

Chapter 7

Optimizing Content Delivery in QoS-Aware Multi-CDNs

Nadia Ranaldo
University of Sannio, Italy

Eugenio Zimeo
University of Sannio, Italy

ABSTRACT

Broadband network technologies have improved the bandwidth of the edge of the Internet, but its core is still a bottleneck for large file transfers. Content Delivery Networks (CDNs), built at the edge of the Internet, are able to reduce the workload of network backbones, but their scalability and network reach is often limited, especially in case of QoS-bound delivery services. By using the emerging CDN inter-networking, a CDN can dynamically exploit resources of other cooperating CDNs to face peak loads and temporary malfunctions without violating QoS levels negotiated with content providers. In this chapter, after a wide discussion of the problem, the authors propose an architectural schema and an algorithm, based on the divisible load theory, which optimizes delivery of large data files by satisfying an SLA, agreed with a content provider, while respecting the maximum budget that the delivering CDN can pay to peer CDNs to ensure its revenue.

INTRODUCTION

The increasing desire of using rich media services is often not satisfied due to bandwidth limitations, high latency and low quality of service. In fact, even if broadband technologies are significantly improving the bandwidth of the edge of the In-

ternet, backbones remain limited and congested because of the huge traffic generated by service providers for replying to user requests. Therefore, smarter infrastructures are necessary to improve quality, bandwidth availability and profitability by using edge technologies, such as Content Delivery Networks (CDNs) (Hofmann & Beaumont, 2005; Pallis & Vakali, 2006).

DOI: 10.4018/978-1-4666-1794-0.ch007

A CDN is a collection of cooperating servers, distributed across multiple locations, optimized to deliver content, such as static Web pages, applications files, stored multimedia objects and real-time audio/video. Content providers, such as large enterprises, Web service providers, media companies and news broadcasters, exploit CDN services to deliver content and applications with high performance and reliability to a wide and geographically distributed audience.

The basic idea of a CDN is to replicate contents on servers (called replica servers) and deliver them to end-users from “the best” server with the aim of minimizing user-perceived response time and of maximizing overall CDN performance with respect to some metrics, such as bandwidth consumption, server utilization and reliability. The redirection of the user’s request to the best replica server is performed using a request routing mechanism, such as the Global Server Load Balancing (GSLB), DNS-based request routing, dynamic metafile generation and HTML rewriting (Pathan & Buyya, 2008a).

From a business point of view, deploying a large-scale CDN for Internet-wide available content is complex and requires high operation costs. It, in fact, involves the availability and management of distributed resources with large storage capacity connected among them by means of high throughput networks. Against sustained costs, CDNs obtain their revenues by service contracts stipulated with content providers, which typically define fees related to the amount of stored and of transferred content. Moreover, Service Level Agreements (SLAs) (Bouman, Trienekens, & van der Zwan, 1999), used to formally define a service contract, agreed between a customer and a service provider, are becoming more and more popular in the context of content delivery services. Commercial CDNs, in fact, are recently starting to adopt SLAs to agree a desired degree of quality of service (QoS) against a penalty to pay to the customer if an SLA is not satisfied; however, at our knowledge, only service availability is taken

into consideration by existing providers (Amazon Web Services, 2010).

CDNs could negotiate further QoS parameters for becoming more attractive to content providers, such as maximum response time for dynamic Web transactions or maximum download time for large multimedia data. However, satisfying SLAs for content delivery services introduces an additional complexity in CDN infrastructures management. This leads to further increase the economic investments against low revenue margins.

An approach that can be exploited to foster QoS-aware CDNs is *CDN internetworking*, also called *content peering* (Day, Cain, Tomlinson, & Rzewski, 2003). This emerging approach, based on interoperability and resource sharing mechanisms, allows a CDN to optimize its content delivery services by dynamically exploiting capabilities of other CDNs participating to the same virtual organization, called *multi-CDN*. Since content peering allows a CDN to expand beyond its borders and to deliver large-scale (eventually Internet-wide) services without huge economical investments, it is gaining popularity for improving scalability and network reach of CDN providers and to tear down the initial investment of new ones.

Content peering allows a CDN to avoid the violation of negotiated QoS levels with content providers by dynamically addressing unforeseen situations caused by malfunctions of CDN infrastructures or peak loads typically generated by global and limited-time interest events, such as worldwide sport competitions and historical significance actuality and political occurrences. As an example, in conjunction with the Soccer World Cup 2010, during which Akamai delivered live and on-demand Web content for 24 global broadcasters in 65 countries for one month, a peak of more than 20 million visitors per minute was registered on June 24th, over twice higher than the World Cup 2006 (Young & Faris, 2010).

In the context of a business multi-CDN, CDNs offer their content delivery services with fees to content providers negotiating SLAs with them. If

23 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:
www.igi-global.com/chapter/optimizing-content-delivery-qos-aware/66997

Related Content

Internet of Things in Disaster Logistics Productivity

Buket Karatop, Bura Takanand Cemalettin Kubat (2020). *Internet of Things (IoT) Applications for Enterprise Productivity* (pp. 91-111).

www.irma-international.org/chapter/internet-of-things-in-disaster-logistics-productivity/250724

A Cyber-Physical System Testbed for Security Experimentation

John Hale, Abraham Habib, Rujit Raval, Ryan Irvinand Peter J. Hawrylak (2020). *Cyber Security of Industrial Control Systems in the Future Internet Environment* (pp. 175-209).

www.irma-international.org/chapter/a-cyber-physical-system-testbed-for-security-experimentation/250111

I-BAT: A Data-Intensive Solution Based on the Internet of Things to Predict Energy Behaviors in Microgrids

Antonio J. Jara, Luc Dufour, Gianluca Rizzo, Marcin Piotr Pawlowski, Dominique Genoud, Alexandre Cotting, Yann Bocchiand Francois Chabbey (2020). *Securing the Internet of Things: Concepts, Methodologies, Tools, and Applications* (pp. 630-654).

www.irma-international.org/chapter/i-bat/234966

On the Performance of Content Delivery Clouds

Mukaddim Pathan, James Brobergand Rajkumar Buyya (2012). *Next Generation Content Delivery Infrastructures: Emerging Paradigms and Technologies* (pp. 29-54).

www.irma-international.org/chapter/performance-content-delivery-clouds/66992

Cultural Heritage and Digitalization in City Branding

Oya Yildirimand A. Celil Çakici (2022). *Handbook of Research on Digital Communications, Internet of Things, and the Future of Cultural Tourism* (pp. 467-483).

www.irma-international.org/chapter/cultural-heritage-and-digitalization-in-city-branding/295518