

The Art of Splitting an ART

CHRISTIAN ESKE BRUUN, Release Train Engineer, Ikano Bank
JESSICA FRASER-DARLING, Release Train Engineer, Ikano Bank
SVEN CARLSTEDT, Agile Coach, Ikano Bank

When organisations use the Scaled Agile Framework® [*Scaled Agile Framework*] they may come to the point where their Agile (ART) Release Trains have grown too big. This report is about our journey of splitting an ART using ideas from Team Topologies and the principle of self-selection.

1. INTRODUCTION

Scaled Agile Framework® (SAFe®) [*Scaled Agile Framework*] is a popular choice when organisations want to start an agile transformation. Many organisations start by forming one or a few Agile Release Trains (ARTs). As time passes, more ARTs may be launched and/or more teams are added to existing ARTs. Another scenario is that an ART may grow so large that it needs to be split; imagine running a Scrum of Scrums or trying to demo all the work each sprint with 25 teams in your ART. Then the question arises:

“How do you split an ART?”

On paper it can sound easy; You decide the number of teams in each ART and move on. What we realised pretty quickly was that this was not easy, not easy at all. Many questions quickly arose: What are the responsibilities of the new ARTs? How do you minimise dependencies between ARTs? The original ART and the architecture are most likely not designed with a potential split in mind and how do we ensure buy-in from the members of the original ART?

When looking at the literature and online research, you can find plenty of results on how to start an ART but very few on how to split one. In this experience report we want to tell our story of how we split one large ART of 16 teams into 3 ARTs plus Shared Services using ideas from Team Topologies [*Skelton and Pais*] and the philosophy of self-selection.

About Ikano Bank AB (publ)

Ikano Bank creates possibilities for better living by offering simple, fair and affordable services, enabling a healthy economy for the many people. Our offer includes savings and loan products for consumers, sales support services for retailers, and leasing and factoring solutions for businesses.

We conduct business in Sweden, Denmark, Finland, Norway, the UK, Poland, Germany and Austria. Ikano Bank is a part of the Ikano Group who owns 51% of the Bank. Ingka Group, a strategic partner in the IKEA franchise system owns the remaining 49% of Ikano Bank. Ikano Bank's head office is located in Malmö, Sweden and the company is registered in Älmhult, Sweden where the business was once founded.

2. BACKGROUND

We decided 3 years ago to transform our digital solutions and technologies as part of our journey to become one sustainable bank. The first step was to create a series of digital finance solutions on this new technology using standardised modules and more flexible front ends. The first target with this solution was to deliver a new product for the UK market. It would be an interest-free retail finance offering where people would have the opportunity to apply instantly using a mobile app or browser. The person applying would - if approved - be issued with a virtual card which they could use to spend in store or online straight away. The main target was IKEA customers.

Christian Eske Bruun, Ikano Bank, Hyllie Boulevard 27, 215 32 Malmö; email: eskebruun@gmail.com
Jess Fraser-Darling, Ikano Bank, Hyllie Boulevard 27, 215 32 Malmö; email: jessfd@hotmail.co.uk
Sven Carlstedt, Ikano Bank, Hyllie Boulevard 27, 215 32 Malmö; email: sven.carlstedt@meqify.se, MEQIFY AB
Copyright 2021 is held by the author(s).

Together with this new product, we chose to begin our agile transformation using SAFe® [Scaled Agile Framework]. One ART was formed and launched in January 2020 which consisted of 2 teams and around 25 people. Almost two years later we had grown to 16 teams and close to 200 team members across at least 7 countries. The teams forming this ART were organised as component teams, a decision made since a completely new architecture was being built from scratch. It resulted in each team being responsible for a few software modules in line with the architecture for transforming into a new banking infrastructure. What this meant in reality was that almost no teams could deliver value to the end user independently.

By the spring of 2021 it started to become obvious that our single ART had grown too large; our Agile ceremonies were a struggle to maintain within a reasonable timebox, the number of team members and teams that needed to coordinate became complex and created bottlenecks and our flow of delivery and lead time was decreasing. It was decided by the leadership team and management that something needed to happen, and the recommendation was to investigate splitting the ART.

The task to split the ART was given to the two RTEs (Release Train Engineers) and one of the agile coaches - the authors of this paper.

3. OUR STORY

When we began our journey to split the ART, we chose to use Kotter's 8 step model [G. Hole] as a guiding principle (Figure 1)



Figure 1. Kotter's 8 step model

We found that it was not a strict linear path through the eight steps; we were working on more than one step at a time and even found ourselves going back and forth between steps. In the following section we will describe our story, thoughts and what we did in each of the steps of the model as well as some of our learnings.

3.1 Part 1: Pre-Split

Step 1 – Create Urgency

Early on in 2021, it became apparent that our ART (The “Pay ART”) had become too big and cumbersome. The lead time for features was growing and people were starting to get lost in the many teams. There were also signs that developing a component sometimes took precedence over getting overall customer value out. The component organisations of the Pay Art also meant that features needed to pass through several teams before being delivered to the end user. In turn, lead time and WIP (Work in Process) were growing and Product Owners (POs) were spending a lot of their time as feature coordinators. Ikano Bank had an ambitious strategy to launch more products in multiple markets which meant during the summer of 2021 the ART would need to ramp up with even more teams.

Leadership at many levels started to notice that something needed to change. An extra RTE was hired and there started to become a buzz in the teams that the Pay ART would be split. At the Pay ART’s last PI (Program Increment) Planning the Global Head of Product went on stage to explain how important the product delivered by the Pay ART was to Ikano Bank and how we now wish to become faster in our time to market. Now the urgency for a change was starting to form in all parts of the Pay ART.

Learning: In hindsight it is very easy to see the importance that the Global Head of Product’s message had on the ART. Having a senior manager so clearly and explicitly stating that something needed to happen was alpha and omega to create our sense of urgency. We should have made this urgency even more clear for the members of

the teams. “Why is it so important to get our lead time down?” was a question we did not answer well enough for the team members. Providing this clarity would have made later steps easier to overcome any resistance to the change.

Urgency needs to be established across the entire organisation - from leadership to developers

Step 2 – Form a powerful coalition

From the beginning of the change, it was clear that this initiative would be led by the RTEs with dedicated support from the agile coach. The magnitude of the change meant that we needed the support from several key stakeholders in Ikano Bank. Some of these key stakeholders were the managers of the Product Owners and the Architects.

Ikano Bank has a very strong community of architects who really have put their heart and soul into creating a solid modern architecture. The split of the ART would impact the architecture and the way POs and Architects work. By forming a strong cooperation with the managers of POs and architects, we were able to use them as mediators and facilitators in the change management process. At the same time, they were able to provide valuable input to guide us towards the solution for how we would ultimately split the ART.

Learning: The key to our success here was that we built a relationship based on trust where we could ask for help and get feedback, something which is reflective of one of our company values “Working Together”. Whilst we identified some of the key persons we would need help from, one coalition we should have focused more on early in the change process was with our lead/key engineers. This would have made some of the change management within the teams easier and ensured that we got a faster buy-in from the team members.

Step 3 – Create a vision for change

It is no surprise that when splitting an ART a lot of questions pop up and need to be answered. We decided that the change we were making was so big that we couldn’t create the entire vision in one go. We therefore went for an iterative approach, starting very far out and then zooming in. For every iteration we asked for feedback and made adjustments.

The ARTs on an overall level

The first step was to define the boundaries for each of the ARTs. We tried to answer the question: **“Why does this ART exist?”**

Product Management performed an analysis to understand how they saw the future ARTs supporting existing and future products and the features within those products. We decided to utilise a splitting strategy based on the expected user journey of our products. By following the natural steps in the user journey, the existing ART was split into 3 new ARTs and a Shared Services function (see Figure 2) all of which fell under an umbrella “Solution Train”. A Solution Train set up is not something that is a necessity when splitting an ART, we use this as a loose organisation structure rather than a true SAFe® [Scaled Agile Framework] Solution Train.

A very important thing to underline is that in this split it was clearly defined where one ART’s responsibility started and where the next one would take over.

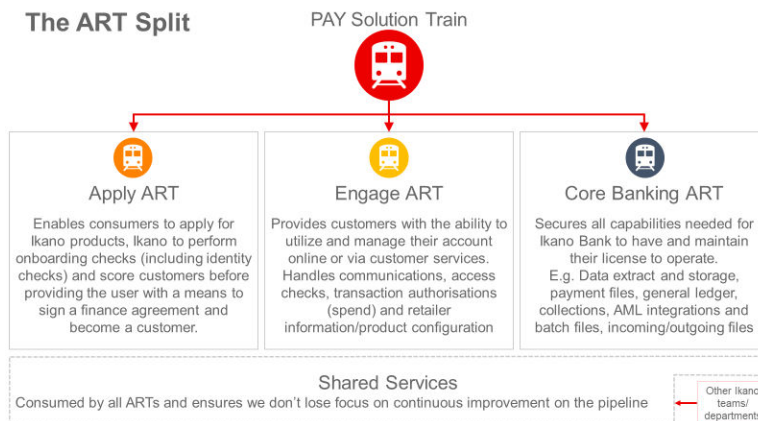


Figure 2. The Overall ARTs

When this analysis was done, we simulated the set-up using features we knew were already in the future roadmap to see how they would fit into this overall ART split. This exercise revealed that the overall structure of the ART was good enough but that our very heavy component setup would create numerous inter-team dependencies. Armed with this knowledge, we knew that we needed an approach to change our component setup if we were to be able to split the ART.

Learning: We used a simulation to provide an initial validation of our ART hypothesis, testing the flow of a few features and getting fast feedback. This enabled us to understand how much the architecture would affect the ART split and consider how we would mitigate this.

Defining the ART leadership teams

Once the overall structure of the ARTs was in place it was important to define the leadership of these ARTs. In SAFE® [Scaled Agile Framework] an ART is headed up by an RTE, a Product Manager and a System Architect.

We believed that if the ART leadership teams were founded later then the Product Manager and System Architect might not buy-in to the way the ARTs were to be split. It was therefore essential to get these leadership teams formed so we flagged to key stakeholders that this should happen sooner rather than later.

Stakeholders from Product Management and Architecture came with a list of candidates and those candidates were involved in discussing which ART they would be able to contribute best to. The PM and SA from each ART were then given the task to select the RTE they found most suited to their ART.

Team topologies

The original ART was organised around components. This meant that every team was responsible for a few modules of code and in order to deliver end user value a feature would have to be worked on by multiple teams. It was architecturally concluded that some of the components needed to be in more than one of the ARTs. Being so strongly organised around components also meant that features would have many dependencies i.e. handovers between teams.

To address these challenges, teams with increased cross functional skills were seen as a way forward. Since the organisation had no experience of a team delivering end-to-end value to the users, a cross functional pilot team was tried out during one PI to start building our knowledge of the opportunities and challenges.

Encouraged by positive feedback from this team, we studied some external material on leadership and since it seemed a good fit with our situation, we agreed to try the approach described in the book "Team Topologies" [Skelton and Pais]. After a vote by the newly formed ART leadership teams and Product Owner community, it was decided to go ahead with this suggestion.

Learning: It was good to have a reference model that gave us a common language, naming the different types of teams and meaning the people and stakeholders could find books and articles to learn more.

ART skeleton

At this point in time, we had a vision for the overall ART split, and what each of the ARTs should cover from a business perspective (see Figure 3).

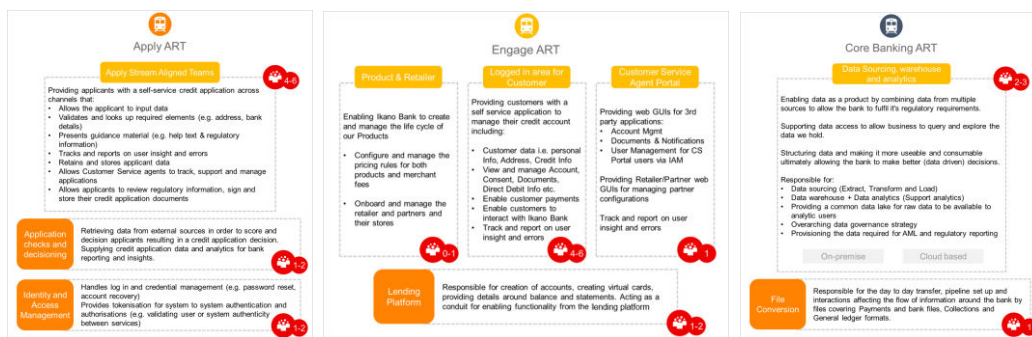


Figure 3. The ART Skeleton

Increased focus on value and delivery is something we wanted to achieve. We had formed leadership teams for each of the ARTs and the decision to use team topologies created a mutual understanding and language of

the type of teams the ART could consist of. Now we, together with the POs and Scrum Masters (SM)s, decided on how many teams and the type of team per ART. The vision of the change was completed.

Step 4 – Communicate the Vision

We always knew that communication would be vital if we were to succeed. This meant that no matter which step we were on in the Kotter model, we would communicate extensively. From the beginning, we decided to be open and transparent about the change, how far we have come and what the next steps would be. We were also open about the fact that we did not have all the answers upfront. To ease the communication, we wrote a series of blog posts or “ARTicles” about the change for all Pay ART members and interested parties to view and comment upon. This was the list of ARTicles:

- **The Why** This was the first article which explained why we needed this coming change. We linked back to the statement from Global Head of Product’s message about needing to get to market faster.
- **Target timeline for the split** To be transparent, set expectations and remove fear.
- **The structure of the New ARTs** Explaining the overall structure of the new ARTs and how they were designed around the end user journey.
- **Team Topologies** and how we would use this to form the new setup of teams. It was here we started introducing new terminologies which we would use going forward.
- **Introducing the Leadership teams** Providing an explanation of what an ART Leadership team was and who were the people in those roles in each of the new ARTs.
- **Updated timeline** Our original timeline started to slide so we were very transparent about this explaining why the timeline was sliding and our best guess on the revised one.
- **The ART skeleton** Here we explained the number of teams in each ART and which team type they would be according to Team Topologies
- **The self-selection** This article introduced the concept of self-selection and how we would do it as well as the playing field we would be working within.
- **Adjustments after self-selection** After the self-selection we performed some inspect and adapt on the output, this resulted in some changes being made and these were communicated to all the ARTs to ensure transparency.

In addition to the ARTicles, we held ART-wide townhalls where we explained what was going to happen and gave team members the opportunity to raise questions and risks with an online form for posting their inputs and thoughts. With the change we also decided to adjust the terminology used in the ART. We started to very actively use words and terms like ‘Stream aligned team’ from Team Topologies. This was a very deliberate tactic as we wanted to ensure people understood that we were starting something new and therefore moving away from the component setup we had before. This change in terminology was the correct decision, it gave a common new language which helped the change and helped foster an idea that there might be ways to be organised other than the original one in the old PAY ART.

3.2 Part 2: The Split

Step 5 – Empower actions

With the ART skeletons in place it was time to move closer and start forming the new teams which the new ARTs would consist of. Forming 25+ teams is easier said than done. To truly engage with the team members and allow them to be a part of the change, a self-selection approach with constraints would be the way forward. We also knew that the collective intelligence of the PAY ART was by far larger than that of the team driving the change. If we were to achieve new ARTs which would truly work and not only make sense on paper, we would need the wisdom of the team members.

We decided to split the self-selection into two parts:

- One where the POs and Scrum Masters chose which ART and which type of team they would like to have.
- The second a self-selection for all team members - First they chose the ART they felt they could add the most value to and then the team which they felt would be the best fit for them given their competence and the needs of the team.

Product Owner and Scrum Master self-selection

There were several reasons why we went for the two-stage approach. We needed to form a strong coalition with the POs and SMs and we needed their knowledge. They were working daily in the teams delivering our product and they would be able to spot any weaknesses in our thinking before we took it further to all team members. Before the workshop we formulated boundaries for the PO/SM self-selection. The boundaries, called playing field, included items such as:

- A PO/SM can only be in one ART
- Teams should be able to deliver value
- Team size limits
- The teams should follow the ideas from Team Topologies [*Skelton and Pais*]

Together with the forthcoming ART leadership teams, we formulated a suggestion of how many teams would make up each ART and if they were 'stream aligned' or 'complicated subsystems'.

We began the workshop by repeating the why, the terminology we would be using and then presenting the suggestion to the type and number of teams. At this point a lot of discussion erupted. The POs had serious concerns about the suggestion and whether it would fly, especially how we would handle the fact that our product exists with several front ends. We deviated from the plan in order to let the POs explore several different paths. The workshop grew from one day to two. After day one we had three competing suggestions so we opened up for a vote on which one we should use. Once we had agreed on which suggestion to use, we asked the POs and SMs to self-select and choose which ART and which team they would like to join. After the initial selections were made, we reviewed the result and where necessary made small adjustments. This self-selection process went surprisingly smooth.

Learnings: The playing field which we had defined upfront was a key component in our facilitation. Without it, the self-selection would have been too chaotic, and we would have risked ending up with a solution which we as change agents would have found hard to support. Several times during the workshop we revisited and referred back to the playing field.

"No plan survives first contact with reality" and when this happens do not be too proud to change your plan. The fact that we took the time to explore and address people's concerns meant we achieved a better result and a stronger support from the PO and SM community.

The final learning for this type of workshop is that being a team of facilitators gives an enormous advantage. We were three facilitators which meant we had a much easier time to improvise in a controlled fashion. We had an open chat going on behind the scenes where we could plan the next steps and support each other.

Team members self-selection

We had the team types, number of teams per ART and the PO and SMs agreed - It was now time to fill the teams with members.

In line with our company value "On Fair Terms", we believed that the team member self-selection should be performed by both our employees and consultants. This would allow us to gather support, buy in and gain the most potential. The collective intelligence would create better teams and by having all members deciding which team they would join, we hoped to mitigate dissatisfaction and demotivation. The concept of self-selection promotes people to form teams with those they have an easier time working with, thereby making the road to building performing teams easier.

We encourage ourselves to dare to be different and this was no exception, self-selection for around 200 people distributed across several countries is a daunting task. The workshop, just like the one for the PO and SMs, would be conducted online using digital collaboration tools. After proving valuable before, we again created a playing field (see Figure 4) which gave boundaries during the self-selecting process. It was important for people who could not join to appoint a proxy who could make decisions on their behalf and who would know the person's preference and wishes. The SMs had helped collect some information about every team member in advance. The information included: Role (tester/engineer etc.), Experience, #PIs in Ikano Bank, previous team(s). The information was to give other people an idea about each individual and his/her background.

We started the workshop by explaining the why, what would happen and the playing field. We also outlined the new ARTs; the number of teams and types (stream aligned, complicated sub-system, enabler or platform team) per ART articulating some of the logical boundaries within the ART skeletons.

The self-selection was split into two parts.

1. People selecting which ART they would like to join.
2. People selecting which specific team they wished to join.

As a check in activity, all attendees would locate a post-it with their name which included the information collected by the SMs. If a person was a proxy for someone, they would state both the team member they represented and his/her own name. This exercise made it much easier for us as facilitators to have an overview and ensure everyone in the workshop was active and present from the outset.

We asked team members to perform the first self-selection - Select which ART they would like to join (excluding the POs and SMs). When they selected which ART, they would also need to indicate if they were front end specialist, back end specialist or both. There was also an option to choose that you did not have any preference regarding which ART to join. This self-selection gave a starting point to judge if we had a good enough distribution between the ARTs. If we had seen that too few people wanted to join one ART, we would have stopped and inspected why.

In the second part of the self-section people were asked to choose which team within an ART they wanted to join. People were not restricted to their initial ART selection, but in reality, very few people changed their ART choice.

The online whiteboard was designed so each team had a space (see figure 5). In this space you could see the PO and SM for the team. The teams were also given an arbitrary name. In our case we chose to name the teams after fruits, this was to break any association with our old component team structure.

In the team you could see that there was space for a total of 7 engineers. There were an additional 3 spots for people who wanted to join. We made it clear that after the self-selection we would inspect and adapt each team. This was not a race for you to select a team so whether you had a waitlist list spot or a team member spot it made no difference. If there were too many people in one particular team, we would talk about it afterwards. We limited the self-selection to 15 minutes, during this time some people moved teams several times but eventually the map settled.



Figure 4. The Playing Field

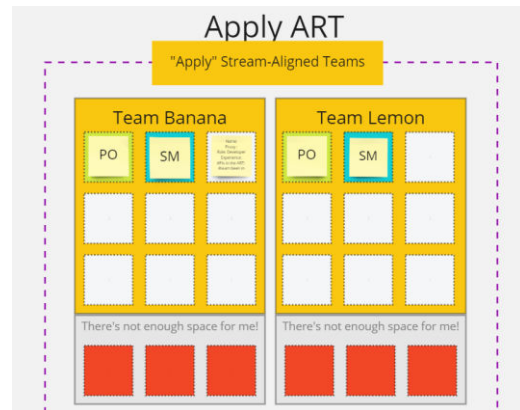


Figure 5. Example Team Space

After this first round of self-selection people began to tire and our time schedule had slipped. We decided to pause and continue the following day which we now believe was the right call. We needed people to be engaged and not just wear them down to the point where they would accept anything to get out of the meeting.

We started the following day with a recap of the teams so far. Using the collective intelligence, we verified that the team was in line with the playing field, if it would function in practice, if there were too many people wishing to join and how we would solve any concerns. Finally, we validated all three ARTs together to ensure they would function. During this team review, several topics and concerns were raised, these were documented and we took them away as actions we needed to work with.

Learning: Preparation is very important, but you need to be pragmatic. It was incredible to see the self-selection process in action and that it actually also works in a large-scale setup. The facilitation of the self-selection workshop was harder and took a lot more time than expected but it proved to be incredibly valuable and was appreciated from the team members' side.

Our practical learnings from these workshops is that it worked best to perform them online using digital tools. Our teams are distributed around the world and it was not possible to fly everyone in. It would have been very difficult to have conducted these workshops in person, even just finding a room which can support close to 200 people would have been tough. We as facilitators would in a physical setting easily have lost the overview, this way it was easy to zoom out in our digital whiteboard to get the holistic picture.

3.3 Part 3: Post-Split

Step 6 – Create Quick Wins

After self-selection the PO and SM started working with these new teams and a few days after the selection a few issues surfaced which meant we needed to make some minor adjustments. These included gaps in competence within particular teams or an imbalance of experience throughout the teams in an ART.

Learnings: We learnt it was much better to make these minor adjustments sooner rather than later. It would create a larger commitment from the teams as they felt we were listening to them and recognising that this is a long change process.

After the split we got a little caught in the trap of jumping directly into execution mode. We didn't take enough time to celebrate the wins we saw from the split. There was too much focus on all the things that now did not work. We could have been better at finding the quick wins and showing them off which would most likely have eased some of the resistance we met after the split.

Step 7 – Build on the change

We have now executed one PI in the new ARTs and as we are writing we are finishing the second PI. When we made the split, we didn't add or remove any members of the ARTs. During execution of the first PI it became very clear that the split had exposed some problems which before had been hidden under the surface and that we still needed to inspect and adapt more on the teams' setup. The challenges we saw/see are:

- Stream aligned teams do not have all competencies needed in the team:
In some of the stream aligned teams we had to add new people with certain skills. We are still norming and storming in some of the teams. The SMs are working to ensure we do not create teams within teams.
- Single person dependencies:
The split revealed some single person dependencies. We are now actively working to ensure more redundancy and T-shaping in order to minimise operational risk but also to make us more effective.
- Ownership of the code when people are now in different ARTs:
Before the split, the code module was owned by a single component team. Now this ownership is sometimes distributed across several teams and ARTs. This is still a challenge we are working to find a good solution for.
- Features need to describe value realised by the new ARTs and team setup:
We have realised that some of our features are still written for component teams. These features are hard to execute on for a stream aligned team without having many external dependencies. The POs and product management team are now focusing on writing Features following standards which are a better fit for our new setup e.g. the INVEST scheme.

Step 8 – Make it stick

We are currently in this stage of Kotter's 8 step model and we are starting to see the gains and benefits of the change we have made. The teams are settling into and forming their new identities with a renewed focus on providing value to our customers.

3.4 Part 4: Reflections

One thing we often have asked ourselves is whether we did too much of a big bang. Would we have achieved the same results if we had taken a more iterative approach? These are questions which are hard to answer.

A more iterative approach would have allowed us to gain greater learnings along the way and given us time to perform more extensive change management. On the other hand, the constraints from our old ART and the architecture were so many and so intertwined that we were perhaps better to just cut the 'Gordian' knot.

One thing is very clear, it was the correct choice to split and descale the Pay ART. We have surfaced underlying problems which were harder to identify in a larger ART. The split has created ARTs which are more focused and the teams in them feel a stronger belonging. We are convinced that with time the split will also provide much higher effectiveness.

4. ACKNOWLEDGEMENTS

There are several people we would like to acknowledge. The members in the new ARTs for trusting in us. The Scrum Masters and Product Owners for honest and fruitful feedback, support and help. Our managers for trusting us and helping us to handle management expectations. Chapter Leads for the Product Owners and Architects. Our Domain Engineering Leads & our Solution Train Engineer for the heavy-going re-organisation work in the background to align contracts, budgets and filling competence gaps identified in the split. Finally, we wish to thank Steve Adolph, our shepherd for his help and moral support, providing us with great feedback and of course encouraging/nagging us to write this report.

REFERENCES

[*Scaled Agile Framework*] Scaled Agile Framework® (SAFe®) 2022, © Scaled Agile, Inc., <https://www.scaledagileframework.com/>
(Accessed May 2022)

[*Skelton and Pais*] Matthew Skelton and Manuel Pais. "Team Topologies" IT Revolution, 1st edition, 2019

[*G. Hole*] The Management Philosopher - Dr. Glenn Hole, "How we used Kotter's eight step model for change and succeeded within a turnaround case of a Nordic BPO suppliers", Online Blog, April 2016, <http://www.dr-glennhole.org/how-we-used-kotters-eight-step-model-for-change-and-succeed-within-a-turnaround-case-of-a-nordic-bpo-suppliers/>