DotSlash - Creating Content Distribution Networks on Demand

Henning Schulzrinne and Weibin Zhao

Department of Computer Science
Columbia University
New York, NY 10027
{hgs,zwb}@cs.columbia.edu

Traditional content distribution networks, such as Akamai, are well-suited for static web services that routinely experience large traffic volumes. They are unsuited for active content, i.e., content generated by scripts from databases, and web sites that are unlikely to receive significant number of requests. However, a few such sites will invariably experience their "fifteen minutes of fame", typically by being mentioned on a high-volume news site such as SlashDot or CNN. Such flash crowds or "Slashdot effect" will routinely cause single-server web sites to collapse.

We have designed an autonomic web replication system, called DotSlash, that allows a web site to expand its capacity on demand. DotSlash allows an origin web server to draft and release rescue servers fully automatically based on its load conditions, without user intervention. An origin web server can discover suitable rescue servers via wide-area service location, either among peer servers or from a dedicated pool of rescue servers, allocate them for temporary use and redirect client requests to them. DotSlash also allows a rescue server to serve the content of its origin web server on the fly without the need of any advance configuration. Our system is completely transparent to clients and does not require URL rewriting or other client modifications.

We have implemented DotSlash as two versions. The first version deals only with static content, e.g., HTML pages or media objects, where a rescue server serves as a reverse caching proxy for its origin web server. The second version also supports dynamic replication of scripts, where a rescue server retrieves the scripts from its origin web server on demand, caches the scripts locally, and accesses the corresponding database server directly. We have prototyped our system for the common LAMP (Linux, Apache, MySQL and PHP) configuration, and shown that a common benchmark for bulletin boards can be replicated without code changes, yielding capacity increases bounded only by the database server. Since many such systems, including most blogs, are bottlenecked by the web server, our system can significantly increase their capacity, and can work well even for extremely rapid load increases.

We are currently investigating how such systems can be further extended by increasing the database capacity of read-mostly systems with loose consistency constraints.

Keywords: CDN, web replication, flash crowd, Slashdot effect, autonomic system, LAMP, active content