Orbital Organik

White Paper

Published: July 1997

for the latest information, see http://www.orbital-tech.com

Transforming Information into eKnowledge

Productivity Gains from IT

Spending on IT in 1996 in the US amounted to $213 billion on computer hardware alone - when software, networks and IT services are included the investment rises to $500 billion in the US and $1 trillion worldwide. And yet productivity growth does not seem to be rising in line with this investment. In fact, it has fallen precipitously over the past 30 years from an average of 4.5% growth per annum in 1960 to about 1.5% in recent years. Surprisingly, knowledge-based companies, who also happen to be the biggest spenders, have been hit the hardest in the productivity slow down.

The productivity gains experienced have tended to be confined to relatively low value transaction processing applications such as trade automation, order processing, stock control etc.. As valuable as it is to automate these activities, and most businesses could not function now without these, there are diminishing returns on further IT spending in these areas. Cost reduction in these areas will very soon reach the point where the incremental costs of the new equipment will exceed the marginal return.

Organisations are now turning their attention to applications of IT that will result in significant returns on investment and enhancement of their competitive edge. These new mission critical applications of IT are in areas of white-collar commercial activity, for example in sales, marketing, management, professional services and so on.

These commercial activities, however, have proved to be very resistant to automation or support from IT. Where support from IT has been provided it has generally been at the task level, using word processors, spreadsheets etc., to enable an individual to accomplish a specific task faster and with, hopefully, higher quality. This task support doesn't come without significant costs though. Gartner Group has calculated that, once all PC costs are aggregated, companies are spending, on average, $13,000 per year per PC. Clearly, there is enormous scope to raise the levels of return on this IT investment.

Activities of the Knowledge Worker

Higher returns can be achieved by focusing on the ways in which workers spend their time when not occupied at their PC engaged in their personal task support activities. These activities, not directly supported by IT, account for much of the white-collar worker's time and result in most of the creation of value for the business. Obviously these activities vary from person to person

and from industry to industry, but a great deal of time is spent in simply finding answers to questions. Given the processing power on users' desks and the fact that many of these computers are networked together there is enormous potential for increasing productivity simply by connecting the right people with the right information at the right time.

Since these collaborative, information gathering activities account for such a significant proportion of a worker's time it is precisely on these activities that attention should be focused. If the time taken to locate and obtain the right information could be significantly reduced then more time could be spent on value-added activities. Additionally, if the information located is also of higher quality and relevance then less time will be spent on remedial activities due to the use of incorrect or inappropriate information.

Organisations now understand that to derive the greatest possible returns on their IT investment they need to invest in applications to improve white-collar workers' productivity. This is where most of the value is created in a company and also where most of the costs reside. There are enormous returns awaiting those organisations that can improve white-collar productivity at the same time as reducing white-collar costs.

Enhancing Competitive Edge

When focusing IT investment on the activities of knowledge workers, as opposed to clearly definable and automatable transaction processing applications, organisations are creating the opportunity to extend their competitive lead. After all, a company exists and is successful because it carries out its business in a manner different to its competitors. It is this unique *modus operandi* that distinguishes one company from another, generates the competitive lead and, in comparison with the competition, results in higher quality products produced at a lower cost, more effective sales and marketing and better customer service.

However, as much as companies would like to closely define all aspects of its operation this is neither possible nor desirable. The actual operation of the business is a result of a collaborative effort by many co-workers, who may be guided by procedures and operations manuals but who have also established their own informal ways of getting the job done. It is here that the intelligence and innovation of workers can be brought to bear on the company's operations and used for the company's benefit.

The worker's knowledge of these informal processes should be captured by the company since it would enable all workers to access the accumulating knowledge base of the company and find out what they need to know quicker. Additionally, it would enable the company to be more resilient to staff turnover - when a key worker leaves the company's employ their knowledge is not lost. Since this knowledge is a key corporate asset it is highly desirable to retain and share this knowledge.

Capturing Corporate Knowledge

Whilst much of this process knowledge and expertise still resides in people's heads, a lot has already been captured electronically in many forms. PCs and networks installed in many organisations have already, albeit in a rather *ad hoc* fashion, acted as repositories for significant amounts of the company's expertise. Networked applications such as Lotus Notes have also accumulated vast amounts of information about the company's activities. Company email systems also act as repositories for much corporate

knowledge, although, unfortunately, this information is very likely to remain locked away forever in private email accounts.

Added to these information repositories, over the last couple of years many companies have implemented intranets or internal web sites and have also given some, if not all, of their workers access to the Internet. The Internet also contains information that might be relevant when determining the solution to a problem. However, a key distinction must be drawn between information on the public network and information in the private, intranet. On the Internet, the truth or accuracy of the information must be viewed as suspect. Unless the information has been derived from someone known to the organisation or from the web site of a professional body or other reputable organisation it should be treated with extreme caution.

Although the information worker spends a great deal of their time trying to find the right information, their efficiency is not necessarily improved by simply making more information available. Very often, it is not a lack of information that constrains productivity but a superfluity. Without the right tools it is difficult for the worker to extract from that information haystack the needle that would answer their question.

Extracting Value from Corporate Knowledge

In the worker's quest for the 'right' answer to their question they turn to both human and electronic information sources. Unfortunately, knowledge of the company, its products and operation that only exists in peoples' heads is not very shareable and information that exists electronically is under-exploited because there are no efficient aids to enable the worker to sort the wheat from the chaff.

Search tools have been developed to enable organisations to extract some value from this mass of unstructured electronic information. However, merely furnishing a user with a search tool to locate information from the information resource by a keyword search, providing that the information can be indexed, is unlikely to deliver acceptable results. The reason is that the user is asking a question to which they do not know the answer. By merely searching for information in a brute force manner they will be returned a plethora of matches but will not necessarily be able to determine which piece of information contains the answer to their question.

Without the ability also to communicate the context in which the question exists a dumb search will return both irrelevant and inappropriate results. Occasionally the user may derive some value from the search results but will very often have to resort to the tried and tested method of calling someone on the phone, emailing them or walking over to their desk.

Consulting an Expert

Indeed it is this manual process of 'consulting an expert' which proves to be the model which, when applied to both human and electronic information repositories, would enable the worker to easily obtain the answer to their question.

Current Situation

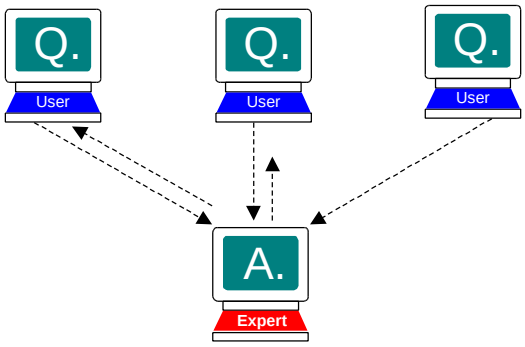

The worker may need to consult an expert in two possible forms. Firstly, the worker should be able to ask a question of an electronic expert who understands who they are and what would be a likely answer to their question given where they work and their question/answer history. Secondly, if this electronic expert should prove not to be able to provide an appropriate answer then the services of a human expert should be employed.

The term 'expert' is used to describe someone in the company to whom people would informally turn when faced with problems of a particular nature. Certainly these people don't have 'Expert' in their job title but people know who they are. They may have achieved this status because of some specialisation at work or experience derived from a previous job, or they may simply have a specific interest or skill in this area. For whatever reason the expert has achieved this status, their expert knowledge is informally used by co-workers.

Once the 'expert' has provided an answer which meets the questioner's satisfaction criteria, it should be recorded in a knowledge base along with the question and appropriate characteristics of the questioner for the benefit of future users. Unless this is done, the question/answer pair, together with relevant context, is not able to be shared and used by other workers with the same or very similar questions.

Thus an answer base or knowledge base would be accumulated which, over time, will automatically meet a greater and greater proportion of the questions asked by the users. This would enable the users to get their questions answered accurately, and this accuracy would increase over time since the expert would 'learn' what are likely to be appropriate and relevant answers. Less time would be spent chasing down the expert to obtain the answer and more time would be spent on more productive activities. The experts would benefit also, since they would only be required to answer a question once as their expertise would be captured by the knowledge base. This would also release more of their time to be used for the benefit of the business.

Organik

In order to achieve the twin aims of capturing corporate knowledge and increasing the productivity of white-collar workers, Orbital Technologies has developed Organik.

Organik is a product which enables the capture and accumulation of expert knowledge whilst providing the ability for users to exploit that information in their day to day activities.

Organik is an example of a new breed of software written 100% in Java. It is an open product supporting all the key Internet standards, including HTTP, HTML, JDBC, POP, IMAP, NNTP, CORBA etc.. This ensures that Organik can be seamlessly integrated into your exiting (or projected) Intranet with ease.

Organik is a server application that can be accessed by either a purpose-built client supplied with Organik or access to Organik can be integrated into any browser-based application. And since Organik is written in Java it is inherently scalable since it runs on any computer, from PC to mainframe, that supports the Java Virtual Machine.

Organik in Action

When a user needs the answer to a question they submit their question in natural language to the Organik server. This question may be submitted either through the purpose-built Organik client, or the custom interface integrated into their browser-based application or it can be submitted by email.

The identity of the user is known at this stage and is used to determine which answer out of the total possible set of answers is likely to be appropriate. For example, if Organik was employed in a power generating and distribution

Query into Organik

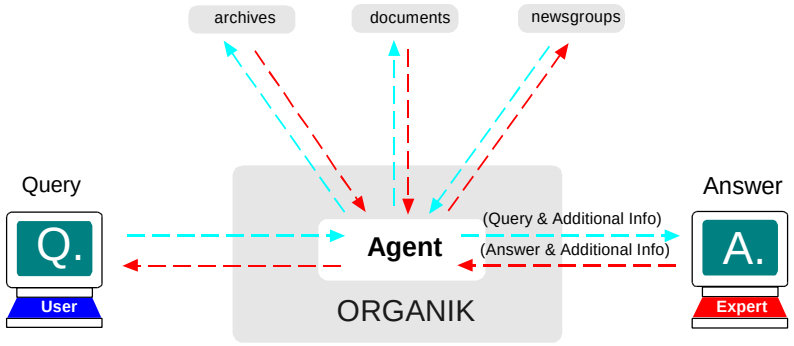

Diagram 1 Information flow through Organik

company and a question was asked about electricity pylons, the suitable answer for a user in Research & Development would be quite different to the answer for someone in the Customer Services department. This relevance rating is crucial if users are to achieve the potential productivity benefits and develop trust in Organik.

The security level of the questioner is also taken into account before answers are delivered to the user. This security level is determined either by the source information repository's security or by a security rating assigned by the expert when the answer is initially formulated. For example, it may be entirely appropriate for a senior manager to ask Organik which were the most

productive departments in 1997 but this same question asked by a less senior worker should not receive any answer.

If a suitable answer is available in Organik it is returned to the user who would then rate the relevance of the responses. This is used to improve the accuracy next time a similar question by a similar type of user is posed.

However, if there is no answer that satisfies the user then the help of an expert is enlisted. This expert receives the request to provide an answer but is also provided with a set of information resources that Organik believes would be useful in developing an answer. These information resources would be composed of potentially multiple distinct data sources, from existing Organik information, to Usenet newsgroups, mailing lists, links to network file server-based files, Lotus Notes databases etc.. However, although multiple sources are available, Organik has either indexed the content itself or makes use of the indexing engines in the source information repository such as the Lotus Notes full-text search engine. This capability liberates the corporate knowledge locked away in various disparate repositories and formats and provides an integrated view over all available information.

The expert then uses the information resources supplied by Organik to answer the question for the user. The question/answer pair together with relevant context is recorded within Organik's object store. The accumulation of knowledge means that the experts will receive progressively fewer and fewer questions. Indeed, by using Organik they will immediately notice a reduction in the repeat questions they are asked since users have access to the existing question/answer pair in Organik

An Organik Network

The ability to chain Organik servers together enhances the scalability and resilience of the knowledge base of the organisation. This allows an Organik server to operate specifically at a departmental level, which increases its ability to locate a relevant answer. However, a departmental Organik server would be able to enlist the help of other Organik servers in the network in order to determine the right response to a local user's question.

Organik's services can also be exposed to your customers via your World Wide Web site or by your Internet email connection. The customer can pose natural language requests to Organik either directly through their web browser or by email. If the interface used is the web browser, their question can potentially be answered immediately without using the company's customer service resources, thus allowing them to service non-web enabled requests more efficiently.

Organik Administration And Maintenance

Comprehensive administration and maintenance tools are provided which may be accessed from any web browser. These administration tools allow you to

- Integrate user directories into Organik from any directory service supporting LDAP in order to access user lists.
- Inform Organik of the user's department and location.
- Assign access rights to users and groups of users in the directory. These access rights are used to determine the security level of information created in Organik - information outside of Organik will have its own security mechanism. Organik will adhere to these existing access controls.
- Review existing answers to ensure that expired content can be either updated or removed.
- Report on question activity by topic/department or other appropriate categorisation in order to identify knowledge deficiencies and training requirements.

- identify external information repositories, (e.g. file store documents, Notes database, Exchange folders, mailing lists, Usenet newsgroups, specific web sites) which should be indexed directly by Organik.
- identify external search engines which should be employed by Organik when searching for relevant content for the support of the expert

What's The Difference Between Organik And A Search Engine

A search engine will index all information on the Internet or an intranet. It will not provide access to non-Internet or intranet information. This means that all the knowledge currently locked away in your existing files and databases are inaccessible to the search engine. This information has been developed over many years and represents a significant investment and significant value to the corporation. A search engine does not allow you to liberate this value.

Search engines also do not take context into account. It is context that allows irrelevant answers to be discarded. Without the ability to record and accumulate contextual information, the volume of irrelevant answers returned will increase over time as the volume of indexed documents grows. This will make it onerous for the user to sift manually through the answer set. The user is also asking a question to which they don't know the answer - so how is it possible for them always to determine which answer is relevant given a wide choice of possible answers? Without context the user will inevitably stop using the search engine.

If the search engine does not return a document that answers the user's question the user is forced into an inefficient (although extant) method of answering questions. Even when the user has received the answer, unless suitable mechanisms are available, this question/answer pair cannot be shared by co-workers.

For further information, please contact:-

Orbital Technologies Limited
90 Corbiehill Crescent
Edinburgh
EH4 5BG
Scotland

Tel: +44 131 538 8850
Fax: +44 131 538 8851

http://www.orbital-tech.com
info@orbital-tech.com