

# Portfolio Kanban

## - Seeing the Bigger Picture

---



***Sandy Mamoli***

NOMAD8

In 2012 TradeMe began introducing Portfolio Kanban alongside workflow and team restructuring to solve problems it was having managing its project pipeline. This paper describes the rationale, process and learnings from that initiative.

**SANDY MAMOLI**, WELLINGTON, NEW ZEALAND



[sandy@nomad8.com](mailto:sandy@nomad8.com)



[@smamol](https://twitter.com/smamol)



# Portfolio Kanban - Seeing the Bigger Picture

---

*Sandy Mamoli, Nomad8*

In 2012 TradeMe began introducing Portfolio Kanban alongside workflow and team restructuring to solve problems it was having managing its project pipeline. This paper describes the rationale, process and learnings from that initiative.

## **1. INTRODUCTION**

TradeMe is New Zealand's equivalent of eBay; a large transactional site where Kiwis buy, sell and trade everything from cars and antiques to clothes, crafts, property and farm gear.

It's also a Kiwi success story, having grown over the past 15 years to a unique position where it commands two-thirds of New Zealand's internet traffic with 1.5 billion page serves a month – and more than 3,000 chickens bought and sold every four weeks!

The company is growing fast. A year ago there were 110 employees in the Technology department, now there are 125. In all, there are 350 people working at TradeMe.

While growth is a good thing, it's not easy. With it comes growing pains and by 2012 delivery had started to slow down.

When I came to give a presentation on multi-tasking and personal Kanban to the Tech team in December that year, my message struck a chord. The Head of Projects, David Mole, invited me to come on board as a consultant.

## **2. OUR WORLD A YEAR AGO**

### **2.1 The Problem**

TradeMe didn't really know how to define the problem they were asking me to help them solve. They knew they had a problem with delivery, resource management and workflow, but they struggled to get clear insight into their pipeline of work.

Like any good consultant, I walked in, shut up and listened – a long-proven strategy for me to find out what's going on and gain an understanding of an organisation. But with TradeMe it was different; everything seemed to be constantly on the move and in a state of flux.

I tried to find out which projects were being worked on, which were running, which were on hold and what was coming next. I wanted to know who was working on what, and whether anyone was working on a particular project at all. But I just couldn't pin it down.

The one thing I could observe was that while there was a lot going on, nobody could tell me exactly what was going on. So I started to ask around to find out whether it was just me who was confused (was I missing something obvious?) or whether other people were in the dark too.

To get more insight, the Head of Projects and I conducted a survey of key Business people at TradeMe and found they had the following perceptions:

*“Projects take a really long time to get finished.”*

*“We don’t know where our projects are at.”*

*“We can’t trust your estimates.”*

It became apparent that the lack of clarity was not just my perception but endemic to the whole organisation, and that it was causing friction between the Business and Technology teams.

We wanted to investigate why this was happening and figure out the root causes.

## **PROJECT STICKIES**

Our first step was simple and, in retrospect, obvious but something that hadn’t been done before. We got into a room, locked the door for three hours and put the name of each project on a sticky note.

Next we drew a large portfolio life-cycle diagram across several sheets of A1 paper and placed each sticky note along the life-cycle according to the stage it was in. (I will describe the individual stages later in this paper.)

In fact, this was our first Portfolio Kanban board and we had just done Step 1: Visualising our Workflow. But we called it Project Stickies at the time.

This process allowed us to see that we had 49 projects in progress. This in itself was not necessarily a problem, except that we had a Tech department of 110 people and it was immediately obvious that 49 projects was too many to be worked on at the same time.

### **2.2 The Causes**

We asked ourselves, “How could this happen?” “How could we not know this?” We knew we had a lot of projects, we knew we had a delivery problem, but we had never realised the extent of it.

Our next step was to look at cause(s) and we found three main ones:

#### Cause Number 1: Lack of Visibility

The main reason, it turned out, was that we just couldn’t see the problem. We had been using an Excel spreadsheet to keep track of our projects. The spreadsheet was updated and maintained and people would add to it as projects arose and developed – but it was effectively hidden. The document wasn’t centrally

visible, and it was only possible to see what fitted on your screen at one time, which meant we were missing the full picture.

This lack of visualisation was what made it possible for us to run too many projects at once without realising we were doing so.

## Cause Number 2: Staffing and Structure

Our Technology staff were organised around professional disciplines with, for example, a BA team, design team, development team and test team and we'd have handovers between disciplines.

As in most matrix organisations, projects were resourced in an ad-hoc fashion according to whoever was available at the outset. For each project there was an allocation of people but people were often part of several project teams and multi-tasking on 3-4 projects at a time.

As a result projects would run in a start-stop-start-stop pattern with a lot of wait time between stages and a long lag before projects went live. Too often, for example, graphic designers and UX experts produced beautiful designs but because no developers were available when they finished, the design languished for so long that eventually other things became more important and the designs ended up in the “design graveyard”. The same frequently happened with code, which would start rotting because all our testers were still busy on other, often delayed projects. These cascading delays caused a tremendous amount of waste.

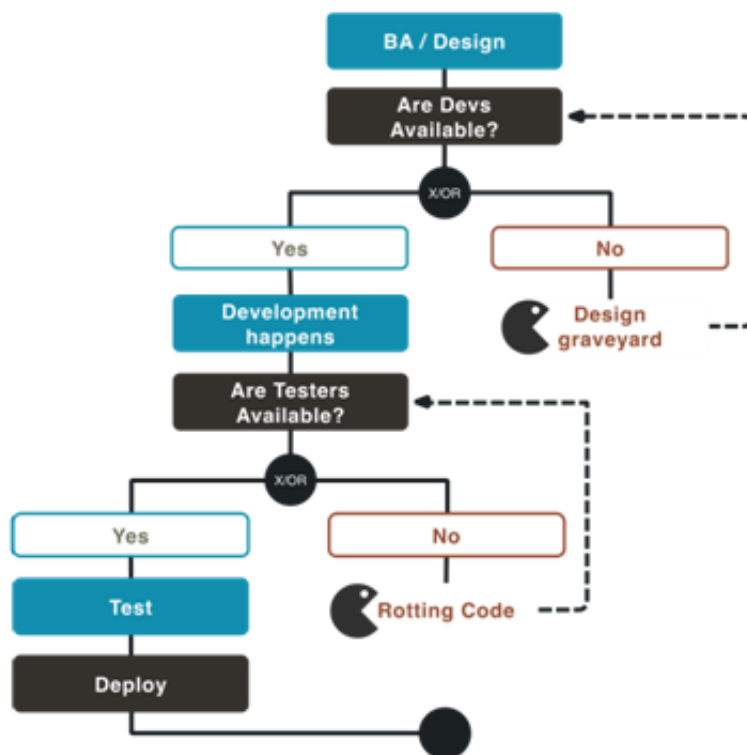


FIGURE 1: A START-STOP WORKFLOW LED TO CASCADING DELAYS AND A TREMENDOUS AMOUNT OF WASTE.

While projects lagged, the business in the meantime would keep moving, change priorities and introduce new projects – which further aggravated the problem.

So even though we had two deploys a day we didn't get any of the bigger features out without it taking a very long time (ca. 9-12 months).

### Cause Number 3: Growth

With around one new person joining the Trade Me team each week, growth was becoming something of an issue and the organisation started to resemble a giant game of Jenga.

The way we managed this growth – from 20 to 30 people, then 30 to 40 people and so on – was to move people around - effectively as you move Jenga pieces in the game.

A manager would move a person from one project to another. New people (each a new Jenga piece) were being added to Trade Me regularly as new developers started, then designers, testers, and BAs, and we would carefully add them to our matrix of people and projects in a way that kept the structure stable.

Just as in the early stages of a Jenga game, this approach initially worked really well for us and the careful balance of people and projects was maintained. But over time we approached a tipping point whereby adding more people didn't mean we got more done; in some ways it was making us even more unstable.

The organisation was growing – and we had to keep growing to serve the needs of the company – but apart from trying to balance people in teams, we had so far made no changes to our structure, project management or communication.

While we never reached a point where the whole thing fell over, we firmly believe we would have done if we hadn't redesigned the game.

## **3. OUR TRANSITION**

### **3.1 THE SOLUTION**

As noted earlier everything seemed to be changing all the time and no one really knew what was going on. By visualising the work in progress and looking at the way work was assigned and managed, we gained insight into the causes of our problems.

We knew that we had to stabilise the system and to do so we needed to focus on two moving parts: the people and the work.

#### **3.1.1 THE PEOPLE**

To tackle the problem of people not being available when needed, we decided to have people work on one project at a time. By having people working sequentially, we could stop the evils of multi-tasking and

shorten the time required to finish projects.

We redesigned our staffing structure to consist of small, stable teams, with each person (in most cases) being part of one team, and one team only, to avoid waiting for people and resources.

### **3.1.2 THE WORK**

We also wanted stable priorities. While we didn't need to consider planning a year ahead, we did want to plan 3-6 months ahead and finish one piece of work before starting another. We wanted that single-piece workflow to keep time lag to a minimum and eliminate start-stop holdups.

“Project Stickies” showed us how powerful it was to visualise our projects. Since Project Stickies was essentially a Portfolio Kanban visualisation, we decided to extend the approach and try full Portfolio Kanban.

## **4. PORTFOLIO KANBAN**

We implemented each of the 6 core practices of Portfolio Kanban:<sup>1</sup>

- Visualise Your Workflow
- Limit Your Work in Progress
- Manage Flow
- Make Policies Explicit
- Feedback Loops
- Improve Collaboratively

I'm mostly going to focus here on 1 and 2.

---

<sup>1</sup> Kanban core practices - [https://en.wikipedia.org/wiki/Kanban\\_\(development\)](https://en.wikipedia.org/wiki/Kanban_(development))

## 4.1 Visualise Your Workflow

Our first step to visualizing our workflow had already happened with Project Stickies, which mapped existing projects to the portfolio life-cycle. Figure 2 shows our initial Portfolio Kanban board.



FIGURE 2: OUR ORIGINAL KANBAN BOARD. THE LEFT-HAND COLUMN REPRESENTS THE DIFFERENT BUSINESS AREAS OF TRADEME AND THE FOUR RIGHT-HAND COLUMNS REPRESENT HIGH-LEVEL STAGES OF A PROJECT'S PROGRESS.

### Ideas / Anything Goes

While the first step in any project is the idea, we chose not to include an Ideas column on our Portfolio Kanban board because we have a separate “Ideas Process” for managing the flow of ideas to realization of concepts.

We think it's important to record ideas and their evolution – we use a tool called Fogbugz to keep track of thousands of ideas generated by people who work at TradeMe and our customers – but we didn't need to focus on those details for this part of our process.

## Stage: Might Do

In this column are concepts that have bubbled up from the thousands of ideas in Fogbugz. Once an idea has developed a bit more substance, it becomes a “Might Do” item.

We were using a “Gateway” session to decide which concepts from “Might Do” would go on to get done. But when we started this visualization process we realised that we started almost everything, and consequently almost everything moved to the next stage.

## Stage: Inception

The next stage in the process is Inception, which culminates in a One Pager (a modified Lean Business Model canvas) about the project. We don’t develop full specifications, requirements or detailed designs at this stage. Rather, we undertake a project inception process which includes a Lean Business canvas and most parts of Jonathan Rasmusson’s project inception from the Agile Samurai <sup>2</sup>.

## Stage: Creating

When work starts on a project, it moves to the Creating column and stays there until the last major deploy is done. This stage includes concurrent Design, BA, Development, Testing, DB work etc. The name reflects the state and that the process is a creative one.

This is where elements are built, in iterations or as prototypes, and where we get potentially deployable code.

## Stage: In The Wild

Once a project has been deployed it moves to the In The Wild column until it is closed. A project is considered closed once we have enough stats and analytics to close the build-measure-learn cycle.

We usually close a project once we know whether we have achieved the desired business impact and success criteria as defined in the One pager.

---

<sup>2</sup> Jonathan Rasmusson, *The Agile Samura*, *Pragmatic Programmers*, 2010 - <http://pragprog.com/book/jtrap/the-agile-samurai>



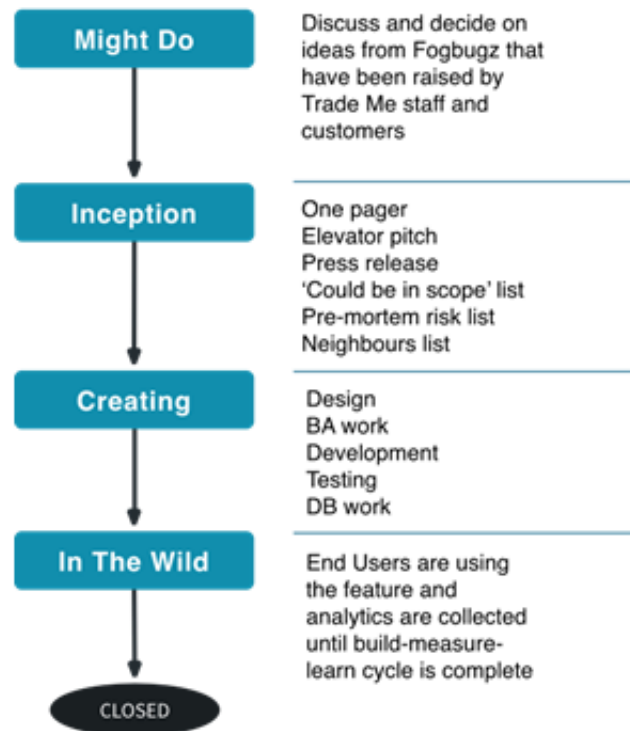


FIGURE 3: THE FOUR STAGES WE TRACK PROJECTS THROUGH ON OUR PORTFOLIO KANBAN BOARD.

So far so good, but just visualising a process changes your understanding but it doesn't change the process or solve the problem! It's relatively easy to convince people of the need for change at this stage. Visualisation creates an immediate connection to the work and to the scale of the projects in progress.

It is during the next step, when you have to decide what to cut from your work in progress, that people start to feel the price they will pay for change.

## 4.2 LIMIT YOUR WORK IN PROGRESS

Limiting work in progress (WIP) simply means reducing the number of projects being worked on at any one time and limiting the number of projects that can be in any one stage of the portfolio life-cycle.

It sounds simple enough, but for us this was the hardest part.

## Deciding on good WIP limits

Determining limits was complicated by the fact that we were implementing Portfolio Kanban at the same time as we were creating stable teams. Since we were dealing with a mixture of ad hoc and stable teams, we couldn't just uniformly set a limit of one project per team from the outset.

As a starting point we talked to the heads of each business unit (Marketplace, Motors, Property, for example) and together we decided on ideal WIP limits. As we had no data or science to rely on, we needed to come up with a way to decide what the ideal WIP limits should be. As a rule of thumb, we assumed that one project team would likely average 5 people and that the number of people working in one business area divided by five would be a good starting point to determine the WIP limit for a business area.

As this was an assumption, and because people weren't always clearly assigned to one business unit or another, we knew that we ultimately had to rely on the gut feeling of intelligent and experienced people. We also knew that the limits would change but at least we had a good starting point.

Initially some people struggled to cope with the concept, so we allowed them to have bigger limits to start with and decided to deal with reducing the limits at a later stage. Our rationale was that it was more important to introduce the concept of limits: as long as the new limit was an improvement on the current situation it was a good starting point and well in line with our pursuit of continuous improvement.

We imposed a WIP limit on all stages except for In The Wild. Notably, we imposed a WIP limit of 8 on the "Might Do" column since we had a history of spending hours and hours talking about the "Might Do" even when we had no capacity to start anything new at all. We used to devote entire project meetings to talking about "Might Do". Now, once the "Might Do" column hits its WIP limit, we move on to talking about projects in other stages.

## Selling the rationale for WIP limits

Imposing a WIP limit on projects from scratch is relatively easy. However, if you're introducing limits into an existing, over-extended workflow it's not so simple.

There are essentially two ways to approach it:

Wait until all projects in progress have finished before starting any new ones. This approach is likely to take a really long time and runs the risk that your organisation and teams will lose momentum and enthusiasm and nothing will really change at all.

Stop working on all projects that are over your agreed WIP limit. This approach is more like ripping off a band-aid. The advantage is that change will be much quicker and you can reap the benefits of lower WIP much earlier.

We chose option 2 and decided to block all projects over our agreed limits.

Preparation was extensive. We had to talk to a lot of people to make this happen. Particularly difficult conversations were held with Business people. We needed to explain to them why it was better for TradeMe if some of their projects were blocked (i.e. that no more work was done on them). We met with everyone affected, we did full company presentations, and we sent a lot of emails.

To make our approach understandable we used the analogy of a traffic jam on the road. It is common knowledge in queuing theory that if utilisation gets close to 80% traffic slows down and eventually comes to a standstill. We needed everyone at the back of the traffic jam to stop their engines, leave their cars and come to the front to help us move the cars out front.

Our path was made smoother because TradeMe was in some pain and Business people weren't entirely happy with Tech teams so they were open to change.

Everyone understood the rationale for stopping work on anything except for the highest priority projects and working on the next most important ones after the high priority ones had gone live.

## Prioritising projects

Our next step was to come up with a prioritised list of projects so we could determine which would go ahead and which would be blocked. This was to be one of the greatest challenges we encountered.

We knew that this prioritisation had to happen across the whole company. We decided to have a workshop with all the leads for the Business units, the CEO, CTO and the Head of Projects: basically everyone who had a stake in what was being built.

Our objective was to work out the 10 most important projects for the year and to rank them. That meant we needed selection criteria in place and a process for how to facilitate the workshop.

Since we had a company-wide strategy in place, we boiled the criteria down to one main question: How does this project align with our strategy? We also decided that any project to be considered for selection should have a One Pager (a Lean Business Model canvas) written about it.

We asked everyone to show up on the day with a list of their top 5 projects for the year and a One Pager for each of them.

For this to work, we needed everybody's buy-in – especially that of the CEO and the Executive Team – to make sure people took the process seriously.

Getting buy-in from the CEO was no problem as we had done enough upfront work to convince people of the necessity of blocking some projects and stopping the excessive multi-tasking. The CEO fully supported the process. In fact, he communicated to the Business people that any project that didn't have a One Pager would not be considered for prioritisation – even those projects already in the Might Do or Inception stages. In the spirit of “finishing things you've started is important”, however, we agreed that everything in the

Creating column would be completed and would not have to go through the prioritisation process.

During the workshop each Business lead introduced their 5 top projects for the next 6-9 months. They wrote project names on sticky notes and placed them on a whiteboard. Then everyone discussed which projects would be most important for TradeMe and the group had a vote.

The outcome was clear and there was consensus around which were the most important projects. People voted not just for their project but for what they honestly thought was best for TradeMe. I was impressed by their engagement and alignment with the company.

Some heads of Business areas pointed out that if we were to work on a company-wide list of priorities they would not be able to fulfil their individual KPIs for their Business units. This made us realise that a lot of people's KPIs favoured local optimisation rather than the much-needed overview and optimisation of the whole. We knew this would have to change and fortunately our CEO and Exec team agreed that this was fair enough and that people could renegotiate their KPIs.

The outcome of the workshop was:

- A ranked list of 14 projects
- A shared understanding of all projects
- An agreement to do a prioritisation workshop every 3-6 months
- Acceptance that things will change

We deemed the session a success. We had a ranked list of projects from 1 to 5, then we had four projects in rank 6, and the rest deemed lower priority. We were happy with that since we thought it was fine to defer the decision about what to do next until priorities 1-5 were complete.

Notably, we found there was almost no overlap between our workshopped list of strategically important projects and what was on our Portfolio Kanban board. We realised that we had been distracted by focusing our work on shiny, cool projects, "fires", and people's pet projects.

This was very useful knowledge that allowed us to correct course and make sure we got the most important projects done. Also, it made us realise how important it was to have a formalised prioritisation process rather than a process reliant on personal relations and shoulder tapping to get things done.

We used all this as input on deciding which projects to block and which to do next.

## Implementation

After the workshop, people who were working on projects deemed lower priority were told their projects were going on hold and that their efforts were needed elsewhere – a hard thing to hear for people who’ve invested a lot of effort, energy, and emotional attachment into what they’re building.

Understanding the rationale for change is one thing; it’s quite another when it’s your project that gets blocked.

Many people were disappointed that their project was put on hold and sometimes reality sank in quite a lot later or people were in denial about what was actually happening.

We began to realise that putting a ‘blocked’ sticky note on the Kanban board was easy; actually stopping anyone from working on the blocked project was way harder. We had underestimated the work involved in chasing people to ensure that what was on the Kanban board reflected reality.

Often the reality of our intention was not properly understood at the beginning and we had to repeatedly explain the benefits of blocking and reiterate that by blocking a project we really did mean no work. There were quite a few, “I know this is meant to be blocked but I’ve just...” and “If I could just get some test time I could...”

There is a delicate balance between education and enforcement and I believe that if we were to embark on this process again we would be more aware of how much work is involved in enforcing intentions. “I’d also look into an underground spy-system to find out when people were working on blocked projects so we could stamp it out!” says Head of Projects, David Mole. Looking back, we probably erred a bit too much on the side of education.

During this time we used visible feedback to illustrate how close Business units were to their WIP limits. We displayed how many projects a unit was still over the ideal limit and adjusted the counter every time a project either finished or got blocked. This allowed us to celebrate along the way and sent a signal to everyone that we were keeping an eye on achieving our goal of reducing the multi-tasking.

### 4.3 MANAGE FLOW

In the past, we had routinely tracked time spent in development and testing but had let those quite detailed measures distract us from the bigger picture: namely how long it was taking us to bring a project from starting point to end users in the wild.

So we changed our focus to what actually mattered: delivering end-to-end. As outlined earlier, we adopted a strong focus on finishing one project before starting the next and had project teams stay together to finish the project before any of its team members could join the next one.

We came to see that as time spent on projects went down, the Business stopped trying to insert new

projects into our existing workflow, which eased the whole process.

We didn't record how much time projects spent in every stage of our portfolio life-cycle, only those that were most problematic. We didn't, for example, record time spent on Inception, since we knew we didn't have a problem there as it generally only took a few days. Inception tasks could only be done when the entire project team was available to do it, and it was up to them to decide how long they spent on it. A rule of thumb was to spend up to one week in Inception on a big, 6- 9-month project.

The single most important aspect of flow management for us was to repeat again and again the importance of managing flow by finishing a project before embarking on a new one and warn repeatedly against the evils of multi-tasking.

#### 4.4 MAKE POLICIES EXPLICIT

It was important that we all had a shared understanding of how to use the Kanban board and exercised discipline in using it. So we wrote and visualized explicit 'exit policies' that dictated when a project could move from one stage to another in the life-cycle.

The exit policies were posted on top of the Kanban board and we used them as a 'definition of done' for each column.

For a project to leave the Inception stage, for example, it needs to have:

- a One Pager (a Lean Business Model canvas)<sup>3</sup>
- an Elevator Pitch<sup>4</sup>
- a Press Release or Product Box (optional)<sup>5</sup>
- a "Could be in scope but is not for now" list (optional)
- a Pre-mortem Risk List<sup>6</sup> (optional)
- a List of Neighbours or people to keep informed

---

3 *Lean Business Model Canvas* - [https://en.wikipedia.org/wiki/Business\\_Model\\_Canvas](https://en.wikipedia.org/wiki/Business_Model_Canvas)

4 *Geoffrey Moore's Elevator Pitch* - <http://www.elevatorpitchessentials.com/essays/CrossingTheChasmElevatorPitchTemplate.html>

5 *Working Backwards: Amazon's Press Release* - [http://www.allthingsdistributed.com/2006/11/working\\_backwards.html](http://www.allthingsdistributed.com/2006/11/working_backwards.html)

6 *Project Pre-mortem* - <http://nomad8.com/how-to-minimise-risk-with-a-project-premortem/>

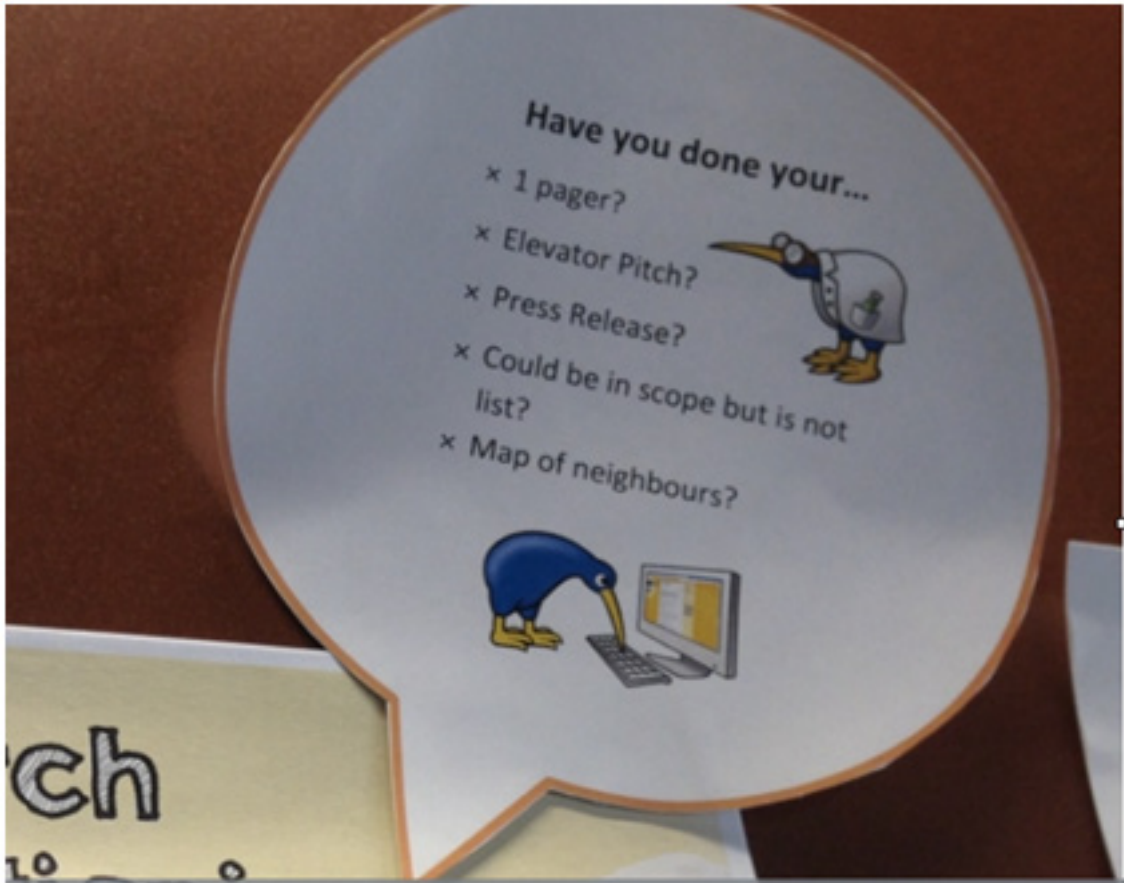


FIGURE 4: EXPLICIT 'EXIT POLICIES' MAKE IT CLEAR WHEN A PROJECT CAN MOVE TO THE NEXT STAGE

We have a few explicit ground rules as well, designed to keep process from getting in the way of speed and creativity. The ground rules are:

Move stickies left to right

Adhere to the WIP limits

Move through every column (unless it doesn't make sense)

It's okay to break the rules – but only if you can convince at least two other people it's worth doing.

#### 4.5 FEEDBACK LOOPS

An important part of introducing such major changes was measuring the impact we were having. We needed to be able to assess whether what we were doing was working and that the situation overall was improving.

We knew that choosing the right metrics would be important because “you get what you measure”. We knew that measures drive behaviour and that it's human nature for people to game them. And we knew our measures didn't need to be pinpoint accurate so long as they reflected trends.

## What We Measure

The following is what we determined was important to TradeMe:

- Get stuff out fast
- Have high quality
- Have happy clients (Business people/end users)
- Have happy employees
- Build the right thing

We spent some time deciding how to measure and visualise the above and ultimately settled on a number of different ways. Three are outlined below.

## How We Measure

We began measuring 'Get stuff out fast' using Cycle Time.




FIGURE 5: A DIAGRAM ILLUSTRATING THE CYCLE TIME FOR PROJECTS FROM OCTOBER 2013 TO JUNE 2014

We introduced Portfolio Kanban in April 2013 and by January 2014 some huge projects that we'd been working on for more than a year became unstuck and finally went live.

Our first read of the graph made us think that things had got worse but in reality we had just delivered some of the big projects that had been going on for a while.

In fact, we went from 49 projects in April to 32 in October 2014. Within a few more months we were down to just 18. Our average cycle time went from 250 days to 55 within less than a year.

## Qualitative Surveys

We also capture trends in customer satisfaction and employee and business happiness through regular qualitative surveys of the business and teams.



## Business Happiness Survey

This survey is conducted at regular intervals and asks:

On a scale of 1 to 5 do you feel...

- you have confidence that the problem you have will be solved, that it will be solved well and it will be solved quickly?
- your projects are hitting the promised dates?
- the amount of time to get a project delivered is right and reasonable?
- project progress is highly visible and you always know what will happen next?
- you know how you can help and also influence successful delivery yourself?

## Employee Happiness Survey

Questions for this survey are based on Daniel Pink's Drive: autonomy, purpose, and mastery.

On a scale of 1 to 5 do you feel...

- you are doing meaningful work that comes to fruition on TradeMe's sites and apps?
- you are allowed to do what's best for your project by focusing on one thing at a time?
- you can and do have a direct influence on how we solve problems and deliver the work?
- you get all the support you need from the people around you? (If you are working within a squad, this question refers to your Squad and Chapter)<sup>7</sup>
- you communicate well with Business people and you know what they want to achieve and why?

Those measures give us a pretty good overview of how things are going for us and whether we're moving in the right direction.

### 4.6 IMPROVE COLLABORATIVELY

We positioned the Portfolio Kanban board in the middle of the office and it has become a place where Tech people go to find out what other teams are working on and where Business people go to see what else is going on. It shows at a glance what is being worked on and what is in the queue, and prompts follow-up discussions and questions.

Our Kanban board has become a communications tool and a place around which we have discussions. We

---

<sup>7</sup> *Scaling Agile at Spotify with Tribes, Chapters, Squads and Guilds* - <http://blog.crisp.se/2012/11/14/henrikkniberg/scaling-agile-at-spotify>

have our weekly projects team meeting in front of the board, which really helps people to co-ordinate and work out technical and/or business dependencies.

The board also helps us to identify opportunities for continuous improvement now that everything is so visible and any bottlenecks, project or capacity issues are highlighted straight away.

Notably, the visibility and discipline of Portfolio Kanban has allowed us to have a process for evaluating “shiny new things” in the context of others. With the visibility Portfolio Kanban gives us, we can assess great new ideas in the context of other great ideas. We’re not wasting hours discussing whether an idea is cool or not when we have either no capacity or other more promising ideas in the pipeline. Also, non-important or rogue projects are being discovered and stopped earlier.

With the information we gather through Portfolio Kanban and the high visibility the Kanban board provides us with, we now have a lot of information and data that we use to run experiments. In general we put a lot of effort into deliberate learning and improvements. Among the tools we use for deliberate learning are regular retrospective sessions across all Tech, across business units and on a company level. We have also developed a taste for Toyota improvement theme Katas.<sup>8</sup>

## 5. WHERE ARE WE NOW?

A year and a half into the process of using Portfolio Kanban and operating teams with a focus on finishing projects before starting new ones, we have seen great benefits:

- In the first six months alone we reduced the number of projects we had in progress from 49 to 18.
- Our average cycle time went from 255 days to 55.
- We now have visibility of work in progress and work coming up. I don’t think we could lose overview again, certainly not to the extent that we did before.
- We are now working on the right stuff. Since we’ve stabilised some of the moving parts we don’t have to fight so many fires which means we focus on strategically important work.
- We are delivering more, faster and with less multi-tasking.
- There is no more start-stop.

The increased visibility we now have has also helped bridge the gap between Business and Technology. “I think Tech gets more empathy now that it’s all visible,” says Head of Projects, David Mole. Our tracking also shows how Business people’s perceptions about Technology have changed:

---

<sup>8</sup> Improvement Theme Katas - <http://blog.crisp.se/2013/05/14/jimmyjanlen/improvement-theme-simple-and-practical-toyota-kata>



FIGURE 6: EVERYTHING IS TRENDING THE WAY WE HOPED

## CHALLENGES

Of course we still have challenges. When things are going well people have a tendency to get sloppy. We sometimes find, for example, that we get busy and forget to update the Kanban board and it quickly gets out of date and loses some of its value. We found that it is very important to have a designated owner, in our case the Head of Projects, and a regular interval at which is it updated – for us this is at the weekly projects team meeting.

It is also sometimes difficult not to get too attached to the current layout of the board. We had a phase where the workflow was represented by beautiful pictures made by the design team which meant we all felt reluctant to change our workflow because the board was so pretty.

Another challenge is that we have three offices in three cities. We don't want to use a non-tactile tool for fear we would lose the focus point for discussions and meetings in our main office. So we decided to keep the physical board and messages but send photos to our satellite offices on a weekly basis. We hope to convince them in the future to maintain a parallel board.

## 6. CONCLUSION

Portfolio Kanban has given us visibility of problems we didn't even know we had. It formed the basis for being able to solve our problems and improve our ability to deliver, which in turn has made it possible for us to focus on how to take Trade Me to the next level.

## ACKNOWLEDGEMENTS

I'd like to say thank you to the following people and companies:

- Trade Me for their openness, generosity and courage in letting me share experiences and insights into the company.
- Nanette Brown for being a great shepherd and pointing out lots of areas for improvement.
- David Mole, Trade Me's head of projects and my partner in crime for pulling off revolutionising Trade Me with me.
- Julie Starr, flow diagramme genius and kick arse writer for helping with her journalism background and heavy editing.
- Brenda Leeuwenberg and Anthony Boobier from Nomad8 for reading, editing and advising.

## REFERENCES

**Kanban core practices** - [https://en.wikipedia.org/wiki/Kanban\\_\(development\)](https://en.wikipedia.org/wiki/Kanban_(development))

---

**Jonathan Rasmusson, The Agile Samura, Pragmatic Programmers, 2010** - <http://pragprog.com/book/jtrap/the-agile-samurai>

---

**Lean Business Model Canvas** - [https://en.wikipedia.org/wiki/Business\\_Model\\_Canvas](https://en.wikipedia.org/wiki/Business_Model_Canvas)

---

**Working Backwards: Amazon's Press Release** - [http://www.allthingsdistributed.com/2006/11/working\\_backwards.html](http://www.allthingsdistributed.com/2006/11/working_backwards.html)

---

**Project Pre-mortem** - <http://nomad8.com/how-to-minimise-risk-with-a-project-premortem/>

---

**Scaling Agile at Spotify with Tribes, Chapters, Squads and Guilds** - <http://blog.crisp.se/2012/11/14/henrikkniberg/scaling-agile-at-spotify>

---

**Improvement Theme Katas** - <http://blog.crisp.se/2013/05/14/jimmyjanlen/improvement-theme-simple-and-practical-toyota-kata>

---