pureinsights™

ENTERPRISE SEARCH EXPERTS

Search is More Than Software

Three essential strategies to make your search bar work like Google

WHITEPAPER

Defining the future of search

The status quo won't cut it

Ask any knowledge management expert to describe how enterprise search is evolving, and you'll likely hear some version of the following sentiment: search is both more essential and more complex than ever before.

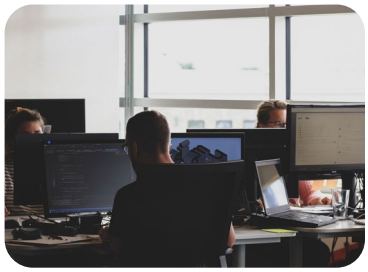

Over the past two decades, with Google leading the way, search has transformed from a novelty of the internet age into a cornerstone of modern life. Search engines are now the primary interface for countless everyday activities like getting directions to the grocery store, discovering new music, and even finding a job.

Individuals have learned to rely on these search engines to tackle both trivial and elaborate tasks. Now search is considered a utility rather than a feature. It's a non-negotiable element of any website or application.

But here's the thing: search is also getting better. With Google once again setting the curve, queries are leading to faster and more accurate results than ever before. People expect trustworthy and contextualized answers rather than a list of results ranked by relevance. They don't just go to search engines with simple, fact-based questions. Now people trust these platforms to help them address everyday concerns in plain natural language, like searching for "the best sushi restaurant in San Francisco" or identifying "the best basketball player of all time."

That combination of efficiency and quality is now the expectation for every search experience—but it is rarely found beyond the major search platforms. Workplace intranets are often slow and incomplete, leaving team members to try different combinations of keywords and sort through pages of results to get what they need. E-commerce search bars might successfully field specific product searches, but they fail to showcase related products or accurately interpret a broad query.

As search becomes more central to customer and employee experiences, leaders can no longer afford to deprioritize innovation in knowledge management and search technology. The consequences are too severe. External search engines directly impact revenue and customer experience, and internal search engines powerfully influence employee satisfaction and efficiency. As an enterprise leader, if you want to stay ahead of the curve, you must find a way to match Google's performance for your internal and external search bars.

From Search Engines to Question-Answering Systems

At Pureinsights, we describe the transformation of search technology over the past decade as keyword-driven search engines evolving into AI-driven **Question Answering Systems.** (Some experts prefer to call them "insight engines" or "cognitive search" services.)

No matter your phrase of choice, these modern search experiences all share similar architectures. They are designed as two-sided platforms that integrate artificial intelligence to simultaneously digest content and serve users. These platforms effectively interpret queries from users and intelligently deliver the most valuable content based on that interpretation. Once deployed, they are constantly digesting new data, making new connections, and analyzing user behavior to improve their predictive capabilities.

On the content-facing side of the platform, a fully optimized Question Answering System should be able to understand concepts, extract facts (and connections between facts), and quickly identify the most relevant concepts and facts to address a query. On the user-facing side of the platform, this system must understand and predict intent in a user's query, present content snippets and related info, and directly answer their question with the most relevant response.

Two-Sided Search Platforms

CONTENT	USER
Understand what a document says	Understand a user's query
Infer concepts in the document	Predict intent of the query
Extract facts and connections	Present snippets and facts related to the query
Quickly identify the most relevant concepts and facts	Quickly display the most relevant concepts and facts

Keeping up with Google

So if the speed and accuracy of a Google-like Question Answering System is the new bar for enterprise search, and you know that you must integrate AI technologies to meet users' expectations, what steps can you take to get there?

The fact is that Google—the very experience you're trying to emulate—actually operates in a simpler environment than most enterprise organizations. It pulls data from one source (the internet) into one public interface (Google.com). Google also operates with a small variety of document types wrapped up in a single data format (HTML) that makes it easier to ingest massive amounts of information. Enterprises, by contrast, must manage multiple data sources, in numerous formats, and feed those into both public and private interfaces that are constantly changing.

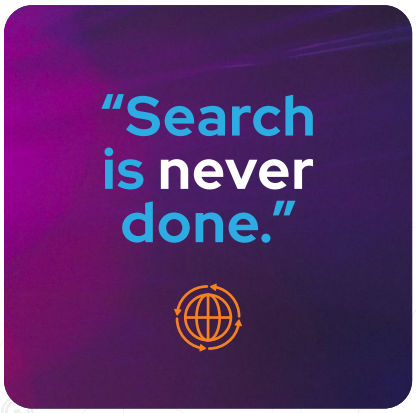

Thankfully, **we now understand how Google powers their search experience,** which means that your organization can match up to Big Tech and delight your employees and customers.

This white paper provides a roadmap for three core strategies to consider as you build your new Question Answering System: **Technology, Talent, and Tuning.** It also discusses existing solutions for each strategy. Finally, this paper will introduce SearchOps, a new managed services model for enterprise search that can vastly improve the ease of running your Question Answering System while maximizing its performance.

Technology: The right AI tools to make search smarter

Defining "artificial intelligence"

In the past, the true north metric for any search solution was keyword-based relevance. Today, most search engines rely on the same open source library (Apache Lucene) to establish keyword-driven relevance rankings, but genuine relevance — the kind that powers a well-oiled Question Answering System — demands new technology.

Over the past few years, leading technology firms including Google, Microsoft, and Amazon have revealed the specific AI innovations that power their world class search experiences: Natural Language Processing (NLP), Machine Learning, and Knowledge Graphs.

These three technologies work together to augment both the content-facing and user-facing sides of a search platform. They turn search systems into powerful tools for understanding queries and data rather than simply matching keywords. And because AI technologies are developing so rapidly, enterprises who integrate these tools into their own search stack will benefit from their ever-strengthening problem-solving capabilities.

The integration of AI is fast becoming table stakes for enterprise search. In IBM's press release announcing their selection for the 2021 Magic Quadrant for Insight Engines, Gartner defines the gold standard for enterprise search vendors as those that "combine search capabilities with artificial intelligence to deliver actionable insights derived from the full spectrum of content and data sourced within and external to an enterprise."

Because NLP, Machine Learning, and Knowledge Graphs influence different aspects of a search platform in unique ways, we like to use the "Three I" framework to help our clients understand each technology's impact on the search experience: NLP **interprets** the meaning of data and queries; Machine Learning **infers** the intention of queries and extracts concepts; finally, Knowledge Graphs **inform** users by making connections between data points and surfacing related information.

Three strategies to keep up with Google.

Technology:
The right AI tools to make search smarter

Talent:
The right people to deploy and manage your system

Tuning:
The right plan to maintain performance

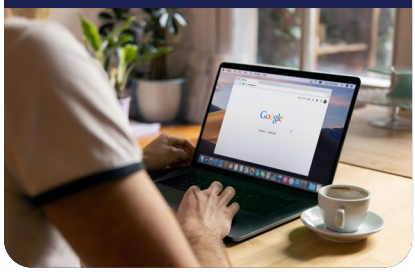

The "Three I" Framework for Enterprise Search

Natural Language Processing	→ INTERPRET
Machine Learning	→ INFER
Knowledge Graphs	→ INFORM

INTEPRET

Interpreting with Natural Language Processing

NLP includes technology that enables computers to understand humans in text and in conversation. It's existed in some form since the 1970s, but recent advancements have increased the visibility of NLP technology, notably through chatbots and virtual assistants.

The premise of NLP for enterprise search is quite simple: by interpreting ambiguous language, search systems can better identify the right piece of information for a user's needs. NLP also improves a system's understanding of indexed content, enabling more accurate delivery of results. Once users experience the question answering capabilities of a search system powered by NLP, they expect similar performance from every search bar they come across, which leads them to enter more subjective and ambiguous queries. If your search platform isn't set up to understand them, you're setting your users up for disappointment.

INFER

Inferring with Machine Learning

Machine Learning affects multiple elements of a modern search solution. In its most common application, this technology trains search systems to understand concepts rather than just objective facts. For example, a system would learn over time that the phrase "happy baby" means the same thing as "pleased child" and may also refer to a common yoga pose.

When parsing queries, this learned knowledge helps a system understand what a user intends, even if they don't use a precisely familiar set of words. When surfacing relevant information to address a query, that knowledge helps the system avoid false negatives.

Machine Learning also enables search systems to understand that two different people searching for the same thing might need different results. A marketing manager searching for "past customers'' on an intranet might need case studies about those engagements. However, a software engineer searching for "past customers" may be interested in the technical details of each engagement. Machine Learning will utilize historical search logs to identify those patterns and adjust results accordingly.

Ultimately, Machine Learning is about automating contextual clues that could be entered manually but become impossible with scale. The more content a Machine Learning system ingests, the smarter it gets, and the better your results become.

INFORM

Informing with Knowledge Graphs

Knowledge Graphs provide a vast pool of knowledge that search systems can use to inform users with specific facts instead of general links or documents. This ability becomes even more important as user behavior continues to shift towards plain language questions.

In practice, Knowledge Graphs are databases of facts about people, places, and things. They also store the relationships between these facts to enable the fast retrieval of context-driven information. Google's Knowledge Graph stores billions of facts that allows it to directly answer questions like "when was the Empire State Building constructed?" with a knowledge card at the top of the search results, followed by answers to related questions that "people also ask" about the building. Without a Knowledge Graph, the results would simply list webpages related to the topic at hand.

Adding AI to your search stack

Designing and deploying a new Question Answering System that leverages AI requires paying careful attention to where new technologies are added in your stack. When building your system, you'll still need all of the standard components you've used for years: tools like data connectors, staging repositories, data repositories, and of course the search engine and search interface itself. So where does AI fit into your existing stack?

NLP is most clearly integrated at the query processing level. Knowledge Graphs serve as powerful tools in the middle of your stack to store facts, entities, and relationships after they're processed and before they're requested. Machine Learning is a workhorse that intelligently processes data when it's ingested; it also works alongside NLP to infer meaning in queries and predict users' intentions.

Existing technology solutions

The first step on your journey to building a Question Answering System is to establish a basic keyword-driven search engine. There are a number of open source tools available to accomplish this, including **Elasticsearch**, **Solr,** and **OpenSearch.** These products make it incredibly easy to spin up search bars on websites, applications, or intranets. Many also integrate elements of NLP, Machine Learning, and Knowledge Graphs to enhance your performance.

However, although most of these "Search as a Service" platforms address hosting and maintenance of technology at enterprise scale, they don't provide hands-on support or easily accessible expert knowledge. That proficiency, which we'll address in the next section, is an essential part of installing and managing search technology over the long-term.

The Modern Search Stack

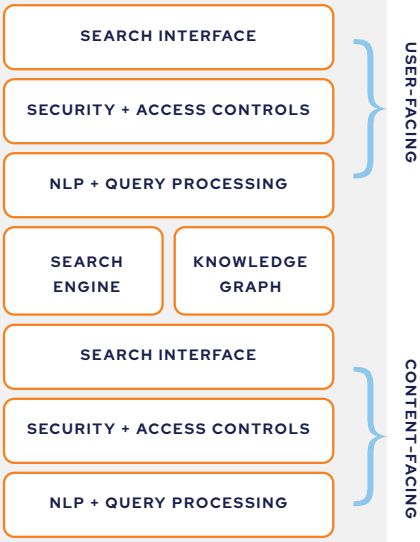

For your consideration

The Pureinsights Discovery Platform™ provides and integrates most of the tools required for a modern search stack. We bring our own proprietary solutions and work with clients' pre-existing technology to preserve investments and reduce costs.

Learn more at pureinsights.com/technology.

Talent: The right people to deploy and manage your system

Don't search for a band-aid when you need a surgeon

If you're putting together a cutting-edge Question Answering System for your enterprise, the best technology is only a starting point. You need knowledge management and enterprise search experts on call to deploy your new system effectively.

Ten years ago, you might have been able to get away with delegation to your existing IT team. However, as discussed in the previous section, enterprise search systems are no longer straightforward, plug-and-play operations. They incorporate a variety of separate connectors, tools, and technologies that all need to work in harmony to answer questions for your employees and customers. In other words: they demand specialized expertise. We find that implementation is the most common obstacle that prevents enterprises from realizing the full potential of their search systems.

The installation process

Installing search technology will look different at every organization, but most enterprises share common tasks. These include:

- Assessing current search applications
- Measuring and improving search quality
- Identifying key objectives for a new search experience
- Selecting new technologies for integration into the search stack
- Determining a plan and schedule for implementation
- Setting up scalable cloud infrastructure
- Integrating data connectors, staging repositories, and data processors with content sources (third party software, websites, file systems, and databases)
- Deploying AI-enabled services (NLP, Machine Learning, Knowledge Graphs)
- Processing and indexing documents to optimize them for search
- Maintaining compliance with security standards

Experienced search experts will save time at each step of this process and improve the performance of your Question Answering System. These individuals play crucial roles ingesting and processing content, implementing the latest updates, and populating knowledge graph databases.

Of all the tasks listed above, ingesting information is one of the most difficult challenges. Modern enterprises install software and produce new information every day. To make their search bar work, they must find a way to connect CRMs, web properties, project management systems, collaboration platforms, and countless other tools and documents to a unified search solution with suitable security and indexing. To add to the challenge: most enterprises haven't consistently indexed documents across their organization. This unindexed content can hamstring the quality of a search system by leaving gaps in data sources and decreasing the accuracy of responses.

To mitigate those combined risks, search teams are left with a triple threat of responsibilities: indexing a backlog of historical documents and data, processing current data, and ingesting new data sources as they're added to the system.

The talent crunch

Who will you trust to guide you through the labyrinth of decisions and processes detailed above? If you're building a basic website search function, you can probably rely on your IT team. But if you're working towards a Question Answering System that gives Google a run for its money, you'll need real domain expertise.

This is where many leaders run into their next hurdle. As enterprise search becomes more prominent, more valuable, and more complex, experienced search engineers are getting harder to find and more expensive to retain.

The real cost of building a team

Research from **Bridge In** shows that IT services and software development industries see 13% average turnover per year, and up to 21.7% for embedded software engineers. According to **LinkedIn,** the cost for each employee lost can be up to 250% of their annual salary.

Let's say you're up for the challenge—what might it take to build a team from scratch? Based on our team's experience, it takes approximately 12 months to train a capable search application developer, 3-5 years to train a lead search developer, and 4-6 years to develop a seasoned search architect.

Existing talent solutions

There are two primary strategies to acquire talent in enterprise search, each with pros and cons.

In-house search teams provide the greatest level of coverage and security. They spend every day managing and iterating on your search application. Of course, given the shortage of qualified experts, building this team will require significant compensation and benefits or a rigorous training program to develop your existing IT talent pool. And even if you've developed a world class team of search experts, you'll likely need outside help when troubleshooting or deploying advanced features. Generally, an investment of this scale makes the most sense for organizations that require search as a core feature of their product or service.

On the other hand, search consulting groups provide expertise on an ad hoc basis to assess search performance, tune search platforms, or build and launch new systems from scratch. The chief benefit of this approach is flexibility, which also translates into cost savings. Hiring third party consultants can be an effective way to optimize your search bar at regular intervals. However, it generally results in a reactive stance—responding to problems as they arise rather than proactively building a better solution. As search increases in complexity and importance, this impromptu strategy eventually diminishes the quality of your search experience. We'll explain why this degradation occurs in the following section.

Tuning: The right plan to maintain performance

Search is never finished

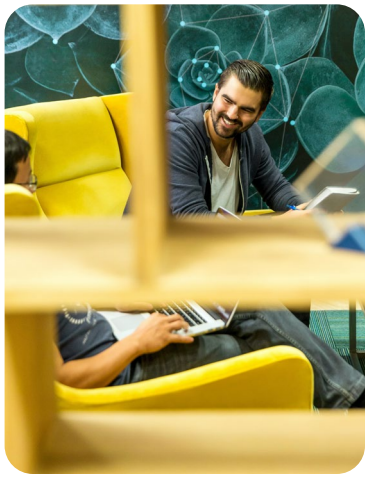

Modern search demands constant attention. The best Question Answering Systems succeed because they are constantly learning and evolving to meet their users' needs and expectations.

As leading search engines continue to innovate, you must be prepared to adjust your own technology stack to keep up with changing expectations. And as your data sources evolve and expand, you must test and refine your platform to maintain peak performance.

It can be tough to justify investment in the management of a search system when other priorities so clearly demand attention. But enterprise search falls into the trickiest category of business operations: those that are only apparent when they stop working.

Nobody questions a search bar when it meets their expectations. After all, it's doing what they expect it to. However, when that search bar starts to produce less helpful results, is missing crucial data, or appears to misunderstand a question, it inevitably leads to a decline in satisfaction and usability. Whether you're operating an internal or external search tool, that decline has very real implications for your bottom line.

Testing and tuning your search system

Search experts refer to the ongoing management of search systems as "tuning." It's the final strategy we included in this white paper because it is absolutely essential for leaders who are serious about designing a search bar that works like Google.

As outlined in our section on technology, modern search stacks contain no less than e ight elements that all work together to produce fast and accurate answers to questions. Optimizing that system means tuning the overall platform as well as its component parts.

To tune your technology, you must first understand how it's performing. Is it consistently delivering the results that users expect? (At Pureinsights we follow a proprietary process we call **search relevance and engine scoring.** You can also read our **"Search Engine Scoring for Dummies"** article on KMWorld for more information.)

Any scoring process includes gathering historical results and normalizing the performance score. Once that benchmark is established, you can confidently design individual tuning experiments to test changes to individual technology tools and processes. By regularly running these experiments in an ongoing process of simulating, iterating, and measuring, you'll always know exactly where you stand—and how to get closer to the Google-like experience.

Existing tuning solutions

In our experience collaborating with enterprise organizations, we see the same linear development cycle play out time after time. An executive is notified about a problem with search performance—often a customer complaint or employee feedback—and they call a meeting to review the organization's search strategy. Then they budget hundreds of thousands of dollars to hire a consulting firm to deploy an updated platform and ingest fresh data with high expectations for its future performance. Inevitably, once the third party has completed its work, and negative feedback has stopped, the organization returns to business as usual, letting the platform run on autopilot for the next few years.

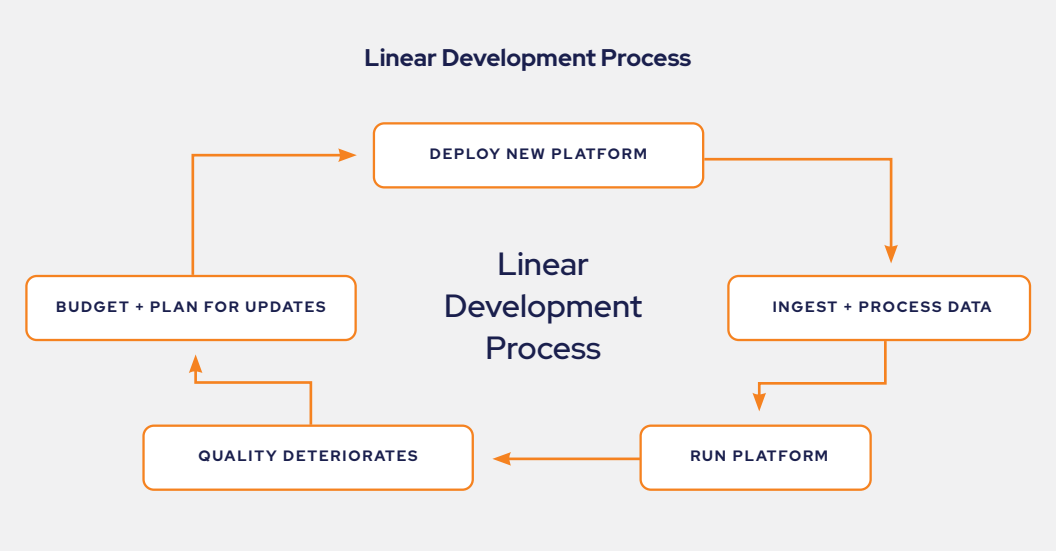

Let's be clear: this linear development process isn't inherently bad. It works well enough to keep your enterprise from slipping too far behind the curve. But it puts you in a reactive posture, responding to problems instead of crafting solutions. It will never get your search bar to a point at which it's constantly beating your users' expectations.

The way to mitigate this risk is to shift to a dynamic development process with testing and tuning baked into your standard operating procedure. It's the more difficult path, but it keeps user satisfaction high, eliminates urgent problem-solving meetings, and ultimately saves money by avoiding the need for brand new platforms every 3-5 years. If your goal is effective search that maintains its performance, you need to take a proactive approach.

SearchOps: A managed service model for enterprise search

Your business isn't search

Developing a Question Answering System that keeps up with the world's largest tech companies is finally possible. Using the right combination of technology, talent, and tuning, any enterprise can deploy a search bar that delivers contextual information to their employees and customers. You *can* truly do it on your own.

Still, enterprise search is more complex than it's ever been. When search was a supplemental technology, it made sense to keep it as a low-priority task for IT teams. Now that search is a foundational element of organizational performance, it demands thoughtful attention and deep expertise that many enterprises can't manage on their own.

In many ways, enterprise search parallels back-of-house business operations like accounting. Even the largest organizations hire third-party firms to manage their accounting and financial concerns, because their time is better spent on their core business. We counsel our clients to think of enterprise search the same way. If you can't justify an in-house team, hire an external team—just make sure someone is managing it on a daily basis.

Taking a proactive stance

In the previous section we discussed the prevailing linear approach to the enterprise search development cycle. We also found that it fails to address the ever-shifting needs of a modern Question Answering System, because such platforms demand constant attention.

After decades helping enterprise clients develop and maintain their search bars to keep up with Google, we've developed a new model for enterprise search management. This dynamic approach positions testing, tuning, technology updates, data ingestion, processing, and security management as ongoing responsibilities rather than tasks completed every few years. It also makes it easy to scale search alongside an organization's growth.

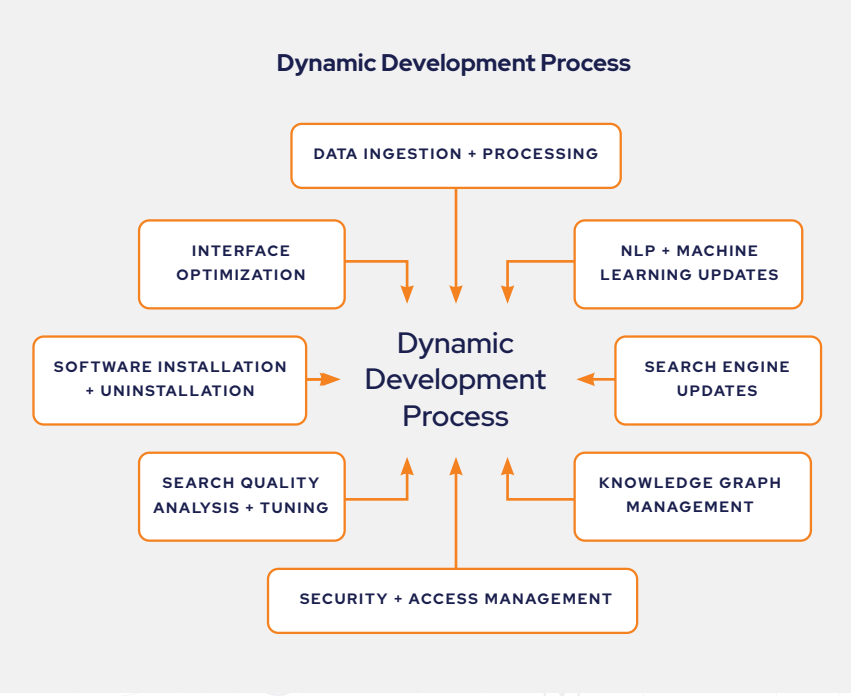

We've seen that dynamic development makes it easier to address many of the obstacles raised in this white paper. By constantly updating your technology stack, you ensure that your search experience is delighting users. By making search management a daily priority, you define clear roles for internal or external talent. Lastly, by conducting regular search quality analysis, you ensure that you'll always have a roadmap to tune your system.

Search as a managed service

We're convinced that existing enterprise search solutions lag behind the curve. Too many software products claim to be turnkey when they really require hands-on integration with an ever-growing search stack. Traditional enterprise search consultants (including our own team) help with one-off deployments, but don't offer the long-term support necessary to take a dynamic management approach.

That's where SearchOps comes in. SearchOps means taking a managed service approach to enterprise search. Just like your favorite accounting firm, a SearchOps provider can handle spikes of seasonal work (during installations or deployments) and manage the daily tasks necessary to keep your Question Answering System firing on all cylinders.

The right enterprise SearchOps provider should be an organization with which you feel comfortable sharing your most sensitive data. They should operate as a full-service solution to execute your technology, talent, and tuning solutions on your behalf. Ultimately, SearchOps is the shortest path to a search bar that works like Google.

Benefits of SearchOps:

- Preserve the best of what you already have
- Continuously improve your tech and tuning
- Experience less downtime and fewer malfunctions
- Decrease your overhead costs
- Stay focused on your core business

SearchOps Services:

- Flexible coverage options
- Full ownership of the search stack, from content ingestion to the user interface
- Counsel from enterprise search experts
- A modular search platform to enable constant iteration
- Scalable cloud infrastructure
- Integration of the latest NLP, Machine Learning, and Knowledge Graph technology
- Constantly search quality analysis and tuning

About Pureinsights

At Pureinsights, we're building a business around our belief in the power of enterprise search and our anticipation for what the future holds. We spend every day helping enterprise leaders figure out how to integrate AI technology and optimize their search systems to surprise and delight their users.

With an average of 15 years working in enterprise search, our team offers consulting, best-in-class technology, and a SearchOps model to help our clients maximize their search experience. If you're looking for a partner to take your own system to the next level, we're available to tackle your challenges one at a time or take full ownership. In other words, if you give us your search bar, we'll make it work like Google.

> We needed a partner that could educate us in the art of the possible and help build our roadmap for the future.
>
> **Steve Dineen**
> Founder / President
>
> fuse

> Pureinsights started the project with a comprehensive and well organised discovery phase. They took the time to properly explore our needs and, crucially, the needs of our clients.
>
> **Steve Nutall**
> Global Head of Product Development
>
> MINTEL

Fiercely vendor agnostic. Infinitely scalable.

Enterprise Search Consulting

Technology Platform

SearchOps

Contact our team at info@pureinsights.com

pureinsights™ ENTERPRISE SEARCH EXPERTS