
Understanding and Implementing Workflow- first Design (Workflow Sprints)

*Why you should ditch the
technology-first lens and think
about workflow.*

Technology! Workflow!

In this article I'll be looking at a critical and under-examined step required to build publishing technology - optimizing workflows.

I am in the business of building publishing technology. However, I actually believe the core of what I do is designing and optimizing publishing workflows. Technology is just an (optional) outcome of that process.

To get to good technology, it is very important to start the process by not thinking about technology at all. If you frame your thinking through the lens of technology, you will more than likely end up with an expensive failure. What we are really interested in are good workflows and, later, designing technology to encapsulate and enable good workflow.



To help us understand this problem we need to define workflow as it means many different things to different people - by workflow, I mean the complete set and order of actions an organization (usually a publisher) undertakes to prepare an object (book, journal article, preprint review, question set etc) for publishing or sharing.

For example, workflow is everything that is done from start to finish that falls within a publisher's control, to prepare a book or journal article for publication. Or, every action required by a community to prepare and share reviews on preprints.

Note: I use the term 'publisher' very liberally (see definitions at bottom).

What publishers want is to improve their publishing processes, but they tend to think purely in terms of technology as both the problem and the solution. As a result of this 'technology first' thinking the process for many publishers looks like this:

1. Choose a new technology
2. Use the technology

It is worth noting that there are some good motivations for wanting to change technology which do not originate in trying to improve workflow. Liberating an organization from expensive licensing of proprietary technology, for example, is an increasingly common driver.

Replacing unstable legacy technology is also a good motivation.

However, regardless of the motivation, if you are determined to change technology, you should also see this as an opportunity to improve how you work - to improve your workflow.

As mentioned above, the “technology first” approach can be problematic. Three common outcomes of technology-first thinking are:

1. You can make an expensive technology change without improving the publishing process.
2. You may create new problems and poorer outcomes for the publishing process.
3. A ‘rush to adopt’ can come with unforeseen tradeoffs that later may prove very costly.

These undesirable outcomes are better avoided by adopting a workflow-first framework.

For new entrants and incumbents alike, a workflow-first process looks like this:

1. Understand the current (or speculative) workflow
2. Optimize the workflow
3. Design technology to enable the workflow
4. Build (or choose) the technology (optional)

In these two articles we will look at steps 1 and 2 only. Steps 3 and 4 might require a thesis.

Step 1: Understand the current workflow - is necessary to generate an understanding of what you are actually doing right now. This step needs to drill down into the nitty gritty of the daily workflow so you can get a complete picture of what actually happens ‘on the floor’. High level executive summaries are not what we are after. We need to know what actually happens.


Without understanding exactly what you are doing now, improving what you do will be based on false assumptions. As poet Suzy Kassem has said:

"Assumptions are quick exits for lazy minds that like to graze out in the fields without bother."

When examining workflow we need to bother people. To ask hard questions and get to the bottom of things. I've often heard folks say they know their workflow, or someone within their organization knows the workflow. But after asking a lot of exploratory questions I inevitably find that no one person understands exactly what happens at all stages of the workflow. Consequently, understanding your current workflow, to the granularity required to move to the next steps, requires a process in itself. But we will discuss this in more detail in the next article.

However Step 1, while important (and important to get right), is just prep for the most critical step of the entire process:

Step 2: Optimize the workflow Optimizing the workflow is what every publishing technology project should aim to do. If you don't optimize the workflow, if you haven't improved what you do, then what have you achieved? It would make sense then, to prioritize thinking about improving workflow before you start to think about technology.




Optimizing the workflow is, in itself, a design problem and requires dedicated effort. Unfortunately this step is often skipped or paid only a passing tribute. Conflating steps 2 and 3, for example, is typical of the popular publishing RFP (Request for Proposals) process (see <https://upstream.forcell.org/can-we-do-away-with-rfps/>) :

1. Understand the current (or speculative) workflow
2. Design technology to enable the workflow
3. Build (or choose) a new technology (optional)

An RFP process may be motivated by a few big ticket items that may need improving (or licensing issues as per above), but often lacks a comprehensive design process aimed at optimizing the workflow. Instead, technology is chosen because it fulfills the requirements of the current workflow audit (Step 1) with a few possible fixes - this is because most RFP processes focus on updating technology rather than improving publishing processes. Consequently there are minimal, if any, workflow gains and a failed opportunity to identify many optimisations that can deeply affect a good technology choice and improve how you publish.

To avoid these problems and to improve how you publish, it is really necessary to focus dedicated design effort to optimizing workflow.

The kind of process I'm talking about was once brilliantly described to me by a friend Martin Thompson (Australian Network for Art and Technology) a long time ago. Martin was explaining to me the principles of video compression for streaming media, but the principles can be applied to workflow optimisation. Imagine a plastic bag full of air with an object in it (let's say a book). The air represents all the time and effort required to produce that book. Now suck out the air, and we see the bag and book remain the same but the volume of air has reduced. We have reduced the effort and time to produce the book, while the book remains unchanged.



This is a simple but clear metaphor for what we are trying to achieve. In summary, optimizing workflow is about saving time and effort while not compromising the end product.

Workflow optimization requires an examination of every stage of the process (as discovered in step one of our workflow-first process) and asking if that stage can be improved or eliminated. We are also asking some other questions such as -

- Can anything be improved to help later stages be more efficient?
- Are there 'eddies' in a workflow that can be fixed?
- Can operations occur concurrently rather than sequentially?

When asking these questions we will discover many efficiencies and we can slowly start drawing out a complete improved workflow. Some steps will disappear, some will be improved by mildly altering the order of operations, others will be improved by small tweaks. Some changes are more significant and may require greater change.

Once we have been through every stage looking for efficiencies, once we have our complete and optimized workflow, we are ready to start thinking about what we need to enable these changes. Technology may be part of the answer, but it is never the only answer, and in some cases it is not the answer at all. If we do decide new technology is required then only now should we start on the path to designing or choosing the tools to support our optimized workflow.

Next we will look at what a methodology for understanding and optimizing workflows can look like - Workflow Sprints.

What is a Workflow Sprint?



A Workflow Sprint is a fast and efficient methodology for helping organizations understand and optimize their publishing workflows before making technology choices.

I designed the methodology after witnessing many publishers making these common mistakes:

1. Investing in slow and costly 'requirements research'

2. Looking at the problem exclusively through a 'technology lens'

Understanding and optimizing workflow need not be a slow or expensive process. Too often, these processes are lengthy research projects that interview individual team members and produce immense amounts of narrative and flow diagrams. Such outputs are seldom clarifying. On the other hand, Workflow Sprints produce crisp clarity in a process that can be counted in days, not months or years.

A Workflow Sprint first investigates and describes the current workflow and then produces an optimized workflow. From these outputs, your organization will be in a good position to consider technology choices.

Principles of Workflow Sprints

Workflow Sprints are a series of facilitated sessions that can occur remotely or in real space. There are some fundamental principles of Workflow Sprints:

Principle 1: Operational Stakeholders are best placed to describe and optimize workflow

The best people for describing an existing (or speculative) workflow and optimizing that workflow are the people actually doing the work now - your team (operational stakeholders - more on this below).

Principle 2: Group discussion simultaneously expedites the process and adds clarity

Group discussion involving the operational stakeholders is the most effective way to understand and optimize workflow. Having these experts in the same room eliminates guesswork and team members can drill down into an issue together in quick time to tease out the fidelity needed.

Principle 3: Facilitation is key

Experienced facilitation conducted by someone external to your organization is critical for keeping the momentum going and managing any potential and problematic team dynamics.

Principle 4: Keep the documentation light

If the documentation of the workflow is not something decision makers can easily grasp, then you are doing it wrong. For this reason, I have designed Workflow MarkUp (WFMU) to aid in expediting the documentation process while producing clear, human-readable documents.

Now that we have the principles covered let's take a brief look at what a Workflow Sprint actually looks like.

What does a Workflow Sprint look like?

Photos included are from actual Workflow Sprints (arXiv, Open Education Network, Ravenspace, Érudit, Organisation for Human Brain Mapping, Wormbase/Micropublications.org).



A Workflow Sprint is a facilitated, structured conversation between the organization's primary stakeholders. The stakeholder group should comprise those people that are part of the current workflow (constituting all parts of the workflow) and decision-makers.



Operational stakeholders are the people who do the work. These are the most important people to have involved because they are the ones that understand what is done, on the floor, at this time. Sometimes, if automation is involved, a trusted technician that knows the limits and possibilities of publishing technology might be useful. No Workflow Sprint can occur without operational stakeholders.



Strategic stakeholders are the organizations decision makers. The decision-makers may, or may not be part of the process, however, they will certainly utilize the outcomes to guide them in making decisions going forward. It can be a good idea to include strategic stakeholders to build internal buy-in.



The Workflow Sprint draws these two types of stakeholders together in a series of remote or real space sessions. If held in realspace the event can occur over one day. For remote events (which we have found to be more fatiguing), the Sprint should be broken down into multiple shorter sessions over several days.



Workflow Sprints are fun, efficient, cost-effective, and fast. The process is a good way to generate additional internal buy-in for a project and to quickly help everyone understand the broader goals and challenges.



We have found that Workflow Sprints are also extremely effective team-building events. If held in real space, we advocate good food and coffee! If done remotely, we sometimes include some less formal moments to help people relax and get into the flow of the event.



Now let's look at the process in more detail. The following is not comprehensive, but it should give you a good starting point. If you want to drill down further into the process, feel free to contact me.

Speculative vs Extractive Workflow Sprints

Before we get started, a quick word about two types of Workflow Sprints that I have already hinted at throughout this document. The two types are:

Speculative - Speculative Workflow Sprints imagine new workflows. Typically new entrants or start ups require speculative sprints. Inevitably the speculative variety can reveal a lot of contradictory requirements and 'blurry vision' that need to be discussed and examined in detail. Sometimes, the outcomes may also lead to compromises the stakeholders may find difficult to accept. Consequently, Speculative Workflow Sprints often take longer and require a higher degree of expert facilitation.

Extractive - Extractive Workflow Sprints start by examining a known (existing) workflow. Generally speaking, these descriptive processes are easier to facilitate and quicker to document.

Real Space vs Remote

A word regarding remote vs realspace events. In-person sprint processes are definitely more effective and quicker. Remote sprints are possible, however, the dynamics and format change somewhat. Documentation for remote events is usually done with shared virtual whiteboards, and the process is usually split into several shorter events spread over a number of days.

Hybrid Workflow Sprints (a mix of virtual and in-person participants) are also possible but are harder still, mainly due to the fact that the tools used for

documentation are necessarily different, which makes it harder to share ideas across a team split into real and virtual spaces.

The Team

When hosting a Workflow Sprint the first question you need to answer is - who do I invite? There is a simple rule - invite enough operational stakeholders, so you have every part of the workflow covered. Having strategic stakeholders may be (not always) a good idea but do keep in mind the more people involved that don't do the work 'on the floor' the longer the process will take. Having said that, strategic stakeholders stand to learn a lot during the process if they participate.

In the case of 'speculative' (new) workflows, you need to gather the people with the vision and some folks that know about the kind of work you have in mind.

The Space

If the Sprint is held in real space, it is good to have a single table that can seat everyone involved. In addition, break-out spaces with enough whiteboards or paper and markers to go around are necessary. Post-it notes are welcome but not required.

Facilitation

It is very important to have a neutral facilitator. The facilitator needs to ask hard, bothersome questions, and consequently, it is much easier to do this if the facilitator is not subject to internal power dynamics and existing interpersonal relationships. In addition, the facilitator must:

1. **Facilitate conversations** - The facilitator must hold the space for the group to come together and share their knowledge and ideas.
2. **Manage the process** - The facilitator must keep the process moving forward.
3. **Balance group dynamics** - The facilitator constantly reads the group, preempting obstacles to collaboration and productivity and adapting the strategy accordingly.

Facilitation is a deep art and there is way too much to go into here (yet another thesis...). However I strongly advise hiring (or 'borrowing') an experienced facilitator from outside your organization.

The Steps

As mentioned above, the process generally follows a similar pattern but should be adapted before or during the event if necessary:

Introductions

Quick introductions of everyone in the room.

High-level description of the goal

The facilitator outlines what the group is trying to achieve and, in quick detail, describes the Workflow Sprint process. A strategic stakeholder may be introduced to give a 5-minute validation of why the group is gathered.

High-level discussion

It is always necessary to start with a very high-level discussion about the workflow and what the group is trying to achieve together. In general, the Workflow Sprint process starts with a very wide (sometimes meandering) conversation and very quickly begins to focus on specifics.

Documentation of the existing workflow

If the group is big, the facilitator may break it down into smaller groups to drill down into specific parts of the workflow. Each group will then present their documentation back to the larger group. At the completion of this process workflow should be documented in its entirety in front of the entire group (using whiteboards or large pieces of paper or shared virtual whiteboards if working remotely). Documenting with everyone present helps validate the workflow and ensures any vagueness or 'gotchas' are caught.

Open Discussion

A general discussion on what people think about the existing workflow. It is good to allow this to slip into a little brainstorming in anticipation of the next step.

Designing an optimal workflow

The facilitator generally breaks the team into smaller groups. Each group chooses some part of the workflow (or the entire workflow) they wish to re-imagine. Documentation is often captured with drawings (there is no preferred format - let each group decide). Each smaller group's work is presented back to

the entire group and discussed. From this, the facilitator can document (also in front of the entire team) a proposed optimized workflow.

Report

Generally, the facilitator will take all the documentation and distill it into a short report. Documentation of the outcomes usually uses workflow notation (see below) as well as pictures of all the artifacts created during the event. The facilitator generally adds further information to tease out some questions for the organization to consider (contradictions/compromises etc).

Stage 6 : Designing an optimal workflow is the most important step. You will find **Step 4 : Documentation of the existing workflow**, if done well, greatly aids the group to think about how to solve some of the problems and create further efficiencies.

Workflow Notation

Steps 4 and 6 above require us to document the workflow. The most effective way to describe workflow is through a numbered list of who and what for each step. I call this WorkFlow MarkUp (WFMU - actually WFMU isn't really a Mark Up Language in the formal sense, but WFMU is one of my favorite radio stations and I just couldn't resist - hi to Dave Mandl if you are out there...).

There is no need to have complex flow diagrams and extended narratives - most of the time, these do not add clarity and sometimes mystify.

The three elements of WFMU are:

1. **Who** - usually the role of a team member or group (eg author, or editor). We use the naming convention of the host organization for the role names. In some cases, the 'who' is an automation 'role' and refers to a machine (computer/s).
2. **What** - usually a single verb that describes what actually happens (eg review', or 'write' etc).

The third part of documenting each step - **When** - is usually understood simply by the position in the numbered list.

An entire publishing process can be described in WFMU. Here is a short excerpt from a Workflow Sprint I facilitated with the Organization of Human Brain Mapping (note this article uses very simple illustrative examples):

1. **Author** creates new Submission
2. **Author** completes Submission

3. **Author** Submits
4. **Editor** reviews Submission
5. **Editor** invites Reviewers
6. **Reviewers** accept the invitation (or reject)
7. **Reviewers** write reviews
8. **Reviewers** submit reviews
9. **Editor** reads reviews
10. **Editor** writes decision 1. If accept → publish (go to 15) 2. If revise → Author revises, return to (5)

Documenting workflow like this is straightforward to create and easy to read.

WFMU enables us to easily identify where problems exist. For example, if we saw this:

4. **Editor** reviews Submission
5. **Editor** invites Reviewers
6. **Reviewers** accept the invitations (or reject)
7. **Editor** considers a possible desk rejection

We can easily see a problem – we have just spent a lot of time inviting reviewers for a submission (Step 5) that might be rejected (Step 7) before reviews are written. These kinds of problems are not uncommon and are clearly exposed with WFMU.

I use WFMU to describe what currently happens in a workflow, to help identify problems, and to document optimized workflows.

Conditional Logic

As you can see from the above example, we can easily capture conditional logic by using nested bullets and conditional statements eg:

10. Editor writes decision
 1. If accept → publish (go to 15)
 2. If revise → Author revises, return to (5)

Documenting Automation

Automation can also be described using WFMU. For example, the following is from a Workflow Sprint with the NIH:

1. Author Submits
2. System checks document validity
 1. If all passes -> create workspace
 2. If fail -> generate report for Author

Note: When designing new or optimized workflows that involve automation we must avoid magical thinking. To this end it is good to have trusted technicians in the room to help the group understand the promise of technology.

What are the outcomes of a Workflow Sprint?

There are two direct outcomes of a Workflow Sprint:

1. Documentation of the existing workflow - A description of your workflow as it exists now
2. Documentation of the optimized workflow - A description of the ideal, optimized, workflow

In addition, depending on how much your facilitator understands about publishing, they may wish to generate a third outcome: A report - A report with the above two items plus narrative about the realities of achieving the optimized workflow and any other issues that may need to be considered.

From this report, the organization should be in a very good position to consider technology options. A publisher can now do a gap analysis to evaluate existing solutions or start the design process to build new technology.

Note - as mentioned above the WFMU examples are very simple workflows. I chose these deliberately to help illustrate how to use the notation. Often the workflows documented with WFMU are far more complex.

Where did Workflow Sprints Come From?

The Workflow Sprint methodology was not born in a vacuum. I designed the process by adapting a previous method I had designed – Book Sprintstm (photos below).



A Book Sprint (<http://www.booksprints.net>) is a facilitated process that leads a group through the process of producing - including conceptualizing, structuring, writing, editing, illustrating, designing and typesetting - a book, from zero, in five days or less.

In the words of the Book Sprint website:

"A Book Sprint is a unique experience, relying on maximum commitment from the participants, a dedicated working space, and strong facilitation. All this and a non-negotiable deadline."



Book Sprints is a 'radical optimization' of the book production process, and the method has been applied to a huge array of use cases for a broad variety of clients, including Cisco, the World Bank, the United Nations, Netflix, VMWare, RedHat, Department of Commerce (USA), California Digital Library, University of Hawai'i, Dell, Haufe, the Government of Burundi, and others.



Following the success of Book Sprints, I designed Workflow Sprints with a similar aim – to radically optimize a legacy slow and expensive process through facilitation.

Summary

Understanding and optimizing workflow is critical before making technology decisions. I have seen many organizations progress down the path of technology-first solution seeking without first analyzing workflow.

Understanding and optimizing workflow is critical to do first and it does not have to be a long, labored, and expensive process. Workflow Sprints is presented here as an example methodology that optimizes for speed and efficiency - producing clear documents quickly, in days, not months (as per legacy processes). If you are unsure about the process I recommend trying it as it is so quick that you could front-load a sprint before pursuing more traditional methods. There is very little downside. If nothing else everyone that participates will learn a great deal from other team members about how things actually work and how their work affects others downstream. This, together with team building, is a good reason in itself for trialing a Workflow Sprint.

But my experience is that this process greatly assists organizations thinking about changing their technology and/or improving their publishing processes.

I designed Workflow Sprints to be lightweight and effective. There is a lot more I could write here about how to get the most out of the process but I hope the above gives you some ideas about how to reduce the burden of this kind of analysis. There are also some deeper items I haven't addressed in this article. After reading a draft of this article Ken Brooks (<https://www.linkedin.com/in/kmbrooks/>) suggested that the very scope of what is considered to be 'the workflow' needs to be addressed up front. He is

absolutely right. If a publisher breaks down the workflow into disconnected components and focuses on optimizing each separately - editorial, publication and digital workflows for example - then you have embedded many assumptions into the process already. These assumptions might be the very conditions that prevent higher levels of optimisation. The solution would be to consider the entire end-to-end publishing path as the subject for a Workflow Sprint (or similar process). For further thoughts on this perhaps consider reading the article on Single Source Publishing I wrote last year.

As you have probably guessed I also facilitate Workflow Sprints so if you need a facilitator, then give me a call! If you want to try Workflow Sprints without me, please go for it! I am happy to have a remote coffee and discuss the process if it will be of help to you. Many thanks for taking the time to read this short article!

Adam Hyde
adam@coko.foundation

Some definitions

As mentioned above, by workflow I mean the complete set and order of actions an organization (usually a publisher) undertakes to prepare an object (book, journal article, preprint review, question set etc) for publishing or sharing.

For example, workflow is everything that is done from start to finish that falls within a publisher's control, to prepare a book or journal article for publication. Or, every action required by a community to prepare and share reviews on preprints.

'Actions' that occur in this process are events performed by both : humans (e.g. reviewing, formatting, Quality Assurance, copy editing etc) computers (e.g. semantic analysis, format conversion etc).

I use the term 'object' to refer to a generic publishable entity e.g. a book, or journal article, dataset, test question or questions, or preprint review.

Lastly, I refer to any organization that prepares information to publish or share as a 'publisher'. I think many of the folks I work with do not see themselves as (capital P) Publishers. They may instead (for example) identify as self-organized communities that share information (like the preprint communities mentioned above). However, for simplicity's sake, I refer to the wide spectrum of various organizational models that prepare objects to publish or share as (small p) publishers.

Credits

Many thanks to Raewyn Whyte for copy editing. Also thanks to Ken Brooks for some great feedback. Thanks to Henrik van Leeuwen for the illustrations.