



Making Search Work

April 2017

Martin White



Contents

- Introduction
 - Making information invisible
 - Find, not search
 - Information deconstructed
 - Information quality
 - Search cannot be perfect
 - Enterprise application support
 - Managing the search experience
 - The critical success factor
 - Quantifying the search team investment
 - Search team or search technology – which comes first?
 - SharePoint search is no different
 - Monitoring search performance
 - The return on search team investment
 - Recommendation
 - The author
-

Introduction

In most organisations there is frustration that the enterprise search function does not work as well as Google. There are some technical reasons why this should be the case but the main reason is that Google has over 72,000 employees who are dedicated to making search work well. The constant evolution in the Google search interface is not a result of ad hoc decisions but of the analysis of 3.5 billion searches a day and constant collaboration between the technology teams, the search performance teams and the business development teams to improve the search experience.

Inside the enterprise there may be just one person with a responsibility for search management, and the evidence suggests that this may not be a full-time role and that the person with the responsibility does not have the set of skills needed to tackle the complexities of search.

Organisations appreciate that when an investment is made in technology, whether in IT or in other areas, there is also a requirement to invest in the people who will make it work. That appreciation does not seem to cover search applications. This may partially be a result of the way in which a Google search appliance could be plugged into a server rack to deliver instant search. No one ever tried to assess whether it was good search - after all it was Google!

The objective of this report is to demonstrate that without an appropriate level of investment in a search support team, a search application will never provide users with a satisfactory search experience. Not only is that a waste of investment in hardware and software but also in the investment the organisation has made in creating information. Outside of employees in a manufacturing environment everyone else is being paid to create information. If employees cannot find this information, then the cost to the organisation will be uncomfortably close to the cost of salaries.

The charts in this report are reproduced with permission from the [AIIM 2014 Search and Discovery Survey](#) and the [Findwise Enterprise Search and Findability Survey 2016](#).

Making information invisible

Most companies recognise that being able to find information and knowledge is business-critical. However, the evidence from global surveys carried out by [Findwise](#), [AIMM](#) and [NetJMC](#) indicates that employee satisfaction with the performance of corporate search applications remains very low. These surveys indicate that in four out of five companies, employees find it difficult, or very difficult, to find the information they are looking for.

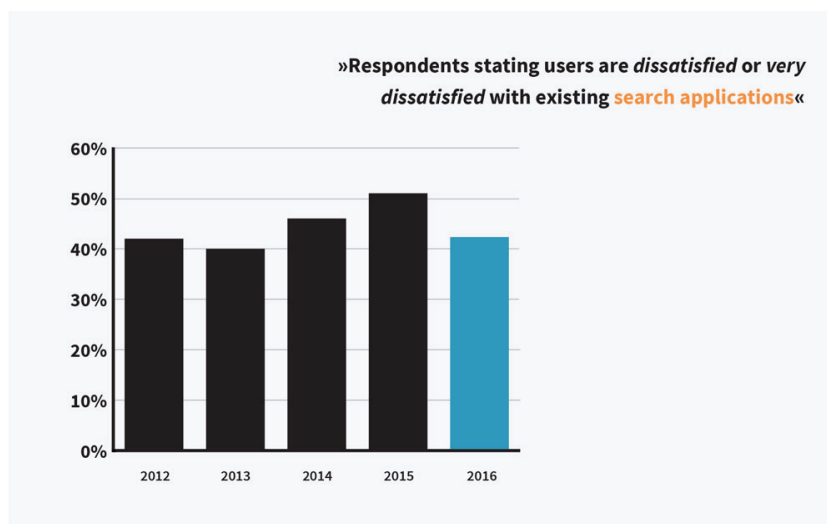


Figure 1 Search application dissatisfaction - © Findwise Enterprise Search and Find-ability Survey 2016

It is very difficult for users to assess the extent to which their experience is unusual and that in other companies it is very much better. However, all the surveys and discussions at conferences suggest that poor quality search is both the norm and is 'acceptable'. Comments are often made that enterprise search is dead and that keyword search has no future, but it seems likely that these comments come from people who have never experienced high quality search applications.

It does not help that it is very difficult to work out why search does not seem to work and that the information required urgently to inform a decision seems not to exist. The reasons could be one or more of the following:

- The information does not exist
- The server on which it exists is not indexed by the application
- The information is new and has not yet been indexed by the application
- The information exists but the filters and facets are not designed in a way that will enable the user to reduce the initial number of results and not run the risk of excluding the information in doing so
- The information has been indexed but the ranking algorithms have pushed it so far down the list of relevant results that the user has given up looking for it
- The user does not have the appropriate security clearance
- The information has been indexed but is not in the language of the query

Whatever the reason or reasons the user loses trust in the search application, conscious that the inability to find the information (or know for certain that it does not exist) could have an impact on the performance of the company and their own career and reputation.

It is important to appreciate that the problem lies with the search technology so long as it has been updated in line with vendor recommendations. All the search applications on the market will deliver a very high level of user satisfaction but only if one crucial success factor is recognised and addressed. This report is about this critical success factor, but before disclosing what this is some background context is essential.

Find, not search

It can be very useful to consider why people use a search application and what their expectations are. This is part of a discipline of information behaviour research and since the late 1980s many frameworks have been developed that recognise that information seeking is a very complex process and is fundamentally a dialogue between the user and the search application.

For the purposes of this report there are three important use cases.

Learning - where the user may not be quite sure what the 'best' query is and will expect the search application to guide them through features such as auto-suggestion. There is usually no time pressure for this learning process, and the user may return to the search application several times to accumulate all the information they need. This is an example of Exploratory Search and the user expects the search application to recall most (and ideally all) of the relevant documents.

Task-completion - where the user wants to find either a specific application to complete a task (initiate the recruitment of a new employee) or a document that provides detailed guidance on the process. The process may be different in different countries but the user expects that either the application or the guidance document will be at the top of the results list even though they may have used a very short query such as [new employee process] and not included their location in the query. This is a very high precision query and is an example of what might be regarded as a Known Item Search.

Reassurance - where the user is about to make a decision and wishes to make sure that they have found not only the most relevant documents but also any documents which might be less comprehensive but are much more recent. They may have the sales for Q1 and Q2 but need to make sure that Q3 data is included before making a decision. This requires currency through constant indexing - a very substantial challenge for a search application.

These use cases highlight the three primary metrics of search performance.

Recall is a measure of the number of the documents presented in the search results as a percentage of the total number of relevant documents that have been indexed.

Precision is a measure of how many relevant documents have been presented in the search results as a percentage of the number of results presented.

Currency is a measure of how long it takes from a document to be added to a repository to when it is indexed and searchable.

For reasons outside the scope of this report only under exceptional circumstances is it possible to tune a search application to give both very high recall and very high precision in response to a query. The challenges of achieving high precision are usually solved pragmatically by promoting content as a 'best bet'.

Information deconstructed

In its 2016 survey, Findwise asked respondents to state how important different categories of information were to the business.




Figure 2 Information priorities - © Findwise Enterprise Search and Findability Survey 2016

The information environment in all organisations is very complex. Figure 2 highlights two important issues. The first is the range of information sources that employees feel they need access to. In the past, the focus has tended to be on finding documents. Now the requirement is for information, ideally as a search card that brings together information from a range of sources (often in real time) rather than a list of documents. Because the source of the information is not usually displayed on the search cards the importance of delivering the highest possible quality of information is very important.

The second is that in this overall priority there seem to be low importance information categories. Across all employees this may be the case but for a smaller group of employees some of these categories, such as product information, are of very high importance. Search has a value to every employee, including for example those in a manufacturing plant. Their safety may well depend on the plant manager being able to find information on chemical hazards.

The challenge for search managers is how to ensure that each of these information categories is managed effectively. Someone searching for 'employee self-service' almost certainly does not want to see a dozen documents on the application but wants to be led to the application URL. However, after trying to use the application the documents may become more important, and to others these documents are just what they are looking for.

Information quality

One of the frustrations about searching for information is that the search identifies both low quality and high quality information with the same rapidity. A search may be deemed a failure by a user because the information is not of the required quality. PowerPoint files are a good example. The titles are usually vague (“Steps towards the horizon”) and consist mainly of charts and diagrams. Searching for one of these diagrams is usually very difficult even for attendees at the meeting at which the diagram was presented.

Information quality covers seven generally accepted dimensions:

- **Accessibility:** is the information easily retrievable?
- **Accuracy:** is the information free from error and unambiguous?
- **Believability:** does the information come from reputable, trustworthy sources?
- **Completeness:** is the information comprehensive?
- **Consistency:** is the information objective and free from personal bias?
- **Relevance:** is the information fit for purpose?
- **Timeliness:** is the information timely for use?

Information quality has to be addressed through a company-wide information management policy. It cannot be fixed by metadata for two reasons. First it is not practical to go back through all the items that have been indexed and add metadata. Second metadata will not improve the quality of the content. Indeed, it may just improve the ability of the search application to find low quality information.

Search cannot be perfect

Search technology cannot be perfect. The nearest you can get is to use specialised e-discovery applications managed by experienced e-discovery specialists to find close to all the information relating to a particular fraud in an email archive. That takes a great deal of time. The reason that search cannot be perfect is that every user has their own view of what is relevant. Two users may look at the results pages from an identical query; one will be very satisfied with the information listed and the other may have the view that search is broken.

Employees take a realistic view of a search application. They know from their experience with Google and Bing that sometimes the outcomes of a search are poor. But a web search is rarely business-critical and usually there are other ways to find the information. Within a company, the search application may be the only way to find the information employees need to make business decisions. The bar to satisfaction is therefore much higher. For the sake of argument, they will expect search to work well on 8 out of 10 queries.

The solution to poor search performance has often been to replace the search application. This may result in an initial increase in use and satisfaction but this soon reverts to the original level (if not below it) as the project implementation team is disbanded. Job done!

Enterprise application support

With all enterprise applications, it is the ongoing support costs that represent a substantial component of the total investment. The US industry research company Computer Economics specialises in tracking the levels of support that companies invest in these applications. Figure 3 shows the range for an Enterprise Resource Planning (ERP) application.

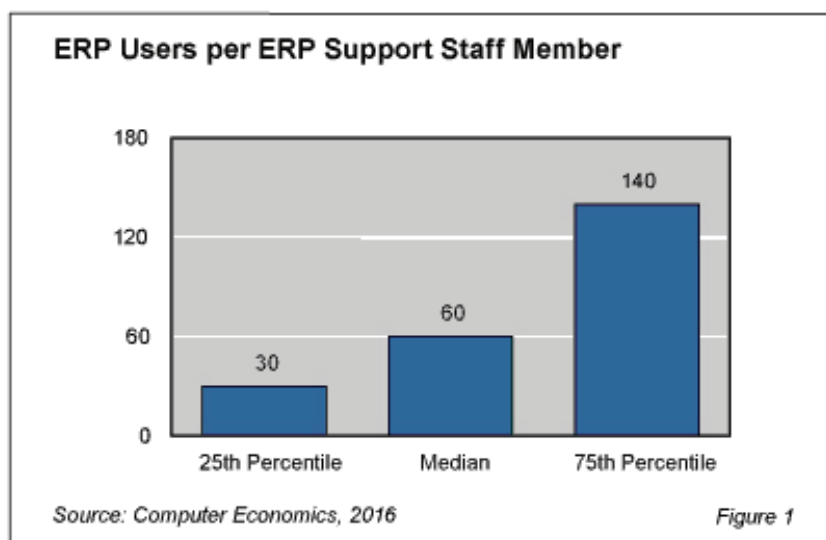


Figure 3 Support ratios for ERP applications - © [Computer Economics 2016](#)

Here “user” refers to “named users” (individuals who have a username and password for the system) and not “concurrent users” (the number of users logged into the system at one time). Included in the definition of ERP support staff are programmers, business analysts, project managers, database administrators, help desk personnel, systems programmers, ERP end-user reporting specialists, and training and documentation personnel, among others. The data includes support staff in the IT organisation as well as support staff in the user organisation and includes temporary contractors in addition to permanent employees. Even in the 75th percentile ‘worst case’ situation there are 140 users for each member of the support team.

Information on the size of search support teams is a feature of the Findwise surveys, and Figure 4 shows the trend over the last five years.

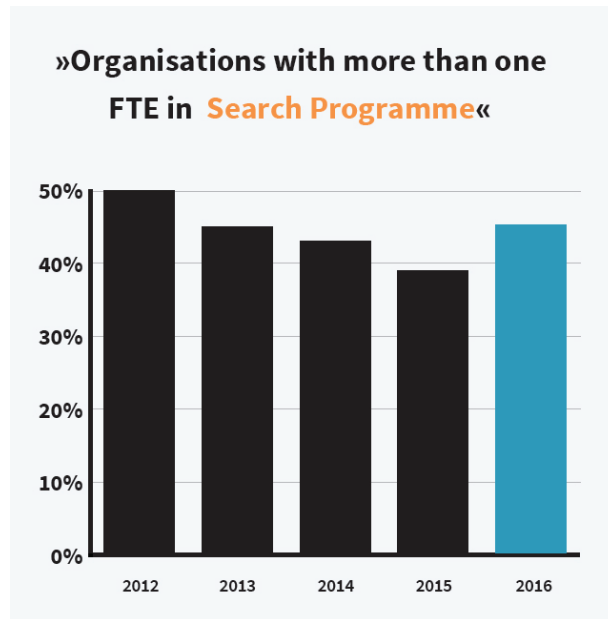


Figure 4 Organisations with more than 1 FTE - © Findwise Enterprise Search and Findability Survey 2016

Managing the search experience

Search development tends to happen on an ad hoc project basis in response to concern from one or more senior managers. Although there seems to be an increase in the number of companies that have a search strategy, it is still below 50% in major companies and perhaps even lower in mid-sized and smaller companies. Search must be seen as at least a programme of projects, with the programme office being a de facto centre of search excellence.

On an almost continuous basis the search management team needs to be aware of:

- The current and planned business objectives of the company
- How search could play a greater role in supporting these objectives
- The capabilities of the current search application(s)
- How they could be enhanced with the current development and search resources
- What could be achieved with short-term consulting and development support
- Potential roadmaps for technology investment, both commercial and open-source
- The risks of not having an appropriate level of search performance

The scale of the work involved is illustrated in Figure 5.

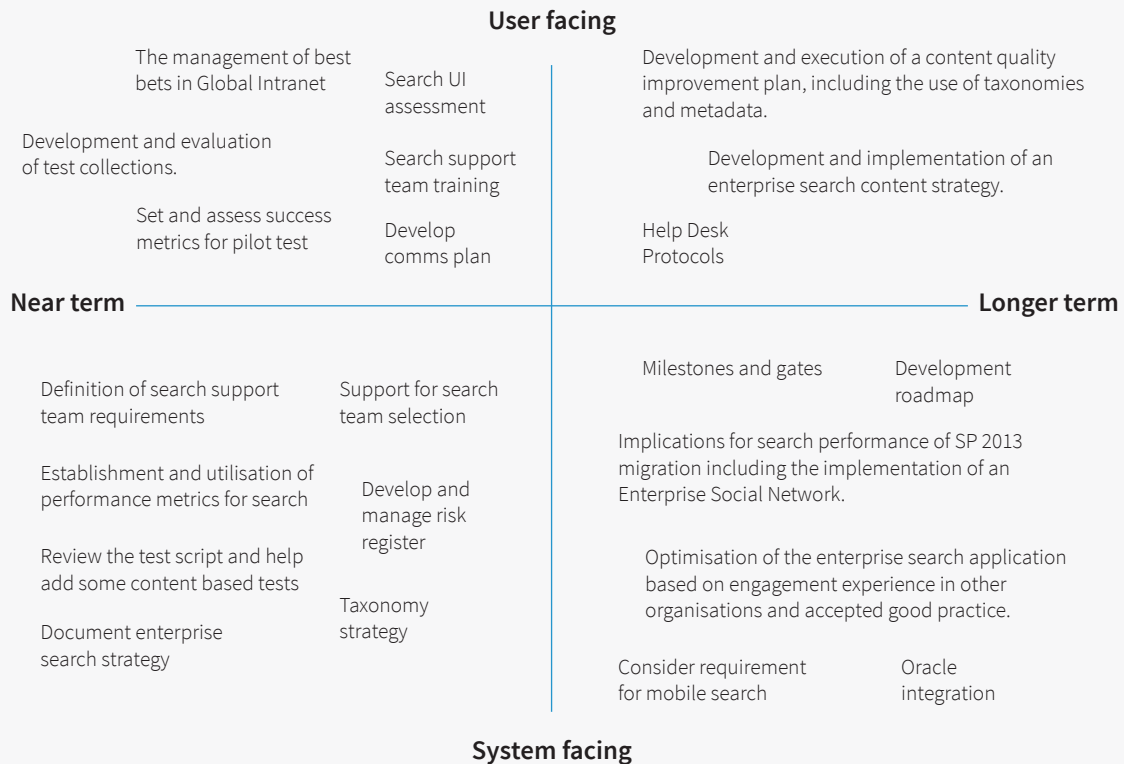


Figure 5 Search roadmap schematic - © Intranet Focus 2017

This matrix is adapted from one developed for a global professional services firm with around 4000 employees, and illustrates the range of actions that needed to be prioritised and resourced.

To illustrate the scale of the work involved in managing an enterprise search application in a global business, we are currently advising a client with over 120,000 employees on the enhancement of their SharePoint 2013-based search application. This company recognises the value of search and the roadmap leads directly into planning for a digital workplace for the company. The workplan for the Q2 2017 enhancements runs to almost 100 individual tasks. The company has invested in a search team of five in its corporate headquarters, regional representatives in major regional operations and support from three search systems integration and development companies.

Even then this team is fully stretched to respond to challenges and opportunities from the business and from users. Without a team of this size with the skills and expertise needed to optimise the search application it is highly likely that poor decisions will be made in a company that has a significant global reputation. It is a risk that the company is not willing to take.

Figure 6 below summarises the elements of a search application that can be used by a search team to enhance user performance. The challenge is not only knowing which ones to use to solve a user challenge but knowing what the potential impacts might be on the way in which search is delivered to other users. These impacts need to be assessed on an on-going basis.

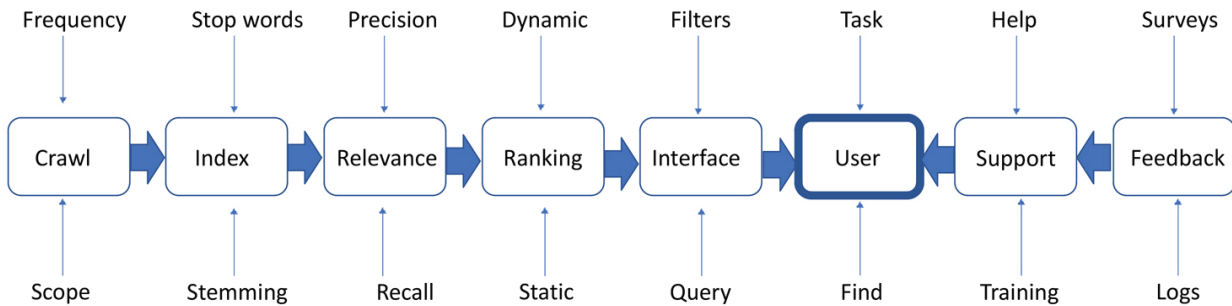


Figure 6 Optimising search - © Intranet Focus 2017

The critical success factor

Achieving high levels of search satisfaction is not a technology challenge but a people challenge. The critical success factor is this.

User satisfaction with search performance is a function of the number of people in the search support team.

Unless there is at least one full-time search manager, even for a small company, then search satisfaction will gradually decrease even from its low current level. The skills needed to support search are not just those related to server maintenance and support. Search applications are computationally quite complex, a result of being at the intersection of computational linguistics, applied probability and information science. None of these three elements are taught in computer science courses to any significant level. There is also a requirement for specialist skills in user interface design. Many companies therefore make use of specialist search integration companies to provide technical expertise as required, especially in the work involved in tuning the relevance ranking which requires a very deep knowledge of the search applications. These companies are also able to support the development of search-based applications.

This chart (Figure 7) comes from a survey conducted in 2014 by AIIM and shows the different categories of support that are needed for a search application.




Figure 7 Resource requirements for search - © AIIM Search and Discovery Survey 2014

There are four primary roles that need to be supported in a search team:

Search Development Manager

- Develops and owns the corporate search strategy
- Sets out standards for service delivery which need to be met by the technical architecture of the search applications
- Sets out the scope and structure of the search evaluation programme
- Leads the search team
- Maintains close liaison with the business through the governance committee
- Takes responsibility for application development

Search Information Specialist

- Defines business and user group requirements and responds to issues raised through search logs and user feedback
- Establishes taxonomy and metadata schema
- Defines information quality standards and guidelines in line with an information management strategy

Search Analytics Manager

- Responsible for defining search analytics requirements
- Integrates the outcomes of search analytics with other user feedback
- Reports on search performance on a regular basis

Search User Support Manager

- Undertakes training and usability testing
- Responsible for the search help desk
- Manages search communications through social media (blogs, wikis etc)

The support requirements are significantly greater when enterprise search is rolled out globally. There is likely to be a need for an Information Specialist for each major content language to identify any issues arising from poor stemming performance and inappropriate metadata tagging. This may not be a full-time position but certainly the expertise needs to be available to the search team. For similar reasons a good case can be made for an analytics specialist for each business area in a highly diversified global corporation. The search terms used for healthcare products will be different from those for pharmaceuticals.

Finding people with the skills in metadata and taxonomy management can be very challenging as these are very specialised skills. Often these skills are needed on a periodic basis rather than on a day-by-day basis, which is why creating a virtual search centre of people with these skills can be a very beneficial investment for an organisation of almost any size.

Ideally there should be Search User Support Manager in each major country, or at least each region (Europe, Asia/Pacific, North America) and language issues must be borne in mind. Although people may well speak several languages in business situations, they will prefer to search in the language in which they have the best command, so Spanish search and support in South America is very important.

Faced with this situation companies are beginning to appreciate the benefits of establishing a Centre of Search Excellence (CSE) and in principle this can help in bringing together staff with specialist expertise that may not be available in all business centres. Although staff in Information Specialist and Search Analytics roles may not need to be located in the countries they are supporting, Search Support Managers should be. Larger companies are now moving towards a two-stream programme for search development.

Stream 1 (Operations) supports the ongoing operational effectiveness of search. Multiple search evaluation techniques are used to identify potential search issues, ensuring that an early warning is gained of potential search issues. In addition, Stream 1 also assesses and prioritises user and business requirements.

Stream 2 (Development) where requirements identified in Stream 1, but which are beyond the skills and time of the search team to address, are put out to specialist consultants and contractors.

The Global Search Manager acts as the coordinator of the two Streams and reports to the budget holder for search and to a Search Governance Committee.

Quantifying the search team investment

It is quite difficult to define specific team sizes because much will depend on the diversity of the business in terms of customers, lines of service and global reach. The chart below (Figure 8) proposes an investment quantification based on the number of employees, excluding those working on back-end support for the search architecture. This is because this technical staff will be working across several different applications. This work can be outsourced to contract search system integration companies.

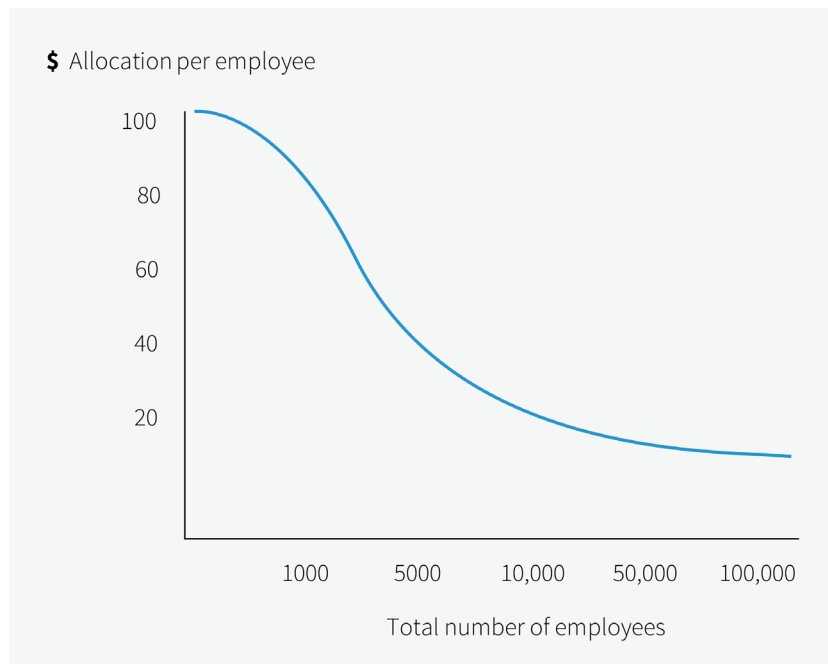


Figure 8 Indicative spend per employee on search team resources © Intranet Focus 2017

Experienced search managers are not easy to find. There are very few training courses available and most computer science courses only include a short module on search technology, making little reference to business aspects of search implementation and management. As a result, salary expectations can be high. In the USA and Europe the salary band for an experienced search manager could be approximately €100k to €150k.

Search team or search technology – which comes first?

The evidence suggests that almost all search applications will work substantially better when there is a search team with the skills and experience to enhance search performance to meet business and user requirements. These skills are so important that any company considering upgrading or replacing their current search applications should not do so unless there is at minimum a Search Development Manager to define the technical and user requirements and manage the selection and implementation process.

SharePoint search is no different

SharePoint is never acquired primarily because of its search capability. Often it is not until SharePoint 2013 is implemented and live that the requirements for search management become obvious. This is especially the case when there is no experience of running the powerful FAST Search Server for SharePoint 2010 application in the Enterprise License version of SharePoint 2010. Without the same attention to resources for search as for the document management and collaboration functionality, users may quickly become frustrated with the performance of the search application. The downward descent into search dissatisfaction can happen very quickly and is very difficult

to arrest. Decisions about search implementation in SharePoint 2016 and Office 365, including the cloud/hybrid/on-premise decisions, should not be made based purely on nominal performance and cost reduction.

Benjamin Nialun is a SharePoint expert. The section below comes from a longer blog post he published in 2015. The post is worth reading not only for Benjamin's advice but for the long string of comments that it generated from the SharePoint community.

<https://en.share-gate.com/blog/no-you-are-wrong-sharepoint-doesnt-suck>

The assumption that simply running SharePoint Search without any optimisation or configuration for your organisation will somehow make it instantly as good as "Google" or "Bing". Yes, because those two magically work, they turn it on and somehow it does what it does, no one is actually working to make it better each day. All right, I know, it's unlike me to get angry but as I mentioned I love SharePoint Search and this needs to be addressed.

If your Search Service Application configuration still includes the default Content Source to crawl and nothing else or if you've never heard of or created a [Managed Property](#) in your farm, then yes you're not going to find what you're looking for.

In all honesty, if it were possible I would tell you to have someone working on the SharePoint Search at all times. To make sure people can find what they are looking for, you need to make sure the Search can crawl it properly and display it properly as well.

This means:

- Configuring Crawls and Content Sources properly based on priority and needs
- Creating the right Result Sources to allow users to scope their searches
- Result Types allow results to be adapted visually based on business conditions. "Show Invoices very differently than regular Word documents"
- Managed Properties are shells that include more than one property found in your environment, allowing you to tell the Search that "Customer and Client" means the same thing regardless. Also used in Search Web Parts to display or filter by. It's like understanding columns, you NEED to know how this works.
- Display Templates are HTML designs you build, completely reusable, that show Search Results in whatever design you wish. Look at my [animated menu using Search Results and a Display Template](#).

These are just a few things that you need to manage if you want the full Search experience, then I promise you that users will not only find what they are looking for, but love using it.

©Benjamin Nialun 2015

Monitoring search performance

The complexity of search applications and of the range of content they are managing means a more thorough approach to monitoring search performance is required. It is quite common to find that the business owners of search, usually in the IT department, have little information about the extent to which search applications are meeting even the most basic requirements of employees. Few companies invest in gathering the information that is essential to assessing and then enhancing search satisfaction.

Table 1 lists nine categories of search evaluation techniques. A combination of most, if not all, of these is essential if employees are going to feel confident that they will find the information they need to make good business decisions.

Parameter	Techniques	Benefit
Quality	Identification and removal of redundant, obsolete and trivial content (ROT) Explicit titles on all documents Date and author information Depth and consistency of meta-tagging	Users will be able to put greater trust in the information as they make decisions.
Capacity	Crawl and index schedules and crawl consistency Connector performance Latency on query performance Latency on result presentation	Although it may be very difficult to respond as quickly as Google that will be the expectation of users.
Usability	Accessibility review Filter and facet performance Heuristic and expert walk throughs One-on-one usability testing	This will ensure that all the filters, facets and other search aids are seen as intuitive in action and outcome to users.
Performance	One-on-one usability testing Precision vs recall assessments Repeat query tests Session length Session stopping strategies Test collections User feedback	Users need to feel that they are getting a search experience which meets their personal expectations even though in reality trade-offs are being made.
Analytics	KPIs Promoted content click-through Search log analysis	Search logs highlight issues and opportunities to improve the search experience. Key Performance Indicators validate that the search application is meeting organisational objectives.
Assessment	After action reports Feedback from users Focus groups Job evaluations	Although these are qualitative metrics gained from a small number of users that have specific search requirements they complement the outcomes of search log analytics in achieving a high level of search satisfaction.
Satisfaction	Employee engagement surveys Feedback Focus groups User surveys	These high-level satisfaction surveys give a good indication of year-on-year trends in requirements and satisfaction.
Impact	Decision analysis Exit reviews KPIs Narratives Onboarding reviews	These techniques ensure that search is supporting the achievement of business objectives.
Benefit	Narrative KPIs	Stories of search success and failure can be very powerful when combined with high-level KPIs to enable the leadership team assess search effectiveness.

Table 1 Search evaluation techniques - © Intranet Focus Ltd 2016

The return on search team investment

In the 2015 Enterprise Search and Findability survey Findwise correlated user satisfaction against the maturity of search management in the organisation. In their view a mature organisation was one in which there was an enterprise search strategy and a search team that had the skills, experience and budget to:

- Develop and manage taxonomies and metadata
- Define and support information quality guidelines
- Monitor search performance on a regular basis
- Respond effectively to changes in user requirements
- Define and report on Key Performance Indicators

Overall the percentage of organisations which were at levels 3 – 5 in terms of search satisfaction totaled 44%. In the case where the above characteristics were met the level was 88%. In reality the gain was more than a doubling in the level of satisfaction as the All category included the mature organisations as well, so inflating the satisfaction scores.




Figure 9 Benefit of adopting a mature approach to search support © Findwise Enterprise Search and Findability Survey 2015

Recommendation

If you are in any doubt about whether search works for your organisation then ask your employees. It is not just a question of how many times they use search. In the case of search 'hits' is an acronym for How Idiots Track Success. It is about whether they feel they can trust the search application to deliver the information they need to make informed decisions. If the answer is in any way negative then you are putting your organisation's performance and reputational risk on the line and wasting a substantial amount of your annual salary bill in creating information that no one can find.

Compared to the investment your organisation has made in creating information the investment in a search support team will be very small. The return on this investment will be very significant indeed.

The author

Martin White

Martin White is an intranet, enterprise search and information management strategy consultant. He founded [Intranet Focus](#) Ltd in 1999. He has undertaken projects or managed business operations in nearly 40 countries.

He has written books on information consultancy; the selection of content management software; intranet management; and managing enterprise search applications. His latest books include the second edition of [Enterprise Search – Enhancing Business Performance](#) (published by O'Reilly Media) and [Managing Expectations – Building Client/Consultant Partnerships](#) (published by Intranätverk).

An information scientist by profession, Martin has been a Visiting Professor at the [Information School, the University of Sheffield](#) since 2002. He writes a monthly column on enterprise search for [CMSWire](#).

Prior to setting up Intranet Focus he held senior management positions with Reed Publishing, International Data Corporation and Logica plc.

12 Allcard Close, Horsham, West Sussex, RH12 5AJ,
United Kingdom.

