

Peer-to-Peer Knowledge Management: Overcoming Bottlenecks and Improving Information Quality

Andreas Neus Philipp Scherf

IBM Unternehmensberatung GmbH

Hamburg, Germany

{andreas.neus, philipp.scherf}@de.ibm.com

Abstract

Access to current, complete and relevant knowledge is a key competitive differentiator in the present economic market space. But most knowledge management today only shifts the traditional, expert-based KM processes onto electronic media. This paper discusses a paradigm shift from an expert-centric to a peer-to-peer approach to knowledge creation and management. Leveraging the lowered transaction costs provided by Internet technology, methods and tools of collaboration that have been invented and refined by the Open Source and Free Software community over the last two decades are now being adopted by early movers in the Knowledge Management space. This new approach, based on a peer-to-peer approach and open collaboration, has shown the potential to revolutionize the way knowledge is created, developed and managed. We outline the characteristics of the two opposing paradigms and present ways in which the peer-to-peer knowledge management approach is already being successfully used in practice today. We address how the quality of information is kept high without a traditional review/quality-check role by using a revision control system and distributing the task to all interested practitioners. Finally, we discuss four key challenges for introducing the new paradigm within companies.

1. Introduction

*“Science is the belief in the ignorance of experts.”
(Richard P. Feynman, Nobel Laureate, Physics [4])*

Several upheavals in the competitive landscapes in recent years have demonstrated that a defining characteristic of electronic business is the disruption of both current market equilibriums and current economic models by technology radically lowering transaction costs.[2] Phenomena like peer-to-peer file-sharing, networks such as Napster and Gnutella [12], or the success of Open Source software such as Linux (an “impossible public good” according to traditional economic theory) [16] [10] prove that the Internet radically changes the way we communicate and collaborate by reducing the cost associated with communication and collaboration

This paper discusses how a paradigm shift from an expert-centric to a peer-to-peer approach to knowledge creation and management has been enabled by using concepts, tools and methods developed in the open source and free software community over the last two decades.

Information today is mostly created, quality-controlled and disseminated in a closed, centralized fashion with a clear distinction between producers and consumers – just like software used to be.

Yet several high-profile initiatives using a peer-based, open source approach to collaboratively create and develop knowledge have already been launched, mainly in academic, educational and research environments (i.e. MIT OpenCourseWare, Harvard OpenLaw, Wikipedia [9] [5] [17]).

While these environments are obviously conducive to such an open model (with the scientific principle based on openly publishing results to enable replication and peer-review) this new approach can also be used inside a commercial organization that is challenged with a need to rapidly develop, update and share information. It is especially applicable when the knowledge about an issue is not (yet) a solid, coherent, structured body of information that can be easily taught, but rather still an evolving system of bits and pieces of knowledge, that has yet to reach maturity. Some people would argue that with the world around us growing more complex and changing ever faster, most of the knowledge fits the latter description.

Knowledge with these still “imperfect” characteristics is created, developed and managed in essentially open systems. We will discuss the four key obstacles for the realization of a Peer-to-Peer Knowledge Management system in companies with regard to:

- Information quality
- Issues concerning confidentiality
- The need for transparency and accountability
- The necessity to implement organizational and cultural changes

2. Peer-To-Peer Knowledge Management

Presently, the idea of Peer-to-Peer Knowledge Management is mostly a grass-roots movement, much the way the idea of open source software first entered corporations below the radar screen, based purely on its practical utility in solving certain problems.

One problem that has been known for some time is the aging of information both between the creation and initial publication and the “aging” of published information that becomes outdated because the world it describes relates to has changed.

The first challenge, which is especially felt in the rapidly developing fields of knowledge such as

biotechnology or particle physics, has been met by researchers publishing their papers first on pre-print servers, in order to get the information out in time to be relevant. The traditional publication, months later, is more pro-forma, as the active researchers have already read the electronic version. While this process speeds up the time to publication, it does not address the problem of information becoming outdated quickly. A method that has been developed to address this latter problem exists in so-called “Wiki Webs”s [8] for supporting collaborative, peer-to-peer knowledge creation and management, which have been increasingly appearing in research, but also corporate environments over the last few years. (Fig. 1)

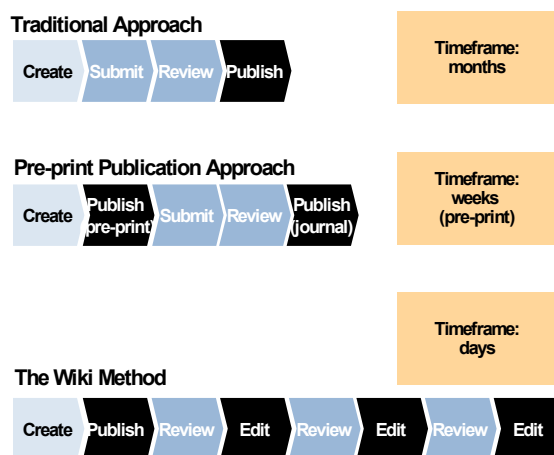


Figure 1: Three approaches to publishing knowledge.

A “Wiki” (Hawaiian for “quick”) is a collaborative knowledge management environment that eliminates the distinction between expert writers and lay readers, *producer* and *consumer* of information by making all users information “*prosumers*”.

In such an open, peer-to-peer context, every user can create and edit information without the bottleneck of a moderator or web-master having to approve or effect the changes. Quality is managed by keeping all changes in a detailed revision control system so that any unappreciated change can be undone by subsequent readers/editors. As an additional disincentive to low-quality contributions, the identity of the author of every change is stored and available to all users.

While such an open approach is certainly not a solution for legally binding or official information, it can be a good, evolving breeding ground for quickly documenting and formulating ideas together.

3. Paradigms: Brooks’ Law vs. Linus’ Law

In order to understand the issues connected with implementing such an open, peer-to-peer approach to knowledge management inside a company – and the related tricky organizational challenges – we compare and discuss several aspects of the two competing paradigms. These are exemplified by a number of opposing concepts. (Fig. 2)

Our traditional paradigm of collaborative development, such as *Brooks’ Law* (named after Frederick P. Brooks, author of “The Mythical Man Month” [1]) can be paraphrased as “*Many cooks spoil the broth*”. In accordance with this belief, only a small and select circle of designated “experts” should be allowed to create and improve high quality information, relegating the vast majority of practitioners to pure consumers of the information produced by the experts. On the other hand, *Linus’ Law*, named after Linus Torvalds, a term coined in the context of Open Source development [16], can be paraphrased as “*The more, the merrier*” and predicts that the more people are involved with (and spend their attention on) an information product, the higher the quality will be.

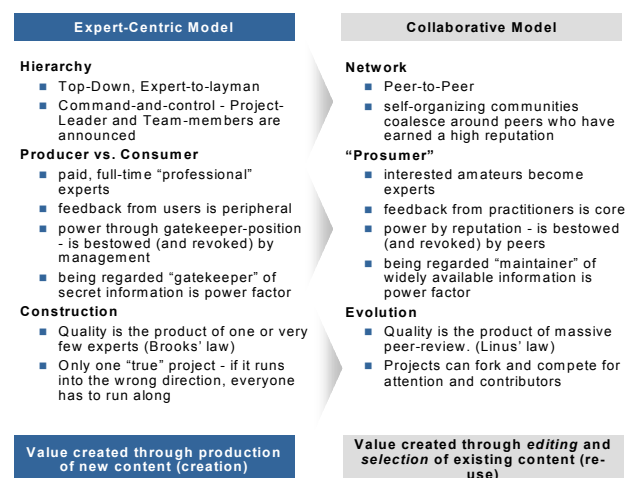


Figure 2: Expert-centric vs. collaborative model

In essence, the traditional, expert-based approach to Knowledge Management is a “construction” or “perfection” approach. Consistent with the high transaction costs for information in previous years, a static, “perfect” document was to be created before the expensive reproduction process and effective publication commenced. However, given the negligible cost of electronic publication, the open, peer-to-peer approach relies on “evolution” or “improvement” over continuing releases rather than initial perfection [11]. Eric Raymond calls this policy “release early, release often” in his famous essay *The Cathedral and the Bazaar* [14].

In the context of knowledge, Ives & Jarvenpaa predict a “revolution in knowledge creation”, as the review process becomes *more open* and the documents *less static*. “*In the past, journals were archived in research libraries where they remain unchanged. But living web documents are considerably more volatile. Simple errors, typographic and otherwise, can be quickly repaired before they are inaccurately cited.*” [7]

Other opposing concepts in this context are:

Traditional approach vs. Open, Peer-to-Peer approach

Brooks' Law	vs.	Linus' Law
Hierarchy	vs.	Network
Experts	vs.	Peers
Teams	vs.	Communities
Cathedral	vs.	Bazaar
Construction	vs.	Evolution
Perfection	vs.	Improvement
Information scarcity	vs.	Attention scarcity

3.1 Hierarchies vs. Networks

Let us examine some of these concepts in greater detail. Hierarchical structures are very good at getting well-defined tasks done in a “divide and conquer” paradigm. But the key challenge of knowledge management today is not *dividing the work*, but rather *sharing the knowledge*. Yet for facilitating the free exchange of knowledge, a network structure (such as peer-to-peer) is inherently better equipped than a hierarchy. From an information sharing standpoint, a hierarchical, tree-like organization is a worst-case scenario because it is a collection of bottlenecks: There is only one “official” path between any two nodes in the graph and the likelihood of people sharing information can drop as a function of their distance in the corporate org chart if alternative venues (physical or virtual) are not provided to support knowledge sharing.

Because the bottlenecks can be overworked, on vacation, or simply disinclined to pass on information and instead decide to keep it for themselves as a strategic advantage.

This does not mean that we would advise to attempt to run a company in a peer-to-peer fashion, but for the special tasks of sharing and improving the collective knowledge, a peer-to-peer network poses far fewer problems in terms of bottlenecks and re-inventing the wheel.

3.2 Experts vs. Peers

In a world of readily available information, the traditional distinction between “expert” and “laymen” blurs and in many respects, we become “peers” with people even on “their” home turf. For example, with the amount of medical information available online, many sufferers of chronic diseases have joined together in virtual communities to exchange information on their common ailment. Medical doctors have been faced with increasingly well-informed patients over the last few years – which is not surprising: If you have the time and obvious motivation to dissect and digest all available new information on your particular disease with a community of like sufferers, spending hundreds of hours educating yourself on all aspects – who is going to be the expert on your disease when you go to your general practitioner? Is it sensible to label the MD “expert” and you “layman” with respect to this niche of knowledge?

This is something we are seeing in many areas. Another interesting clash between “expert” and “layman” is the famous public debate between Tanenbaum and Torvalds over whether Linux had any future held on Usenet in 1992.

Professor Tannenbaum, the prominent and recognized expert on operating systems, predicted that Linux would never be portable, i.e. run on another processor than the x386. Linus Torvalds, the computer science student, defended his idea against the expert. Today, Linux is known as the most portable operating system of all times – running on everything from embedded systems, PDAs, PCs, up to supercomputers and mainframes [3].

In this case again, the traditional distinction between expert and layman is not helpful to decide who had the more worthy input. Note that in the traditional, pre-Internet world, the layman would never have had the chance to publicly debate with the expert in a way that gives both an equal footing in the debate.

3.3 Teams vs. Communities

Some people have taken to calling teams “communities” because this term has become a focus of attention in the knowledge management space, especially in the context of *Communities of Practice*. But the two terms are not interchangeable. A team is simply a group of people who have been assigned a certain task together. A team usually has a formal, appointed leader who can wield authority over the others. A team is given goals by which it is measured. It is clear at any time who is a member of the team and who is not – the boundaries are clear. And, most importantly, team members are selected by management.

A community, as in a Community of Practice, however, is a different story. Communities usually do not have a formal leader with authority over the community members. That does not mean that all members are equal – far from it: A community usually has very different roles. A coordinator may exist, and leaders who are valued as advisors by their colleagues will usually either emerge, or may have been the community’s attractor in the first place. But any authority they wield is usually not formally bestowed by management. This is because of another important attribute: Members of communities are *volunteers*. They *choose* to interact, to share information on a common set of problems or interests. They decide their level of activity or inaction. Members are not primarily measured by their activities in the community. And communities do not have clear boundaries. There will be different roles with different levels of involvement, such as community elders, “newbies”, “lurkers” (those who read but rarely contribute), FAQ maintainers, etc. Also, the level of members’ activity changes over time to accommodate their knowledge sharing needs.

3.4 Information vs. attention scarcity

Information used to be expensive and difficult to access. Those who had access to privileged information, through professional associations, a corporate library or at university, were at a clear advantage.

But today, using internet technology, being denied access to information is no longer a problem in most areas. You can get the most current, cutting edge research articles in physics from pre-print servers on the internet half a year

before they are ever printed and put into a library. Today the problem is the huge information overflow.

Thus the scarce resource is no longer the available information, but the available limited attention we can spend on the information glut. This has been likened to trying to drink from a fire-hose. Therefore, the focus of knowledge management is no longer on storing and indexing the few “holy scriptures” for a given subject, but rather on helping us to sift through the millions of pieces of information and putting them together in a way that makes sense. It is less the creation of new knowledge from scratch, but rather the combination and smart re-use of existing information. This is also why Communities of Practice play such an important role: Every member serves as eyes and ears for the thematic community and helps filter out the relevant pieces from the information flood. These pieces can then be assembled by the community of peers to gain and document new knowledge. The Usenet FAQs are a good example of this. They direct our limited attention and allow us to optimize our time.

3.5 Key benefits

The key benefits associated with an open, peer-to-peer knowledge management approach are

- much higher speed of content creation, capture and development
- broader collaboration and use of skills
- getting know-how out of people’s minds and desktop computers and into the open
- increased transparency and accountability
- more efficient re-use of existing assets (less “re-inventing the wheel”, more “standing on the shoulders of giants”) and
- dramatically reduced bottlenecks for the flow of information.

4. Overcoming obstacles to P2P KM

But reaping the benefits of the peer-to-peer paradigm on a larger scale inside a company requires overcoming a number of concerns and obstacles.

4.1 Information Quality

Electronic business relies on quick, easy and cheap exchange of information. It is an often raised argument that by empowering regular employees to be not just a *consumer*, but also a *producer* of information, the quality of information is diluted. Only the experts, the argument goes, should be allowed to create, review and update information on a corporate intranet, in order to assure high quality. But quite on the contrary, it has been demonstrated numerous times that *massive peer review* and participation in reviewing and editing information can actually lead to higher information quality than traditional approaches based on the review of only a few designated experts. This has not only been shown in the area of Open Source development and documentation, but also regarding traditional journal publication (i.e.

PublicLibraryOfScience.org [6]).

Indeed, there has been a recent case of large-scale scientific fraud, where a prominent researcher was able to publish 16 different papers in traditional, respected science journals, without the fraud ever being noticed during the 16 review processes. Only a tip-off by a peer brought the house of lies tumbling down [15].

4.2 Confidentiality

Confidentiality and intellectual property issues are part and parcel of corporate life. I.e. a consulting client may require strict non-disclosure agreements to protect his assets. Of course, this makes a peer-to-peer approach to knowledge management more challenging. However, corporations have dealt with confidentiality issues for a long time and have developed proven methods for coping with them, e.g. judicial, administrative and technical solutions. The real task at hand is questioning the basic notion that knowledge created in one team or department should be kept secret from another at all, as this creates corporate fiefdoms, silo thinking and leads to re-inventing of wheels.

Of course, in day-to-day reality, “keeping information secret” is not usually the result of a conspiracy or stubbornly denying access, but a brew of lack of transparency, lack of (default) access rights and non-publication of pointers to the information. You can *theoretically* track down the information if you know it exists – but the chances of finding it serendipitously are very slim.

4.3 Transparency

It appears that many of the concerns and fears regarding quality and confidentiality are in fact a reaction to the greater *transparency* associated with an open, peer-to-peer approach to knowledge management. This may be disconcerting to some people. With an open, peer-to-peer approach to knowledge management, people outside the designated “expert” team get to review, comment on and improve the available information. To quote Linus Torvalds [16]:

“In that regard, open source – or open anything, for that matter – is unforgiving. It shows who can get the job done, who is better. You can’t hide behind managers.”

This new degree of accountability can understandably appear threatening to some who are relying on traditional, non-transparent approaches to sustain their expert status.

But it is precisely this transparency that creates an incentive to share rather than hoard information. The decision for knowledge workers to share or not share their information in a situation where their decision is not transparent to a large number of their colleagues, can be modeled as a simple prisoner’s dilemma. Defecting (using information provided by colleagues but not providing any information yourself) is the rational strategy, as it maximizes your gains. The adverse effect to your reputation is negligible, if your behavior is only apparent to one colleague at a time. This is typically the case when

there is no “shared space” with easy write-access for exchanging information, but rather an official knowledge database that is filled with information through a standard process and used mostly in “read-only” fashion by the knowledge workers.

The picture changes dramatically when we introduce true shared spaces, such as newsgroups, discussion boards, team-rooms or Wikis. Non-compliance with a request for sharing information will result in significant damage to one's reputation, if both the request for help and the non-cooperation can be witnessed by many colleagues. This will lead them to one of two conclusions: A) the self-proclaimed “expert” refuses to cooperate or B) the self-proclaimed “expert” has no valuable knowledge to share.

The first conclusion lowers the expert's social standing, the second lowers his professional standing with his peers. In such a “shared space” setting, it is likely that the role of experts as valued advisors will quickly pass from those who hold an “official” title to those who are actually willing and able to help you in your day-to-day work. This creates an internal marketplace for information, where an expert reputation has to be earned, rather than bestowed.

Our prisoner's dilemma is also transformed: From the simple prisoner's dilemma to the iterated prisoner's dilemma, which is played with the same partners over extended periods of time and where the history of previous actions of all partners are visible and available to all players. In this setting, cooperation becomes the rational choice, as the reputation based on previous defection or cooperation behavior becomes a significant asset. Nobody is going to play cooperation with a known defector, whereas the known cooperators can optimize their mutual gain in the long run.

4.4 Organizational and Cultural Change

Therefore, the concerns of the employees have to be mitigated by a culture that encourages experimentation and failure over inertness and perfection. Only in such a “forgiving” environment will employees feel confident to experiment with and vent new and “half-baked” ideas, which can then be developed and brought to perfection in a collaborative approach. A culture will need to be established in which the incentive for sharing information is greater than the benefit of hiding it.

This is possibly the hardest part of implementing a peer-to-peer approach. An organization is a very complex system of human actors, with their different needs, motivations and expectations set within org charts, processes and both written and unwritten laws governing behavior. Any change is likely to be rejected unless introduced, motivated and managed properly. And the transparency awarded by an open approach to knowledge management is very different indeed from the way most organizations organize knowledge management today – as a form of electronic library, focusing on storage and retrieval rather than creation and development.

5. Conclusion

The application of peer-to-peer concepts to knowledge management has the potential to improve the quality, currency and comprehensiveness of a company's readily available internal knowledge by an order of magnitude.

With Internet technology at our disposal, we have the necessary tools in various different flavors. But the key challenge for the success of a peer-to-peer knowledge management is the creation of a cultural foundation which encourages and rewards transparency over information hiding. It requires the sponsorship and political will to “rock the boat” and to threaten the positions of information gatekeepers where they are more of a burden than a boon. Only an environment where every practitioner of information can also easily internally publish without going to a gatekeeper, only when the child (intern) can publicly call the emperor (expert) naked is the necessary feedback in place to keep improving the internal knowledge at a pace that keeps up with the needs of knowledge workers.

Tools can make this transition easier, but the key component of modern knowledge management is the support of the practitioners themselves. You cannot force a knowledge worker to share his knowledge. Today, we see the focus of Knowledge Management moving beyond IT into organizational and cultural change issues. The solution of which, while very challenging, also hold yield the rich rewards for an organization in terms of speed, adaptability and competitiveness.

References

- [1] Brooks, F. P. *The Mythical Man-Month*, Addison Wesley, 1975.
- [2] Coase, R.H. *The Firm the Market and the Law*, University of Chicago Press, 1988.
- [3] DiBona, C., Ockman, S., & Stone, M. (Eds.). *Open Sources: Voices from the Open Source Revolution*. O'Reilly, 1999.
- [4] Feynman, R. P. *The Pleasure of Finding Things Out*, Perseus Publishing, 1999.
- [5] Harvard OpenLaw Project. <http://eon.law.harvard.edu/openlaw/>
- [6] Karow. “Publish Free or Perish.” *Scientific American*, April 2001. <http://www.scientificamerican.com/explorations/2001/042301publish>
- [7] Ives, B. & Jarvenpaa, S. L. “Will the Internet Revolutionize Business Education and Research?” in: Smith, D. E. (Ed.) *Knowledge, Groupware and the Internet*. Butterworth-Heinemann, 2000.
- [8] Leuf, B. & Cunningham, W. *The Wiki Way: Quick Collaboration on the Web*. Addison-Wesley, 2001
- [9] MIT OpenCourseWare Project. <http://web.mit.edu/ocw/>
- [10] Moody, G. *Rebel Code: The Inside Story of Linux and the Open Source Revolution*, Perseus Publishing, 2001
- [11] Neus, A. “Managing Information Quality in Virtual Communities of Practice,” In: Pierce, E. & Katz-Haas, R. (Eds): *Proceedings of the 6th International Conference on Information Quality at MIT*.

- [12] Oram, A. (Ed.) *Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology*, O'Reilly & Associates, 2001.
- [13] Public Library of Science Project. <http://www.PublicLibraryOfScience.org>
- [14] Raymond, E. S. *The Cathedral and the Bazaar*. O'Reilly & Associates, 2001. <http://www.tuxedo.org/~esr/writings/cathedral-bazaar/>
- [15] Stieler, W. "Tiefer Fall: Skandal um Datenfälschung erschüttert Bell Labs", *C'T Magazin für Computer und Technik*, 21/2002, 66-67.
- [16] Torvalds, L. & Diamond, D. *Just For Fun. The Story of an Accidental Revolutionary*, Harper Business, 2001.
- [17] Wikipedia: The Free Encyclopedia. <http://www.wikipedia.com>
- [18] WikiWikiWeb. <http://c2.com/cgi/wiki?WikiWikiWeb>