8 Managing the modelling

INTRODUCTION

In Chapter 2 we looked at the business of modelling a single process from the 'technical' point of view. But the process modelling activity itself needs to be managed if it is to be successful. This chapter gives guidance on actually doing the process modelling work to ensure that it gets to answers as efficiently as possible. We look at modelling from a procedural point of view (i.e. as a process itself), and, to a lesser extent, from the 'soft' point of view taking into account some of the sociological and political and people issues. We shall step through a basic process modelling procedure – in particular, a modelling workshop – addressing the soft issues as we go. We'll take a generic view of the modelling process for now, leaving it to subsequent chapters to describe how to adapt it for specific situations: discovery, definition, design, diagnosis and enactment. The 4Ds and an E.

There can be no hard and fast rules for the procedure we should follow. It can be affected as much by political and logistical issues as by technical modelling issues. However, here's the basic scheme:

1 Decide on the objectives of the modelling.
2 Brief ourselves by getting an overall picture, no matter how coarse, from a variety of sources.
3 Run one or more interactive workshops of those involved to draw up a RAD that meets the objectives of the modelling.
4 Use other appropriate sources of information.
5 Review, revise and validate the model using other inputs.
6 Use the model.

No surprises there. Remember that we have already prepared our process architecture.

Before we look at each of these steps in detail, there is one topic that we have touched on already but which we need to consolidate: the difference between abstract and concrete process models.

Abstract and concrete models

There is an important choice that we must make at the start of any session: do we need an abstract model or a concrete model, or one that is a mix? When we examined the different process concepts in Chapter 2 – role, action and interaction – we saw how we could choose to model abstractly (concentrating on intent) or concretely (concentrating on mechanism), or

both. Let's remind ourselves of the options and then look at two models of the same process: one abstract and the other concrete.

 If we are modelling how a process operates *now*, we might wish to capture, for instance, the division of labour between people and computer systems. Or we might decide that it is not important initially whether an action involves people or machines, and decide instead to capture the 'essential' content of the process, not worrying how it is manifested physically. Yet again, we might decide that in future we want a particular action to be done by a person, or automated and given to a machine to do, or done by a person with the aid of a machine (such as a database system), and we might choose to model this.

For example, we might prepare a model of how you get a purchase order form from someone who keeps a stock of blanks, fill it in with the details of something you want to buy, pass it to me, get me to authorise the purchase by signing the form, and return it to you. We might represent this process fragment as some actions in our respective roles and interactions between them. Those interactions have a purchase order form as the gram passing between us.

The abstract process might simply be 'You ask me to authorise a purchase,' a single interaction without any gram. What we capture is not the mechanism but the intent. Whether we choose to model a process concretely or abstractly depends on what we are trying to achieve by modelling.

If we are involved in process improvement or re-engineering we might find it useful to

☞ draw up the current concrete process model;

☞ 'abstract' it to yield the current abstract model;

☞ find a better way of implementing it in a new concrete model.

Remember that we are not thinking here of preparing detailed as-is concrete models, and then deriving complete abstract models, and then complete to-be concrete models. The suggestion is only that by 'moving' in some sense from the concrete to the abstract and back to the concrete we will gain insights. That could mean simply thinking it through in rough sketches or on white-boards.

This route also makes sense when we are modelling a process as a prelude to providing computer support to some or all of it.

Concrete vs abstract actions

When we draw an action on a concrete RAD we will want to show how the action is done; we shall want to talk about the mechanism. In an abstract RAD, on the other hand, we shall talk about the intent or purpose of an action. So, in Figure 8-1, we see a very concrete action – *Complete form 21b* – whose name is explicit enough to be a work instruction, but gives us little idea as to what is achieved by taking this action except that a form has been filled in. The figure also shows the action expressed in abstract terms – *Prepare sales analysis for the month* – which tells us what we want to achieve but not how we do it. Of course, we are quite at liberty in our models to give an action a name that tells us both things – *Prepare sales analysis for the month by completing form 21b* – but it is as well to be aware of the two contrasting styles.
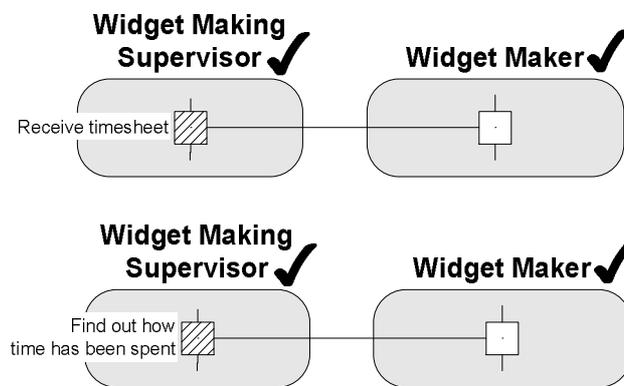
Figure 8-1 – A concrete action and its abstract counterpart

Concrete vs abstract interactions

Precisely the same applies naturally to interactions, as illustrated in Figure 8-2 (which we saw earlier in Chapter 2).

Figure 8-2 – A concrete interaction and its abstract counterpart

When we model interactions concretely we shall see verbs like 'send', 'sign', 'copy', 'pass', 'receive', 'get' – mechanisms. When we model abstractly we will use words like 'request', 'delegate', 'authorise', 'approve', 'report' and 'agree' – intents.

Concrete vs abstract events

The differentiation between mechanism and intent carries naturally into events. We might label an event as *Form CC received from customer* or as *Customer makes a claim*.

Concrete vs abstract roles

Finally, we can do the same thing for roles, something we discussed a lot in Chapters 1 and 2. When we take a concrete view of a process, we shall be very likely to choose as our roles things recognisable in the organisation: posts, departments and computer systems. When we take an abstract view we shall be more likely to identify areas of responsibility without reference to the way they are allocated to organisational entities.

So, in the top part of Figure 8-3, we talk in terms of areas of responsibility called *Project Managing* and *Purchase Approval*. In the bottom part, where we are taking a more concrete view, our concrete model shows the job title of *Project Manager* and the post of *Finance Director*.

Figure 8-3 – Abstract roles and their concrete counterparts

Project Managing

Purchase Approval

Obtain purchase approval

Project Manager

Finance Director

Obtain purchase approval
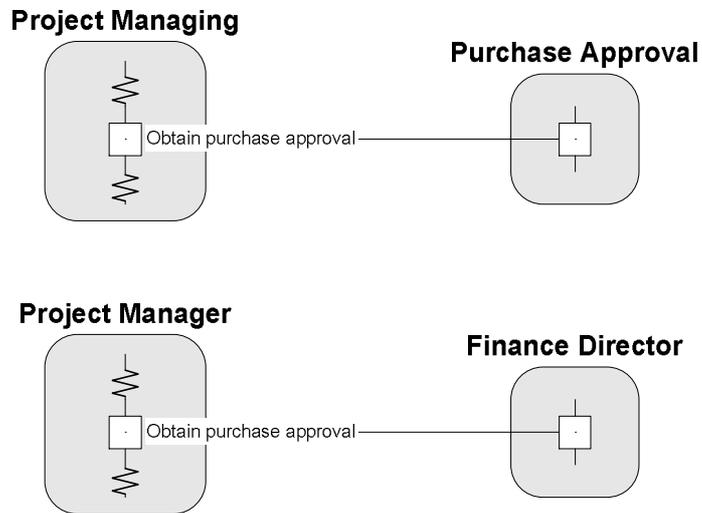
An example

In Figure 8-4 we have a process described in very concrete terms. Real posts and departments appear as roles; actions involve physical objects and physical actions on them; interactions are expressed in terms of the paperwork involved; and there is even a computer system taking part as a role. This would make an excellent RAD for telling people exactly what to do. But it would be of less use for getting an overall picture of why things are being done.
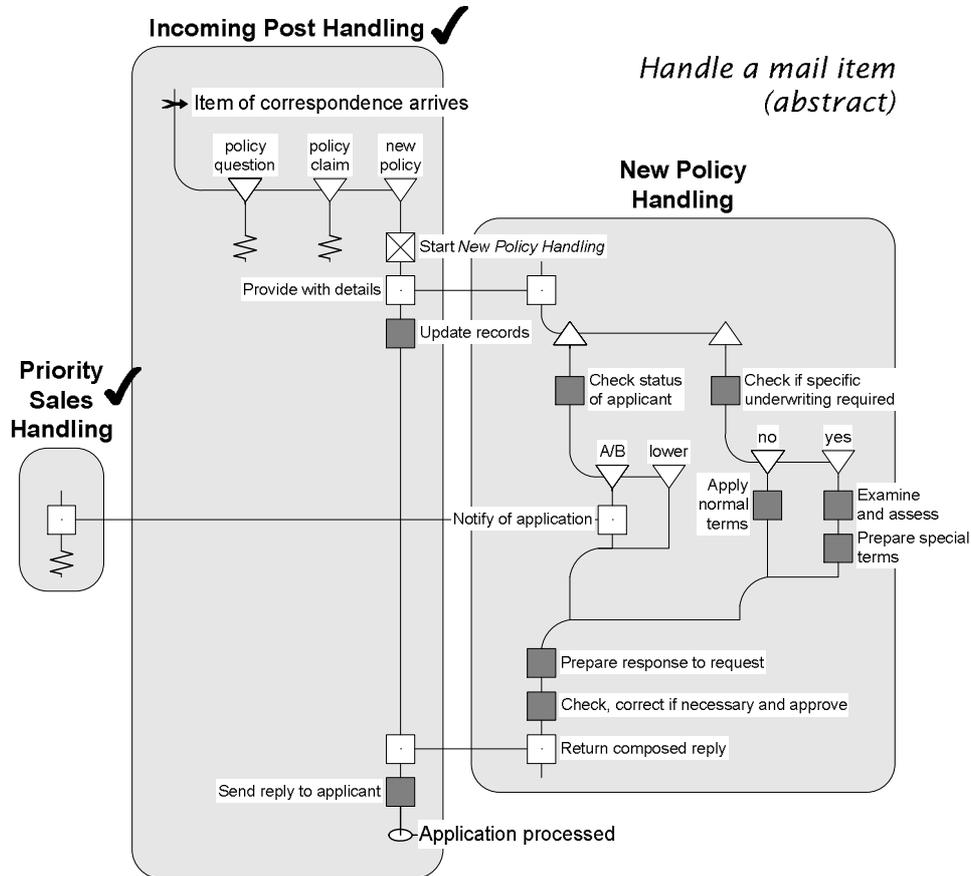
Figure 8-4 – A concrete model of the handling of a mail item
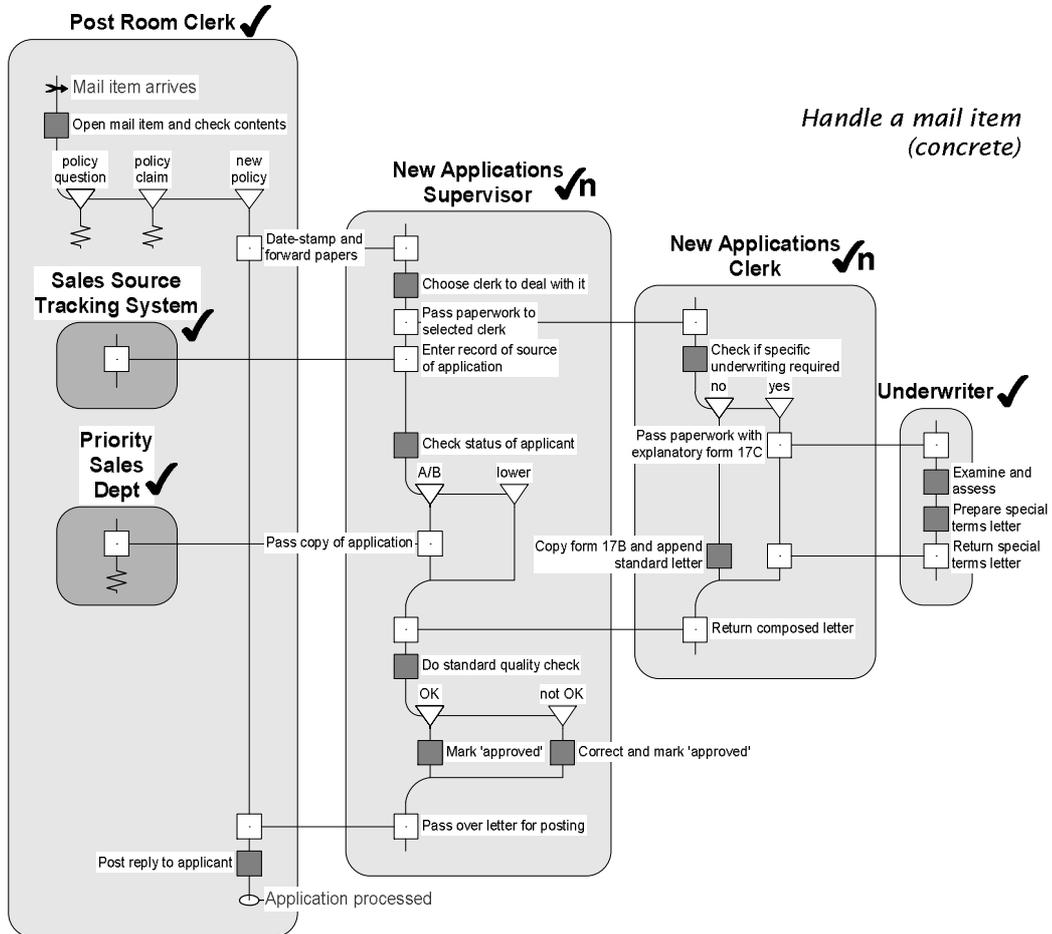
Suppose now that we rework this model, trying to move towards a more abstract representation of the process. We might draw something similar to Figure 8-5.

Figure 8-5 – An abstract model of the handling of a mail item

Incoming Post Handling ✔

Handle a mail item (abstract)

Item of correspondence arrives

policy question

policy claim

new policy

New Policy Handling

Start *New Policy Handling*

Provide with details

Update records

Check status of applicant

Check if specific underwriting required

Priority Sales Handling ✔

no

yes

A/B

lower

Apply normal terms

Examine and assess

Prepare special terms

Notify of application

Prepare response to request

Check, correct if necessary and approve

Return composed reply

Send reply to applicant

Application processed

Tutor: What do you notice about the change?

Pupil: Well, the abstract version is much simpler. We've removed a lot of stuff that's shown in Figure 8-4 only because of the way they've chosen to implement the process: posts, computer systems, communications mechanisms, paper flow etc.

Tutor: Right. The abstract model is getting at the *essence* of the process, at what we are trying to do.

Pupil: Would it be fair to say that it's therefore a better model, because it's simpler?

Tutor: NO! It's vital to remember that abstraction isn't some sort of *summarising*. Our aim is to model the process in terms of intent or purpose, rather than mechanism. A side-effect might well be that the model is pictorially simpler, because we will ignore 'implementation detail'. But a simpler model is not necessarily a better model. Hiding implementation detail might be the last

thing we want to do. It might be precisely our aim to demonstrate that a process is a shambles of crazy activity! We might model all that madness and end up with a crazy RAD – and that will be exactly what we want management to see to prove to them that they need to do something about it! Showing them a tidy model will not scare them into action. Our second model might be an abstract version, as the first step to simplifying the way the process is carried out. But then we have a different purpose in mind for the model.

KEY POINTS

Before modelling we must decide if we are working with an abstract or a concrete model.

An abstract process model shows intent.

A concrete process model shows mechanisms.

A model can be a mix of abstract and concrete perspectives.

The type of model must be determined by the purpose of the model.

A simple model is not necessarily a better model.

If it's a muddle, we might need to model the muddle.

STEP 1: DECIDE ON THE OBJECTIVES OF THE MODELLING

The importance of this step can never be over-emphasised. Over the years, process modelling has acquired a bad name. This is because all too frequently it has been done badly, has lost the plot, wasted everyone's time and money, and yielded little except a doorstop of pictures. Part of the problem has been the poor methods that have been used: methods that rely on decomposition (which all too easily results in unstoppable modelling), or that provide no clear direction (with random models as a result). The other part of the problem has been that modellers frequently forget why they are doing it, and get stuck in a belief that they are working on *the* model – whereas of course there is no single model, only the model you find useful.

Once we have answered the question 'Why are we doing this?', I suggest we have it printed in red on a large sheet of paper that is displayed whenever and wherever the process is being thought about. We must avoid analysis paralysis. Then, to answer questions such as 'Do we need more detail?', 'Where should we stop?', 'Do we need to cover such-and-such?', just look at that large sheet of paper and rephrase the question: 'Have we enough detail to do what we are trying to do, to answer the questions we have posed, to achieve the goals we have set?'

If we do not know why we are modelling and do not have a clear idea of the outcome we are looking for or what we want to be able to do with the model, the modelling activity will be slow and undirected at best, and at worst will fail. We are more than likely to end up modelling aimlessly. We have already seen that there is no single model of a process, so when we start to model we need to know what perspectives we should be taking.

Early in this book we identified a number of different reasons for drawing a model of a process. In subsequent chapters, we shall look at these different aims:

☞ To discover a process or to define it to the degree we want to enforce it – Chapter 9.

☞ To map the as-is process for diagnosis and to improve it – Chapter 10.

☞ To design a brand new process from scratch – Chapter 11.

☞ To provide a basis for the design of an information system – Chapter 12.

☞ To build an enactable model – Chapter 13.

STEP 2: GET AN OVERALL PICTURE

Now that we know where we are headed, our aim in the second step is to map out the ground, in particular to identify the boundaries of the process model, and to identify the perspective or perspectives that will help us get to answers. Our aim is to get ourselves briefed before the group workshops and individual interviews, so that we have a good idea of what we can expect to hear: the roles we can expect to hear about, where difficulties and tensions lie, where the process starts and finishes, and so on.

What is our starting point?

Our first input is of course the process architecture. We might have produced it as part of step 1, when we fixed on the organisation that we are interested in, or it might be in place already. Either way, it will have chunked the activity of the organisation that we are looking at, and we shall have chosen one or more processes that we want to capture, or design, or diagnose, or improve, or enact.

The architecture gives us first-cut answers to the following questions. Write those answers down and publish them for the process you are dealing with.

☞ What is the 'organisation' that we are concerned with? The process architecture defines the organisation in terms of the units of work it deals with, of course, but let's come up with a simple characterisation such as 'All of the Marketing Department's activity, excluding TV campaigns,' or 'The interface between clinicians and the pharmacy,' or 'The work of the group that deals with vehicles abandoned in the street.'

☞ Are we looking at a case process, a case management process, or a case strategy process?

☞ How is the process activated or triggered?

☞ What relationships does it have with other processes? In particular, does it generate new cases of any units of work itself?

Roughly what happens?

The process architecture only draws a boundary round a process. It tells us nothing about what is inside it, except that we know the sorts of things that will appear in a case management process, for instance. Our next step is to get an appreciation of the contents of our process. My experience is that this can be done efficiently by interviewing someone who has a good grasp of the whole process even though they might only operate a part of it. They might not be a

senior person in terms of rank, but they might well be senior in experience in that organisation, having worked in many parts of it, for instance. The essential thing is that they should have a broad view. If an organisation has previously done some sort of investigation into how things work (or don't work), the person who led that could be a useful starting point. This work can take a couple of days of intensive discussion, using lots of informal pictures.

At some point there might be enough information about the process to take a first cut at a RAD. But remember that this is a private RAD: we are sketching it as part of our briefing prior to getting close to the action.

What roles can we expect?

Early on, we can start to list the roles that we are likely to come across as we get into the modelling sessions. Natural sources are

- ☞ posts from the relevant parts of the organisation chart,
- ☞ job titles,
- ☞ departments, branches, etc,
- ☞ roles identified in existing written procedures,
- ☞ committees, task forces, working groups, etc,
- ☞ the customer(s) of the process, and perhaps their customers,
- ☞ suppliers to the process,
- ☞ regular meetings.

Using this list we also need to decide whom to invite to be sure of having the right people in the room to answer the question we have set. I have a subversive suggestion at this point; unless they have a role to play in the process, do not invite managers … more of this later.

What terminology is used?

Now is the right time to start a glossary of terms. If, as analysts, we are new to the business we are modelling, we have what is both an advantage and a disadvantage: we don't know the terminology of that business. This is a disadvantage because at first we shall be slow to understand what people are saying when they use terms that have a special meaning for them. This is an advantage too because, in trying to find out exactly what they mean by those terms, we shall start to uncover ambiguities in the organisation's views about itself: two groups might use the same word but with different meanings. Such ambiguities can be indicators of misunderstandings or even conflicts.

What do senior people think?

Working our way down through the organisation chart will generally be the politically correct approach: involve department heads early on. We should find out whom they regard as authoritative about what goes on in their groups, get their commitment to the use of their staff's time, manage their expectations about possible outcomes, etc. What is politically correct and what is necessary to get the facts will determine what order we speak to whom. This is all bound up with the way in which the larger project – for radical change, or incremental change, or the introduction of a Quality Management System, or whatever – is itself being handled.

Most such initiatives require senior level backing and sponsorship for success, a topic that is adequately covered in the literature on these larger topics and one which I do not address here. But we should be aware of that problem.

STEP 3: RUN ONE OR MORE INTERACTIVE WORKSHOPS

Having been briefed in the process we're interested in and having cleared the way with the appropriate senior people, we now reach perhaps the most important step in the process modelling project: the group session in which we will model the process that we are working on.

But let's first take a small digression here to worry about whether group sessions are a good idea at all. If you are an ethno-methodologist you will believe that the one true way of finding out how an existing process works is to sit and watch it, perhaps even take part in it. We must ask however whether, given the questions we are trying to answer and the challenges we are facing, we are actually interested in pure discovery. If our aim is to design an improved process, we might not be in the slightest bit interested in capturing how things are done today: we know the process is wrong – why draw a picture of it? If our aim is to define a process so that we can standardise it – 'this is the way we do things round here' – then we won't want to capture all the ad hoccery of actual practice: we want to produce a work instruction, not a depiction of reality. So, whilst we can appreciate the different dimension that ethno-methodology can bring to the depiction of existing operational processes, for the most part that is not our interest.

Our interest *is* in getting a group to engage with their process – for definition, diagnosis, improvement or enactment. If we are not designing a process in a green-field situation, in other words if there is an existing process out there, we shall aim to get in one room a representative from each of the roles that we have identified during our briefing. It may be the first time those people have been in a room together, so ingrained might be the functional silos of the organisation. For the first time they might be seeing how their work fits with the work of others, and facing for the first time, from a general perspective, the areas of conflict or stress between groups in the process.

The output of the session will be a RAD. But here is an important point: we shall not introduce the group to the RAD notation. We use it almost without mention, drawing the process as we go, using the symbols we have become familiar with. Our job, as the person at the front holding the marker pen, it to translate what we hear into a picture – a RAD. Each time we use a new symbol – the first part refinement, say – we shall simply make a point just to say 'When we have a number of parallel threads starting I'll draw this …' My experience is that any group readily takes to the notation, provided that the analyst actually does the initial drafting for them. It's sufficiently transparent for people to work on their *process* during a modelling session, rather than to work on (struggle with) the notation.

Starting the workshop: laying out the ground

This is not a book on facilitating a workshop. Much of what follows is just a *Riva* slant on everyday facilitation. However, I like to start by agreeing the timetable: 'Here we are. It's 9.30. We will finish by 12.30. If you commit to giving me your undivided attention during the next three hours, I commit to finishing this workshop no later that 12.30.'

When people enter the room, make sure they find on the wall that large sheet of paper with the purpose of the session on it. Spend some time getting their focus on it.

Next, spend time agreeing the bounds of the process to be modelled. From our briefing work, we already have a good idea of what to expect, of course. But besides scoping people's expectation and thinking, this exercise will help each individual warm to the area and the challenge written up there in red. Questions that help this exploration include:

☞ **What's the process called?**

This may sound a stupid question, but it often reveals different perspectives or agendas in the group. To get people to think I replay their name for it:

What's it called?

It's called Accept delivery of raw materials

I visualise some raw materials that have just arrived and you are going to accept them – yes?

… No, we don't always accept them: sometimes we reject them. We're really checking them to see if they can be accepted.

OK. I visualise some raw materials that have just arrived and you are going to check them to see if they can be approved. So perhaps it should be called Check delivery of raw material?

… In fact, we check who they've come from: if they come from an approved supplier we accept them immediately; otherwise they are turned away.

And so on.

Our aim should be to bring the group towards the name that we were led to by the process architecture, **Handle a delivery of raw materials**, for instance. If things won't go in that direction then perhaps there is a message there: that the architecture is faulty in some way, or we have the wrong group here, or there are such different viewpoints in the room that we need to back off and deal with the discrepancy.

☞ **For a case process, what does this process deal with? What does it handle? Or what does it produce?**

In *Riva* terms, what's the unit of work? We must avoid getting into the input/output style of thinking, of course: this doesn't help in scoping the process and can too easily have people focusing on the way the process works.

☞ **How does the process start?**

In *Riva* terms, what events trigger action? A case process will generally only yield one trigger, but the group might identify triggers for interactions with other processes: **Make a product batch** has to respond to enquiries about progress from management for instance. A case management process should generate plenty of triggers, as we saw in Chapter 5.

☞ **How do we know when the process finishes (if it ever does)? What are the goals of the process? What are the possible outcomes of the process?**

Groups often need prodding to recognise that a process can have more than a successful outcome: rejection, or failure, or handing off, or escalation can all be alternatives. We must

rephrase the question in many ways to winkle them out: *How would I know it has finished?*

What's the last thing that happens?

Can the process fail in some way?

If that's success, what's failure?

How many different ways can it fail?

☞ Who are the people or groups involved?

What roles will we expect to see appearing in the model? People will shout out all the job titles and posts that they can think of. At the edge of the process, they will identify larger groups, perhaps entire departments that they have connections with. Write them all up. There may be dozens – I have seen forty. Keep writing. We will probably find that only a handful actually appear on the final RAD, but the list will prove a useful memory jogger.

Identifying all the roles is not always as easy as it might seem. People think naturally in terms of departments and named individuals and this is a perfectly good starting point. In his work with a leasing company, my colleague Tim Huckvale worked with a group who identified Kate as the person who did such and such – Kate had always done that. Kate's name went on the flipchart.

☞ What areas do we want to ignore today? What areas are definitely in the discussion?

The process architecture has chunked the organisational activity, so we might know that the processes for certain units of work are being dealt with in separate workshops and so we only need to go as far with this process as the boundary concerned. If we are dealing with a case process, we know that case management is outside our work. If we are dealing with a case management process, we can ignore how individual cases are dealt with. And so on.

Our prior briefing has told us what to expect to hear and where to probe. We shall put all the answers to the questions onto flipchart sheets and stick them on the walls as reminders whilst we work on the process itself.

For the actual modelling, a large whiteboard is essential. Preferably two … or three. We need a large area to work in and we will be drafting, correcting, changing, and rubbing out a great deal – space and flexibility are key. Do not use flipcharts for the RAD – you cannot rub things off. Above all, do not attempt to use a specialised drawing tool on a PC. It is very distracting, and very inflexible. People must be focussed on the process. Icons, and buttons, and windows, and all the flummery of a software application simply divert attention from the real job. If you have a printing whiteboard or a digital camera, so much the better. If the whiteboard is a projected virtual one that you can sketch freehand on directly into a computer, better still. But please leave your favourite drawing tool at home! If we find we have to sit down at the end of the session and copy the whiteboard contents onto paper or into a laptop, then that is a small price to pay for the complete concentration of the group. Technology distracts.

Doing the modelling

Now comes the moment when the modelling has to start.

If the process has a natural start-to-end flow about it – perhaps it is a case process – we can take advantage of this: we draw the main triggering event for the case at the top left-hand corner of the white-board and put it into the role where the process starts. We find out which role that is by asking 'Who notices this?', 'Who first gets to hear that something is needed?', 'Who deals with it in the first place?' We look at our brainstormed list of candidate roles for inspiration. Not surprisingly, this first step can take a worrying amount of time: all the questions about whether this is a concrete or an abstract model, whether we are modelling organisationally or in terms of responsibilities, whether the process starts here or before or after, etc have to be answered. We shall be patient and work through this. These first decisions will set the tone for the entire session, so we must be happy that we have these decisions right before we move on and make more. We should expect to go into the room knowing the appropriate answers, of course – that is our responsibility as leader.

We start by drawing the goal (i.e. the desired outcome) of the process as a state somewhere in some role at the bottom right of the whiteboard. The rest of the workshop is now about filling in the process between these two points. Roles appear as they enter the process, spreading across to the right. Don't start by writing all their names along the top of the board – we shall only use a fraction of the ones that were identified earlier.

If the process does not have a neat flow, things are less easy, and it becomes necessary to do the same thing for each of the threads that exist.

Either way, as the RAD develops, we shall draw it and redraw it many times as the group explores the process. Few will have thought about their daily lives in this way, so there is an (enjoyable) element of exploration for them. There will be problems that we have to solve as the modelling proceeds. Just how much detail do we want to get into? Shall we ignore that role's contribution at this stage? Is it sufficient to summarise that set of interactions as just one for now? Do we want to separate those two roles or treat them as one at the moment? Shall we collapse all that activity into just one black box? Should we regard the work going on in that other process as outside our boundary and simply capture it as an external event or two? There are, as ever, no stock answers to these questions. *It all depends why you are modelling: look up at the words in red for an answer.* That said, we must get to the end of the process in the available time: we made that commitment. Going away with at least a rounded – if not 'complete' – model has a value: we can always explore further detail in later sessions either with the same group or with smaller groups and individuals.

Given that a process will generally have many threads and that they cannot all be explored simultaneously, we shall need to be careful to note where threads still have to be closed off. And it is useful for the group to know that this care is needed so that they are encouraged to point out unfinished threads that will have to be returned to at some time. Simply drawing a 'spring' at the appropriate point on the RAD is usually enough to indicate 'unfinished business'.

Should we look at the 'normal' situation first and then come back and add the exception condition handling and abnormal situations later, or should we try and deal with them all on a single sweep? There is some virtue in the first approach in order to build a framework on which everything can be hung. The danger is of course that exception and abnormal situation handling can easily be forgotten if put off 'until later', and it is often those parts of the process

that reveal areas for improvement or suggest the possible use of computer systems to reduce the likelihood of error.

Work-arounds – additions to the 'approved' process in order to make it work – are a fruitful source of ideas about what is going wrong and what could be done to remedy things. It is not unusual to hear something like 'Then I go and get approval from the monthly Management Meeting … well – I say that – in fact sometimes I can't wait that long, so I check it out with the chairman Brian and then get it rubber-stamped at the next Management Meeting, otherwise nothing would get done.' That work-around is a clue. Indeed, we shall need to probe for the existence of such work-arounds:

> What do you do if you don't get the stuff in time?
>
> Do you ever get on with that even though you haven't had authorisation?
>
> How long do you wait for that?

As the modelling proceeds, some of the roles brainstormed earlier will find their way onto the RAD. People will start to abstract away from named individuals – 'Well, Mary does do that, but she's signing it off in her capacity as Site Safety Officer' – and the roles appear. Quite often roles that are not on the 'main' stream of the process are missed initially. They might only be involved for a single interaction, but of course it can be that interaction that holds things up, simply because it is some form of approval that is required from an 'outside' role: *Get Health & Safety to sign off the risk management plan*, *Get the plan signed off by QA*, *Get Finance to agree to the budget*. Equally, we must be prepared to strip out roles that don't materially contribute to the process or its understanding as the modelling proceeds.

This is part of the trick of knowing what to put in and what not to put in – it's a modelling decision. The boundary needs constant validation:

> We've got this role in here – do we care for now?
>
> Is it worth looking at what happens before this trigger or not?
>
> Is this really the goal, or is there actually something earlier/later that we are really interested in?

and so on.

Standing, pen in hand, at the whiteboard, we have an important task: eliciting the process from the group, getting it onto the whiteboard, allowing the group to own what is drawn and to buy into what is drawn, steering the modelling, and bringing it to a conclusion. The result of their work is that RAD on the whiteboard. But many other important items of information and clues will have cropped up during the discussion and debate, things that are spoken or just hinted at, and for this reason we shall find it very useful to have another person simply keeping a record of things such as:

☞ avenues that were explored but backtracked from;

☞ any decisions to ignore certain detail that would need to be picked up later;

☞ concerns about the way the process works or doesn't work currently;

☞ suggestions as to how it could be improved;

☞ situations where errors frequently occur;

☞ points of stress in the process;

☞ judgements about the relationship with other processes/departments and their effect on the process under discussion.

People will say a lot of things during the modelling that could act as pointers to inefficiencies, problems and solutions. We must record these for later analysis. Our note-taker, perhaps more than anyone standing at the whiteboard, has to be sensitive to what is said:

> When these forms arrive, the applicant's policy number is rarely filled in and we end up having to go back to the originator to get the information.

> We generally don't have enough time to handle that fully and we only go back to it when we get a quiet period.

> Couldn't that be sorted out at the weekly meeting rather than waiting for the next management review?

All these signals need to be noted for future use if we are not going to explore them there and then.

Our aim is that, by the end of the session, the process on the whiteboard is *their* process, and the model is *their* model – they have after all drawn it, albeit with the help of the analyst who held the pen. This element of ownership is as ever vital for subsequent work.

Closing the workshop

Having made that commitment at the start of the session to finish on time, we shall stick to it. Our aim – whatever happens – is to get to the end of the process, even if we have to gloss over some areas, recognise that we have not adequately dealt with others, and so on. We must cover the *entire* process to some degree. We can always come back and go over it again. But we cannot put it down half way through and expect to pick the threads up again a week later – not least because we probably won't have the same people in the room.

So, remembering our deadline, we shall also aim to leave enough time at the end of the allotted time to revisit all of the brainstormed triggers, outcomes and roles that we put up on the flipchart. Did we cover them all? If not, why not? Are we clear why not and happy with that? Although we might well have started with twenty or thirty candidate roles, we'll probably find that we only have half a dozen on the RAD – this is quite typical. For instance, the brainstormed list probably had all the individual job-titles in the Accounts Department, but – for this model – we only needed to show the role *Accounts*.

The clock says 12.25 and we wind up with a final question: 'We are going to take this information away, draw it up tidily, and circulate it to you. Do you think we have finished our discovery/diagnosis/design today?' The group needs to decide if its work is done. We should always aim to get them back at least once for a review session. One workshop is not enough. In some cases, four or five have been necessary for a group to get to a final process design that it is happy with, or that meets the design goals set for it, or that it thinks it can operate.

As our group disperses, we gather up all the material that has emerged and transfer the pictures and all the other information to a more portable form. We now have to organise what

we have heard. Our first task will be to draw up the RADs 'properly' – at last we can use our favourite drawing tool. This inevitably reveals unfinished threads, missing detail, doubts, misgivings, questions about terminology, and so on. We collate these and add all the signals and messages the note-taker heard during the session. We will use these to guide further information gathering and as input to the review workshop to follow and subsequent workshops that continue the capture, design, or diagnosis that we are doing.

KEY POINTS

A good basic procedure for a group modelling workshop is:

1   Arrange to involve representatives from all the likely roles.

2   Ensure the workshop room has plenty of drawing space: whiteboards and flip-charts. And working pens!

3   Get the group together for an agreed period of time.

4   Put the purpose of the overall exercise on the wall.

5   Brainstorm triggers, goals and outcomes onto flipchart sheets.

6   Brainstorm an initial list of roles onto flipchart sheets.

7   Walk through the process from start to end, picking up all the triggers.

8   At the end, revisit the roles, triggers, goals and outcomes and ask if they have all been adequately covered.

9   Record all issues and concerns separately as they are raised.

10 Finish on time whatever it takes.

When is a process model 'complete'?

If we examine an island such as Anglesey or Martha's Vineyard to answer the apparently simple question 'How long is its coastline?', we come up against a problem. If there is a road around the island we could measure its length and decide that that is the length of the coast-line. But the road doesn't trace round each inlet (there might be a bridge over it), or round each peninsula (it cuts across the base). To get a 'more accurate' measurement we might decide to walk around the coastline with a pedometer, walking into each inlet and around each penin-sula along the coastal path. We will get a greater distance than we obtained from the length of the coastal road. To get a yet 'more accurate' figure we might decide – having much spare time – to run a tape-measure along the edge of the water as we pace the coastal path. We will obtain a yet larger distance. And so on.

This is a good metaphor for processes. There is always more detail if you want to look for it. Whether the detail is useful and justifies the expense of collection, only the process modeller can determine – there is no simple rule that can tell you 'You have finished!' Completeness is in the eye of the modeller. The answer is in those words in red on the wall.

So, one session might not be enough. We might need to reconvene the group and revisit those parts of the model that we have not finalised. Our note-taking and post-workshop review

will tell us what remains to be done. And it is when we all find ourselves in the room together again that we realise the value of having got to the end of the process at the first session: we have the whole thing laid out in front of us – right or wrong – and we can take a view of it in the cold light of the new day. Having slept on it, when they see it again, our group might decide it was all a terrible mistake: 'That won't work – it's too complicated,' or 'It still doesn't feel sufficiently responsive to the customer,' or 'Will that really work with the new database that is being put in next quarter?' We will not be afraid of tearing it up and starting again if this should happen. Hopefully, things won't be that bad, and the new session will be about filling in gaps, refining where necessary, removing unnecessary detail in places, and general reworking.

By the time we have finished, our RAD might look a tangle, and once again our job is to take it away and rearrange it so that it becomes clearer to read and understand. (Note that I am only going to rearrange it diagrammatically: I am not suggesting that we simplify it in any way.)

STEP 4: USE OTHER SOURCES OF INFORMATION

When we capture or diagnose or design a process with *Riva*, we work primarily with workshops, an approach which has obvious drawbacks: those involved can tell us untruths, they can forget to tell us about interesting things, they can tell us what they think we want to hear, they might conceal things they don't want us to hear, and so on. How can we deal with this? Where possible, we must deal with it at the workshop itself, but there are some other routes that we can use.

Examining existing documents

A document almost invariably gives solid form to an interaction somewhere – it is after all a way of collaborating. There are potentially four groups of people involved with a document, and they represent roles interacting for some reason:

☞ *The author(s).* They have some reason for producing it: to inform, to instruct, to report, etc.

☞ *The reviewer(s).* They provide quality control on the document and its contents.

☞ *The authoriser(s).* They are approving or authorising the publication of the document for some reason: they are the budget holder; the information is being released in their name; they are responsible for public statements; they are agreeing to certain aspects of it, etc.

☞ *The recipient(s).* They are presumably expected to act on the document. The recipient might receive the document 'for information only' and not act on it, but each 'copied to' role represents a potential interaction, whether or not it serves a useful purpose. When we look at the 'copied to' list we might see a list of job functions or positions – 'Finance Director', 'Marketing', 'QA' – or a list of names. In the latter case we have the task of determining which role the recipient is acting when they receive the document concerned.

Documents often record the path that a case has taken through the case process. Good document control keeps a record of the history of a document, making it possible to see what process has been applied to it in reality.

Examining existing documents describing processes

Our organisation may well have documented some of its processes in the past. Documented processes can take a number of forms including procedures, manuals, work instructions, and Quality Manuals. Where these exist they will clearly be an important source for us, in that, in theory at least, they should describe the process in some fashion. But there is a danger here too in that 'in theory': written procedures and practices are not always followed scrupulously, or even at all, especially if they can only be made usable and efficient by 'adapting' them. In a modelling session, people often ask 'Do you want us to tell you how we actually do this, or how we're supposed to do it?'

Highly regulated industries are more likely to have documented procedures. They are also more likely to follow them if there is the constant expectation of an auditor arriving unannounced at the front desk.

Examining existing terms of reference, personal objectives, etc

We have seen how a role groups a set of responsibilities within a process. Such responsibilities are often communicated to the incumbent as written terms of reference, or objectives, or a job description. These can give us clues about

☞ the desired outcome of the work of a role,

☞ the way the job is to be carried out,

☞ the things the role is to produce,

☞ targets that the role must achieve,

☞ resources that the role can use,

☞ interactions that the role must have,

☞ who the role's customers are.

Identifying regular meetings and their purpose

We saw earlier how boards and committees that have regular meetings can play a role, in the sense of having defined responsibilities.

We can represent a meeting simply as an interaction between the roles that are represented at the meeting, or we can regard the group that meets – indeed the actual meeting – as a role in itself. Which view we take depends on whether the group that meets has some responsibility of its own in the process, or is simply a way for the individual roles to get together.

I have modelled a process in which the *Weekly Development Meeting* was shown as a role. The meeting had responsibility for making a joint decision in the process (in theory at least), so we decided to show it as a role. In practice, the meeting was often unable to make decisions, as we discovered when we attended one, because representatives were not always empowered to make decisions on behalf of their departments. As a result the meeting became more of an updating session – an interaction that could be carried out in many simpler, less expensive ways.

When we look at meetings we should ask questions such as

Who attends?

What roles are they playing when they attend?

Why do they attend? Are they there for reporting, receiving information, authorising, or taking decisions on behalf of themselves or the group they represent?

Or does the group that meets play a role itself and have its own responsibilities?

How does the outcome of the meeting get propagated? How does it cause subsequent activity in the process?

By which roles?

Does the meeting report its outcome to other roles who did not attend?

How are they supposed to react?

One way we can answer these questions is to actually attend and observe what happens.

Interviewing individuals

Some items on our 'Issues arising' list might best be addressed by interviewing an individual. If they have not seen a process model before, then we must decide whether walking them through the model is the best way of doing things, as opposed to a simple question and answer session. We need never feel obliged to show the RAD to the interviewee.

Once again, I prefer to have two people carry out an interview: one questions while the other records. It can be beneficial to hold the interview at the interviewee's normal place of work. Very often, in order to explain something to us, the interviewee will say 'Let me show you an example … ,' reach into their filing cabinet and produce an illuminating document. This has to be balanced against the usual problem of interruptions to the interview and hence everyone's concentration, but overall interviewing people on their home territory seems most effective.

A two-hour interview is about the most that both sides can take. The interviewee becomes drained, and the interviewers overloaded with information. We should budget about half a day for the two analysts to go over the information gleaned, in particular working it back into the RADs, and recording new questions and issues that will need to be referred back to the interviewee or on to subsequent interviewees or the next group session.

Setting the scene at the interview is key. Time is limited and there is much to cover, and so we should spend a few minutes covering a number of points with the interviewee. They go roughly as follows:

☞ Thank the interviewee for their time.

☞ Ask how long we have actually got for the interview.

Although we might have asked for a two-hour interview, we are probably starting late, and the interviewee will have subsequently agreed to give someone else the second of our two hours! Agreeing at the outset how long the session will last means that we can pace our

questioning and ensure that we cover the key points rather than wasting all the precious time on smaller issues; and the interviewee makes some sort of commitment to the time that's agreed.

☞ Outline the purpose of the project.

The interviewee might well have heard of the project and have some idea of what is going on. We should describe the project overall, and then place our activity in that framework. Being open about our motivations helps the interview along. It is generally the case – and we should stress it – that our work is non-judgemental: we are not there to observe and then say : 'Aha, there's wastage, that's inefficient, why on earth are you doing it that way?' Our role as facilitators is to bring the organisation to these sorts of statements from its own observations and judgement.

☞ Describe how we are doing our work.

We want to position this interview in the larger scheme of things. Why are we interviewing people, and this person in particular?

☞ Say whom we have already spoken to.

This helps to prevent repetition and to make it clear either that this person is very important and is being seen first, or that we have already spoken to this person's boss and hence are here with some authority. There are of course sensitivities here and dangers too: repetition can be a good thing if it reveals differences of opinion about a process, and a different viewpoint often provides new detail or insights. Also we don't necessarily want subordinates to feel that they have to toe the party line and say what their superiors would want them to say.

☞ Explain how far we have got.

How much have we found out so far? What areas do we think we have some grasp of, and which do we think we are struggling with? People generally like to tell you what they know, so admitting ignorance at this point encourages them to tell.

☞ Describe what we are doing now.

Are we trying to establish the ground? Or do we have a good model already and are now trying to flesh out detail?

☞ Tell the interviewee how they can help us now in this interview.

☞ If we plan to use a RAD with the interviewee, tell them, and add that we shall explain the notation as we go along.

This leaves us ready to get to the core of the interview.

STEP 5: REVIEW, REVISE, VALIDATE THE MODEL

Throughout the modelling project, our RAD will be under constant change as more information is obtained, other information is discarded, the perspective is altered, and so on. The task is all about using our noses, chasing things to ground, following leads, and backing off them. There is no simple procedure for successful process modelling whatever its purpose: discovery, definition, diagnosis, design or enactment. Like any such activity, the skill of analysis is with the analyst as much as with the analytical method.

However, constantly taking the RADs back to the process actors and replaying new versions is clearly a major part of our work. A feedback session is one way of doing this. It naturally involves a group of people who have some stake in the process. It is not unlike the group modelling session, except that the model is now being replayed to the people who originally had a hand in producing it or who, though not originally involved, play a part in the process.

Such a session can serve several purposes:

☞ It is a way of validating the models that have been constructed.

☞ It provides a way of letting the process actors work through and explore potential improvements for themselves, say.

☞ It can be an important part of change management in a radical or incremental change programme, by providing a communications channel from the change management team back to the organisation.

The RAD provides a vehicle for description, discussion, and decision whatever the situation. There is a risk that a new group will want to redesign the process or change the model for some reason – we must be clear whether they have the power to do that and, if so, on what terms and with what safeguards for the model's consistency.

I always reckon to keep each new version of a RAD as it develops. Small changes can be made to the current version ad lib, as the process is clarified or the perspective is clarified. But when a major change is made – a major realignment of perspective or a major restructuring – we should store the current version (call it version N) and copy it to a new one ($N+1$) on which the major change is made. Keeping old versions serves two purposes. The first is the pragmatic one that, even though a major change feels right today, in the cold light of tomorrow it might not seem such a good idea and it is nice to be able to return to a previous version without pain. The second purpose is that the record of how the RAD changes will itself be a useful teaching tool: we can see how the work went, went wrong, was put back on the right tracks, diverged, returned, and so on; all this information will help us to understand the modelling process itself and do it better next time.

Quality control of a RAD

As the modelling proceeds, we must maintain tight quality control over the RAD: being tidy helps enormously, especially since we are showing the RAD to people constantly and there is great value in their seeing a consistent notation and usage in what we show them.

As in any language, good style helps communication. I am a stickler for accuracy in representation: RADs provide a concise and unambiguous way of saying things, and there is no point in throwing that advantage away by being fast and loose with the notation – all the rules below are there for a reason. Get into the habit of using them even when sketching a RAD, rather than trying to 'correct' the RAD once it has been captured.

☞ Label each action in verb-object format.

For instance, *Prepare the monthly report*, *Classify client request*, or *Assemble business case for approval*.

☞ Identify the nature of each interaction with an appropriate verb.

For instance, *Agree …*, *Approve …*, or *Delegate …*

☞ Label each interaction against the initiating role and word it from the point of view of that role.

For instance, *Hand over monthly report*, *Receive classified client request*, or *Approve business case*.

☞ Label important states so that just the interesting part of the state is briefly described. Describe the state with a sentence in the present tense.

For instance, *The monthly report has now been prepared*, *The client request is new and exceeds £10,000*, or *There is no approved business case*.

☞ Annotate the RAD with a highlighted text block when you need to make a comment about the accuracy of the RAD itself.

For instance, *What happens now if the application is refused? What happens to the documentation of a refused proposal?* or *Regional Manager step omitted for simplicity*.

KEY POINTS

To avoid analysis paralysis, ask frequently 'Have we enough detail to do what we are trying to do, to answer the questions we have posed, to achieve the goals we have set?'

A process model is complete when it is sufficient to address the declared goal.

Complement the workshops by

☞ looking at documents that suggest process,

☞ examining written procedures (whether or not they are actually followed),

☞ looking at people's terms of reference,

☞ observing regular meetings,

☞ doing follow-up interviews with individuals.

Quality control your RADs constantly.

Keep old versions of RADs – you may need them! And they hold valuable lessons.