

# Globule: a Platform for Self-Replicating Web Documents

Guillaume Pierre  
gpierre@cs.vu.nl

Maarten van Steen  
steen@cs.vu.nl

Vrije Universiteit, Division of Mathematics and Computer Science,  
de Boelelaan 1081, 1081 HV Amsterdam, The Netherlands

**Keywords:** Replication, adaptive systems, World-Wide Web

## Abstract

Replicating Web documents at a worldwide scale can help reducing user-perceived latency and wide-area network traffic. This paper presents the preliminary design of Globule, a platform which automates all aspects of such replication: server-to-server peering negotiation, creation and destruction of replicas, selection of the most appropriate replication strategies on a per-document basis, consistency management and transparent redirection of clients to replicas. To facilitate the transition from a non-replicated server to a replicated one, we designed Globule as a module for the Apache Web server. Therefore, converting normal Web documents should require no more than compiling a new module into Apache and editing a configuration file.

## 1 Introduction

Large-scale distributed systems often address performance issues by way of caching and replication. In the Web, attention has traditionally concentrated on caching and in much lesser extent to replication. Recently, Web hosting services such as Akamai [6] and Digital Island [3] have started to emerge as the solution to achieve scalability through replication. In this approach, content is replicated to places where user demand is high. Web hosting services have the advantage over static mirroring of Web sites in that decisions on *where* to replicate content can be made automatically. However, current hosting services apply a single system-wide policy concerning which consistency is needed and how it is to be realized.

Many protocols have been proposed to achieve caching or replication, each of which presents specific advantages and drawbacks. We have shown in a previous article that no single policy is efficient for replicating all documents. Significant performance improvements can be obtained by replicating each doc-

ument with the policy that is most suited to its access pattern [9].

To associate a replication policy with each document separately, we consider a document and all its replicas as a single replicated distributed object. All replication issues are encapsulated *inside* the object. Such an architecture allows a document to monitor its own access pattern and to dynamically select the replication policy that suits it best. When a change is detected in access patterns, it can re-evaluate its choice and switch policies on the fly [10].

Although we have demonstrated the feasibility of this approach, we have not yet addressed its integration in the existing Web infrastructure. There are two problems that need to be addressed. First, servers need to be adapted so that they can support adaptive per-document replication policies. Second, servers need to cooperate to allow replicas to be dynamically installed and removed, and to redirect clients to the nearest replica. One additional requirement is that adaptations should fit into the current Web infrastructure, requiring minimal modifications to existing servers.

This paper presents the preliminary design of Globule, a platform for hosting adaptive Web documents. It is designed as a module for the popular Apache server [1]. Converting normal Web documents should require no more than compiling a new module into Apache and editing a configuration file. Globule handles all management tasks: discovering and negotiating with remote sites for hosting replicas; replicating static and some dynamic documents; and transparently redirecting clients to their closest replica.

The paper is structured as follows: Section 2 describes our document and server models; Section 3 details the architecture of the system, and Section 4 shows how such a system can be implemented as an

Apache module. Finally, Section 5 presents related work and Section 6 concludes.

## 2 General Model

Our system is made of servers that cooperate in order to replicate Web documents. This section describes our document and server models.

### 2.1 Document Model

In contrast to most Web servers, we do not consider a Web document and its replicas as a collection of files, but as a single distributed entity. Each distributed document is implemented as a single distributed object replicated across the Internet. Each such object implements a standard interface (`put()`, `get()`, ...) to allow for delivering and modifying its content.

Our system is based on Globe, a platform for large-scale distributed objects [14]. Its main novelty is the encapsulation of issues related to distribution and replication *inside* the objects. We designed documents that dynamically select their own replication policy. Our “documents-are-objects” model is also a key to replicating dynamic Web documents.

#### 2.1.1 Adaptive Replicated Web Documents

We have shown in previous papers that significant performance improvements can be obtained over traditional replicated servers by associating each document with the replication strategy that suits it best [9, 10]. Such per-document replication policies are made possible by the encapsulation of replication issues inside each document.

The selection of the best replication policy is realized internally to each document by way of trace-based simulations. Replicas transmit logs of the requests they received to their master site. At start-up or when a significant access pattern modification is detected, the master re-evaluates its choice of replication strategy. To do so, it extracts the most recent trace records and simulates the behavior of a number of replication policies with this trace. Each simulation outputs performance metrics such as client retrieval time, network traffic and consistency. The “best” policy is chosen from these performance figures thanks to a cost function. More details about adaptive replicated documents can be found in [10].

#### 2.1.2 Replicating Dynamic Web Documents

Many documents are not made from static content, but are generated on the fly. For each request, a Web server executes a request-specific piece of code whose output is delivered to the client. This code can in turn access external resources such as databases, execute

shell commands, and issue network requests for generating a *view* of the dynamic document.

Replicating dynamic documents requires replicating the code as well as all data necessary for its execution (databases, etc.). This can be done by encapsulating the code and the databases in replicated distributed objects. Every request sent to the database is intercepted by the encapsulating object before being executed. If the request is likely to modify the internal state of the database, then it is propagated and applied to the other replicas in order to maintain a consistent state [14].

Our object model does not differentiate between static and dynamic documents. Therefore, dynamic documents are considered as objects implementing the same interface as static documents, but they differ in their implementation: whereas static documents use always the same implementation to access various internal states, dynamic documents differ in both their internal states and method implementations.

The main issue with respect to dynamic documents arises when converting existing dynamic documents into objects. The Web server must be able to determine automatically which resources, files or databases are accessed by each dynamic document in order to include them inside the object. This can be difficult for documents such as CGIs, where the server delegates request handling to an arbitrary external program. However, a large portion of dynamic documents such as PHPs and ASPs are in fact scripts interpreted by the Web server itself. In this case, the server knows the semantics of the document, and can often automatically detect which resources are required by the document.

Dynamic documents that the server can not analyze cannot be encapsulated into objects. Therefore, they are not replicated in our approach. Client requests are directed to the master site.

## 2.2 Cooperative Servers

One important issue for replicating Web documents is to gain access to computing resources in several locations worldwide (CPU, disk space, memory, bandwidth, etc.). On the other hand, adding extra resources locally is cheap and easy. Therefore, the idea is to trade cheap local resources for valuable remote ones. Servers automatically negotiate for resource peering. The result of such a negotiation is for a “hosting server” to agree to allocate a given amount of its local resources to host replicas from a “hosted server.” The hosted server keeps control on the resources it has acquired: it controls which of its clients are redirected to the hosting server, which documents are replicated there and which replication policies are being used.

Of course, servers may play both “hosting server” and “hosted server” roles at the same time: a server may host replicas from another server, and replicate its own content to a third one. We use these terms only to distinguish roles within a given cooperation session.

Servers communicate with each other in an NNTP-like fashion. Each server is configured to communicate with a few other servers, which in turn communicate with more servers. This network of servers allows information to be spread efficiently among the whole group of cooperating servers: messages are propagated from host to host through the entire network. Each message carries a unique identifier, so that duplicate message transmission can be avoided.

This communication channel allows servers to broadcast data such as their name, location and current publicly available resources. When a server searches for a new hosting server in a certain area, it uses its local list of server locations to choose a site which would be suitable. It then contacts it directly and starts a negotiation phase. This negotiation can, for example, determine the price asked for accessing the required resources [2, 15].

Other types of data that flow through the network of servers are information about disconnected nodes (which may be caused by a failure of simply by a server deciding to leave the group). That way, hosted servers are informed when their hosting servers become unreachable, so they can stop redirecting clients to them.

## 2.3 Security Issues

Such a system presents obvious security risks. A malicious server could accept to host Web document replicas, and deliver modified versions to the users. It could also accept replicas and refuse to answer any request directed to them, causing a “denial of service”. More subtle attacks could be, for example, to allocate less resources to replicas than was negotiated.

We propose to solve this problem by using a peer-to-peer trust model. Each server associates a recommendation to a number of nodes, representing its trust it has that these sites are honest. Recommendations can be either positive, meaning that the recommended site is probably honest, or negative, meaning that the site is probably evil. Recommendations are made public thanks to the information propagation technique presented above. Based on the full set of recommendations, each node autonomously decides whether it trusts or not each of the cooperative servers.

Although this scheme cannot provide full certitude about a node being good or bad, it generally allows to discriminate good nodes from bad ones. It also tolerates a relatively high number of wrong recommenda-

tions (good recommendations for a bad node or *vice versa*).

Discussing the full details of our trust model is beyond the scope of this paper.

## 3 System Architecture

Figure 1 shows Globule’s architecture. It provides several distinct features: negotiating with remote servers for resource allocation and resource management, document replication and consistency, and automatic client redirection.

### 3.1 Delegated Resource Management

When a server notices that replicating its documents would improve the quality of service for its clients, it identifies one or more locations where replicas would be needed. Those locations are determined, for example, as the Autonomous Systems from which most of the requests originate.<sup>1</sup> It then locates suitable hosting servers in those locations and negotiates resource allocation, as was discussed in Section 2.2.

When a hosted server has acquired appropriate resources on a hosting server, it can use them to create replicas, and redirect clients to them. Its only obligation is to not exceed the amount of resources that have been allocated. This rule is enforced by the hosting server: to prevent excessive use of storage space, a replacement module similar to those of caches is associated to the set of replicas from this hosted server. If a hosted server tries to allocate more resources than allowed, the hosting server will automatically delete other replicas from the same server.

A hosting server, having allocated resources for several different hosted servers, manages each resource pool separately. Therefore, it must have several replacement module instances, each of which enforces resource limitations to one of the hosted servers. This mechanism is similar to those of partitioned caches [8].

Even though resource limitation is enforced by the hosting server, the hosted server remains in control of its allocated resources. It does so by attaching priority flags to its documents, indicating how important each replica is. When making replacement decisions, the hosting server will take these flags into account in addition to standard parameters such as the frequency of requests and replica size. The more important a replica, the less likely it is to be removed.

Similar mechanisms can be setup to enforce limitations in resources such as bandwidth usage [7]. Basically, one can associate bandwidth usage limitations to the replicas from a given hosted server. When the

---

<sup>1</sup>Usually, a small number of origin Autonomous Systems account for a large fraction of the requests, as shown for example in [9].

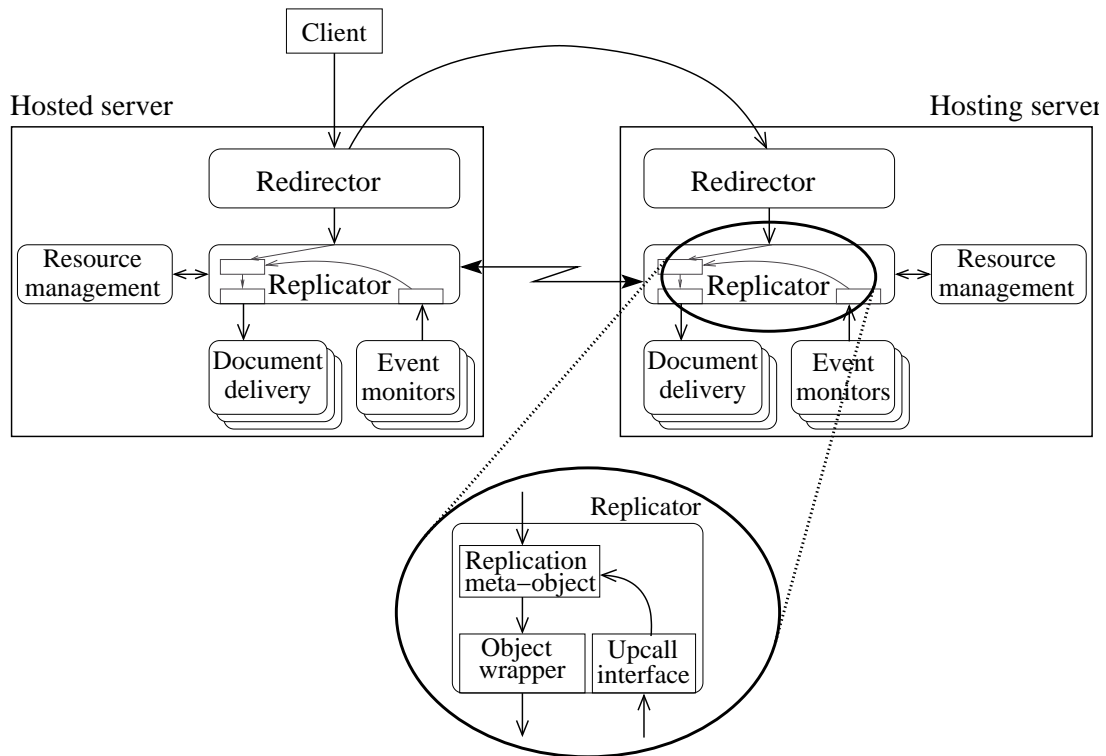


Figure 1: General Architecture

request stream to this set of documents exceeds the authorized bandwidth, the transmission of responses is slowed down so that they share the authorized bandwidth. This algorithm is very similar to reserved bandwidth in routers.

This bandwidth management technique seems adequate to smooth occasional load peaks. However, slowing down connections is acceptable only as long as it remains exceptional. If a server must slow down connections on a regular basis, it will trigger a re-evaluation of the placement and replication policy, probably leading to creating more replicas in its neighborhood (see Section 2.1.1).

### 3.2 Document Replication

As is also shown in Figure 1, a replica is made of two separate local objects: a document's state which is available in the form of delivery components capable of producing documents, and a replication meta-object which is responsible for enforcing the document's replication policy. Each object can either reside in memory or be marshaled to disk. State transitions between memory and disk are dictated by the per-hosted-server resource limitation module described in Section 3.1.

All replication meta-objects implement a standard interface, but they can have various implementations depending on the replication policy they represent.

They maintain information about the object's consistency, such as the date of last modification and the date of the last consistency check.

Each time a request is issued to a document, the server transmits the characteristics of the request to the replication meta-object. Based on its implementation, the policy object responds by indicating how to treat the request: reply immediately based on the local replica, or require to first check for freshness, etc. The server is in charge of actually performing the operation.

Once the replication meta-object has authorized for it, the Web server uses one of its standard document delivery modules to answer the request. These can be modules that deliver static documents, or modules that generate a document on request.

Certain replication policies require taking actions at other times than request time, such as periodically checking for a document's freshness, sending invalidations to replicas when the master copy is updated, and processing incoming invalidations. To do so, meta-objects can register to local services for being invoked when certain events take place. For example, a meta-object can request to be woken up periodically or when a given file is updated.

### 3.3 Client Redirection

Each document is assumed to have a home server. This server is responsible for automatically redirecting clients to their most suitable replica. Knowing the location of clients and replicas, such a selection can be reduced to a shortest path problem [11]. Each incoming client is associated with its Autonomous System. Knowing the locations of replica sites and the map of network connexions between autonomous systems, each client is directed to a server such that the number of autonomous systems on the path between the client and the replica server is minimized. If necessary, more criteria can be used in the selection of a replica site, such as the load of each replica site and the network distances within a given AS.

Two mechanisms can be used to effectively redirect clients to replicas:

- HTTP redirection: when receiving an HTTP request, the server sends a redirection response, indicating from which URL the document should be retrieved. This scheme is very simple, but it is not transparent. That is, browsers display the URL of the mirror site instead of the home site. This may become a problem if, for example, a user bookmarks a mirror page. Later on, when he tries to access the page again, this mirror may have been removed.
- DNS redirection: before accessing a page, a browser needs to request the DNS to convert the server's name into an IP address. After locating an authoritative server for the given domain, the client's DNS server contacts it for actually resolving the name. DNS redirection requires the authoritative server to send customized responses depending on the location of the client [13]. Small TTL values are associated to responses, so that client DNS caches are updated often. A customized authoritative DNS server is necessary, but no other DNS server needs to be modified. This method is fully transparent to the user, since URLs do not need to be modified. On the other hand, it has a coarse granularity: it is not possible to replicate only part of a server, since all requests to this server will be sent to the mirrors.

We decided to use DNS redirection, as did most of the content distribution networks, such as Akamai [6]. However, our system's architecture does not depend on this decision; we may later decide to use another mechanism, for example when the HTTP standard will feature more powerful redirection mechanisms. As we discuss in next section, we will implement a customized DNS server *inside* the Globule Apache module. This allows for easier installation as well as

for tight coupling between the replication and redirection modules.

## 4 Integration into the Apache Web Server

In order to allow for easy deployment, we decided to develop Globule as an Apache module. This way, turning a normal Apache server into a replication-aware server would require only compiling a new module into the server and editing a configuration file.

The Apache Web server is built from a modular design, which enables one to easily add new features. It decomposes the treatment for each request into several steps, such as access checking, actually sending a response back to the client, and logging the request.

Modules can register handler functions to participate in one or more of these steps. When a request is received, the server runs the registered handlers for each step. Modules can then accept or refuse to process the operation; the server tries all the handlers registered for each step until one accepts to process it.

Many third party modules have been developed to extend Apache in a number of ways. Among the most popular ones are server scripting languages such as PHP.

The architecture of Apache provides us all the tools necessary to implement Globule: we can, for example intercept a request before being served by the standard document delivery modules to let the replication meta-objects check for consistency. Servers can communicate by sending HTTP requests to each other.

Although Apache is a server for the HTTP protocol, one can write modules to implement other protocols. So, for example, one could turn Apache into a DNS server. We plan to use this feature for redirecting clients to mirrors.

## 5 Related Work

Many systems have been developed to cache or replicate Web documents. The first server-controlled systems have been push-caches, where the server was responsible of pushing cached copies close to the users [5]. More recently, content distribution networks (CDNs) have been developed along the same idea [3, 6]. These systems rely on a large set of servers deployed around the world. These servers are normal caches, configured as surrogate proxies. The intelligence of the system is mainly concentrated in the DNS servers which are used to direct clients to a server close to them. Consistency is realized by incorporating a hash value of document's content *inside* its URL. When a replicated document is modified, its URL is modified as well. This scheme necessitates to change hyperlink references to modified documents as well. In order to deliver only up-to-date

documents to users, this system cannot use the same mechanism to replicate HTML documents; only embedded objects such as images and videos are replicated. Globule presents three major differences with CDNs. First, since its consistency management is independent from the document naming scheme, it allows to replicate all types of objects. Second, consistency policies are selected on a per-document basis so that each document uses the policy that suits it best. Finally, the system does not require one single entity to deploy a large number of machines across the Internet: Globule users automatically trade resources with each other, therefore incrementally building a worldwide network of servers at low cost.

Globule can be seen as a peer-to-peer system. Most of these systems, such as Gnutella, are used to locate files at user's locations without the need for a centralized server [4]. Other peer-to-peer applications allow for distribution of parallel computations on users' machines [12]. Globule qualifies as a peer-to-peer system because it makes nodes from different administrative domains cooperate to provide a service that no single server could ever achieve. Moreover, its architecture is entirely decentralized, which means that no central server is necessary. The main difference between Globule and other peer-to-peer systems relies in the nature of the inter-node cooperation. Resources being exchanged in Gnutella are files whereas Globule servers exchange storage space and network bandwidth.

## 6 Conclusion

We have presented Globule, a platform for Web document replication. Globule integrates all necessary services into a single tool: dynamic creation and removal of replicas, consistency management, and automatic client redirection. Globule will be implemented as a module for the Apache server.

The architecture presented in this article is still work in progress, but we hope to release a first prototype soon. Two problems have been left for future work. First, security: the administrator of a server would like to make sure that remote servers which accepted to host his replicas will do it without modifying documents, for example. We plan to use a trust model to solve this. Second, details of server-to-server negotiation still have to be figured out.

When the Globule project is completed, we expect to provide a free cooperative platform for Web document replication that will match the ever-increasing quality of service expectations that users have.

## References

- [1] *Apache Web server*, <http://www.apache.org/httpd/>.
- [2] Rajkumar Buyya, David Abramson, and Jonathan Giddy, *An economy driven resource management architecture for global computational power grids*, Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000), June 2000, <http://www.csse.monash.edu.au/rajkumar/papers/GridEconomy.pdf>.
- [3] *Digital Island*, <http://www.digitalisland.net/>.
- [4] *Gnutella*, <http://gnutella.wego.com/>.
- [5] James Gwertzman and Margo Seltzer, *The case for geographical push-caching*, Proceedings of the HotOS '95 Workshop, May 1995, <http://www.eecs.harvard.edu/~vino/web/hotos.ps>.
- [6] F. Thomson Leighton and Daniel M. Lewin, *Global hosting system*, United States Patent, Number US6108703, August 2000.
- [7] *mod\_bandwidth – bandwidth management module for Apache webserver*, [http://www.cohprog.com/mod\\_bandwidth.html](http://www.cohprog.com/mod_bandwidth.html).
- [8] Cristina Duarte Murta, Virgílio Almeida, and Wagner Meira, Jr, *Analyzing performance of partitioned caches for the WWW*, Proceedings of the 3rd International WWW Caching Workshop, June 1998.
- [9] Guillaume Pierre, Ihor Kuz, Maarten van Steen, and Andrew S. Tanenbaum, *Differentiated strategies for replicating Web documents*, *Computer Communications* **24** (2001), no. 2, 232–240.
- [10] Guillaume Pierre, Maarten van Steen, and Andrew S. Tanenbaum, *Self-replicating Web documents*, Technical Report IR-486, Vrije Universiteit, Amsterdam, February 2001, [http://www.cs.vu.nl/~gpierre/publi/SRWD\\_ir486.php3](http://www.cs.vu.nl/~gpierre/publi/SRWD_ir486.php3).
- [11] Lili Qiu, Venkata Padmanabhan, and Geoffrey Voelker, *On the placement of Web server replicas*, Proceedings of the IEEE Infocom conference, April 2001, <http://infocom.ucsd.edu/papers/839-4067851989.ps>.
- [12] *SETI@home*, <http://setiathome.ssl.berkeley.edu/>.
- [13] Wenting Tang, Fan Du, Matt W. Mutka, Lionel M. Ni, and Abdol-Hossein Esfahanian, *Supporting global replicated services by a*

*routing-metric-aware DNS*, Proceedings of the 2nd International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems (WECWIS 2000), June 2000, pp. 67–74.

- [14] Maarten van Steen, Philip Homburg, and Andrew S. Tanenbaum, *Globe: A wide-area distributed system*, IEEE Concurrency **7** (1999), no. 1, 70–78.
- [15] Rich Wolski, James S. Plank, John Brevik, and Todd Bryan, *Analyzing market-based resource allocation strategies for the computational grid*, Tech. Report CS-00-453, University of Tennessee, Knoxville, 2000, <http://www.cs.utk.edu/~rich/publications/CS-00-453.ps.gz>.