the RA(β) but also a connection putting the RA(α) – operated so far by the organization that operates the remaining IA – inbetween them as a proxy. Furthermore, we augment the certificate policy of the superseding IA so that it is compatible with the policy of the RA(β). We also discuss an applicability of existing CA profiles such as MICS (Member Integrated Credential Service) profile and its extension.

International Symposium on Grids and Clouds 2016
13-18 March 2016
Academia Sinica, Taipei, Taiwan

*Speaker.

PoS(ISGC 2016)015

1. Introduction

International academic research projects need an authentication and authorization infrastructure (AAI) for secure sharing of data generated by large research facilities or to make joint use of supercomputers. Although such infrastructure crosses countries, the AAI should be offered to (only) the academic research community centering around a world-wide joint-use facility. Therefore, in an infrastructure based on the public-key infrastructure (PKI) technology, the representative institute in each country that participates in the research collaboration often builds by itself a certificate authority (CA). Each CA is then operated according to a commonly-agreed profile, so that a trust federation is created on the distributed infrastructures for academic research. The Interoperable Global Trust Federation (IGTF) [1, 2] is one of such trust federations.

To keep CA trustworthy, it is necessary to maintain the CA equipment and to hire specifically trained operators. The CA equipment typically contains expensive equipment such as a certified hardware security module (HSM). For academic research institutes whose essential duties are not the CA operation, it is getting harder every year to maintain and operate a CA due to the high cost of its operation. To make the CA operation sustainable, cost reduction – mainly by increasing the efficiency of the operation – is an important issue. However, it would be difficult to further raise the efficiency of the operation of a single CA in itself.

This paper discusses how to make the operation of CAs more efficient by integrating jointly more than one CA into a PKI federation. There is a method to straightforwardly integrate CAs as follows: current top-level "root" CAs are turned into intermediate CAs and a new root CA is built. However, this integration model would not reduce each CA's operational duties such as issuing, revoking, and identity vetting. In addition, building the new root CA that covers different academic research communities would actually raise the total costs of CA operations. It is important to provide a suitable integration model to reduce the CA operation burden as a whole.

This paper proposes an integration model to increase the efficiency of CA operations as a whole. We take note of the distinct roles of issuing and registration authorities constituting a CA, and consider the following integration model: it integrates the issuing duties, and each operating organization carries out the registration duties as before. In this model one issuing authority (IA) takes over the duty of the other IA. The duty of each registration authority (RA) basically remain unchanged. However the RA(β)[1] operated by the organization that closes its IA(β) must be connected to the superseding IA for conveying requests such as certificate signing. We discuss how to connect the superseding IA with the RA(β). We also consider its certificate policy in the CA integration model.

The reminder of this paper is organized as follows. In Section 2 we describe a basic idea about CA integration and issues to be addressed based on a typical CA architecture concerned. Section 3 presents solutions to the connection issue between RA and IA. Section 4 makes discussion about the proposed integration model. Section 5 refers to related work. Finally, Section 6 concludes the paper.

[1] We use the Greek alphabet to distinguish IAs and RAs as well as CAs.

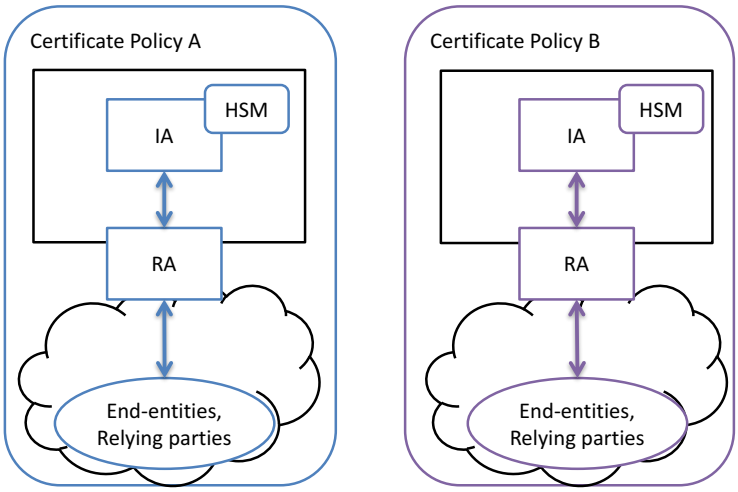

Figure 1: Typical CA architecture.

2. Basic Idea and Issues

This section describes a basic idea about CA integration and issues to be addressed. First, we present a typical CA architecture in an academic PKI trust federation. After that, we discuss the integration of such CAs and adopt a CA integration model. Finally, we detail the issues related to the realization of the integration.

2.1 CA Architecture

We consider CAs that obey an authentication profile provided by IGTF: such authentication profiles [5, 6] managed by IGTF characterize CAs in e-Science infrastructures.

Figure 1 illustrates two PKI domains. PKI participants in each domain are the CA (the elements in the black rectangle), end-entities (natural persons and network entities), and relying parties. The PKI domain is characterized by a certificate policy. The certificate policy – in addition to certification practice statements (CPS) – is formulated by means of the framework defined by RFC 3647 [3]. The left-hand side of Fig. 1 obeys the certificate policy A, and the right-hand side the certificate policy B. In general, CA can be divided into an IA and an RA. The IA and RA servers are usually set up on separate machines because the IA handles the private key that is used to sign the public key of an end-entity. The IA server is located in a private network in order to securely protect the private key of the CA, and communicates with only the RA server. The square drawn with dashed-line in Fig. 1 denotes a private network segment. This separate IA-RA model is one of recommended models in the on-line CA guidelines [7] offered by the IGTF. The IA server usually uses an HSM to store the private key of the IA. The RA server also communicates with clients via the Internet and receives the requests from clients such as certificate signing request (CSR). The interactions between the IA, RA and clients are often prescribed by the Certificate Management Protocol (CMP) [4] or its variants.

2.2 CA Integration

Let us suppose that there are two CAs operated by different organizations respectively and

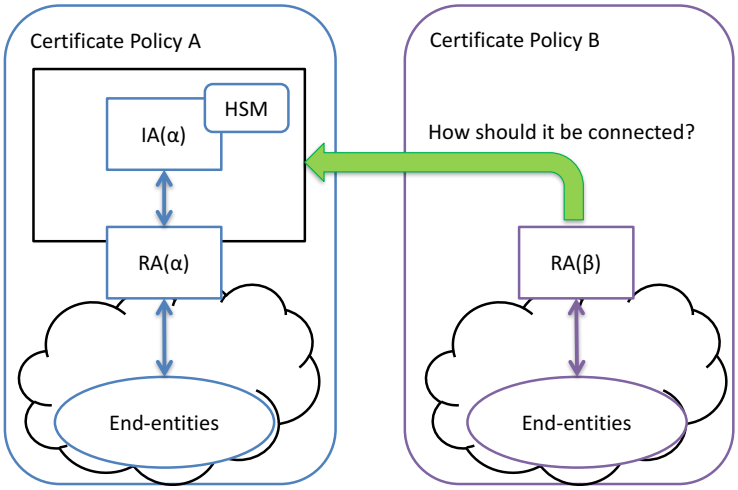

Figure 2: CA integration model.

that each CA covers different research communities (Fig. 1). We assume the following as premise of the argument: we do not discuss how to build *from scratch* a CA that covers each research community. If CA-operating organizations are joined in a merger CA building from scratch may be meaningful. However, such a merger would not frequently happen. Also, we do not consider that multiple CA-operating organizations outsource CA duties or part of them to a commercial CA vendor and so share the expenses in the CA outsourcing.

We discuss how to integrate existing CAs *without being forced to a drastic change* in order to reduce the cost of the CA operations. CA integration model should fulfill the following requirements:

1. From a user's point of view, CA service procedures for users should remain unchanged in as far as possible.

2. From an operator's point of view, the changes of CA operational procedures should preferably be small.

We investigate the integration possibility of each component of CA in order, namely RA integration and an IA one.

On the RA side, it is difficult for one RA to vet user identities in the other community because the identity vetting is an effort-intensive duty. Basically, the research community should have the responsibility to vet the identities of the users who belong to the community. If each RA is independently operated as before, interference in certificate policies would be kept down to a minimum.

On the IA side, issuing operations are limited exclusively to management of the CA private key and to respond to the requests from the RA. It would be unnecessary to operate IA at one's own expense if one RA can rely on another IA. Such IA system outsourcing does not prevent the CA operating organization from providing an authentication infrastructure based on PKI.

The integration model considered in this paper is as follows: one IA is terminated, the other IA is charged with both issuing operations, and each one's registration duty basically remains

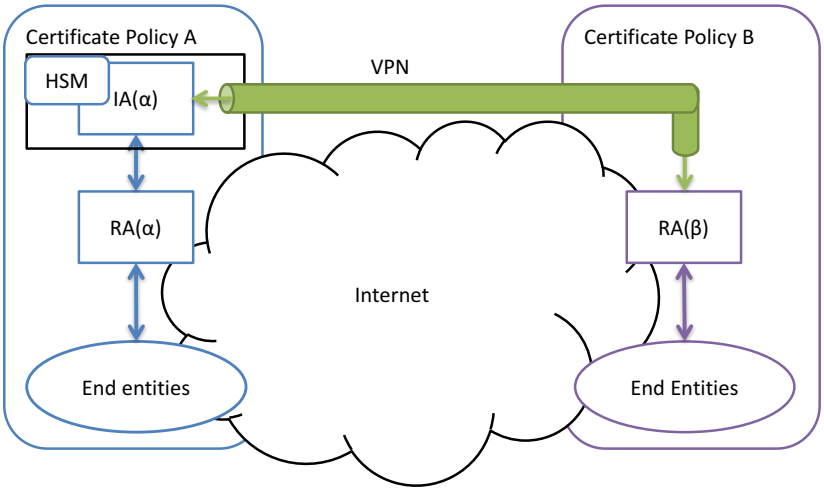

Figure 3: Direct connection.

unchanged as before. In this sense, the proposed model is a partial integration of both CA. To illustrate the model, we suppose that the organization operating a CA as per the certificate policy B closes its IA and the other organization – operating as per the certificate policy A – takes over the issuing duties of the closed IA (Fig. 2).

One of the issues to be addressed is how it connects RA with outside IA in order to send material such as certificate signing requests. In Fig. 2, it is shown how to connect the RA(β) with the IA(α). We discuss solutions to this issue in next section. In addition, both certificate policies should be revised so that the IA(α) can issue certificates to end-entities in the right-hand side of Fig. 2 and the RA(β) can send the issuing requests to the IA(α).

3. IA-RA Connections in the CA Integration

This section presents two types of solution to the RA-to-IA connection issue mentioned in the previous section. There are two methods of connection, logically direct and indirect connections between IA and RA.

3.1 Direct connecting

To establish a logically direct connection between RA(β) and the outside IA(α), a virtual private network (VPN) such as Layer 2 VPN is needed because the IA(α) is located in a closed network for security reason mentioned in Sec. 2.1. In this model (Fig. 3), the RA(β) authenticates an end-entity client and checks the validity of the request as usual. After that, the RA(β) sends the request to the IA(α) via the VPN. The IA(α) receives the request from the RA(β), processes it, and then returns the result to the RA(β). The RA(β) receives the result from the IA(α) and finally returns it to the end-entity client. Note that the interaction between the RA(β) and end-entities remains unchanged.

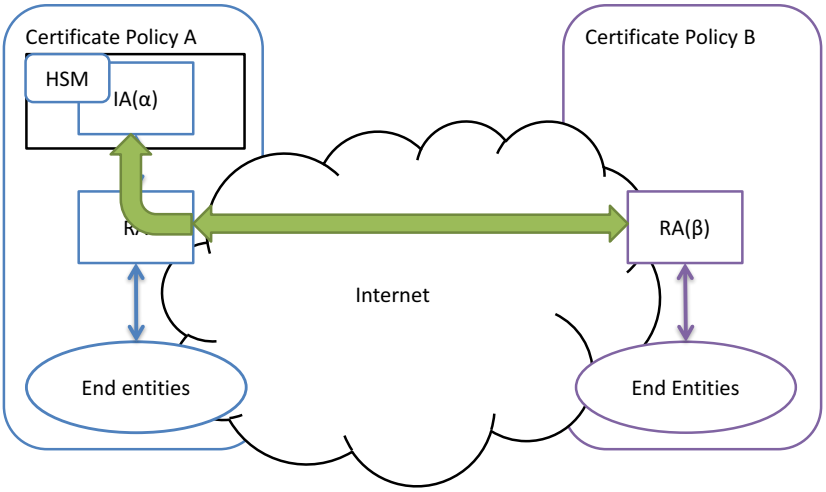

Figure 4: Relaying RA connection

3.2 Relaying

As another solution, we consider logically indirect connection, namely a relaying RA. Figure 4 illustrates a relaying RA model. In this model, the RA(β) sends a user's request not to the outside IA(α) directly but the RA(α). The RA(α) conveys the request to the IA(α). After receiving the result of the request the RA(α) returns it to the RA(β). The communication between the RA(β) and its supporting end-entities is the same as the direct connecting model. On the connection between the RA(α) and RA(β), two kind of the requests from the RA(β) and from the end-entities in the left-hand side of Fig. 4 are essentially the same for the RA(α). Therefore, the connection between two RAs does not necessarily need to establish a L2VPN link. A connection on the same level with the connection between the RA and end-entities is sufficient for the connection between RAs. This model also keeps each interaction between RA and end-entities of both domains unchanged.

We consider the design of the relaying RA model in further detail. Figure 5 shows the technical design of the relaying RA model. In Fig. 5, IAd(α) denotes a daemon program via which issuing functions are implemented on the IA(α) server. Similarly, RAd(α) denotes a daemon program on which registration functions are implemented on the RA(α) server. This design of the relaying RA model sets up a new component "RA daemon compatible with IA" to relay requests from the RAd(β) to the IAd(α). This "RAd-compatible" service behaves like an IA daemon to the RAd(β) but is not an IA daemon itself. Specifically, the RAd-compatible service does not actually carry out certificate issuance. Pretending to be the IA, the RAd-compatible service receives the request when a issuing request was sent from the RA(β), relays it to the IA(α), and returns the result to the RA(β). Thanks to the behavior of the RAd-compatible service, the RAd(β) does not need to modify the request message to IA and can use the same message format as before.

As mentioned above, a dedicated link such as L2VPN is not the only choice of the connection between the RA(β) and RA(α). A connection with mutual authentication on Transport Layer Security (TLS) provides enough functions to establish a secure connection. The initial establishment of the TLS connection would be done in advance with an out-of-band method such as a face-to-face meeting with the CA personnel concerned.

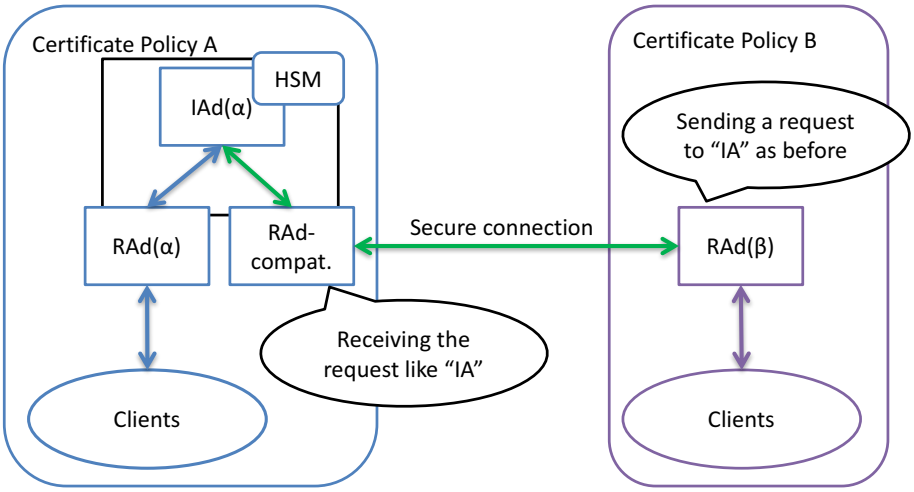

Figure 5: Design of relaying RA connection.

Finally, we summarize the specifications for the IAd and RAd-compatible service in the relaying RA model:

- the IA daemon can handle the requests from multiple RA daemons as per each certificate profile if needed.
- the RA daemon compatible with an IA daemon can authenticate the outside RA daemon and relay a request from the outside RA daemon to the inside IA daemon.

The other daemons, RAd(α) and RAd(β), do not need to be changed.

4. Discussion and Future Issues

In this section we discuss the IA-RA connections in the proposed integration model described in the previous section. We also examine four patterns of combination of two certificate policies provided by the IGTF.

4.1 IA-RA Connections

Using logically direct connections such as an L2VPN link, it has the following advantages: it keeps unchanged the RA interface to end-entities in the research community and needs basically no software development. The end-entities in each research community can interact with the RA as before, including for the initial vetting of identity. Since the RA connects with the outside IA with the L2VPN link, operating organizations can select a preferred software implementation to build a PKI as long as software package or alternatively the IA-RA interface specification used by each organization is the same one. It will be easy to establish a VPN between CAs in the same domestic region because the organizations concerned can expect a support from National Research and Education Network (NREN). A lot of time and operational effort may be needed to connect CA components across countries via L2VPN because there are more arrangements or negotiations between NRENs or network operation centers than those of the domestic connection.

Using a logically indirect connection such as a relaying RA, the RA interface to end-entities in the research community also remains the same, and there is no problem on network topology because existing network links can be used as it is. However, this approach depends on the software implementation of CA because the relay function of the RA needs to be newly developed.

On designing the relaying RA model in Sec. 3.2, we introduced an RA daemon compatible with the IA interface. As another design, we can set up an RA daemon that behaves like an end-entity client to the outside RA. Concretely, the RA(β) in Fig. 4 receives a request from an end-entity client as usual and sends the request to the outside RA(α) as a client. If the RA(α) is not appropriately adapted, it does not distinguish the requests from any end-entity client requests in the left-hand side of Fig. 4 and those from the RA(β) client, and verifies the digital signature of the certificate signing request. However, this verification has been done when the RA(β) received the request from the end-entity client in the right-hand side of Fig. 4, hence the verification at RA(α) is redundant. To remove this redundancy the RA(α) has to recognize the request sent from the RA(β). Therefore, the design modifying the RA(β) so as to behave like a client is not as simple as the design introducing the compatibility service at RA(α) with IA(α).

With both logically direct and indirect connections, the number of issued certificates of the IA(α) will increase. However there will be no problem of the performance of certificate issuance if the total number of end-entities is of the same order of magnitude because the IA can automatically sign the verified CSRs with software. The proposed integration model reduces the total costs of CA operations because the model dispenses with the maintenance and operation concerned with the IA(β). Since the IA(β) is closed, the trust anchor of the end-entities in the right-hand side of Fig. 2 obviously changes from the IA(β) to the IA(α). However no problem occurs because the relying parties rely on the PKI trust federation to which both CAs belong.

Since the above model is still at the design stage, an implementation of the proposed model should be put into practice as a future work. As a software package for building a PKI, we will choose NAREGI-CA software [8]. The NAREGI-CA software provides IA and RA daemon programs as well as client one. The on-line interaction between CA components is defined with Lightweight Certificate Management Protocol (LCMP) [9] that is a variant of CMP. We will make an implementation of the proposed integration model to the NAREGI-CA software package in the near future.

4.2 Policies

To discuss the feasibility of the CA integration model proposed in this paper, we consider four cases in IGTF-accredited CAs such as Classic and MICS CAs.

The Classic Profile represents a CA as per the "Authentication Profile for Classic X.509 Public Key Certification Authorities with secured infrastructure" [5]. The MICS Profile represents a CA as per "Profile for Member Integrated X.509 Credential Services with Secured Infrastructure" [6]. Both authentication profiles describe the minimum requirements on each CA. A CA based on Classic or MICS profile can issue long-term (at most one year and one month) certificates to end-entities. There is a big difference between Classic and MICS in RA duties, namely the initial vetting of identity for any end-entities. While RA in Classic is responsible for identity vetting, the RA in MICS has delegated the responsibility for it: in MICS a primary Identity Management (IdM) system is responsible for identity vetting. Any primary authentication service can be regarded

as a MICS IdM system as long as the primary authentication service satisfies the requirements of the MICS profile. The MICS IdM system may be operated by the organization independent of the MICS RA. According to the MICS profile, the initial vetting of identity for any entity in the primary authentication system that is valid for certification should be based on a face-to-face meeting and should be confirmed via photo-identification and/or similar valid official documents. This requirement is the same as the identity vetting rules in the Classic profile.

As mentioned in Sec. 2.2, the proposed CA integration model in this paper requires revising each certificate policy so that the IA(α) can issue certificates to end-entities for whom the RA(β) is responsible in terms of the initial vetting of identity, and for which the RA(β) will send the issuing requests to the IA(α). The certificate policy A that the IA(α) is based on should be revised regardless of the authentication profile which the IA(α) is subject to. Whether the certificate policy A can be revised depends upon the objective of the representative organization operating the IA(α). For example, if the representative operating organization supports a nation-wide academic information infrastructure, it may be possible for the IA(α) to issue certificates to end-entities in the scientific research community that the RA(β) covers.

For the RA(β) to send the requests to the IA(α), the RA(β) should be one of RAs that communicate with the IA(α). The differences between the certificate policies A and B are the target audience of the IA, and the type of certificates the IA issues. There is no difference between kinds of certificate profile for each community because those are in the same PKI trust federation. Therefore, the difference between the policies is just the community that end-entities belong to, so that the policy for the RA(β) can coexist with the one for the RA(α) as long as the IA(α) communicates with the RA(β).

In addition, it might be necessary to conform the RA(β) to the authentication profile that the RA(α) follows. To express possible combinations, we use the symbol "(Profile, Profile)". For example, the symbol (Classic, MICS) stands for the case that the IA(α) and RA(α) obey the Classic profile and the RA(β) obeys the MICS one. Finally, we consider four combinations below.

(Classic, Classic) Since the authentication profile for the RA(β) is the same one for the RA(α) in this case, the RA(β) could become one of distributed RAs in the Classic profile.

(Classic, MICS) In this case, the MICS RA(β) could not be formally regarded as the Classic one because the RA(β) can be responsible for RA duties without identity vetting. It is necessary to construct a framework that enables the whole components including the RA(β) and its IdM system to be considered Classic RA.

(MICS, Classic) In this case, the Classic RA(β) could become the MICS RA and IdM. Namely the identity vetting in the RA(β) could be regarded as the one of MICS IdM, and the remainder could be regarded as MICS RA.

(MICS, MICS) This case is similar to (Classic, Classic). The RA(β) could become one of distributed RAs in the MICS profile.

Currently the integration with combination (Classic, MICS) is not straightforward. However such situation is considered as a rare case because the MICS profile was created later than the

Classic one. It is feasible to integrate CAs with the other combinations. However, for the integration with (MICS, Classic), it is necessary to create guidelines for the MICS Identity Management System as future work.

5. Related Work

This section refers to a framework called RPS and to existing CAs concerned.

5.1 RPS

RPS is the abbreviation for Registration Practice Statement. Standards for RPS has already been discussed in the IGTF. RPS can be considered as a subordinate document to the CPS. Reference [10] suggests that separating RAs from the CA function has benefits that are useful for more efficient trust processing of the overall system. Therefore it is worthwhile applying the RPS framework to the proposed integration model in this paper.

5.2 Preceding CA transitions following this model

AusCERT PKI Certificate Service [11] powered by QuoVadis [12] and Comodo [13] offer Australian and New Zealand education and research organizations certificates for a wide variety of uses. AusCERT behaves as an RA of QuoVadis and the RA covers not only Australia but also New Zealand. Although this paper excludes the situation such as AusCERT, this can be considered as one of CA models that realize efficient CA operations.

Academia Sinica Grid Computing Certificate Authority (ASGCCA) [14] obeys the Classic profile. The RA of the ASGCCA covers Indonesia, Philippine, and Vietnam as of August 2015. The ASGCCA delegates the authentication of individual identity to RAs and prescribes registration procedure to be an RA of the ASGCCA, so that the ASGCCA can issue certificates to end-entities in the other countries as well as Taiwan. Hence the ASGCCA can be considered as the CA integration with the combination (Classic, Classic) discussed in Sec. 4.2. This paper discusses the integration with not only the combination (Classic, Classic) but also combinations including MICS. The ASGCCA is a noteworthy activity of CA in Asia Pacific because the ASGCCA is introducing the RPS framework.

6. Summary

In this paper, we consider an integration model of certificate authorities in a PKI trust federation such IGTF. Our contribution in this paper is as follows:

- We proposed a suitable integration model without being forced to a drastic change in order to reduce the cost of the CA operations.
- We evaluated two connection types between IA and RA in the proposed integration model.
- We considered CA integration with four combinations of Classic and MICS profiles provided by IGTF.

We will implement the proposed relaying RA model to the NAREGI-CA software package and perform demonstrative evaluation.

References

[1] Interoperable Global Trust Federation, https://www.igtf.net/

[2] D. Simmel, S. Rea, and A. Stolk, *An Introduction to The Americas Grid Policy Management Authority (TAGPMA) and the International Grid Trust Federation (IGTF)*, http://www.tagpma.org/files/CLCAR-Paper15-Simmel-Rae-Stolk.pdf

[3] S. Chokhani, W. Ford, R. Sabett, C. Merrill, and S. Wu, *Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework*, https://tools.ietf.org/html/rfc3647

[4] C. Adams, S. Farrell, T. Kause, and T. Mononen, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*, https://tools.ietf.org/html/rfc4210

[5] D. Groep (Ed.), *Authentication Profile for Classic X.509 Public Key Certification Authorities with secured infrastructure*, Version 4.4 (2012), https://www.igtf.net/ap/classic/

[6] D. Simmel (Ed.), *Profile for Member Integrated X.509 Credential Services with Secured Infrastructure*, Version 1.3 (2013), https://www.igtf.net/ap/mics/

[7] D. Groep (Ed.), *Guidelines for On-line PKI Certification Authorities*, Version 1.0 (2014), https://www.eugridpma.org/guidelines/online-cas/

[8] NAREGI-CA development, https://ca-dev.naregi.org/

[9] E. Sakane, K. Aida, and K. Motoyama, *An Improvement in On-line Interactions between Public Key Infrastructure Components at Certificate Renewal*, in proceedings of *the 1st International Workshop on Cloud Computing & Applications*, (2012) 37.

[10] Scott Rea, *Standards for Registration Practices Statements*, http://agenda.nikhef.nl/getFile.py/access?contribId=6&resId=0&materialId=slides&confId=1890

[11] AusCERT PKI Certificate Service, https://cs.auscert.org.au/

[12] QuoVadis, https://www.quovadisglobal.com/

[13] Comode, https://www.comodo.com/

[14] Academia Sinica Grid Computing Certification Authority (ASGCCA), http://ca.grid.sinica.edu.tw/