
Workflow for the Information Worker

Keith Swenson, MS2 Inc., USA

OVERVIEW

Workflow technology has a clear application in the domain of highly repeatable business processes. There is, however, a large body of business activity that so far has proven difficult or impossible to support with production workflow systems. Information workers find that the objects of their work change form substantially from case to case, and that the procedures that they use are highly dependent upon the details of the current case. Because of its non-routine qualities, this work is often an organization's most important and most costly. While more and more of the routine work is becoming automated, we find more people doing less routine activities; organizations are finding an increasing part of their workforce performing information work. This paper considers the needs and requirements of information workers in order to best support, if not automate, their daily activities, and how these requirements differ from those for production workflow.

EXAMPLES OF WORK

Production workflow applications help to support well-known and well-defined business processes. Three such jobs are briefly described below followed by three information worker jobs for comparison.

PRODUCTION WORKFLOW EXAMPLES

A bank teller performs many mission-critical activities that can be supported by production workflow. There are a set number of procedures that are done at the counter. Each procedure triggers follow-up actions by others within the organization. It is critical to the correct operation of the bank that a teller's every action is well defined and consistently performed across a large number of tellers.

The processing of insurance claims is the second example. To process the claim, a number of people need to do a number of activities in a well-defined manner; the actions must be performed in the right order, without skipping any steps or unnecessarily repeating any steps. Since each person does a small part of processing independently, each part needs to be rigidly defined.

The third example is loan processing. Again, being financially oriented, there are rigorous rules to be followed, both in terms of accounting practices as well as laws that define how such organizations are allowed to operate. Since the product supplied, such as money in the case of the loan application, is identical from various sources, the market favors those companies that can handle a large number of

loans accurately and correctly. Production workflow helps by assuring the repeatability of a process that is done by many individuals.

INFORMATION WORKER EXAMPLES

The term "information worker" is used here to describe a collection of jobs that are less rigorously formalized yet are done regularly and repeatedly by a large number of people. While all jobs of this nature might not be accurately called "information work", it is the case that information work inherently tends to fall into this category.

The first example is the development of a *press release*. Organizations produce press releases on a regular basis to inform the market about their doings. Each release involves a number of people who specialize in different aspects; the originators provide the initial content; the legal department keeps the company from violating laws or being sued; corporate marketing assures that certain corporate standards are being met; product divisions assure that a consistent message is presented across a range of products; it is reviewed from many points of view to assure that it is clear and concise; and there are press relations specialists to assure that the form of the release is suitable for the channels through which it will be released. If the press release involves multiple companies there are even more reviews as many roles from each company need to get involved.

One can see a lot of the elements of a business process: there are requirements to be met and different roles to be played by different individuals. This is hard to automate because the exact process depends so much upon the information in the press release; some releases have extremely sensitive information requiring a lot of special handling; press releases involving a dozen companies must be handled very differently than a single company press release; some press releases are more momentous than others, and require special considerations.

A press release about a product needs review by the product team. A press release about an event needs review by those that set up the event. A press release from finance will have very different review steps than one from the research labs. At an abstract level, press releases go through a process with similar steps, but the details reveal that the process is highly dependent upon the content.

The second example is that of a *market competitive analysis*. Every company developing products or services will do a competitive analysis to determine how much chance there is of selling a particular product, to whom to sell it, and how to position the product. While a single analysis may not involve a large number of people, with most of the work often being done by a single individual, the results must be read and understood across the company. It is the need for consistency that drives an organization to set up a standard process of review and approval.

Like other examples of information work, the steps necessary, and the people who need to be involved in the review, are highly depend-

ent upon the type of product and type of market. It is very likely that the actual activities performed will vary greatly from one instance to another. To accomplish this, the successful analyst draws upon training and experience of the past, and applies it on a case-by-case basis.

The third example is *product development*. How best to define and develop successful products has been the subject of countless studies, and the literature is filled with a rich variety of recommended approaches. At a very high level, there is agreement on the essential form of such work: generally, the process starts with creation of a product requirements document. From that, a functional specification is created and this is followed by a detailed design specification. A documentation plan may be required. Quality Assurance may produce a test plan. Ultimately some form of Release Plan and a Lifecycle Plan will be required.

While there is agreement on the basic abstract process, there has been little success in automating this process with production workflow. This is because the specifics of the process depend so heavily upon the specifics of the product being produced. Even the definition of what constitutes a requirements specification varies greatly not only across industries, but even within a single organization from case-to-case. Large projects involve many more people, and more specialized review, than small projects. Products aimed at specific sub-markets require special considerations unique to that market.

CONTRASTING THE TWO

In the first three examples, we see processes that can be rigorously defined and will remain the same across a large number of instances. The number of exceptions that need to be handled outside of the process are relatively few. There is generally a large scale of use, and there is a lot of value to be gained by executing the steps of the process consistently every time.

In the three examples of information work, agreement exists at a high level that a process is to be followed, but the details of the process depend heavily upon the specifics of the case at hand. Part of the job is to make a *plan of action* on how to complete the rest of the work. It is this self-referential quality of information work which, more than anything else, makes it hard to automate with production workflow.

PROCESS SUPPORT NEEDED

While the need for information workers to follow a process is widely recognized, such processes are not widely followed. Surveys of over 200 product developers over the past five years found that fewer than five percent had a completed specification before beginning production design. Only 58 percent of requirements were specified before design activities began on the average.[1] Another study found that business units that applied process management to production development reduced their average development times by 30 percent to 50 percent.[2] Workflow might have a significant impact on the effi-

ciency of an organization because it can be a way to improve communication and make status more visible.

Organizations value creativity highly in the execution of information work. These jobs tend to be more highly paid, and require a higher skill level than production jobs. A formal process that forced these jobs to be done in the same way every time would actually hinder the ability information worker to be successful. Are we then to conclude that there is no hope for supporting processes for information workers? The answer lies in understanding what it means to *support* a process.

ELEMENTS OF COORDINATION

The underlying problem that workflow attempts to solve is the coordination of a group of individuals while they work concurrently. In order to coordinate their own work, a person needs to understand (1) the goal of the work, (2) what others have done, and (3) what he or she needs to do next. These three needs are the same whether you are talking about either production or information work.

Let's use the term *item* to denote a discrete unit of coordinated work. For the bank teller, an item might be an account transaction. For the insurance company, a claim would be considered an item. The press release is an item. For product development, a project is an item that is composed of smaller items: the requirements document, the functional specification, etc. An item is a conceptually meaningful unit to the organization around which people coordinate their work. In many cases, the item is a document of some sort—but not all cases. People can discuss the status of an item. Using workflow terminology one would say that an item is a *process instance*.

In a production work case, because the process is well defined, people can work in relative isolation of each other. The bank teller needs not personally know the back office auditor, nor vice versa. For claims adjustment, the fact of being in a particular state of the process tells a participant a lot about what has happened before and what needs to happen next.

For the case of information work, the lack of a well-defined process, along with the need to make room for creativity, means that there is a much greater need to communicate details about what has been done and what should be done. Given that the goal of the item is clear, there are two strategies for meeting the other two needs. The first is to provide to participants a lot of information about what has transpired on a particular case; this way the participant can determine what to do next based on experience and training. The second strategy is to have an active coordinator who is aware of the history, and provides fairly detailed guidance to others on what needs to be done next.

This second strategy is employed most often. Each item being developed has an explicit *owner* who takes it through all the steps needed. The owner is responsible for reaching the final goal, and therefore

defines what needs to be done. This is sharply different from the production workflow where there may be no identifiable owner for a given item.

It is interesting to note that while the process for information work is abstract, it is perfectly suited for status discussions. It is common to talk about a particular document that is "being reviewed" without being specific about what this really means. "Being reviewed," means very different things for different documents.

It would seem then that in these organizations, workflow might be valued for the ability to produce a report showing all the items along with each item's status. This gives people an overview of what the organization is doing so they can better focus their resources on long-term goals.

Talk about status is meaningless without an understanding of the process, so the second value that workflow brings is a way to communicate the process. As a learning organization grows, its processes change. By providing a way to communicate the changes in a process, a workflow system can actually help promote a learning organization.

BALANCING FLEXIBILITY AND CONSTRAINT

In order to support communications regarding the status, and to support the change of process over time, there needs to be a representation of the process.

What does a workflow process for information workers look like? Because the organization values creativity, and is using skilled people to determine the right thing to do at the right time, the process needs to allow flexibility in its usage.

One potential approach might be to make a formal workflow process that maps all possible courses of action, with decision points at all appropriate points for the owner. This approach is flawed because of the complexity of such a process.

Complex processes have two problems; one is that they need a lot of context data in order to work correctly. If mistakes are made entering the context data into the system, the process can take an incorrect turn. The second problem is fragility; if laws or corporate rules change, the cost to rework a complex process can be significant. It is almost impossible to make a process definition that would be complete enough to handle all the situations, and if you did make such a process definition, it would be so complex as to be difficult to use accurately.

A second possible approach might be a workflow system that provides an approximate representation of the process, and which the owner can change on the fly to fit the needs of the instance. Experience with the Regatta Project [3] at Fujitsu in the early 1990s suggests that process definition is difficult and time consuming for the average information worker. Even a simplified visual process defini-

tion language requires thinking about what is to happen in a formal way that is unfamiliar to most people. Those who learn the skill of defining the process find that keeping the process up to date is time consuming beyond the value. For practical reasons you cannot expect an information worker to define the processes that they are using—they are too busy doing the real work and there is not enough return from the investment.

For the information work, the process is not concrete enough for formal process support. The organization does agree on a set of abstract activities to be performed, but the representation of these activities must be oriented less toward automating the work, and more toward supporting the owner in communicating the status of an item. Instead of a formal process representation, it might be better to simply list the abstract activities such as that shown in Figure 1.

Status	Activity	Expected	Actual
<input checked="" type="checkbox"/>	Requirements Gathering and Validation	13-Mar-00	26-Mar-00
<input checked="" type="checkbox"/>	Release Planning	20-Mar-00	26-Mar-00
<input checked="" type="checkbox"/>	Authorization to Proceed	21-Feb-00	27-Feb-00
<input checked="" type="checkbox"/>	Feature Planning	27-Mar-00	05-Apr-00
<input checked="" type="checkbox"/>	Design	10-Apr-00	26-Apr-00
<input checked="" type="checkbox"/>	Development	29-May-00	02-Jun-00
<input checked="" type="checkbox"/>	Integration Test	09-Jun-00	09-Jun-00
<input checked="" type="checkbox"/>	Alpha Test	05-Jun-00	09-Jun-00
<input checked="" type="checkbox"/>	Beta Test	15-Jun-00	30-Jun-00
<input type="checkbox"/>	Release Approval		
<input type="checkbox"/>	Released	04-Aug-00	
<input type="checkbox"/>	End of Life Notification		
<input type="checkbox"/>	End of Life		

Figure 1: Process for Product Development

The owner of an item has a lot to gain from a simple list of the normal activities in the order that they usually appear. Ideally, the process should help guide the owner toward the normal course, without constraining him or her to stay on it. Such a checklist can be a reminder of what needs to be done. The owner might check off activities as they are completed. As the item progresses through the states, other participants can see the current progress.

The determination of when a particular activity needs to be started or completed is a judgement call of the owner. Therefore, traditional process logic that controls the progress from step to step is not a help, but rather a hindrance to the owner. You might say that the owner provides the process logic. From a process enforcement point of view, there is very little control of the process, which in fact must be to allow the information worker to do the job. Because the activi-

ties are abstract, it would be impossible for a workflow system to determine if an activity was completely performed or not.

Even without the process enforcement, the process still helps the group coordinate by making people aware of the current state. For example, consider a typical item, such as a design specification, that starts out in the Draft state as shown in figure 2. While in draft state, the author is assumed to be the only person working on the document; others tend to avoid doing anything with the document, even reading it. Later, if the document is in review state, then the author is doing relatively little with it, and others know that they might now read and give feedback on it. Even later, when it reaches the Approval state, people know it is too late for changes. The presence of a list of activities helps the whole group to be better informed about the status.

Status	Activity
	Draft
	Review
	Approval
	Complete

Figure 2: A typical document process

PLANNING AND STATUS REPORTING

Manual control of when the process moves forward is necessary but not sufficient. The owner will need more degrees of freedom when the process depends in significant ways on the subject of the work. The process itself is often emergent, not known beforehand and only discovered as it progresses. In order to support these situations the workflow process needs to support some special cases.

It will happen that the owner needs to get some activity performed on the item that is not in the process, and quite possibly has never been done before. For example, imagine that an owner determines that the item needs something we can call “Special Processing.” One might imagine a system where the owner could add the activity into the middle of the process as shown in figure 3.

Doing so for a single item would be misguided, because the way that the abstract activities are used as indicators of current status. To say that the item is in the “Special Processing” state is meaningful only to those people who understand what the owner means by this activity, and where this activity occurs in the process. Because the information worker is using the process primarily as a way to communicate status, one can conclude that only if there is general consensus within the organization about an activity can one be added to the process. This strongly implies that the information worker may identify extra activities to be done, but must handle these informally outside of the process.

Status	Activity
<input checked="" type="checkbox"/>	Draft
<input checked="" type="checkbox"/>	Review
<input checked="" type="checkbox"/>	Special Processing
<input type="checkbox"/>	Approval
	Complete

Figure 3: Adding a “Special Processing” activity is not meaningful unless there is broad agreement about what it means.

It is more reasonable to expect that the process will include a large number of possible activities from which the owner may select what is to be done. An owner will decide that certain activities do not need to be done for this particular item, and can mark those activities as being skipped.

Figure 4 shows such a process for a product feature where the double down arrows are used to indicate unneeded activities. It might be that the process was intended to represent all the options for a large product feature but in this case, it is a very small product feature.

Status	Activity
<input checked="" type="checkbox"/>	Requirements
<input type="checkbox"/>	Functional Specification
<input checked="" type="checkbox"/>	Design
<input checked="" type="checkbox"/>	Development
<input type="checkbox"/>	Review
<input type="checkbox"/>	Feature Test
<input type="checkbox"/>	Integration
<input type="checkbox"/>	Integration Test
	Complete

Figure 4: A process with skipped steps

The execution sequence needs to be flexible. It is common for information workers to start an activity before the previous activity is finished; the workflow system must allow activities to overlap in execution. Activities will usually be done in order, but not always. The dates in Figure 1 show that the third activity “Approval To Proceed” was expected and achieved before the other activities. This, too, must be allowed by the workflow system.

It is also possible that due to a fundamental change in the item, the process may need to be sent back so that some of the activities may be performed again. One example of this is in a Press Release when the legal review turns up a fundamental problem that requires rework; the revised form of the release often needs to be reviewed again by participants that were earlier in the process. From this we can

conclude that the owner needs the ability to reset activities in order to perform them a second time.

In summary, the requirements are that the process needs to be relatively fixed, but the list of activities can be treated as a cookbook from which the item owner can choose and use at discretion. This is quite different from production workflow. One might even question whether this really is workflow, and whether it is possible to interoperate with other workflow systems.

WF-XML AND INFORMATION WORK

Knowing what workflow for information workers would look like; we can consider how it fits with the existing standards for workflow as defined by the Workflow Management Coalition (WfMC).

Workflow for the information worker is consistent with the WfMC architecture [4] in a number of ways:

- The process consists of a set of activities
- Each activity can be started and completed
- The items are essentially process instances
- The information associated with the item is process application data.

The biggest difference is simply that the process logic is essentially missing—the activities must be manually started and completed by the owner; yet, the WfMC architecture does not prohibit such trivial process logic. If the workflow for information workers fits the WfMC architecture, then we know it might be able to interoperate with other workflow systems using WfMC interoperability standards.

Wf-XML [5] is a format for encoding the current status of a workflow process instance. If an external system wished to monitor the status of a particular “Product Development” item, a Wf-XML representation of the current status could then be returned as the result of an automated request. Wf-XML also provides a way to manipulate the process instance. It can be used to start new items, or to terminate existing ones.

It will be possible to link the activities of one workflow system to external processes. The starting of an activity in one system might automatically cause a process in another workflow system to start. Such linkage is defined in the OMG Workflow Management Interfaces [6] and is expected to be part of future versions of Wf-XML.

One might imagine this capability to be used to link together two groups of information workers. For example, one group might own the overall product development item for a particular product using the process shown in Figure 1. The first step is “Requirements Gathering,” which might automatically start a Requirements Specification item. If this item is the responsibility of a different group, as it often is, then it would be possible using Wf-XML to communicate to that group’s workflow server, and to start the item process there. Since

Wf-XML is designed to support notifications and status monitoring, the originating information worker would still be able to track the progress of the Requirements Specification, and even to receive the final copy of it when complete.

This sort of integration could prove to be very useful for information workers in order to interface to more production-oriented parts of their organizations. The information worker may create an item and then, at a particular step in the process, hand the document over to another part of the company by submitting it to a production workflow process via a Wf-XML interchange.

SUMMARY AND CONCLUSION

It stands to reason that as the more formalized tasks become automated, more people within organizations will spend their time doing less structured “information” work. By 2006, nearly half of all U.S. workers will be employed in industries that produce or intensively use information technology, products, and services, according to U.S. Department of Commerce projections. Over the century, mass-production occupations have been steadily replaced by office-worker and service-provider occupations. Indeed, virtually all of the jobs that were lost in goods production and distribution since 1969 have been offset by office jobs. [7]

In order to support the kinds of jobs that information workers do, workflow technology will be required to:

- Represent the organizationally agreed upon processes as a list of relatively abstract activities;
- Use the activities as an indication of the state of progress of an item;
- Allow for a “strong ownership” model where the owner of an item is relatively free to choose:
 - Which activities are performed,
 - Which are skipped,
 - Which order they are performed in,
 - And which activities overlap;

Such workflow will be viewed as supportive of information work, aiding the communication necessary to coordinate across an organization, without being constraining and restrictive, which would get in the way of the flexibility and creativity needed to do the job.

Such workflow systems can still be made to be consistent with the WfMC workflow architecture, to use workflow standards, and to interoperate with other workflow systems.

The result of using workflow for information workers will be more visible business processes, better communication about status, and

better coordination of work among team member, making the organization more efficient, with fewer time-consuming status meetings.

REFERENCES

- [1] Stefan Thomke, Donald Reinertsen, "Agile product development: Managing development flexibility in uncertain environments", California Management Review; Berkeley; Fall 1998
- [2] Paul S Adler, et al. "Getting the most out of your product development process," Harvard Business Review; Boston; Mar/Apr 1996
- [3] K Swenson, et al. "A business process environment supporting collaborative planning" Journal of Collaborative Computing, March 1994
- [4] Workflow Management Coalition, "The Workflow Reference Model," WfMC-TC-1003, Version 1.1, Jan 1995, available online at <http://www.aiim.org/wfmc/standards/docs/tc003v11.pdf>.
- [5] J Hayes, E Peyrovian, S Sarin, M Schmidt, K Swenson, R Weber, "Workflow Interoperability Standards of the Internet" IEEE Internet Computing, May/June 2000
- [6] "Workflow Management Facility" Object Management Group Joint Submission, revised submission, July 1998, available online at <ftp://ftp.omg.org/pub/docs/bom/98-06-07.pdf>
- [7] US Department of Labor Report, "Futurework: Trends and Challenges for Work in the 21st Century," September 17, 1999

Bio and Affiliation

Keith D Swenson

MS2 Inc.

2440 El Camino Real,

Mountain View, CA, 94040

+1 650 623-2329

kswenson@ms2.com

Keith is currently Director of Engineering / Software Architect at MS2 Inc, producer of a collaborative Product Lifecycle Application known as “MS2 Accelerate”. Previously he was Software Architect at Netscape for the internet workflow product “Process Manager” and the developer tool “Visual JavaScript”. Before that, at Fujitsu, he started the Regatta Project, which produced the software products “TeamWARE Flow” and “iFlow”. Before that at Ashton Tate he managed the development of Framework IV. From 1995 to 1997 he served on the ACM Special Interest Group for Group Support Systems. He has participated in the development of a number of workflow standards and, in 1996, was elected a Fellow of the Workflow Management Coalition.