BA Insight

Modernizing enterprise search with hybrid index and search time merge

How BA Insight SmartHub brings web-like search to the enterprise

What is Federation?

> **Federated search retrieves information from a variety of sources via a search application built on top of one ore more search engines.**
>
> **Alexis Linoski and Tine Walczyk**
> Federated Search 101
>
> *LibraryJournal.com*

Federated search, according to Alexis Linoski and Tine Walczyk in their *Federated Search 101* article, "retrieves information from a variety of sources via a search application built on top of one ore more search engines."

Users start by making a single search request. That search is then extended to various search engines and data sources included in the federation. From there, the federated search pulls all the results together, presenting them to the user. Federated search can be used to integrate disparate information resources within a single large organization or for the entire web.

From a technology perspective, there are two approaches to search labeled as index-time merging, and query-time merging. In general, vendors and analysts in the enterprise search/insight engines market have advocated index-time search and have positioned query-time merge as something that can be used but is not recommended.

So, the approach has been an either or paradigm where they define use cases in which one is the right solution. They are effectively heavily-biased to index-time merge, which in reality is what the vendors provide or at least promote heavily. It creates the view that one approach should be chosen over the other.

This school of thinking removes a valuable approach from potential implementation strategies that has been proven to be the correct approach in many scenarios. Single index search is a solid strategy; however, it is not the only viable solution.

The reality is that with the continued growth of data and availability of new, open and flexible search engines such as Azure Cognitive Search, Amazon OpenSearch, and Elasticsearch, along with modern applications, and robust APIs, it is possible to have a hybrid approach as opposed to an either/or approach.

9hrs

how much time employees spend weekly searching for information.

Today, the one-size-fits-all approach that has worked for the last decade won't work in many situations. There are times that business requirements are best served using a single-index approach and others that will be best served using a hybrid indexing approach, where the solution combines the index time and query time merging. Our view is that hybrid indexing is becoming the defacto approach for deploying enterprise search.

An interesting paper to read is Google's research paper on the journey around indexing the world wide web. Abhishek Das and Ankit Jain of Google summarize:

> **"Why can we not have multiple indices—bucketed by rate of refreshing? We can and that is what is standard industry.**

What Google establishes as standard industry practice for the web is the concept that search indices should be based on the business drivers and needs of the content that will go within them.

Each index serves a purpose and works in tandem with all other indices, also serving specific purposes. Purposes range from speed of refresh, necessary metadata enrichment, geographical distribution, performance spread, etc. Within enterprises, the same reality exists.

Each application has its own index for its own purposes, and where possible we should make them work in tandem with each other without requiring the crawling and indexing with the infrastructure and support costs it brings.

Why does Query-Time Function have such a bad reputation?

Historically, organizations that attempted this type of federated search via query-time merge could run into challenges that resulted in poor results.

As an example, performance of query results is often raised as an issue of federation. A decade ago, query-time indexing of multiple search indices caused performance and user satisfaction issues. This still holds true for older applications, but not so for many modern applications, especially search engines like Azure Cognitive Search, Amazon Kendra, Amazon OpenSearch Service, Elasticsearch, Microsoft Search, SharePoint Online SharePoint On-Premise, and Solr.

Outdated federated approaches waited for the slowest index and then returned all combined results.

Modern federated approaches now return results in managed blocks that are incorporated as they are presented. The slow index causing these difficulties can be tuned or its computing environment addressed (or both) to raise performance.

Alternatively, asynchronous web technology can interject results as they become available, almost like a real-time dashboard of information. Modern technological approaches have turned a weakness into a strength.

Another overcome challenge is relevancy quality. True relevancy quality can only be achieved at the layer where the user interacts with the search interface - identifying user properties, preferences, and query intent - and then translating that to the relevant search indices that are best suited to answer that specific question.

No one relevancy algorithm within a single search index could address those needs. By using a hybrid federated search approach, the search results are always fresh, resulting in improved end-user satisfaction and productivity.

Traditional solutions can still be applicable and may be the best solution where simplicity is required. Tabbed, categorized, or bucketed results can still be used to provide users easy-to-navigate and recognizable interfaces.

Different search engines can provide varying levels of query sophistication. In the past, federation technologies at query time were limited by a least common denominator approach to align all queries to all search engines with the capabilities of the least capable search engine driving those queries, meaning results returned were controlled by the least capable index.

New sophisticated query parsers, cush as SmartHub, pass bespoke queries to each search index that are specifically crafted at runtime to optimize and align for each individual index.

How SmartHub addresses historical limitations to Query-Time federated search

Historic Objections	Modern Technical Realities
Query-Time Merge introduces complexity, redundancy, and administrative burden	• Query-Time Merge simplifies the process, you don't have to crawl systems to keep the central index updated. The search results always contain the latest content since it's searching the original source. • Metadata changes made in the federated system are immediately available in the search solution without re-indexing new data. Previously, re-indexing took days or weeks. • Query-Time Merge allows you to pinpoint similar documents across sources. The same document can be surfaced from different sources with different metadata without the index automatically trimming the duplicates.
Federation prevents correlation of signals	• Omit use cases are driven by capabilities such as analytics, learn-to -rank, document similarity and natural language query, where you can curate content from different sources to get suggestions for relevant documents, regardless of the source index.
Query-Time Merge can be brought down by the weakest link	• Administrators decide which sources get interleaved and which are surfaced in their own containers. This negates the issue of performance. They have the flexibility to separate less performant sources into side-sources displayed in their own container that respond in their own timeframe.
Query-Time Merge prevents cross-source relevance calculation	• Based on previously discussed cross-source correlation of signals, query-time boosting of relevant data does not rely on crosssource relevance. Federated results can identify the best source and rely on the individual relevance returned instead of deciding which document(s) across all sources are the most relevant. • The ability to normalize relevancy values returned by the sources allows interleaving the top documents from each source in a consistent and applicable manner. The best results are interleaved in the correct order due to the normalization.

56%

of employees, without a robust internal search, report never finding the content they're looking for.

How BA Insight makes Query-Time and Index-Time merge intelligent with SmartHub

The BA Insight SmartHub Query Engine supports Federated, Single Index Search and Hybrid Search, where the solution is a combination of the two. It transforms a user's query and broadcasts it to a group of disparate systems, each with its own index with the appropriate syntax.

It merges the results collected from the indexes, re-sorting them in a succinct and unified format, allowing the user to work within the returned merged result set using various tools.

An example of this is sorting or using facets to help the user quickly find information.

This powers organizations to go beyond the traditional single index approach, enabling the implementation of powerful enterprise search solutions that can integrate results from multiple search engines into a single, interleaved result.

Users can effectively use any technology within the organization to obtain the best search results.

SmartHub Architecture

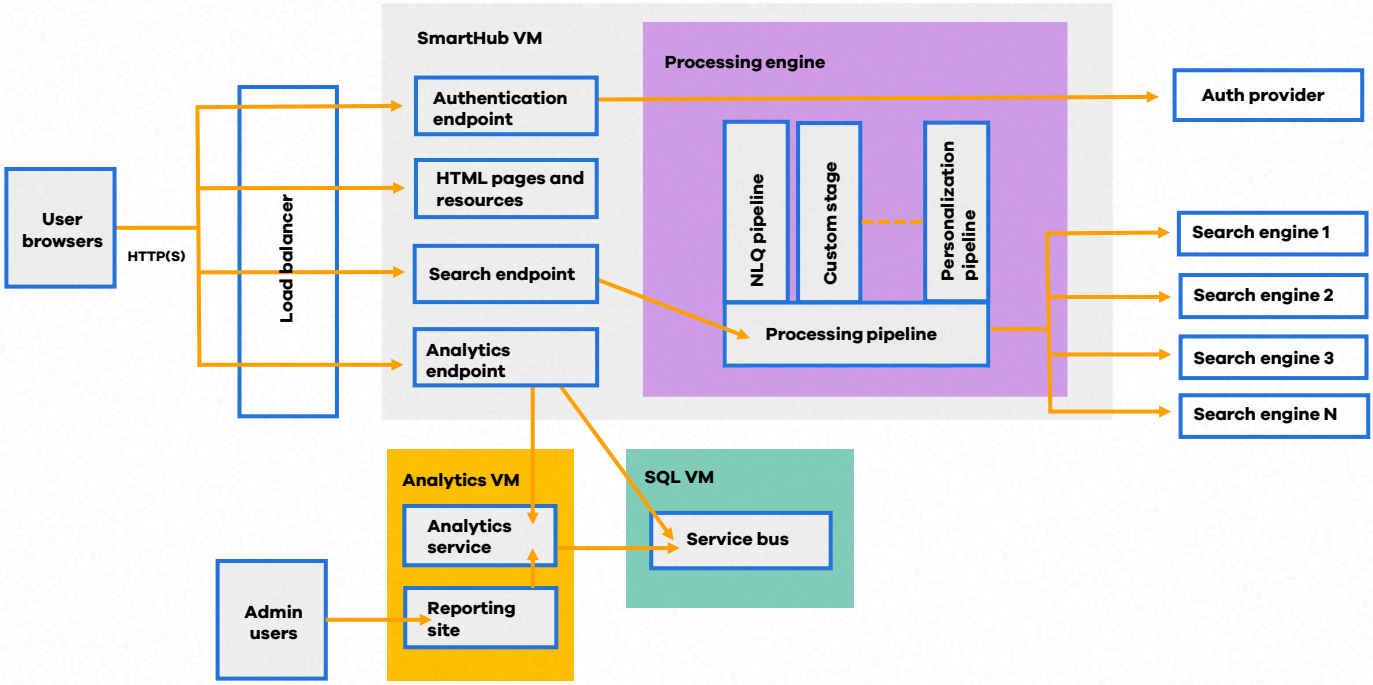

Orchestrating results across search engines

The BA Insight SmartHub Query Engine provides intelligent brokering capabilities between multiple indices. SmartHub can be configured to retrieve results from any combination of the supported search indices via a simple configuration screen.

When the SmartHub Query Engine receives a search query, it transforms and dispatches the query to the appropriate back-ends to obtain the best responses. It recombines the responses collected from various systems into a unified set of results and makes them available to the user via the intuitive search user interface. The high level SmartHub architecture is shown in the diagram above.

SmartHub has been designed with flexibility and adaptability in mind. As part of deployment, an enterprise can add its own logic and set of transformations by adding stages to the query and result pipelines, effectively enabling the creation of smart answers, similar to internet search engines. Each stage allows a distinct transformation to the query or results to be applied.

For example, synonym expansion, stemming, rank boosting, query rewrite (i.e. white paper federation could be rewritten as: “White Paper” as a phrase and the product metadata must be equal to SmartHub, resulting in much higher accuracy), translating results to English, injecting real-time inventory status for a product or partor check-in/check-out status for a document.

SmartHub Query Engine in the Real World

Examples of the many use cases in which federated enterprise search can address what a single index

Purpose-driven indices

A pharmaceutical company had purpose-driven separate indices and was seeking a way to deliver a single interface for all use cases.

One index was scientific, providing answers to scientists, utilizing things like chemical compound search, related information, etc. It had metadata and relationships that only made sense to the research and development community. Another index was corporate-related, containing policies, HR data, and other intranet-driven content.

With SmartHub, this organization was able to bring together both indices, provide their users a single source of truth, leverage personalization to understand users, use AI integrations to determine the intent of their query, and then deliver the correct set of data, regardless of which index it was stored in.

Geographic data restrictions

A global organization faced data export restrictions: data generated in a specific country was not allowed to be stored outside of that country. However, end users did not want to go to a separate search site for each individual country to find the information they were looking for.

SmartHub provided a single point of entry to each country-specific index. Additionally, through language detection capabilities in SmartHub, they were able to support their multi-lingual end users and multi-lingual documents. SmartHub analyzed the language used in the user's query and returned the best result in the correct language for their query. This removed the need for separate language-driven interfaces or indices and saved the company even more time.

Supplementing cloud and multi-tenant applications

The proliferation of multi-tenant cloud applications has created a problem for many organizations. These applications become an integral part of daily activities, but the data is stored and managed by the cloud provider, limiting users' ability to relate it to other organizational data.

As an example: a law firm is storing all of their matter-related work within iManage and their matter and client information in another system. They want to be able to leverage this client and matter data while searching for related documents. They want to find example lease agreements for clients over $100M in revenue with offices in the US and EU. With SmartHub, this law firm was able to integrate the iManage cloud index with their own supplemental client and matter index to seamlessly search for and leverage that exact data.

90+

total connectors available for BA Insight users.

Personal use indices

One large problem with a single index approach is determining the data to add to the index. Often, organizations choose not to include personal data stores like OneDrive, Box, or personal folders. Organizations assume that since this data is usually only available to one user, they shouldn't invest in indexing it.

However, this instantly creates a problem for the end users who now have to go to two places to search: one, for their own data and another for corporate data. With SmartHub, the personal use index can be queried at the same time as the corporate index, and a single result set is returned to the user without infrastructure or indexing cost.

Very large data sets

There are a couple of use cases around extremely large data sets that pose a problem for a single index approach.

1. Large data archives

Many organizations have hundreds of terabytes of archives spread across systems like Box, SharePoint, etc. These data sets are often unnecessary for primary searches, but they do need to be searched for regulatory queries, deep data analysis, or for archival and retrieval purposes.

Content that is very large in size or count and not queried on a regular basis often places a strain on the central index infrastructure and cost requirements. With a federated approach, the archival systems can be queried directly, when needed, and have results returned along with results from the primary index.

This allows organizations to maintain their access and regulatory requirements while removing the complexities and costs associated with indexing this data.

3.5 M+

more than 3.5 million users worldwide rely on BA Insight for superior search experiences.

2. Purpose-built systems

Purpose-built systems like NetDocuments in the legal market, store 100+ million documents which are actively used by legal professionals.

Since the rate of change is in millions per week, trying to even incrementally crawl and create a single index causes data to lag behind, resulting in bad search results.

We partnered with NetDocuments to create a hybrid approach where we crawl and create a single index for SharePoint or any other search platform and perform query-time merging between that index and NetDocuments index. This ensures that data shown is alway new and up to date.

Summary

Federated search at query time is now as good as federated search at index time. In many cases a hybrid approach is a better approach than either of the two options and should be considered as a eay of modernizing enterprise search or as part of digital transformation projects.

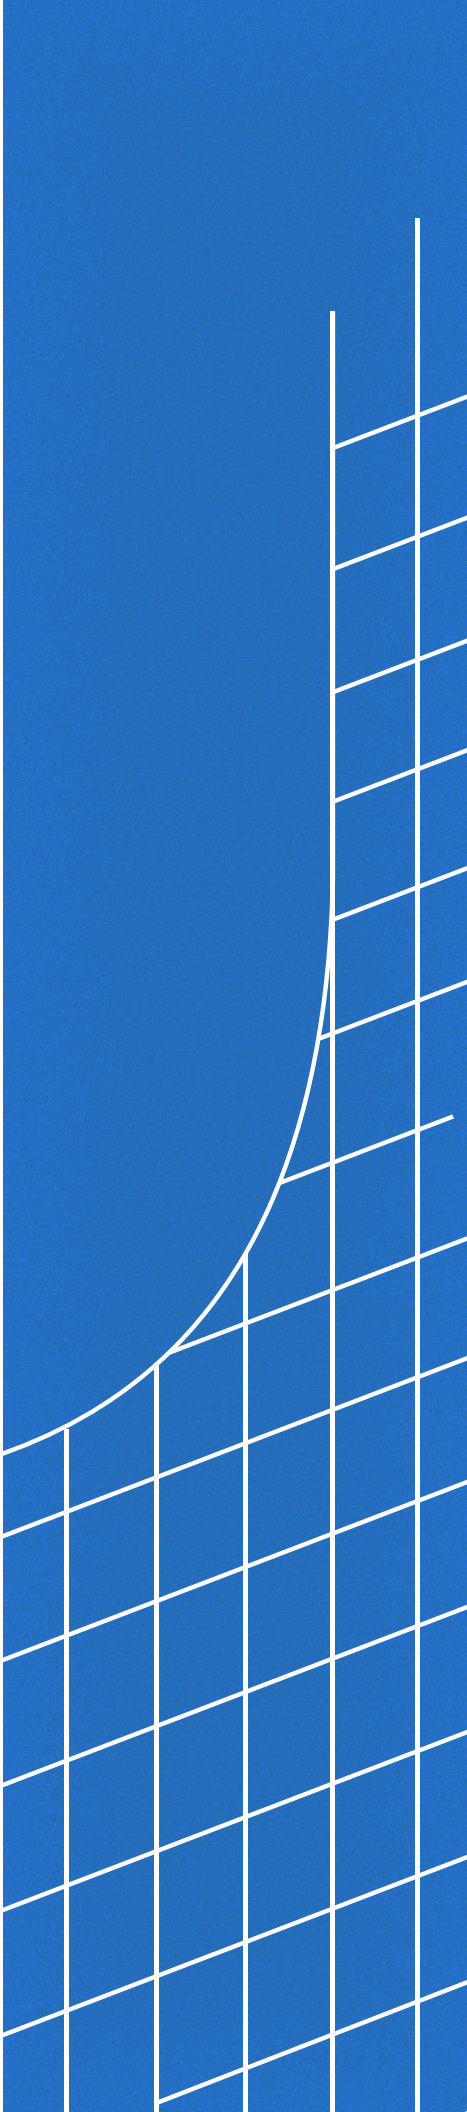

BA Insight

BA Insight leverages AI technology to deliver a web-like advanced search experience with results that are relevant, personalized, and actionable. The connector-based technology works with enterprises, customer portals, and websites, turning searches into actionable insights, regardless of where content or users reside.

Want to learn more?

Contact us.

bainsight.com
phone: +1.339.368.7234