Paper LI07

Landing Moonshots

Ross Farrugia, Roche, UK

ABSTRACT

In 2024, at Roche/Genentech we achieved an ambitious goal of moving 90% of our teams and 1,000+ data scientists to a completely new platform and way of working with our inhouse built OCEAN system and git-based workflows. When the target was originally set few believed it could be achieved, but we ended up hitting 91%! Often such goals can feel like an added weight and burden of expectation and pressure on people. However, this talk will discuss the strategies used to mobilise and embed such large-scale change, through using the moonshot goal as a catalyst for the “in it together” spirit and collaboration that was needed from all involved. The talk will also include our 2025 equivalent around the adoption of R as our primary backbone language for clinical submissions via the usage of open source tools and AI/LLM solutions, and how we’re gamifying this movement to promote friendly competition.

STAR-GAZING: BACKGROUND

Coming towards the end of 2023, at Roche/Genentech we had recently launched a new, in-house built platform intended for over 1,000 data scientists. This system introduced a whole new git-based workflow and opened up multi-lingual potential under one roof.

Early adoption feedback from a small handful of teams was promising, but it was clear we had a long way to go. Our people were in equal measures excited and anxious of the scale of the change standing before us. In past such junctures we had often chosen caution with steady and gradual adoption of the new solutions over many years, knowing that at our large company such foundational change can often be akin to trying to steer the Titanic.

This time we instead channelled our more ambitious childhood spirit of the past and aimed for the stars, by setting a moonshot goal. Jointly with leadership, we decided to set a department-wide target: **In the next year, we will onboard 90% of all teams!**

What were we thinking?! Our platform was still being built iteratively, and our workflows were still developing and there were gaps galore. Yet, we needed a target that everyone could rally behind as one. We wanted a catalyst to accelerate this transformational journey, and most importantly, we truly believed in the benefits of the direction we were moving towards and the potential we would unlock for our users.

In this paper, I’ll discuss the very human-centered strategies we used to target this moonshot, and less on the tech stack specifics (though for the curious, there’s a great public webinar available on this: https://www.youtube.com/watch?v=nqJsLSLd39A).

INITIATING THE SPACE RACE: SETTING THE 'WHY'

John F. Kennedy famously said of the original moonshot, *"We choose to go to the Moon... not because they are easy, but because they are hard"*. Back in the 1960s, the political motivation was to win the space race, justifying spending over 4% of the entire federal budget on NASA - a huge figure! That famous line helped rally the public behind the goal.

When we set our own moonshot, we had to clearly articulate *our* 'why' and set a vision that people at all levels could get behind, even though we knew the journey would be bumpy.

We branded the move as more than just a platform lift-and-shift. It was an opportunity to revolutionize the way we had long worked by introducing new technology, languages, and workflows. When users saw it as more than just a cost-cutting exercise, they started to lean in and co-create their own future. We branded it "our NextGen" - emphasis on "our". Building this platform in-house gave us the ability to include users right from the early design stages, allowing us to directly address previous frustrations like fragmented architecture or a lack of adherence to FAIR principles. Ownership and a sense of purpose were key.

BUILDING MISSION CONTROL: THE ADOPTION SQUAD

Every successful mission has its unsung heroes in mission control. Ours were formed of a group we called the Adoption Squad.

This team diligently led our department through the transition, bringing together a powerful mix of end users, proven leaders who had user experience, and change management expertise. They helped to design and deliver a staggered yet accelerated user journey for 1,000+ people, heading toward a destination that was still evolving, all while ensuring no real study deliveries were ever impacted.

It was, at times, the equivalent of training astronauts to get to the moon, only to realize midway through the journey that we actually needed to get them to Mars, requiring continuous adaptation!

The most important characteristics of this team were incredible empathy (we had to constantly put ourselves in the user's shoes) and genuine credibility. The team wasn’t made up of only generic change managers or communication experts, many were people who had lived and breathed data science roles. This built strong trust and two-way respect with our end users.

The squad was organized into dedicated leads focusing on:

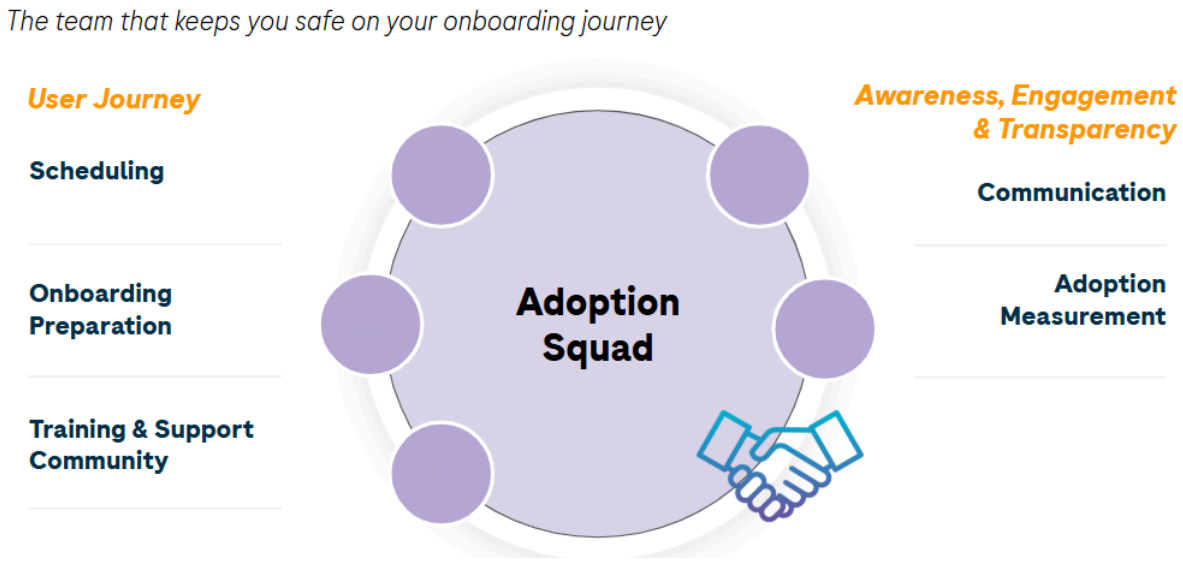

1. **Scheduling:** Prioritizing who joined the journey when. This included balancing early adopters (those with low risk and high passion), non-core user groups, and working around each team's busy reporting periods.
2. **Onboarding Preparation:** Planning the entire user journey and coordinating with a dedicated migration squad that helped move necessary files, through helping teams strategically choose what to archive versus what to migrate.
3. **Training & Support Community:** Ensuring every element of the user journey was delivered consistently. Training plans and support methods had to resonate cohesively between the platform and tools, preventing the various development teams from throwing out different, confusing expectations of the users.
4. **Communication:** Keeping people informed of only what they needed to know, exactly when they needed to know it, thereby controlling against information overload.
5. **Adoption Measurement:** Checking that the new solutions actually landed and got embedded as we hoped, with the predicted value realization, and creating a continuous 360 degree feedback loop connecting users back to our tech development teams.

THE ASTRONAUT TRAINING PROGRAM: MAPPING THE USER JOURNEY

To handle this change of such magnitude, we created a detailed user journey, giving everyone a clear map from the "current state" to the "future state". It was meticulously planned, though still flexible enough to cope with the occasional detour to re-fuel.

The journey steps (with rough timings shown) consisted of:

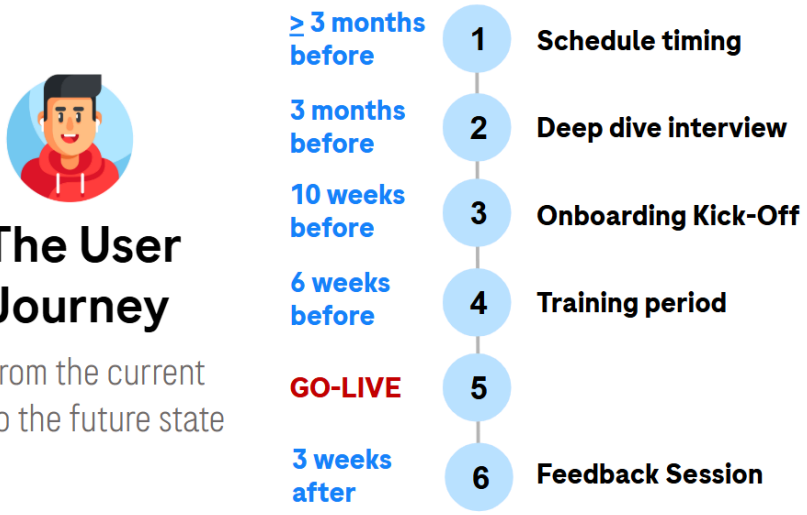

- **Schedule Timing (>3 months before):** Early on, we'd identify the right time for each molecule team to come onboard, always ensuring a significant notice period.
- **Deep Dive Interview (3 months before):** This helped us to understand each team's detailed needs and analytical strategy. From this, we could develop bespoke training and migration plans, and ensure the platform status was actually ready for them.
- **Onboarding Kick-Off (10 weeks before):** Team members received a detailed run-through of the journey ahead and, crucially, were connected with their **onboarding coach**. This was a key role, filled by a mix of members from our tech teams and early adopter teams, whose mission was to support teams and make them feel safe. We found that giving each team a personal guide meant they were more open to share any struggles or concerns, and they'd never feel embarrassed of asking the potential "silly questions" as they would have in larger settings. It also allowed for approaches to be tailored to work best for each team, e.g. considering geographical split, experience levels, training gaps, etc. During this phase, any pre-requisite training (like core skills in UNIX, git, and R) was made available here in case needed.
- **Training Period (6 weeks, ~1 day/week expectation):** This period helped new users acquire the skills for the new environment, through the following ways:
    - **Tiered Learning:** Trainings were split into basics and advanced levels. Users could master the basics first to successfully go-live, and then pick up more advanced topics later once they felt comfortable.
    - **Safe Practice Space:** We provided a sandbox training environment with test data and lots of practical exercises, allowing trainees to practice the new workflows safely at their leisure.
    - **Team Hackathon:** At the end of the training period, we held a team hackathon - an end-to-end challenge to see if the team could make a dummy delivery covering most of the basic skills required.
    - **Community:** Training as a team (as opposed to individual by individual) fostered community-based support, as team members helped one another alongside their onboarding coach.
- **GO-LIVE:** The big day - each team's unique "launch". However, we often reiterated that this was just the start of the voyage, as in truth it's the skills and confidence you pick up over time when working in production that really cement the learnings.

- **Post Go-Live (3 weeks after):** Onboarding coaches stayed with the team for the first three weeks to ensure everyone landed safely. After that, they moved on to support a new team, having passed the torch.
- **Feedback & Pay-It-Forward:** We closed the loop with feedback collection, ideation, and an open ask for them to help support those following them later on the journey.

To enable the above user journey, we had to make certain strategic choices, especially around training and support philosophies. We built these around user empowerment, as follows:

- **Adaptive Training:** We made training as self-service and flexible as possible to allow for different pathways and a gradual learning curve. For example, not everyone needed to jump into R immediately. This enabled more potential for just-in-time upskilling. For core skills like git, we complemented self-learning with in-person workshops to solidify understanding.
- **Community Support Model:** We sought to embed a community-focused model to spread subject matter expertise far and wide, ensuring we didn't rely only on a small group of technical experts. We wanted users empowered to help one another. We adopted an internal version of Stack Overflow for Q&A, plus an informal user chatroom to promote knowledge sharing. Our high-priority project teams were supplemented with extra technical support to quickly overcome any short-term learning blocks, and to ensure never any risk to key deliveries.

KEEPING THE PUBLIC ENGAGED: TRANSPARENCY

As with John F. Kennedy's approach to engaging the public, we felt extreme transparency was critical to building trust and managing anxiety during such a major upheaval. People fear what they don't understand, especially if that fear relates to their everyday activities.

We maintained a central monthly news bulletin to keep everyone - even those who wouldn't be joining for some time - up-to-date on the latest platform news. We also had a tracker that was open to all, that showed the schedule of upcoming teams joining and those who had already onboarded. This helped create a feeling of community across teams experiencing similar changes, and helped manage expectations for other upstream and downstream system teams and wider stakeholders.

We openly communicated platform gaps and features that were planned for future releases. This showed users that this was an ever-evolving, agile and iterative project where their feedback could genuinely have an impact. The key tenet we always returned to was user co-creation - having people lean in to create a platform they truly wanted to work on.

To keep spirits high and the goal visible, we shared a thermometer visual in our news bulletins, showing the status "heating up" towards our 90% target. Every onboarded team was called out and recognised for their contribution towards this shared goal, especially those teams who were willing to take the dive into the unknown during the early stages.

FINDING OUR NEIL ARMSTRONG: THE POSTER CHILD

Thinking of those early teams, we knew right from the start that we would need a "Poster Child" team. A brave team of early pioneers, who would take on every available change and deliver the ultimate proof-of-concept: a successful end-to-end R and open source based global submission on the new platform. A team who could prove any doubters and resistors wrong, and show that anything is possible in the new world. We needed our very own Neil Armstrong and crew!

We were incredibly lucky to find such a team who embraced the risk and whose members were passionate about the new platform and moving towards R/open source. This team became an extension of the tech development teams, giving crucial, little-and-often feedback and supporting key design decisions. Most importantly (even though they needed a lot of hypercare along the way), their successful filing made everyone else believe the moonshot was truly possible!

This team even presented a wonderful webinar on their experience of filing using an end-to-end R-based solution, primarily leveraging open source packages. They highlighted how cross-industry efforts like pharmaverse (a PHUSE working group - see https://pharmaverse.org/) made a world of difference to their groundbreaking achievement. You can learn more about their experiences by watching this here: https://www.youtube.com/watch?v=BlJNILSoZlM.

ROCKET FUEL: USING AI

You can't have a talk/paper nowadays without mentioning AI, so let me share a couple of ways we used artificial intelligence to assist our users going through this massive change.

1. **AI Assistant/Chatbot:** We developed an AI assistant/chatbot trained on all our technical and user-facing documentation, as well as our internal Overflow Q&A pages. This chatbot greatly reduced the level of user support needed and helped people feel more comfortable asking basic questions.
    - **Honesty is the Best Policy:** In early iterations, where internal information was still scarce, we included a traffic light system for the answers. This helped users understand where the AI might be clutching at straws because the documentation didn't exist, making hallucinations more likely.
    - **Utility Growth:** This chatbot quickly grew into providing extra useful utilities such as for debugging pipelines, searching for re-usable code snippets, and providing help for working with git. Since inception, it has assisted users with over 100,000 queries!
2. **Multi-Agent R Solution:** As part of the journey was a move towards R, we also developed a multi-agent AI coding assistant. This provided the user with bespoke agents dedicated to each specific R package (including any company-specific closed source packages). This ensured users got the most accurate possible answers to help with coding challenges. In the first month of launch alone, over 250 colleagues used this AI solution for more than 1,000 questions!

We also had the option to embed commercial AI solutions, such as GitHub Copilot, directly into our development environment.

MISSION ACCOMPLISHED

So, with all these well-intended adoption plans, did we make the goal?

Of course we did! It would make for a horribly anti-climactic paper if we hadn't. We actually hit 91%, even though the original 90% was itself seen as an audacious, stretch target.

Plus none of this was achieved with a smooth ride. During the year, we had leadership changes, resourcing impacts, new acquisitions - there couldn't have been more hurdles to overcome. Through all this though, people stepped up, and the "in it together" spirit we instilled from the start is what carried us through. Our study teams were truly heroic in balancing an influx of learning alongside keeping the critical business deliveries going for our patients.

When we eventually hit the target, we celebrated across our whole data science department. This was everyone's collective achievement.

THE NEXT FRONTIER

After such a taxing user journey, we had to be careful not to rush into setting another moonshot. But, like all progress in science, success only unlocks looking towards the next realm of possibility.

In 2025, we supported the final 9% of teams to complete their full onboarding and migration. Here, we learned to never underestimate the effort required to move the straggler teams - there are often many complex reasons as to why they ended up last in the queue!

Now, given our rapid onboarding and training, we are working to achieve mastery in our users. This involves reinforcing learnings and continually simplifying workflows until we can reach a stable state of operations.

Furthermore, we now have the solid foundation to accelerate our adoption of R and open source. We are achieving this through a mix of soft targets, leading to mandates for new studies, and we've even tried some gamification methods here to introduce more competitive elements into molecule teams making the switch. After all, it's important to still make things fun for our teams.

At this stage, the next frontier we've been undertaking is augmenting AI into all of our workflows to help our users maximize their time and focus. We found that landing our first moonshot has only propelled us to explore further and achieve more, so I'll end with a beautiful thought from sci-fi author Arthur C. Clarke: *"The Moon is the first milestone on the road to the stars"*.

Thank you for reading.

ACKNOWLEDGEMENTS

Thank you to the exceptional Adoption Squad and to each and every person in the Roche/Genentech data science department who lived through this change.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

ross.farrugia@roche.com