A PRESCRIPTION FOR KNOWLEDGE MANAGEMENT

WHAT HOFFMANN-LAROCHE'S CASE CAN TEACH OTHERS

Innovation in Action

About Patricia Seemann:

This article reflects the thinking and work of Dr. Patricia Seemann, who was formerly Director of Knowledge Systems for Hoffmann-LaRoche. Since the events described here, Seemann has left that firm to accept a position with Ernst & Young's European management consulting practice.

Seemann is a medical doctor by training, but has spent her career in the pharmaceutical industry. Prior to joining Hoffmann-LaRoche, she restructured and led two departments for a large biotechnology firm. Contact her at 106424.2323@compuserve.com

Managers thinking about embarking on knowledge management projects might benefit from the experience of Hoffmann-LaRoche, the international pharmaceutical firm. Management there credits a single knowledge initiative with making a significant difference in the profitability of new products. Yet the initiative was a fairly modest one—it called for no huge new information systems, no army of information processors. In fact, the lessons it teaches are all about focus—on the right problem, at the right level, and on the right goals for the business.

Targeting the Right Problem

For Hoffmann-LaRoche, as for every pharmaceutical company (and many other types of companies) much depends on the speed of new product launches. Industry observers estimate that development of a new drug takes, on average, five to eight years and costs over $250 million. Firms that can expedite that process stand to gain tremendously. First, they recoup their development costs faster—and generate higher profits. At the same time, shorter development time means more new product ideas can be placed in the pipeline—helping to hedge the risk of any of them fizzling (as many do). In an industry like pharmaceuticals, where a firm's market standing is only as good as its current patents, fast and sure drug development is the key to survival.

Accelerating New Product Development, pg. 55

"At the outset, we wondered 'where should we start?' And my view was that we needed to go for a part of the business that was truly relevant to its strategy. Which meant that it would get senior management attention, and that we would have tangible economic and business results. If we had said, 'let's start with something in human resources'—like how we deploy ex-patriots globally, for instance—that would not have been the right move."

The competitive intensity is similar in other industries, but for pharmaceutical firms there is even an additional complication. At a computer company, for example, the *Eureka!* may have barely passed the scientists' lips before marketers are making their first sales. No such luck in Pharmaceuticals. Great science only gets you to the next step in the process: the FDA. Without approval from the US Food & Drug Administration (and its counterparts around the world) the product goes nowhere—and that's a step that can take months or even years.

When Patricia Seemann took on the role of Roche's Director of Knowledge Systems in 1993, the firm's track record on new drug applications (NDAs) was mixed. On many occasions, new products sailed through the approval process and enjoyed a prolonged marketplace advantage. Other times, though, NDAs got hung up by requests for more information or additional trials—or were approved for more limited usage than hoped. Dr. Seemann's responsibility in general was to determine where better knowledge management might make a real difference to Roche. Looking across the company's operations, her eye fell naturally on this spot.

Goal-Setting: A Balance of Ideals and Conservatism

The opportunity seemed almost too good to be true. Here was an area where the company had proven its abilities again and again. The problem was not its capability but only its consistency. Clearly, this was a case of harnessing the knowledge the company already had and applying it more uniformly. It was also an area that promised huge payback. For every day gained in the market availability of a new drug, Roche had determined, the company stood to gain a million dollars. If only its occasional height of achievement could be made the norm, Seemann considered, Roche could leapfrog its competitors handily.

The potential benefit of accelerating NDAs was further compounded by the fact that there were over 30 new drug projects in progress at the time—a typical number for the firm. If better knowledge management could pare down average approval time by only a fraction, the return on investment would be immense. Scanning the opportunity, it seemed resonable to Seemann to expect at least a three-month reduction in major new drug approvals. At $1 million per day, that would represent some $90 million in additional profitability—*per product.* At the meeting where her team presented their plan, says Seemann, "we had a reasonably captive audience."

Creating Fertile Ground . . . , pg. 34

"A question we asked ourselves at Hoffmann-LaRoche was: when we develop a product, what are we really developing? We're not developing the tablet. We're developing the *knowledge* around a disease and its treatment. That is what we're putting out on the market. All of a sudden, knowledge is no longer just a tool to do your business. It *becomes* your business."

article abstract

Defining the Knowledge Challenge

As Seemann's team looked further into the new drug approval process, it became more and more clear that this was truly a challenge of knowledge management. First, the product involved—the application document—was purely and simply a knowledge product. And second, compiling it drew on the work and insight of literally hundreds of knowledge workers. These two basic observations led Seemann to design a knowledge management initiative with two thrusts: 1) to help product teams prototype the knowledge required for their new drug applications; and 2) to produce a comprehensive "map" of the knowledge sources in the company that might contribute to their completion.

Prototyping Knowledge

The first—and possibly most important—step at Roche was to perceive that its primary product is knowledge. This is particularly true in the new drug approval process, where the "customer" is the regulator. Regulators, after all, do not need or want the drug made by a company, but only the information about it. This shift in perspective, considering the regulator as customer and the new drug application as product, was a crucial one. It became evident immediately that the NDA was a product that, unlike Roche's drug products, was not manufactured for efficient or effective performance. To prove the point, Roche could reflect on one particularly ineffective NDA. In that document, Roche's collaborators had failed to emphasize evidence that the drug was effective even when taken just once a day. As a result, the FDA approved it at a twice-a-day dosage, negating what would have been a significant advantage over competitors' products.

In general, the problem with NDAs was that they were treated as data dumping grounds. Often amounting to tens of thousands of pages, they buried the most important information (like possible side effects) with the most available (like how often the lab animals' cages were cleaned). Picture a cold remedy with 10 milligrams of active ingredient delivered in a tablet the size of a suitcase. That was how Roche was delivering its most important knowledge.

Prototyping knowledge requirements meant understanding better the real needs of their customers. To get that insight, the Roche team assembled a group of ex-regulators who had real insight into the approval process. They helped Seemann's team by observing that regulators are really only trying to answer three crucial questions: Is the drug safe? Does it work? And is it of sufficient quality? The myriad data requirements specified by regulators' guidelines were simply extensions of those questions. The team also worked to compile the various guidelines issued by the separate regulating bodies of Roche's top 20 markets worldwide.

"The goal that we set for ourselves at Roche was the following: We wanted to make sure that people would be able to *consistently* access and *effectively* access the organization's knowledge. We included three areas that people really had to have access to. One was customer knowledge. The second was experience knowledge. And the third was process knowledge."

article abstract

The prototypes that emerged from this work represented the total accumulation of knowledge required for each NDA. For one thing, they covered in one place the requirements to get a drug approved around the globe—something that previously would have been considered a miracle. Even more importantly, they focused directly on the three questions that had to be answered—and on communicating Roche's key messages relating to each. Yes, these prototypes still responded to the minutiae of the regulators' requirements—but they did so by way of presenting a logical and compelling response to what those regulators were really asking. In other words, these were knowledge products potent and beneficial enough to bear the Roche name.

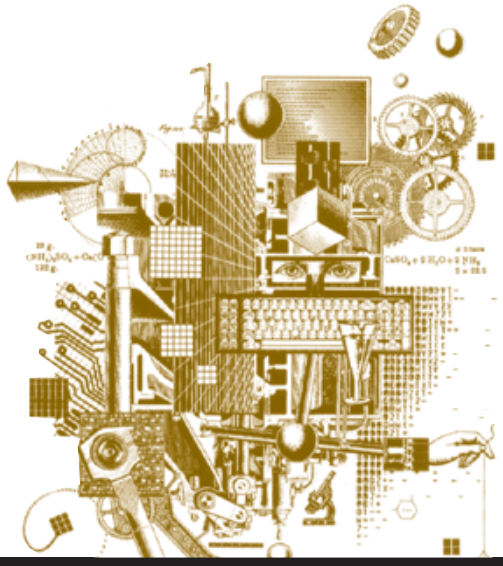

Mapping the Knowledge of the Organization

If the prototypes represented where the product teams needed to be with their knowledge, Seemann still needed to provide a map of how to get there. Prototyping knowledge raised as many questions as it answered: Where would all the supporting information come from? When would it be available? Who should contribute at what point? To clarify the sources and flows of knowledge related to Roche's new drug development, Seemann's team constructed a knowledge map.

In its simplest form, a knowledge map is a straightforward directory pointing people who need access to knowledge to the places where it can be found. Usually, such maps recognize both explicit and tacit knowledge—that is, knowledge that has been captured in documents and databases, and knowledge that resides only in the heads of experts. Seemann's map certainly included this "Yellow Pages" functionality, but it went further. She incorporated features that would tie the map more directly to the completion of an NDA, and make it a better tool for accelerating the launch of new drugs.

Creating Fertile Ground . . . , pg. 34

"When people don't share their knowledge, in the majority of cases, it's not because they're malevolent. It's not because they're not being rewarded for it. And it's not because the structure of the organization somehow gets in the way of it. It's because they don't know *why* and *when* and *how* they should share it."

article abstract

The first of these features was a Question Tree. Any map needs an organizing framework, an architecture of the content it will cover. Seemann's map took as its starting point the rewritten guidelines developed with help from former regulators. Again, these focused on the three basic questions Roche had to answer about every new drug: Is it safe? Does it work? And what is its quality? Branching out from those questions, the map showed the more detailed sub-questions that followed. Under "is it safe?", for example, were things like: "at what dosage level is it safe?", "does it have side effects for the patient?", and "is there any effect on the patient's offspring?" For each question, the map pointed to the source of its answer. It simultaneously detailed what knowledge would have to be gained for the typical new drug approval and where to go to find it.

The second key feature Seemann added to the Roche map was a specification of "Knowledge Links." It wasn't enough, she believed, to create a purely passive map. That would serve people who needed knowledge, but not people who had knowledge to offer. Knowledge links are road signs that show with whom and at what point a person or group should share their knowledge. For example, the team conducting animal testing was not in the habit of immediately sharing what was learned about dose-response with the person who would eventually direct the clinical (human) tests. Their work was too far separated by time, geography, and hierarchy for this to happen naturally. But having those findings when they became available would help clinical testers design more efficient studies more quickly, and perhaps collapse the drug approval time frame significantly. With the addition of Knowledge Links, the map became a more active driver of knowledge-sharing. To specify the links, Seemann's team simply had to get people from various parts of the organization to sit down together and talk about what they would like to know and when.

Finally, the knowledge team added a great deal of content to the map to show how specific questions in it had been answered successfully in the past, or how they might be better answered in the future. With all these valuable features, the map had become a powerful—and sometimes controversial—thing. Part directory, part process description, part best-practice repository, it was precisely the knowledge tool the organization needed.

Condition Cured—Plus Some Beneficial Side Effects

Roche's first indication that its knowledge management effort was succeeding came early on, when it asked the ex-regulators to react to its knowledge map. The group responded with high praise. Had they been equipped with the insight it promised, they said, they could readily have approved a new drug.

Could it be true? More to the point, could Roche's knowledge map really deliver more effective NDAs? It wasn't long till Roche had a real test, and this time the feedback was definitive. In an application for new indications for a drug, Roche expected its filing time to consume 18 months. With the new knowledge tools in place, it took just 90 days. Approval from the FDA, projected at three years, came within nine months.

Perhaps even more important in the long run are some of the less measurable—and in some cases, unexpected—benefits of the initiative. Seemann reports, for example, a surprising outpouring of gratitude from people using the knowledge map (it was rolled out to over 3,000 employees). For the first time, many of them had a clear picture of the multi-year work underway, and where their current work fit into it. Many, too, were gratified to be included in the map's "yellow pages" as important sources of knowledge. This recognition had nothing to do with hierarchy and everything to do with value to the organization. These two benefits—greater transparency of Roche's workings and a motivating recognition of knowledge sharing—will yield much greater returns than any single new drug approval.

Seemann also believes that making Roche's knowledge and knowledge processes more visible has helped to break down some vexing interdepartmental barriers. Every type of organization has its factions and its walls. In pharmaceutical companies, there is a traditional rift between the MDs and the PhDs. When every group has a greater perspective on how their work dovetails, greater tolerance and cooperation results. This works across geographic/cultural barriers, as well. After the knowledge map was issued, Roche's animal testing department, which had wrestled with globalization for years, was able to find a globalizing structure within a few months. It seemed to help, too, in achieving a rapid assimilation of Syntex after that firm was acquired by Roche.

Whether due more to its tangible or its intangible benefits, the knowledge management initiative was embraced by Roche. In fact, it is generally agreed that it is the first management initiative to become embedded in the organization.

"In very large and complex organizations and projects, it is very difficult for people to understand 'where are we in all this?' When the knowledge map was implemented, what people were most grateful for was the possibility of truly understanding what they were doing—where they were in the company, and how it all fit together."

Lessons to Take Away

Even this quick telling of the Hoffmann-LaRoche story makes some important lessons clear. First, managers hoping to make a difference through better knowledge management should start by focusing on the right problem. Patricia Seemann chose a spot that was closely tied to the strategy of the business, and a driver of the firm's future growth. She also focused on a process that was undeniably knowledge-intensive, ensuring that the impact of knowledge improvements would be great. This points to a second lesson: set definitive goals for what the effort will achieve. Preferably, as at Roche, these can be stated in terms of ultimate increases in profitability.

The third lesson to take away from Roche's story is that knowledge management need not be technology-intensive, and should not be technology-driven. Tools like prototypes and knowledge maps can be surprisingly low-tech. They don't require people to buy into major infrastructural overhauls up front and on faith—they simply get a job done, and win converts along the way.

Finally, Roche's success teaches a lesson about bringing together the right project team. A mix of twenty-five Roche people and a variety of outside consultants, Seemann's was small enough to move fast, but big enough to bring a variety of perspectives to the table. Most importantly, every member of the team was drawn from the best and the brightest Roche had to offer. Too often, Seemann knew, internal projects are staffed with employees who have time on their hands. Unfortunately, they may be free for good reason—they are not the organization's most valued contributors. Getting the best benefits a project on two levels: it gets the work done faster and better, and it makes a very visible statement about the project's importance to top management. In Seemann's words, "Do not divest knowledge management to your deadwood. Knowledge is something that is so dear to the company that only the best and brightest can actually bring it out."

Epilogue: Why Knowledge Management Matters

Better knowledge management certainly made a difference for Hoffmann-LaRoche. It also made a difference for the rest of us. Recall that Roche is a company dedicated to developing treatments for people with life-threatening diseases. When the time came to test its new knowledge tools, one major drug in question was Invirase™ , and the disease it targeted was AIDS.

Knowing the Drill: Virtual Teamwork at BP, pg. 14

"The creation of value is coming increasingly from the collaboration of groups—the combination of their experience and skills. So, that combining and recombining is becoming the true challenge for companies. The point is no longer to manage the silos, but to bring together around a problem the right group of people with the right knowledge."

article abstract

Invirase™ is the first compound in a new class of anti-retroviral agents called proteinase inhibitors. They work by interfering with the activity of proteinase, an enzyme that is critical to the replication of HIV, the virus that causes AIDS. Hoffmann-LaRoche submitted the NDA for Invirase on August 31, 1995. It was reviewed by the Antiviral Drugs Advisory Committee of the US FDA on November 7, 1995.

Today, the drug is being used widely by HIV-positive patients in ongoing studies of its safety and efficacy. Thousands are receiving the drug free of charge from Roche through its "Compassionate Treatment" program. Most are experiencing fewer negative side effects than with former treatments. And the treatment is the most effective to date in the fight to eradicate this terrible disease.