Challenges of Federating National Data Access Infrastructures

Margus Freudenthal and Jan Willemson

Cybernetica, Ülikooli 2, Tartu, Estonia
{margus,janwil}@cyber.ee

Abstract. X-Road is a secure and scalable database access middleware originally developed in Estonia in early 2000s. In 2014, a decision was taken to also deploy X-Road infrastructure within Finland, hence facilitation cross-national federation. Even though being very close both geographically and culturally, the legislation, technology and best practices used by the two nations differ. This paper discusses the nature and implications of these differences in the context of federated installation of the infrastructure.

Keywords: Secure database access, cross-national security infrastructure federation

1 Introduction

By late 1990s, the level of computerization in both public and private sectors had reached the stage where large-volume digital data exchange between organizations became both feasible and necessary. Various government registries implemented electronic interfaces that could be used to query data from the registry. However, these interfaces suffered from two problems.

First, each registry implemented the interfaces independently, often using a proprietary protocol implied by the technology used. Hence, when an organization needed interfaces to several other registries, new interfaces had to be implemented for each one of them almost from scratch.

Second, as digital data exchange became more widespread, security aspects of the queries required more and more attention. On one hand, registries often contain confidential personal data, hence access to it must be tightly controlled. On the other hand, the client requests the data to make a (possibly costly) decision based on it. Hence it must be possible to verify the integrity and authenticity of the received data. However, the registry interfaces had varying levels of security, depending on the implementer.

Thus there was clear need to enhance standardization both from interoperability as well as security point of view. To achieve this goal in Estonia, a unified data exchange middleware called X-Road was launched in December 2001 [9, 7]. During the next years it has evolved with addition of features and evolutionary changes to protocols and data formats, reaching version 5 in 2010.

By the end of 2016, X-Road had 1789 connected services by 246 service providers. Altogether, 975 member organizations exchanged roughly 575 million transactions per year[1]. For comparison – the population of Estonia is slightly over 1.3 million which gives more than 430 transactions per inhabitant per year.

Development of next generation of X-Road started in 2014. It was based on product prototype developed two years earlier by Cybernetica AS, the developer and maintainer of the original X-Road software. One of the goals for new version was better support for international deployment and cross-border electronic services. The new version, version 6, was a fresh start and did not use the same technical solution and protocols as the previous versions 1 to 5.

X-Road version 6 was also licensed to Finnish government and the source code was published on Github[2]. Currently it is maintained in cooperation by Estonian and Finnish governments. X-Road version 6 is being implemented in both Estonia and Finland.

There is also a commercial branch of X-Road developed by Cybernetica, the company responsible for development and maintenance of the previous X-Road versions. It is called Cybernetica UXP® (Unified eXchange Platform) and is based on the same product prototype as X-Road version 6. UXP includes improved versions of X-Road components as well as additional components that simplify implementation in other countries. As of 2017, this product has been installed in Haiti, Namibia and also in a pilot environment in the United Kingdom. UXP uses the same data formats and message standards as X-Road, maintaining full service level compatibility.

Since X-Road technology is already deployed in several countries, one major prerequisite for implementing cross-border digital services is fulfilled. However, actual implementation of such services introduces an extra layer of complexity even for countries that are culturally close and otherwise friendly (such as as Estonia and Finland). For example, there are still noticeable differences in the legal systems and trust levels pro-

[1] `https://www.ria.ee/ee/x-tee-statistika.html`
[2] `https://github.com/ria-ee/X-Road`

vided by the trust service providers (like certification authorities) may be incompatible.

The aim of the current paper is to describe technical implementation of X-Road federation and explore the problems that may arise as a result of the differences between federated installations.

2 X-Road infrastructure

General structure of X-Road infrastructure as deployed in one country is described in detail in [8]. In this paper, we will provide a concise overview; see Figure 1.

An X-Road installation is managed by a *Governing Authority*. This is the body responsible for determining the legal status as well as overall policies concerning the data exchange. The Governing Authority manages the members of the X-Road installation. In case of federation, Governing Authorities will also serve as national contact points, establishing bilateral agreements and everything else needed to ensure interoperability.

The Governing Authority is responsible for setting up and maintaining a *Central Server*. The Central Server contains a member directory and other data. This data is distributed to organizations as global configuration (see Section 2.1 for a more detailed discussion).

All the messages exchanged over X-Road are digitally signed to provide both integrity and non-repudiation properties. Hence, X-Road assumes a Public Key Infrastructure (PKI) to function properly. The main PKI components required by the system are Certification Authority and Timestamping Authority. The Governing Authority specifies a list of the trusted PKI service providers and this list is distributed as part of global configuration.

In X-Road interoperability layer, the *X-Road member organizations* communicate directly with each other for data transfer. The communication is structured as synchronous service calls. The data exchange uses mutually authenticated TLS as the transport protocol. All messages carry proof value that is created by signing and timestamping all the exchanged messages.

X-Road members are very different organizations varying from small companies to large governmental institutions. Accordingly, their IT capabilities also vary. Together with different technologies used by every organization, achieving a standardized set of well-implemented security measures is very difficult. Instead of specifying the security protocols and relying on the members to correctly implement it, X-Road uses standard

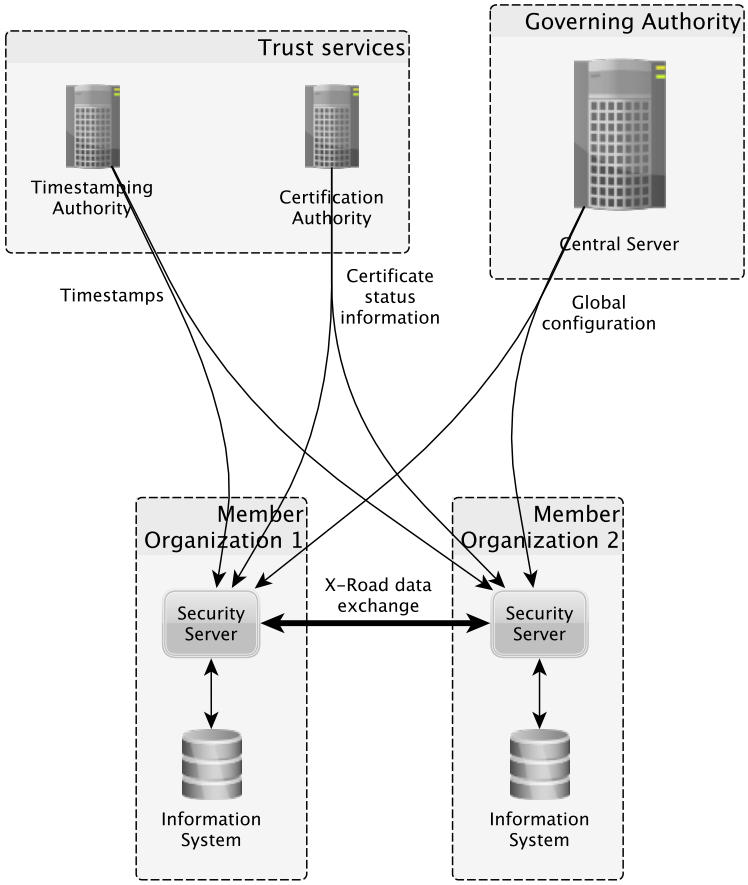

Fig. 1. X-Road infrastructure

components called *security servers*. Security servers encapsulate the X-Road security protocol and ensure that it is implemented properly. They act as gateways between the organization's information system and the X-Road infrastructure (see Figure 1). Security server software is developed centrally and distributed to the member organizations.

Whereas data exchange takes place directly between the member organizations' security servers, the decentralized system is governed by a central Governing Authority. Besides maintaining registry of the members and security servers, the Governing Authority defines, distributes

and enforces policies for the whole system, e.g. security policies. The security policy of an X-Road instance consists of the following items:

- list of trusted certification authorities,
- list of trusted timestamping authorities,
- some tunable security parameters such as
  - maximum allowed lifetime of an OCSP response (how often the certificate validity information must be refreshed),
  - maximum allowed granularity of time-stamping confirmations (how much the time in the cryptographic timestamp can differ from message time).

Security policies are, together with some other management information (like the X-Road member directory), distributed as part of *global configuration*. Since all the trust that X-Road members have towards the whole infrastructure relies on authenticity and integrity of this configuration, its distribution is the most security-critical operation during the initialization process.

2.1 Configuration management

Global configuration is distributed by the central server to X-Road members in a set of signed XML files. However, the corresponding verification key can not come as a part of this configuration, but must enter the system from a different, *a priori* trusted source. In case of X-Road, this trusted source is established by loading the verification keys manually to the security servers in form of *configuration anchor* files.

A configuration anchor file contains a URL that can be used to download the global configuration, and a set of public keys to verify authenticity of the downloaded files. The configuration anchor is distributed via out of band means and loaded into the security server on initialization. The anchor is then used to verify the downloaded global configuration (containing approved CA certificates) that, in turn, is used to verify certificates used by security servers.

In case of federated X-Road installations, configuration is typically split into two parts (see [6]):

- *private parameters*: set of parameters that are used only by members of this X-Road instance (for example, addresses of certain management services), and
- *shared parameters*: set of parameters that are used by members of this X-Road instance and other federated instances (for example, member directory, list of trusted CAs).

For added flexibility, private parameters can include additional configuration anchors. These anchors can either refer to shared parameters of the same instance (typically both configuration files are served from the same URL endpoint) or some other X-Road instance. This mechanism is used to set up federation relationships between X-Road instances.

Figure 2 shows an example configuration with two X-Road instances. Here the security servers are initialized with a freestanding configuration anchor. Using this anchor, they can download and verify a private parameters file. The anchors in the private parameters file can, in turn, be used to download shared parameters files of both local and remote X-Road instances. Note that the configuration anchors cannot be chained – the security servers only trust anchors found from the private parameters file of their own instance.

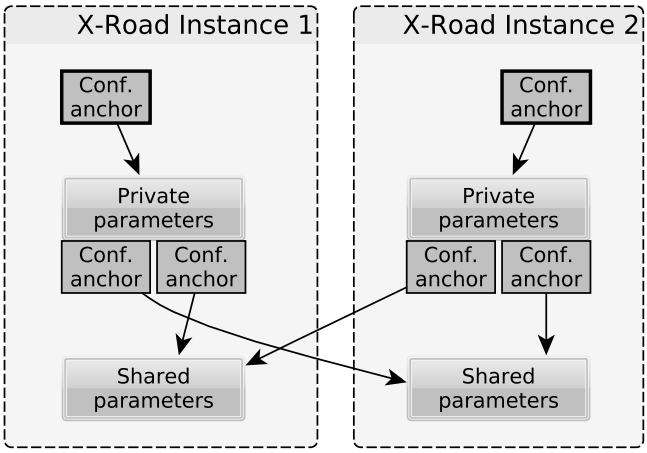

Fig. 2. X-Road configuration management

3 Implementing X-Road federation

The options for federating different X-Road installations were first studied by Ansper and Willemson in 2008 [11]. Three possible strategies were proposed:

1. A new higher level is defined having all the present X-Road infrastructures as its descendants.

2. To facilitate international queries, a new cross-border X-Road instance is established in parallel with the existing ones.
3. All nations have their own X-Road infrastructures, and no additional ones are defined. In order to allow international information exchange, bilateral agreements are made between the existing governing institutions.

This paper explores more closely the third option that is also selected for federating Estonian and Finnish X-Road instances.

X-Road federation infrastructure is depicted in Figure 3.

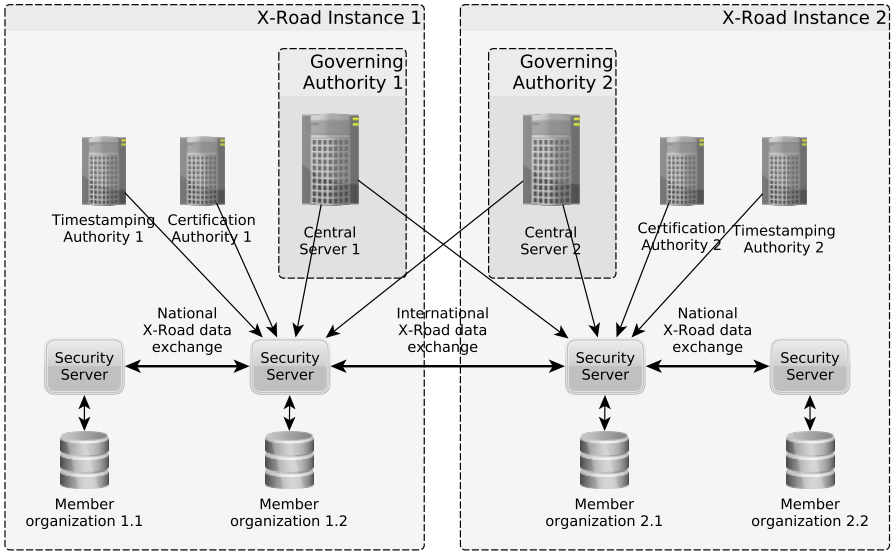

Fig. 3. X-Road federation

Here the two governing authorities enter into a bilateral agreement. Based on this, they exchange the configuration anchors pointing to their respective shared parameters file. Each governing authority copies the anchor of their partner into their private parameters file so that all the security servers can download and verify the configuration of the federated instance. X-Road uses the PKI in such a manner that each security server only interacts with the Certification Authority that issued its own certificates. Every time a security server receives a certificate, it comes with all the necessary information (OCSP responses for all the certificates in the certificate chain) needed to verify it. This method is also used in

the federated setting – security servers do not make any requests to trust services of another federated instance; they only use list of trusted CA certificates from global configuration to verify certificates.

4 Legal challenges

Reaching bilateral agreements between two governing authorities is always non-trivial. Even though huge efforts are put to trust service level unification on the European Union level, there is still a lot of room for discrepancies.

For example, the Regulation No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services (known as eIDAS regulation [1]) defines several trust levels of digital certificates. In the Estonian X-Road, the signatures created by security servers comply with the requirements for a qualified electronic seal. This, among other things, means that qualified signature creation devices are required for all the X-Road members [4].

However, obtaining such devices and certifying the corresponding public keys is both costly and troublesome. For this reason, the Finnish installation of X-Road uses more relaxed requirements on certificates and signature creation. Effectively, every member of the Finnish instance can freely choose one's own signature creation mechanism, and consequently the level of assurance its messages carry. This causes an inherent asymmetry between the levels of trust one can have in the messages coming from Estonian and Finnish X-Road instances. eIDAS does not provide a mechanism for communicating the level of assurance required by the receiver of the signature. The only choice the receiver has is not to accept messages that carry a lower level of assurance, but this comes with a price of rejecting all the communication with the party that uses low-level certificates. This clearly contradicts the whole idea of X-Road federation.

It is also the case that existing international legislation (like eIDAS) covers only a fraction of all the required legal context. For example, dispute resolution must take place under some jurisdiction. In case of Estonian-Finnish federation, Estonian legal system is selected to be the primary dispute resolution context [5]. However, it is an open question what kind of a legal framework is appropriate for larger federations.

Another complicated legal issue is caused by possible incompatibility of certification policies in different jurisdictions. Even though on the technical level certificate interoperability can be rather well achieved by following the same standard (typically X.509v3 [2]), not all the aspects

of trust are established using technical measures. Every Certification Authority also follows a certification policy that specifies a set of requirements and best practices that the CA follows when issuing the certificates.

A certification policy determines how certificate holders' identities are verified, how certificate life cycle is managed, how validity information is distributed, how all the processes are be audited, etc. A typical certification policy comprises of dozens of pages of loosely structured text. Making sure that two certification policies coming from different sources are in some sense compatible is a highly non-trivial task.

5 Technical challenges

There is a number of technical parameters determining the service and trust level of X-Road infrastructure and messages exchanged over it. To ensure meaningful interoperability, these parameters must be comparable between different federated instances. Table 1 summarizes the parameter values in case of Estonian and Finnish deployments [10, 3, 4].

Table 1. Parameter comparison for Estonian and Finnish X-Road instances. The values marked with $*$ are not formally regulated

Parameter	Estonian instance	Finnish instance
Validity of OCSP responses	8 hours	23.5 hours
Validity of global configuration	6 hours	72 hours
Minimal required hash function	SHA-256	SHA-256*
Minimal required asymmetric algorithm	RSA2048	RSA2048*

Note that in the case of Finnish X-Road instance, the minimal required security level of cryptographic algorithms is not defined by a formal regulation, but by simply listing the corresponding TLS cipher suites as allowed as part of the security servers' configuration.

From Table 1 we see that the biggest discrepancy between the Estonian and Finnish instances is in the validity periods of OCSP responses and global configuration. In a regular day-to-day operation these differences should not matter much. However, the core idea of X-Road is to provide reliable data for decision making, with the option of holding the data source responsible if incorrect data can be proven to be the cause of an incorrect decision.

For example, in the case of OCSP responses it is possible that a certificate has been compromised and revoked, but one of the OCSP responses

with a longer validity period can still be accepted by the members of one X-Road instance. Who should be held responsible in case a questionable decision has been made as a result is currently an open legal issue.

There are also other operational differences between the Estonian and Finnish X-Road instances. The original Estonian X-Road was built to support decision-making process based on the data obtained from other parties. In order to be able to later prove rightfulness of the decisions, the data needs to carry long-term evidentiary value. For this reason, all the X-Road messages are signed and timestamped.

Note that signing and time-stamping alone do not guarantee long-term proof value. The messages also need to be stored for later verification. This is why security servers in the Estonian X-Road instance support extensive message logging and archival.

In case of the Finnish instance, however, logging facilities are not utilized. This creates a potential situation where a Finnish X-Road member takes a decision based on the data obtained from, say, some Estonian collaborator, but will later be unable to prove the correctness of its actions. It is impossible to predict the outcome of the following disputes.

Another challenging aspect is authorization within member organizations. X-Road queries must be initiated by, and the results should eventually be used or interpreted by someone. X-Road infrastructure only deals with access control on the organization level. End-user authorization and access control within the service client's information systems are left as responsibilities of the service clients themselves. If these management practices are lax, an X-Road member sharing its data sets is risking a potential privacy leak due to a careless employee of the partner organization.

For an X-Road member, it is very hard to impose formal access control requirements on another organization, or verify that these requirements are fulfilled. This problem is even more serious concerning an organization in another country.

In principle, this problem should be solved by service use agreements – when gaining access to a service, the service client agrees to implement the required controls for authentication, access control and managing the received private data. The service provider opens the service only if it is satisfied with the level of security implemented by the client. However, for service providers with many clients, the case-by-case approach does not scale. It is infeasible to audit all the client information systems for compliance with the requirements. Thus, instead of treating each client separately, the service provider can require that all the clients implement

a common security standard (assuming, of course, that the standard complies with the requirements of the service provider).

Both Estonia and Finland have established frameworks for assessing and ensuring security levels of governmental information systems. In case of Estonia, a baseline security system ISKE, a derivative of German BSI, is established[3]. Its Finnish counterpart is called VAHTI[4]. However, both frameworks are extensive (for example, if printed out, ISKE threat and countermeasure catalogues span over 3000 pages). To the best of our knowledge, there has been no thorough comparative analysis of the two, so from the viewpoint of a member of one X-Road instance, it is very hard to tell what security level can be assumed from a member of another instance.

It should be stressed again that these frameworks are only compulsory for governmental information systems. Private companies can implement these requirements, too, if they choose to, but there is no such obligation. Consequently, it is harder to state anything about the security level of data handling practices for private X-Road members.

6 Conclusions

Federating infrastructures like X-Road is inevitable in order to utilize data across national borders. However, there are still more question than there are answers.

Even though Estonia and Finland are close both geographically and culturally, there exist many differences in legislation, technical solutions and best practices. These differences have the potential to cause non-matching interpretations of various events, which in turn may lead to an unclear state of possible disputes. This contradicts the overall ideology of introducing X-Road in the first place.

This paper pointed out some of the most urgent problems that need to be addressed in both legal and technical aspects. However, even though there exist similar technical and operational standards in the two considered countries, not all the implementation aspects have been nor can be fully aligned.

While X-Road federation is still in the planning state, these issues are easier to fix than on the running system. On the other hand, actual severity of the identified problems is rather hard to assess without ob-

[3] `https://www.ria.ee/en/iske-en.html`
[4] `https://www.vahtiohje.fi/web/guest`

serving them in practice. Thus this line of research needs to be continued throughout the life cycle of federated X-Road implementation.

Acknowledgements

The research leading to these results has received funding from the European Regional Development Fund through Estonian Centre of Excellence in ICT Research (EXCITE) and the Estonian Research Council under Institutional Research Grant IUT27-1.

References

1. Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC, http://eur-lex.europa.eu/eli/reg/2014/910/oj
2. Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile (2008), iETF RFC 5280, https://tools.ietf.org/html/rfc5280
3. Infosüsteemide andmevahetuskihi Eesti keskkonna tingimused (2016), https://www.ria.ee/public/x_tee/Eesti_keskkonna_tehnilised_tingimused.pdf
4. Infosüsteemide andmevahetuskihi usaldusteenuste tingimused (2016), https://www.ria.ee/public/x_tee/usaldusteenuste_tingimused.pdf
5. Trust federation of Estonian X-tee and Finnish Palveluväylä. General Agreement between Estonian Information System Authority and Finnish Population Register Centre (2016), https://esuomi.fi/wp-content/uploads/2016/09/RIA_VRK_Trust_federation_Agreement_.pdf
6. Annuk, S., Nõgisto, I., Freudenthal, M., Mattila, J., Kallio, S.: X-Road: Protocol for Downloading Configuration (2017), https://github.com/ria-ee/X-Road/blob/develop/doc/Protocols/pr-gconf_x-road_protocol_for_downloading_configuration.md
7. Ansper, A., Buldas, A., Freudenthal, M., Willemson, J.: Scalable and efficient PKI for inter-organizational communication. In: 19th Annual Computer Security Applications Conference, 2003. Proceedings. pp. 308–318 (December 2003)
8. Freudenthal, M., Hanson, V., Nõgisto, I., Kromonov, I., Annuk, S., Seppälä, I.: X-Road Architecture. Technical Specification, version 1.3 (2015), https://github.com/vrk-kpa/xroad-public/tree/master/src/doc/Architecture
9. Kalja, A., Vallner, U.: Public e-Service Projects in Estonia. In: Proceedings of the Baltic Conference, BalticDB&IS 2002 – Volume 2. pp. 143–154. Institute of Cybernetics at Tallinn Technical University (2002)
10. Kivimäki, P.: Trust federation of Estonian X-tee and Finnish Palveluväylä. General Agreement between Estonian Information System Authority and Finnish Population Register Centre. Annex X – Environments (2016), version 0.1
11. Willemson, J., Ansper, A.: A Secure and Scalable Infrastructure for Inter-Organizational Data Exchange and eGovernment Applications. In: 2008 Third International Conference on Availability, Reliability and Security. pp. 572–577 (March 2008)