

Introducing...

Web browsing is too slow.

poor performing servers, network latencies, bottlenecks,
sudden surges (slashdotting).

Proxies are of limited value.

Content Delivery Networks (CDNs) offer more effective solution.

replicate documents, control document consistency, analyze global access pattern,
push replication, browsers need no-proxy setup, etc.

However, most CDNs are enterprise based, requiring control over all servers.

Alternative: end-users construct their own CDN using Globule, where all its users' machines (incl. desktop PCs) form the CDN.

- provide (cheap) local resources;
- gain access to (valuable) remote resources;
- content owners remain in control of their documents.

Globule is a module for Apache.

- easy to upgrade from non-replicated server;
- current version 1.2.0 available from www.globule.org for Linux, Unix and Windows*.

Advantages of Globule replication:

- accessing nearby replicas increases client perceived performance;
- multiple replica servers help overloaded sites (scalability);
- fault tolerance to keep site always available;
- adaptive replication policies for optimal usage.

(*) Development of the Windows version is sponsored by NLnet.

Adaptive replication

Documents have different access characteristics.



- size, read and update frequency, spatial and temporal access.
- no one-size-fits-all approach to replication

Globule does not apply the same policy to all documents:

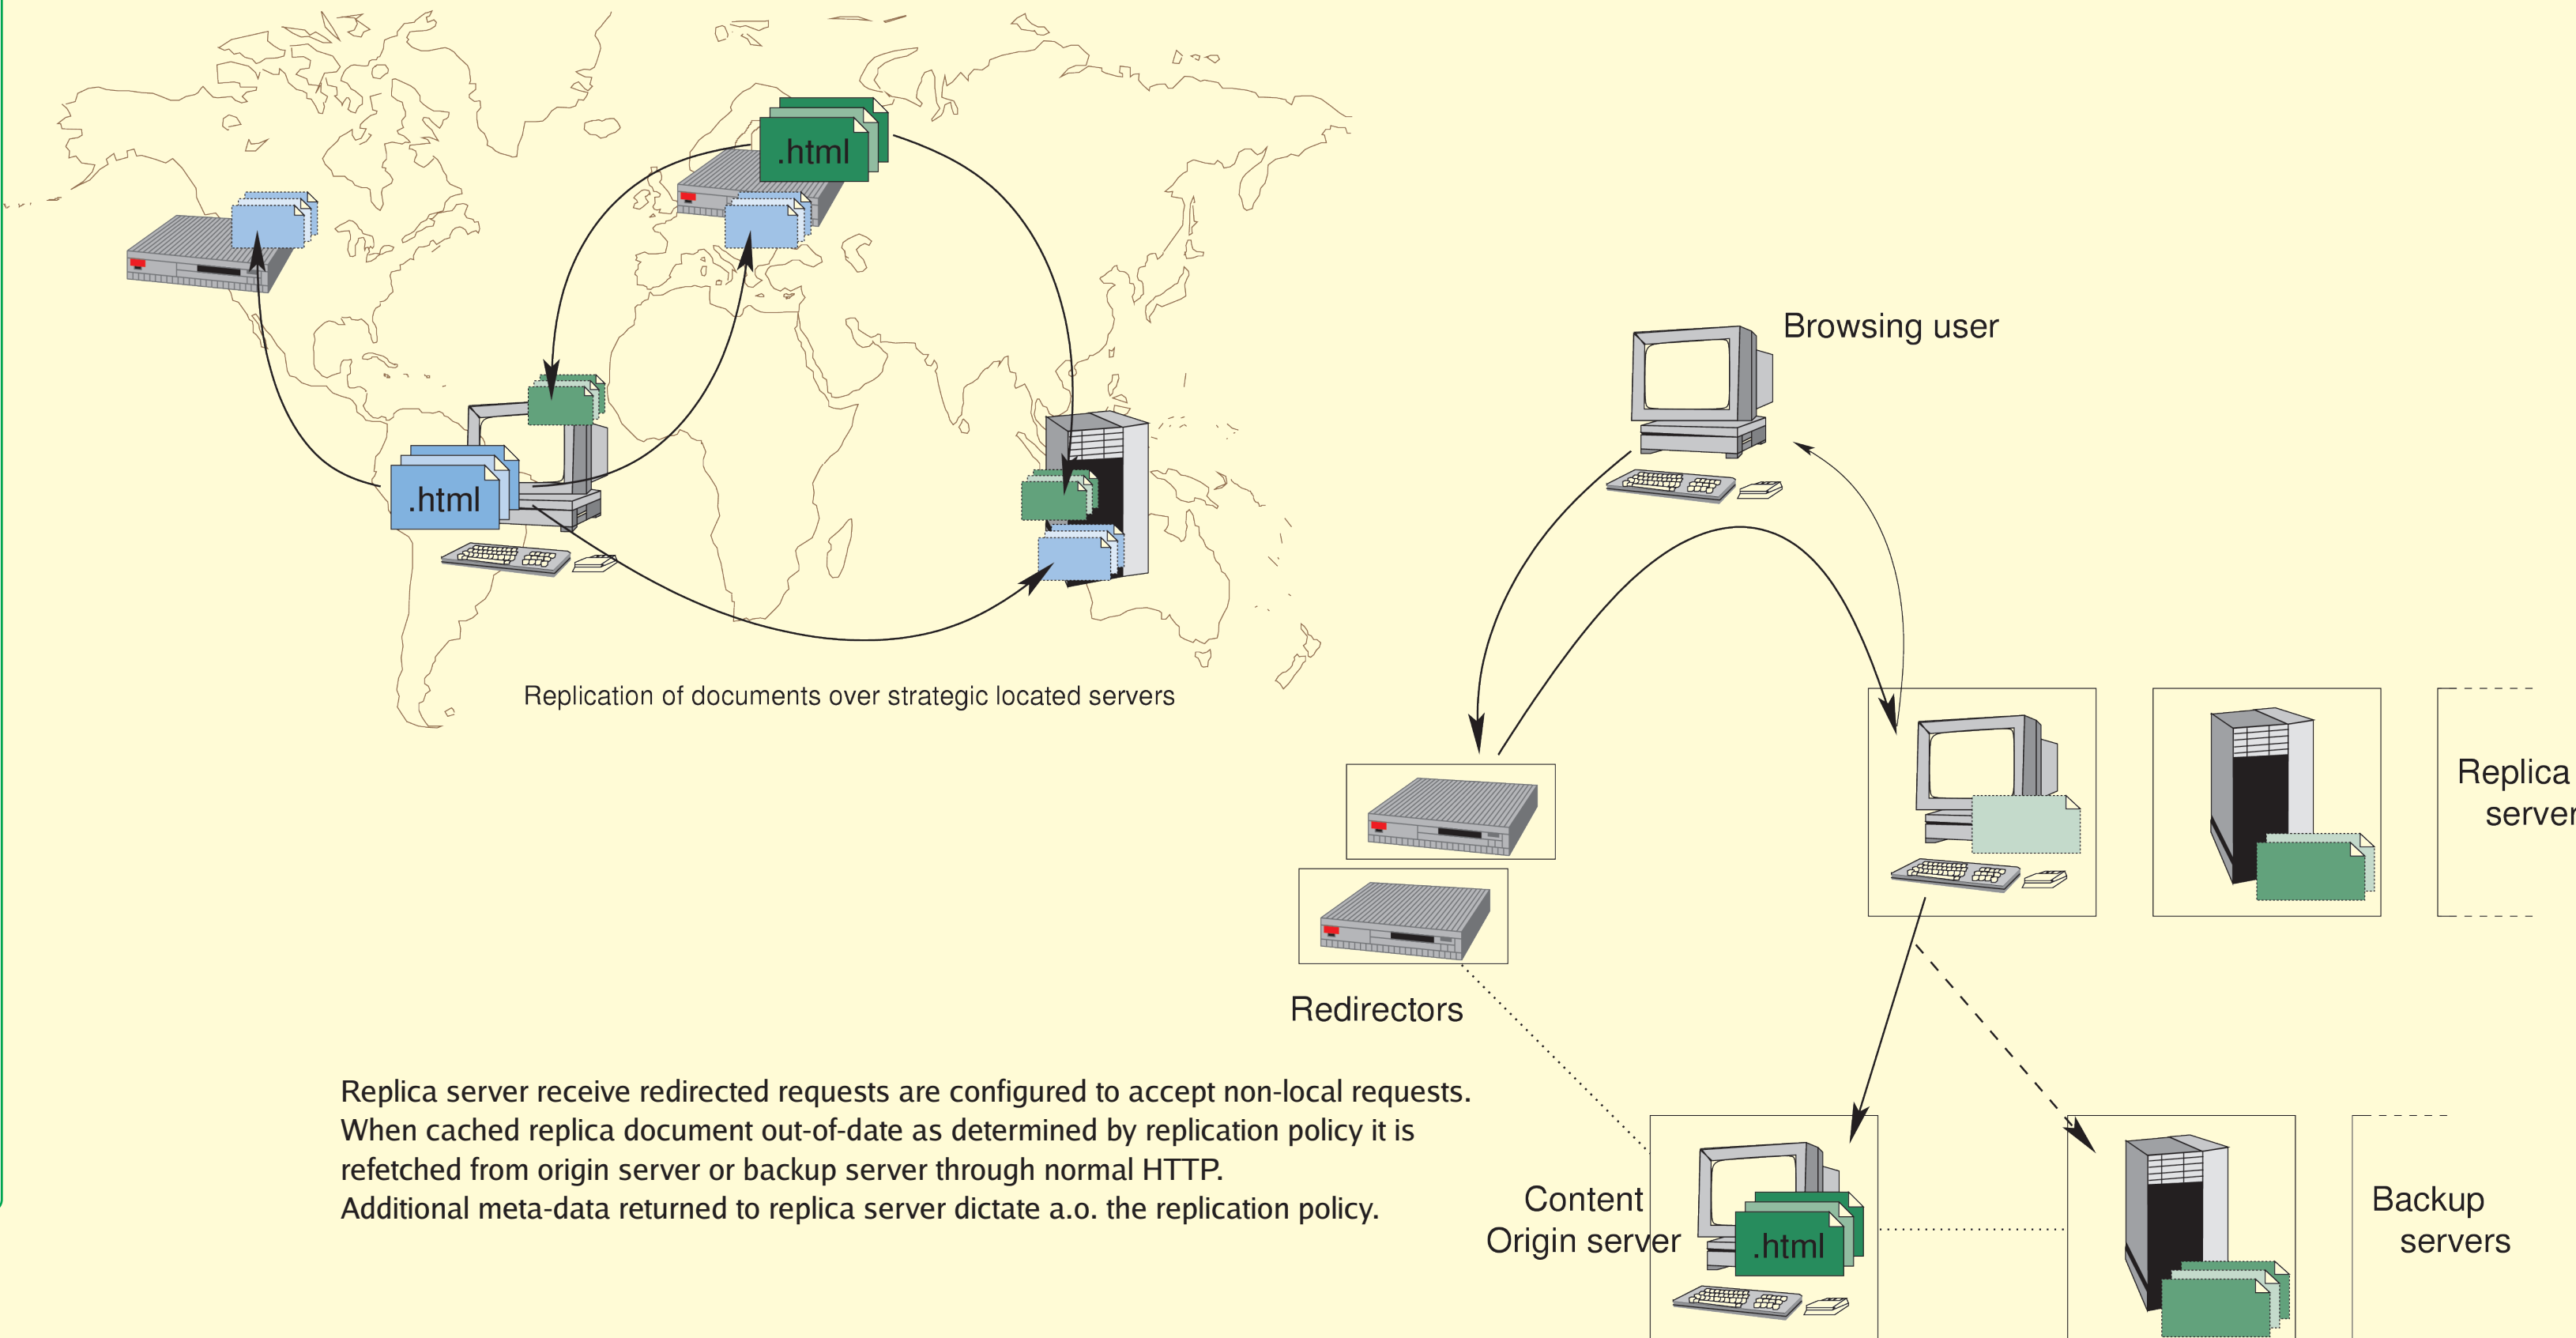
- Each document is replicated with the policy that suits it best;
- Choice of policy is done automatically;
- Globule dynamically changes policies if access patterns change.

Evaluated by computing best performance of each possible strategy over past period, depending on several metrics (e.g. network latency, consistency tolerance, update traffic)

Server of document owner dictates selection of replicas, consistency strategies.



Guillaume Pierre, Berry van Halderen, Swaminathan Sivasubramanian
Vrije Universiteit, Amsterdam



Replica server receive redirected requests are configured to accept non-local requests. When cached replica document out-of-date as determined by replication policy it is refetched from origin server or backup server through normal HTTP. Additional meta-data returned to replica server dictate a.o. the replication policy.



www.globule.org

Browser redirection

Web clients use standard unmodified browsers, without special proxy settings.

Globule supports two *redirection mechanisms*:

- HTTP redirection, initial GET requests returns a HTTP_MOVED_TEMPORARILY. All future requests of that client are then redirected to a replica server.
- DNS redirection, redirecting clients to different IP address when looking up host in URL.

Redirection policies determin where to redirect a client to.

Varying policies from simple to complex:

- round-robin;
- proximity based using AS-path lengths (hops);
- based on latency distance estimation, see below (planned).

Disconnected operation

-- *fault tolerance*

Globule is attractive also because of its fault tolerance aspects.

What happens if one server is unavailable?

tolerate host failure, maintenance, network problems, overloaded network or servers.

Content owners monitor their replica servers.

- Redirect clients only to available servers

What if the content owner becomes unavailable?

- Backup servers hold master content as secondary storage for replicas if point of origin in down.
- Globule allows for out-sourcing of responsibility of redirection to multiple third-party servers.

Out-source tasks to hosts which provide a higher availability service.


Note that tasks are out-sourced, but the content owner remains in control by dictating different policies.

Latency estimation

Scalable Cooperative Latency Estimation (SCoLE).

obtain latency to other nodes without need to always measure.

- Each client machine is associated to an N-dimensional coordinate.
- Latency between two machines is estimated as the distance between the coordinates.



- select N+1 "landmark" nodes;
- measure latencies between landmarks;
- set coordinates such that $\text{latency}(x,y) \approx \text{distance}(\text{coord}(x), \text{coord}(y))$;
- other nodes measure N+1 latencies to landmarks and triangulate distance determining their coordinates;
- latency between any two nodes is the distance in N-dimensional space.

Dynamic document replication

Planned

Most large web sites use dynamic document technology (e.g. PHP, CGI, ASP) to generate web pages. Stored data used by these scripts is being accessed from local databases.

Replication of dynamic documents requires replication of both code and data.

In principle straight forward, however:

- code modifies underlying data thus the system must maintain consistency among replicas;
- scripts should transparently run on different hosts than from which they originate;
- replica servers should be protected from replicated scripts trying to access local state.

Globule will support incorporating on-demand replication of PHP scripts and application data. Replicate partial data only on servers that access them, data is kept consistent at low cost.

Replica Placement

Planned

How do you select the right servers to host your content?

1. locate where your clients are.
2. identify hot regions where most clients reside.
3. pick a server located inside the region.

Replicas are strategically placed nearby clients.

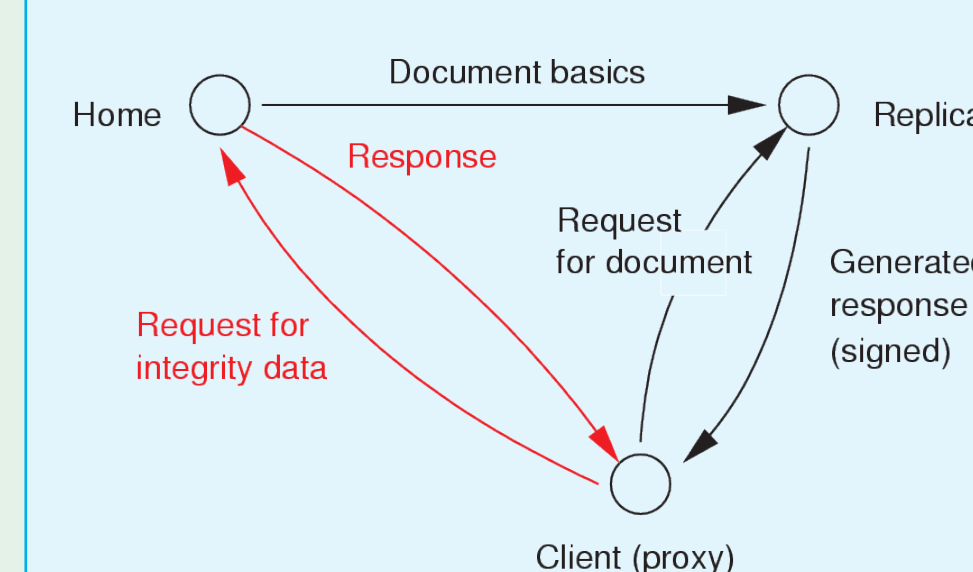
- Minimizing client-to-replica latency.
- Providing region relevant data distribution.

Internet latencies are modeled as an N-dimensional space using our SCoLE research (see browser redirection).

Security

Unreleased

How do you make sure that replica servers deliver your content? Need mechanisms that verifies the integrity from responses of untrusted replica servers.



Place replicas only at servers that you trust.

Arbitrary requests are checked by comparing hash value of content between replica and document owning server. Violation results in removal of the list from active replicas.

