

**Talk given at the Agile Business Conference (ABC2012)
10th October 2012
held at the Inmarsat Centre, London**

***Agile Project Management for Government
Case study:
The Success of the FBI Sentinel Project***

Brian Wernham

© Brian Wernham 2012 CC BY-NC-ND

Good afternoon to you all!

There are some important questions which we will address together today. Agile promises to deliver technology projects faster, at better quality and with better budgetary control than the traditional project management approach.

- So what are the factors that are stopping us from using Agile project management more widely in Government?
- Indeed, what evidence do we have that it has something to offer on Government projects?
- How can you convince your colleagues that Agile should be more widely used?
- How can your organisation use Agile on large, mission-critical projects?

There is now an extensive literature available on Agile. Qualifications in Agile are available, for example, from the APMG and PMI. Conferences, such as this one, are drawing large numbers of attendees.

So, we have Agile processes defined, and we can train people in their use. But this is not enough.

Most of the Agile literature out there is based on a leap of faith - it has not been evidence-based. What many

people in your position have told me they need, is a credible, evidence-based argument that puts the business case for Agile to top management.

I have spent the last year researching large-scale successful Agile projects in governments around the world to find this evidence. Many of these projects have been in the US but also in the UK, Australia and elsewhere. Are there common factors in these success stories that you could apply in your projects?

My research has been focussed on critical, main-stream projects. There are Agile success stories out there: the Ministry of Defence, US Veteran Affairs, the Government Digital Service, Housing Benefits in Australia – all over the world these pockets of excellence demonstrate that government projects can be Agile.

As an example: in New Zealand the government got a disaster compensation system working within three days of the Christchurch earthquake. The project team used an Agile approach to control their work. Releases of working software were scheduled on a daily, half-daily and sometimes even hourly basis. The system was not small - it paid

out more than two hundred million dollars, and ensured financial continuity in the face of a natural disaster.

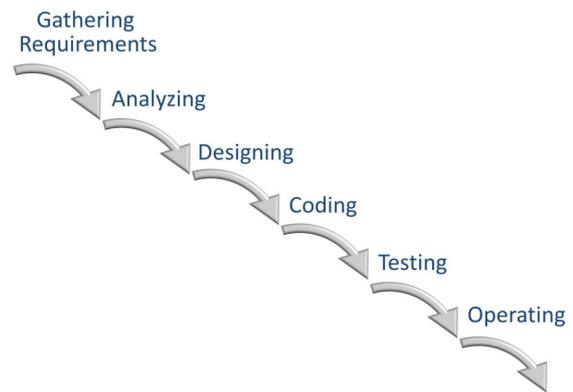
My research findings are that we need more than methods, processes and best practice guidance.

We need people to lead their projects, not manage them. What we need is Agile leadership. And we need proof that Agile leadership will bring success.

Many people in your position face barriers to adopting Agile in their organisations. Perhaps that is true for some of you. The research shows that commitment to waterfall processes is widely spread. And the tendency to start up mega-projects fails to save money, often increases risk and seldom delivers on time, if at all.

What is a 'waterfall process'?

Waterfall projects segment work into discrete steps, each dependent on the completion of the previous step, but without feedback or iteration. When using a waterfall approach, one cannot start a step until the previous has been completed. This leads to a series of one-way 'Gates'.



Once one has committed to swimming downstream, it is impossible to return to an earlier stage without a lot of effort – similarly difficult to attempting to swim up a waterfall.

In contrast to doing just enough design, a waterfall approach requires a grand design in detail before any solution building commences. A waterfall approach is appropriate for some civil engineering projects that are monolithic in nature, such as building a skyscraper, but in technology projects a waterfall approach will tend towards what Kent Beck called 'Big Design Up Front' (BDUF) when describing a fundamental problem of the waterfall lifecycle – that it relies upon pinpoint accuracy and perfect logic at every step if it is to produce a workable solution. Kent's argument, and one that I emphasize here, is that we should aim for Enough Design Up-Front (EDUF), not BDUF.

An example from popular culture!

You may recall an episode of the 'Simpsons' where a Detroit car company decides that being an average American, Homer is the perfect person to design a new car. Homer is given entirely free rein in the design and specifies a car with every feature he could ever want. A bubble dome, a Rolls Royce radiator - and huge tail fins!



The car turns out to be totally unusable, and too expensive to produce.

You can help your organisations think differently about how objectives are agreed up front on projects. You can influence your colleagues in procurement, and those that carry out project audits and assurance work to adapt and harness an Agile - not waterfall, approach.

One of the things I want to explore with you this afternoon is how certain Agile leadership behaviours in government projects bring success.

Project success depends on leadership behaviours, not processes. And successful Agile projects are led, not managed.

In particular I want to tell you a story of Agile success - the FBI Sentinel project. It is one story of many where a large-scale project, with the objective of developing mission critical technology has succeeded where a previous traditional waterfall attempts did not. Let's consider why two previous attempts at Waterfall failed at the FBI, and the consequences. Then let's investigate how the Agile approach was made to work. Then let's identify the leadership behaviours that were critical to the eventual success when the FBI went Agile.

I find the success of the FBI Sentinel project a fascinating illustration of the weaknesses of the waterfall approach and how Agile leadership delivered where traditional approaches failed.

There are three reasons why this case

study is relevant to you.

Firstly, this was a large project - hundreds of millions of dollars.

Secondly, this Agile project succeeded where two previous waterfall projects using the same technology had failed.

The third reason that this case is relevant, is that this was a high-profile and mission-critical project - in government. Precisely the environment and scale at which, I am sure you will agree, Agile is often misconstrued as being inappropriate.

Back in year 2001, the FBI was using old technology to record evidence and other intelligence information. They relied upon ad-hoc processes to share documents, photos, and other electronic media. The handling of the Oklahoma City Bombing case highlighted the deficiencies of the FBI's old technology and business procedures.

That year, just one week before the scheduled date of the execution of the Oklahoma City bomber, the FBI revealed that it had not disclosed

over 700 documents to the defence attorneys. The FBI had simply forgotten to send materials and in many cases had lost evidence.

The legal process was thrown into turmoil, a stay of execution of the bomber was granted - the FBI came under severe criticism. An independent investigation showed that the combination of the old computer system and manual processes were to blame. The computer system simply did not support the critical operational processes, it was over complex, and difficult to use.

The FBI started to set up a project to build a new Virtual Case File system (VCF). This VCF project was to be let as one massive contract at a total cost of \$379m. It was planned as a classic waterfall project - with a grand design being drawn up before work would start on the development of a monolithic system. Testing would be carried out at the end, and the whole system would go live at once - a classic big-bang implementation approach. Science Applications International Corporation (SAIC)

The 9/11 attacks increased political pressure for better homeland security and data sharing between

agencies.

Responding to this pressure, the FBI made promises to bring forward deployment of the new VCF case management system by 6 months - to December 2002. Then they received an additional \$78m of funding and promised to chop off another 6 months from the schedule.

The classic symptoms of waterfall project failure started to reveal themselves. Project plans were found to be unrealistic, and the oversight of project spend was inadequate.

It became obvious that the project would not meet its accelerated deadlines. A commitment to using unproven thin client technology had been made - and the design for a web-like access to a centralised database was deeply flawed - BUT: up-front contracts with suppliers bound the project to this technology and the testing that could have revealed these flaws came too late to allow a change of direction.

After a year of development, it became evident that a total re-write of the software was required.

Even with the additional \$78m of funds, the project missed its milestones. Audit reports took a

traditional view of what was wrong: more discipline was required - more detail and planning needed.

Work continued, and each year the deadline was put back by another 12 months. Every year a new project executive was appointed, but eventually the project was cancelled in 2005.

Of course, the need for better integration and co-ordination of FBI intelligence remained - it was obvious that a new project had to be set up, and lessons needed to be learned.

Many reports were written analysing the failure of the VCF project. Unfortunately, they applied a waterfall perspective in analysing why this, first attempt at a replacement system had failed. Their comments pursued the line of thought that if more detail had been planned upfront, with a more strict set of waterfall processes, then failure would not have occurred.

Examples of their conclusions listed many possible reasons for the failure:“(The project had) poorly defined design requirements, a lack of mature management processes, high management turnover, poor oversight, the lack of a mature

Enterprise Architecture ... a lack of specific completion milestones, review points, and no penalties (for suppliers) if milestones were not met.

However, one factor that these reports did not consider was whether more upfront planning, design and top-down control could fix a broken waterfall model.

Attempt 1

- Virtual Case File (VCF) Project
 - Monolithic BDUF design
 - Waterfall approach
 - Budget \$379m (later \$475m)
 - Big promises
 - 2002: no delivery
 - 2003: no delivery
 - 2004: no delivery
 - 2005: cancelled

Anyway, in 2005, with the FBI still relying on its increasingly outdated old case management system and complicated manual procedures, plans were drawn up for a second project - to be called "Sentinel".

The aim, as with the cancelled VCF project, was to create a web-enabled case management system and to develop it using a waterfall approach based on a Big Design Up-Front. Just as with the previous VCF project, Sentinel would take years to develop.

So, the first phase of the Sentinel project gave users a web interface to

hide the screens of the old, difficult to understand mainframe system.

But, despite a better search facility, no new data capture or sharing functions would be added until the whole system was ready.

In other words, the initial planning for the Sentinel project promised no net benefit to the users until the whole system was developed.

Jack Israel, the Chief Technology Officer at the time explains that:

"The new screens just allowed agents to interact with the old system through a sexy-looking Web browser. Some called it lipstick on a pig!



Web browser access to old system
- No new features
- Many critical functions missing

Yes! - And at sixty million dollars it was expensive lipstick!

But, the substantive deliverable of the Sentinel project - a new case-handling system, was planned to

take several years, and be completed by 2009. Initially, delivery was broken into phases, each about a year long.

Everyone agreed that the new Sentinel project would address the problems of the previous VCF project and deliver an up-to-date case management system. A strict waterfall lifecycle described in a "Life Cycle Management Directive" would reduce risks.

Much time was spent carrying out a beauty-parade of suppliers. A desk-based exercise scored suppliers' proposals against theoretical statements of work and hypothetical project schedules.

The FBI finally awarded the contract to Lockheed Martin to develop the Sentinel system over three years to 2009. The total project budget was \$425m.

Now, although \$305m was budgeted for Lockheed Martin, over \$120m was allocated for the FBI to run a massive program office to carry out detailed and prescriptive oversight of the work.

That's one quarter of the budget being spent on planning and control of the contractor!

Project control was setup in a traditional waterfall fashion. The specification was fixed at the beginning and spend was allowed to vary. The initial estimates could not be validated, so a contingency was allowed for. This 15 per cent contingency was to prove wildly inadequate.

Everyone was convinced that risks could be averted. So much so, that very little contingency planning was carried out. With this much control in place, and external auditors crawling over every detail of the project, everything would run according to the detailed plan - surely?

As part of the first phase, the "lipstick on the pig" features were delivered. They were 2 months late, but most worryingly lacked 57 promised critical features.

After a few months or so of using the phase one system its use was being abandoned. There were many features which had not yet been catered for. The old system still needed to be used for the opening and closing of cases, so - many users found it easier to stick with the old system rather than switch back and forth.

No data cleansing activities had been carried out. The data on the old system was still in a mess, and the team had not yet migrated it over into a test database. This was a key task that would have flushed out the technical difficulties in advance of phase two. Despite this, the FBI paid Lockheed Martin as if phase one was 100% delivered.

But: Audit reports continued to be optimistic because of the apparent comprehensiveness of controls and processes that had been put in place.

In December 2008, Chad Fulgham was appointed as the new CIO.



Chad Fulgham

Fulgham, who came from Wall Street, and brought a business mentality with him that favoured quick results rather than drawn out planning. He saw that little had been delivered, and that key tasks were behind schedule. Allowing the team

to take a year to deliver the next phase of work was very risky.

The new CIO decided to carry out a strategic replanning - Fulgham now planned for outputs every 3-6 months. He also noticed that the paper specifications for Sentinel did not take into account the recent organisational and process changes. So he now made sure that these changes were incorporated into the grand design. This large-scale change, of course, took time and money - plans slipped another 6 months.

Fulgham was publically optimistic about the prospects for delivery, but the project still remained bound to detailed specification and plans that were very inflexible, and the idea of delivering every few months soon proved to be unworkable.

As functions were delivered, the users found that they did not meet their requirements, and the technical approach needed to be reworked again and again.

The planned end-date of phase two slipped, the cost increased by another \$18m, and the security and authentication functions were still not working.

Then, it became evident that the new system actually needed a faster network than was expected. Some users had reported that it could take up to half an hour just to login to the system, so an additional \$39m was spent to improve and streamline the network.

The end-date of the Sentinel project slipped again to September 2010.

The rigid and hierarchical project reporting structure was large, unwieldy, and exhibited a huge optimism bias in status reporting. The bloated \$120m budget for the PMO was spent on inexperienced project managers with general administration backgrounds. They had received basic training, but had little or no background in technology development.

Throughout the project, the reports produced by this PMO, despite being detailed and full of statistics, never reported even one sub project as being in trouble, even as the project was obviously out of control. As late as December 2009 the FBI was still “expecting to provide capabilities to users sooner than originally planned”.

Users rejected Segment 3 of phase two during testing, even though it was theoretically compliant with the FBI’s specifications. They required a complete redesign of the screens. Despite these problems, the FBI Project Management Office (PMO) remained optimistic with project status reports showing “a horizontal thermometer, which expressed the project’s overall status in red, yellow, or green. From meeting to meeting, the temperature never changed—it was always yellow, trending toward green.”

However, the more difficult tasks were left to the end. Important tasks, such as developing the migration processes, were left until fourth and final segment of phase two. Migration was known to be a key problem. Names, addresses, and phone numbers in the old case management system did not match the format written down in the designs for the new system. In the end, the data migration processes and interfaces took two years to create, and when delivered in 2010, they were still not adequate.

Sentinel Phase Two Is Stopped

Then on March 3, 2010, the FBI

decided to reject the deliverables from the fourth and final segment of phase two because of continuing usability, performance, and quality problems. The FBI issued an order to Lockheed Martin to stop development on future phases until the problems were resolved. The FBI could now not be sure that the system would meet user requirements, and could not agree with Lockheed Martin how the project was to proceed. The viability of the September 2010 end-date was called into question.

Not only were some of the essential functions still missing, there were also significant performance issues. These were not just due to poor network infrastructure, but also to poor quality in the coding of the software. In some cases, users could create and use fake identities when signing documents electronically.

More functions were delivered to FBI agents to use, but the system still only had the capability to process four of the 18 forms, and these only partially.

Because of Sentinel's delays and cost increases, in July 2010 the FBI issued a complete stop-work order.

An independent report estimated

that completing Sentinel under the current development approach would cost at least an additional \$351m on top of the \$405m already spent, and take another six years. In addition, the risks of working to an outdated specification now loomed. Some of the redesigned BPR processes were now six years old, technology had moved on, and there had been significant changes to the FBI's work processes that made them outmoded.

The Sentinel system that had been implemented so far was little used by FBI staff. Where it was, it merely resulted in duplication of effort, because data still had to be double-keyed into the old case management system. Confidence in the system was so low that its use was officially optional. Between July and August 2010, only 1% of new cases were handled using the new system.

By now, FBI agents should have had a case management system with workflows for managing their work – instead they were continuing with the same time consuming paper-based case management processes that had threatened the Oklahoma Bombing judicial process. The promise of electronic information sharing both within the

FBI and to and from other federal agencies to “connect the dots” between cases and suspects had still not been realized.

Attempt 2

- Sentinel initiated - starts on a waterfall path...
 - (Even more) monolithic BDUF design
 - Waterfall approach
 - ‘Group-think’ = over-optimistic status reports
 - Budget \$425m (1/4 on a PMO function!)
 - Big promises - BUT:
 - 2006: ‘(stuck on a pig’ delivery
 - 2007: no meaningful delivery
 - 2008: no meaningful delivery
 - 2009: Fulgham pins down milestones into smaller phases: no meaningful delivery
 - 2010: \$405m spent so far
 - no meaningful delivery
 - cost to complete estimate = \$351m
 - Lockheed Martin checked out
 - PMO disbanded
 - Agile approach adopted...

Sentinel Recovers Using Agile Approach

In September 2010, the FBI announced that it would take direct management of the development of Sentinel and use an Agile project approach:

“The FBI made a difficult but sensible decision to develop an alternative plan for completing Sentinel. We examined several options in detail, and selected an approach based on what is known as “Agile development” method. This approach will reduce our reliance on traditional contractors and allow for cost-savings by dealing directly with product experts.”

The migration of the 8.3m live, cold, and closed cases from ACS to the Sentinel database was in doubt. If it did not work, then ACS would have to remain alive for many years, and FBI agents would have to work with two separate case management systems at the same time. An automated facility to “join the dots” between new and cold cases would not exist.

The existing requirements were analysed, prioritized, and sequenced to focus on the most valuable requirements with the greatest benefits to agents and analysts.

Within a month, the FBI took direct control of development, removed all Lockheed Martin personnel from development work on the project, and started to supervise the sub-contractors directly. Fulgham reduced it from over 125 heads to a team of 55.

The project adopted the Scrum Method, with a Scrum Master coordinating the development team. This is a role different from that of a project manager in a waterfall project. The Scrum Master leads and enables the team, rather than ‘managing’ it. They empower a self-organizing team, rather than imposing structure on it.

The original, monolithic requirements document was modularized into 670 separate user stories. The team prioritized each user story in a product backlog, each one describing just one end-to-end process that the system needed to do.

Work now started to develop these user stories incrementally. Each cycle of work (or sprint) was two weeks long. At the end of every sprint, all testing had to be complete. The software had to be demonstrated to project stakeholders, and ready for deployment to users if required. 21 sprints were planned to develop all the user stories. Although there was concern that a continual churn of changes could ensue, the brevity of the 10-day sprints kept the danger of uncontrolled changes to what the FBI called “just 9 days of risk”. Previously, arguments over change control and scope creep took up much more time and effort than that.

At the start of each sprint, the development team identified which stories they were to develop during that sprint – these formed a work plan called the sprint backlog. At the end of each sprint, regardless of whether all work was complete, the development team had to test and

demonstrate the system. The team could only claim those stories that passed tests as completed. Where a test failed, that user story was placed onto the product backlog for rescheduling into the next or some other future sprint.

The amount of work required to develop each user story was initially estimated as a number of story points. These story points were a relative measure of difficulty and size. As work progressed, the team could see how fast they were working, and could start to calibrate their efficiency. After a few sprints, it became possible to forecast the rough timescales and start to plan the dates for incremental implementation. This was to be in two increments, with about half the user stories implemented in September 2011 and the rest in November.

However, concerns were raised about the first implementation being near to the tenth anniversary of 9/11 – potentially a time of heightened security. Therefore, the team carried out additional testing which showed that although Sentinel now had adequate functionality and usability, there were still concerns about its performance and availability. The

implementation was then planned to be phased in alongside the standard five-year refresh of computer hardware.

In the full year to 2011 (FY 2011) only 52% of the much reduced Agile development budget of \$32.6m had been spent to build 88% of the system. Jack Israel later commented on this success:

“Agile is not just a method or a process, it’s a way of being. You don’t do Agile. You are Agile. The FBI has arranged to loan their Scrum Master to other teams to get them trained. Increased transparency has kept stakeholders in sync. Further, stakeholders would modify their expectations, based on the increased visibility of the process.”

By June 2012, the revised technical targets had been met, and two releases had been achieved. There was a substantial increase in information sharing of case management information and a resolution of IT problems. \$46m had been spent on making progress over the last 12 months, and most importantly, agents were now using the system on real cases. 13,268 agents created 623 documents and made 92,546 searches in the first quarter 2012, against target of

11,000, 550 and 77,000 respectively. The key operational target of 13,200 leads per quarter in the full year 2012 was missed by a whisker (1% short of target).

In the first release, seven functional areas had now gone live, including allowing different user roles, storing attachments for sharing, and automatic routing, workflow and notifications of urgent actions needed on cases. The second release was a fully functional pilot of more functions at selected FBI field offices. User feedback was positive, and full Operational Capability was achieved in May 2012. We will not know the actual business benefits from the new operational processes until the new Sentinel system beds in after the final user of ACS logs off in 2013.

Attempt 3 – Agile!

- Sentinel recovers...
 - BDUF design broken into 670 user stories
 - Iterative development
 - 2 week long ‘sprints’
 - Self-organising teams involve the users
 - Actual spend \$114m

Conclusions

The cost of the initial failed Trilogy VCF project was \$170m. The cost of

the written off work of Phase 1 and 2 of Sentinel up to the firing of Lockheed Martin was \$427m.

The total spend of these failed attempts to replace the ACS system was \$597m and wasted 10 years. The Agile project, which is now delivering a solution, will only cost \$114m for a three-year long project.

Attempt 3 – Agile!

- Sentinel recovers...
 - BDUF design broken into 670 user stories
 - Iterative development
 - 2 week long ‘sprints’
 - Self-organising teams involve the users
 - Actual spend \$114m

What lessons are to be learned?

My research has identified nine leadership behaviours which enable Agile success. Two of these were particularly salient in the case of the recovery of the Sentinel project.

9 Agile Leadership Behaviours

- Satisfy the customer
- Harness change
- Be incremental
- Get the business and technical people together
- Create trust through leadership and process: ‘Light-Tight’ discipline
- Encourage face-to-face conversations
- Set targets and reward real progress towards a working solution
- Pursue simplicity, not complexity
- Give your teams the space they need to excel

Firstly: Being incremental. When Fulgham was appointed the FBI Chief Information Officer, he had to deal with a non-performing BDUF contract. By breaking the work into shorter sub-phases, each 3 months long, he placed a forensic spotlight on the supplier's failure to deliver. But this was not enough. It merely flushed out the symptoms of the problem - it did not solve it.

Fulgham had to take direct control of the project, and focus the team on very short increments of work - each fully integrated and rigorously tested, and approved by stakeholders. He helped the team to claw its way back to success - two weeks at a time.

The second Agile leadership behaviour that Fulgham exhibited, was 'light - tight' management discipline. Let me explain. 'Light - tight' management ensures light management of the team, whilst senior management follow tight disciplines.

The light management of the team is more than just delegating decisions. It is about making it clear that decision-making belongs by default with the experts working at the coal-face.

Tight management discipline was

exhibited at senior levels, and in implementation planning:

- In stakeholder engagement
- In project board supervision on spend and budgetary control
- And in planning and executing a realistic data cleansing, conversion and migration strategy.

So to go back to the two questions I posed at the beginning of my talk:

- The factors that are stopping us from using Agile approaches

more widely on government projects are related to key Agile leadership behaviours - not reliant on a search for the holy grail of a perfect Agile method.

- And secondly: We need case stories of success in government. I hope that I have provided you with one such story, of many from my research that makes a convincing case for wider use of Agile on large, mission-critical projects in government.

Thank you!

Brian Wernham's new book, "Agile Project Management for Government" was published this summer by Maitland and Strong.

brian.wernham@gmail.com