PRIVACY IN CONTENT DISTRIBUTION NETWORKS

A framework description

R.J. Hulsebosch
Telematica Instituut, PO Box 589, 7500 AN, Enschede, The Netherlands
e-mail: Bob.Hulsebosch@telin.nl; phone: +31 (0)53 4850498; fax: +31 (0)534850400

Abstract: A framework is proposed for secure delivery of personalized content in the presence of a transparent intermediary while the privacy of the user is being preserved. The framework includes delegation of authority to the intermediary content distributor allowing him to act on behalf of the content provider. To maximize privacy, the framework enables users to control their profile and the ability to manage which personal information gets disclosed in the presence of a transparent intermediary.

Key words: Privacy, Proxy Certificate, Content Distribution, Identity management.

1. INTRODUCTION

Network edge caching proxies have been successfully deployed to accelerate web content delivery towards the user and reduce the load on origin web servers of the content provider (CP). In these so-called Content Distribution Networks (CDNs [1]), the intermediary edge caching proxies act as a gateway between web users and CPs and are nowadays more and more utilized with intelligent services other than simple caching. User-demanded services like virus scanning, content adaptation for personalization, and insertion of advertisements are suitable for such proxy servers and fit very well in the CDN model. The increased need for intermediaries to acquire access to security-related information, personal information about individual users and organizations, and proprietary information belonging to users and CPs, however, introduces new risks.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-0-387-35691-4_52

D. Gritzalis et al. (eds.), *Security and Privacy in the Age of Uncertainty*

Several problems have to be tackled for delivering personalized content to the user in a CDN. The intermediary CD is dealing with both the CP and the user. For personalized content adaptation the CD must obtain information about the user and must use this information in delivering and adapting the content and services that the user chooses. The main problem is related to the fact that the CD often operates transparently for the user and on behalf of the CP. The user does not have any knowledge of the existence of the CD. Normally, a user enters the URL of the CP in his web browser. If this CP has an agreement with a CD, the request will be (almost) transparently redirected to this intermediary. The user usually is not aware of this and shall only allow the CP insight in some of his profile attributes in return for content, not the CD. Therefore the CD requires credentials of the CP to convince the user he is acting on behalf of the CP in order to obtain the user credentials.

A solid framework for authentication and authorization is required to allow end-points to delegate authority to intermediaries so that they can adapt content on behalf of them. The framework must also facilitate secure handling of user information required for content adaptation services. In this communication we describe a certificate-based framework for secure content adaptation by the intermediary CD in a CDN. The framework enables delegation of authority to the CD allowing him to act on behalf of the CP. A Profile Manager (PM) is introduced for management and selective dissemination of the personal profile of the user. A scenario for anonymous browsing while personalized content is being delivered is discussed.

2. FRAMEWORK FOR SECURE PERSONALIZED CONTENT ADAPTATION AND DELIVERY

We propose a framework that is an extension of the IDsec concept for virtual identities on the Internet [2,3]. IDsec introduces a Profile Manager (PM) that allows the user to manage and control the distribution of personal credentials to the CP. Digital certificates are used as a tool for access to these credentials. Session Certificates (SCs) are introduced to provide for a pointer to the profile located at the PM. These SCs do not contain any other information. Access to profile information is based on the elements in the digital certificate of the CP. The high-level basic framework constituting the actors, their relations, and the most important attributes is shown in Figure 1.

The user logs in at the PM and alters his profile (1). When he logs off, a SC is returned (2). When the user requests for content, a SC is included in the request (3). The CD receives the request. Prior to sending the content, the CD wants to adapt the content he has obtained from the CP. In order to do this, the CD needs to know the preferences and capabilities of the user, i.e., his profile at the PM. For this reason the CD has to identify himself and reassure the PM that he is working on behalf of the CP (4). If this is true, the PM will grant the CD access to the credentials the user has given the CP access to (5). The CP provides the original content (6). Once the CD has obtained the credentials of the user, content adaptation can be performed (7). The adapted content is returned to the CD (8), and forwarded to the user (9).

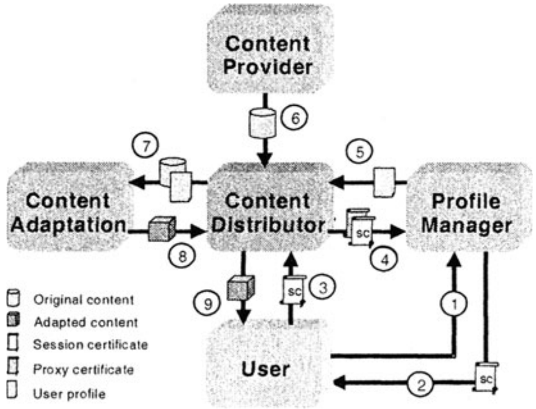

Figure 1: Framework for personalized content delivery in a CDN environment.

The fourth step in this scenario is the most critical one. Here the CD has to assure the PM he is operating on behalf of the CP in order to obtain access to the profile of the user. For this to work out we propose the use of proxy certificates (PC, [4]). The use of so-called proxy credentials for impersonation is a common technique to allow entity A to grant another entity B the right for B to authenticate with others as if it were A. In other words, entity B is impersonating entity A. The CD can use a PC to authenticate to the PM on behalf of the CP. The outcome is that the CD obtains access to the profile attributes of the user, which are stored at the PM. The Grid Security Infrastructure (GSI, [5]) for instance has proven the viability of impersonation as a basis for authentication. The delegation process and the X.509 PC profile used in our model have been extensively described in an IETF Internet Draft of the PKIX working group [6].

In the next sections we will have a closer look at the different processes that take place during content adaptation and delivery.

Profile Management

The PM allows the user to manage his profile attributes and the access to these attributes. The user - PM interactions are as follows (see Figure 2):

1. The user enters the PM via secure channel. Mutual authentication between the user and the PM is essential. Both the user and the PM need to be sure that they are dealing with the real party and not a bogus one.

2. Via the PM the user is able to structure his profile and to control access to his profile attributes. This means that the user can disclose a specific portion of his profile without revealing any other information.
3. A user-authenticated asymmetric session key pair is generated. This key is pair is used during the Internet session and has a short lifetime. The private key is securely stored while the public session key is sent to the PM via the secure connection. The PM generates a SC including the public session key of the user.
4. The PM returns the SC to the user. The SC identifies/represents the user on the current terminal but does not expose any user identification.

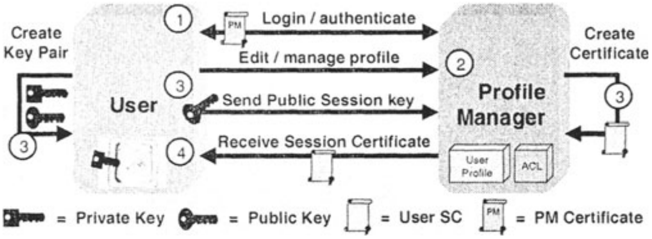

Figure 2: User profile management.

Users will be required to have multiple attributes for a number of different CPs. To assign CPs to the profile attributes they need the user is able to create Access Control Lists (ACLs). Access to the attributes is based on the distinguished name of the entity to which the certificate relates (e.g. the CP) *and* the digital signature of the certificate issuing party [2].

Delegation of content adaptation rights

The CD must be authorized by the CP in order to adapt the original content. In other words, the CP must delegate the content adaptation rights to the CD. Moreover, these delegation rights have to be used by the CD to obtain the profile of the user. Since access to the profile attributes is based on the distinguished name of the CP [2], the CD must convince the PM he is operating on behalf of the CP. The CD therefore must perform several operations that allow him to approach the PM server on behalf of the CP, i.e. the CD must obtain a PC of the CP. The delegation of content adaptation rights from the CP to the CD is depicted in Figure 3.

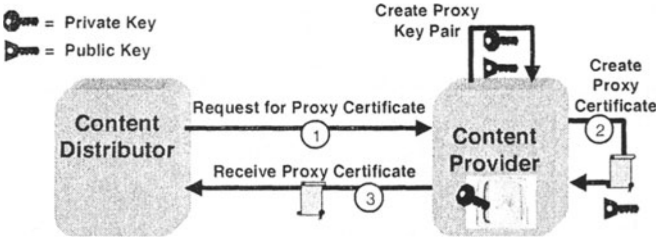

Figure 3: Delegation of content adaptation rights.

Content delivery

Figure 4 explains the process of personalized content delivery via the CD:

1. The SC is sent to the CD together with the URL request that the user makes.
2. To obtain the user profile attributes the CD must enter the PM. He uses his personal certificate for authentication at the server.
3. The PM authenticates to the CD and a secure channel is established.
4. The SC, PC and the CP certificate are sent to the PM via the secure channel. Both the SC and PC are used by the PM to determine the access rights of the CD. The CP certificate is used for verification of the PC signature. The signature of the CP certificate determines whether the PM trusts the CD - CP relation. If the PM trusts the signature, the SC allows him to perform a lookup for the session and user in order to retrieve the attributes of the user profile that the CP has access to. Access to attributes is controlled by the user and is based on the distinguished name of the PC since the ACL contains these names [2].
5. The profile is returned to the CD and is used to adapt the content. Naturally, the CP has stated the boundaries for the CD to adapt his content. A negotiation process between the limitations posed by the CP, the capabilities of the CD, and desires posed by the user for content adaptation is started.
6. The adapted content is sent to the user. If confidential information needs to be exchanged between the CD and user, the CD may use the public session key to encrypt the content sent. Only the user will be able to read the messages since he owns the private session key for decryption. If required by the CD, user authentication can be done via the private key of the SC. This would dynamically prove that the user is the owner of the SC.

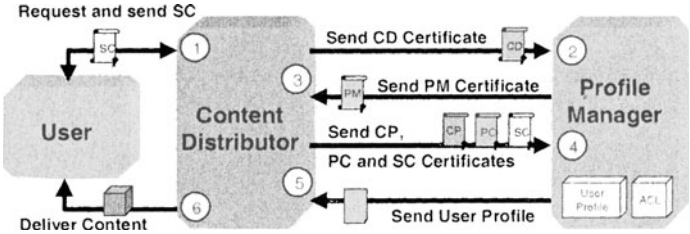

Figure 4: Personalized content delivery based on SCs.

Note that in this scenario the user has stayed anonymous for the CD and CP.

3. CONCLUSIONS

Personalized and localized content adaptation services are hot on the Internet. The rise of CDNs and the proposed value-added edge services fulfill and facilitate this need [7]. As a result, users on the Internet are often privy to large amounts of personal information, and should be very careful to prevent unintentional leakage of this information to other sources. Therefore, the parameters required for content adaptation need careful handling by an intermediary party. Though a CD typically operates transparently, the intermediary value-added services provided by a CD must therefore not be transparent: they have to be authorized by either the content requestor or the provider, corresponding to whom the service is being provided for. CPs may have a say on what adaptation is allowed, while on the other hand, the content consumers have the right to authorize services on their behalf. Authentication and (delegation of) authorization are two crucial aspects in establishing a trusted relationship in this complex situation.

We have presented a viable framework for a transparent intermediary CD to securely deliver personalized content to an anonymous user. The model allows a user to provide personal information while maintaining privacy by keeping control over its release. As a result the CD can offer services matching the need of each user. The model uses a PM for secure online management and administration of the private credentials of the user and issuing of SCs and encryption keys. The PM allows the user to determine what information gets released to which site. Thus, the responsibility of balancing access and privacy lies ultimately with the user. Since access is based on the distinguished name of the CP, delegation of authority is necessary to allow the CD access. Delegation of authority means that the authenticated identity of the CP must be carried over to the CD. This is accomplished by giving the CD a so-called PC to further perform operations on behalf of the CP.

REFERENCES

1. B. Hulsebosch, R. Brussee, H. Eertink, W. Huijsen, M. Rougoor, W. Teeuw, M. Wibbels, H. Zandbelt, Content Distribution Networks State of the Art, Telematica Instituut, 2001, https://doc.telin.nl/dscgi/ds.py/Get/File-15534/cdnsota.pdf.
2. H. Zandbelt, B. Hulsebosch, H. Eertink, IDsec: Virtual Identities on the Internet, Internet draft, http://idsec.sourceforge.net/draft-zandbelt-idsec-01.txt.
3. R.J. Hulsebosch, J.F. Zandbelt, E.H. Eertink, Virtual Identities on the Internet, submitted to the Proceedings of the Twelfth International World Wide Web Conference, 20-24 May 2003, Budapest, Hungary.
4. B.C. Neuman, Proxy-Based Authorization and Accounting for Distributed Systems, Proc. 13th International Conference on Distributed Computing Systems, Pittsburg, 1993.
5. Grid Security Infrastructure description: http://www.globus.org/security.
6. S. Tuecke, D. Engert, I. Foster, V. Welch, M. Thompson, L. Pearlman, C. Kesselman, Internet X.509 Public Key Infrastructure Proxy Certificate Profile, draft-ggf-gsi-proxy-03, 2002, http://www.gridforum.org/security/ggf5_2002-07/draft-ggf-gsi-proxy-03.PDF.
7. A proposed edge server technology can be found on the web site of the IETF OPES working group, http://www.ietf.org/html.charters/opes-charter.html.