A Secure and Scalable Infrastructure for Inter-Organizational Data Exchange and eGovernment Applications

Jan Willemson
Cybernetica, Aleksandri 8a, Tartu, Estonia;
University of Tartu, Liivi 2, Tartu, Estonia
Email: janwil@cyber.ee

Arne Ansper
Cybernetica, Akadeemia tee 21
Tallinn, Estonia
Email: arne@cyber.ee

Abstract

As more and more information becomes accessible via on-line databases, more public services can be provided and more complex queries involving several registers become feasible as well. However, not all of the digitally stored data is public and thus strict access control mechanisms must be enforced. At the same time, in order to take full advantage of on-line data sources, high availability must be achieved as well. This paper describes an infrastructure developed in Estonia to satisfy these somewhat contradictory requirements. This infrastructure (called X-Road) enables different organizations to access each other's data and rely on it when taking legally binding decisions. We discuss technical details of X-Road together with the issues arising when the infrastructure is to be implemented on national or international level.

1. Introduction

Rapid ICT developments in early and mid 1990-s affected both private and public sectors. Many services were made available to citizens via Internet, many databases were computerized and put to information systems with online connections. Still, many of these systems were developed independently without paying too much attention to interoperability issues. On European Community level, this problem was acknowledged in 1995 by the European Council's decision no. 95/468/EC on Community contribution for telematic interchange of data between administrations in the Community (IDA) [1]. In 1999, this document was further extended by decisions no. 1719/1999/EC [2] and 1720/1999/EC [3], and as a result, many countries across Europe started developing their own interoperability networks in late 1990-s and early 2000-s (UK [4], France [5], Germany [6], Denmark [7] etc.). Besides these general projects, there were also others aiming more specifically at secure document exchange up to the level where the documents received over the exchange infrastructure could be used to take legal decisions [8].

Similar developments took place in other countries as well, including Estonia (not yet being part of the EU back then). By the end of 1990s, public administration in Estonia had reached the level where most of the state registers were computerized. However, interoperability of different registers was still very weak because of several reasons. First, there were many legal entities responsible for different registers and this caused several organizational issues. Besides legal aspects, there were also major technical problems to be solved – most of the registers had been developed independently, using different database backends and having non-standardized interfaces. Thus there was a definite need for an infrastructure that would serve as a bridge between different agencies and registers, enabling the data to be stored and maintained in one place and to request it over the Internet by other agencies only when needed. Of course, many state registers contain private data and access to such information must be restricted accordingly. Thus, high security requirements had to be enforced during the whole project.

An infrastructure addressing the issues above was developed in 2001–2002 by the Estonian government and called X-Road (Crossroad) [13, 9]. Later, several improvements and extra features were added. The current document describes the state X-Road project has reached by the end of 2007, but this is by no means the end of the road. There are several aspects yet to be considered, most notably interoperability of different X-Road infrastructures.

The paper is organized as follows. Section 2 gives an overview of X-Road's requirements and technical solution, Section 3 describes the current state of organizational matters, Section 4 states the routines for nationwide implementation of X-Road in a new jurisdiction and Section 5 discusses the problems arising when several national X-Road infrastructures are to be interconnected. Finally, Section 6 draws some conclusions and sets directions for future work.

2. Description of the X-Road System

2.1. Background and Requirements

In Estonia, there are roughly 500–1000 different state registers (depending a bit on the definition of a register), all very different by their architecture, managed and developed by different organizations and financed separately.

The variety of potential users is even larger – in principle, every person, every company, every government and non-government organization has right to access certain kinds of information stored in state registers. Most of these companies and organizations are rather small, they have only limited budget and limited knowledge of security; all of this being even more true for private individuals.

The security requirements, at the same time, are rather high. Registries often contain personal data that is in some cases used to make high value decisions and in some cases needed in real time.

The initial analysis showed that the priorities of security requirements for the X-Road infrastructure were the following. First, all applications required **authenticity** and **integrity** of the data. Additionally, the data received over X-Road was required to have **evidentiary value**, i.e. the receiver had to be able to prove the origin of the data. Second, it was envisioned that X-Road would be used by time-critical applications, like verifying identities on the border or performing police operations. Thus X-Road had to provide high **availability**. And finally, **confidentiality** was required in most, but not all cases.

While satisfying these security requirements, the main goal of X-Road still had to be achieved – exchanging information between different agencies had to be as smooth as possible. Even more, X-Road system architecture had to take into account the fact that information systems of the organizations participating in X-Road were very different. It was unrealistic to assume that many of the systems would have been substantially reorganized.

2.2. Security Mechanisms

Evidentiary Value and Integrity To ensure evidentiary value and integrity of the data, all messages going out from the registers are signed. In order to be able to sign anything, certified signature keys are needed and those are provided by a special third party – X-Road Central Agency – that acts as a certification authority.

In order to preserve evidentiary value of the messages over longer time periods, just plain signing is not enough, because signature keys can become revoked or compromised. To overcome such problems, X-Road infrastructure includes logging and time-stamping facilities. All messages received over X-Road are logged and the logs are linked together using a cryptographic hash function. The intermediate hash values are periodically time-stamped by the X-Road Central Agency. This allows detecting the message log tampering attempts. Later in case of disputes, message receiver can use the logs and central time stamps to prove the origin and time of received messages.

Availability In order to achieve high availability of the infrastructure, X-Road is built as a distributed system with minimal number of central services. Besides that, measures must be taken against temporary unavailability of some services and possible Denial-of-Service (DoS) attacks.

There are three kinds of security-related central services needed by X-Road infrastructure and these are provided by the X-Road Central Agency:

- **Certification** is an off-line process and hence not really vulnerable to availability threats.
- **Time-stamping** is used only for log management in X-Road and is thus not time-critical either.
- **Directory service** is used to distribute addresses and certificate validity information, both being security and/or time critical in many scenarios. In X-Road, directory service is built on top of Secure DNS (DNSSEC) [10, 12, 11]. This well-proven approach provides very robust, scalable directory service with built-in caching and redundancy. Security extensions of DNS (signed zones) ensure that the data cannot be tampered with.

There are actually more organizational services provided by the Central Agency that will be discussed in Section 2.4.

X-Road protocol supports redundant servers and load sharing. When there is more than one server that offers some service, they will be used in random order by the clients.

X-Road servers also have some protection mechanisms against DoS attacks. Critical resources (i.e. CPU time and file handles) are shared between different clients in a fair manner, e.g. clients submitting many requests are not given high preference.

Confidentiality Most of the data that is exchanged via X-Road is not public or has some special access rules that must be followed. All major confidentiality threats can roughly be classified as external (e.g. network sniffing) or internal (e.g. gaining unauthorized access) attacks.

To combat external attackers, all data exchanged over X-Road is encrypted and standard SSL protocol is used for encryption.

In order to prevent internal attacks, two-level access control mechanism is used in X-Road, the two levels being inter- and intra-organizational.

X-Road system core only deals with inter-organizational access control, where one organization grants access rights for some service to another organization as a whole. It is the responsibility of the other organization to ensure that only the right people can use this service, by using whatever technical and organizational means it finds appropriate.

By isolating the details of authentication and access control mechanisms used internally by the organizations greatly reduced impact to the existing systems. This was one of the key success factors of X-Road.

2.3. Technical Solution

If an agency wants to connect to the X-Road infrastructure (either as a service provider or a service consumer), it mostly needs to install a dedicated X-Road security server (called just X-Road server later on) in its premises. The only exception can be made for service consumers with limited requirements and capabilities, see Section 4.3 for a more detailed description. The X-Road server acts as a firewall between the agency's information system and Internet. It routes messages to their recipients and secures messages both for transport (by encryption) and for long-term validation (by signing and logging).

On the service provider side, an additional adapter server is required. Adapter server transforms the requests obtained through the security server into the format understandable to the register and translates the responses the other way round. In terms of client information system, the security infrastructure is completely transparent – the application level protocol used between the client information system and its security server matches exactly the protocol used between an adapter server and its security server.

Agencies exchange information using XML-based query-response protocols. In the first version of X-Road, XML-RPC was used. SOAP support with two-way transliteration to XML-RPC was added in the second version. The third version added support for SOAP attachments and asynchronous operation, where X-Road servers queue the messages targeted to other organizations. Starting from January 1st 2008, XML-RPC support is considered obsolete and it is dropped in favor of SOAP completely.

In addition to the real services provided by participating organizations, X-Road also provides some meta-services that can be used to find out the structure and properties of the system. E.g., it is possible to obtain the list of organizations providing and consuming different services, the list of provided services and their formal descriptions in WSDL or proprietary XML format.

For a client, X-Road server functions as a proxy to all services provided by other organizations, so all services are accessible through the same URL. The parameters of the requested service are specified in the SOAP header. They are also protected by digital signature and time-stamping mechanisms.

2.4. Central Agency

As noted above, X-Road has a Central Agency that manages centralized tasks and is responsible for the X-Road infrastructure to stay operational. Since data received over X-Road can be used to take legal actions, one of the most important responsibilities of the Central Agency is to ensure legal status of the information exchanged via X-Road by enforcing the stated policies. Organizationally, Central Agency is also responsible for steering the further development of X-Road and maintaining its consistency and integrity.

Besides security-related features described in Section 2.2, Central Agency also provides some technical services, like the **monitoring service** that monitors all the servers in the system and the **web-based portal** for accessing the X-Road services in a simple and centralized way. The portal is intended to be used for citizens and smaller organizations without sufficient IT capability.

2.5. X-Road Platform

X-Road servers are built on GNU/Debian Linux, one of the most stable, free and best maintained Linux distributions. All the additional software is packaged as Debian packages that can be installed and maintained using standard Debian tools. To simplify X-Road server installation by participating organizations, a self-contained installation CD is provided that installs and configures a GNU/Debian Linux system and X-Road software with minimal user intervention.

X-Road servers have a simple and clean graphical user interface for maintenance and configuration tasks. There is also a built-in patching system that allows to distribute and update X-Road software in a secure way. All keys, including the top-level certification keys, can be changed on the fly without interruptions to the system operations and with minimal user intervention.

3. Current State

X-Road is currently used by the Government of Estonia and private companies. X-Road is the preferred way for connecting governmental agencies and is also used by private companies to exchange data with the government and with other organizations.

By the end of 2007, there were 69 service providers and 408 institutional service consumers connected to X-Road. Almost 216,000 private individuals had used the system by this time. The record of 217,282 X-Road queries made in a single day was achieved in December 3rd 2007.

Table 1 summarizes the overall numbers of queries made throughout the years 2003–2007. Considering that the population of Estonia is only 1,35 million, there were approximately 30 queries per peron made in 2007.

Table 1. Queries made in 2003–2007

Year	No. of queries
2003	590,337
2004	7,802,582
2005	14,208,028
2006	29,728,546
2007	≈40,000,000

In order to get a better idea what kinds of services are provided over X-Road in Estonia, consider the following examples.

- **Parent benefit in Internet** In order to process parent benefit applications, information from five different registers and information systems is needed and seven different forms must be processed by the civil servants. By using X-Road, most of this processing and information retrieval can be done automatically and no paper forms are used.
- **Document repository** Document management systems of the public sector have an interface with the central document exchange point. They periodically send and receive documents for and from other systems. There is no more need for traditional post or scanned documents, all documents with metadata in XML format are exchanged over X-Road.

4. Nationwide X-Road Implementation

If X-Road infrastructure is to be implemented on a national (or even international) level, several legal and technical problems must be solved first. This section gives a short overview of the respective prerequisites.

4.1. X-Road Central Agency

If X-Road implementation is planned in a new jurisdiction, first and foremost, an organization is needed that acts as X-Road Central Agency. Due to its tasks and great responsibilities, this organization must have both legal and IT capability. After a suitable body has been chosen or established, the following steps are necessary.

1. A set of policies and rules for operating the X-Road infrastructure is needed. It would be a good idea to align the security measures within the existing security frameworks (e.g. BSI in Germany[1]) and define the required security levels via existing terms. X-Road Central Agency is responsible for ensuring that only the organizations meeting the stated security requirements can connect to X-Road.
2. Next, X-Road Central Agency must develop its own security policy and operation procedures to ensure proper handling of security critical data and achieving the required legal status of the information exchanged over X-Road.
3. Finally, X-Road central services must be installed and configured. Due to their importance, the servers must be located in secured premises. Hence, some extra effort might be needed for securing the server room. It would be a good idea to have several, physically separated locations for central services.

4.2. Service Provider

If some organization wants to start acting as a service provider in X-Road infrastructure, it has to take the following steps.

1. Service provider must ensure that it has sufficient security measures in place in order to join X-Road. Security policies and operational procedures may be revised and updated according to the legal requirements and/or policies enforced by the X-Road Central Agency.
2. X-Road server(s) must be installed and configured in the premises of the organization.
3. Existing or new SOAP services must be developed according to X-Road specifications.
4. After making contracts with service consumers, access rights to use the service are granted to the client organizations.

4.3. Service Consumer

If some organization wants to start using services provided by other parties of X-Road, it must first ensure that it has sufficient security measures in place in order to join X-Road. Security policies and operational procedures may be revised and updated according to the legal requirements

[1] `http://www.bsi.bund.de/`

and/or policies enforced by the X-Road Central Agency. One of the most important steps is to ensure that sufficient user authentication and access control mechanisms are in place.

Service Consumer has two options for using the X-Road services – either by integrating them into its information system or using the X-Road portal. We will consider both options in a bit more detail.

Integration of the Services This is the preferred way to use X-Road services and is accomplished by taking the following steps. First, a separated X-Road server is installed in the organization's premises and configured. Second, the information system of the organization is modified to use X-Road services proxied by the X-Road server. Finally, after making the contract with service provider, the user access rights are set up according to the contract.

Portal X-Road Central Agency may choose to run X-Road portal. In this case the client organization does not need to have X-Road security server or information system. This is very convenient for private individuals or smaller organizations without remarkable IT capability.

The client organization must still satisfy some security requirements. First, there must be organizational and physical security measures in place ensuring that the information obtained via X-Road is handled in a safe manner. Second, the organization must appoint the User Manager who is responsible for setting up the user access rights in the X-Road portal. Finally, after making the contract with service provider, the user access rights are set up by the User Manager according to the contract.

In Estonia, there is a special Citizen Portal run by X-Road Central Agency. Citizen Portal acts as an agency information system that can be used by all citizens to access information about themselves. All queries that are accessible to Citizen Portal take citizen's personal identification code as an argument (usually it is the only argument) and return information stored in the registry concerning this particular citizen. Although almost all Citizen Portal queries are informational, some queries are developed which allow citizens to send documents to state agencies.

5. International X-Road Implementation

Future developments of the infrastructure and policy standards are needed when X-Road is to be implemented internationally. Both legal and technical systems need amendments.

If there are two (or more) countries/jurisdictions whose organizations wish to start sharing information over X-Road, there are in principle three possible approaches to extend the infrastructure.

1. A new higher level is defined having all the present X-Road infrastructures as its descendants. This higher level would get a new root key that would be used to sign all the descendants' root keys, thus creating a trust path to verify for the agencies participating in separate infrastructures (so-called *bridge-certification* [14]).

2. There exist both national and international X-Road infrastructures that are basically independent. An agency is allowed to belong to several infrastructures at the same time. This approach would allow not to change any existing infrastructures, and joining the international X-Road could even be done invisibly to the agencies (by issuing new certificates for the existing public keys).

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 Central Agencies and the respective governmental bodies acknowledging legitimacy of the data received from the other infrastructure. Such bilateral agreements can later develop in a natural way into multilateral ones. On the technical level, the root key of one X-Road infrastructure is used to certify the root key of another infrastructure and vice versa (so-called *cross-certification*).

The first and the second approach both suffer from similar shortcomings. Establishing a new international infrastructure with a new Central Agency would need major political agreements between different countries, e.g. the question in whose premises this new Central Agency would reside still needs an answer. It is not clear how countries with contradicting interests would reach such agreements.

In a way, both of these solutions also act against the spirit of X-Road. The major motivation for different organizations to join the infrastructure is to simplify and unify its communications with other parties. When a new higher or parallel infrastructure is introduced, the participating organizations potentially need to undertake another process of joining, together with the implied bureaucracy and general management issues. This may decrease the organizations' motivation to join the international X-Road considerably.

It is currently the belief of the authors of this paper that the third approach (cross-certification) hits the best balance between the obtained benefit and new problems caused. This approach avoids the politically sensitive problem of international Central Agency and lets each country to stay on top of its own infrastructure instead. The main technical problem to be solved is integrating different DNSSEC directory services. On one hand, directory service is the most security critical component of X-Road. On the other hand, information contained in the directory of one

participating infrastructure (keys, addresses) must be available for other infrastructures, too. Other services (log time-stamping, monitoring, web portal) do not require such level of integration – e.g. each Central Agency will still be providing time-stamping service for its own infrastructure.

From the legal point of view, some changes may be needed to the legislation of participating countries to give the Central Agencies enough power to make the required agreements. Among other things, it should be decided how potential disputes are solved.

However, software used in X-Road servers does not currently support any of the approaches described above, since the services are meant to run in a single flat infrastructure only. This implies the need for extra development efforts, but those can be made only after the corresponding political decisions have been taken.

6. Conclusions and Further Work

In this paper we presented an infrastructure developed in Estonia for secure inter-organizational data exchange. It is fair to say that X-Road project has been a real success and it is used heavily to smoothen everyday bureaucracy both in Estonian public and private sectors. Every day more and more services are offered and consumed over this infrastructure.

Even though running very well in one jurisdiction, X-Road still has to stand the challenge of becoming international. There are both technical and political obstacles to tackle, the political ones being probably much harder to overcome. In the paper we argued that cross-certification is the easiest solution for international X-Road implementation, but even this approach assumes that the underlying national implementations are rather similar. However, even in the days of European Union, there is still a lot of distrust between different countries and many of them try to solve similar problems in similar, but non-compatible way.

We conclude that the future of the project greatly lies in the hands of politicians and not that much of the technical community. Still, there are several technical issues to be solved for X-Road. Right now X-Road infrastructure supports only one Central Agency and assumes that all the agencies are its descendants in the certification hierarchy. When cross-certification with other hierarchies will be used, this will no longer be the case and hence, X-Road servers must be updated accordingly. Another major problem is integration of different DNS-SEC directory services, since most of the security features of X-Road rely on them. This problem will currently remain the subject for future development as well.

7. Acknowledgments

This research is supported by Estonian Science Foundation grant #7081. The authors would also like to thank Ahto Kalja and Martin Undusk for valuable input to the paper.

References

[1] Council Decision 95/468/EC of 6 November 1995 on a Community contribution for telematic interchange of data between administrations in the Community (IDA). Official Journal of European Communities, 11 November 1995.

[2] Decision No 1719/1999/EC of the European Parliament and of the Council of 12 July 1999 on a series of guidelines, including the identification of projects of common interest, for trans-European networks for the electronic interchange of data between administrations (IDA). Official Journal of European Communities, July 1999.

[3] Decision No 1720/1999/EC of the European Parliament and of the Council of 12 July 1999 adopting a series of actions and measures in order to ensure interoperability of and access to trans-European networks for the electronic interchange of data between administrations (IDA). Official Journal of European Communities, July 1999.

[4] e-Government Interoperability Framework v 3.0, October 2001.

[5] Common interoperability framework for public information systems, version 1, January 2002.

[6] SAGA Standards and Architectures for e-Government Applications, version 1.1, 2003.

[7] Architecture for e-Government in Denmark – Challenges and Initiatives, June 2004.

[8] Secure exchange infrastructure for e-government presented in France. eGovernment News, http://ec.europa.eu/idabc/en/document/3357/357, October 2004.

[9] A. Ansper, A. Buldas, M. Freudenthal, and J. Willemson. Scalable and Efficient PKI for Inter-Organizational Communication. In *Proceedings of the 19th Annual Computer Security Applications Conference ACSAC 2003*, pages 308–318, 2003.

[10] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirements. RFC 4033, March 2005. http://www.ietf.org/rfc/rfc4033.txt.

[11] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Protocol modifications for the dns security extensions. RFC 4035, March 2005. http://www.ietf.org/rfc/rfc4035.txt.

[12] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Resource records for the dns security extensions. RFC 4034, March 2005. http://www.ietf.org/rfc/rfc4034.txt.

[13] A. Kalja and U. Vallner. Public e-Service Projects in Estonia. In Hele-Mai Haav and Ahto Kalja, editors, *Databases and Information Sustems, Proceedings of the Fift OInternational Baltic Conference, Baltic DB&IS 2002*, volume 2, pages 143–153, June 2002.

[14] W. T. Polk and N. E. Hastings. Bridge Certification Authorities: Connecting B2B Public Key Infrastructures. NIST, 2000 September.