

# Search Application Processing From the Edge



# Lucene

powered by Akamai EdgeComputing

## Run Search On Your Web Site With No New Infrastructure

Over the past few years, the amount of information on the Web has grown almost exponentially. While the increased amount of content has made the Web far more useful to the average user, it has also created a virtual haystack of information users must sort through for the needles they desire. Dozens of search application vendors have responded to the problem and many enterprises use their technology to make their sites more usable. While the search application market continues to mature, businesses still face a few key problems when it comes to “search” as an application on their Web sites:

- **High Cost** The software for many of the more common search applications is expensive—with installations costing anywhere from \$100,000 to \$1 million, not to mention ongoing software maintenance fees. For sites that generate a low to medium amount of traffic, the ROI may be difficult to justify. Additionally, for sites that do not do e-commerce, many of the special “bells and whistles” drive up the price, but add little useful functionality.
- **Managing The Infrastructure Is A Hassle** Since “search” tends to be resource intensive, enterprises adding “search” to their sites typically face the choice of either deploying new search-dedicated hardware or degrading performance on existing applications. Additionally, IT personnel often need to closely monitor usage on search in order to handle expected or unexpected surges in demand.
- **Difficult Implementation** Often, to fully leverage the functionality of a search application, the IT staff must overcome a steep learning curve. The time delay lengthens the time to market and damages the ROI the enterprise is looking for.

Lucene powered by EdgeComputing solves these problems by offering search as a turnkey on demand service. Cost is reduced since the offering is based on Jakarta Lucene, an open-source search application created by the Apache Software Foundation. It leverages EdgeComputing’s On Demand Distributed Computing model to run the software entirely on Akamai-managed infrastructure, obviating further hardware expenditures for the enterprise. Finally, even though Lucene is already simple and well-documented, Akamai’s professional services team is available to assist in implementations, ensuring that customers go live within weeks, not months.

## How Lucene powered by Akamai EdgeComputing Works

Lucene is a full-featured text search application library built entirely in Java. It is an open source project of Apache Jakarta. Akamai has optimized Lucene to run as an EdgeComputing

application on Akamai’s EdgePlatform, enabling users to run Lucene in an on demand fashion on managed infrastructure.

*continued on the next page*



## Benefits To Your Business

Lucene powered by EdgeComputing enables users to search a Web site for products or documents using logic that is deployed entirely onto Akamai’s EdgePlatform. Key benefits to the business include the following:

- **Zero Infrastructure Required**  
Because Lucene powered by EdgeComputing is an on demand service, companies needn’t deal with the hassle of provisioning and amortizing hardware. Since search tends to be relatively CPU-intensive, it either severely taxes existing infrastructure or requires procurement of new hardware. Deployed on Akamai’s EdgePlatform, Lucene can scale to virtually any level of traffic that an origin site might see without increasing the load on the origin site.
- **Zero Dependency On The Origin**  
Deployed on Akamai’s EdgePlatform, Lucene can be run wholly as an on demand service without any run-time requests to the origin. Hence, users can continue to perform searches using Lucene even if the origin goes down. This capability makes Lucene a valuable addition to Akamai’s e-Business Continuity solution.
- **Reduced Cost**  
Whereas most search implementations require costly expenditures for software and ongoing maintenance, Akamai is able to offer Lucene at far more competitive pricing. Rather than needing to pay for software licenses, customers can take advantage of EdgeComputing’s pay-as-you-go model.
- **Full-text Search**  
Using Lucene, searches can be performed for any word that is present in any document that has been indexed. This ensures that relevant documents are not overlooked.
- **Simple Deployment And Management**  
As with any other EdgeComputing Application, Akamai provides in-depth insight and control over Lucene while it is running on the EdgePlatform. Click-to-deploy provisioning ensures that customers can modify and re-launch the application when necessary. In-depth reporting capabilities of the EdgeControl<sup>SM</sup> Management Center provide traffic and usage statistics.
- **Fast Time To Market**  
Time is money. Whereas some search options take months to install, Lucene can be implemented within weeks. This capability makes Lucene useful for product launches, marketing campaigns, contests, and other “temporary” Web sites that need to be implemented quickly and easily.



continued

## How Lucene powered by Akamai EdgeComputing Works

There are a few simple steps involved in implementing Lucene powered by EdgeComputing:

### 1. Indexing The Site

The customer creates an "index" file of all documents and other content that it would like to be potentially included in search results. Akamai can help recommend the appropriate parser (or parsers) that the customer can then use to scan through documents, extracting keywords and adding them to the index. (This process can typically be automated for simpler ongoing maintenance.)

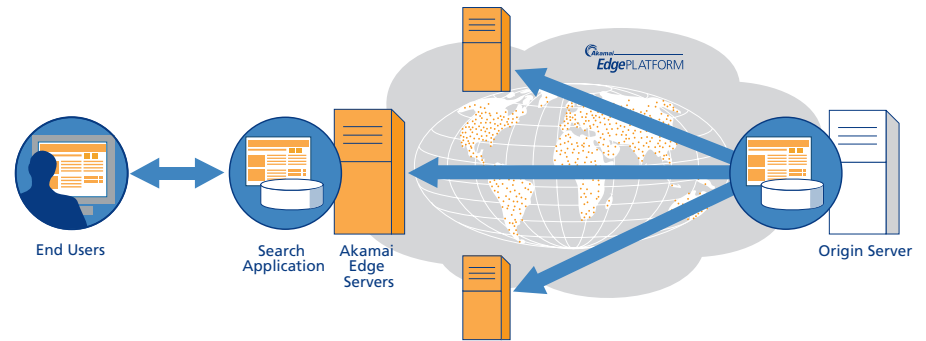
### 2. Building The Result Pages

The customer creates a page or pages to return search results to users. Typically, this is done using one or two simple JSPs, which format the results returned by search queries run against the index file.

### 3. Packaging and Deploying The Application

Using Akamai's EdgeControl Management Center provisioning tools, the customer deploys the application, along with the index, to the edge to be run without any further dependency on the origin.

On an ongoing basis, the customer runs an automated process to generate an index file and package it for deployment to the Akamai EdgePlatform. End-user search queries are then processed entirely on the EdgePlatform, which generates and formats the results for the user.



When users search against the site, the search queries are intercepted by the application and run against the index. The resulting "hits" from the index are reformatted and returned to the user. The user can then click on any of the hits to retrieve the underlying page or document associated with any "hit". In a normal setup, some of these documents may be requested from the origin site and passed through an Akamai server to the user.

As a way of ensuring high availability of not only the search application itself but also these underlying documents, a customer may choose to utilize Akamai's e-Business Continuity Solution. The Site Snapshot Tool, a component of e-Business Continuity, is useful in automating the process of replicating such content onto the EdgePlatform<sup>SM</sup>, helping ensure protection from origin downtime.

## About Akamai

Akamai®—The Business Internet, is the world's largest on demand distributed computing platform for conducting profitable e-business. Overcoming the inherent limitations of the Internet, Akamai's services ensure a high-performing, scalable, and secure environment for organizations to cost effectively extend and control their e-business infrastructure. Headquartered in Cambridge, Massachusetts, Akamai's industry-leading services, matched with world-class customer care, are used by hundreds of today's most successful enterprises and government agencies around the globe. For more information, visit [www.akamai.com](http://www.akamai.com).



**Akamai Technologies, Inc.**  
 U.S. Headquarters  
 8 Cambridge Center, Cambridge, MA 02142  
 Tel 617.444.3000  
 Fax 617.444.3001  
 U.S. toll-free 877.4AKAMAI  
 (877.425.2624)

**Akamai Technologies, GmbH**  
 Heisenbergbogen 2, 85609 Dornach, Germany  
 Tel +49 89 94006-0

[www.akamai.com](http://www.akamai.com)

© 2004 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are federally registered marks. EdgeSuite is a registered trademark. EdgeComputing and EdgeControl are Akamai service marks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice.

APP-LUCENE0404