

SOFTWARE FOR BIG SCIENCE: **THE INSIDER'S GUIDE**



CONTENTS

Case Studies

Unleash the Power of Scientific Data

Neutron Facilities Embrace an Open Approach

Community Platform Drives Data Discovery

The Insider's Guide

Better Software Delivers Better Science

Over three decades Tessella has had the privilege of working with some of the world's most advanced organisations to deliver innovative answers to their most complex challenges. Nowhere is this truer than in the field of science research, where our work has helped create robust and efficient informatics systems that support some of the most important and ambitious projects of our time.

These projects are all huge in scale, and often tend to require specialist software. Recently, we've also seen that the complexity research organisations must conquer is exacerbated by the incredible possibilities that new developments in technology offer, which generate vast amounts of data, but require sophisticated analysis and management to deliver new levels of insight.

Right now we're also finding that many facilities want to really strengthen their position in competitive markets and collaborate more effectively with partners. For example, we've partnered with STFC's ISIS Neutron and Muon Source for several years in the development of Mantid, a powerful data reduction and analysis application, which has now been adopted by an increasing number of world leading research facilities to help them grapple with huge volumes of information. More recently, we helped ISIS to manage the transition from its existing platform for instrument control to an Open Source environment. In the process we enabled ISIS to share and use off-the-shelf instrument drivers and collaborate more closely with other facilities.

The purpose of this book is to build on the knowledge we've gained through this kind of practical experience and pass it on to the science research community. We've been working with many of our clients in this sector for very many years, and I'm delighted to say they've openly contributed a great deal of insight on their real-world challenges and how they overcame them. I extend my thanks to all those who have kindly contributed their time in compiling this book.

We also worked closely with the Institute of Physics and the editor of Physics World to frame our questions and carry out the interviews. This means we're able to offer you invaluable and impartial advice, covering all the key issues that need to be considered when embarking on a new project. The report also includes case studies from facilities that have already been through this process, which means you can really learn from their experiences.

Whether you're a senior decision-maker or involved in managing complex software systems in Big Science facilities, I hope there's something here for you.

We hope you find the book insightful and will give you all the information you need to make the strategic decisions that will influence the long term success of your project.



Alan Gaby, Chief Executive, Tessella Group



UNLEASH THE POWER OF SCIENTIFIC DATA

The growing trend towards open data will require a more disciplined approach to managing the software systems used to acquire, process and store experimental results.

Ever since the Royal Society was founded in 1660, generations of researchers have embraced the scientific method: design an experiment that tests a particular hypothesis, and measure the result in a way that other scientists can reproduce and verify. The experiments of the past required only physical equipment, which is relatively easy to replicate, but today's scientists must also ensure that the software used to acquire and manipulate their experimental data always yields the same outcome.

"For small projects it may be sufficient to design the software to ensure that the experiment would produce the same result if it was run again," says Neil Chue Hong of the Software Sustainability Institute (SSI), which was set up in 2010 to help researchers to follow best practice when developing their own code. Without proper oversight, says Chue Hong, small changes to the software might cause the same experiment to produce different results – but no-one would know why.

Chie Hong explains that the SSI runs courses to provide researchers with the basic skills to avoid such problems, such as proper version control, understanding the dependencies between different pieces of software, and developing effective documentation. "For scientists the goal is to get the results out as quickly as possible, which doesn't help reproducibility," he says. "But it's vital to know where the data came from."

Understanding the provenance of experimental data is becoming even more important as pressure grows on researchers to open up their datasets to the broader scientific community. The UK has already mandated open access to data acquired through publicly-funded research, and it seems likely that Europe will follow suit in the not too distant future. But, Chue Hong points out, just making the data available is not the same as making it easy for others to access and re-use. "The emphasis is now shifting to making sure that datasets can also be exploited by other people," he says.

"The test is whether the experiment has been described well enough to enable people to use the data for a different purpose"

That means storing the data in a structured format with additional information, known as metadata, that describes how, where and when the data was acquired. As an example, a dataset of ocean temperatures is meaningless without knowing where the data were recorded, at what time of day and year, and in what units. "The test is whether the experiment has been described well enough to enable people to use the data for a different purpose," Chue Hong continues.

Large facilities such as those run by the UK's Science and Technology Facilities Council (STFC) already see the value in opening up the data obtained by visiting scientists. "The astronomers have already got this," says Andrew Taylor, Executive Director of the National Laboratories run by the STFC. "They know that their observations are valuable to other people, so they curate their data and include metadata to make it easier for other people to use."

Taylor is aware that scientists who use the STFC's facilities may be reluctant to make their hard-won experimental results so readily available to other research groups. He explains that astronomy data is typically kept private for six months before being made available to the wider community. "Here at the STFC our approach will be to keep the data private for the duration of a PhD project, and then open it up," he continues.

"Scientists often want to write their own software, but their focus is on the experiments and they don't always take the time to write the code properly"

Taylor also recognizes that the shift towards open data dictates a more professional approach to software development. "Scientists often want to write their own software, but their focus is on the experiments and they don't always take the time to write the code properly," he says. "We want to combine the creativity of the scientists to design new experiments with the discipline of proper software engineering."

Scientists at the STFC have been working alongside software engineers from Tessella to ensure a more structured approach to software development. "Tessella is focused on scientific software engineering," says Robert McGreevy, director of the ISIS neutron source, which has worked with Tessella on two major projects. "They bring professional expertise but they also understand the scientific environment."

Chue Hong advocates a new career path, which he calls software for science, that would provide greater recognition for people who understand the science and also have the skills to create a professional software solution. "There can be a sense that the software is less important than the science, and that needs to change," he says.

For Taylor, the drive towards open data will play a crucial role in increasing the research impact of the facilities run by the STFC. "We need to optimize the use of our data by ensuring it is repeatable and reusable," he says. "By getting more out of the data, and by introducing more automation and remote management, our goal is to double the volume of research carried out at ISIS by 2020."



A stylized atomic model is visible in the background, featuring a central nucleus and several elliptical orbits with small spheres representing electrons. The entire image has a green color scheme with a grid pattern.

NEUTRON FACILITIES EMBRACE AN OPEN APPROACH

The bold decision to replace a myriad of separate software systems with a single open-source framework could yield a new standard for data processing at neutron sources worldwide.

Software systems have a habit of growing haphazardly: one piece of code is written to achieve a specific function, then another is developed to achieve a different objective, and so on. Each individual step might be a logical progression, but the result is a fragmented collection of distinct systems rather than a coherent solution.

That was exactly the problem confronting the ISIS neutron source, operated by the UK's Science and Technology Facilities Council, when it embarked on a major upgrade of its instrumentation in 2008. ISIS is responsible for 25 different experimental stations, and each one requires software to clean up the raw data to extract the information that's valuable for scientific analysis – a process known as data reduction.

“Previously, the software had been developed by the scientists working on each of the instruments,” says Debbie Greenfield, head of the instrumentation division at ISIS. “They have a detailed understanding of the algorithms they want to implement, but they aren't professional software engineers and so had developed the software in their own way.”

“ISIS needed a more organized, more professional solution, and realized that they needed to put proper resources into it”

The result was separate pieces of code with little documentation or version control, which made it difficult to support and maintain – and in some cases also led to duplication of effort. “ISIS needed a more organized, more professional solution, and realized that

they needed to put proper resources into it,” says Nick Draper, a project manager from scientific software company Tessella who has worked with ISIS on the project. “The goal was to replace all the individual pieces of software with a single framework that could be implemented on all of the existing instruments, and could be easily extended to operate with new experiments.”

ISIS was also keen to develop the platform as an open-source solution that could be adopted by other neutron facilities around the world. “We knew we needed a single platform, and we made the decision to make the initial investment ourselves,” says ISIS director Robert McGreevy. “But it was always our intention to get other facilities on board.”

At the start of the project, remembers Draper, it was difficult to devise a single software solution that would work effectively with all the different instruments. “But then we started to break down the operation of each instrument into smaller tasks, and found that everyone was essentially trying to do the same thing,” he says. “By creating smaller building blocks, it was possible to build a flexible solution that could be adapted or extended as needed.”

The resulting platform, known as MANTID, has since been implemented across most of the experiments at ISIS. And the gamble of developing an open-source solution has paid off: the system has been installed at the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory in the US, and will also be implemented at the European Spallation Source in Lund, Sweden.

“We can share expertise and share development, but we still have the freedom to develop our own software”

For McGreevy, the long-term benefits of such an open-source approach are obvious. “Our requirements are roughly the same, and by working with the SNS we are effectively doubling the number of people we have working on the project,” he explains. “We can share expertise and share development, but we still have the freedom to develop our own software.”

Greenfield explains that there is a single master copy of the MANTID system, which the two sites jointly develop. A steering committee composed of scientists from both facilities agrees the priorities for further development, with each partner contributing to the ongoing operational costs. The same approach will be exploited when the ESS implements the system, and McGreevy is keen for other neutron sources to join the consortium. “It will be great for new facilities like the ESS because it will be very quick for them to implement,” he comments.

Crucial to the success of the project has been a professional approach to developing the MANTID system, which has been introduced by Draper and his colleagues at Tessella. At the beginning of the project ISIS had no internal IT specialists, and so Tessella provided a complete cohort of project managers and software developers. Since then Greenfield has established an internal team of software engineers to reduce the facility's reliance on external

expertise, and Tessella engineers now work alongside the ISIS developers in a fully integrated team.

“Tessella can offer expertise and experience in software engineering, while the internal developers have better links to the instrument scientists,” says Draper. “MANTID provides a basic framework that can be extended to work with specific applications. It means that the instrument scientists can spend less time coding, and more time on the science.”

More generally, MANTID clearly shows the power of collaboration. Although ISIS took the initial lead on the project, the UK facility has also benefitted from techniques pioneered at the SNS. One example is a new way to handle event-based data, which was initially developed at Oak Ridge National Laboratory but has since been adopted at ISIS with minimal adaptation.

“A common platform operating at different sites is much more efficient, and much better for the researchers who use the facilities”

The success of MANTID within the neutron science community could mean that the open-source approach becomes the standard approach for software development at large scientific facilities. “A common platform operating at different sites is much more efficient, and much better for the researchers who use the facilities,” concludes Greenfield.





COMMUNITY PLATFORM DRIVES DATA DISCOVERY

Opening up software and data to a broad scientific community can spur progress through greater collaboration, but also demands accessible and reliable software solutions.

Making predictions about the atmosphere – whether forecasting the weather three days ahead, or trying to understand long-term climate change – is a notoriously perilous task. Large-scale computer simulations of weather and climate have become more complex and more granular as computing power has grown, but the accuracy of the models depends on the availability of reliable observations to calibrate the models and validate the results.

Atmospheric and earth observation scientists have found ways to capture data describing everything from ocean temperatures through to the water content of clouds, but bringing these diverse datasets together has proved a challenge. “Scientists have been fighting with different technologies and different data formats,” says Philip Kershaw of the UK Science and Technologies Facilities Council (STFC). “A consistent set of tools is essential for them to record and store their data.”

That’s the idea behind the Community Inter-Comparison Suite (CIS), an open-source toolkit developed in a collaboration between the STFC’s Centre for Environmental Data Archival and the University of Oxford. “Researchers had developed their own techniques for converting raw data sources into a format that enabled direct comparison, but they weren’t consistent,” says Nick Cook of scientific software company Tessella, which provided professional support for the project. “The solution was to develop a generic algorithm for each type of data, which ensures reproducibility because the same analysis technique is always used for a particular type of data.”

The CIS framework enables data recorded in different formats to be visualized and analysed with the same set of tools. “If a new type of data becomes available, it’s easy to write a new plug-in that works within the overall framework,” continues Cook. Plug-ins can also be developed to address new analysis techniques, such as a co-location plug-in that allows datasets covering the same area and time period to be mapped onto each other for easy comparison.

“We’re opening up the software developed by the scientists to wider community by making it accessible, reliable and repeatable”

Working with such a diverse community of users presents its own unique challenges. Anyone can access the toolkit and download it to their local computer, which means that reliability and ease-of-use are of paramount importance. “We’re opening up the software developed by the scientists to wider community by making it accessible, reliable and repeatable,” says Cook.

Key to the success of the project, says Kershaw, has been the agile development process introduced by Tessella. In this methodology the project proceeds through a series of two-week-long steps, known as “sprints”, with clear objectives for each sprint agreed and reviewed with the user community. “The agile process has provided a mechanism for us to engage with users during the project, allowing us to respond to their input and adapt the objectives as necessary,” explains Kershaw. “Without this contact and ongoing dialogue, it would be very easy to build something they don’t want.”

A crucial element of the agile approach is that testing is built into the process, rather than being seen as an afterthought. A test is identified for each objective before development even starts, ensuring that new pieces of code are verified before being deployed into the overall framework. This ongoing testing process ensures that all data collected through the CIS framework is always consistent, and that the software is always reliable for the end-user community.

Kershaw explains that the CIS toolkit is part of a larger project, called JASMIN, that offers scientists a shared infrastructure for accessing and analysing the results from large-scale simulations. “As computer models of the climate system become more accurate and more granular, they produce more and more data,” he explains. “With JASMIN we are bringing together the results from computer models with observational data to provide a single analysis platform for multiple organizations.”

At the heart of the JASMIN system is a supercomputer capable of petascale data processing and storage, hosted by the STFC but available to researchers anywhere in the world via web-based access protocols. According to Kershaw, one key challenge was ensuring that users who have never worked in a high-performance computing environment can find and exploit data held within the system.

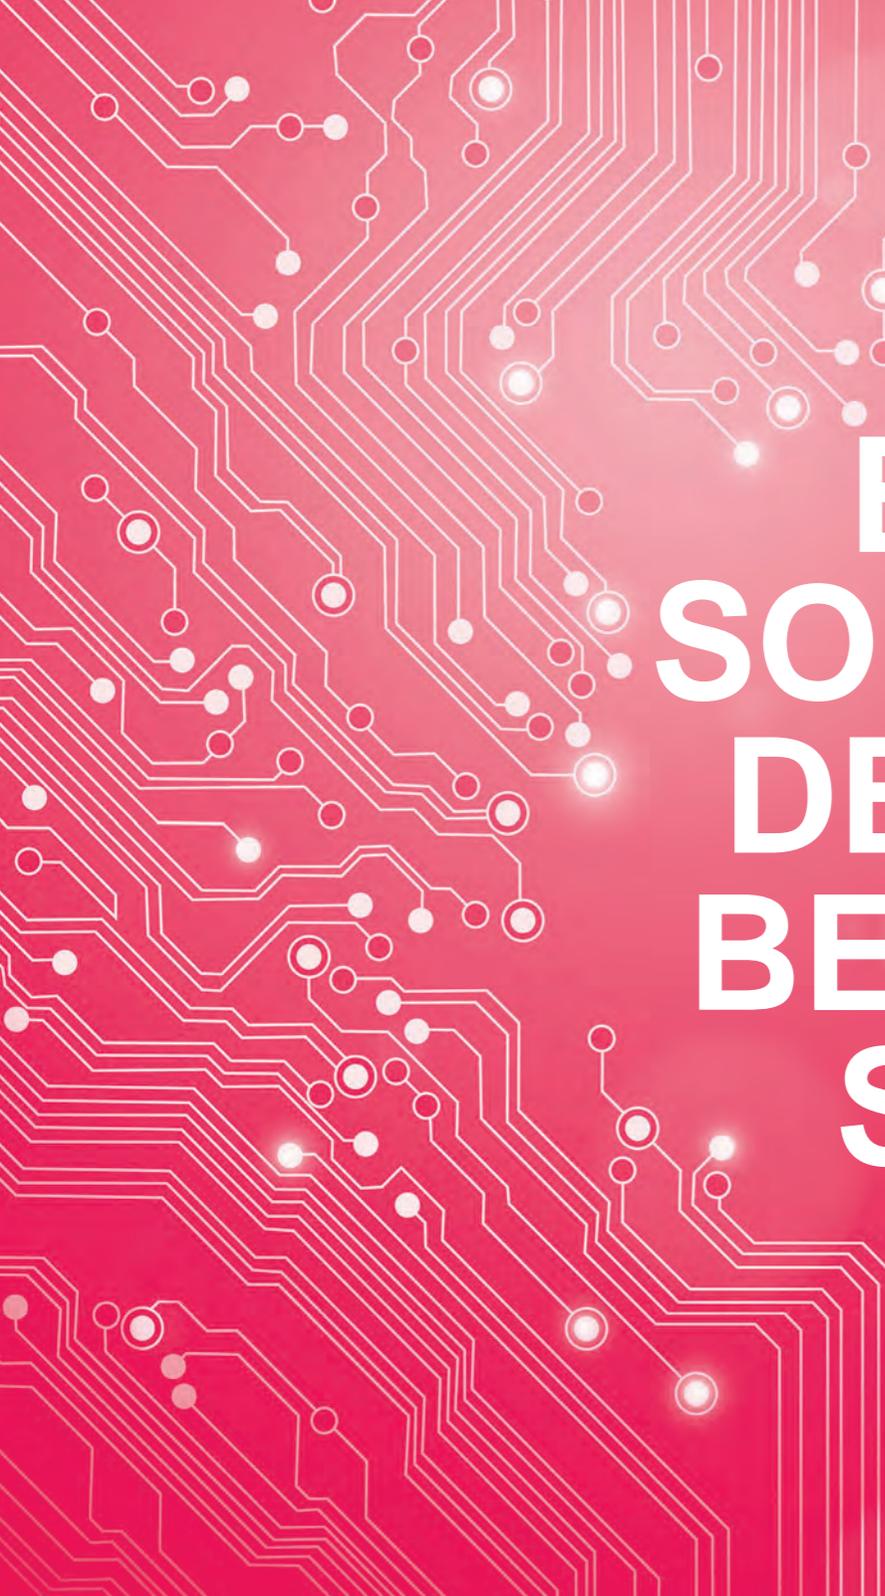
“We wanted our wide spectrum of users to have the same experience as they would with a desktop computer”

“Most scientists are computer literate, but not everyone is familiar with the technical processes normally needed for high-performance computing,” he says. “We wanted our wide spectrum of users to have the same experience as they would with a desktop computer, while also providing access to the computing power and datasets available through the JASMIN system.”

JASMIN is already showing its potential for enabling collaborative working within the climate science community. For example, it holds a 400 TB dataset produced by the UPSCALE project, in which scientists at the UK Meteorological Office and the University of Reading simulated the global atmosphere at resolutions down to a grid size of 25 km – the same as is used for global weather forecasts. The simulations, which ran over a 27-year time period, required an unprecedented 144 million processing hours on the HERMIT supercomputer in Germany, and the resulting dataset is so large that any analysis must be run through the JASMIN infrastructure.

“JASMIN will be transformational for the people in this community,” says Kershaw. “We have already demonstrated some excellent results to users, and we are now upgrading the infrastructure to offer new services to the community.”





BETTER SOFTWARE DELIVERS BETTER SCIENCE

Software has become an integral part of modern scientific research, and putting the right solution in place from the start will yield long-term efficiencies as well as more valuable experimental data for the research community.

Embarking on any major software project is always exciting, but also tinged with trepidation. At the outset it's possible to focus on the outcomes: the new possibilities the software will allow, as well as the improvements in speed, performance and accuracy that will be achieved. But there's always a worry that the project will exceed its budget, take longer to complete than expected, or – most scary of all – that it won't deliver its promised outcomes.

And yet software is now a critical element of any scientific project. It is a vital tool for controlling instrumentation, acquiring experimental data, and analysing and visualizing the results to gain new scientific insights. What's more, large projects and facilities are now generating huge datasets from their experiments, which in turn are becoming an increasingly valuable resource for the wider scientific community.

“We need to combine the creativity of the scientists to design new experiments with the discipline of proper software engineering to ensure the data is reproducible and reusable”

“We want to maximize the exploitation of our data,” says Andrew Taylor, Executive Director of the National Laboratories, operated by the UK's Science and Technology Facilities Council (STFC). “We need to combine the creativity of the scientists to design new experiments with the discipline of proper software engineering to ensure the data is reproducible and reusable.”

As a result, experimental facilities such as those run by the STFC are placing more emphasis on understanding the software and data requirements for their projects, as well as considering how both the software and the data might be used in the future. This is driving a more professional approach to software development, making it easier for the facilities to manage their systems and deliver new projects on time and on budget – while also extending the possibilities for scientific enquiry.

We asked experts from leading research organizations, along with the software professionals with whom they have partnered, to share their secrets for success. Taken together, they offer a blueprint for ensuring that your next project delivers on its promises.

1. Open up your software

It may seem counter-intuitive, but making your investment in software development available to a broader community can deliver major gains in the long term. Take the example of the ISIS neutron source, which has developed a single framework for data reduction and analysis that can operate across all 25 of its instruments. Originally set up as an internal project, the MANTID system was developed as an open-source platform that has since been adopted by the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory and will also be implemented by the European Spallation Source now being built in Lund, Sweden.

“It allows us to share expertise and share development resource, while also allowing us to introduce our own software if we need to”

“We made the decision to put the money in, but it was always our intention to get other facilities on board,” says ISIS director Robert McGreevy. “It allows us to share expertise and share development resource, while also allowing us to introduce our own software if we need to.”

McGreevy says that both facilities have benefitted from the open arrangement. As an example, SNS was the first to develop techniques for handling event-based data, and ISIS was able to introduce the same system with only minor modifications. The ongoing efficiencies gained through such open-source development makes it an increasingly attractive option for large scientific institutions to develop new software solutions.

2. Ask for advice

Before embarking on a major project, it's always worth talking to the professionals first. “We look at what they're trying to do, what the obstacles might be, and how they work with software,” says Kevin Woods, a project manager for scientific software company Tessella. “In some cases we find that other issues need to be resolved before starting on any new software solution, and it's best to be candid if that's the case.”

“Progress can be slow if software projects aren't specified properly and don't have enough resource”

On the flipside, it's crucial for research organizations to consider the IT infrastructure and data requirements as an integral part of the planning process. “Progress can be slow if software projects aren't specified properly and don't have enough resource,” says McGreevy. “It's really important to understand how much effort is needed to deliver the solution you want.”



3. Data is king

As computational analysis becomes more and more important for scientific progress, it's crucial that experimental data is collected and analysed in a reliable and repeatable way. "For small projects it may be sufficient to design the software to ensure that the experiment would produce the same result if it was run again," says Neil Chue Hong of the Software Sustainability Institute, which aims to help researchers to follow best practice when developing their own code. "But the emphasis is now shifting to making sure that datasets can also be exploited by other people."

"The test is whether the experiment has been described well enough to enable people to access and re-use the data"

That means storing the data in a structured format with additional information, known as metadata, that describes how, where and when the data was acquired. "The test is whether the experiment has been described well enough to enable people to access and re-use the data," Hong continues.

And this isn't mere altruism. Large facilities, such as the ISIS neutron source run by the STFC, could have a much greater research impact by opening up access to their data. "We need to optimize the use of our data through proper curation," says Andrew Taylor. "By getting more out of the data, and by introducing more automation and remote management, our goal is to double the volume of research carried out at ISIS by 2020."

4. Take a professional approach

Most research scientists have the skills to write their own software, but as experimental data becomes more valuable it's becoming ever more important to adopt a disciplined approach to software engineering. "It's easy to write code that does the job," says Tessella's Nick Cook. "But without proper oversight the software becomes difficult to support or maintain, and can even produce inconsistent results without proper software management."

Those problems are compounded in a large facility, where separate pieces of software often evolve over time to operate with different instruments. At the ISIS neutron source, for example, the single MANTID framework is replacing software that has been developed independently for 25 separate instruments. "The instrument scientists have a detailed understanding of the algorithms they want to implement, but they are not professional software engineers," explains Debbie Greenfield, head of the instrumentation division at ISIS. "Introducing a single system across all our instruments required a more formal approach to ensure proper documentation, testing and version control."

"The instrument scientists have a detailed understanding of the algorithms they want to implement, but they are not professional software engineers"

5. At the interface

While scientists once only produced code for their own internal projects, it's becoming more common for research institutes to produce software for commercial or collaborative projects. For example, Paul Holligan at the STFC's Central Laser Facility is implementing a control system for a turnkey laser system that will be delivered to the HiLASE project in the Czech Republic. "It is a commercial product so we want to provide a professional solution," he says. "As part of the project we have updated the user interface so that it works on tablets, which allows remote control of components and subsystems."

"Most scientists are computer literate, but not everyone is familiar with the technical processes normally needed for high-performance computing"

And the JASMIN project, developed in a collaboration between the STFC's Centre for Environmental Data Archival and the University of Oxford, aims to provide climate researchers anywhere in the world with access to a supercomputer capable of petascale data processing and storage. "Most scientists are computer literate, but not everyone is familiar with the technical processes normally needed for high-performance computing," says the STFC's Philip Kershaw. "We wanted our wide spectrum of users to have the same experience as they would with a desktop computer, while also providing access to the computing power and datasets available through the JASMIN system."

6. Choose your partner with care

If you think your project needs external support, it pays to find a software company that will work effectively with your organization. Since most scientists have the skills to write their own software, they may well get frustrated with external software engineers who impose a solution on them. For Robert McGreevy, it's crucial to work with an organization that specializes in scientific software development. "They understand what we're about, they understand the scientific environment."

External partners must also be able to adapt their approach to meet the needs of the organization. "A very formal approach to software development doesn't usually work well in a research-driven environment," says Tessella's Kevin Woods. "We can introduce techniques that will help them, while they can focus on their research."

"We're on the same wavelength, and we understand what they're trying to achieve"

For shorter projects in particular it can feel as though consultants are being parachuted in, and so it's even more crucial to win the confidence of the scientific teams right from the start. According to project manager Nick Cook, it helps that the software engineers at Tessella are also trained scientists. "We're on the same wavelength, and we understand what they're trying to achieve," he says. "It helps that we can talk to them about the issues that matter to them, such as the challenges they have with funding and the science they're trying to solve."

7. Take a flexible approach

Many scientists aren't comfortable with a formalized approach to software project management, particularly when it requires them to plan every detail months in advance. The solution in many cases is to adopt an agile approach, in which the project proceeds through a series of small steps – but always with an ultimate goal in mind.

“We work closely with the instrument scientists to introduce the control system in an adaptive, step-by-step process”

At the STFC's ISIS Facility, for example, an agile process has been exploited to introduce the controls framework across 25 different instruments. “Our approach has been to focus on one or two instruments at a time,” says Woods. “We work closely with the instrument scientists to introduce the control system in an adaptive, step-by-step process, which minimizes the impact on the operational system and allows us to troubleshoot any problems in a more localized way.”

An agile approach has also helped the STFC's Paul Holligan keep his laser project on track. “When capturing initial requirements for components and subsystems the instrument scientists ask for everything, but that can be a problem when we have a fixed budget and a fixed deadline,” he explains. “Through the agile process we have worked with them to agree a minimum baseline - the functionality that is crucial for the system to work - and then additional features can be added through a series of iterations.”

8. Keep control

Sometimes an organization might need to draft in a complete team of software specialists because it lacks the internal expertise or resources. That was certainly the case at ISIS at the start of the MANTID project, as Debbie Greenfield remembers. “We involved Tessella because there were no professional software engineers working at ISIS,” she says. “Tessella provided the project manager, the software project manager, as well as a team of developers.”

“By working together, the internal teams benefit from our experience”

But this is mission-critical software, and over time Greenfield has established an internal team of software engineers to reduce the organization's reliance on external knowledge. “It's important for the facility to have internal IT specialists who understand how the system works,” says Tessella project manager Nick Draper. “By working together, the internal teams benefit from our experience, and also gain efficiencies by having an external resource that is dedicated solely to this project.”

9. Keep talking

Ongoing communication with the end-user community is crucial for success. An agile approach builds regular feedback and review into the project, which makes it easier to clear up any misunderstandings as the project proceeds. “The agile process provides a mechanism for us to engage with users during the project, allowing us to respond to their input and adapt the objectives as necessary,” says the STFC's Philip Kershaw. “Without this contact and ongoing dialogue, it would be very easy to build something they don't want.”

“Regular communication is crucial to manage the expectations of the scientists, and make sure they are kept aware of the process”

Paul Holligan also advocates regular dialogue, particularly when implementing such “under the hood” software as control systems. “Regular communication is crucial to manage the expectations of the scientists, and make sure they are kept aware of the process,” he says. “Users don't necessarily see the progress with control systems, and regular updates keep everyone aware of progress and the current priorities for the project.”

10. Partnership counts

Although it can sometimes help to draft in additional resource just to get through the work, working in partnership with software specialists generally yields better results. In the case of ISIS, Tessella engineers work alongside the internal developers, creating a completely integrated team. “You wouldn't know who works for ISIS and who works for Tessella,” says Debbie Greenfield.

“Working on site allows us to understand problems and discuss possible solutions”

An obvious advantage of such integrated teams is that the external developers build up knowledge about the project and communication becomes much easier. “Working on site allows us to understand problems and discuss possible solutions,” says Nick Draper. “The Tessella developers can contribute to the design in a way that enhances the MANTID framework, while the internal developers have better links to the instrument scientists.”



Tessella would like to gratefully acknowledge and thank the following contributors:

Neil Chue Hong

Director, Software Sustainability Institute

Debbie Greenfield

Head, ISIS Instrumentation Division, Science & Technology Facilities Council

Paul Holligan

Leader of Electrical Section, Central Laser Facility, Science & Technology Facilities Council

Professor Robert McGreevy

Director of ISIS, Science & Technology Facilities Council

Philip Kershaw

Technical Manager, Earth Observation, NCEO/Centre for Environmental Data Archival,
RAL Space

Dr Andrew Taylor

Executive Director, National Laboratories, Science & Technology Facilities Council



www.tessella.com

Copyright © July 2015 Tessella all rights reserved. Tessella, the trefoil knot device, and Preservica are trademarks of Tessella.