lucid
IMAGINATION

Thinking Lucene ◥ Think Lucid.

E-Commerce Search Strategies: How Faceted Navigation and Apache Solr/Lucene Open Source Search Help Buyers Find What They Need

A Lucid Imagination White Paper

Executive Summary

A successful store merchandiser knows his/her products and customers. A robust search platform is an essential tool in applying this knowledge to online merchandising. The ability to quickly and easily tailor your search solution for your customers means they will quickly find the exact products they want to buy—a distinct competitive advantage that keeps customers on your site and increasing conversion rates.

Apache Solr/Lucene—the leading open-source search technology—offers sophisticated capabilities that help business analysts and online merchandisers improve online sales. This includes:

- Tools and capabilities that improve *relevance*, helping shoppers find what they're looking for.
- Full use of keywords and other metadata, including Boolean logic and phrases, and easily update this information to accommodate new products and new search strategies.
- Helping customers refine searches to discover the products they want using *faceting*.
- *Language analysis* tools to help improve search results.
- *Multilingual support*, a requirement in the global world of e-commerce.
- The ability to promote popular and high-margin products with *best bets*.
- Hinting and assistance techniques, such as *auto-suggestion*, *related searches*, *more like this*, *recommendations*, and more.

Solr, the Lucene Search Server, also offers several technical advantages. These include the ability to get up and running quickly, rapid prototyping and integration, and efficient utilization of servers and systems. Because Solr/Lucene is open source, software licensing costs remains constant—even as your site's search load increases. New requirements from the business side are easily accommodated; their effect on discoverability/findability can be effectively monitored and analyzed using market-leading tools.

This paper examines some of the key factors to consider when configuring search and discovery on your e-commerce site. Solr, the open-source search server, is leading customers to products at major e-commerce sites around the world. Solr offers ready-to-use features such as faceting, suggestions, and hints, language localizations, and more. Open-source, a worldwide community, and dedicated support from Lucid Imagination mean you can customize it to get the relevant results your customers want and need for a compelling e-commerce experience.

Contents

Introduction

Where a classic brick-and-more storefront, even a warehouse-style retailer, may offer hundreds of choices, online stores offer thousands of choices per product line, often with subtle variations, along with an inventory of hundreds to millions of SKUs. Given the size and scale of the top online shopping sites, merchandising strategies and techniques are vital to online commerce. Consider the following:

- Online sites offer a huge diversity of products, and shoppers need to find what they're looking for, despite not always knowing how to ask for it. Shoppers expect relevant results.
- Online commerce sites can have tens of millions of queries per day, and more during peak times such as the holiday season. Improving search—by even a small percentage—such that more queries deliver what people are looking for will increase conversion, resulting in a significant increase in revenue.
- As shoppers click around your store, they expect near-instantaneous response through your product selection. Speed matters—poor performance pushes potential customers to competitors.
- Suggested selling techniques, such as recommendations, more-like-this, product images, and so on can increase revenue, but different products require different approaches that can be rapidly prototyped, tested, and evaluated.
- Many online commerce sites are adding metadata beyond the basic price, titles, short/long description, and so on. This includes manuals and support information, extended products descriptions, and tags. When properly incorporated, they help drive conversion; otherwise, they can end up bloating the search infrastructure.
- Spelling is a contributor to poor search results. Product marketers, striving to make products stand out, sometimes make things harder because they use branding terms that are not real words. Is the product name one word or two? Is the first letter capitalized or not? Are numbers spelled out with letters? Is the search term generic, or specific?

These factors, among others, contribute to how easily your customers can discover and find what they're looking for, which in turn improves conversion rates. Successful e-commerce sites use a variety of techniques to improve discoverability and findability. Your site can present information such as related searches, similar items, recommendations, and auto-completion. Your customers

can use tools to narrow or broaden results. Improving findability and discoverability helps streamline the user experience and increase sales.

Properly executed site search is a cornerstone of online merchandising, guiding shoppers through a seemingly unlimited selection to the exact product they want to buy. Optimizing and improving the customer shopping experience is a continuous effort, requiring not only a comprehensive set of search capabilities but also an agile IT environment.

In the past, e-commerce vendors were limited to costly, proprietary commercial site-search software, with costly customizations tied to the software vendor. Not only did this approach involve high-priced licenses, but pricing often scaled with the amount of data searched, penalizing growth.

Major e-commerce sites around the world, including AT&T, Buy.com, Macy's, Sears, Zappos, and many other household names, have worked with Lucid Imagination to move to implement their e-commerce search with the Solr[1] open-source search server, leading millions of customers to the products they want to buy. Solr's comprehensive features helps improve discoverability and findability by offering ready-to-use functionality such as faceting, suggestions, and hints, language localizations, and more (as outlined below). Beginning with the attractive economics of open source, along with formidable depth in customization, Solr is supported by a worldwide community, and by the dedicated experts at Lucid Imagination. This means you can both customize your search to tightly fit your business process and consistently give your customers the relevant results they need to make their purchase decisions on your site.

In this paper, we'll review the major capabilities that search with Solr brings to the table for e-commerce. We've also provided a summary table of Solr's key features and their application to e-commerce at the end.

[1] Solr is the Lucene Search Server. It presents a web service layer built atop the Lucene search library and extending it to provide application users with a ready-to-use search platform. Most search applications are best built with Solr. See the Appendix for further detail

E-Commerce Search Checklist

There are multiple factors that contribute to a successful e-commerce site. But no matter how clear your copy or design, how clutter-free your web page, how prominent the shopping cart, or how well-priced the products—if shoppers cannot find what they are looking for, then you won't get the sale. Successful search on e-commerce sites goes beyond simple matching of product names or price points, though that is important, too. This section describes essential factors to consider when implementing search for an online site.

> 80% of online shoppers use the search function to find the products they want.
>
> *Jakob Nielsen, UseIT.com*

Relevance

Good search results find every single item that is relevant, and no documents that are not relevant. Precision search is the ultimate goal, and placement within the "first ten" (the first page) is the goal. This is difficult to achieve—an Aberdeen Group study[2] showed that "best in class" sites only get the most relevant results on the first page 67% of the time, and lower-rated companies only 42% of the time.

There isn't a universal magic formula for search, as it is highly dependent on each user (or class of users) at each site. Practically speaking, there are many factors that go into relevance beyond the strict definitions. In designing your system for relevance, consider the following:

[2] http://www.informationweek.com/news/internet/search/showArticle.jhtml?articleID=220300901

- Do your shoppers prefer accuracy, or would they rather have as many feasible matches as possible?
- How important is it to avoid embarrassing results?
- What factors matter besides pure keyword matches? Do they prefer newer results over older?
- Do they prefer specific brands?
- Do users want results for items that are close to them physically (aka "local" search)?
- Do recommendations help sway them?

Ultimately, the most important consideration is whether your search capabilities simultaneously meet both user needs and business goals. Solr/Lucene offers the ability to use different search algorithms, such as Normalized Discounted Cumulative Gain, Precision/Recall, and others tobest meet your unique goals. This enables you to be proactive about your relevance and more directly apply the knowledge of what your customers are doing into delivering the search results they want to see. For example, Netflix, which uses Solr as its search engine, was able to customize Solr and implemented a search measurement system that gauges the effectiveness of its search results. Known as Mean Reciprocal Rank, or MRR, it gives one point for a click through to the first-ranked item, 1/2-point to the second-ranked item, 1/3-point to the third-ranked, and so on. This provides a very nice aggregate picture of how well users are finding what they are looking for. A good benchmark, or stretch goal, according to Walter Underwood, who developed the Solr-based search system at Netflix: 0.5 MRR, with 85% of users clicking on the first result. Ultimately, this search algorithm improved relevance and increased customer revenue.

> An important consideration for the e-commerce platform of the future: as smart phones become even more popular, your online store should accommodate searching with them as well. With smaller screens and lower available bandwidth, accuracy may be more important. You want to make every effort to ensure the top results are relevant after typing in a keyword. Users on a small device can't or won't scroll through pages of results to find what they needed since the small screen places a premium on putting the right results at the top.

Keywords

Successful searches often with keywords, one of the basic foundations of any search strategy. You will need to continually refresh this aspect of your search strategy over time, adding new terms to reflect new products, or any changes in the way your customers are looking for products.

Figure 1: **Proper keywords and search strategy are essential to finding products that customers want to buy. Searching for the phrase "best selling ipod" produced only these two results.**

Faceting/Discovery

Faceted search—also known as guided navigation—is the dynamic clustering of items or search results into categories that let users drill into search results (or even skip searching entirely) by any value in any field. Faceted search provides an effective way to allow users to refine search results, continually drilling down until the desired items are found. You should consider whether users can select single or multiple options in refining their search, and whether using graphics as part of the results would be beneficial.

Figure 2: Faceted search helps users select products by features or price points that are important to them.

Flexible Language Analysis

Language analysis tools improve the speed and accuracy of your site's e-commerce search. Your product offering may require searches with letters and numbers, mixed case, dashes, spaces, and more. Capabilities such as stemming, case sensitivity, and protected words or phrases improve the results when your search engine can intelligently produce variations on tokens that are being produced, and properly index them.

Multilingual Support

Most retailers have a global reach, with multiple languages in each region. The ability to accept queries and return results in multiple languages enables your e-commerce site to move into new markets. Even if you focus on one country, some portion of your potential customer base may be more comfortable in another language besides the default, or misspell words with non-English characters. Additionally, non-native English speakers may think about products differently than English speakers, requiring additional adjustments to your search capabilities. For example, in some geographies, shoppers may use "bespoke" while in others, "custom."

Frequent Incremental Updates

Suppliers and vendors are constantly updating their product line, including pricing, descriptions, new models, limited-time sales, and close-outs. User-generated content, such as mini-reviews and

ratings, also changes frequently. All this means you may need frequent incremental updates to your index to make these changes available as soon as possible.

Best Bets

Best-bet capabilities enable you to return a specific result even if the calculated result is something different. Typically tweaking metadata or the search configuration doesn't always ensure that the most relevant results always appear at the beginning of the list. The concept of "best bets" is gaining in popularity, and many users are coming to expect it. It can be influenced by a user's past behavior, on popular searches or items, or a combination. Sales can be increased by overriding search results with a popular item that has not yet become the most relevant in your search infrastructure. Care must be taken to monitor the results over time, and keep from overuse.

Hinting and Assistance

There are several ways that site search can improve the results, and the likelihood of a transaction. This includes:

- **Auto-suggest, or auto-complete**: A popular feature of most modern search applications is the auto-suggest or auto-complete feature where, as a user types their query into a text box, suggestions of popular queries are presented. As users type in additional characters, the list of suggestions is refined. This can help overcome any spelling issues or finding non-standard terms—users can see what others have used for searches.

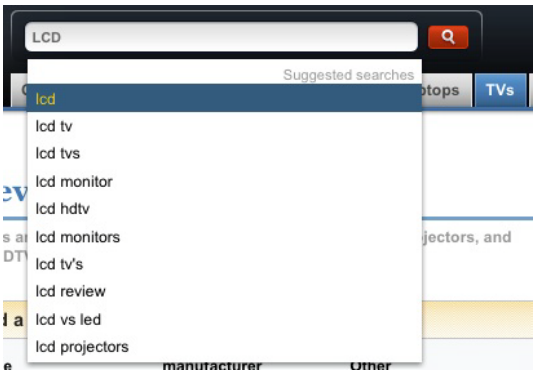

Figure 3: Auto-suggest

- **Did you mean?:** All modern search engines attempt to detect and correct spelling errors in users' search queries. A good "Did you mean?" capability can use linguistics and lexicons as well as past user behavior to suggest/verify that the user submitted the correct query and overcome issues related to typing errors, phonetics, compound or separated words, and singular or plural words.

- **More like this:** Shows what customers who bought this item also bought. This can include similar items, such as competitive products, or complementary items, such as accessories, or both.
- **Related search:** Shows the results of queries similar to the user's. For example, related search to "TV" would show "LCD TVs," "plasma TVs," and "mp3 players" would show "iPod," "Zune," and so on.
- **Recommendations:** Similar to related search, but surfaces what other users bought or recommended. Results are often based on the collective interactions of previous users, but might also consider the content itself as part of the determination.

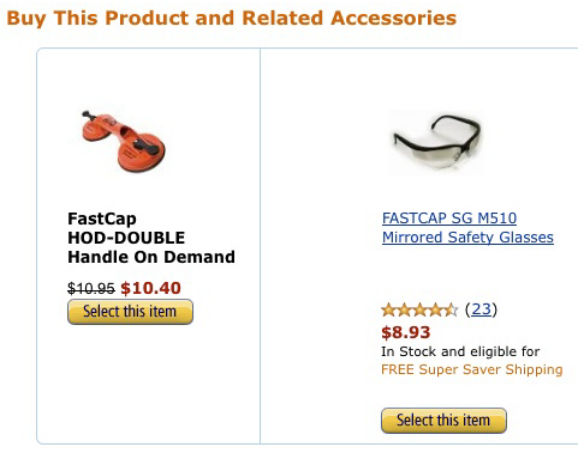

Figure 4: **Related products and accessories, based on browsing history.**

Business and Administrative Capabilities

There are several search capabilities that are useful to site operators. For example, editorial relevance controls enable a business analyst to provide higher ranks to products with higher margins, or promote products where there is excess inventory while still meeting user needs. Effective tools allow analysts to give higher ranks to higher margin products, or products with excess inventory, and also help monitor activity.

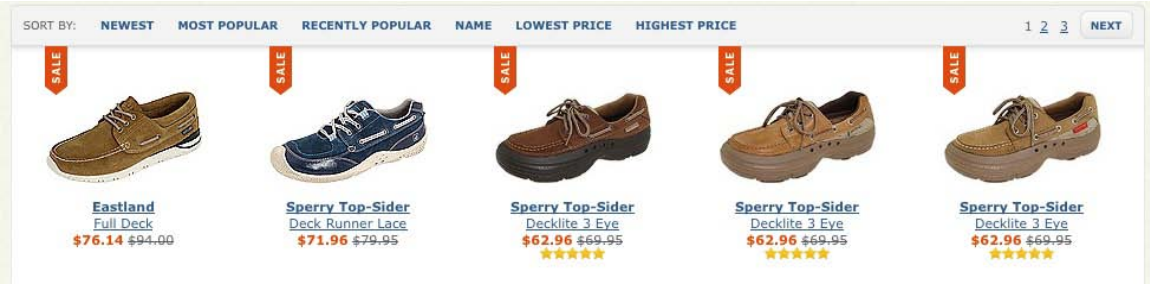

Figure 5: **Business analysts may want to affect search results, for example, by listing certain products first. Here, a search for "deck shoes" listed products on sale at the top of the results.**

Business and analytics tools should also be able to provide testing and monitoring of business activity. These provide feedback that what you are doing is moving you toward the goal of maximizing sales (or other business goals, such as reducing customer service and support calls). Analytics should be able to examine logs and highlight what users are doing, how often they are doing it, and what results they are getting. Business tools should also allow for experimentation (split-testing or A/B testing for instance) and insight into the engine's choices in order to properly try out new approaches and fix any issues that might arise.

Any site search solution should also have robust administrative tools that facilitate operations. This includes easy set-up and configuration, the capability to scale up quickly and effectively to accommodate rapid growth, fault tolerance, and high availability to ensure 24x7x365 access.

E-commerce and Solr

Solr is uniquely powerful in its combination of the adaptive resilience desired by e-commerce marketing leaders with a very robust technical platform.

While relevancy and speed are the core requirements for any search solution, increasing findability improves the customer experience with your site, and increases sales. Solr drives site search on some of the most high-traffic, high-value sites on the web, including Netflix, Zappos and CNET—pushing the boundary on many of the key features described above, including auto-correct, suggested selling, localization and more. Solr is cost beneficial, feature rich and highly flexible and adaptable, increasingly becoming the number one choice for powering site search for online retail.

Solr provides e-commerce application builders a ready-to-use search platform on top of the Lucene search library. It ranks among the top 10 open source projects, with installations at over 4,000 companies. Lucene/Solr have seen such tremendous rates of adoption for powerful reasons.

Solr's features are derived everyday from users who are improving and evolving e-commerce capabilities far faster than any proprietary solution. Solr offers state-of-the-art search capabilities, including excellent performance, relevancy ranking, and scalability. Most importantly, it can be customized and tuned to your site, guiding your customers to your products. (For a more detailed technical description of Lucene and Solr, see the Appendix.)

Solr is an open-source enterprise search server based on the Lucene Java search library, with XML/HTTP APIs, caching, replication, and a web administration interface. Apache Solr is used by top e-commerce sites because it offers out-of-the-box functionality, is highly configurable and customizable, and delivers best-in-class features and functionality. The Lucene library that Solr delivers in "serverized" format was originally released several years prior, but today Solr is the go-to search application development platform for high-performance, highly-customizable, highly scalable e-commerce search. Open source Solr can change the way your site search works as fast as you need it to, enabling you to respond to dynamic market conditions. You can try new ways to improve findability, more easily adapt metadata from new and different suppliers, and otherwise have complete control over a core driver of your site's revenue potential—your site search. For example

- The structure of the search data is separate from the search application, enabling you to search the way you want, and change it as needed. The search strategy can be changed while in production.
- Instead of a "one-size-fits-all" strategy, you can have multiple strategies to optimize results by user class or product line.

Solr has superior capabilities in relevancy ranking, performance, and scalability, all of which help turn shoppers into customers. Solr offers many advantages, including:

- **Widely-accepted technologies.** Built using popular standards such as Java, RESTful APIs, and XML, Solr simplifies configuration and operation, and narrows the skill set required for and modifications. This means it's ready to go in Jetty, Tomcat, and other leading servlet containers, and is easily modified to suit your purposes.
- **Built-in Lucene best practices.** Solr offers an enormous set of search features, including caching filters, queries, or documents, spell checking, hints and suggestions, and performance enhancements such as background warming Searchers – features that applications using Lucene java libraries directly would have to implement themselves.
- **Infrastructure services.** File operations, memory management, I/O configuration, administration, and many more platform capabilities are already there, so you don't need to write them yourself.

Solr has flexibility and offers plenty of customization to meet with the tremendous amount of growth that internet-scale ecommerce applications need. That you must know your customer and how they think about your products is well understood, but tuning your search engine to address this is key to a successful e-commerce site. For example, the terms "notebook" or "laptop" when referring to portable computing devices may seem totally synonymous to some, while to others they may be separate products altogether. (The Top 1000 site buzzle.com offers on article on when to buy a notebook versus a laptop[3].) Additional complexities, such as regional dialects, different languages, emerging slang all affect how your customers search for products, and how relevant your results will be. Without a good grasp of subtleties such as these, you are missing sales.

Effective search is key to a successful e-commerce operation. It is the online equivalent of merchandising, guiding shoppers to the exact items they are looking for. Because of the large number of SKUs and potential customers only rapid, relevant search results will turn shoppers into sales. Solr offers market-leading search performance and features. Optimizing Solr for your business requires mission-critical capabilities.

[3] http://www.buzzle.com/articles/notebook-vs-laptop.html

Summary E-Commerce Solr Feature Checklist

Solr Feature	Description
Full-Text Search	Solr uses the Lucene library for full-text search, and provides excellent results because it compares every word, and not just an abstract or set of associated keywords.
Advanced capabilities	Solr offers a robust set of advanced search features and capabilities: ▪ Faceting and multiselect faceting: Narrow search results by one or more sets of criteria. ▪ Accurate probabilistic ranking: More relevant documents are listed first. ▪ Phrase and proximity searching: Searching by exact or similar phrases. ▪ Relevance feedback: Improves ranking and can expand a query, find related documents, categorize documents, etc. ▪ Structured Boolean queries: Use AND, NOT and other expressions. ▪ Wildcard search: Substitute defined characters in words to expand search results. ▪ Spelling correction: Offer spelling options as user types query to improve results. ▪ Query assistance: Solr helps users find what they needs with more like this suggestions, auto-suggest, sounds like, and best bets. ▪ Enable configuration of top results for a query, overriding normal scoring and sorting. ▪ Highlighting: Hits can be highlighted to help users see the results of their queries.
Language analysis	Solr offers a rich set of customizable language analysis capabilities, including analyzers, filters, tokens and token filters, and stemmers that can evaluate, extract, and manipulate text to improve findability.
Currency/freshness	Simultaneous search and update, with immediately visibility for new documents. Solr offers several optimizations and techniques for ensuring that the search index is current, including the ability to add new documents or information without the need to modify unchanged segments.
Extensibility	▪ Plug-ins quickly and easily expand functionality and administrative capabilities. Available plug-ins include Terms for auto-suggest, Statistics, TermVectors, Deduplication. ▪ Flexible and adaptable with XML configuration (schema.xml), such as data schema with dynamic fields, customizable Request Handlers and Response Writers, and more.
Rich Document processing	Solr generally uses XML documents corresponding to the schema structure of your schema, but other formats can also be used: PDF, HTML, Microsoft OLE 2 Compound Document (Word, PowerPoint, Excel, Visio, etc.), and other formats such as zip and Java Archive (JAR).

Solr Feature	Description
Administration	▪ Out of the box, runs in a servlet container such as Tomcat or Jetty. ▪ Ready to scale in a production Java environment. ▪ In Solr 1.4, replication Is abstracted and implemented entirely at the Java platform layer; it works the same wherever the Java platform runs. ▪ Uniformly configure replication across multiple Solr instances. ▪ Replication does not require a backup and the index is copied from one live index to another. ▪ Backups can be performed in the same way on a Solr instance, regardless of hardware or operating system. ▪ Server statistics exposed over JMX for monitoring ▪ Web-based administration tool, or interface to enterprise system management tools. ▪ Monitorable logging
Web scalability	▪ Solr is optimized for demanding web environments, and provides results on 10s of millions of queries per day at leading sites. ▪ Efficient caching and replication, incremental updates. ▪ Integration with other open-source technologies ▪ Hadoop: Sharded index across multiple hosts ▪ Mahout: Scalability for reasonable large data sets, including core algorithms for clustering, classification, and batch-based collaborative filtering implemented on top of Apache Hadoop using the map/reduce paradigm.
Performance	Maximum throughput and minimum response time enabled by a high-performance architecture, including sharding (index split across multiple servers), multithreading, optimized libraries, and more. For example, caching moves frequent search results into memory from disk, and caches can be warmed in background, or autowarmed.
Standards-based, open APIs and widely-accepted technologies	XML, HTTP, Java, RESTful, JSON, PHP, Ruby, Python, XSLT, RESTful APIs, C#, C.
Language support	Solr supports over 25 languages.

About Lucid Imagination

Lucid Imagination can help you use Solr to get the most from your ecommerce search applications. Lucid Imagination has the world-class expertise, resources, support and services needed to cost-effectively architect, implement, and optimize Solr/Lucene-based solutions. We provide commercial-grade support, training and consulting and by offering certified, tested versions of Lucene and Solr. Lucid Imagination's goal is to serve as a central resource for the entire Lucene community and marketplace, to make enterprise search application developers more productive. We also provide access to Solr/Lucene experts, well-organized information, and documentation.

We've have helped hundreds of companies get the most out of their search infrastructure. Customers include AT&T, Buy.com, Cisco, Ford, Macy's, Sears, Shopzilla, The Motley Fool, Verizon, Edmunds.com, GSI Commerce, Zappos (Amazon), and many other household names. Lucid Imagination is a privately held venture-funded company. The investors include Granite Ventures, Walden International, In-Q-Tel and Shasta Ventures. To learn more please visit

http://www.lucidimagination.com

http://www.lucidimagination.com/solutions/services

For more information on what Lucid Imagination can do to help your employees, customers, and partners get the most out of your e-commerce efforts:

- Support and Service inquiries: support@lucidimagination.com
- Sales and Commercial inquiries: sales@lucidimagination.com
- Consulting inquiries: consulting@lucidimagination.com

Or please call: 650.353.4057

Appendix: Lucene/Solr Features and Benefits

Lucene and Solr are complementary technologies that offer very similar underlying capabilities. In choosing a search solution that is best suited for your requirements, key factors to consider are application scope, development environment, and software development preferences.

Lucene is a Java technology-based search library that offers speed, relevancy ranking, complete query capabilities, portability, scalability, and low overhead indexes and rapid incremental indexing.

Solr is the Lucene Search Server. It presents a web service layer built atop Lucene using the Lucene search library and extending it to provide application users with a ready-to-use search platform. Solr brings with it operational and administrative capabilities like web services, faceting, configurable schema, caching, replication, and administrative tools for configuration, data loading, statistics, logging, cache management, and more.

Lucene presents a collection of directly callable Java libraries and requires coding and solid information retrieval experience. Solr extends the capabilities of Lucene to provide an enterprise-ready search platform, eliminating the need for extensive programming.

Solr provides the starting point for most developers who are building a Lucene-based search application. It comes ready to run in a servlet container such as Tomcat or Jetty, making it ready to scale in a production Java environment.

With convenient ReST-like/web-service interfaces callable over HTTP, and transparent XML-based configuration files, Solr can greatly accelerate application development and maintenance. In fact, Lucene programmers have often reported that they find Solr contains "the same features I was going to build myself as a framework for Lucene, but already very well implemented." Using Solr, enterprises can customize the search application according to their requirements, without involving the cost and risk of writing the code from the scratch.

Lucene provides greater control of your source code and works best in development environments where resources need to be controlled exclusively by Java API calls. It works best when constructing and embedding a state-of-the-art search engine, allowing programmers to assemble and compile inside a native Java application. While working with Lucene, programmers can directly control the large set of sophisticated features with low-level access, data, or state manipulation.

Enterprises that do not require strict control of low-level Java libraries generally prefer Solr, as it provides ease of use and scalable search power out of the box.

As functional siblings, Lucene and Solr have become popular alternatives for search applications; the two differ mainly in the style of application development used. Key benefits of search with Lucene/Solr include:

- **Search Quality: Speed, Relevance, and Precision** Lucene/Solr provides near-real-time search and strong relevance ranking to deliver contextually relevant and accurate results very quickly. Tailor-made coding for relevancy ranking and sophisticated search capabilities like faceted search help users in sorting, organizing, classifying, and structuring retrieved information to ensure that search delivers desired results. Search with Lucene/Solr also provides proximity operators, wildcards, fielded searching, term/field/document weights, find-similar functions, spell checking, multilingual search, and much more.
- **Lower Cost and Greater Flexibility, Plug and Play Architecture** Lucene/Solr reduces recurring and nonrecurring costs, lowering your TCO. As open source software, it does not require purchase of a license and is freely available for use. The open source code can be used as is, modified, customized, and updated as appropriate to your needs. Solr is easily embedded in your enterprise's existing infrastructure, reducing costs of installation, configuration, and management.
- **Open Source Platform for Portability and Easy Deployment** Because Lucene/Solr is an open-source software solution, it is based on open standards and community-driven development processes. It is highly portable and can run on any platform that supports Java. For instance, you can build an index on Linux and copy it to a Microsoft Windows machine and search there. This unsurpassed portability enables you to keep your search application and your company's evolving infrastructure in tandem. Lucene, in turn, has been implemented in other environments, including C#, C, Python, and PHP. At deployment time, Solr offers very flexible options; it can be easily deployed on a single server as well as on distributed, multiserver systems.
- **Largest Installed Base of Applications, Increasing Customer Base** Lucene/Solr is the most widely used open source search system and is installed in around 4,000 organizations worldwide. Publicly visible search sites that use Lucene/Solr include CNET, LinkedIn, Monster, Digg, Zappos, MySpace, Netflix, and Wikipedia. Lucene/Solr is also in use at Apple, HP, IBM, Iron Mountain, and Los Alamos National Laboratories.
- **Large Developer Base and Adaptability** As community developed software, Lucene/Solr provides transparent development and easy access to updates and releases. Developers can work with open source code and customize the software according to business-specific needs and objectives. Its open source paradigm lets Lucene/Solr provide developers with the freedom and flexibility to evolve the software with changing requirements, liberating them from the constraints of commercial vendors.

- **Commercial-Grade Support for Mission Critical Search Applications from Lucid Imagination** Lucid Imagination provides the expertise, resources, and services that are needed to help enterprises deploy and develop Lucene-based search solutions efficiently and cost-effectively. Lucid helps enterprises achieve optimal search performance and accuracy with its broad range of expertise, which includes indexing and metadata management, content analysis, business rule application, and natural language processing. Lucid Imagination also offers certified distributions of Lucene and Solr, commercial-grade SLA-based support, training, high-level consulting and value-added software extensions to enable customers to create powerful and successful search applications.