

Digital Certificate Management for Document Workflows in e-Government Services

Abstract. This paper presents a proper solution for a medium enterprise or public institution that enables easier management of the digital documents library and eases the common document workflows. The main problem addressed by the proposed project is the complexity of document workflows in public administration. Documents that need to be filled out and signed are always around us and often can cause problems and delays when poorly managed. With its characteristics, our solution eliminates all the inconvenience of the document workflows helped by the document library and workflows, while keeping the security part, now represented by hand signatures with the implementation of the digital signatures. The main benefit it brings to the client is that it automates the signing and approval process to any kind of document it uses inside or outside the company. The signature system allows signing on multiple levels (counter-signatures) and multiple signatures per level (co-signatures) for perfectly mimicking a plain document.

Keywords: e-Government, Electronic Services, Digital Certificate, Document Workflows, Public Key Infrastructure.

1 Introduction

In our days information security is a very delicate matter because more and more information is sent electronically and some of it requires special handling and quality standards. Public Key Infrastructures (PKI) offer the services required to meet those security and quality standards required [1]. PKI based applications bring those services to the end users enabling them to use digital documents as securely as old, plain authenticated documents.

The strength and security offered by system using PKI is proven by the multiple examples of existing solutions that are built the same way and are around us for some time. When reading an email you can be very sure that you are reading exactly what the other person wrote before hitting the send button, or you can be sure that no other person but you has read the information, features enabled by the PKI products, the user digital certificates. When accessing an online-banking account to operate a payment or money transfer or simply when buying anything on the internet, you are sure that your most sensitive credit card information is secure, thanks to the PKI enabled SSL certificates. When enabling a VPN connection to a remote place, all the network traffic flow is encrypted using digital certificates [2]. All the above examples stand as proof that the PKI is a viable option for our design [3].

For a public institution, implementing the proposed solution would certainly bring a boost in quality services. Currently, most trips to a public institution are viewed as a nightmare mostly because of the confusion around all the paperwork that needs to be filled out, the tight schedules of the institutions employees or even the waiting time at the counter especially around certain deadlines (income declaration deadline). Enabling citizens to digitally fill out all paperwork with good guidance and examples all available at a public digital library would help everyone a lot [4]. From the institution's point of view it would mean less paper document handling, susceptible loosing or damaging, less people at the counter bringing important cost savings.

For an enterprise, implementing the solution enables better management for internal documents. Again, dealing with digital versions of a document is much desirable than handling paper documents that are very prone to being lost, damaged and even forged [5]. Digital signatures evolved to a point where it is easier to forge a hand signature than forging a digital signature produced even by a medium security PKI environment in the same time eliminating the human factor in deciding if the signature is indeed authentic. Last but not least there is a small ecological point of view in the entire project by reducing paper usage in the company.

The rest of this paper is structured as follows. We first present the related work and critical analysis of similar solutions. In the next section we present implementation details of the pilot solution, presenting the technical issues and solutions. Next we present several obtained results and, in the final section, we give the open issues, improvements, future development and conclusions.

2 Related work

In this section we present results for a solution similar to what we are trying to develop. We also present some options for a PKI deployment and compare them with Microsoft Certificate Services from Windows 2003.

Microsoft Office SharePoint Services (MOSS). This is the big brother of the actual chosen solution [6]. Unlike it, MOSS is not free and the extra features it brings do not compensate for the extra cost and complexity increase of both deploy and manage. One of those features that we could have been interested in would be the integrated document approval workflow and form providing out-of-the-box solution for the main type of workflow needed in our application. However, being only able to work with .doc files represents a very big draw-back and ultimately led to not choosing this solution. A critical analysis of MOSS is presented in Table 1.

OpenSSL. It is an open source project, focused on developing a free, open, toolkit for SSL and TSL protocols and cryptography. At the core of the toolkit stands the “openssl” command line application [7]. Even if it’s not featured as a CA, “openssl” can provide the services of a Certification Authority with the help of the integrated cryptographic library: Creation and management of public/private keys, creation of X.509 certificates, CRL files and PKI cryptographic operations (see Table 2) [8].

Table 1. Critical analysis of MOSS

Plus	Minus
Easy and quick to implement (out-of-the-box solution)	Only works for “Doc” documents.
Good workflow management, alerting and task management	Not free, unlike the simpler version SharePoint Services
	A lot more complex and more complicated to use.

Table 2. Critical analysis of OpenSSL

Plus	Minus
Free to use.	Designed and optimized for SSL and TSL protocols
Open Source	No user interface, hard to manage.
Still under development, offers not regular but often updates.	No OCSP supported
	Not really designed to be used as a CA in a PKI environment.
	Cannot deploy multiple tier architecture.

CertSign certificates and signing application: CertSign is a company that offers digital certificates for any person that wants a digital identity. Included in their offerings package there is an application that can be used for digital documents management. Implementation of this pseudo-solution in an enterprise would mean

buying certificates from CertSign for every employee and use the application to perform operations on digital documents. In Table 3 there is a list with the summarization of this solution, pointing out the strengths and weaknesses.

Table 3. Critical analysis of CertSign

Plus	Minus
Good, secure PKI, trusted.	No workflow management
Good application for document signing, producing signatures in well known formats (PKCS#7) that can be processed with any other application	Expensive to implement and maintain (certificates need to be renewed annually at a certain price)
Offers certificates on USB tokens – provides good security	No control over issued certificates

CryptoBOT e-Workflow: This is a commercial application part of a bigger solution, also including a signing application (e-Crypt) and a Certification Authority server. It also works with other externally issued certificates. The e-Workflow solution is composed from two main components: The server component is simply an application running on a designated computer that listens for workflow requests and forward alerts to users that need to take action in a workflow. Also interacts with the database for logging [9]. The client part is a more complex application where the user logs in and can see a list of documents and running workflows. The application seems complicated and difficult to use and manage and does not have an attractive look and feel. The critical analysis is presented in Table 4.

Table 4. Critical analysis of CryptoBOT e-Workflow

Plus	Minus
Can sign and integrate any type of document in a workflow	Server component is a simple desktop application.
Can define levels of urgency to documents in a workflow	Client application is difficult to use, does not have an attractive look and feel
Ability to insert comments along with the signature using the e-Crypt application	The alerting system is not based on email but on alerts being sent from the server to the client application, requiring the client app to be started at any time.
	Only works with e-Crypt application forcing you to buy the whole package.
	Saves the signature files in a proprietary format thus not allowing the signatures to be verified and validated outside the environment.
	As the web-site shows, it appears the solution has not been updated for a long time.

EJBCA. Another open source PKI Certificate Authority built on J2EE technology. It is advertised as a strong, flexible, high performance CA with a lot of features, ready

to be implemented. One of the greatest advantages of EJBCA is that it is platform independent, being built on J2EE [10]. Also it supports most cryptographic algorithms and formats for certificates, revocation lists and it supports OCSP responders [11]. Actually, this is a very good alternative to Microsoft's Certificate Authority and could be a valid option in our infrastructure if the client requires it. The main problem with it, and the reason that we could chose Microsoft over it, is that being an open source project, there is no real official support or security patches that can be produced in little short time in case of an emergency.

Table 5. Critical analysis of EJBCA

Plus	Minus
Free to use	Management is done with a web application that depends on a web server.
Open source	No official tech support
Implemented as PKI in some word-class enterprises.	

3 Public Key Infrastructure Architecture

In this section is presented the proposed PKI architecture including both the hardware structure, with network design and components, and the software that all relies on. The target market for this project is made up by medium-large companies or public institutions that need or want to upgrade their digital document handling capabilities by proving means to safely and securely process documents without printing them. No doubt that every deployment site has a computer network in place and most probably has a way of managing users and services in that network and it is likely that it uses an Active Directory to achieve that.

As stated above, the solution focuses on providing digital signatures and/or encryption to internal or external documents, thus there is only one scope for all certificates issued and only a few types of certificates to be issued. Considering this information the best architectural decision would be in favor of a two-tier Certification Authority (CA) architecture consisting of one Root (Self-Signed) CA and one Issuing (Subordinate) CA. The layer missing is the policy layer, which is encapsulated in the Root CA. Three-tier architecture would have had sense if different issuing policies were needed: Issuing certificates for employees for internal use such as domain authentication requires a different policy than issuing certificates for clients and internal staff for document signatures and communication, resulting in a 3-tier architecture with one policy CA for each scope. If dealing with a larger scale institution multiple Issuing CA's can be added to the PKI for load balancing or availability issues.

Certificate Services relies on Active Directory Domain Services (AD DS) [12] providing multiple services such as: storage of configuration information, certificate publishing, and policy and authentication base. Even though it is possible to deploy the PKI in a non AD environment (by using only stand-alone CAs) it s highly

discouraged, as it lowers the security standards. Certificate Services will work on most AD environments.

If the client's network does not contain an AD then we will deploy a single-forest, single-domain AD based using 2 Domain Controllers (Main and Backup). However if the company already has an AD environment, there are several cases we need to take into consideration. The simplest situation is a one-forest, one-domain AD, in which we simply deploy our CAs in the existing domain. Next, for a single-forest, multi-domain case, we need to decide, after consulting with the IT manager, where to deploy the CA. The most difficult scenario is the multi-forest, multi-domain AD, where we need to deploy one CA for each forest in the domain because the Enterprise CA cannot issue certificates to users outside the domain.

A very important point in the infrastructure's design is the physical foundation it has, the actual network that links all the systems and components together. The proposed network design consists of 3 separate sub-networks that define the modules described above. Security is the first and most important aspect when implementing the PKI (see Figure 1).

The attention has to be set on the CA network because it holds the most important piece of information: the Root Authority Private Key. The Root CA system is connected only to the Issuing CA (or CAs if more issuing CAs are used). The link is temporary and is used only when the Issuing CA has to renew the authority certificate or when the Root has to issue the Certificate Revocation Lists. Attached to the Root CA is the Hardware Security Module (HSM) [13], a piece of hardware specifically designed to safely store the Authority private key. The same module can be used for private key archiving, mandatory when the PKI is issuing encryption certificates. Other security measures regarding the Root CA range from disconnecting the server from the network to locking the physical machine locked in a safe, depending on the security needs of the particular implementation.

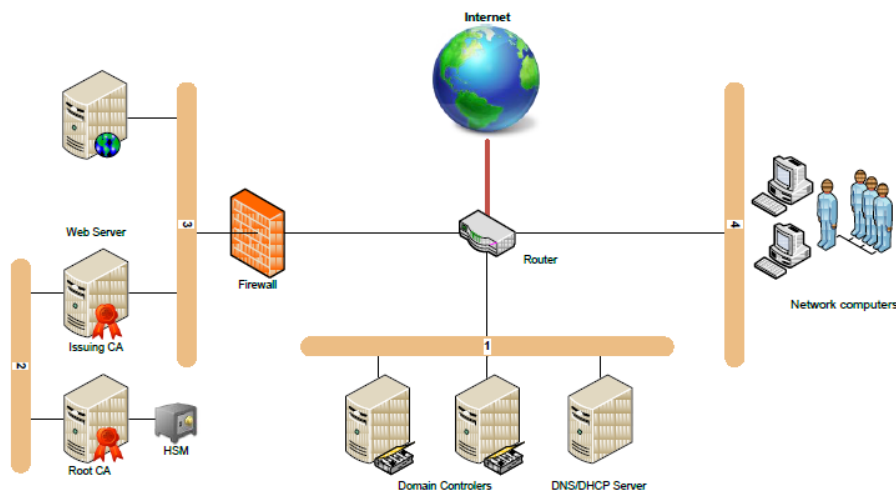


Figure 1. Solution network design

Along with the CA computers, the sub-network could contain an IIS driven web-server responsible with certificate and Certificate Revocation List (CRL) distribution and certificate enrollment. Most interesting task is the certificate web enrollment scenario. This allows certificate issuing using Microsoft provided API using a customizable web page. This feature is active and usable by default when deploying Issuing CA server. It is not uncommon that the IIS server to be installed on the same machine as the Issuing CA, but for some cases, the separate machine is recommended.

The entire sub-network is protected by a hardware firewall, blocking all incoming connections excepting certificate or certificate validation requests. Requests for certificate enrollment are permitted only from private hosts and only for the specific computer running web certificate enrollment. The firewall becomes a single point of failure in this scenario but problems resulted by a failure can easily be overcome. If the firewall fails, two important services cannot be fulfilled by the CA network: certificate enrollment, certificate validation and distribution. First problem is easily solved by manually issuing a certificate based on a certificate request file and using the integrated Certificate Services administration console available on the Issuing CA server. Access to certificate distribution and revocation lists can be quickly regained if 2 separate machines are used for the Issuing CA and IIS server, and it's done by simply connecting the IIS server directly in the domain network. However, make sure to disable web enrollment (do not allow issuing certificates using web pages) while outside the protected network (see Figure 2).

If extreme security is needed, and if certificate issuing is not required to be real-time a different issuing method could be employed. This implies using a second web server, placed in the AD or services network, with the sole responsibility to receive the user requests for certificates, named Registration Authority (RA).

The users apply for a new certificate using the pages offered by the server. Upon completion, the server will create a certificate request and stores it. At specified time intervals, the Issuing CA polls the RA for the certificate requests, processes the requests internally and issues or denies the certificates. They are sent back to the RA and the users can pick up their certificates. The main advantage of this solution is that it offers further protection to the CA network by configuring the firewall to deny all packets originating from outside the CA network unless they are a response to a CA started connection.




Figure 2. External enrollment method

Main elements in the network are the Domain Controllers that make up the Active Directory environment. Using a single domain controller would create a very disruptive single point of failure, in a place where down-time is not acceptable as all

security policies, user accounts, network services are provided by the domain controllers. The easiest and cheapest way to reduce the risk of failure in this area is creating and maintaining a second, backup domain controller. Other services such as DNS, which are required by AD, will be installed in the network, if not already present in the institution.

The proposed architecture and network design benefits from two main architectural characteristics: Modularity and Extensibility. It is modular as it can be easily integrated in any network design the client might have, and it's extensible because it allows for quick upgrades or additions to any component.

4 A pilot implementation

Deploying the some of the basic components such as the domain controller, DNS server DHTP server or IIS server, from the proposed design is a common, if not trivial, task for all system administrators.

We propose an implementation totally based on Microsoft solutions. Before choosing it we carefully analyzed the pros and cons of it because this is a major architectural decision that would be hard to change.

The main disadvantage identified is that all components of the infrastructure are sealed (not open-source) and most of them require licensing (they are not free). On the other hand, being a security related project, it is important to keep up to date with any improvements or upgrades that appear in the field. Microsoft's constant development and upgrades of its products along with a high variety of good technical documentation are very good reasons to choose it. The good part about paying for it is that along with the product, you are also offered official support and guidance that can be extremely useful in emergency cases. Also the fact that a lot of companies use Microsoft technologies for they IT infrastructure made a decisive point in choosing Microsoft as a base for the infrastructure.

At the base of the PKI there is the "Certification Services" component from Windows 2003 Server. Certificate Services provides customizable services for issuing, managing and revoking digital certificates.

Even if it might seem like a trivial task, deploying a certificate authority requires several steps of planning and configuration even before installing begins. As the base operating system has been settled to Windows 2003 Server, we now know need to plan for the characteristics of the authority we are deploying.

After the installation is complete, we must check that everything is ok and complete the server's configuration by reviewing and updating registry information. All the steps for this configuration can be done via the CAs graphical interface, but using command line commands, enables building a batch that creates a good automation for a future configuration. The most important settings that are applied in this step are the CRL Distribution Points (CDP) and Authoritative Information Access (AIA) locations. The information will be added as an extension to every certificate issued by the Root CA, enabling end-applications using the certificate to locate the required information about the CRLs and parent Certificates. All commands are sent using the "certutil.exe" command line utility using the "-setreg" modifier.


```

certutil -setreg CA\CRLPublicationURLs
1:%windir%\system32\CertSrv\CertEnroll\%3%8%9.crl
15:ldap:///CN=%7%8,CN=%2,CN=CDP,CN=Public Key Services,CN=Services,%6%10
6:http://portal.cad-ca.net/crls/%3%8%9.crl

certutil -setreg CA\CACertPublicationURLs
1:%windir%\system32\CertSrv\CertEnroll\%1_%3%4.crt
3:ldap:///CN=%7,CN=AIA,CN=Public Key Services, CN=Services,%6%11
2:http://portal.cad-ca.net/crls/%3%8%9.crt

```

The commands above use define 3 locations for CDP and AIA parameters. Contents for each location is defined using a number resulted from binary representations of different types. The actual location address is defined after the “.” symbol and uses %x variables, which take their values according to the Certificate Services root certificate name and Active Directory names.

An Issuing certificate authority, unlike the Root Authority does not self-sign the base certificate. It uses a certificate signed by the parent authority (the Root CA in our case), thereby being granted the full trust of the Root authority.

The installation of the second and last CA server in our PKI, the Issuing CA consists of the same steps as the installation of the Root CA in consequence we not repeat them. There are however some subtle differences in the configuration process that we highlight and most important there is an additional step required for finishing the installation.

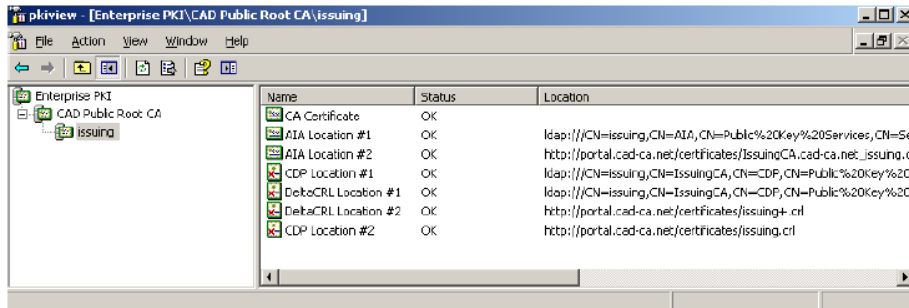
Only one additional step required in the post-installation configuration, consisting in creating a script that issues the Certificate Revocation Lists and then deploys them to all CRL distribution points defined. Deployment to active directory is made automatically when issuing them so in our case we still need to make them available on the website as stated in the configuration. To manage this with a script we need a public share on the web server. With the share in place, the script would look like the one presented below:

```

certutil -crl
copy /y %windir%\system32\certsrv\certenroll\*.cr? C:\_CA-Services\CertROOT
copy /y %windir%\system32\certsrv\certenroll\*.cr? \\portal\certificates

```

After all configuring of the CA is finished; we can start to plan for managing it. A great tool to quickly assess the PKI status is “PKI Health” available on Windows 2003 Resource kit available for free download on Microsoft’s site. After installing it, it can be run with the command “pkiview.msc” from the Run menu. A healthy PKI should look something like the picture below.



The application verifies if all the CA's in the organizations are working properly and if all CRL files are available at each location specified.

Another issue to be taken care of is configuring the CA for key archival in case of issuing certificates for key agreement purpose (data encryption). This is important because a lost private key will render all encrypted documents useless because there is no other way to decrypt them. In order to activate key archival, you need to create a Key Archival template, a Key Archival Manager and issue a certificate. The certificate is then set as Key archival certificate for the issuing CA enabling users who request encryption certificates to apply for key archival.

5 Open issues

For Digital Certificate Management the open issues refer to main possible extensions of our application.

Security updates and improvements. Security is a sensitive matter and has to be treated with care. For the PKI area, improving the security policies for the CA basic functions, writing and maintain the practices to be followed in any situations should be under constant watch.

Better key archival methodology. When a PKI starts issuing certificates for encryption purposes it is mandatory that proper private key archival to be set up. The PKI currently employs basic key archival for those certificates but this is a certain area that could be improved.

Add OCSP responder support. The PKI currently relies on CRL for certificate status distribution which has some major disadvantages. Adding OCSP support would greatly improve this area. A good option to take into account could be the migration of the Issuing CA to Windows Server 2008 that has an integrated OCSP responder.

Analyze migration of Certificate Services to Windows 2008. One of the reasons for which we chose Microsoft's platform for the PKI implementation was the good support and updates it provides and mentioned how crucial they are in this field of work. So, migrating to the latest version of Certificate Services would be the normal course of action. We already have the first white ball to Windows 2008 in the matter of the default OCSP responder but before making this important step all scenarios must be carefully analyzed.

Implement customized portal to issue certificates online. Issuing certificates online is a good feature of the current solution. But, the default user interface enabling the online certificate issuing could use a nice redesign so creating a new web-site, maybe integrated in the SharePoint portal could be a good future improvement point.

Automatic CRL issuing and distribution batch. The PKI administrator is now responsible with issuing the CRL certificates every time they are close to expiration and publish them to all CDP locations. The job is already half automated in the fact that all the operations are gathered in a command batch. To make it complete, a time schedule component should be added to the batch in order to remove all human factor from the equation. Having expired CRL's will cause all certificate validation to throw an error, stating that they cannot evaluate the certificate's validity, so the problem of always issuing timely CRL's is almost of critical importance.

6 Conclusions

Deployment of the PKI infrastructure based on Windows Server technology, based on a deployment plan designed by me following the first of all requirements and best practices presented in specialty books and articles. Architecting, designing and implementation of the cryptographic application used for enabling security to any kind of document, developed with Microsoft .NET Framework following guidance from several reference books in this domain. Finally we chose SharePoint for the online document library and developed and implemented a couple of workflows for testing and exemplification.

We consider that implementing the solution in a public institution would greatly improve the quality of services offered by eliminating the high waiting times for signatures and validations that increases dramatically with the number of people involved. Also, the project could improve document travel times, rendering document mailing useless in the same time as winning points for being a green project because it reduces the paper consumption.

References

1. Price, G. 2006. Public Key Infrastructures: A research agenda. *J. Comput. Secur.* 14, 5 (Sep. 2006), 391-417.
2. Bou Diab, W., Tohme, S., and Bassil, C. 2007. Critical vpn security analysis and new approach for securing voip communications over vpn networks. In *Proceedings of the 3rd ACM Workshop on Wireless Multimedia Networking and Performance Modeling* (Chania, Crete Island, Greece, October 22 - 22, 2007). WMuNeP '07. ACM, New York, NY, 92-96.
3. Staff 2009. E-Government Developments. *IWAYS* 32, 3 (Aug. 2009), 136-187.

4. Chou, C., Cheng, W. C., and Golubchik, L. 2010. Performance study of online batch-based digital signature schemes. *J. Netw. Comput. Appl.* 33, 2 (Mar. 2010), 98-114.
5. Zhao, C. and Gao, F. 2008. The business process model for IT service management. *WTOS* 7, 12 (Dec. 2008), 1494-1503.
6. Antonovich, M. 2008 *Office and Sharepoint 2007 User's Guide: Integrating Sharepoint with Excel, Outlook, Access and Word (Pro)*. APress.
7. Haron, G. R., Siong, N. K., and Bee, T. F. 2008. Translation of RFC 3820 to software codes using OpenSSL primitives. In *Proceedings of the Fourth IASTED international Conference on Advances in Computer Science and Technology* (Langkawi, Malaysia, April 02 - 04, 2008). S. K. Sahni, Ed. International Association Of Science And Technology For Development. ACTA Press, Anaheim, CA, 109-114.
8. Nambiar, V. P., Khalil-Hani, M., and Zabidi, M. M. 2009. Accelerating the AES encryption function in OpenSSL for embedded systems. *Int. J. Inf. Commun. Technol.* 2, 1/2 (Jun. 2009), 83-93.
9. Sun, H. and Ding, Y. 2008. An Adaptable Method of E-Workflow Composition Based on Distributed Awareness. In *Proceedings of the 11th Asia-Pacific Symposium on Network Operations and Management: Challenges For Next Generation Network Operations and Service Management* (Beijing, China, October 22 - 24, 2008). Y. Ma, D. Choi, and S. Ata, Eds. Lecture Notes In Computer Science, vol. 5297. Springer-Verlag, Berlin, Heidelberg, 503-506.
10. Pantaleev, A. and Rountev, A. 2007. Identifying Data Transfer Objects in EJB Applications. In *Proceedings of the 5th international Workshop on Dynamic Analysis* (May 20 - 26, 2007). International Conference on Software Engineering. IEEE Computer Society, Washington, DC, 5.
11. Muñoz, J. L., Esparza, O., Forné, J., and Pallares, E. 2008. H-OCSP: A protocol to reduce the processing burden in online certificate status validation. *Electronic Commerce Research* 8, 4 (Dec. 2008), 255-273.
12. Policelli, J. 2009 *Active Directory Domain Services 2008 How-To*. 1st. SAMS.
13. de Souza, T. C., Martina, J. E., and Custódio, R. F. 2008. Audit and backup procedures for hardware security modules. In *Proceedings of the 7th Symposium on Identity and Trust on the internet* (Gaithersburg, Maryland, March 04 - 06, 2008). IDTrust '08, vol. 283. ACM, New York, NY, 89-97.