

GridPP

UK Computing for Particle Physics

Bottom up Design of Change Process for the RAL Tier-1

Rediscovering ITIL from scratch

Andrew Sansum

Presented by John Gordon

25th April 2012



Science & Technology Promotion Council
Rutherford Appleton Laboratory

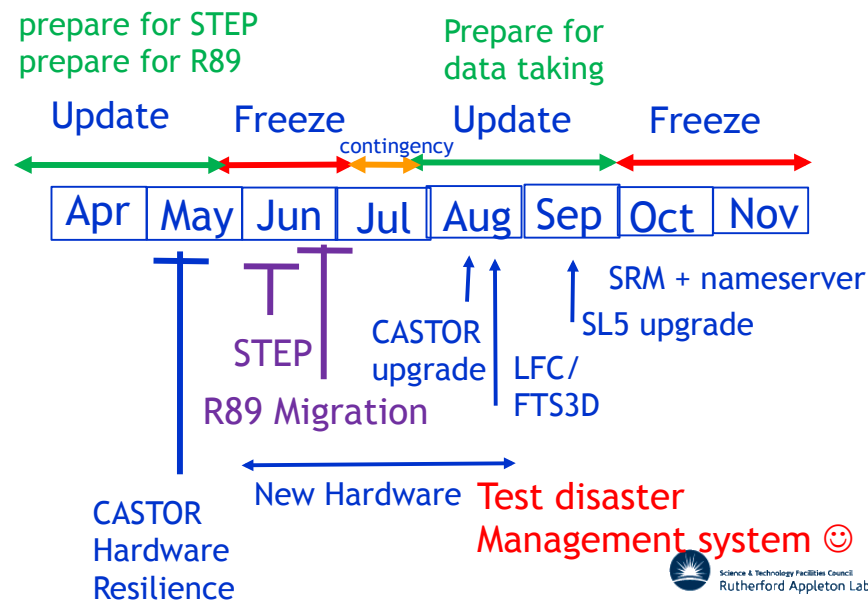
- Last Service Challenge and first LHC data
- Coordinate changes at weekly operations meeting - informal
- Risk Unaware/Risk Averse
 - No changes during sensitive data taking periods (Freeze periods)

- No central control over teams
- Significant number of overruns/failures.
- Component interdependence not fully recognised leading to unrecognised knock on effects
- Inter team communication failures.
- Freeze not sustainable for 2010



GridPP
UK Computing for Particle Physics

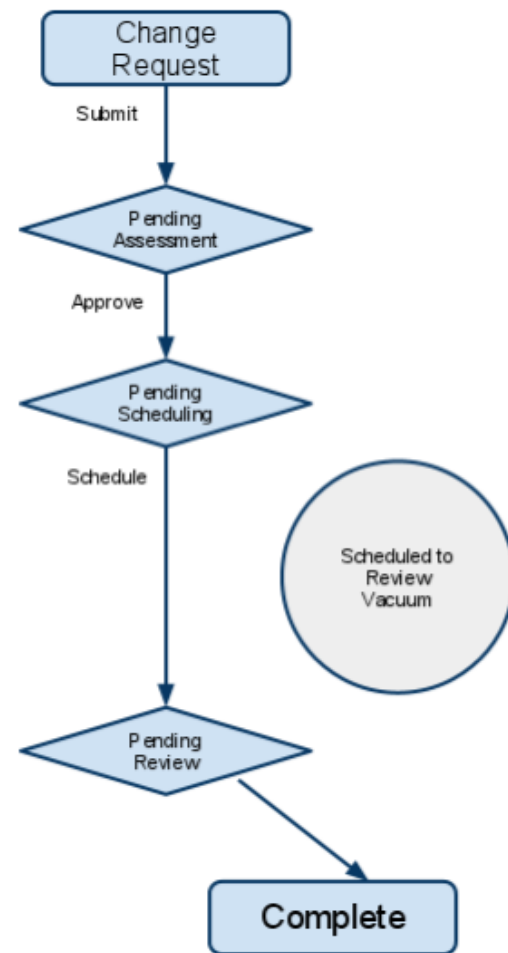
The Plan



- Need to be able to carry out changes during data taking without messing up.
- Must team change culture from “development culture” to “LHC production culture”.
- Some antipathy within teams:
 - To central control ... “you don’t trust me”.
 - Fear of bureaucracy ... “filling in forms is too much hassle”
 - Fear of delays “but I need to do it today”
- Consult widely to obtain team buy in.
 - Teams leads all agree what makes a good change and what causes changes to fail.
- Design
- Hearts and minds to reassure team members.

- Be risk aware not risk averse
- Change form structured to drive culture of good practice
- Skeptical approval process designed to “push back”
- Senior technical staff work together to improve changes.
- Approval by consensus of change team not change manager
- Process should ensure good communication between teams
- Delegate authority to lowest appropriate level
- Don't inhibit routine changes they should already be optimal
- Accept we need to respond quickly to emergencies.
- Separate change approval from change scheduling. Change team approves changes. Production Manager schedules changes

- Change request not needed for routine/standard changes or post hoc for emergencies.
- Team member submits form
- Team Leader assessment
 - Likelihood*impact=risk
 - - if risk = low then APPROVE
- If risk not low -> review meeting
 - If happy approve, else request improvements.
- After implementation review outcome and score change.



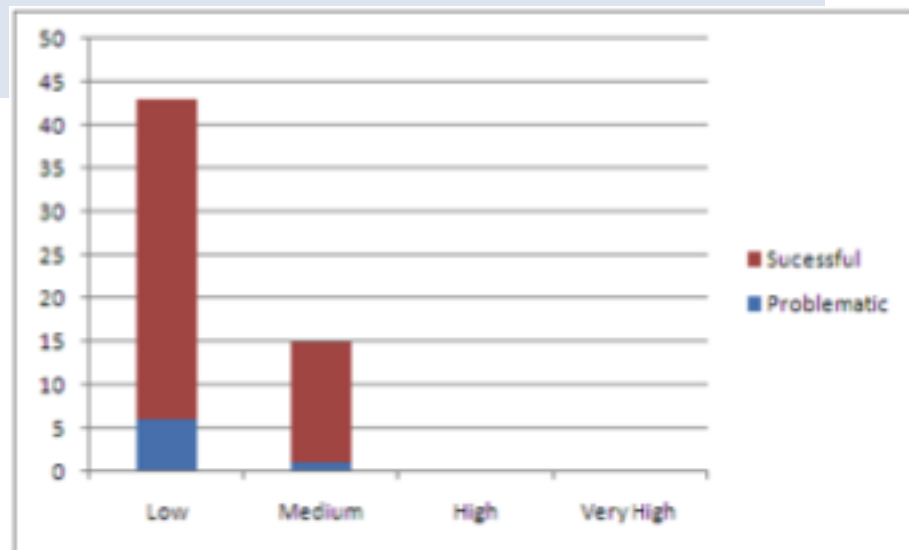
- Process working very well. Hearts and minds mainly won.
 - 58 changes of which, 6 “problematic” (lowest problem level) and 1 overran scheduled downtime. Happy with this level.
- Main problem in data, likelihood assessment.
 - Negative correlation between fault likelihood assessment and fault occurrence.
 - No medium or high likelihoods!

•Poll team: “what does low likelihood mean”

- Range from .0001 to .3!

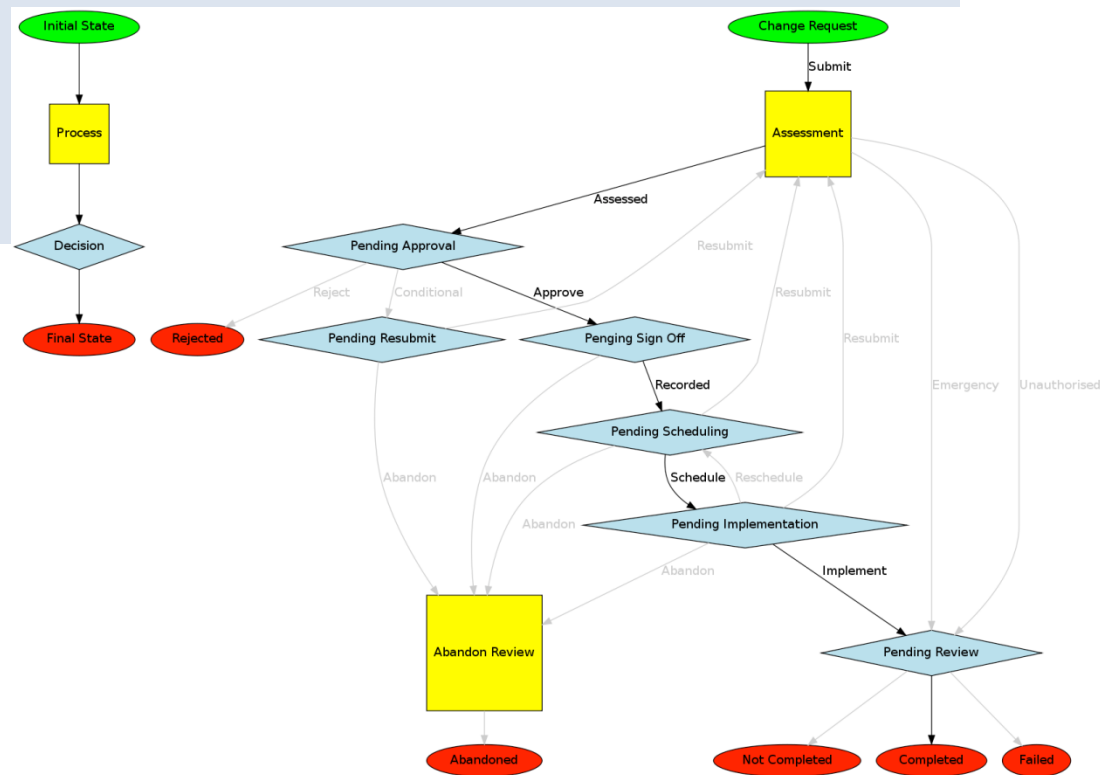
•No appetite to identify changes as high likelihood!

•Lack of enthusiasm to score outcomes as bad or failed!



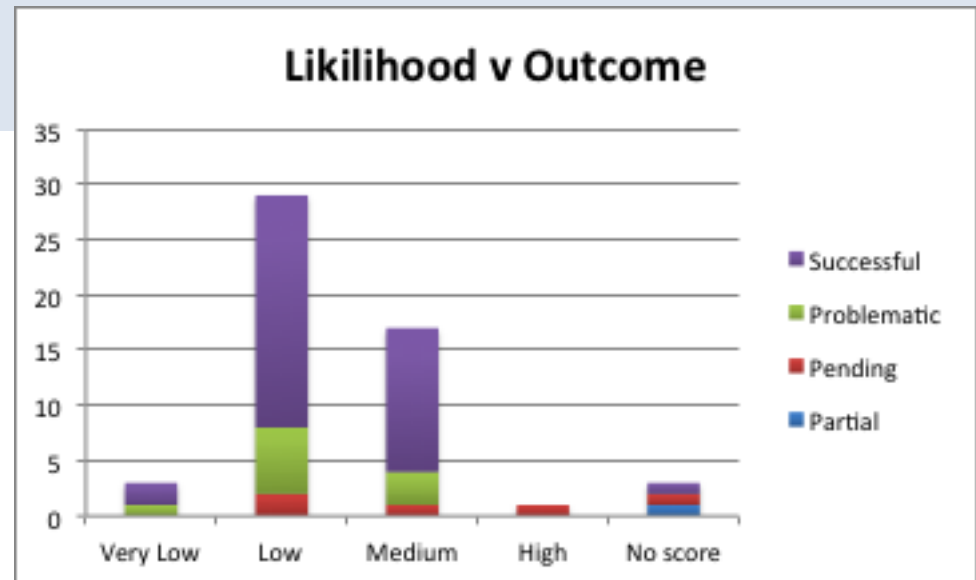
- Good overlap! We have similar concepts:
 - Standard changes, normal changes, emergency changes
 - Risk assessment
 - Scoring and review

- Use ITIL to enhance our workflow
- New end states (failed, incomplete, abandoned, rejected)
- Emergency handling
- Concept of post hoc capture of unauthorised changes (for statistics purposes only).
- Back pressure: resubmit



- Reviewed 2011 changes.
 - 53 changes of which, 12 “problematic” or overran (4 overruns)
- Main issue remains risk assessment not strong predictor
 - Likelihood estimate is not a good predictor of problem occurrence
- Most changes very successful but we have become less tolerant of problems.

• Definition of “problematic” has drifted to encompass “did not go according to plan but outcome OK”



- Almost all team members recognise benefits of the system.
- Has changed culture. Better testing and phase in and better reversion plans. Communications cockups now rare!
- Hearts and minds success owing to:
 - Emphasis on improving changes not approving changes.
 - Ensuring fast response for urgent issues ensuring best chance of success even at short notice.
 - Be phlegmatic when truly urgent emergency changes are carried out without approval (recognised by workflow)
- Beginning to percolate out of Tier-1 to other e-Science projects as team members wish to use it elsewhere.
- Change form beginning to be used within teams as method of capturing work plan even when out of scope.



- The system is fundamentally fine already. Some benefit from improvements, but main benefit already exists.
 - Likelihood assessment not working. Either give up, improve guidance to assessors or (probably better) use actuarial method .. score according to change properties (eg complexity, component, testability, phase-in, reversibility, who asked, day of week, role of dice ..).
 - Inconsistent impact assessment. If the component is critical but the change can be backed out quickly is impact low or high etc etc ... guidance needed.
 - Lack of desire to score bad outcomes - better guidance and hard line.
 - Identified risks often weak. Restructure form and use risk identification to drive risk mitigation plan.

- **Infrastructure**
 - Critical services (and cooling) supplied with diesel backed UPS. Multiple power feeds to site, multiple physical network paths
 - Building up modest “on-site” DR centre remote from main machine room (ATLAS Centre). National and international critical services, configuration management etc. Some cover but more to achieve here.
- **Risk**
 - Management of risk has long been part of our standard project management by our funding project (GridPP) and corporately by the organisation (STFC)
 - Risk to IT services is part of this.

- Processes
 - Disaster Management Process
 - Formal management and escalation process for issues that have the potential to affect Tier-1: “Safety, Reputation, Finance, Operations).
 - Has contingency plans and escalation processes but main emphasis is on catching problems early not when they become disasters
 - Exception Handling (like many sites)
 - On-call gives many benefits beyond out of hours cover. Particularly ensures well developed remote operations infrastructure. Regularly carry out (whole team) work at home days to ensure remote operations work when needed. No problem to handle unexpected site closures (eg for snow).

- We have built a change control process that works pretty well.
- Although designed from the bottom up it has many similarities to ITIL change management model.
- We updated process to look more like ITIL, but new features actually rarely used.
- If we had started with ITIL and imposed a solution
 - Team “buy in” may have been less
 - We may have generated too complex a solution
- We are still challenged to identify which changes are likely to fail.
 - Although tempting to give up on numeric risk assessment, being able to differentiate between low and high risk changes is vital ... further work needed here.