

# Security for Replicated Web Documents

B.C. Popescu  
I. Kuz  
M. van Steen  
A.S. Tanenbaum

**Internal report IR-498, June 2002**

## **Abstract**

The WWW is experiencing explosive growth and an increasing number of security-sensitive applications make now use of it. To achieve worldwide scalability and reduce latency in handling user requests, many of these applications make extensive use of data replication through caches and Content Delivery Networks. However, such replication mechanisms place data on untrusted hosts, which introduces various security problems. In this paper we present an architecture that combines data content, replication strategies and security in one unified object model and offers integrity guarantees for Web documents replicated on non-secure servers.



*vrije* Universiteit

**Department of Mathematics and Computer Science**

# 1 Introduction

The World Wide Web is increasingly experiencing scalability problems. When a document suddenly becomes very popular (a phenomenon known as *flash crowds*), clients experience long delays in retrieving it. The single hosting server simply cannot cope (CPU-wise or bandwidth-wise) with the sudden high demands. A possible solution to this problem is data replication.

A common way of replicating Web documents is through content delivery networks (CDNs). These are normally run by companies that set up a large number of servers around the world, each server dedicated to hosting Web document replicas. The owners of these documents usually have no control over the replicating servers. Thus, in the CDN scenario, clients are retrieving data from hosts that are controlled neither by them, nor by the content providers. This situation is by no means specific only to CDNs; caching, presents the same problem.

Retrieving data from untrusted hosts introduces a series of security issues, which can be grouped in two categories - **document integrity** and **secure naming**. Document integrity deals with ensuring that a document retrieved by a client from an untrusted host has not been tampered with by that host. Secure naming deals with establishing secure associations between Web documents and the real-world entities in charge of them.

This paper introduces a security architecture that guarantees data integrity and secure naming for Web content even when this content is replicated on untrusted hosts. While read-only secure storage systems have already been proposed [6], the main contribution of this paper is to describe a design where data content, replication strategies and security are all integrated in one unified object model.

The rest of the paper is organized as follows: in Section 2 we introduce a new Web document model with a quick discussion on how such a model can be implemented in order to use most of the current Web infrastructure. Section 3 describes our security architecture, and how this architecture tackles secure naming and document integrity issues. In Section 4 we discuss related work, while in Section 5 we give a quick overview of our implementation effort, as well as our plans for future work.

## 2 Globe Web Documents

In the past years we have been working on the design and implementation of Globe [16] - a middleware architecture that allows the development of distributed applications based on replicated shared objects. One such distributed application is GlobeDoc [17]. GlobeDoc objects are the building blocks for our new Web docu-

ment model.

In the GlobeDoc model of the Web, a Web site is composed of related Web documents. A Web document itself is a collection of logically related Web resources. These Web resources are referred to as a document's **page elements** and can be anything that is accessible over the Web (e.g., HTML files, text files, images, audio files, video files, applets, etc.). The relation between the resources contained in a Web document is generally stronger than that between the documents contained in a Web site. For example, an organization's Web site contains a collection of Web documents that are somehow related to that organization, while a Web document representing a news story would contain only the elements directly relating to that story (e.g., the HTML page plus any icons and images relating to the story or the page layout).

In GlobeDoc, a Web document is encapsulated in a Globe distributed shared object which contains that document's elements in its state. Such a GlobeDoc object offers methods that allow clients to access and modify its state on a per-element basis. The hyperlinked structure that is normally provided by Web pages is maintained in GlobeDoc. A relative hyperlink contained in some GlobeDoc object's element refers to another element in that same object. Likewise, an absolute hyperlink may refer to an element contained in another GlobeDoc object.

Every GlobeDoc object is identified by a unique object ID (OID). This is a 160-bit number which does not contain any location information and is not human readable. The GlobeDoc architecture, therefore, contains two services that facilitate locating objects. The Naming Service [3] maps human readable object names onto OIDs. The Globe Location Service [15] maps OIDs onto contact addresses, which contain information about where and how to contact a GlobeDoc object.

By virtue of it being a Globe distributed shared object, a GlobeDoc can distribute (replicate) its state over multiple physically separated address spaces (or machines). As such, requests for an object's state will be distributed over these various machines, thereby spreading the total load generated by the requests and preventing any single machine hosting the GlobeDoc from becoming overloaded. Similarly, by strategically replicating a GlobeDoc object's state so that it is close to large concentrations of clients, the traffic at each replica will have a local character, increasing responsiveness for clients and decreasing overall network traffic and saturation. This architecture is ideally suited to the creation of (peer-to-peer) content delivery networks.

Also, because Globe makes the state distribution transparent to clients and because the distribution policies can be determined independently per Globe object, GlobeDoc makes it possible to apply distribution strategies on a per-document basis. In this way GlobeDoc allows replication of Web documents without imposing any single global replication policy on all documents. In [13] we have shown that

applying per-document distribution strategies can lead to better efficiency than the application of one-size-fits-all strategies.

## 2.1 GlobeDoc Services

A client accesses a GlobeDoc object using a regular Web browser. However, standard Web browsers cannot directly connect to GlobeDoc objects and invoke its methods to retrieve the page elements. For this task, browsers rely on a GlobeDoc proxy server to intercept requests to GlobeDoc objects and manage the interaction with the object (i.e., connect to the object and retrieve its elements). Because regular Web browsers do not understand GlobeDoc object names, **hybrid URLs** are used. These are just regular URLs that start with a distinguishing prefix (which in our case is *http://enter.globeworld.org*). GlobeDoc and page element names are then embedded in these hybrid URLs.

GlobeDoc proxies connect to GlobeDoc objects by binding to them. The result of binding to a GlobeDoc object is that a local representative of that object is placed in the address space of the binding process. A local representative is a local part of the GlobeDoc object. It can be a simple object proxy, forwarding method invocations to other replicas, or it can be a full replica containing a local copy of the GlobeDoc object's state. The client, however, is unaware of this and simply invokes methods from the local representative as though it was a local object.

Binding itself consists of two distinct phases: (1) finding the object, and (2) installing the local representative. This is illustrated in Figure 1. Finding an object is separated into a name-lookup and a location-lookup step; installing the local representative requires that a suitable contact address and implementation be selected.

### 2.1.1 Naming Service

To find an object, a process must resolve the object's name to an OID. This is done by passing the object name to a naming service. The naming service returns the object's OID. Whereas an OID uniquely identifies an object (i.e., an object always has exactly one OID), an object may have multiple names that resolve to its OID. An object name is simply a human-readable string that represents an object. It is up to the naming service to interpret this name and resolve it to an OID.

### 2.1.2 Globe Location Service

Resolving an object's name leads to an OID. An OID, however, does not identify the location of an object and its replicas. To find the actual object, the OID must first be passed to a location service, which returns one or several contact addresses

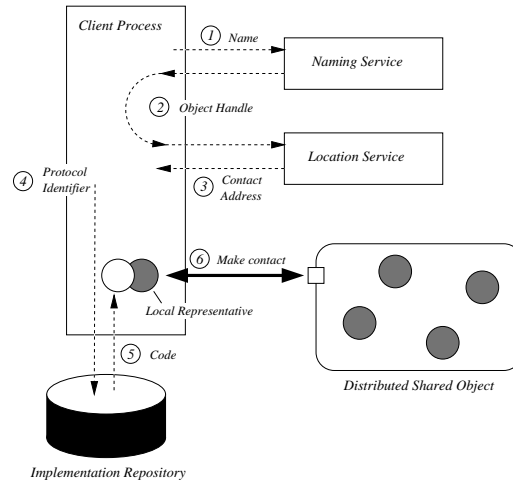


Figure 1: Binding to a GlobeDoc object

which represent GlobeDoc object contact points. The location service is responsible for storing contact addresses and resolving OIDs to these contact addresses. Besides looking up contact addresses for objects, the location service also allows addition, deletion and modification of contact-address mappings.

The Globe Location Service is implemented as a distributed search tree. In this tree, the world is divided into a hierarchical set of domains. At the lowest level there is one domain per site; a collection of sites form a region, etc. An object is recorded at each site where it has a contact address, and recursively in each enclosing region up to the root of the tree. Initially, a record at the site level contains the actual contact addresses while records at higher levels contain pointers to the next lower level. Recording an object at multiple levels allows searches with expanding rings: a search starts at the local site, followed by the local region, then the next higher level region, etc., and eventually followed by the root.

### 2.1.3 Object Server

A GlobeDoc object's replicas are implemented as stateful local representatives hosted on object servers. An object server is a process that provides an **address space**, **contact points** and **runtime services** to the local representatives that it hosts. The object server also manages a local representative's access to local resources such as local disks and networking resources. Besides simply hosting LRs, the Globe object server also has a remotely accessible interface that allows other local representatives, other Globe object servers, or administrators to request ser-

vices from it. These services include the creation and destruction of GlobeDoc objects and their replicas.

### 3 The GlobeDoc Security Architecture

In this section we will describe the mechanisms used to make GlobeDoc objects secure. We have been using well-known security and cryptographic techniques as building blocks, and combine them in an architecture that can offer secure naming and data integrity guarantees.

Following the model of the Web master for regular Web documents, we require that behind each GlobeDoc object there is a person or organization - the **object owner** - that is in charge of it. The owner is the one who creates the object and is responsible for providing permanent storage for it, updating it, and setting up a replication and security policy for the object. Furthermore, we require each GlobeDoc to have a unique public/private key pair associated with it, which is generated by the owner when the object is created. As we will describe next, the object owner uses the object's private key to sign the object's state (files) before it replicates it, while clients use the public key to perform integrity checks when retrieving parts of the object's state from untrusted hosts.

#### 3.1 Secure Naming

##### 3.1.1 The Problem

Secure naming deals with creating secure and trustworthy associations between Web resources and qualified names, which in turn convey information about the real world entities that are behind these resources. In the current Web, name binding is accomplished through the Domain Name System (DNS) [12], a hierarchical distributed database that translates between human-readable resource names and their IP addresses. Data integrity requirements were left out from the original DNS design, which made it vulnerable to various types of masquerading attacks. To prevent these, a secure DNS (DNSsec) framework has been proposed [5] and is currently being deployed.

All DNSsec does is establish secure associations between resource names and sets of IP addresses. However, clients cannot authenticate IP addresses, so their attempts to retrieve data from servers at these addresses are vulnerable to man-in-the-middle attacks. To prevent this, secure Web servers have public/private key pairs which they use to authenticate to clients. The server's public key needs to be certified by some Certificate Authority (a trusted third party, Verisign for example) through a digital certificate that binds the public key to the server's name (which

presumably can be traced to the qualified path name solved through DNSsec). Having this public key, securely associated to the server's name (by the CA), and to the server's IP address (by DNSsec), the client can then authenticate the server. In this way, the man-in-the-middle attack becomes impossible, and the name binding and secure channel establishment can be safely combined.

The problem with this approach is that it does not work well with dynamically replicated Web documents. Although DNS supports mirroring Web sites by allowing multiple addresses associated with the same name, the basic assumption is that mirroring is more or less static. Dynamic active replication of Web documents would put serious strain on the resource-record caching that makes DNS so efficient.

In conclusion, the approach as followed in DNSsec for secure naming of Web documents will not work well for the GlobeDoc model we are proposing. For GlobeDoc objects we need a naming mechanism that both:

- Securely associates objects with the real-world entities behind these objects.
- Supports efficient lookup operations even when objects are massively replicated and replicas' network addresses change frequently.

### 3.1.2 Our Solution

As described in Section 2, each GlobeDoc is identified through a unique 160-bit object ID (OID). Secure name binding requires creating a secure association between an OID, the object's public key, and the real-world entity (individual, company, organization) that is in charge of the object.

First, we will examine how to securely link the object's public key to the OID. This can be accomplished by having the OID be the 160-bit secure hash *SHA-1* [1] of the object's public key, in fact creating a **self-certifying OID** for the object. The *SHA-1* secure hash function has the property that it is computationally intractable to find two different inputs with the same hash. Therefore an OID obtained in this manner is securely linked with the public key of the object. This approach of using self-certifying OIDs is similar to the one taken by the designers of the SFS system [11], who have pioneered the concept of self-certifying resource names [10].

Now, the only thing left is to securely associate self-certifying OIDs with the real-world entities in charge with their corresponding objects. As stressed out in the SDSI document [14], we believe that establishing trust in a remote entity (trusting that entity to be who it claims to be) is such a security-sensitive task that it is better to give most control over it to the users themselves. To facilitate this, GlobeDoc objects have a special security interface that can be used by clients (in fact by their proxies) to retrieve any certificates the objects can provide to prove their identity.

The users themselves can specify a number of CAs they trust, and store their public keys with their user proxy. When the user requests an element part of a GlobeDoc, its proxy requests identity certificates that match the user's list. For the first match found, the proxy displays the naming information in the certificate. The user examines this information and can then make a decision how much trust to put in that object.

The CA-mediated secure name-binding mechanism we just described is appropriate when Web objects are used for highly sensitive applications such as e-commerce or e-banking. For less sensitive applications, users may find the name-binding information provided by a DNSsec-like name service sufficient. The good news is that DNSsec can be extended to support self-certifying OIDs by storing them in the resource records instead of IP-addresses [3]. The great advantage of this would be that DNS would store only **location-independent** data. The location-dependent information associated with each object (the addresses of its replicas) is retrieved in an additional step from the Globe Location Service. This two step secure name resolution allows us to overcome the problems DNS has when tracking dynamically replicated documents.

One point we want to stress here is that users do not have to trust the information stored in the Location Service. In fact, using the Location Service is not even mandatory for GlobeDoc objects; a per-object dedicated directory replica keeping track of all the other replicas' contact points would accomplish the same task. A malicious Location Service server can return false contact points to its clients, making these clients bind to replicas which are not part of the objects they want to contact. However, as we will see in the next section, clients can easily detect when the data they are retrieving is not part of the GlobeDoc object they want to browse, so the most harm a malicious Location Service server can do is a temporary denial of service.

As a conclusion, secure naming for GlobeDoc objects is achieved through self-certifying OIDs, which can be stored in DNSsec resource records. When additional security guarantees are necessary, object-provided certificates, signed by trusted CAs can also be employed. Because only location-independent information is stored, this mechanism has the two properties we required in Section 3.1.1: it securely associates Web objects to the real-world entities in charge of them, and can efficiently support massively replicated objects even when replicas network addresses change frequently.

## 3.2 Document Integrity

### 3.2.1 The Problem

Most of the data transfer in the current WWW is insecure. Clients simply connect to Web servers and request the documents stored there. This approach is clearly vulnerable to man-in-the-middle attacks, not to mention malicious caches. In such an attack scenario, an active attacker intercepts the client's request, and answers with his own document. Although such attacks are infrequent, that is probably as much due to the lack of determination on the part of the attackers as to the inherent security of the underlying network architecture. Furthermore, due to the possibility of such attacks, the security of an HTTP request is downgraded to the security of the weakest network link/router on the request path.

In the current WWW, the most common protection against such attacks is through TLS (Transport Layer Security) [4]. TLS uses public-key cryptography to authenticate servers and establish secure channels between servers and clients.

The main problem with TLS is that it requires servers to be trusted. The secure channel between the client and server does not help at all if a malicious server sends bogus data over it. For this reason, TLS allows documents to be replicated only on trusted servers, which greatly restricts the set of acceptable hosts. In addition to this, TLS introduces a significant delay for the connection establishment due to the computationally expensive public-key operations that occur in this phase. For these reasons, only highly security sensitive Web applications (such as e-commerce, or e-banking) make use of TLS.

Turning back to our GlobeDoc objects, we can see that a mechanism like TLS is clearly not suited for ensuring data integrity. As mentioned in Section 2, GlobeDoc objects dynamically place their replicas on (possibly) untrusted object servers close to where their client requests are coming from. Our assumption is that most of these servers are honest, but we need to consider the possibility that some of them may try to replace the documents they host with fake data. Therefore, we need a security mechanism that enforces the following three properties on the replicated state of a GlobeDoc:

- **Authenticity** - the document the client receives from a server has indeed been created by the object's owner. No attacker or malicious server should be able to pass one of their own documents as being part of the object.
- **Freshness** - the client is guaranteed to receive the most recent version of a document part of a object. No attacker or malicious server should be able to pass old versions of a document and convince the client they are fresh.

- **Consistency** - the client is guaranteed to receive a document, part of the object, that is consistent to what she has requested. No attacker or malicious server should be able to replace the requested document with another fresh document part of the same object.

Our aim is to come up with a security design that enforces these three properties, and at the same time is efficient and lightweight (especially on server load), so that it can be employed with every type of Web application.

### 3.2.2 Our Solution

As we mentioned in Section 2, a GlobeDoc consists of a number of page elements. These can be HTML source files, images, Java applets and so on. We preserve the integrity of the object's state by having an **integrity certificate** associated with each page element (file) part of the object. As shown in Figure 2, this is a digital certificate signed by the object's owner that contains a *MD5* hash [8] of the page element, the element's name, a validity interval, and some optional revocation information.

Integrity Certificate
Page Element Name
Page Element Hash
Validity Interval
CRL Directory Address
CRL Validity Interval
Signature

Figure 2: Integrity certificate for a GlobeDoc page element

Every server that hosts GlobeDoc replicas is required to store all of the object's page elements **and** their corresponding certificates. As we mentioned in Section 3.2.1, page elements retrieved from untrusted servers should be authentic, fresh, and consistent with the user request. The integrity certificate associated with each such element allows the user to check for these properties in the following steps.

- Using the object's public key, the client verifies that the signature on the integrity certificate has been generated using the object's private key (authenticity).
- The client computes the *MD5* hash of the page element and makes sure it is the same as the one in the certificate (authenticity).
- The client checks the time of retrieval to ensure it falls in the validity interval specified in the certificate (freshness).

- The client checks the “element name” field in the certificate to ensure it is the same as the element name she has requested (consistency).
- If the certificate requires revocation verification, the client queries the revocation directory to make sure the certificate has not been revoked (freshness).

The optional revocation information stored in the integrity certificate consists of the freshness interval (how often a certificate revocation list - CRL - is published) and the address of the on-line directory where this list will be stored. The CRL is issued and signed by the object’s owner. One possibility is to store this CRL on the server that hosts the object’s master replica, so that client proxies can retrieve it from there (CRL pull). Yet another possibility is that the object’s master replica periodically distributes the CRL to all the hosting servers, which send it to the clients together with the data and integrity certificates (CRL push).

### 3.3 Putting the Pieces Together

The security architecture for GlobeDoc objects results from combining the various security mechanisms described so far. The key advantage of this architecture is providing integrity guarantees for end-users, even when the Web content they are browsing is replicated on untrusted servers. Figure 3 shows how secure browsing through GlobeDoc objects works.

The user starts with a hybrid URL (as described in Section 2), with an embedded object and page element name. The user can either directly type the hybrid URL in the Web browser, or get to it through an external link in some other GlobeDoc object. When the user’s proxy receives this URL, it will resolve the object name by contacting the secure naming service, thus obtaining a self-certifying OID. The proxy then queries the Location Service for that OID and finds the closest object replica. The Location Service is not trusted, so there is a chance that the address it returns may point to a replica which is not part of the object. However, this can cause at most denial of service for the user, since she can always check the authenticity of the data retrieved from the replica. Our assumption is that the Location Service provides accurate information most of the time; if attackers are able to corrupt some of the LS’s servers, this can be easily detected, and appropriate measures (rebooting servers, restoring the original data content from backups, etc.) be taken.

Once the user is connected to a GlobeDoc object replica, her proxy first retrieves the object’s public key, takes its *SHA-1* hash and makes sure it matches the self-certifying OID. As an extra security check, the proxy can request some additional identity proof from the object, in the form of a name certificate signed by one

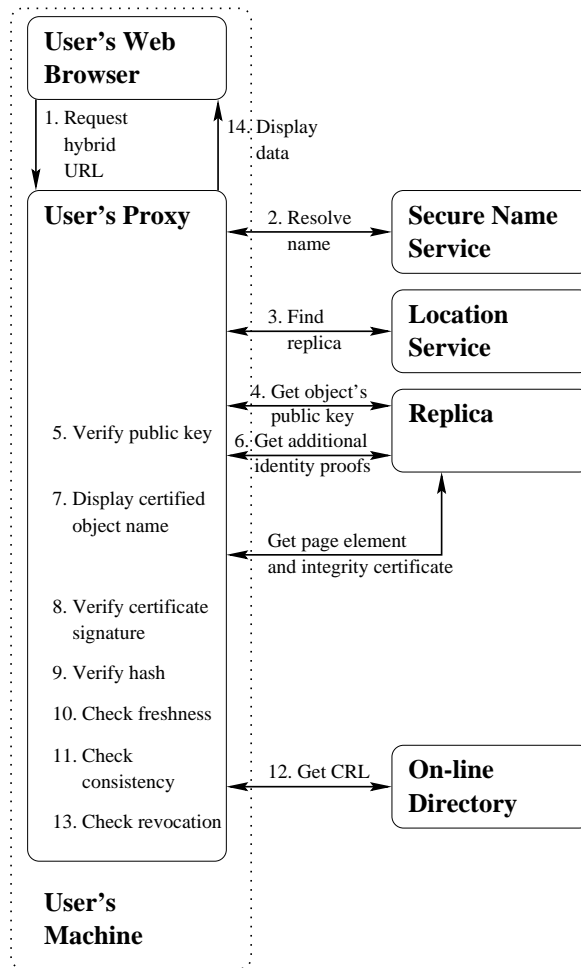


Figure 3: Secure Web browsing through GlobeDoc objects

of the CAs trusted by the user. If such a certificate is found, the naming information in the certificate is displayed to the user in a “*Certified as:*” window. At this point, the secure binding between the proxy and the object is complete; the proxy now requests the page element specified in the URL and the element’s integrity certificate; once it has them, it performs the authenticity, consistency and freshness tests discussed in the previous section. Next, if the integrity certificate contains any revocation information, the proxy retrieves a fresh CRL (from an on-line directory, or from the replica itself), and ensures the certificate is still valid. Finally, if all these tests are successful, the page element is sent to be displayed in the user’s

browser, otherwise a “Security Check Failed” HTML document is generated.

## 4 Related Work

The WWW is experiencing explosive growth, and an increasing number of security-sensitive applications make now use of it. Because of this, mechanisms that can make the Web more secure are a hot topic of research and a number of solutions have been proposed. What makes the GlobeDoc approach distinctive is the fact that it considers not only security but other aspects, such as data content aggregation and replication, creating a unified Web document object model.

The current “state of the art” for secure Web browsing is the combination of DNSsec [5], certified Web servers’ public keys and TLS [4]. In Section 3 we have described how these mechanisms can be combined in order to achieve secure naming, server authentication and integrity guarantees for the data received by the clients. However, this approach does not support document replication on untrusted hosts. Furthermore, since DNSsec records are used to store location-dependent information (IP addresses), this approach does not scale well for dynamically-replicated documents.

Like GlobeDoc, the read-only Secure File System (r-oSFS) [6] focuses on securely replicating data on untrusted servers. The basic architectural element for r-oSFS is the file system; because of this it is possible to use r-oSFS as a middleware and implement various distributed applications on top of it, a secure Web infrastructure being one such possible application [7]. In order to guarantee the integrity of data replicated on untrusted hosts, r-oSFS constructs a hash tree by applying a secure hash function (*SHA-1*) on the data blocks and i-nodes of the file system. This approach is very efficient, since only the root of the tree needs to be signed by the owner, but has the drawback that only one global (per-file system) consistency interval can be supported, instead of allowing per-file freshness constraints.

Although r-oSFS file systems can be replicated on untrusted hosts, there is little support for the actual replication. In contrast, each GlobeDoc comes with its own replication policy which is **part of** the object itself; this allows for very fine-grained (per page-element) replication policies to be defined, which has been proven [13] to greatly improve performance. Following the same logic, the GlobeDoc security architecture uses per page-element integrity certificates, which allow owners to set per page-element freshness constraints (which is not possible with r-oSFS).

The Globe Distribution Network (GDN) [2] is another distributed application built on top of the Globe middleware. GDN aims at providing secure distribution of free software packages. Like GlobeDoc, GDN also relies on untrusted servers

for data replication, and employs similar cryptographic techniques to guarantee data integrity. However, GDN focuses more on data content traceability to prevent the distribution of illegal material, and is less concerned with the freshness and consistency aspects of replication.

Finally, OceanStore [9] is a project that aims at using untrusted hosts to provide a “*utility infrastructure designed to span the globe and provide continuous access to persistent information.*” To accomplish this ambitious goal, the designers of the system make use of peer-to-peer technologies, such as associative storage, distributed routing, and probabilistic query algorithms. Although both make use of untrusted storage, OceanStore and GlobeDoc have slightly different goals: OceanStore assumes that all the storage is untrusted, and focuses on high data redundancy to prevent loss due to malicious hosts or catastrophic events. GlobeDoc on the other hand assumes that each document has access to some secure storage provided by its owner (the traditional Web document model), and relies on untrusted hosts for replication in order to improve performance. Although we recognize the many revolutionary ideas OceanStore introduces, we believe that the GlobeDoc Web document model is more appropriate for the next generation of secure WWW services.

## **5 Implementation Status and Future Work**

We are in the advanced stages of the implementation process, and we are making rapid progress because we are able to re-use some of the security functions developed for the GDN project. We have developed a software tool - *tika* - that takes a number of page elements and creates the corresponding GlobeDoc object. *Tika* is to be used by object owners on their (secure) machines; it uses the Java 1.4 security library functions to generate the GlobeDoc object’s public/private key pair; the private key is then safely stored in the owner’s Key Store. *Tika* then generates the integrity certificates for all the object’s page elements; the object’s state will include the page elements, their corresponding certificates and the object’s public key. We have also used Java to implement the GlobeDoc proxy, which now fully supports object binding and integrity certificate verification.

Currently, our prototype does not offer any support for integrity certificate revocation, which we plan to integrate in the near future; another feature we need to integrate is the CA mediated authentication mechanism, as well as an interface that would allow users to input the public keys of the Certificate Authorities they trust.

## References

- [1] Secure Hash Standard. FIPS 180-1, Secure Hash Standard, NIST, US Dept. of Commerce, Washington D. C. April 1995.
- [2] A. Bakker, M. van Steen, and A. Tanenbaum. A Law-Abiding Peer-to-Peer Network for Free-Software Distribution. In *Proc. IEEE Int'l Symp. on Network Computing and Applications*, Cambridge, MA, February 2002.
- [3] G. Ballintijn, M. van Steen, and A. Tanenbaum. Scalable user-friendly resource names. *IEEE Internet Computing*, 5(5):20–27, 2001.
- [4] T. Dierks and C. Allen. The TLS protocol version 1.0,. IETF RFC 2246, January 1999.
- [5] D. Eastlake. Domain name system security extensions. RFC 2535, March 1999.
- [6] K. Fu, M. Kaashoek, and D. Mazieres. Fast and secure distributed read-only file system. In *Proc. 4th USENIX Symp. on Operating Systems Design and Implementation*, pages 181–196, San Diego, CA., Oct. 2000.
- [7] M. Kaminsky and E. Banks. SFS-HTTP: Securing the web with self-certifying URLs. [citeseer.nj.nec.com/470041.html](http://citeseer.nj.nec.com/470041.html).
- [8] C. Kaufman, R. Perlman, and M. Speciner. *Network Security*. Prentice Hall, Upper Saddle River, NJ, 1995.
- [9] J. Kubiawicz, D. Bindel, Y. Chen, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. OceanStore: An Architecture for Global-scale Persistent Storage. In *Proc. 9th ACM ASPLOS*, pages 190–201, Cambridge, MA, November 2000. ACM.
- [10] D. Mazieres and M. F. Kaashoek. Escaping the evils of centralized control with self-certifying pathnames. In *Proc. ACM SIGOPS*, Sep 1998.
- [11] D. Mazieres, M. Kaminsky, M. F. Kaashoek, and E. Witchel. Separating Key Management from File System Security. In *Proc. 17th Symp. on Operating Systems Principles*, pages 124–139, Kiawah Island, SC, 1999.
- [12] P. V. Mockapetris and K. J. Dunlap. Development of the domain name system. In *SIGCOMM*, pages 123–133, 1988.

- [13] G. Pierre, M. van Steen, and A. Tanenbaum. Dynamically Selecting Optimal Distribution Strategies for Web Documents. *IEEE Transactions on Computers*, 51(6), 2002. To appear.
- [14] R. L. Rivest and B. Lampson. SDSI – A Simple Distributed Security Infrastructure. Presented at CRYPTO’96 Rumpsession, 1996.
- [15] M. van Steen, F. Hauck, P. Homburg, and A. Tanenbaum. Locating Objects in Wide-Area Systems. *IEEE Commun. Mag.*, pages 104–109, January 1998.
- [16] M. van Steen, P. Homburg, and A. Tanenbaum. Globe: A Wide-Area Distributed System. *IEEE Concurrency*, pages 70–78, January-March 1999.
- [17] M. van Steen, A. Tanenbaum, I. Kuz, and H. Sips. A Scalable Middleware Solution for Advanced Wide-Area Web Services. *Distributed Systems Engineering*, 6(1):34–42, March 1999.