Managing Risk by Measuring Outcomes

The 18th April event at the British Library Every CIO's Guide to Risk Management (outline at http://www.thehubevents.com/every-cio-39-s-guide-to-managing-risk/prod_94.html) drove a number of timely and relevant discussions pertaining to the management of risk in the real world.

Most delegates agreed that formal risk management, with upfront analysis to identify risks and develop countermeasures, is the ideal – but also raised the point that many organisations simply have neither the time nor budget to carry out a formalised risk management programme.

The formal definition of "a risk" often results in heated debate – but for practical purposes a risk can be defined as a possible event which, if it occurs, will have a negative effect on one or more desired outcomes.

Some formal risks will always be identified and included in a project's risk register: events such as supplier failure or a new technology failing to live up to expectations can, and shoud, be planned for in advance. Other events may still occur, but not be identified early enough to watch for the cause.

How does a project, a programme or an organisation guard against the effect of the unknown?

Think about the definition again: "an event which has a negative effect on one or more desired outcomes". Also consider that not all negative effects are absolute, such as a vendor ceasing to trade or support a particular product, but may manifest themselves over time as may be the case when a new software component doesn't perform as described or as expected, leading in turn to an effect on delivery timescales.

Measuring actual progress against the planned or expected timeline for achieving outcomes will provide an early indication if some unplanned event is occurring and having an effect. Effort can then be focused on identifying and countering the cause.

This is the approach used in many Agile frameworks: as an example, consider the SCRUM burndown chart below. A series of delivery outcomes (story points) are planned to be achieved in a time frame (a number of sprints). Progress is measured from the start of the x-axis (time) with 100% of work outstanding at the beginning, and 0% remaining (all work completed) at the end.
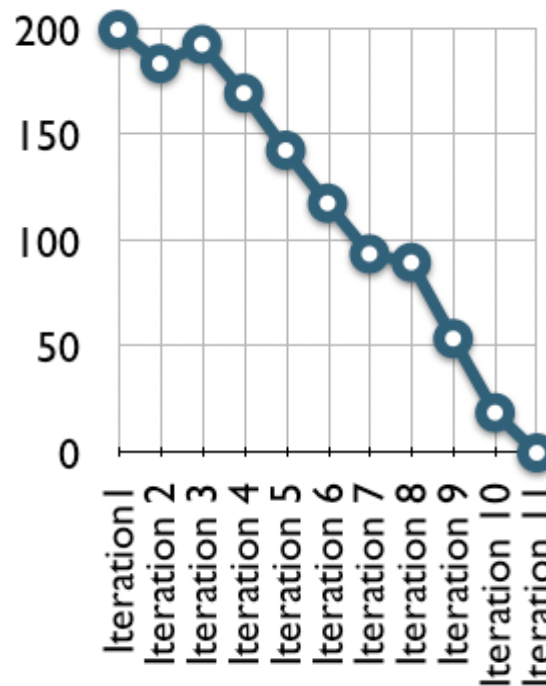
Example of a SCRUM burndown chart

Two measures indicate whether or not known (or unknown) risks are affecting progress:

- The actual trend of the burndown chart (rate of change) tells us if the rate of completion of work units is sufficient to finish the work by the planned end date.
- As estimates are revised (usually as iterations refine the team's understanding of the work to be done, or the addition of new story points/requirements) any <u>upward</u> trend in the graph of work remaining over time indicates a serious problem.

The example chart shows an upward change in the estimated amount of work remaining at iteration 3. Something has happened, or started to happen, which is affecting the desired outcome of completing all the work units on time.

This could be an expected risk, such as a change to the stated requirements, or it may be an unexpected risk – perhaps a chosen technology proved more difficult than expected to implement in the first two iterations. Either way, the project manager is made aware of the effect of such a risk by its effect on the completion rate.

The burndown chart is a specific case of a generally effective approach to monitoring the effect of external factors on the achievement of goals. To make this approach work, we need four things:

- A clear definition of goals/outcomes
- A set of data points with which to measure completion
- A timescale in which to achieve completion
- An understanding of dependency: which tasks must be completed to allow others to be started or completed

Progress towards outcomes can then be measured:

- Are tasks being completed on time?
- Is the rate at which tasks are being completed within the planned timeframes? Do estimates need to be revised?
- If tasks are not being completed on time, what's causing the deviation?

A good estimating framework is key:

- are the timescales allocated to tasks accurate?
  Close attention needs to be paid to the actual completion versus estimated completion of tasks (or work units, or story points).
- Are estimates based on a framework or method, or are they simply guesses (or even worse, optimistic guesses!)?
- Is there a process in place to review and revise estimates based on completion rate and lessons learned during the project?
- Are requirements clearly defined, or are they changing?

What are the human factors telling us about the project? Does each person:

- Have a clear understanding of what they need to do?
- Have all the information needed to complete their immediate task?
- Have the skills needed, or does time need to be factored in for learning about / familiarisation with new tools or technologies?

Mitigate Risk: Act Early

Simply identifying the effects of risk on a project isn't enough. The early warning signs must be acted upon. The temptation to fix systemic problems "on the fly" or "in a future iteration" must be resisted – even if it means an apparent stalling of progress while the problem is investigated and resolved.

We will be re-visiting this subject at our next Risk Management seminar in Manchester on 20th September. Details are on http://www.thehubevents.com/

© 2012 Gavin Martin