

Simplicity and enterprise search

A new model for managing your enterprise information



SIX WAYS THAT POOR SEARCH WASTES COMPANY TIME AND RESOURCES:

1. Time lost to ineffective search
 2. Time and money lost to administration of search systems and data (both IT staff time and maintenance contracts)
 3. Time spent tweaking and weighting documents to satisfy the requirements of complex systems
 4. The lost value of missing company information
 5. The lost value of undocumented employee knowledge
 6. Revenue lost through delays in time-to-market
-

In today's business environment, time and information are arguably our two most precious resources. Yet most businesses squander both on a daily basis.

- According to research firms IDC and Delphi Group, the average knowledge worker spends about a quarter of his or her day looking for information.
- Add to that the time spent by IT personnel and various specialists to manage company information – and the technologies used to store, organise and locate it.
- Figure in the time spent (or not spent) by employees to create documents that will match the exacting criteria of high-overhead information retrieval systems.
- Consider the wasted value embedded in documents that have been forgotten, mislabeled or put where no one can find them – at least, not fast enough to do what needs to be done.
- Think of the information that stays locked in employees' heads – or on their desktops – because it's "too much work" to publish where colleagues can find it and put it to use.
- Finally, tally the lost revenue resulting from delays in time-to-market.

The high costs to a company of not finding information, or of finding it too late, include faulty decisions, duplicated efforts, lost productivity and missed opportunities. Their impact can cascade throughout an organisation. The personal frustrations involved are also considerable – as anyone who has tried and failed to find a slide or a spreadsheet in time for a vital meeting can testify. The costs can be even higher in fields such as health care, the pharmaceutical industry and life sciences, where up-to-date information and time-to-market are competitive essentials.

The solution: As easy as 1, 2, 3

The symptoms are complex, but the prescription is simple: Managers and administrators need tools to connect employees with relevant information quickly and easily. Three things are required:

- 1. Fast, accurate search results.** To be successful, enterprise search must be powerful enough to deliver the most relevant information, consistently and efficiently, whenever and wherever it's needed.
- 2. Minimal administrative overhead.** Enterprise search must be quick enough to deploy and easy enough to manage that the cost of installing and maintaining it won't exceed the benefit.
- 3. An intelligible user interface.** Enterprise search must be simple and effective enough that users will actually use it.

“Make everything as simple as possible, but not simpler.”

– Albert Einstein

Simplicity without sacrifice

Search quality: deliver the goods

To fully realise the value of the information assets your business creates:

- Information must be readily and reliably accessible to everyone who’s entitled to view it.
- The information delivered must be current and relevant (the user needs the right document, usually in the most recent version).
- A clear, accurate ranking system should guide users swiftly and accurately to the data they need.
- Your intranet search should put your whole organisation on the same page, providing a consistent view of information across your company, while keeping sensitive documents secure.

Usability: keep your users happy

Unless your employees actually use the search tools you provide, your search technology – and the information assets your employees would have found and used – are both wasted investments. To deliver the increased productivity you seek, the solution you deploy must win acceptance; to accomplish that, it must be both effective and easy to use. People use tools that get results and avoid tools that don’t deliver. So you need power under the hood – even if users never open it.

As Mark Gallagher, manager of the intranet communications team at Bank One Corp., told InformationWeek, “If employees search for a term and don’t get what they’re looking for right away, they write off the search engine.” In deploying the Google Search Appliance at Bank One, Gallagher found that adoption was driven both by the simplicity of the interface and by the relevance of the results. As the article explained:

What makes the Google appliance especially effective is that it works like the consumer version – the Web’s most popular search engine – right down to a Google branded search window on the company intranet. The tool’s power is clear: Since Bank One went live with the system in September, the number of employee searches per day rose from about 4,000 to nearly 7,000. (InformationWeek, January 20, 2003)

Here the stereotypical “choice” between simplicity and performance turned out to be a false dilemma: Both were essential. Simplicity, familiarity and trust drove initial adoption; delivering results that matched users’ high expectations kept usage high as well.

Complex problems, simple solutions

Innovation and simplicity are often the best way to attack a complex problem. The first Google search engine – built with 30 off-the-shelf PCs running the free Linux operating system – is a case in point. Google’s design coupled innovative algorithms with a clustered approach to hardware infrastructure that capitalised on the falling prices of PCs, disk drives, memory, bandwidth and data centers, and on the availability of continually faster, cheaper processors. The open source Linux operating system was chosen for similar reasons: It was well supported and reliable, could be customised at will and cost nothing to use. Grid computing enabled the modular, scalable framework into which these elements fit.

THE CRITERIA FOR GOOD ENTERPRISE SEARCH:

Fast, accurate search results.

To be successful, enterprise search must be powerful enough to deliver the most relevant information, consistently and efficiently, whenever and wherever it's needed.

Minimal administrative overhead.

Enterprise search must be quick enough to deploy and easy enough to manage that the cost of installing and maintaining it won't exceed the benefit.

An intelligible user interface.

Enterprise search must be simple and effective enough that users will actually use it.

This whole infrastructure provided a robust ecology in which Google's data indexing and retrieval algorithms could thrive. The result was simple, powerful, flexible and highly scalable – as evidenced by the fact that Google's architecture remains essentially the same now as then, though with about 1,000 times as many machines.

Google in a box

To extend this search platform to the enterprise, Google once again used simple solutions to address complex problems. The technology that powers Google.com was put to work in a straightforward package: a plug-and-play search solution integrating hardware, software and support.

The Google Search Appliance crawls all platforms without distinction, capturing data on highly distributed, heterogeneous networks in a single coherent view. Intelligent algorithms automatically detect network settings; heuristics recognise date formats, languages and spelling mistakes.

This innovative, automated approach achieves an important operational and financial benefit: Even as it simplifies search for end users, it takes the burden of organising information off the administrator. Optimised for compatibility and easy, rapid deployment, "Google in a box" presents a simple face to users and administrators alike – without sacrificing the power and sophistication of Google search technology.

Three pitfalls to avoid

"First, do no harm," Galen advised his fellow physicians. His dictum applies equally well when choosing an enterprise search technology. The last thing your users need is a new set of obstacles masquerading as a solution. As you seek to expedite the flow of information in your company, here are three ways to avoid making things worse:

- Choose a search solution that will index documents without adding overhead, either for document creators or for administrators.
- Remove obstacles to document creation wherever possible.
- Make sure you know what's actually on your servers.

Minimise overhead with come-as-you-are indexing

A come-as-you-are approach to indexing eliminates the overhead of preparing documents for admission to the body of searchable data. In any case, your data shouldn't need a laborious makeover for your search solution to provide relevant results. Manual weighting and tweaking of indexed documents may have its place, but effective search should not depend on it. For one thing, algorithms scale better than humans. For another, software has no issue with boring, repetitive work; people do, and quite often they simply won't do it. Eliminating such requirements reduces the chance that useful information will remain unpublished or unindexed, and thus never become search-accessible at all.

Publish or perish

In fact, when a document is not found by a search user, it's often because it isn't there. It's therefore also crucial to remove obstacles to document creation wherever possible.

Among other things, this means that users should be able to create a document without doing extra work to make it searchable. Otherwise, a large amount of the intellectual value that your organisation generated will remain locked in the heads of employees where no currently available algorithm can touch it. Most

Your data shouldn't need a laborious makeover for your search solution to provide relevant results.

organisations are plagued by this problem to one degree or another – as they discover when laborious makeover for your search solution to provide relevant results. someone in a key role gets sick, goes on vacation, leaves the organisation, or is simply unavailable when crucial data is needed.

Making it simple to publish information where others can find it is a great way of making sure that it's accessible when it matters most. With rare exceptions, good search shouldn't depend on knowledge workers to make their work system-friendly. They should be free to focus on creating useful information and putting it where colleagues can realise its value. Workers whose work has to match stringent indexing criteria are far more likely to publish nothing at all. It's consequently important to choose an intranet search solution that lowers the bar for publication rather than raising it.

Know what's on your servers

Here's a message for your CIO: You need a clear and comprehensive view of the documents on your company websites. If you don't know what's there, you have three important problems beyond the obvious question of productivity:

- **Potential liability.** For instance, inappropriate content on your network can trigger complaints – or even litigation – by employees with “hostile workplace” concerns. Similarly, material on your servers in violation of intellectual property laws can expose your company to liability.
- **Security.** You can't control access to information if you don't know where it is, or whether you have it at all. In a number of instances, companies implementing the Google Search Appliance have discovered hundreds of sensitive documents on their servers of which they were totally unaware.
- **Policy.** Inaccurate or anachronistic information can hurt your business. But to keep obsolete product data and outdated price lists out of your customers' hands, you need accountability. And you can't delegate authority – or insist on follow-through from your team – until you (and they) know what's actually on your servers.

Besides putting information into users' hands at the moment they need it, effective search gives you a detailed record of what's on your system – enabling you to monitor the information available and ensure that sensitive information is only available to those who should see it. An effective search solution will make your documents fully searchable, while restricting the delivery of sensitive search results to those with appropriate permissions.

Administration

Simplicity doesn't only benefit end users. Complex systems, by their nature, are costly and time-consuming to maintain and reduce administrative flexibility through the limitations they impose (often through their incompatibilities with other systems). In addition, the upfront costs of implementing a new technology can be staggering – encompassing, as they usually do, three kinds of overhead:

- Installing, configuring and testing the system itself
- Preparing data and other resources for use by the new system
- Training users and evangelising adoption

HOW MUCH DOES SEARCH TIME COST YOUR COMPANY?:

Multiply your number of users by the figure below that best represents your current level of search:

- £1,641 per employee – least usable intranet
 - £1,116 per employee – average usable intranet
 - £843 per employee – most usable intranet
-

To be cost-effective, it's essential that initial overhead be minimised; that administration and maintenance be automated to the greatest extent possible; and that the system be modular and scalable to avoid elaborate redesign and retrofitting as objectives change, volumes of data increase and the user base expands.

Usability

Usable search increases adoption

You can spend a fortune on technology that purports to optimise the flow of information in your organisation. But how can you be sure that employees and other end users will actually use the systems and processes you put into place? If a system goes unused – or, worse, presents an obstacle to the effective flow of information – your investment is wasted at best.

Usable search boosts ROI

To work as effectively as possible, information workers – and at this point, that means almost everyone – need fast, easy, accurate access to existing data.

“Search is one of the most important functions on any intranet,” says *Designing Usable Intranets*, a November 2002 report by the Nielsen Norman Group based on intranet usability studies of 14 corporate intranets. “Finding information through the navigation should be simple, but when it is cumbersome or fails entirely, users turn to search. It should be easy for users to find the search capabilities, construct a search query, search in the right places, and deal with the results.”

The report underscores the importance of effective search by noting what happens in its absence. Poor search, it observes, was the single greatest cause of reduced usability across the intranets studied. “Search usability accounted for an estimated 43 percent of the difference in employee productivity between the best and worst intranets.”

The study also found dissatisfied search users turning to expensive and inefficient alternatives (“The results are really gibberish. I’d stop now and use the phone”) or giving up altogether – which, for such users, brings the ROI for the unused online assets down to zero. Disappointed by the irrelevant results returned by one company’s intranet search, “Some users had stopped using it entirely. One user said, ‘I’ve tried using the search and I think it’s terrible. I don’t think I’ve ever successfully searched for anything and found what I’m looking for. If I tried anything now, I bet it wouldn’t work.’”

Measuring users’ performance for 16 common tasks across the 14 intranets studied, the report found wide variations in usability, with corresponding variations in cost. After accounting for salaries and overhead, a company with one of the least usable intranets in the study would spend £1,641 per employee annually to cover time spent on the 16 tasks measured, while one of the most usable would cost only £843 per year.

Extrapolating from these and other figures, the report estimates the total annual cost of intranet use at various levels of usability for companies with 10,000 users:

- **Good usability:** £8.4 million annually
- **Average usability:** £11.2 million annually
- **Poor usability:** £16.4 million annually

“Search is one of the most important functions on any intranet. It should be easy for users to find the search capabilities, construct a search query, search in the right places, and deal with the results.”

Designing Usable Intranets
Nielsen Norman Group, Nov 2002

A company with 10,000 users would need to invest about £270,000 to move from one level to another. The report therefore calculates the ROI for improving intranet usability as ranging from a factor of 20 for the lowest group to a factor of 10 for intranets of average usability. “Clearly, the biggest gains are to be had from improving bad and horrible intranets,” the report explains. “These low-quality designs are also the easiest to improve, because they are going to have some very low-hanging fruit.”

With search usability accounting for about 43 percent of the productivity delta between most and least usable, it’s obvious that improvements to search will generate a significant share of that return. The report’s #1 suggestion for improving search usability: invest in a good search engine. The comforting corollary is that if you choose the right one, it will pay for itself.

Motivate adoption

Ideally, employees will use intranet search frequently (because it’s so effective) but briefly each time (for the same reason). Adoption will suffer if your search:

- Is unintuitive, unfamiliar or otherwise difficult to use
- Returns poor, unreliable or inconsistent results

Interface simplicity is a virtue – provided it really translates to ease of use and not just obscurity. Familiarity also promotes frequent use, as the city of San Diego found when it implemented the Google Search Appliance in place of a previous information-retrieval system that wasn’t giving users the results they needed. Partly because city employees were already so familiar with Google.com, reports InformationWeek:

The city opted for a single Google server with a license to search an index of up to 150,000 documents. The result has been a welcome improvement for the city’s 8,000 computer-equipped employees and its nearly 250,000 unique monthly site visitors: Cull says employees are using stuff they didn’t know existed, and citizens are sending email about the search success they’re having.

“When we saw it in action, we found that it was pulling up Documentum files, our Sun One portal, and even database applications,” he says. Now, employees are able to enter “GroupWise” in the search engine and find a link to their web-based email access and find check-stub data by typing in “E-Pay.” And the city has been able to simplify other tasks, such as finding related documents with data that previously would have been duplicated or populating an online calendar by using Google to pull items from the city’s events database. All this, and the server and software were up and running in 30 minutes. Says Cull, “It’s hard to measure the value to the city.” (InformationWeek, January 22, 2003)

Of course, as noted above, a significant part of that value is measurable – and in dollars and cents. Hence the evaluation by Mark Gallagher, Bank One’s intranet communications manager, who calls the Google Search Appliance “underpriced.”

Kaiser Permanente: usability success story

When Kaiser Permanente, America’s largest not-for-profit HMO, developed a clinical-knowledge portal for its 50,000 doctors, nurses and other caregivers, search was a key part of the plan. The portal, available from anywhere in the Kaiser wide-area network, gives medical staff access to diagnostic information, publications and other clinical resources. Putting the right information quickly and easily into caregivers’ hands is essential to the clinical portal’s success.

“Right out of the box, without any tweaking at all, the Google Search Appliance was more effective than the system we’d been working on for a year and a half.”

Brad Hochhalter
Director, Kaiser Clinical Portal

Brad Hochhalter, director of Kaiser’s clinical portal, initially turned to a system that required manual tagging to power the clinical portal’s search. But after spending almost two years trying to optimise the costly, complex system, it still wasn’t delivering the expected results. Manually weighting and indexing the portal’s content to produce relevant search results became a laborious undertaking, and even then Hochhalter found the results iffy. “Because of the enormous challenges of trying to organise medical content in a systematic way, we weren’t getting the results we needed quickly enough.”

The deluge of irrelevant and unreliable results also made caregivers hesitant to use the portal’s search at all. “I was getting soaked on search, and I wasn’t getting results,” Hochhalter notes. “I wasn’t getting anything better for the time and money I was putting in. It just wasn’t worth it.”

“Right out of the box, without any tweaking at all, the Google Search Appliance was more effective than the system we’d been working on for a year and a half,” says Hochhalter. Clinicians now search the site in situations that run the gamut from leisurely research to urgent care, from the exam room to the emergency room. Doctors and nurses use the Google Search Appliance to help them reach diagnoses and specify treatments, check the side effects of new medications, and consult clinical research studies and other medical publications. Physicians can also search the Kaiser intranet from home when they’re on call in order to give immediate guidance more easily.

Hochhalter credits the Kaiser staff’s quick acceptance of the new search partly to its speed and accuracy, and partly to their familiarity with Google. “We didn’t make any tweaks at all to the results page – we left it exactly like the pages on Google.com. Frankly, there’s a certain cachet with Google that helped me get people interested.” In the months since Google search went live at Kaiser, the frequency of searches has risen rapidly. “It’s clearly better than it ever was before,” says Hochhalter, “and the number of searches is going up significantly each month. There was a 30 percent increase from January to February alone.”

Conclusion

In today’s business climate, no organisation can afford wasted time, lost value and missed opportunities. To capture the value and optimise the flow of knowledge in your organisation, you must:

- Remove the obstacles that stand between your team and the data they need – obstacles that waste your information assets and hamstring your productivity.
- Implement a system that’s easy to install and maintain – and remove the burden of organising information from administrators and end users.
- Choose a solution so simple and effective that people will actually use it.

Connecting people to the relevant information they need – quickly, easily and accurately – empowers your users and offers your organisation significant bottom-line benefits. Enterprise search that intelligently integrates usability and power will boost your productivity and put your intellectual capital to work. For once, a complex challenge has a simple, cost-effective solution: Fast, easy and accurate search.



Google, Inc. 1600 Amphitheatre Parkway, Mountain View, CA 94043