

GET MORE OUT OF YOUR DATABASES WITH SEARCH-BASED APPLICATIONS (SBAs)



TABLE OF CONTENTS

1) The Restrictive World of Database Applications.....	1
2) The Search-Based Application (SBA): Liberating Data and Data Access.....	1
3) How Database SBAs Work.....	1
3.1 The Search Engine Index: An Alternate Data Access Layer.....	1
3.2 Querying and Output.....	2
4) Architectural Features & Benefits.....	2
4.1 Infrastructure Level: High Performance with Database Offloading.....	2
4.2 Application Level: Agile Development.....	3
4.3 End User Level: A Web-Style Experience.....	4
5) Applications & Case Studies.....	5
5.1 Information Access Applications.....	5
5.2 Operational Data Access and Reporting.....	6
5.3 Comprehensive Business Applications.....	6
6) EXALEAD CloudView™ for Database SBAs.....	7
6.1 Scalability & Performance.....	8
6.2 Agility.....	8
6.3 Usability.....	8
7) Conclusion.....	9

FIGURES

Fig 1: The Traditional Database Application Model.....	1
Fig 2: A 'Flat' Data Representation Model that Retains Classification Information.....	1
Fig 3: Replace Complex, Pre-Determined Query Forms with Web-Style Search.....	2
Fig 4: A Clean Start, Infinitely Rich Exploration.....	2
Fig 5: SBA Architecture: Benefits at the User, Application & Infrastructure Levels.....	2
Fig 6: Agile SBAs Provide Flexible, 'Forgiving' Information Search and Retrieval.....	4
Fig 7: SBA-Powered Search and Navigation for Rightmove.co.uk.....	5
Fig 8: Search Platform: an Ideal Tool for Improving Operational Decision Making.....	6
Fig 9: Database Offloading and Operational Reporting for GEFCO.....	6
Fig 10: A New Customer Service Portal for a National Postal Service.....	7
Fig 11: Unlimited, Effortless Scaling with Load Balancing, Partitioning & Replication.....	8

1) THE RESTRICTIVE WORLD OF DATABASE APPLICATIONS

For the past 25 years, databases have served not only as the primary storage vehicle for business information, they have also provided the primary means of accessing that information. In the traditional database application model, information is entered into a relational database, and end users access this data via pre-determined SQL (Structured Query Language) queries executed within a commercial-off-the-shelf (COTS) or custom application.

Restrictive for End Users

With this model, if end users' access needs have not been anticipated, or if they are not among the elite minority that knows SQL, they cannot get to the information they need.

Unpleasant for IT

The traditional database application model is also challenging for IT staff, forcing them into an uncomfortable intermediary role, trying to anticipate and 'hard code' users needs into applications that are difficult or impossible to modify, and seeking to avoid system slowdowns and crashes by limiting database access and shifting data updates to off hours, sacrificing the timeliness of available information. The model is also labor and resource intensive, incurring heavy development, licensing, and infrastructure costs.

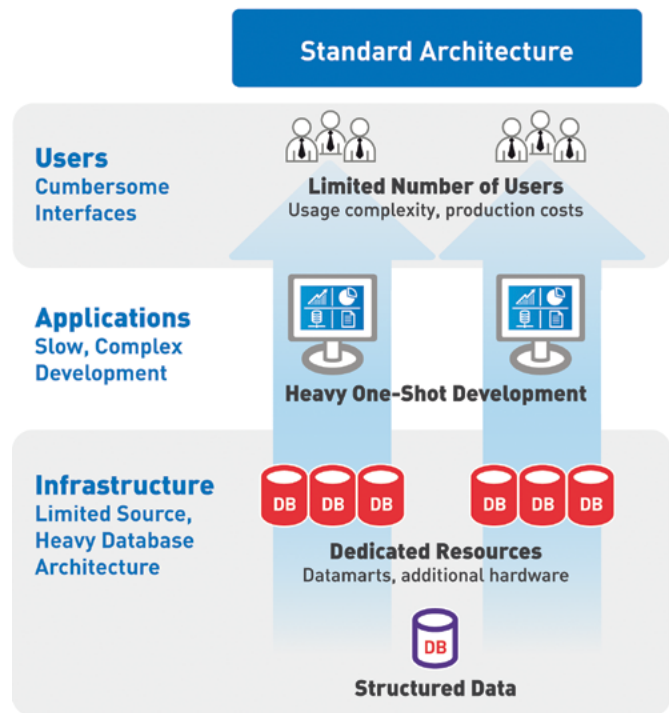


Fig 1: The Traditional Database Application Model

2) THE SEARCH-BASED APPLICATION (SBA): LIBERATING DATA AND DATA ACCESS

Unlike traditional database applications, SBAs provide information access via a search engine index and complementary Web technologies rather than through direct database queries. The result is an information access solution that more closely resembles the Web than an enterprise system:

End User Experience

Characteristic	Traditional Database Applications	Database SBAs
Interface	Complex forms, pre-determined query options	Simple forms, fuzzy natural language search, open exploration
Data Volume	Millions of records	Billions of records
Data Freshness	Day (data warehouse)	Real-time or quasi-real-time
Responsive Time	Minutes	Sub-second

Application Development

Characteristic	Traditional Database Applications	Database SBAs
Time to Deploy	Months, years	Days, weeks
Ability to Evolve	Difficult	Rapid, easy
User Base	Limited (training, performance, cost)	Unlimited
Cost of Ownership	Very high per user	Very low per user

3) HOW DATABASE SBAS WORK

3.1 The Search Engine Index: An Alternate Data Access Layer

In the SBA framework for database applications, a search engine indexes content from a database system (individual databases, data marts or data warehouses), creating an optimized view of the information contained in the database system.¹ End user and/or application queries are then run against this index rather than directly against the database system (employing a 'document model' rather than a 'data model' for information access, and achieving database offloading).

Unlike the familiar world of Web search, where 'flat' indexes can only produce laundry lists of results based on keyword searches, today's advanced semantic search engines produce an index that is fully optimized for fast information access ('read' operations), yet nonetheless can extract and retain the complex categorization and classification information stored in relational data tables.²




Fig 2: Indexes Use a 'Flat' Data Representation Model yet Retain Classification Information

3.2 Querying and Output

Accordingly, index-based access is easy yet multi-dimensional: deeply structured queries, numerical operations, fuzzy natural language search and faceted navigation are all natively supported. Query results can be dynamically output as hyperlinked lists, categorized data, bar charts, maps, real-time operational reports and dashboards—the possibilities are endless. This inherent flexibility is unavailable with both traditional database technologies and traditional search engine technologies.³

TECH NOTES FOR DBAS

1 Feeding/Updating the Index: With an engine like EXALEAD CloudView™, you can choose the method that works best for you: a pull strategy with an ODBC or JDBC connector, a push strategy using a Java API, an ESB messaging layer—you are in control.

2 Structured Queries & 'Flat' Indexes: Semantic indexes put multiple data attributes together in a single field, yet these attributes can still be queried individually, as though they remained in discrete columns. As the DBA, you control the mapping of your database model to the index model.

3 Multi-Dimensional Results Output: In considering how a single query can automatically generate clickable related categories and statistical elements like charts and graphs, think of the search engine as executing on-the-fly and in a single pass not only a SELECT statement but multiple GROUP BY clauses and automatically running any counters and aggregates you may desire.

A Simple Start...

As in the familiar world of the Web, users can begin their search by entering a simple, natural language query in a single text box.




Fig 3: SBAs Replace Complex, Pre-Determined Query Forms with On-the-Fly, Web-Style Search

...Rich Exploration

From this simple start, users can drill down on an unlimited number of data characteristics for reporting and research, a breadth and depth impossible to achieve using pre-determined SQL queries. This is due to the unique capacity of a search engine like CloudView to generate navigational links on the

fly from all of the classification information maintained in the source database.



Fig 4: A Clean Start, Infinitely Rich Exploration
Clickable summary and detail views related to the user's specific query are automatically generated for any and every data characteristic available in the database. There is no need to pre-determine and hard code queries for all available views, summaries and reports.

"Here, you have a clean start but a rich finish in results. When you do a search, what you get back is deeply textured...This is where information access applications are going—clean start, textured results. Don't make the mistake I kept making: be open to new information access strategies."

Whit Andrews, Vice President, Gartner Research

4) ARCHITECTURAL FEATURES & BENEFITS

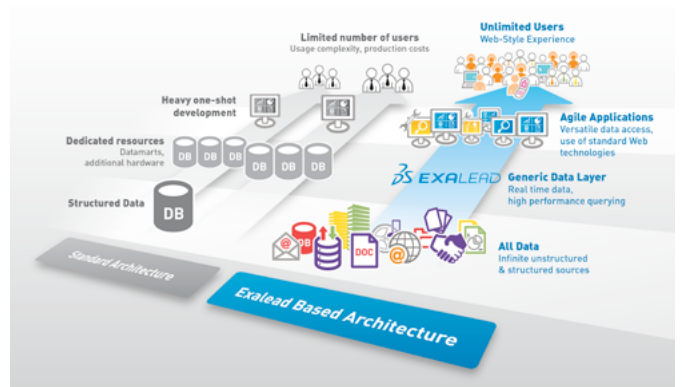


Fig 5: A Database SBA Architecture Provides Benefits at the End User, Application and Infrastructure Levels

4.1 Infrastructure Level: High Performance with Database Offloading

As most database queries entail simple access operations (basic search and retrieval or read operations), shifting these queries to the index dramatically reduces the load on existing database systems while preserving their availability for mission-critical transactions (write operations). This database offloading provides a wide range of benefits at the infrastructure level, including:

- **Reduced IT Costs**

Database offloading enables you to reduce the licensing and infrastructure costs associated with scaling database systems and meeting data integration needs (search engines can not only index a single database, they can extract and synthesize data from multiple systems).

- **Improved Performance**

With the load thus reduced, the performance of existing systems during transaction processing immediately improves. And for information access requests processed by the search engine, sub-second responsiveness is standard—a rate on average 100s of times faster than database query processing.

- **Extended Access & Unlimited Scalability**

Designed from inception to handle high volume data access requests for a virtually unlimited user base, search engines enable you to instantly expand access to the content in your databases without having to scale the database infrastructure itself. Post-deployment, you can scale access even further simply by adding inexpensive commodity hardware to the search platform.

- **More Timely Data**

With dramatically reduced loads on their databases, many companies find they can afford more frequent data updates, sometimes dropping off-hour batch updates altogether in favor of real-time incremental updates, which search engines are intrinsically designed to support.

- **More Secure Access**

An SBA architecture shifts authentication from the application layer to the infrastructure level, providing single sign-on convenience yet natively enforcing security down to the metadata level. SBA platforms can also respond in real time to changes in rights and permissions.

4.2 Application Level: Agile Development

Deployment of an SBA architecture for database applications brings unprecedented agility at the application level as well. It is an ideal architecture for iterative development, a strategy of rapid prototyping and early, regular application releases that maximizes feedback and almost universally results in applications more closely aligned with end user needs. SBA platforms can be deployed in just a few days, and an endless variety of advanced applications built upon them in on average 2-8 weeks. They can also be modified quickly and easily, without altering underlying information systems. This agility is enabled by three factors:

- The independence of the data layer

- The use of a service-oriented architecture and standard Web formats and protocols
- A drastically reduced need to anticipate and plan for user needs

4.2.1 An Independent Data Layer

First, in building the index, the search engine essentially decouples data from the underlying applications that generate it, creating an independent data layer that can be directly searched or exploited by an unlimited number of users and applications. Furthermore, this index build process is completely non-intrusive, meaning it does not require changes to, or interfere with, existing information systems. Maintaining the index is likewise non-intrusive: one can add new data or data sources and construct (or reconstruct) data relationships on the fly without affecting production data. This non-intrusiveness frees IT to experiment, encouraging innovation while ensuring business continuity and preserving existing IT investments.

Non-intrusive SBA platforms do not interfere with existing IS infrastructures

4.2.2 A Service-Oriented Architecture and Standard Web Formats and Protocols

Second, if the search engine is well-adapted for SBA development, it will feature a service-oriented architecture (SOA), which means its core operations (like indexing and query processing) are made available as on-demand 'services' that can be easily tapped by other applications or services using standard Web technologies (SOAP, REST, XML, RSS, RDF, OWL, etc.). It is a 'light' application framework that facilitates prototyping and integration and shortens development cycles.

4.2.3 A Reduced Need to Anticipate and Plan for User Needs

In a traditional database application, data categorization and classification needs must be manually built into the structure of the relational database tables, and the range of possible end user information requests and desired views anticipated in advance and hard-coded in the application programming. This necessitates extensive planning and complex programming. If unanticipated needs arise, modifications are difficult to implement.

An index-based solution eliminates these constraints. Semantic search engines enable intuitive, 'forgiving' search. Users can retrieve data on the fly using either fuzzy natural language search or structured queries, and the engine itself performs the heavy work of interpreting user requests, identifying and presenting related category and classification data, as well as producing

relevant statistical counts and dynamic summaries and reports. This flexibility and automation also significantly accelerates development cycles.

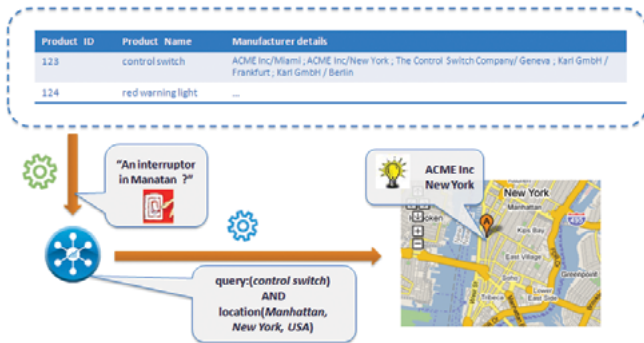


Fig 6: Agile SBAs Provide Flexible, ‘Forgiving’ Information Search and Retrieval Here, a semantic engine understands that an “interruptor” and a “control switch” are synonymous, that the user meant “Manhattan” though they typed “Manatan”, and that Manhattan refers to a borough of New York City. Despite the imprecision of this query, the engine successfully locates a provider of control switches in Manhattan, ACME, Inc., and produces the query results in map form.

4.2.3 Benefits of Agile Development

- Reduced Development Costs**
 Shorter development cycles reduce the staff labor costs associated with traditional development, including labor costs for database administrators, application developers and the end users who must participate in lengthy planning processes.
- Greater Independence**
 The reduced need to anticipate and hard code user needs and the simplicity of developing applications using standard Web formats and protocols means complex database applications that are frequently outsourced can be brought in-house, or at least easily modified internally even if they are still outsourced.
- Faster Time to Market**
 Comprehensive out-of-the-box functionality, easier programming and iterative development provide a competitive advantage in getting new applications to market quickly.
- Increased Responsiveness to Shifting Demands**
 When user demands or market conditions change, a business’s capacity to rapidly adapt is key to surviving in competitive markets. As SBAs are both easy to modify and can easily scale on demand, they significantly aid this capacity.
- Mitigated Risk**
 Because SBAs decouple the data lifecycle from the application lifecycle, new features can be introduced without affecting production data, and data migration projects can

proceed at an infrastructure level without impacting end user information access. This flexibility, along with non-intrusive scaling, helps mitigate risks and ensure business continuity.

4.3 End User Level: A Web-Style Experience

SBAs provide an experience close to that of the consumer Web, an experience that is intuitive, timely, responsive and relevant. This significantly boosts user satisfaction and enables IT to democratize information access.

4.3.1 An Experience Close to the Consumer Web

Intuitive & Engaging

SBAs support queries as deeply structured as relational database queries, yet they also offer alternative access in the form of faceted navigation of information and support for Web-style fuzzy natural language search, making database content available to any user—no matter their technical skills. Because they use standard Web technologies, SBAs also enable the rich Web 2.0 interactivity and collaboration enterprise workers are increasingly demanding of business applications.

‘Zero-training’ SBAs parallel a consumer Web experience

Timely

Due to the batch updates associated with traditional database architectures, business people too often must resign themselves to working with data that is hours or even days old. Having been designed for real-time access to staggering volumes of information, the Web-born SBA architecture enables real-time data access and operational reporting in a corporate environment, boosting competitiveness and agility and more closely aligning enterprise IT with the world of the Web.

Responsive

SBAs also provide the sub-second responsiveness to which Web-savvy users have become accustomed. Captive business users may have no choice but to cope with slow applications, but who would linger on Amazon if queries took 15—or even 5—minutes to process?

Relevant

Because SBAs can automatically generate context-sensitive navigation of related information, migrating to an SBA framework can immediately boost users’ sense of the relevancy of available information, even though it is exactly the same content as that treated by the legacy application.

4.3.2 Benefits of an Enhanced User Experience

- **For Business Applications: Reduced Costs & Added Value**

For enterprise applications, a more intuitive user experience cuts training costs and derives new value from existing information assets through democratized access, high application adoption rates and extensive daily usage. It also reduces the labor costs associated with unsuccessful search: IDC estimates that information workers spend 48% of their time searching for and analyzing information, with one third of that time resulting in failed searches (and re-created work), costing organizations \$28,000 per worker per year (*IDC Predictions 2009, 12/2008*).

- **For Commercial Applications: Increased Loyalty, Higher Traffic, More Revenue**

For online businesses, a better user experience has a highly visible impact on the bottom line, generating more revenue through improved customer loyalty, higher conversion rates, increased site traffic, and/or better exposure of related products and services.

5) APPLICATIONS & CASE STUDIES

The end uses for database SBAs are as varied as traditional database applications, serving needs from online publishing to supply chain management to customer service, though at the functional level, SBAs can be classified as: **1)** information access applications, **2)** operational data access and reporting systems, or **3)** comprehensive business applications. The common structural benefits for all include database offloading, agile development and an improved end user experience.

5.1 Information Access Applications

The independent data layer created under the SBA model is optimized for fast, real-time information access, versus the data mining and transaction processing for which databases were designed. 'Information access' encompasses search and retrieval and content presentation, taking advantage of the engine's built-in capacity for multi-dimensional access and faceted content presentation:

- **Multi-Dimensional Access**

On the search side, SBAs can process Web-style natural language queries from users, as well as application-level structured querying (XML, XPATH, XQUERY, etc.)

- **Multi-Dimensional Facets**

For content presentation, semantic search engines can automatically mine data for inherent categories, classification metadata and relationships, and generate context-specific faceted navigation on the fly.

Case Study: Rightmove



Fig 7: SBA-Powered Search and Navigation for Rightmove.co.uk

Rightmove, the UK's top real estate classifieds portal, wanted to enhance their end user experience, improve system performance, and reduce IT costs. They decided to shift from a classic database application with queries run directly against an Oracle database system to an index-based SBA strategy using EXALEAD CloudView™. As a result, they have been able to:

- **Reduce Costs and Improve Performance through Database Offloading**

Using an index-based approach, Rightmove provides search against 2 million real estate ads for 29 million monthly visitors, with a peak processing throughput of 400 queries per second (QPS) at a 99.99% availability rate. They have achieved this level of performance while reducing their IT infrastructure—replacing 30 Oracle CPUs with 9 EXALEAD CPUs—and reducing their cost of search per 100 queries from £0.06 to £0.01.

- **Achieve Agile Development, with a Rapid Time to Market and Development Independence**

With CloudView, Rightmove was able to take the new platform to market in only 3 months, with minimal professional services support. With data connections handled by a built-in ODBC Connector, and application customization facilitated by open, standards-based application programming interfaces (APIs), Rightmove's IT staff have also achieved the independence they have been seeking, being able to modify or expand functionality entirely on their own. This independence and agility also extends to scaling. They can easily scale the platform on demand simply by adding inexpensive commodity hardware.

- **Improve the End User Experience, with Simpler, More Robust Search and More Timely Data**

With a smart semantic engine under the hood, Rightmove's new CloudView SBA provides search and navigation features that are at once more intuitive and more powerful. The system automatically incorporates data facets that could only be made accessible through manual programming with the legacy application. Users also have access to more timely data, with a data refresh rate of less than 2 minutes.

“Rightmove has already found that EXALEAD CloudView™ has allowed the speedy development of advanced search functionality whilst reducing search costs by 83%.”

Peter Brooks-Johnson, Product Director, Rightmove

5.2 Operational Data Access and Reporting

With a native capacity for indexing staggering volumes of data in real time and built-in data clustering and statistical computation capabilities, SBAs are also ideal for operational reporting. Operational reporting complements analytical reporting by extending real-time information access to the broadest possible user base to better support day-to-day, operational decision making. This approach matches the right tool with the right job: traditional database applications for transaction processing, Business Intelligence platforms for analytical reporting, and search engines for operational data access and reporting.




Fig 8: A Search Platform is an Ideal Tool for Improving Operational Decision Making

Case Study: GEFCO



Fig 9: Database Offloading and Operational Reporting for GEFCO (EXALEAD partner on project: ST Groupe)

GEFCO, one of Europe’s top ten transport and logistics groups, was dissatisfied with the extranet that had been created to help staff, partners and customers locate, track and optimize vehicle transport across 80 countries and 500 international routes. Their three main challenges were **1)** usability (the interface was highly complex), **2)** scalability (cost, complexity and perfor-

mance limited access) and **3)** data latency (with a latency rate of 24 hours due to batch scheduling of database updates).

In redesigning this logistics portal, they considered a traditional database-centered track and trace application and an SBA approach. After viewing the powerful capabilities of a prototype developed in only 10 days with CloudView, they decided an SBA strategy was the right choice, enabling them to:

- **Reduce Costs and Improve Performance through Database Offloading**

The new CloudView-powered vehicle track and trace application significantly reduced the load on GEFCO’s Oracle databases, enabling the company to cut data latency cut from one day to 30 seconds (though they decided 15 minutes was sufficient for their needs), and to significantly expand access to information critical for better operational decision-making. Their new platform uses only three servers (1 each for indexing, search and high availability) to provide search and quasi-real-time operational reporting for 600,000 daily transactions.

- **Achieve Agile Development, with a Rapid Time to Market and Development Independence**

Using the native CloudView ODBC Connector, the application was developed and deployed in only 60 days, a time to market unattainable with conventional application development. Development was also simplified by the engine’s built-in capacity to enforce strong firewalling of confidential client data at the search engine rather than application level. GEFCO now has an agile base upon which internal staff can rapidly construct new applications and services.

- **Improve the User Experience with Simpler, More Robust Functionality & More Timely Data**

In addition to a vastly simplified user interface, the SBA approach enabled GEFCO to offer natural language search and to provide timely operational reporting using the engine’s native classification and statistical computation capabilities. Users simply click on statistics of interest, and visual reporting tools such as pie charts are generated on the fly. And because with this type of application there is no constraint on the number of data attributes tracked, the system can effortlessly and organically scale in dimension and depth.

5.3 Comprehensive Business Applications

The SBA model is also serving as the foundation for comprehensive business applications that include database update functions along with data access operations (i.e., both read and write operations).

In this model, the search engine serves as the main data access

layer and provider of information access and operational reporting services within a framework that parallels a Web application. These comprehensive SBA applications are characterized by a rapid time to market, high performance, high scalability and extensive operational agility, and typically employ:

- A model-driven architecture (MDA)
- An agile development methodology
- Web formats, protocols and scripting languages

Case Study: National Postal Service

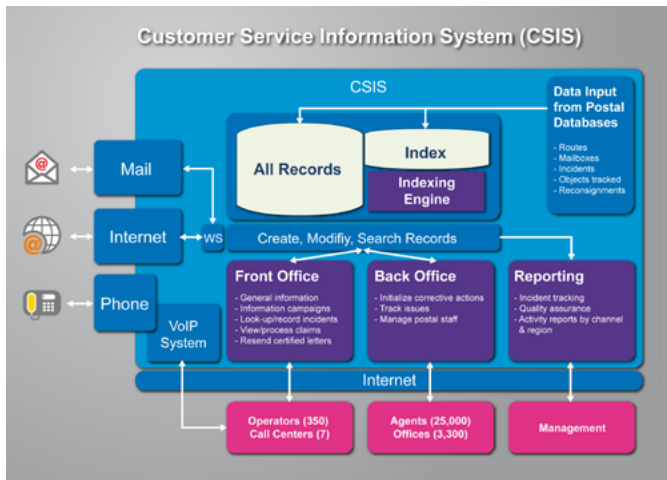


Fig 10: A New Customer Service Portal for a National Postal Service Architecture and schema by EXALEAD partner ARMOSC; systems integration by EXALEAD partner Capgemini

To meet the aggressive customer service objectives set forth in its strategic plan for 2008-2012, a national postal service in Europe needed to create a robust, unified application for responding to and analyzing customer inquiries and requests received by mail, Internet and telephone. Their ideal solution would provide a single platform for use by the Front Office (350 operators in 7 call centers), Back Office (25,000 agents in 3,300 mail facilities) and Management.

This Agency chose an innovative path for accomplishing this goal, with an SBA framework and extensive use of Web models and technologies. Developed using an iterative release strategy and guided by a model driven architecture (MDA), the new platform leverages:

- An open source infrastructure (Linux + MySQL + PHP, i.e., a 'LAMP' architecture)
- A Voice over Internet Protocol (VoIP) for the call centers
- Standard Web formats and protocols for managing data (REST & Atom RSS)
- A search engine, CloudView, as the main data access layer and provider of information access and operational reporting services

An Agile Solution with a Low TCO

The result is an agile, robust, and cost-effective system that provides stunning performance for a vast user base against a high volume of data. The application processes national requests spanning 36,000 postal districts in just a few seconds, with a response rate of less than 500 milliseconds for call center operators.

The SBA model made it possible to deploy an operational prototype in only three months

The agile, iterative development enabled by the SBA architecture made it possible to deploy an operational prototype in only three months, followed by four iterative releases over the next six months. This iterative approach resulted in a better final product, one that incorporates essential features and functionality unanticipated in initial planning.

This agility applies not just to the development process, but to the application itself. With the SBA approach, the Agency was able to meet its objective of leveraging one solution to serve three different audiences with three different needs: simple information search and retrieval and updating for the Front Office, complex queries and updates for the Back Office, and operational reporting for Management.

Operational reporting is easy with an SBA model

Because this latter service, operational reporting, is automatic and effortless under an index-based SBA model, the Agency has now decided to extend similar reporting capabilities to Back Office staff as well, capitalizing on the potential of operational reporting to improve day-to-day decision making by operational staff. In addition, they have decided to deploy the agile, cost-effective model used for their customer service portal (REST + LAMP + CloudView) for other complex, mission-critical projects, including a recycling management application.

6) EXALEAD CLOUDVIEW™ FOR DATABASE SBAs

EXALEAD is the leading infrastructure provider for SBAs. Much of this success as an SBA platform is due to CloudView's unique evolution: CloudView was designed from inception for both the Web and the enterprise. Because of this dual Web/enterprise DNA, CloudView alone provides the balance of corporate

functionality and Web simplicity, scalability and innovation essential for successful SBAs.

On the Web side, CloudView drives a 16 billion-page public search engine and serves 100 million unique monthly visitors through its public and customer installations. On the enterprise front, CloudView couples this performance and scalability with advanced semantic processing of unstructured data, superior structured data handling, and full compliance with existing security systems.

6.1 Scalability & Performance

CloudView is the only enterprise search engine designed from inception for multi-billion document scalability. More importantly, CloudView scales easily and cost effectively. The system is extremely resource efficient, supporting real-time indexing of 100 million documents and processing up to 20 queries per second on a single dual-processor server.

CloudView can index 100 million documents and process 20 queries per second on a single commodity server

Furthermore, with CloudView’s distributed architecture, scaling can be achieved simply by adding commodity processors or servers—no painful migration is required. This provides crucial business agility and continuity, enabling your business to scale serenely no matter how sharply or rapidly demand increases. (For client benchmarking data, please see our whitepaper *The Hidden Costs of Scaling Search*.)

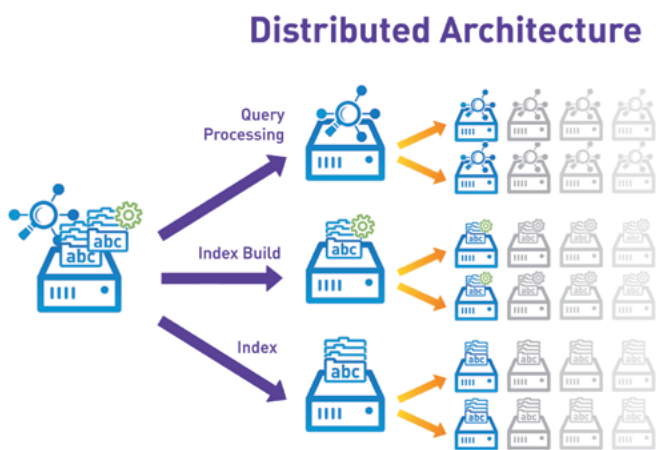


Fig 11: Unlimited, Effortless Scaling with Built-in Load Balancing, Partitioning and Replication

6.2 Agility

CloudView further provides the most agile platform on the market. Its service-oriented architecture (SOA) and extensive application programming interfaces (APIs) ensure:

- **Unlimited Data Flexibility**
The platform can connect to any internal or external source (CloudView provides out-of-the-box support for more than 124 languages, 320 data formats, and 50 types of databases). In fact, CloudView is so adept at manipulating all types of structured data that some clients are finding their new ‘database application’ doesn’t need a database at all: with an SBA, they can efficiently exploit structured data like logs, lists, Excel spreadsheets, XML files, etc. in their existing format, simplifying development even further.
- **Unlimited Access Flexibility**
CloudView supports fuzzy natural language search as well as structured queries (XML, XPATH, XQUERY, etc.), and offers an API for complete presentation versatility (bar charts, graphs, maps, real-time operational reports and dashboards in addition to the default hyperlinked lists and categories).
- **Rapid Time to Market**
CloudView features a non-intrusive ‘plug and play’ design that enables rapid deployment and leaves existing information systems undisturbed. This design is complemented by open APIs and support for standard Web formats and protocols (SOAP, REST, XML, RSS, RDF, OWL, etc.), making it an ideal platform for iterative, agile development and providing easy integration with third-party applications or services.
- **Unmatched Scalability, Performance and Availability**
CloudView’s fully distributed architecture enables easy process distribution, load balancing, index partitioning and index replication, ensuring the platform can grow and adapt painlessly and inexpensively as demands and business requirements change.

6.3 Usability

CloudView’s fully customizable interface combines the speed and simplicity of Web search with the rich output of a structured enterprise application. The result is immediate intuitive use, more successful search, easier content exploration, and exceptionally high adoption rates. Features include:

- Navigation of results by categories dynamically extracted from the database, with automatically generated options for drilling up or down within results.

- Natural, 'forgiving' search, with advanced semantic processors that can interpret imprecise requests and offer spelling corrections, close matches, and related content.
- Generous visual aids with statistical processors that can generate charts, graphs, reports and dashboards on the fly in response to user queries.
- And, for applications merging database and unstructured content, at-a-glance scanning of results with extracts, file type icons, thumbnail images and application-independent content previews with search term highlighting.

Easy to Administer

This ease of use continues into the back end, with a Web-based console for all management tasks: user interface configuration, indexing control, search performance tuning, etc.—tasks additionally exposed via an API for maximum flexibility. CloudView's engineers have also strived to make the platform as self-maintaining as possible:

- All of the natural language processing modules (dynamic categorization, spell-checking, spelling suggestions, etc.) evolve automatically and in real-time as your data evolves.
- Index updates are also automatic, real-time and incremental. Documents can be dynamically added, removed, or replaced in the index, and a single field of a document can be updated without modifying the content of other indexed fields.
- Dictionaries are likewise fully automatic, incremental, and real-time. No hand-built dictionary or manual assistance is needed.
- The index build components are fully distributed and designed to run 24 hours a day, 7 days a week, without any human intervention. They also automatically perform maintenance tasks such as removing references to deleted documents.

CloudView offers an ease of use for end users, developers and administrators that has earned it a 100% customer loyalty rate

These 'self-maintenance' features not only make administration easier and more pleasurable, they significantly reduce administrative labor costs, lowering TCO and augmenting ROI.

7) CONCLUSION

SBA's are redefining information access because they are meeting:

- End users' need for fast, intuitive information access
- Businesses' need for better operational decision making and competitive differentiation
- IT's need to simplify operations, gain agility, and reduce costs

CloudView is the information access platform of choice for SBA's for databases because of its superior handling of structured data, unlimited scalability, low TCO, high performance and open, Web-standards compliant architecture.

CloudView is also the SBA platform of choice for SBA's integrating structured and unstructured content (such as email messages, Office documents, presentations, Web pages, blogs, forums, and RSS feeds). CloudView can automatically structure staggering volumes of such unstructured content and meaningfully synthesize it with structured data.

Enterprises are leveraging this capacity to extend enterprise applications like Customer Relationship Management (CRM), Supply Chain Management (SCM) and Business and Competitive Intelligence (BI and CI) with important emotive and qualitative data from unstructured sources. Online businesses are harnessing it to create innovative, high performance mashups that merge content and functionality from diverse sources such as databases, mapping services, business applications and the Web.

Call us today to discover what a CloudView SBA can do for your business.

Delivering Best-in-Class Products



Virtual Product Design



Virtual Planet



3D for Professionals



Information Intelligence



Realistic Simulation



Dashboard Everything



Virtual Production



Social Innovation



Global Collaborative Lifecycle Management



Online 3D Lifelike Experiences

About EXALEAD

Founded in 2000 by search engine pioneers, Dassault Systèmes EXALEAD® provides search and unified information access software that drives innovation and performance in the enterprise and on the Internet. The company's EXALEAD CloudView™ platform is the industry's most sophisticated and scalable infrastructure for Search-Based Applications (SBAs), with over 30,000 business decision makers, half a million enterprise search users, and 110 million Internet users relying on EXALEAD to make their information universe accessible and meaningful.

EXALEAD EMEA

Dassault Systèmes
10 place de la madeleine
75008 Paris
France

Visit us at
3DS.COM/EXALEAD

