Designing a Governmental Backbone

Arne Ansper

Cybernetica AS, Tallinn, Estonia

Abstract. The presentation is about the design of the backbone of the Estonian governmental information systems - X-Road. The system is already ten years old and has proven to be useful and reliable. The presentation describes the vision of the system, the requirements analysis process and the technical design decisions. The vision was to create a web-services based unified access to all governmental registries. The requirements analysis was guided by the existing legislation and organizational setup of the government. The technical design was pragmatic and based on some unorthodox solutions.

Keywords: case-study, eGovernment, web-services, digital signature, interoperability framework, distributed middleware, DNSSEC.

1 Introduction

X-Road is the backbone of the Estonian governmental information systems. It is a system for securing the inter-organizational data exchange. It is not a physical communication network, but more like a VPN with digital signatures and access control for web-services. We could call it "distributed middleware" or "distributed ESB (Enterprise Service Bus)".

2 History of the X-Road

It is already an old system - it has been in use for almost ten years now. There are more than 600 organizations connected to X-Road, offering more than 1500 services that were used more than 225 million times during 2010. The idea of the X-Road was born in the Department of State Information Systems, sometime at the end of last century. They wanted to offer unified data exchange environment that would reduce money and time spent on integration projects. The prototype was built and tested. The main idea was to use web-services, as a platform neutral protocol for inter-organizational communication and build a central hub that would mediate all service requests and perform access controls.

This was quite an innovative approach, because the web-services were just a couple of years old. The prototype was sound from the functional perspective - it provided unified access to different registries. But it did not address security concerns and did not take into account the legal and organizational framework - basically how the state is functioning.

P. Laud (Ed.): NordSec 2011, LNCS 7161, pp. 1–3, 2012.

3 Requirements for the System

The idea was to come up with a solution that would allow effortless access to the data in state registries, without compromising the security of the data, with minimal impact to the existing systems and without requiring major legal changes. The functions of the government are divided between agencies. They have a freedom to decide how they implement those functions - what are the procedures, systems, etc. On one hand this division creates redundancy and inefficiency. On the other hand, this creates stability and helps to ensure that the principles of the democratic government are followed. When we apply IT, we optimize the system. There is a danger that we over-optimize and create a system that puts too much power into someone's hands.

Each agency is authority in its field and responsible to ensure the rightful usage of its data. Centralization of the data or access to the data would violate this principle. Deployment of such a system would see a great resistance from the agencies. Creation of such a super-database would also create an inviting target for attacks and is bad from the security viewpoint.

The solution must be decentralized and based on collaboration. Central hub is not a solution. The data that is managed by one agency is needed to make some decision in another agency. The agency that makes a decision needs some evidence to prove later why such decision is made. We need a system that would preserve the authenticity, integrity and evidentiary value of the data. If we have such a system in place, it will be used more and more. Many business processes that used to be independent will start using external services. The system must ensure the high availability of the services.

Finally, some data is confidential. It must be protected against external and internal attackers. It is important to notice that the security requirements are prioritized. When people normally talk about "security" they think about confidentiality. In reality, the other properties are more important. In fact, this is pretty standard set of security requirements. Indeed, we can use standard security measures to satisfy them. The important question is how to make all this technology and procedures available to all organizations that need to exchange data. Most of the organizations are very small and without IT and security management capabilities. The solution must be very easy to deploy, maintain and use.

4 Architecture of the System

We designed a distributed system. Each organization runs a security server. Security server is a mono-functional self-contained GNU/Debian based server that implements all security related aspects of the inter-organizational communication. It is an appliance. Security servers are communicating with each other directly. The data flow between organizations is direct.

Security server is basically an application level firewall for SOAP + digital signature creation and verification device for SOAP messages + highly available VPN device.

Local applications see the security server as a provider of all web services offered by other organizations. Remote service requests by local application will be proxied by security server. Security server will sign all the outgoing SOAP messages (requests and responses). Security server will verify the signatures of all incoming SOAP messages, will time-stamp them and archive them. Security servers contain full history of communication. Digital signature mechanism together with archiving ensures the authenticity, integrity and evidentiary value of the exchanged data.

Security servers also control access to the web-services at the organizational level. The organization receiving the service must ensure that only right people can use this service, by using whatever technical means it sees appropriate. This obligation is enforced by service provisioning contract between the organizations. Two level access control isolates the details of organizational authentication and access control mechanisms and minimizes the impact to the existing systems. The balanced use of technical and organizational security measures was an important success factor of the X-Road.

The availability of the system is increased by having a minimal number of central services: only time-stamping and secure directory. Time-stamping is used in a way that makes it non-time critical. The time-stamping is performed asynchronously in the background. Directory service is based on the Secure DNS (DNS-SEC). All the information that needs to be shared is put into DNS zone: IP-addresses, valid certificates, names of the organizations and groups, etc. Well-proven DNS protocol and implementation provide robust, scalable directory service with built-in caching and redundancy. Security extensions ensure that the data cannot be tampered. All X-Road servers run a local caching DNS server that performs strict validation of zone signatures and ensures the availability of directory information during network outages.

There are also mechanisms against DoS attacks. Critical resources (i.e. CPU time, file handles) are shared between different clients in a fair manner. Security servers provide also meta-services that help to discover the structure of the system: what organizations are connected, what are the services, download the WSDL descriptions of the services, etc. Meta-services are used by portals to automate the generation of the user-interfaces for invocation of services.

5 Organization and Procedures

X-Road has central agency that ensures its operation. X-Road central agency was created in parallel with the system development. The operating procedures of the agency, security regulations and rules for organizations that connect with the system were all aligned with the technical solution.

Central agency ensures the legal status of the X-Road and the information exchanged via it, by enforcing the stated policies. It is responsible for steering the further development of the X-Road and ensuring its consistency and integrity. It provides central services like certification, time-stamping, secure directory and monitoring and resolves the potential disputes between communicating parties.