

Work in Progress

A Case for Peering of Content Delivery Networks

Rajkumar Buyya and Al-Mukaddim Khan Pathan • University of Melbourne
James Broberg and Zahir Tari • RMIT University

Service providers often geographically distribute their Web servers' facilities to improve performance, reliability, and scalability. Content Delivery Networks,¹ which first evolved in 1998, replicate content over several mirrored Web servers, strategically placed at various locations to deal with flash crowds and to enhance response time.² A CDN improves network performance by maximizing bandwidth, improving accessibility, and maintaining correctness through content replication.

Unfortunately, although many commercial CDN providers exist, they don't cooperate in delivering content to end users in a scalable manner. In addition, content providers typically subscribe to one CDN provider and thus can't use multiple CDNs at the same time. Such a closed, noncooperative model results in "islands" of CDNs.

Here, we present a model for an open, scalable, and service-oriented architecture (SOA)-based system. This system helps to create open Content and Service Delivery Networks (CSDNs) that scale well and can share resources with other CSDNs through cooperation and coordination, thus overcoming the island CDN problem.

Motivation for peering CDNs

A CDN combines a content-delivery infrastructure, a request-routing infrastructure, a distribution infrastructure, and an accounting infrastructure. Figure 1 shows a typical content-delivery environment, with replicated Web-server clusters located at the edge of the network to which the end users are connected. In such a CDN environment, the origin server fetches the Web content based on user requests and then serves the content to the user from the nearby replicated server. That way, CDNs can offer fast, reliable applications and services.³

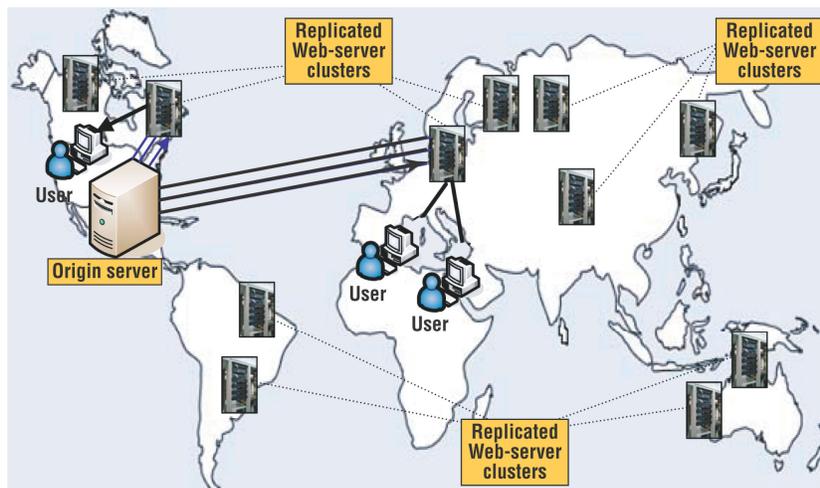


Figure 1. An abstract architecture of a Content Delivery Network.

However, because CDNs are proprietary—individual companies own and operate them—they comprise their own closed delivery networks, which are expensive to set up and maintain. They also have limited scalability. Running a global CDN is even more costly, requiring an enormous amount of capital and labor. Furthermore, commercial CDNs make specific commitments to their end users by signing a Service Level Agreement, which outlines specific penalties if they fail to meet those commitments. So, if a particular CDN fails to provide the client with quality service, it might violate the SLA and end up costing the provider. To cut expenses, and to avoid adverse business impact through SLA violations, CDN providers should partner together.

The CDN provider can accomplish this by establishing peering with other CDN providers that have caching servers located near its client. That way, CDN providers can share each others resources to serve their clients' requests effectively and meet their quality-of-service (QoS) requirements. To ensure sustained resource sharing between CDN providers, peering arrangements must ensure sufficient incentives exist.

To better understand CDN internetworking, consider the scenario in figure 2. Suppose that the ICC Cricket World Cup 2007 is being held in the Caribbean, and www.cricinfo.com is supposed to provide live media coverage. As a content provider, www.cricinfo.com has an exclusive SLA with the CDN provider, Akamai. However, Akamai doesn't have a point of presence (POP) in Trinidad and Tobago (a Caribbean island), where most of the cricket matches will be held. Akamai management might decide to place its surrogates in Trinidad and Tobago, or they might use their edge servers on other Caribbean islands (such as St. Lucia).

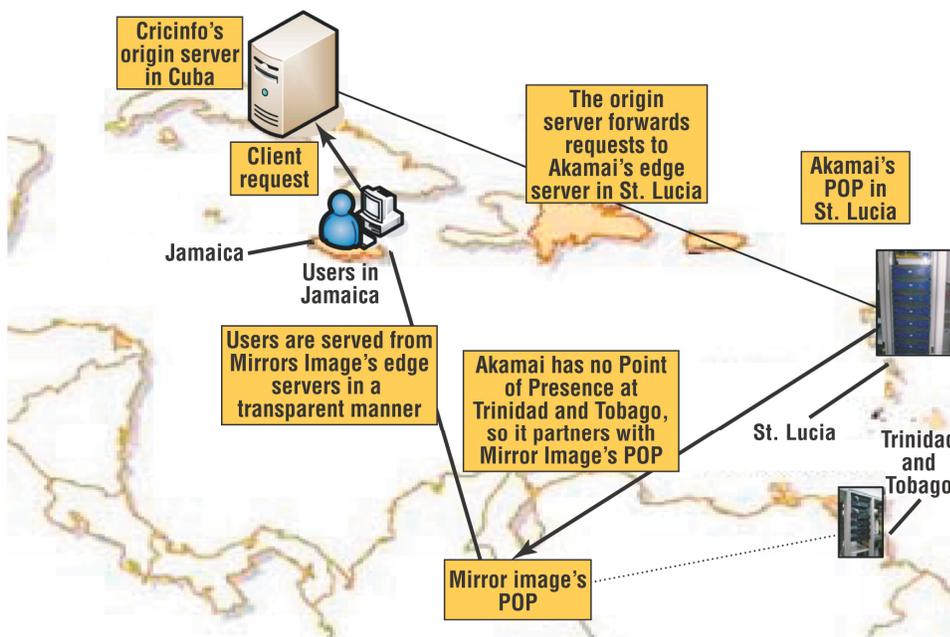


Figure 2. A CDN internetworking scenario.

In the first case, placing new surrogates just for one particular event would cost Akamai too much money and might not be useful after the event. On the other hand, Akamai risks its reputation if it can't provide quality service according to client requests, which could violate the SLA and still cost the company money. If another CDN provider, Mirror-Image, has a POP in Trinidad and Tobago, Akamai could partner with Mirror-Image's edge servers. In other words, by collaborating with another CDN provider, content networks can better satisfy their customers and meet their QoS requirements.

To make content services an Internet infrastructure service, vendors have implemented content service networks (CSNs),⁴ which act as another network infrastructure layer built upon CDNs and provide next-generation CDN services. So, a CSN is another variation of the conventional CDN. This logical separation between content and services under the “content delivery and distribution” and “content services” domains is undesirable considering the ongoing trend in content networking. A unified content network that supports the coordinated composition and delivery of content and services would be much better.

Our model

Our proposed system ensures the quality of services based on SLA negotiation and solves the problem of the logical separation between CDNs and CSNs. We propose a Virtual Organization (VO) model for forming CSDNs that share Web servers not only within their own networks but also with other CSDNs. To encourage sustained resource sharing and peering arrangements between different CDN providers at a global level, we propose using market-based models in resource allocation and management inspired from their successful utilization in the management of autonomous resources, especially in global Grids.⁵ Figure 3 shows such a system’s architecture, and the following sections describe its main elements.

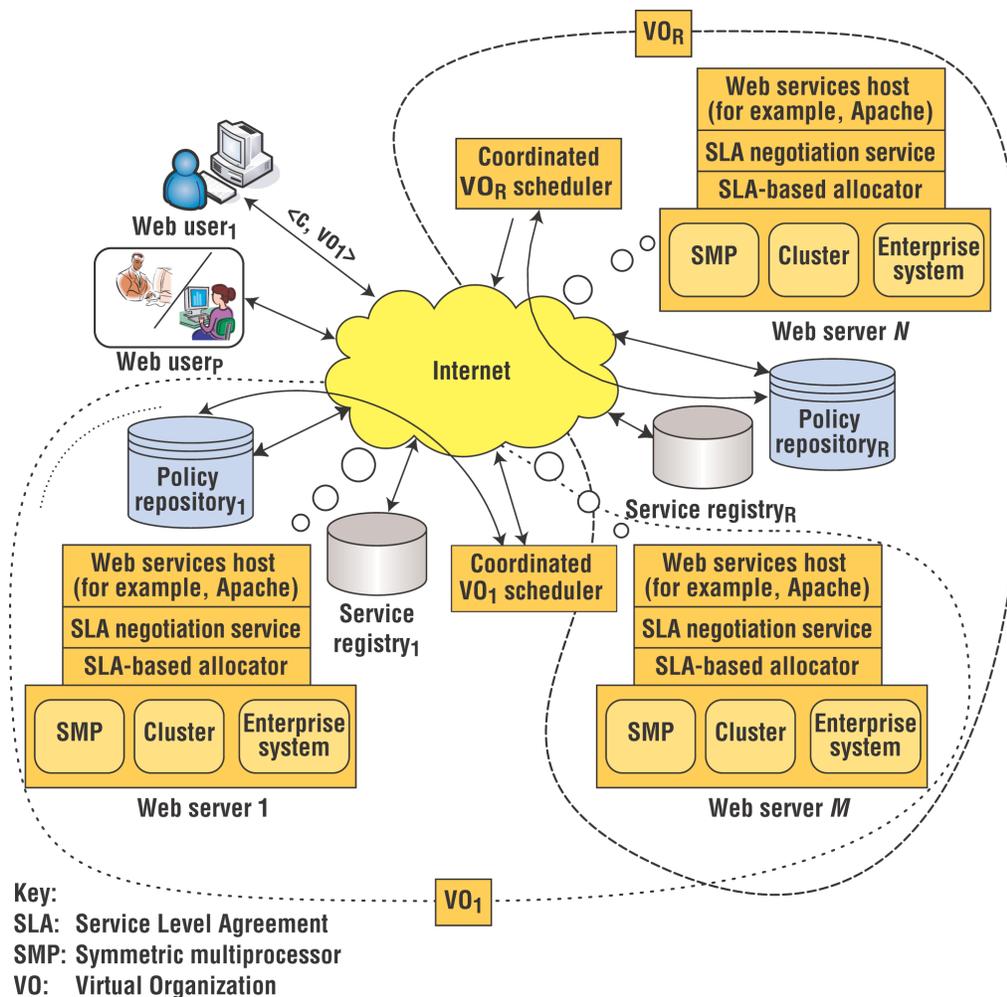


Figure 3. The architecture of open, scalable, service-oriented-architecture-based system to assist the creation of cooperative and coordinated Content and Service Delivery Networks.

Web server

The Web server is a CSDN's most important element. Servers are responsible for storing content and value-added services as infrastructure services and delivering them in a reliable and cooperative manner. Servers within each CSDN can deliver content and services to meet end-user QoS requirements.

We can divide a server's structure into two layers: an *overlay layer* and the *core*. In the overlay layer, a server comprises a Web-service host (for example, Apache or Tomcat), an SLA-negotiation service module, and an SLA-based allocator. The negotiation module, with the help of a coordinated VO scheduler, is responsible for cooperating and coordinating with other servers (located in a local or global CSDN) through SLA-based negotiation. The SLA-based allocator delivers content and services based on the negotiated SLAs with other local or global CSDNs' servers.

The Web server's core consists of high performance computing systems such as symmetric multiprocessors, cluster systems, or other enterprise systems (such as desktop grids). The server's underlying devices and tools must store the content and services and assist in responding quickly and reliably to client requests to meet the negotiated QoS requirements. For content and service location and routing, the Web servers' underlying technologies perform on-demand cooperative caching through coordination with other servers. Efficiently balancing the load across different Web servers is critical to produce the required QoS. Hence, servers are adopted with appropriate load- and resource-distribution strategies.

Coordinated VO scheduler

A coordinated VO scheduler is put in each VO and is responsible for ensuring collaboration and coordination with other CSDNs through policy exchange and scheduling of content and services.

Service registry

A service registry enables CDN providers to register and publish their resources and service details. An SLA-negotiator service and allocator module uses this service registry to discover CDN providers and negotiate QoS parameters and resource allocation to maximize cooperative CSDNs' potential.

Policy repository

A policy repository stores the policies that the administrators generate. These policies are a set of rules to administer, manage, and control access to VO resources. They provide a way to consistently manage the components deploying complex technologies.

Realizing our VO model for forming CSDNs and the policy framework within the VO should be a timely contribution to the ongoing content-networking trend. Our work on peering CDNs is a joint collaboration between the Grids laboratory at the University of Melbourne, and the DSN Laboratory at RMIT University, Australia. For more information, please visit the project Web site at www.gridbus.org/cdn.

References

1. G. Peng, *CDN: Content Distribution Network*, tech. report TR-125, Experimental Computer Systems Lab, Dept. of Computer Science, State Univ. of New York, 2003; <http://citeseer.ist.psu.edu/peng03cdn.html>.
2. M. Arlitt and T. Jin, "A Workload Characterization Study of 1998 World Cup Web Site," *IEEE Network*, May/June 2000, pp. 30–37.
3. G. Pallis and A. Vakali, "Insight and Perspectives for Content Delivery Networks," *Comm. ACM*, vol. 49, no. 1, ACM Press, 2006, pp. 101–106.

4. W.Y. Ma, B. Shen, and J.T. Brassil, "Content Services Network: Architecture and Protocols," *Proc. 6th Int'l Workshop on Web Caching and Content Distribution (IWCW 01)*, Boston Univ., 2001, pp. 89–107.
5. R. Buyya, D. Abramson, and S. Venugopal, "The Grid Economy," *Proc. IEEE*, vol. 93, no. 3, 2005, pp. 698–714; www.gridbus.org/%7Eraj/papers/ieee-grideconomy.pdf (pdf).

Rajkumar Buyya is Director of the Grid Computing and Distributed Systems Laboratory at the University of Melbourne. He also serves as Chair of the IEEE Technical Committee for Scalable Computing. For details on his research work, see www.gridbus.org. Contact him at raj@csse.unimelb.edu.au.

Al-Mukaddim Khan Pathan is a masters student at the University of Melbourne. Contact him at apathan@csse.unimelb.edu.au.

James Broberg is a PhD student at RMIT University. Contact him at jbroberg@cs.rmit.edu.au.

Zahir Tari is a professor and the head of the Distributed Systems and Networking Discipline at RMIT University. Contact him at zahirt@cs.rmit.edu.au.

Related Links

- "Content Delivery Networks: Status and Trends," *IEEE Internet Computing*
- "Globally Distributed Content Delivery," *IEEE Internet Computing*

Cite this article:

Rajkumar Buyya, Al-Mukaddim Khan Pathan, James Broberg, and Zahir Tari, "A Case for Peering of Content Delivery Networks," *IEEE Distributed Systems Online*, vol. 7, no. 10, 2006, art. no. 0610-o10003.