A Business Outcome Driven Approach for Governance of Digital Initiatives

ELMOHANNED MOHAMED, IBM Australia

In 2017, a major utility in Australia embarked on an organization-wide business transformation program to improve efficiency and increase adherence to regulatory obligations. This the story of how they evolved their governance structure from a traditional approach based on the iron triangle to an agile approach focused on business outcomes.

1. INTRODUCTION

Many organizations struggle to focus on achieving the business outcomes intended from their digital initiatives. Stuck with the fallacy of the perfect business case, and the divide between IT and Business, organizations opt to govern digital initiatives using the iron triangle: scope, budget and schedule. Business outcomes are dealt with as an afterthought evaluated after the end of the project or initiative.

This is the story of how a major energy utility in Australia adopted a business outcome driven approach to govern their digital initiatives using a mix of Agile and Lean Startup. While Agile institutionalises an iterative feedback loop to make sure the product can be built, Lean Startup's focus is "Should it be built?". The story covers the details of the first initiative where this approach has been implemented. I share the beginnings, the drivers and the different alternatives that were discussed before the approach was adopted. Governance plays an important role in setting up the model and hence is discussed with appropriate detail. Lessons learned and challenges are shared with a focus on what worked and what did not work.

2. BACKGROUND

Star Energy[1] is an Australian state government owned utility that builds, maintains and operates the electricity transmission and distribution networks for one of Australia's largest states. In business for more than 70 years, Star Energy's network connects more than one million customers over an area larger than the United Kingdom, with power lines spanning more than 100,000 km.

In 2017, Star Energy embarked on a business-wide transformation program driven by an evolving energy landscape, the threat of technology disruption and regulatory pressure for energy market reform. The program aimed to help Star Energy meet its ongoing regulatory requirements, safety obligations and improve readiness for an upcoming market reform. A transformation office was established, and a Big Four consultancy was contracted to help run the office. The office developed a transformation roadmap and 40+ initiatives were drafted. Each initiative crafted to achieve a set of business outcomes. The initiatives were grouped into streams: Asset operations, finance optimization, customer, etc. Most of the initiatives had a significant ICT component designed to achieve the target business outcome. A big technology consultancy was selected as the partner to assist in running the ICT program. At the time of this experience, the author, ElMohanned, was part of the ICT consultancy team. He was appointed to lead the delivery of the stream of finance initiatives. The CIO believed that Agile could help transform the ICT organization. Scrum, with standard 2-week sprints, was selected as the preferred methodology. In parallel with the program initiation, Agile coaches delivered an ICT wide agile training program.

3. GOVERNANCE FRAMEWORK

3.1 Lifecycle of an Initiative

A typical initiative would start as a one-pager under one of the approved transformation streams. The one-pager would outline on a high-level the problem, the solution options and the target business outcome. A team would then be formed to build a lightweight business case. The business case would expand on the one-pager, adding more details and estimates for cost and business benefits. As a rule, all initiatives had a maximum payback period of 2 years to be approved by the transformation office.
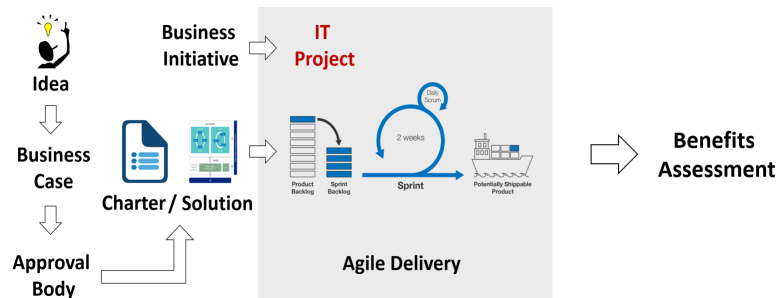

Figure 1. Lifecycle of an initiative

At this stage, the initiative team was formed of a business consultant responsible for delivering the lightweight business case and an ICT architect responsible for the IT side of things. Both would work with a business owner, who was accountable for achieving the target business outcomes. The business owner reported to the sponsor, usually the head of function, who was the ultimate owner of the budget and the ultimate decision maker. Once the business case was initially approved, a product owner and a project manager were assigned. The team worked together to deliver a project charter. Once approved, IT project could kick off. Agile delivery took over from this point on.

The typical team would be formed of the project manager, product owner, scrum master, architect, developers and testers. The project manager was focused on delivering the scope within the cost and schedule boundaries. The product owner was focused on building the best product possible. The scrum master focused on the ways of working.

3.2   Steering Committee

A steering committee was formed as soon as the project kicks off. The steering committee was generally formed of Project Sponsor, Business Owner, Project Manager, Product Owner, Head of Application Development, Workstream Manager (representative of the transformation office) and IT Program Manager. A weekly status report was presented to the steering committee. On a high level, the report covered the following: Progress (using actual / projected story point burn up charts or cumulative flow diagrams), Schedule (planned vs. expected calculated using the burn up projections), Cost (planned vs. actual), Risks, Issues and scope concerns. Albeit using agile tools, story points and burnup charts, the project governance was generally focused on the iron triangle: scope, cost and schedule.

3.3   Fixed Cost, Partially Variable Scope Model

In the spirit of Agile, Star Energy was keen to allow an element of variability in scope. The program strategy was to implement Fixed-Cost Variable-Scope projects / initiatives. For the CIO, this strategy was also a tool to make sure IT programs do not run over budget. The strategy did not appeal to business owners: If the scope is variable, how can I guarantee that the delivery team will deliver what I need to meet the target business outcomes? The transformation office and the CIO adopted a hybrid approach. All project charters were expected to have a backlog of Epics. Epics were labelled using a MoSCoW (Must, Should, Could and Would) prioritisation scheme. Project teams were expected to deliver at least 70% of the epics in the charter. In other words, 70% of the backlog was a Must-Have while 30% was flexible. This provided a middle ground where there is variability in scope and a minimum threshold for delivery.

3.4   Benefits Assessment

Benefits were assessed for achievement during project closure. The project closeout report listed the expected benefits of the project as per the business case and the business owner was expected to review and confirm if the benefits were achieved. The only other time benefits were mentioned during the project execution, was to use it as a negotiation technique. You would hear the sponsor or the business owner say: "If the budget increases by that much, I would need to go back to check my cost benefit analysis", or "If we de-scope this part, I will not be able to achieve the benefits of the initiative".

It was clear for everyone that the current benefit management approach was flawed in many ways. The benefits were not actively managed during the project and they were not objectively measured at the end of

the project. There was no mechanism to track benefits that would take a period, often years, after the project to realise. Realising the issue, the transformation office started an initiative focused on defining and setting up a benefits realisation framework for Star Energy.

3.5 The Business Initiative Versus the IT Project

It is interesting to analyse the transition from a Business Initiative to an IT project. As a business initiative, the focus was on the business outcomes. Business clearly was running the show with support from IT. IT solutions were discussed only from the perspective of achieving the business outcome and to perform cost benefit analysis. As the initiative shifted to become an IT project, the focus shifted to scope, cost and schedule. Business outcomes moved to become an afterthought. While business was still the owner, the relationship between business and IT was now a relation between customer (business) and service provider (IT). IT was running the project and was responsible of delivering the agreed scope within the approved budget and schedule.

The process of delivering the business case relied heavily on desktop research and well-crafted presentations reviewed and approved by experts from business and IT. Once approved, the business case was never questioned again until benefits are assessed after the project. On the other hand, the IT project was expected to run in an agile iterative manner with short feedback cycles to continuously prove delivery capability and acceptance by business and end users.

4. THE PROBLEM

The key problem for the CIO, and the IT organisation, was how to keep projects and programs within agreed budget and schedule while satisfying their business customers. The overall satisfaction among leaders of different business functions was low. Projects often extended beyond their anticipated timelines and budgets. Business complained they are not getting what they need when they need it. From the CIO's perspective, the problem's root cause was ineffective scope management. The IT organisation was unable to manage the expectations of business users that seemed to infinitely demand additional features. Not only did this impact the reputation of IT, but also led to continuously growing cost of the running IT at Star Energy. The CIO's response to this problem was twofold: Outsourcing and Agile.

Prior to starting the business transformation program, the CIO led a restructure program that outsourced the majority of IT and kept a minimal technical governance team within the organisation. Two IT consultancies were selected to run the IT support and IT delivery organisations. The CIO expected that the external consultancies would bring better project management discipline to the organisation. In addition, by setting up fixed price contracts, it was now the vendor's problem rather than the CIO's problem.

Agile was the second pillar of the CIO's plan. The CIO believed that short feedback cycles would bring early insights into scope and delivery issues. Having a dedicated business product owner assigned to the project would enable a better partnership between business and IT. It was anticipated that a healthy tension between the project manager, focused on delivery, and the product owner, focused on the product, would lead to a successful project. A training program was designed and delivered to all the selected product owners. The training focused on the concept of a Minimal Viable Product (MVP). The organisation's definition of MVP was the minimal product to achieve the target business outcome. Coupled with the 70% must-have rule, the IT organisation felt the new measures would get projects under proper control.

The new setup brought some success but did not come close to the desired outcome. Business owners did not appreciate the MVP concept. For them it was a nice way of giving them a product with the same price but fewer features. Product owners, trying to satisfy their stakeholders, continued to push for more features. More features meant more business value. In the absence of an objective assessment of business value, the decision if a specific feature should be part of the MVP was highly subjective. The 70% must-have rule for Epics was partially effective. The epics were too high level. Most of the "growth" in backlog happened at the detail level.

The outsourcing tactic did not work out as expected. Indeed, the vendors were good at managing projects to budget. This meant that pressure from product owners to build more, only led to a surge in change requests. The IT organisation found itself in continuous negotiations for change requests across the various initiatives. Projects' budgets and timelines continued to grow.

The new ways of working did bring some success in the form of more visibility to the problem. It was now clear to the CIO that the real problem was the lack of focus on business value. More features should not be equivalent to more business value. Business benefits defined in the business case were very loosely coupled to the epics and features. The benefits estimation approach used in the business case depended on ROI calculated

on the project level. It was assumed that the whole backlog would deliver the whole benefits. Surely some epics have more impact on business value than others. But how could benefits be broken down to the epic / feature level? And how would we know if we have built enough to achieve the outcome? Was there a way to assess business outcome during the project?

It was time to bring the Agile consultants back and ask the question: How can Star Energy embed business benefits into its delivery process? Agile consultants led several workshops with different parts of business and IT. At the end, they suggested using Business Value Points. In a nutshell, the whole project was assumed to deliver a 100 business value points. The product owner and business SMEs were asked to estimate the business value of the key epics / features as a percentage of the 100 value points. If a certain epic is estimated as 10 value points, this meant it was expected to deliver 10% of the benefit. Foundational must-have requirements—for example, log in, log out, etc.—were excluded from this assessment. Value points had absolutely no relation to story points. I theory, a story could be estimated from an effort / complexity perspective as 1 story point and deliver as 90% of the value.

The technique made some progress. It gave sponsors and stakeholders a sense of how much business value was delivered versus what was yet to deliver. It forced product owners to think harder about the breakdown of business value across Epics. It was a step forward. The CIO was not fully convinced with the approach though. The assumption that the business outcome is automatically delivered by delivering the software was intrinsically flawed. It also did not solve the problem of when to stop. How much of an epic was enough? On the other hand, the value points concept did not appeal very much to business and product owners. For them it was one more intangible concept— in addition to story points. It was subjective, easy to manipulate and shifted the discussion from business problems and improvements to "points".

5. A BUSINESS OUTCOME DRIVEN APPROACH

At the same time the business points pilot was concluding, the customer damage claims initiative was about to start under the finance optimisation workstream. The initiative lead team had prior experience in Lean Startup and they believed it was the answer to Star Energy's problem. Lean Startups excel at solving the type of problem that Star energy was facing. The lean startup methodology aims to eliminate wasteful practices in building new products. A key waste to eliminate was that of building product features that are not proven to add business value. Lean Startup achieves this by using an iterative process to achieve the so-called "Validated Learning". The iterative process starts by articulating the hypotheses underlying the business model of the startup. A build-measure-learn feedback loop is iteratively applied to validate each of the hypotheses. A minimal product is built to assess the validity of each hypothesis, measuring the outcome using an actionable leading metric, then applying the learning to the next iteration.

The concept of the leading metric is a key enabler of the process. Business benefits are usually lagging, for example, increased revenue or decreased cost. This makes them hard to measure as product development progresses. On the contrary, leading metrics are a change in users' behaviour that is assumed to be positively correlated with improving the lagging metric. By measuring the change in users' behaviours, the process enables validating business value as the product is iteratively released.

Although, as the name implies, Lean Startup is designed for startups, the initiative team believed it could be applied to Star Energy. The team presented their proposed approach to the CIO and the key stakeholders in IT and got their endorsement to pilot the new approach.

5.1 Customer Damage Claims Initiative—The Business Problem

The claim handling process was one of the pains of Star Energy. Unplanned power outages often occurred due to various reasons: Failure in equipment, accidents, storms, etc. If the power outage caused damage to customers, they can submit a damage claim and get paid for damage caused. Damage claims were submitted online to be received by a team of claim officers that would assess claim liability then decide to reject or accept and pay the claim. The legacy process of handling claims led to a substantial backlog of unhandled claims, which had accumulated and were increasing. Backlogs persisted despite Star Energy allocating extra short-term resources. The claim lifetime extended in some cases to over a year. The backlogs frequently led to complaints, including complaints to the Ombudsman and the Minister, undermining Star Energy's reputation and brand.

The corporate legacy Public Liability Claims Management System (PLCM) was highly unstable. It consisted of an online static html form that submitted claims via email to officers, and a backend system to manage claims. To make things worse, the primary and secondary experts in PLCM had left Star Energy. As a result,

even relatively simple but not infrequent issues, such as being unable to log into PLCM, takes more than a day to resolve. PLCM lacked integration to back office financial systems. Financial tractions had to be manually triggered on a separate system.

The claim workflow was as follows. A claim was submitted online. A claim officer received a claim via email and logged it in PLCM. The officer uploaded any received attachments to the corporate document management system and then performed a series of manual checks. Fraud checks to check if the claim was a fraud. Liability checks to check if Star Energy was liable for the claim. Planned outages checks to check if claim address / time fit within a planned outtage, hence Star energy was not liable. Sensitivity checks, as claims from people with personal injuries or medical institutions were considered sensitive. Finally, the officer either approved or rejected the claim. Officer submits result of his assessment to team lead for approval

Star Energy received up to 75 claims per week in times of a storm or bad weather that led to several unplanned outages. A back-of-the-envelope estimation of current cost per claim was $500. The claim officers team had two permanent employees and was augmented by three contractors to help manage the backlog.

The benefits stated in the business case were straightforward: Cost reduction and increased customer satisfaction. By reducing the effort needed to resolve a claim, the team size could go back to 2 claim officers rather than 5. This mapped to a saving of 60% in the direct cost of handling a claim. The lifetime of a claim would also reduce, resulting in better response times that should lead to happier customers.

5.2 The Proposed Solution

Three options were explored. Option 1 was to fix the existing PLCM system. This was quickly ruled out as the PLCM technology was unsupported and the experts had already left Star Energy. Option 2 was to purchase and implement a COTS claim management package. That option was ruled out as well as the cost of the licenses and implementation would have been much higher than the expected benefits. Option 3 was to build a custom claim management solution that fit exactly what Star Energy needed. Option 3 was selected.

Option 1 appealed to the initiative team. After all, the issues mentioned about the PLCM system did not seem unfixable. It would have probably been the lowest cost option. However, neither IT nor Business leadership seemed interested in exploring the option. From an IT perspective, this was an old technology, and resources were rare and expensive in the market. It would have been a bad investment to increase its lifetime further. It was preferable to move to a technology that Star Energy had abundance of skilled resources in. Business was also adamant that the PLCM was not fixable. It had to be replaced. It was never explicitly said, but if PLCM was not that bad, then the poor performance of the claims department could have been attributed to people rather than technology. This was not a hypothesis that business was willing to discuss.

Finally, option 3, building a custom claim management solution for Star energy, was selected. The solution consisted of four main components:

- Replace existing static online form with a smarter form providing better usability to end users and quick online checks. This should avoid some cases where the claim officer had to go back to the claim submitter requesting additional information, hence contributing to reduction for claim lifetime.
- Build an auto triage engine to auto approve / reject simple low-cost claims. To fit the criteria, the claim had to be submitted by a valid customer, within the time of unplanned outage, in areas impacted by the outage, and had no suspicion of fraud. A simple claim would be auto approved if it was for less the than the average cost of handling a claim, i.e. less than $500. It did not make sense for Star Energy to spend in assessing the claim more than the requested claim amount. It was cheaper to just approve them.
- Automate manual checks: For claims that could not be auto approved or rejected, the system would help the claim officer by automating some of the checks that the officer needed to do. The results of the auto analysis would be recorded against the claim record to be further reviewed by the claim officer.
- Replace backend system with a new claims management system that was more usable, more stable, had improved search capability, and integrated directly with the back-office finance system for payments.

5.3 Hypotheses and Leading Metrics

The first step to implement the new model was to articulate the hypotheses underlying the proposed solution. The process progressed as follows: Step 1: Define the hypothesis in the form: If we deliver feature X, we expect a specific change in user behaviour, that would lead to a positive impact on the final
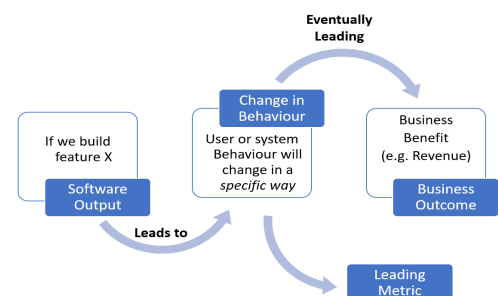
Figure 2. Hypothesis Driven Model

business outcome. Step 2: Define a leading metric that can measure the expected change in behaviour. Step 3: Design an MVP, defined as minimal set of features to achieve the desired change in behaviour outcome and allow validation of the hypothesis.

This process of articulating the hypotheses and designing MVPs, showed the first signs of success of the new model. Immediately more than half of the initially proposed features were knocked out. They were either completely unrelatable to the business outcome or very hard to justify their impact on the users' behaviour to the extent that would then achieve any of the hypotheses. The process ended with a list of MVPs, planned as sequential releases, each linked to a leading metric and final business outcome. Each row in the table below presented a hypothesis with the corresponding confidence level.

#	MVP	Expected Behaviour	Business Outcome	Leading Metric	Confidence
1	Online Form + Auto Triage Engine	Claims officers would be able to handle claims quicker.	- Saving of 400K / year - Higher customer satisfaction	Rate of claim closure to increase from 2 to 5 claims per FTE per week	High
2	Run Auto Triage Engine against backlog of claims	Reduce existing backlog by 60%	Higher customer satisfaction	None	Certain
3	Automate Manual checks	Claims officers would be able to handle claims quicker	- Saving of 200K / year - Higher customer satisfaction	Rate of claim closure to increase from 5 to 7 claims per FTE per week	Medium
4	Further enhancements to backend solution	Claims officers would be able to handle claims almost quicker	- Saving of 200K / year - Higher customer satisfaction	Rate of claim closure to increase from 7 to 10 claims per FTE per week	Low

Table 1. List of Hypotheses and MVPs. Numbers and metrics have been modified from reality keeping proportion.

Selecting the leading metrics was not as easy as it might seem. It took a while before everyone could get their heads around the concept. The business outcomes were focused on reduced cost/time of handling a claim, which was expected lead to faster response and higher customer satisfaction. These outcomes easily mapped to Effort Per Claim and Claim Lifetime metrics. Both metrics look good on paper but were not practical. Average effort per claim could indeed be used as a leading metric, but it was hard to measure. It was hard to segregate the exact effort spent on a claim given it was likely that claim officers work on multiple claims in parallel, each potentially at a different stage in its lifecycle. Average Claim lifetime would take time to change due to the existing backlog skewing the results. Alternatively, we could have used average claim lifetime for new claims only. But this would assume that new claims would take precedence over older claims. Finally, the team agreed on the rate of claim closure per week per FTE. This metric was logically related to claim cost and claim lifetime. It should improve almost immediately with new functionality releases, so it provided a good indicator if the team was on the right track.

5.4   The New Governance—Fixed Outcome, Fixed Cost, Variable Scope

The new governance model was focused on business outcomes. At the onset of the project, the target business outcome and the estimated cost to achieve it were agreed and fixed. The cost was based on estimates of the proposed solution, but also had to fit the cap of "how much Star Energy was willing to pay to achieve the business outcome". Scope and timeline were flexible. There was agreement on initial scope at a high level, but the product owner and the project team had the authority to change it to achieve business outcome.

The steering committee mandate changed. The focus of the steering committee was now to ensure the team could achieve the target business outcome within the target budget. This entailed three measurements / reporting domains: 1. Business outcome measured by leading metrics; 2. Cost measured by actual and forecast team cost versus budget; and 3. Delivery measured by throughput measures. Risks and Issues were discussed as normal. The steering committee that followed each release would focus on the resulting metrics. Based on the results, the committee made one of three decisions: Persist, Pivot or Stop. Persist if the metrics were in the right direction, or it was believed they would pick up. Pivot if the solution direction needed adjustment and the team thought a change in direction was necessary to achieve the business outcome. Stop if no more budget, or

it was clear to the committee that the business outcomes will not be achieved even with a pivot. Both pivot and stop decisions implied that the business case underlying assumptions were partially or totally invalid. In some cases, the metrics after the release were not indicative. They needed more time to mature. In these cases, the results of the metrics would continue to be monitored in parallel with working on the next release.

The new model instituted a different relationship between business and IT. Instead of the customer- service provider model, both now had shared responsibility of the outcome. While delivery was still led by IT, the overall outcome was led by business. Delivery and business outcomes were both questioned at each step of the process. Both business and IT partnered in solution and pivot decisions. The model presented a level of accountability on business during the project that did not exist before.

5.5   How the Project Progressed

Release one hypothesis was quite successful in terms of claim officer productivity. The release focused on the minimum needed to auto approve or reject simple low-cost claims. In addition, the release included some foundational work that served all releases. It was a tax that the first release had to pay. Luckily enough, a storm happened just a week after first release. This resulted in a surge of submitted claims. It was a good test to the logic behind the triage engine. The engine was able to auto resolve 60% of the claims. This helped the claim officers focusing their efforts on claims that needed further analysis and meant 60% of the customers got almost immediate response.

Despite the quick response, customer complaints surged unexpectedly. How did this happen? Well, customers who got immediate rejection complained that their requests were rejected without appropriate analysis. This meant the original assumption around impact of quick response on customer satisfaction was not 100% accurate. Customers whose claims were quickly approved were happier. Others whose claims got quickly rejected were not. Luckily the fix for that was reasonably quick. Timing and content of communicating the rejection were revisited. This was a minor pivot.

Release 2 did not need a leading metric. The result was almost certain. Quick pre-analysis on the existing backlog showed the number of claims that would be resolved by the triage engine. There was no need to validate the hypothesis. The team just had to go ahead and migrate the backlog and run them against the engine. And so, they did. The results were as expected. A good portion of the existing backlog was resolved.

There was reasonable confidence that release 3, automation of manual checks, would lead to better productivity (higher closure rate of claims). However, some people were sceptical if it would achieve the full target improvement of the metric. The implementation progressed as planned. Some of the automated checks turned out to be much harder than expected and the release was delayed. The steering committee questioned whether to continue the release or not, given the increase in cost and the uncertainty of benefits. The decision was to continue, and the checks were released with a delay in schedule and an increase in estimated cost.

The change in leading metrics after release 3 was not immediate. The business argued that the checks needed a bit for time to reflect on performance as the checks did not apply to all types of claims. The steering committee discussed whether to stop or continue with release 4 which had even more uncertain outcomes. The business pushed to continue arguing their confidence in the result and their concern that the full business outcome would not be achieved if the final release was not implemented. As the budget was not fully consumed, the steering committee decided to continue until the full budget was consumed. The final set of features was released.

Near the end of release 4, the business attitude was becoming less transparent regarding the metrics. Numbers were not openly discussed as often and were replaced by generic messages around productivity. With the delivery of release 4, the project was closed within budget but with a delay in schedule. The official message from business was that the productivity goals were achieved, and the target business outcome was delivered. The project was closed on that note.

6.   LESSONS LEARNED

The new model succeeded in shifting the focus to business outcome rather than scope. At many instances during the project, scope discussions were quickly resolved by questioning the potential impact of a requested feature or enhancement on the users' behaviour and the related metrics. It was argued at the beginning that startups were different from corporations and hence Lean Startup was not the right approach. While it is true that startups are inherently different from corporations, Lean Startup still presented a solution to star Energy's problem: How to measure business outcome and reduce the waste of building features not aligned to business outcome.

The first challenge the project met was identifying the hypotheses and defining the leading metrics. While the concept was easy to explain, it was hard to implement. In hindsight, the hypotheses and metrics seemed logical. In reality, it took a few iterations before the team was able to articulate the hypotheses and define good leading metrics. The importance of this could not be underestimated. The whole delivery model was built on top of the hypotheses and target metric. The key questions to help tease out the hypotheses and metrics are the questions around the change in end user behaviour. What change in end user behaviour do we expect as a result of a certain feature? How would it lead to the target business outcomes? Can we baseline the behaviour and measure it as it changes?

Setting the target change in leading metric proved to be non-trivial as well. Would the claim officer's productivity rise to 7 or 10 claims per week? It was hard to tell. In some cases, it was possible to analytically estimate the target. In other cases, it was a pure leap of faith, or even wishful thinking. This is fine as long as there are mini steps toward the goal that can be measured, and the metric can be assessed after each step.

As discussions to scale the approach across Star Energy started, some people argued the approach was not scalable. Their rationale was it was hard and sometimes impossible to define leading metrics for other initiatives. While I acknowledge it is harder in some cases to map business outcomes to leading metrics, I would argue that inability to define leading metrics reflects a lack of understanding of how the target features will lead to a change in users' behaviour and often how the change in behaviour will help achieve the business outcome. This makes the whole validity of the business case questionable.

One thing that was clear during the project is the presence of implicit goals that are not explicitly articulated. For instance, IT had a goal of getting rid of the old technology to be able to consolidate technologies supported. Business had an implicit goal of the claim officers' satisfaction. They had been under pressure for some time and it was important for business to make sure they were satisfied with the new solution. Ideally such goals should be articulated, quantified and measured appropriately. Otherwise you run the risk of not achieving them.

Although the project implemented a new governance model focused on business outcomes, members of the steering committee often switched to the old ways of thinking. This was apparent especially during the delivery delays that faced the project during release 3. The initiative team had to continuously remind the governance body of the governance model in place. That scope is not the target, rather business outcome is. This reflected the need of a wider organisation change program to institutionalise the new model and the new ways of working.

As mentioned earlier, the new governance model changed the relationship between Business and IT from service provider model to a partnership of results. Although, this proved quite successful, it did come with some unexpected side effects. The business was not used to being put under the spotlight in this way and felt the pressure to prove the business case assumptions were valid. The pressure was quite visible especially after release 3 and during release 4 when the leading metrics progress slowed down. It resulted in the lack of transparency from business side near the end of the project. This can be attributed to the lack of culture of experimentation in the organisation. The sponsor and business owner should have been comfortable with letting the project validate or invalidate the assumptions. One thing that aggravated the issue was that the project ran as part of a big transformation that led in some cases to people losing their jobs. This created a defensive finger pointing culture. No one wanted to be blamed or found guilty. This could explain the business lack of transparency near the nerd of the project. This again reflects the importance of a safety culture that does not condemn failure, but rather cherishes learning.

Another learning for the team was that portions of the backlog did not relate directly to the business outcome. Some of it was regulatory requirements that had to be part of the system, like availability of audit logs for all steps of claims analysis. Others were security requirements, technical quality requirements, or foundational components that had to be built to support other functions. It is important to understand the source of these requirements and be clear how much should be implemented. Ideally, they should be implemented after the core assumptions are validated to make sure the investment is not wasted. This is not always possible.

Last but not least, the importance of technical delivery excellence cannot be overemphasised. While the model is focused on business outcomes, nothing will be achieved without successful delivery. The presence of a well-oiled DevOps machine is essential to facilitate continuous delivery and hence enable continuous feedback on the business outcomes. Architecture might need to be modified slightly to enable easily measuring the leading metrics. This needs to be planned early in the process. Overall, the project was a powerful learning experience for the project and Star Energy.