



US 20070156845A1

(19) **United States**

(12) **Patent Application Publication**  
**Devanneaux et al.**

(10) **Pub. No.: US 2007/0156845 A1**

(43) **Pub. Date: Jul. 5, 2007**

(54) **SITE ACCELERATION WITH CONTENT  
PREFETCHING ENABLED THROUGH  
CUSTOMER-SPECIFIC CONFIGURATIONS**

(52) **U.S. Cl. .... 709/217**

(57) **ABSTRACT**

(75) **Inventors: Thomas P. Devanneaux**, Los Altos, CA  
(US); **Laszlo Kovacs**, Foster City, CA  
(US); **Stephen L. Ludin**, Mill Valley,  
CA (US)

A CDN edge server is configured to provide one or more extended content delivery features on a domain-specific, customer-specific basis, preferably using configuration files that are distributed to the edge servers using a configuration system. A given configuration file includes a set of content handling rules and directives that facilitate one or more advanced content handling features, such as content prefetching. When prefetching is enabled, the edge server retrieves objects embedded in pages (normally HTML content) at the same time it serves the page to the browser rather than waiting for the browser's request for these objects. This can significantly decrease the overall rendering time of the page and improve the user experience of a Web site. Using a set of metadata tags, prefetching can be applied to either cacheable or uncacheable content. When prefetching is used for cacheable content, and the object to be prefetched is already in cache, the object is moved from disk into memory so that it is ready to be served. When prefetching is used for uncacheable content, preferably the retrieved objects are uniquely associated with the client browser request that triggered the prefetch so that these objects cannot be served to a different end user. By applying metadata in the configuration file, prefetching can be combined with tiered distribution and other edge server configuration options to further improve the speed of delivery and/or to protect the origin server from bursts of prefetching requests.

Correspondence Address:  
**LAW OFFICE OF DAVID H. JUDSON**  
**15950 DALLAS PARKWAY**  
**SUITE 225**  
**DALLAS, TX 75248 (US)**

(73) **Assignee: Akamai Technologies, Inc.**

(21) **Appl. No.: 11/647,750**


(22) **Filed: Dec. 29, 2006**

**Related U.S. Application Data**

(60) **Provisional application No. 60/755,176**, filed on Dec. 30, 2005. **Provisional application No. 60/755,908**, filed on Dec. 31, 2005.

**Publication Classification**

(51) **Int. Cl. G06F 15/16 (2006.01)**



AKAMAI





Figure 2

200




Figure 3

Table 1. Supported Tags/Attributes

BASE	href
A	href
IMG	src, lowsrc, longdesc
SCRIPT	src
FRAME	src, longdesc
IFRAME	src, longdesc
LINK	href
OBJECT	codebase+[archive classid data]
AREA	href
APPLET	codebase+[archive src code object]
EMBED	codebase+[archive src code object]
INPUT	src, lowsrc
OFORM	action
BODY	background
TABLE	background
TD	background
TH	background

Figure 4

```
<!-- Global fetch limit settings to protect origin; should be included in all
configurations -->

<edgeservices:prefetch>
  <status>off</status>
  <fetch>
    <max-prefetches-per-page>100</max-prefetches-per-page>
    <max-urls-per-page>100</max-urls-per-page>
    <log-r-lines>on</log-r-lines>
  </fetch>
  <fetch.limits>
    <status>on</status>
    <time-scale>1s</time-scale>
    <requests-high-watermark>20</requests-high-watermark>
    <requests-low-watermark>15</requests-low-watermark>
  </fetch.limits>
</edgeservices:prefetch>

<!-- turn prefetch on for content likely to be HTML -->

<match:uri.ext value="jsp html htm asp">
  <edgeservices:prefetch.status>on</edgeservices:prefetch.status>
</match:uri.ext>

<!-- turn prefetch on for directory defaults requested without extension -->
<match:uri.ext value="">
  <edgeservices:prefetch.status>on</edgeservices:prefetch.status>
</match:uri.ext>

<!-- mark prefetchable content as such -- expand this list of extensions as
needed -->
<match:uri.ext value="css gif jpg jpeg js ico mov png swf txt wav wma xml">
  <edgeservices:prefetch.prefetchable-object>on</edgeservices:prefetch.
prefetchable-object>
  <edgeservices:prefetch.status>off</edgeservices:prefetch.status>
</match:uri.ext>

<!-- prefetch even if the HTML is already cached -->
<edgeservices:prefetch.prefetch-on-hit>on</edgeservices:prefetch.prefetch-on-
hit>

<!-- when Tiered Distribution is used -->
<edgeservices:prefetch.indirect-only>on</edgeservices:prefetch.indirect-only>
```

Figure 5

**SITE ACCELERATION WITH CONTENT  
PREFETCHING ENABLED THROUGH  
CUSTOMER-SPECIFIC CONFIGURATIONS**

[0001] This application claims priority to Ser. No. 60/755,176, filed Dec. 30, 2005, and Ser. No. 60/755,908, filed Dec. 31, 2005.

[0002] Portions of this application contain subject matter that is protected by copyright.

BACKGROUND OF THE INVENTION

[0003] 1. Technical Field

[0004] The present invention relates generally to content delivery in distributed networks.

[0005] 2. Brief Description of the Related Art

[0006] A company's Web site represents its public face. It is often the initial point of contact for obtaining access to the company's information or doing business with the company. Public facing Web sites are used for many purposes. They can be used to transact commerce, where end consumers evaluate and buy products and services, and they are often linked to revenue generation and satisfying customer requests. They can be used as news and information portals for supplying the latest content for consumers. A company's Web site can be used as a customer self-service venue, where customer satisfaction is critical to loyalty in getting customers to return to the Web site. These are merely representative examples, of course. As companies place greater importance on the Internet, Web sites increasingly become a key component of a company's business and its external communications. As such, the capability and flexibility of the supporting Internet infrastructure for the Web site becomes mission-critical. In particular, the infrastructure must provide good performance for all end user consumers, regardless of their location. The site must scale to handle high traffic load during peak usage periods. It must remain available 24x7, regardless of conditions on the Internet. When performance, reliability, or scalability problems do occur, Web site adoption and usage can be negatively impacted, resulting in greater costs, decreased revenue, and customer satisfaction issues.

[0007] It is known in the prior art to off-load Web site content for delivery by a third party distributed computer system. One such distributed computer system is a "content delivery network" or "CDN" that is operated and managed by a service provider. The service provider typically provides the service on behalf of third parties. A "distributed system" of this type typically refers to a collection of autonomous computers linked by a network or networks, together with the software, systems, protocols and techniques designed to facilitate various services, such as content delivery or the support of outsourced site infrastructure. Typically, "content delivery" means the storage, caching, or transmission of content, streaming media and applications on behalf of content providers, including ancillary technologies used therewith including, without limitation, DNS request handling, provisioning, data monitoring and reporting, content targeting, personalization, and business intelligence. The term "outsourced site infrastructure" means the distributed systems and associated technologies that enable

[0008] FIGS. 1-2 illustrate a known CDN infrastructure for managing content delivery on behalf of participating content providers. In this example, computer system 100 is configured as a CDN and is managed by a service provider. The CDN is assumed to have a set of machines 102a-n distributed around the Internet, and some or even all of these machines may be located in data centers owned or operated by third parties. Typically, most of the machines are servers located near the edge of the Internet, i.e., at or adjacent end user access networks. A Network Operations Command Center (NOCC) 104 may be used to administer and manage operations of the various machines in the system. Third party content sites, such as Web site 106, offload delivery of content (e.g., HTML, embedded page objects, streaming media, software downloads, and the like) to the distributed computer system 100 and, in particular, to "edge" servers. Typically, this service is provided for a fee. In one common scenario, CDN content provider customers offload their content delivery by aliasing (e.g., by a DNS canonical name) given content provider domains or sub-domains to domains that are managed by the service provider's authoritative domain name service. End users that desire such content may be directed to the distributed computer system to obtain that content more reliably and efficiently.

[0009] The distributed computer system typically also includes other infrastructure, such as a distributed data collection system 108 that collects usage and other data from the edge servers, aggregates that data across a region or set of regions, and passes that data to other back-end systems 110, 112, 114 and 116 to facilitate monitoring, logging, alerts, billing, management and other operational and administrative functions. Distributed network agents 118 monitor the network as well as the server loads and provide network, traffic and load data to a DNS query handling mechanism 115, which is authoritative for content domains being managed by the CDN. A distributed data transport mechanism 120 may be used to distribute control information (e.g., metadata to manage content, to facilitate load balancing, and the like) to the edge servers. As illustrated in FIG. 2, a given machine 200 comprises commodity hardware (e.g., an Intel Pentium processor) 202 running an operating system kernel (such as Linux or variant) 204 that supports one or more applications 206a-n. To facilitate content delivery services, for example, given machines typically run a set of applications, such as an HTTP Web proxy 207, a name server 208, a local monitoring process 210, a distributed data collection process 212, and the like. For streaming media, the machine typically includes one or more media servers, such as a Windows Media Server (WMS) or Flash 2.0 server, as required by the supported media formats.

[0010] The CDN may be configured to provide certain advanced content delivery functionality, for example, in the case where the edge server does not have the requested content (e.g., the content is not present, the content is present but is stale, the content is "dynamic" and must be created on the origin server, and the like). In such circumstances, the edge server must "go forward" to obtain the requested content. An enhanced CDN often provides the capability to facilitate this "go forward" process. Thus, it is known to provide a "tiered distribution" by which additional edge servers in the CDN provide a buffer mechanism to the Web

# Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

## Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

## Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

## Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

## API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

## LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

## FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

## E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.