

Dieselpoint

© 2013 by Stephen E. Arnold, www.arnoldit.com

Dieselpoint offers a Java-based search system. The system implements a quantum random walk architecture, one of the more remarkable descriptions of search.

Author's note: This is an unpublished, preliminary draft of a description originally destined for a report published by Gilbane, now a unit of Outsell.

Dieselpoint is important for three reasons. The company was among the first to emphasize the value of Extensible Markup Language and Java as the core of a search system. For a period of time, Dieselpoint also offered a viable option for those considering the comparatively expensive Endeca and MarkLogic systems. What's more, Dieselpoint has been an advocate of no-frills marketing. The strategy keeps costs down, and the tactic ensures that Dieselpoint is often overlooked by procurement teams. As other vendors sucked in venture funding and spent it on online advertising, for-fee consultant endorsements, and moon shot sales pushes--Dieselpoint has just chugged along.

After its founding, the company rated positive comments from those who were true believers in Extensible Markup Language and its close cousins. The Dieselpoint Web site has been updated, but the "new" cloud service seems to be a no show. When I suggested that the company had gone quiet, the founder contacted me and said, "We are very much in business." Okay. Great. This Xenky description of the system dates from 2000 to 2006. The company's Web site is www.dieselpoint.com. More recent information may be available directly from the company.


This information is a rough draft and is frozen.

Introduction

Chris Cleveland, a political science undergraduate with an MBA, launched Dieselpoint in January 2000. The system evolved from software he had developed at Genesee Development, then a Lincoln Park, Illinois-based industrial business technology firm that Mr. Cleveland had founded in 1990. Mr. Cleveland sold Genesee’s consulting arm in order to set up Dieselpoint and market the firm’s information retrieval system. The system is available as Dieselpoint Search, a not-yet-released cloud-based solution called Engineez, and an open source content acquisition system named Open Pipeline.

The name *Dieselpoint* refers to the strength and power of what the company was trying to create. One engineering textbook describes the 'diesel point' as the point at which the combination of heat and pressure in a diesel engine causes ignition. Dieselpoint intends for its software to ignite possibilities for its customers too.

Table 1: Dieselpoint: A Bird's Eye View

| Product Thumbnail | |
|---|---|
| 1 Search Brand | Dieselpoint Search. The hosted version of the system is Engineez. Dieselpoint offers a Database Appliance Product. |
| 2 OS Supported | Windows, Web Based, Mac, Open Source, Linux/Unix, and any system that runs Java |
| 3 Est License Fee | Dieselpoint Search Enterprise, Starts at \$20,000 for a two CPU installation Engineez Cloud Service, contact the company for availability and prices |
| 4 Functions | Dieselpoint offers Boolean search, categorization, and support for structured and unstructured information |
| 5 Claimed Features | The system permits the generation of “dynamic FAQs” for customer support applications, a “self-learning” function, natural language search, and “guided problem solving.” |
| 6 Downsides | The system requires appropriate resources to perform content processing, index updating, and query processing. |
| 7 Similar To | Endeca, Exorbyte, MarkLogic |
| Product Close Up  | Dieselpoint Search is search and navigation software for enterprise data including document collections, databases, and XML. The system can search unstructured documents, semi-structured XML or fully structured SQL databases. It can work like a Google-style search engine and perform a full-text search and also as a SQL database and implement parametric searches. Coded in Java, the system includes full-text search syntax including linguistic tools. the firm's Web site is www.dieselpoint.com |

The product has evolved into a solution for applications that require full-text, navigational, and parametric search.¹

Mr. Cleveland told me:

“Enterprise search is not the same as Web searching, because it entails all of the nitty-gritty preparation for search – that is, it requires doing all of those things you need to do to get a document and standardize it before indexing. Open Pipeline aims to streamline the preparation process through its innovative document-processing capabilities.” - Chris Cleveland, CEO of Dieselpoint

Providing high performance faceted search and navigation for terabyte-sized datasets with millions of items is what Dieselpoint was designed to do. However, we are often selected because we are 100 percent Java with elegant and open APIs.

Rumor has it that Google did some tire kicking several years ago as part of Google’s initial interest in enterprise search. Then in a Googley way was distracted and moved forward without acquiring Dieselpoint.

Dieselpoint asserts that it has some high-profile clients with plenty of search experience under their belts. The implication seems to be that these customers have licensed the Dieselpoint technology and dropped their licenses for other, less effective advanced text processing vendors.

A Theoretical Swiss Army Knife

Dieselpoint is a Swiss Army knife system. Like the MarkLogic Server product, Dieselpoint can provide search, slice-and-dice, and analytics from content stored in its proprietary data management system. For example, Dieselpoint allows a licensee to manipulate document attributes as well as document text. The system can be used like a facet-capable search and retrieval system. An organization can acquire, index, and retrieve information in unstructured documents, semi-structured XML, or fully structured SQL databases. Like Sphinx Search, Dieselpoint can put a friendly face in front of a user so database stored information can be accessed with simple English queries, not complex command syntax.

The approach, therefore, leverages search, metatagging, indexing, and other rich text processing functions. Dieselpoint is “just like” Autonomy, Endeca, Fast Search & Transfer, and other high-profile search systems. The company delivers scalable search and faceted navigation, which enables search results to be ordered and classified in multiple ways based on matching attributes. Result sets of any size can then be navigated using dynamically generated menus. The “hot links” are generated from the underlying document attributes or metadata. The idea is that these “hot links” give users context-dependent browse capability. A user can see options available at each step in reviewing query results.

Dieselpoint asserts that its software is used in a variety of applications including e-commerce, document search, site search, product life cycle management, enterprise content management, and in other vendors’ solutions. Information Builders integrated Dieselpoint into its data management system.

¹. This bit of jargon surfaced in the late 1990s. The idea is that key word search matches strings. A *parametric search* can query a database engine or index so that the metadata of the underlying content is used for retrieval. By 2010, *parametric* implied a natural language query passed against a database index so that no structured query language syntax was required.

Dieselpoint provides text processing and search for enterprise-class applications. System licensees are currently using Dieselpoint for XML search, PDF search, catalog search, and Intranet search, and building search-based applications.

Management

The lean management of Dieselpoint consists of:

- Chris Cleveland, Founder
- Andrew Beiman, President, formerly the firm's marketing manager
- Irena Matveeva, contractor and author of "Generalized Latent Semantic Analysis for Document Representation."²

Clients

The firm's clients include:

- Albania
- City of St. Johns
- Cognitor
- Cygage
- Datech Soft
- Digi Tech
- Epsilon Systems Software
- FMC Energy Systems
- ITT Industries
- LinkTier Technology
- McGraw Hill Companies
- Newview
- Oceaneering
- Ohio State Health
- PSS World Medical
- SPX Corporation
- Porsche
- Prometheon

². A book version of her dissertation is available at <http://goo.gl/92C8FG>. ISBN-10: 1243453133.

- Quilters
- Risk Management Technologies
- Sony Australia
- Tasmania
- TrustNet
- Valley Forge Financial Information Services.

Partners

Integration services are available from Raritan Technologies. Cognitor appears to use the Dieselpoint search system in its Enquire knowledge management tool.

Dieselpoint in Action

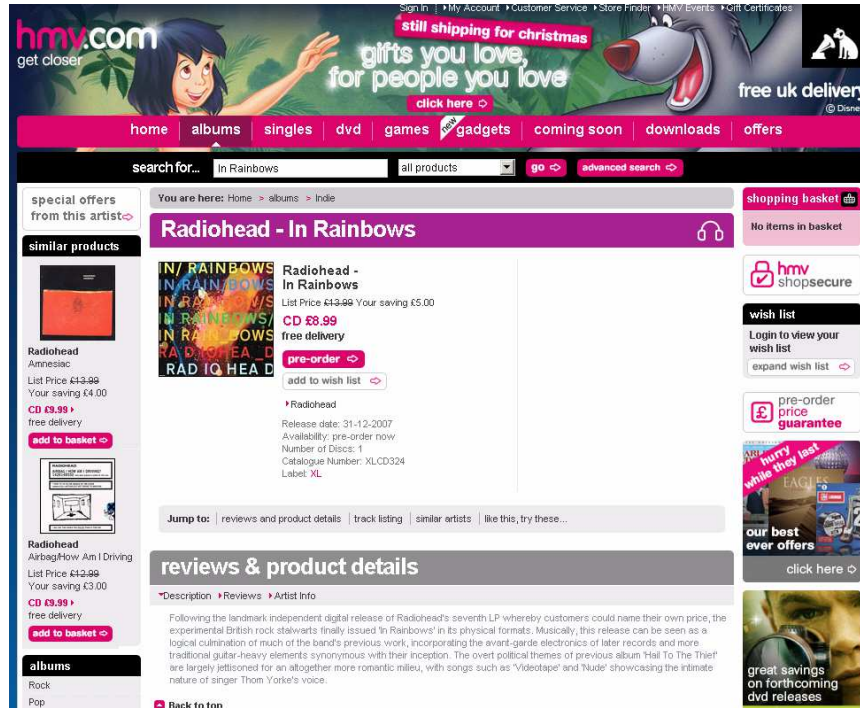
In the 2007 release, Dieselpoint enhanced support for taxonomy-driven searches. In search software terminology, a *taxonomy* is a knowledge model organized into a hierarchy of major and minor concepts.

HMV

One licensee, HMV, a U.K. based retail operation, uses Dieselpoint for “faceted navigation. The idea is that a visitor to an online music site can drill-down through categories (for example, song genre) to find results, rather than just doing a text search. In the HMV implementation, song lyrics might be categorized by genre (for example, rock) and each genre may include sub-categories (for example, soft rock, classic rock, etc.).

HMV users could navigate large information spaces without feeling lost. The HMV interface guides the user toward Amazon-style recommendations. The end result was intended to boost HMV’s online sales. Because Die-

selpoint was written in Java, the HMV implementation was accessible from HMV's mainframe-based point-of-sale terminals.³



The HMV implementation of Dieselpoint includes structured and unstructured data presented in an interface with a search box and “faceted navigation”.

AC Nielsen

AC Nielsen (a unit of the Dutch publishing giant VNU) allegedly shifted from Autonomy IDOL to Dieselpoint in 2007. Daniel Morse of DPM Technologies, Inc., the integrator working on this “rip and replace” project, said at Enterprise Search West Conference in November 2007:

Dieselpoint exceeded our expectations. We implemented facets, clustering, integration, and reporting on time and at a lower cost than the Autonomy solution.

At that time, AC Nielsen searched three million documents in a Vignette content management system repository. The system was implemented quickly in comparison to a failed two month Endeca implementation effort.

³. As we update this rough draft, HMV has faced some financial challenges and made changes to its Web store front.

Northrop, a leading US defense contractor, deployed Dieselpoint's easy-to-implement solution for a large parts database. More than 1,000 users rely on Dieselpoint's system to find information required for projects.

Technology

A Dieselpoint white paper, written in 2004, described the basics of the firm's search technology.

Dieselpoint search is written in Java. The firm's engineers have found a way to deliver performance that "outperforms code written in other programming languages, such as C++."⁴ The company's Java capabilities allow Dieselpoint to embody "advanced full-text and linguistic techniques with high-performance data navigation to yield intuitive user interfaces not possible with other approaches."

Java permits Dieselpoint to be deployed on Intel, Sun, and IBM hardware running Windows, Linux, Solaris, IBM iSeries, and AS400, and Macintosh operating systems. The search system is "fully standards compliant and can rely on a large body of third-party tools and components." The white paper does not address the security and code management issues sometimes associated with Java-based systems.

According to the company, Java simplifies implementation and enables interoperability within a range of environments.

Dieselpoint's performance addresses a number of challenges that are common with some competitors' products, namely, Dieselpoint's system:

- Processes "large volumes of data"
- Operates in a "highly multi-threaded environment" and deals seamlessly with locking and resource contention issues
- Deals with string manipulation issues
- Manages effectively large memory buffers.

The system incorporates a number of performance-enhancing methods; for example, the Java-based system optimizes disc access. If sufficient resources are available, the indexes can reside in memory. By eliminating spinning disc seeks, Dieselpoint delivers more responsive query processing than some competitors' systems.

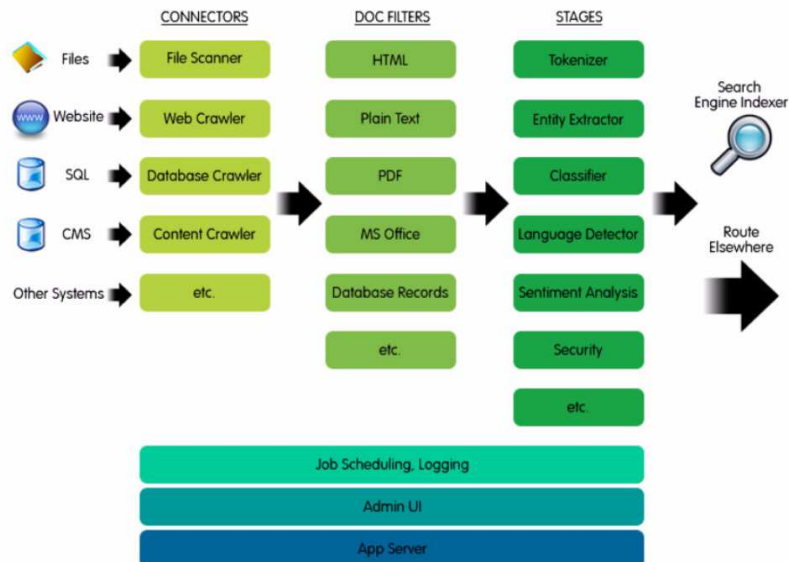
Dieselpoint, like some other XML-capable systems, makes use of data compression. On modern systems, the overhead for compression and decompression is negligible.

⁴. "A White Paper on Java and Performance: Is It Possible to Use Java to Design High-Performance Software Applications?" December 2004, page 2.

The system uses a proprietary “unsynchronization” method to eliminate the overhead associated with less-sophisticated Java synchronization methods. The Dieselpoint system also reuses memory buffers and implements its own heap profiler.

The system delivers near-real-time incremental index updating. The system indexes metadata and field attributes as well as text and documents. Dieselpoint's functionality is similar to that delivered by Endeca and Siderean Software, the latter a vendor of semantic technology that went out of business in 2008. Users are able to explore search results by pointing and clicking on categories.

The system understands numbers, because Dieselpoint permits parametric queries. Therefore, a licensee finds it trivial to allow a user to click on a price range to generate subsets within a specific range and perform additional calculations such as generating a monthly payment estimate.



Like Vivisimo, the company provides federated search to its licensees. Dieselpoint's content acquisition method processes data from different sources. The index makes it possible to run a query across the processed content.

The basic architecture of the Dieselpoint search system is similar to that used in most competitors' software. The key difference is Dieselpoint's use of Java.

Dieselpoint will run in any J2EE-compliant application server. It is designed to be called from a user-written application and its API is designed with such applications in mind. For example, it returns metadata about search results so applications can dynamically create user interfaces relevant to those results. Applications can call Dieselpoint through a Java API, a JSP front

end, JDBC, or XML. For users who do not want to write their own application, Dieselpoint ships with a number of sample applications (including a product catalog application) and a generic, JSP-based user interface that is “suitable for common uses”.

Indexing and Search

Most enterprise search systems were not designed to handle both unstructured and structured information, but Dieselpoint was. As a result, some of the company’s indexing and search capabilities are suited to collections of content consisting of documents, email, and information stored in traditional databases and data management systems.

Indexing

Dieselpoint indexes documents and data specified by the user and then permits queries against those indexes. Index and configuration files reside in a single directory structure, making it easy to move and make backups of indexes.

Indexes are stored in a Dieselpoint-proprietary file format. No external database is required for the system. The file format and index structure is fault-tolerant, requiring no rollback or recovery procedures after a server crash. After an unexpected server reboot, a Dieselpoint Search application will simply pick up where it left off, eliminating the time consuming and error-prone index recovery procedures required by systems like Fast Search & Transfer. Unlike some resource-heavy solutions, Dieselpoint Search integrates easily into enterprise applications and custom Java applications. By default, no configuration is required at all, although several hundred configuration and indexing options are available.

Dieselpoint indexes documents and data retrieved by a crawler from Web sites, directories, and databases. It can index documents (XML, HTML, PDF, Microsoft Office), databases (via JDBC), and flat files (comma-separated, tab-separated, and so on). Data in other formats can be indexed via calls to a user-implemented API.

“Dieselpoint is very intuitive for the users. Some people like to search for keywords, and some like to browse. We accommodate both of those users”
- Chris Cleveland,
CEO of Dieselpoint

The indexing subsystem extracts data in the form of attributes, such as document metadata, XML elements, and database columns. A preprocessor, if implemented, allows user-written code to modify, categorize, or reject items before they are indexed. Dieselpoint can make use of controlled term lists (dictionaries) if available.

Dieselpoint asserts that its approach permits low-latency incremental indexing. The system can index databases as full-text and includes field context.

“The importance of investing in scalable software and hardware redundancy is rarely more apparent than during holiday traffic surges” said Chris Cleveland, CEO of Dieselpoint, Inc. “Dieselpoint Search’s speed, scalability and small footprint enabled HMV to handle a rapid increase in users efficiently and at extremely low incremental cost.” - Chris Cleveland, CEO of Dieselpoint

Search

Dieselpoint uses a proprietary query language, which supports full-text and parametric searching. (*Parametric searching* limits a search to a particular attribute, such as a title, part number, or description.) Search clauses can be joined in any way by AND, OR, NOT, and parentheses, and can include comparisons (=, >, >=, <, <=, <>), wildcards, and regular expressions. Full-text features include stemming, thesauruses, stop words, misspellings, relevance, and hit highlighting. Search results can be returned as a JDBC result set or XML document and can be sorted by relevance or attribute value. The user-facing search results may be customized to meet the needs of the licensing organization.

Search features include searching by element or attribute and by XML path. (The indexer preserves the XML hierarchy.) The query engine can return complete documents or portions of documents most relevant to the user’s query. The system can treat fragments of a document (headed by a particular element name) as separate documents if required. Dieselpoint understands both ECCMA (an XML language for catalogs) and Dublin Core and provides special processing for both. In addition, it can handle XMP metadata (RDF documents) embedded in PDF documents.

One useful function of Dieselpoint is its display of hot links to content related to the user’s query. This use of “faceted navigation” makes it easy for an unsophisticated user to browse content related to the query. When the user enters a free text or Boolean query, the system can offer search suggestions.

In the Dieselpoint results list, the user’s key words from the query can be highlighted. The system can be set up to display document summaries. If the user enters a parametric query (text string with no explicit switches), the Dieselpoint system administrator can configure the system to show the user values and counts for the structured content.

The user can sort results displayed in a table. Dieselpoint uses its index and Dieselpoint-generated metadata to identify relevant data. Non-matching content is not displayed. In theory, the results of a query can be more easily scanned and compared by the user. The system can be set up to perform basic report generation.

Other Dieselpoint Features

Dieselpoint, like other search vendors, claims a laundry list of features, functions, and components. The Dieselpoint Web site lists more than 60 Dieselpoint features. Many of these are characteristic of most enterprise search systems. Of particular importance is the inclusion within the product of an administrative interface, sample applications, and a bundled Java Server Pages servlet container and application server. The Dieselpoint system includes a Web server component and associated container scripts.

Let me highlight a handful of Dieselpoint Search features.

Administrative Interface

Dieselpoint includes an administrator interface for performing such tasks as managing indexes, defining data sources, and scheduling the crawler. It also contains a Web server and servlet container. Many administrative functions can be handled through a graphical interface; for example, relevance tuning and setting indexing aggressiveness.

The administrative interface makes it relatively easy for most licensees to create, update, and search indexes. Wizard-like screens capture requirements and write them out to configuration files formatted as XML. The search engine uses these XML files to index data and define system behavior.

A Wide Open API with No Dependencies

The Dieselpoint application-programming interface allows a developer to create a search-based application or integrate Dieselpoint into other enterprise solutions. The API allows a developer to extend the basic Dieselpoint system. The core Dieselpoint Search engine is implemented as a Java library. The product ships with a pre-built eCommerce subsystem. A developer can examine this component to get ideas for other ways to extend the Dieselpoint system. The source code for this application is included with the basic Dieselpoint package.

Dieselpoint uses a taxonomy data type that allows a developer to exploit interrelationships between and among content attributes expressed as metadata. The system includes redesigned internal indices for taxonomy attributes that are automatically generated.

Analytics

The system includes a range of analytic functions. Standard reports include system usage and query data. Other analytic functions can be integrated via the SDK, or from third-party analytics tools from Visual Sciences (formerly Web Side Story) or other vendors.

Language Support

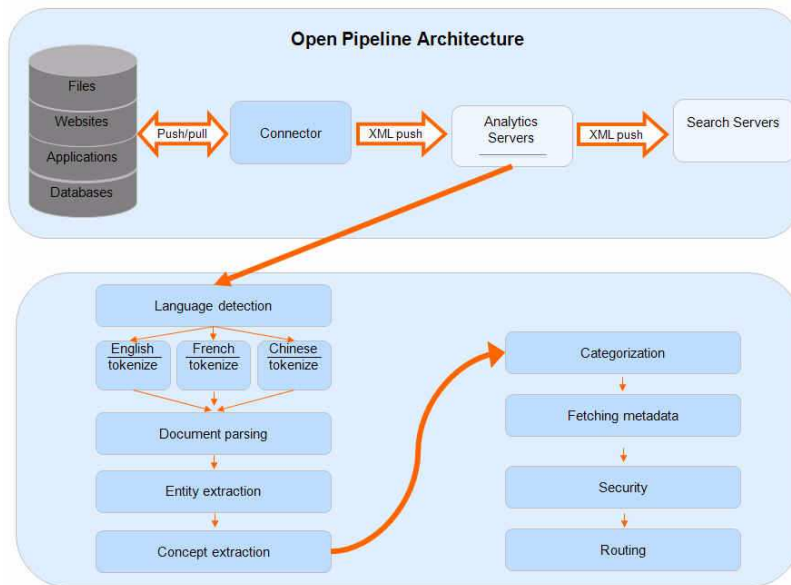
Dieselpoint provides support for more than forty languages and “140 dialects.” In addition to Spanish and other Romance languages, the system can process Chinese and Japanese. The system supports Unicode and can handle queries with or without diacritics. The system includes linguistic stemmers for Danish, Dutch, English, Finnish, French, German, Italian, Norwegian, Portuguese, Russian, Spanish and Swedish. The system includes stop word lists for English, French, Spanish, and a handful of other languages.

Marketing

Dieselpoint survives without the blue-chip marketing efforts associated with venture-funded firms like ElasticSearch and LucidWorks and the recently acquired search vendors like Autonomy and Endeca. Dieselpoint flies below the radar. The company has a basic Web site and the information is focused on identifying features. The company has used open source software as a way to reach potential licensees.

Open Pipeline

Dieselpoint carved out the content acquisition subsystem from Dieselpoint Search and released the subsystem as open source in 2008. The purpose of the open source initiative is to allow those interested in search to solve certain integration problems with Dieselpoint innovations.



The idea is that Dieselpoint uses a services-oriented architecture and adaptors. A content source “plugs into” the Dieselpoint service layer. The content then moves through Dieselpoint’s rich text processing subsystems.

Open Pipeline, therefore, is a way to encourage developers to use the Dieselpoint search engine. Open Pipeline is free, and Chris Cleveland said, “There is no downside. This is a way of getting high-quality software, commercially-sold and supported, while solving the interoperability problem.”⁵

The system makes it possible to tap into the power of Dieselpoint’s content acquisition system and have the option to hook seamlessly to the Dieselpoint search system.

⁵. “Open Pipeline: Open Source Software for Crawling, Parsing, Analyzing, and Routing Documents,” 2008.

Dieselpoint's open source Open Pipeline middleware is agile. Open Pipeline can be used for content acquisition (crawling), parsing, and routing documents through a workflow.

The idea is that a layer of software allows different content sources to be “plugged into” the search system. Some vendors like TIBCO call this an information bus; others refer to it as a framework. Regardless of the terminology, the search system using this technique can crawl data from a variety of sources, process it, and route it quickly and easily. Systems that make it possible to obtain content from different sources are sometimes described as federators or federating search engines. The idea is that the search engine's interface allows access to content from multiple sources from a single interface. Search engines that index a single server's content are finding it difficult to make headway in the crowded search and rich text processing market.

Open Pipeline implements a publish and subscribe model for data feeds. Subject to security and access rules, the system allows an authorized user to subscribe to feeds that support such standards as HTTP, Atom, and RSS.

Throughput

By opening up the process of analyzing, representing, and routing data, Dieselpoint functions as “middleware for search”. One twist in Dieselpoint's implementation of federation is that the company has engineered parallelization into the content consolidation function. This change increases system throughput, which can hit content processing in the 500+ megabytes per an hour range when properly resourced and configured. The Dieselpoint approach also includes replication so that queries do not choke the system when indexing and query processing hit peak loads. Keep in mind that these types of features work only if you are running appropriate hardware with adequate bandwidth.

Connectors

Dieselpoint arrives with a number of software connectors or filters. The system includes a JDBC connector to databases; for example, IBM DB2, Microsoft SQL Server, MySQL, Oracle's database, and PostgreSQL, among others.

| | |
|---|--|
| File Scanner | Scans a file system and processes files |
| SQL Crawler | Crawls tables in a SQL database, handles inserts, updates, and deletes |
| Web Crawler | Crawls websites, follows links, processes pages |
| JCR Connector | Scans the contents of a JCR repository. |
| Sharepoint Connector | Scans the contents of a SharePoint repository. |
| Documentum Connector | Crawls a Documentum repository |
| Exchange Connector | Crawls a Microsoft Exchange server. Returns both documents and email. |
| FileNet Image and Document Services Connector | Crawls a FileNet Image and Services repository |
| FileNet P8 Connector | Crawls a FileNet P8 repository |
| Hummingbird DM Connector | Crawls a Hummingbird DM repository |
| IBM Lotus Notes Connector | Crawls Lotus Notes. |

The current version supports content in repositories or indexes from:

- ANSI standard databases or proprietary databases that can export comma separated value and other common file types
- Documentum
- Linux, OS/400, UNIX, and other common file systems.
- Vignette.

Dieselpoint includes support via the Java Content Repository Standard for:

- FileNet
- Interwoven, now a unit of Hewlett Packard Autonomy
- Lotus Notes
- SharePoint
- OpenText.

Dieselpoint's SDK makes it possible for a licensee to create other adaptors, filters, and connectors as required.

Enginez

A second marketing initiative is the cloud-based version of Dieselpoint called "Enginez." As this draft is moved to the Xenky Web site in October 2013, the Enginez service has a low profile.

In 2011, Mr. Cleveland announced the Enginez service. At that time, the system was positioned as the quickest way to deploy a search-based application in the cloud. The video about the new cloud service began with the statement: Enginez is the most exciting application you have ever seen."⁶

ArnoldIT Opinion

Dieselpoint was visible in the enterprise search and eCommerce markets in 2004 to 2007. The company "went quiet" in 2008 and has maintained a low profile when compared to the high-amplitude marketing from Attivio, IBM Vivisimo, and Microsoft.

In the expansionist phase of the company, Dieselpoint's key advantage was that the system could index content in almost any Extensible Markup Language "flavor." Dieselpoint's parser could make use of tags and index the textual content of an XML document. For its search results, Dieselpoint could be set up to display the portion of the document that was most relevant to the user's query.

⁶. Source: Dieselpoint video at <http://goo.gl/8LZHq6>

Table 2: Dieselpoint Checklist

| Attribute | Dieselpoint Asserts | ArnoldIT Comment |
|--------------------------------------|---|--|
| 1 Platform | Most operating systems | |
| 2 Key word search | Yes | |
| 3 Text mining | No | |
| 4 Automated indexing | Yes | |
| 5 Personalization | No | A personalization subsystem can be implemented via the API |
| 6 Workflow | No | |
| 7 Interface | Graphical and customizable | |
| 8 Hosted service | Yes | The Engineez service |
| 9 Administrative interface and tools | Yes | |
| 10 Application programming interface | Yes | The Dieselpoint system is Java centric |
| 11 Professional services | Check with Dieselpoint | |
| 12 Security | Username/password for crawling restricted sites | Via the API most security methods can be implemented |
| 13 Connectors | Most common document types | Via the API additional content types can be supported |
| 14 Support for structured data | Mainstream databases and MySQL | |
| 15 Relevance ranking | | |
| 16 Video | No | |
| 17 Federated search | Yes | |
| 18 Fielded search | Yes | |
| 19 Content crawler | Yes | The Open Pipeline open source component is the Dieselpoint content acquisition subsystem |
| 20 Price | Begins at \$20,000 for a two CPU installation | |

Before MarkLogic positioned itself as the leader in XML-centric content processing, Dieselpoint offered content acquisition, content processing, and search well suited to XML.

Anticipated Benefits

The core Dieselpoint Search engine is implemented as a Java library. It has a simple, intuitive, but extremely powerful API. Index and configuration files

reside in a single directory structure, making it easy to move and take back-ups of indexes.

The product ships with three ancillary modules: an administrative interface, sample applications, and a bundled JSP/servlet container/ app server suitable for common uses.

Another plus is that Dieselpoint's administrative interface makes it easy to create, update, and search indexes. Wizard-like screens capture requirements and write them out to configuration files formatted as XML. The search engine uses these XML files to index data and define system behavior.

Unlike other resource-heavy solutions, Dieselpoint Search integrates easily into most enterprise applications. If an organization does not have an enterprise application, Dieselpoint's system can be used to support a customized search-based application.

Upon installation, Dieselpoint does not require much, if any, user configuration. Keep in mind that Dieselpoint offers hundreds of configuration and indexing options. The system will run on any system that supports Java. Have an IBM-compatible mainframe? Dieselpoint will run on that system.

Other upsides for Dieselpoint include:

- The system delivers a good balance of flexibility, support for enterprise content sources, and performance
- The Dieselpoint Open Pipeline engine makes it comparatively easy to access content from many disparate sources. Unlike federating tools from such vendors as Vivisimo, Dieselpoint minimizes "script fiddling" or can allow a developer to make extensive modifications. The choice is the licensee's.
- Dieselpoint provides an alternative to the SAP-style approach of reworking the infrastructure of the organization in order to accommodate a new search system.

Possible Drawbacks

Dieselpoint maintains a low profile. As a result, the system may be overlooked by some potential customers or rejected by procurement teams seeking the financial and technical resources of a company like Hewlett Packard, IBM, or Microsoft.

One consideration is that indexes are stored in a Dieselpoint-proprietary file format. No external database is required. The file format and index structure is fault-tolerant, requiring no rollback or recovery procedures after a server crash. After an unexpected server reboot, a Dieselpoint Search application will simply pick up where it left off, giving you peace of mind that your application is reliable.

Dieselpoint's roots anchor in eCommerce and online directory-type applications. The extension of Dieselpoint to other applications is possible, but the licensee may have to write original code.

The downside associated with Dieselpoint includes:

- The company appears to rely on word-of-mouth and presentations by the president to generate buzz
- Some licensees may find that custom scripting is required to support their organizations' security systems
- The Java-centric approach increases flexibility, but to get the most out of the Dieselpoint system, appropriate resources are necessary
- The low profile of Dieselpoint may make it difficult for a licensee to find a local integrator to handle tuning and customizing tasks.

Dieselpoint can perform this work, but the company has a small staff, so you may have to wait to get access to technical specialists. Dieselpoint does assert, however, that they have integrators readily available in both the United States and the United Kingdom, including, but not limited to DPM Technologies, Raritan, and Bluetab.

Conclusions

Dieselpoint is one of those text-processing solutions that delight customers savvy enough to find the vendor that supports XML, embraces Java, and provides a Swiss Army knife's versatility. Dieselpoint has a low profile, which inhibits its market penetration. The company's technology delivers good value when compared to better-known competitors such as Autonomy, Endeca, and Fast Search & Transfer. For ecommerce, parametric, and federating applications, Dieselpoint can deliver comparable functionality, scalable performance, and extensibility at a reasonable cost. The open source content acquisition system is a great value and available without charge.

Stephen E Arnold

Minor edits to a rough draft on October 30, 2013