

Knowledge Work and Unpredictable Processes

Keith D. Swenson, Fujitsu America, USA, and Vice Chair, Workflow Management Coalition

INTRODUCTION

What is the next thing beyond Business Process Management (BPM)? To many this is an unexpected question. Is there anything wrong with BPM? Any reason it seems to be flagging? Over the course of 2009 there were a number of high profile acquisitions of BPM companies, and many say this as an indication of the end of BPM. Others, however, see this as an indication that BPM is mature, solid, and relatively well-defined, and a natural occurrence of a maturing technology area. Either way, it prompts people to wonder what is going to be next.

According to Forrester, we saw in 2009 a surge of interest in case management, so much so, that in December Craig Le Clair and Connie Moore released a white paper on Dynamic Case Management (Moore 2009). A number of companies, including IBM, announced in April 2010 a new offering in Case Management above and beyond their extensive BPM presence.

Ironically, it is precisely the success of BPM that leaves us looking for what is next. BPM has slowly but profoundly affected our business culture to the point that BPM is being used in almost all organizations to support routine work processes. This is not yet the case with *all* routine processes—there are still many more to be implemented—but BPM has freed a lot of people from routine tasks, allowing them to focus more on the harder-to-automate knowledge work tasks. Our working population is about 40 percent knowledge workers, and this percentage is growing steadily.

Can we use BPM to support knowledge work? The answer is “no.” I will give a number of reasons for this below. This is a highly charged question, which I expect to be subject to quite a bit of debate.

This answer is not lightly given.

Before explaining the answer fully, we need to understand what knowledge work is, then the qualities of BPM. To support knowledge work, I will propose that a new kind of technology will be needed: Adaptive Case Management (ACM) (Swenson 2010).

WHAT IS KNOWLEDGE WORK?

Don't think of knowledge work as the manipulation of large amounts of information. The term “knowledge worker” tends to connote someone who is unusually knowledgeable or manages a large volume of information like a professor or a librarian. It is true that a professor does research, and a librarian categorizes books. These tasks are knowledge work. However other tasks of managing a library, such as the routine storing, filing, and managing of books should not be considered knowledge work because they are routine tasks. While professors and librarians can be considered knowledge

workers, they are really not very typical examples. Most knowledge workers are not so explicitly tied to large volumes of information.

A knowledge worker is anyone who draws upon tacit knowledge, evaluating input from many very widely disbursed and rapidly fluctuating sources. Knowledge workers must figure out what to do next as they are working. An excellent example of a knowledge worker is a search and rescue worker. Often, the information on the rescue is not very well known when they start out, and they have to gather it as they proceed. Based on information they gather, they need to make decisions about what is to happen next.

It is probably not surprising that doctors and nurses practice knowledge work when they determine the course of treatment for a patient, always gathering more information, and always tweaking the treatment in light of what they have discovered.

Lawyers and judges are knowledge workers not only because they bring a lot of knowledge to their work, but that the case unfolds in response to what new things are uncovered, and that there are many unpredictable decisions to be made along the way. Designers are knowledge workers. Executives are knowledge workers. Project managers are knowledge workers. Customer support is mainly knowledge work. Product research and development is knowledge work. Knowledge work requires one to think in order to do the job. It is that thinking that separates it from routine work, not just thinking about a result, but thinking about the way the work needs to proceed along the way.

While routine work can be predicted, knowledge work is unpredictable. Knowledge work does not look like a traditional business process where work is performed according to a detailed plan prepared in advance. Knowledge work involves what are called “emergent processes.” In an emergent process, the sequence of actions depends strongly upon the specifics of the situation. An initial plan is drawn up with whatever information is available at the time. It may incorporate elements of prepared procedures. The plan may draw upon roles that individuals have been trained to play. The specifics of the situation will drive a unique plan with a unique collection of individuals to meet the needs of a particular case.

The initial plan is not enough. As the situation changes, often the original plan has to be discarded and a new plan drawn up. This is the nature of knowledge work: as the situational information changes, so too must the plan. It is not simply a matter of assess, plan, then execute—it is an ongoing activity of continuous assessment, continuous planning, and continuous execution. No plan is ever final until the goal is achieved.

Like rescue workers, fire fighters “figure out” what they have to do as they go along. Forest fires are increasingly battled using IT resources (Gogek 2007). Satellite pictures are requested and then analyzed with increasing sophistication to produce an accurate outline of the area on fire. Sophisticated models are used to project the position of the front of the fire. Databases of toxic hazards are accessed and symbols are overlaid onto maps warning workers of dangers. This knowledge helps fire fighters deploy resources as effectively and safely as possible. Detailed instructions flow to the people in the field and detailed status information, including GPS coordinates, come back in near real time.

While rescue and fire workers are dramatic, many more knowledge workers spend their days in the office working for businesses. Consider how often business people describe their day as having been spent “solving problems.” Workers who leverage knowledge and make course-of-action decisions account for 25 percent to 50 percent of the workforce (Davenport 2005). A stock trader goes to work without a predefined script for the day, instead, news from many sources will guide the actions of that stock trader on that day. A medical doctor does not usually have a complete script for the treatment of all patients, but instead, must follow the progression and make decisions about patient care at many points along the way. A social worker does not have a prescribed way to handle all cases but must make decisions along the way to decide what to do next. If you look past the special skills required for a particular line of work, many jobs consist of being prepared for and handling emergencies.

QUALITIES OF KNOWLEDGE WORK

Non-repeated

Knowledge work is rarely, if ever, repeated the same way many times in a row. In some cases, it is not possible to do the same work in the same manner twice, such as the negotiation of a merger of two companies. In other cases, it is simply unnecessary, such as the discovery of a new product idea.

There are always elements of similarity. For example, all news reporters might submit articles by a particular time of the day to make it into that day’s newspaper. But the journalists’ work over the entire life of the article—what is done, who is contacted, and where information is gathered—will have elements that are unique for any given article. A merger of two companies may follow a very familiar pattern: initial contact, hypothetical probing conversations, a getting-to-know-each-other phase, a proposal, a due diligence phase, agreement, and follow-through. This high-level pattern does not say anything specific about what a given individual would be doing at a particular time. A stock trader goes to work every day and buys and sells stocks, but that does not mean that the same work is repeated every day; quite the contrary, the actual trading patterns are adjusted over the course of the day in response to many factors.

Routine work cases may be very similar, but the cases are not entirely identical. A bank may open fifty accounts today with each account for a different person with different personal data. One might argue that every new account opening is unique, since it is for a different customer. These differences, such as the name and background of the customer, do not materially affect the course of actions that the bank must go through. There may be variants of the process, such as loans for high-income applicants might follow a different path than loans for low-income applicants. A loan application process is often visualized as having branches for treating cases differently. But even with the branches, the process can still be thought of as highly repeatable.

The key is the degree of repeatability: routine work will be repeatable enough that it might benefit from a formal description of the process to take. Whereas, knowledge work is not repeatable enough, and a formal description of the process tends to cost more than it would benefit.

Unpredictable

Routine work is predictable; knowledge work is unpredictable. When talking about unpredictable, we need to be clear about the level of unpredictability. A bank branch manager may not know who is going to walk into the branch in a given day to sign up for a new account, but the process of signing a given person up is predictable and repeatable. A checkout clerk will not know ahead of time what items will be scanned, but the task of scanning the barcodes of all of the goods being bought and totaling the amount is very predictable.

We need not be concerned about unpredictability at the micro scale. A checkout clerk will not know exactly which items, or even how many items will be scanned in a particular sale, but it is unimportant to know that level of detail in advance.

When we talk about the course of events being unpredictable, we mean that the sequence of significant human acts is not knowable in advance, and the course may vary greatly from case to case. The course will depend greatly on the details of the situation itself, and the details themselves may change before the work is finished. We will see examples below of a hospital which will start with the relatively routine task of admitting a patient, but what happens next is entirely unpredictable. The care may involve a single procedure taking one hour, or it may involve hundreds of procedures taking many months.

Prediction of the course of events must be based on actual information that the organization has at the time. Omniscience might make knowledge work predictable, but we cannot claim omniscience. For example, care must be provided for a patient that arrives unconscious, without identification, and without retrievable medical history.

Emergent

Stories told about knowledge work have the quality of “unfolding” as they go along. As a knowledge-work scenario unfolds, an early step may yield some knowledge, and that discovered knowledge determines the next step to be taken. The second step yields more knowledge, which in turn determines the third step to be taken, and so on. This iterative unfolding aspect is what makes knowledge work so unpredictable.

If I were to ask you to run to the store and bring back a box of mints, I could probably generate a plan of how to accomplish this to an annoying level of detail. I could predict fairly accurately the route to walk and the number of paces to the car, the route to drive, the approximate place to park, the number of paces to the store, the aisle to visit, the location of the product, the approximate location of the cash register, the amount to pay, and similar details for the return trip. This is because the entire task can be known as a whole and, given sufficiently detailed information, the entire plan can be laid out. In execution, problems may arise—such as needing to detour because a particular road is blocked or having to park farther away because the parking lot is full—but these unpredictable changes cause very little structural change in the overall plan. They remain minor perturbations from the original plan.

Investigation of a crime, which is knowledge work, has a completely different nature. One clue leads to the questioning of one individual. This in turn yields a clue that leads to another investigation and more clues. This may

lead to a laboratory test that yields more clues. The details of the case exert a strong influence on the course of events for that case. A good detective will apply heuristics persistently in the hope that, eventually, enough clues will be discovered to solve the case, but at no point can the detective confidently map more than a few steps into the future. Often, there is a breakthrough that allows the next few steps to become clear—this is also known as an “ah ha!” moment. This sudden realization may change the current course of subsequent action. In knowledge work, even the concept of a plan has a special meaning, since plans are always tentative, contingent upon the next piece of knowledge uncovered. The hallmark of knowledge work is the mention of “Plan B,” as this implies the very unpredictability of the situation.

This sort of work is not exclusively the domain of Sherlock Holmes, but it appears in many lines of work: for example, working together with a customer to find a new product direction, coming up with a new advertisement campaign, increasing the quality of a production line, and finding a source of funding for a public works project.

Robustness in the face of Varying Conditions

One of the most non-intuitive aspects of knowledge work is reliability and robustness in the face of variable conditions. Organizations that need a very high level of reliability make use of knowledge work techniques to ensure that.

The U.S. Marines train their soldiers to know that decisions are made “on the front line.” Marines are not sent into the field with immutable instructions detailed to the finest level, but instead, they have the flexibility to adapt the plan as the situation evolves. This is not meant to imply that they act “willy-nilly” in any regard, nor do they modify the goals of the mission. The original orders represent the strict goals to be achieved, and there exist well-defined regulations on how they can be carried out. But the generals know that they don’t have perfect knowledge of the situation the soldiers will face, and they must count on the behavior of intelligent adaptation as necessary to the specific situation as it unfolds. From the general’s position, the exact details are not entirely predictable, but the goals can be achieved reliably.

This is counterintuitive because we normally think of machines as precise and reliable, while humans are error prone. Some might think that jobs left to the knowledge of the worker would be erratically implemented and error ridden. That is possible, but not usually the case.

Consider instead the reliability of an extremely rigid automated process. By defining the exact process too thoroughly, you can end up with a fragile process that breaks upon encountering the first exceptional situation. Realize that in any formalized work process, there is an assumed amount of knowledge work going on around it, either to handle the exceptional cases or to analyze and improve the process plan itself. Knowledge work exists outside of and all around the more routine practices making sure that cases that fall out of the routine course are picked back up and put back on track.

There are also cases where very well-defined, high-risk tasks are made considerably more reliable by layering knowledge work on top of the routine work. The U.S. Navy nuclear program has an unblemished record because of a practice of engaging every participant in the constant practice of looking for unexpected situations and continually suggesting improvements to the process (Spear 2009). Similarly, Toyota practices the same strategy with

their Lean TPS (Toyota production system) method of developing cars. The knowledge work of suggesting improvements is unpredictable, but it serves to make the routine work processes more reliable.

WHAT ARE THE REQUIREMENTS TO SUPPORT KNOWLEDGE WORK?

Given the qualities of knowledge work, we can then propose that Adaptive Case Management, the category of tools that support knowledge workers, must have the following capabilities (Swenson 2010):

1. Explicit representation of goals and sub goals. While the process is not predictable, the goal is well known, and can be used to drive the work to a successful completion.
2. Rich data model. The various entities are represented directly and often remain behind as a permanent record.
3. Ability to task people directly to do a thing, without needing it to be part of a pre-defined process.
4. Ability to support a process (a sequence of tasks) when custom or law requires a specific sequence of tasks to be done.
5. Collaboration support that allows for threaded discussions among the participants involved in a case, without that communication being tied to a particular action. Some call this “social network” capability.
6. Rich access control model. Many BPM systems control access mainly by controlling who is assigned at a particular time.
7. History of all actions is kept, even meta-actions such as giving access to someone, or revoking access permissions.
8. Ability to start with a blank slate and build everything as you go without any preparation needed. There is no distinction between design time and run time. There is no development lifecycle where processes are developed and then deployed.
9. Case templates which are collections of elements that the case manager can use at run time to compose the case itself.

The above capabilities are tuned to the way that knowledge work presents itself: unpredictable, non-repeating, and yet still very important.

CAN KNOWLEDGE WORK BE SUPPORTED BY BPM?

When you have a hammer, all the problems start to look like nails. What is happening is that we are, at last, realizing that while there are many nails, there are also many non-nails that cannot be solved with a hammer.

Those who worry about the end of BPM are not looking at the entire picture, and are not seeing the real trend that is emerging. Understand that there have been many different meanings for BPM. Many have promoted it vigorously, claiming that it would solve every organizational woe. In many senses, BPM was the answer to everything, or so it would seem. All you need is BPM, and your organization will be transformed into a global front runner.

BPM (like many other high tech topics) has been overhyped, and proponents have said that it will be useful for things that we are now seeing. This just ain't so. There were many people promising that BPMN would be used directly by business people, but what we see is not only that business people shun such diagramming, but that tool vendors make their modeling tools more and more for trained specialists. Even the BPMN 2.0 standard committee is

pushing it to be more and more oriented toward a language you build systems with.

When a “hammer user” encounters a screw for the first time, it appears to be a nail with a very inconvenient edge wrapped around it. After pounding a few in, they conclude it is a very inferior sort of nail.

When we encounter something new, we evaluate it and judge it in terms of what we already know. It is very hard to escape the context bias. Knowledge work is really quite different from routine work, and requires a different approach.

Many people define work as that activity which is done for a purpose, and is repeatable, and produces a measurable result in a fixed amount of time. Work is what laborers do, not what executives do. It is very hard to observe exactly what an executive produces, and hard to define the process they use to achieve it. Many people consider that type of knowledge work non-work: for example, have you heard someone say “I didn’t get anything done today because I spent the whole time putting out fires.” Unless that person is actually a fireman, they probably meant that they spent their time solving problems, making decisions, and otherwise taking care of things that had to be done, but which they did not consider to be work. Part of the problem is that when a person says that BPM can handle all kinds of work, they are often excluding any consideration for handling knowledge work.

When a process automation expert looks around to find work that might be automated, the work of an executive is never considered, for obvious reasons: it is not predictable or repeatable. The job of a facilitator is never considered for BPM support. And so it is with all knowledge work. The question we face today, and increasingly over the next decade, will be how to support knowledge work.

Many BPM defenders will then jump in and say “but our BPM system is flexible and can be used by an executive for knowledge work just as well as routine work.” I have an answer for this, but let’s start by clarifying two key points.

First, I am not asking if a particular product might or might not support knowledge work. Many products go beyond BPM, and I am not in the position to judge particular products. Instead, this statement is based on the most common definition of BPM and what must logically follow from that.

Second, I do not want to get into a lengthy discussion of the definition of BPM. There are many contradictory opinions on exactly what constitutes BPM. Ignoring the exotic definitions, BPM is clearly about “management” of “business processes,” and without any doubt the process is the central concept which is to be managed. Most descriptions of BPM include the idea that a representation of the process (a process definition) becomes an important organizational asset which is improved over time. The purpose of BPM is to collect a library of best practice process definitions, which are optimized either through analytics or simulation.

It follows logically that in order to have a representation of the process, then you must be able to predict the process. But emergent processes such as you see in knowledge work are not predictable, and therefore are not representable as a process definition. In order to support knowledge work, we must use something other than a process definition. For this reason I believe it is clear and unequivocal that BPM, because of its focus on process is in-

appropriate for knowledge work. There are three responses you might encounter.

First, some vendors are sure to say “Our system is so flexible that you can support work without having a pre-defined process definition.” That is great. Yet this capability goes beyond the definition of BPM. This non-BPM capability may be useful for knowledge work. I am not trying to argue that none of the BPMS products can handle knowledge work, but that BPM by its definition is not appropriate for knowledge work and emergent processes.

Second, some may claim that knowledge work is a simple fixed process: (1) figure out what to do, and (2) do it. This extremely generic two-step fixed process is the simplest possible, and more elaborate 3-, 4- or more step generic processes have been suggested. The problem with these generic processes is that the history that results is generic and meaningless. Seeing that 1000 people figured out 5000 generic things to do and did 4000 generic things today is not useful to the organization. To track what knowledge workers are accomplishing, the activities need to be specific, not generic.

Third, some have technology that allows what is needed, and argue that it should still be called BPM because the knowledge worker is directly and individually involved in the management of those business processes. While they don't start with a pre-defined process definition, within the scope of a single case, they define and improve the process. With this narrow interpretation of the terms, the technology that supports knowledge workers might be called BPM, but doing so would do violence to many of the other aspects of BPM that are normally assumed, such as cyclical improvement of processes and finding best practices.

A DIFFERENT STYLE OF MODELING

In the end, if one insists that the technology that supports knowledge workers be called a BPMS, then so be it, but there is a serious danger in doing this. Almost all BPM systems existing today are designed to separate the role of process definers from the end users. This happens from the desire to make the most faithful representation of the process possible. To make a process that runs accurately 1000 times, you need a specialized skill and a complex language. For a process that is created on the fly by the case manager, and then discarded, the requirements on the language are very different.

Consider the restaurant analogy with two kinds of restaurants: a sit-down restaurant and a self-serve buffet. The final dinner plate is analogous to the process. In a sit-down restaurant, you order a meal, and it is delivered completely prepared on the plate, ready to eat. This is like BPM where the end user chooses a process definition and runs it. In a buffet the diner starts with an empty plate. From the buffet table, the diner selects and combines the selected food on the plate, ultimately ending up with a meal to eat. The buffet table does not consist of prepared plates, but instead food in a different form ready for the diner to select from. In ACM, the end user (the case manager) composes the process from parts that are made available in the template.

In BPM, the process is designed by a specialist, and is meant to be used by a user without the user having to compose anything. This makes the processes more elaborate, complex, and fully defined. In ACM, the case manager composes the process while the work is going on. This means that the processes

have to be much less complex, much less complete, but much easier to manipulate. It would be a grave mistake to assume that one style of process modeling will be the same for both domains.

There is quite a bit of interest in presenting a simple checklist representation of the model to a case manager (Swenson 2001). While a linear checklist does not contain the fine details that would be necessary to have a single process definition run for thousands of cases, it is however convenient when the task list is needed only for a single case. An elaborate diagram using BPMN would be too much trouble for a process will be run only once. Also, the checklist metaphor is easy to all people to understand. Adding a task into a list of tasks is uncomplicated. It is possible that for one-off process definitions, a checklist approach may be the best.

A DIFFERENT METHOD OF IMPLEMENTATION

When implementing a BPM solution, the first step is “process discovery” where one attempts to write down the process, either by interviewing the participants, or by mining the log files of what has occurred in the past. This process definition becomes the core of the new solution.

Adaptive Case Management does not start this way. The first thing you do is define the entities involved. For a medical case template, you start with a representation of a patient, including medical history. You might define some roles to represent different kinds of interactions that might be allowed, and access control to that information. There is no necessity to define a process, but often a collection of pre-defined activities is included in the template so that the case manager can bring them in with minimal trouble. You may also have a collection of process fragments which can be brought in after the case has started. A case instance can be started almost completely blank and nothing more than a case file to place things.

CONCLUSIONS

In order to summarize what has been presented, we see that knowledge work is widespread (40 percent of the workforce) but is unpredictable. The four main characteristics of knowledge work are given. Because it is unpredictable, BPM, which is based on defining a process definition ahead of time does not work. Instead, a different approach is needed, called Adaptive Case Management (ACM). ACM has a different way of modeling processes from BPM, and is used as part of a different methodology. It is possible that a single product may contain technology that can support both predictable and unpredictable process support, you should approach such claims skeptically because of the way that knowledge work falls outside of the expectation of what is normally considered “work.”

Twelve experts in this space came together to collaborate on a book called “Mastering the Unpredictable” (Swenson 2010) that focuses on the need to support work which is not predictable. Some of this article contains excerpts from that book, mixed with additional information uncovered during the launch discussions. Mastering the Unpredictable continues to provide 350 pages of description and discussion on how Adaptive Case Management might be used to support unpredictable work. Find out more at <http://www.MasteringTheUnpredictable.com/>

ACKNOWLEDGEMENT

The author gratefully acknowledges the support of Fujitsu America, Inc. for the research on these topics. Fujitsu produces the Interstage family of middleware products, including Interstage Business Process Manager a full featured BPM Suite supporting server to server integration, human routine process automation, dynamic process capability and collaboration platform for knowledge workers. The suite is rounded out with a process analytics server, and a unique Automated Process Discovery capability which can mine the process out of existing application log files to be exported onto the suite platform to streamline and automate the business processes. Visit <http://www.fujitsu.com/interstage> for more information.

REFERENCES

- (Davenport 2005) Thomas H. Davenport. *Thinking for a Living: How to Get Better Performances and Results from Knowledge Workers*. Boston: Harvard Business School Publishing, 2005.
- (Gogek 2007) Jim Gogek. University of California, San Diego, "UC San Diego's High-tech Tools Helped Combat Wildfires: Emergency Advancements Can Aid Disaster Worldwide," News Release, November 5, 2007. <http://ucsdnews.ucsd.edu/newsrel/general/11-07HighTechToolsCombatFiresJG-L.asp/>.
- (Moore 2009) Connie Moore and Craig Le Clair. "Dynamic Case Management—An Old Idea Catches New Fire." Forrester Research (December 28, 2009). http://www.forrester.com/rb/Research/dynamic_case_management_%26%238212%3B_old_idea_catches/q/id/55755/t/2/.
- (Spear 2009) Steven J. Spear. *Chasing the Rabbit: How Market Leaders Outdistance the Competition and How Great Companies Can Catch Up and Win*. New York: McGraw-Hill, 2009.
- (Swenson 2001) Keith D. Swenson. "Workflow for the Information Worker," in *Workflow Handbook 2001*, edited by Layna Fischer, 39–49. Lighthouse Point, FL: Future Strategies, 2000. [http://kswenson.workcast.org/2001/Workflow for the Information Worker.pdf/](http://kswenson.workcast.org/2001/Workflow%20for%20the%20InformationWorker.pdf/).
- (Swenson 2010) Keith D. Swenson. *Mastering The Unpredictable: How Adaptive Case Management Will Revolutionize The Way that Knowledge Workers Get Things Done*. Meghan-Kiffer Press, 2010.