



Neil Ward-Dutton



An IT Strategy Brief for the New Era of Intelligent Automation

September 2018

MWD Advisors is a specialist technology advisory firm that shows how digital technology changes work; helping today's innovation, architecture and technology change leaders accelerate their success and manage risk. Our approach combines flexible, pragmatic mentoring and advisory services, built on a deep industry best practice and technology research foundation.

Top takeaways

1

A new era of automation requires a new strategy

Automation has been on a steady march through business for over 200 years: nevertheless, the past three years or so have seen a huge step-change in the variety of technologies being promoted to help businesses automate aspects of their work.

Given the wide variety of automation-related technologies being pitched to organisations, and the speedy pace of automation technology evolution, your organisation must develop an Intelligent Automation strategy that clearly and deliberately sets out how and where you will look to apply different kinds of automation technology capability – and why you are looking to apply particular technologies in particular places.

2

Think about 'how work works' as your Intelligent Automation starting point

We all know intuitively that not all work is equally amenable to automation: some work is highly suited to extensive automation, whereas the opportunity to automate other work is much less obvious. It's also true that all automation technology – RPA, workflow, case management, AI, and so on – is specialised to deal with particular kinds of situation. There's no 'silver bullet'.

In exploring how to apply different automation technologies to work, start by categorising work according to how much expert discretion is required to perform it. In practice we find that there are three types of work commonly found 'in the wild' in organisations, with different kinds of automation technology playing roles in addressing different blends of features across the three work types. We call these **exploratory work**, **transactional work** and **programmatic work**.

3

Turn strategy into action – at digital speed

Once you've begun to analyse how you might use different kinds of Intelligent Automation technology to address opportunities in the different kinds of work you have in your organisation, it's vital that you think beyond this analysis and put some plans in place to drive that analysis forward into action that will have real impact.

There are four things you particularly need to focus on: foster an automation technology ecosystem; remember that Intelligent Automation is about *business* change; review your assumptions and analyses regularly; and lastly: choose tools and practices that can help you work iteratively, quickly and flexibly.

The automation imperative

It might seem like automation in the workplace is a relatively new phenomenon – there's certainly a huge amount of talk about it in today's corporate meeting-rooms, conference halls and news media. The reality is that automation has been on a steady march through business for over 200 years: nevertheless, the past three years or so have seen a huge step-change in the variety of technologies being promoted to help businesses automate aspects of their work.

From financial services to healthcare, automation is top of the agenda...

Businesses the world over are rushing to explore digital initiatives. Some are driving full-on digital transformation programs; others are working more modestly, looking to apply digital technologies to particular business products and services, capabilities, activities or processes.

What's become abundantly clear is that any organisation thinking about 'digitising' aspects of their business needs to think holistically. Customers' experiences aren't only shaped by customer service teams – they're shaped by the interactions they have online; conversations with salespeople; the quality of a product; the accuracy of bills; and so on. By definition, delivering great customer experiences depends on activities that your organisation carries out across many teams and departments. The "digital outside" of your organisation must be connected to and supported by a "digital inside".

However, that's not the only reason organisations are keen to explore opportunities to automate work. For example, in financial services, compliance pressures – as well as changing customer expectations – mean that the large volumes of back-office work that still require significant manual, paper-based interventions have to be addressed urgently. In retail, volatility in product lines and assortments, and challenges of obtaining decent operational data – together with massive pressures from big online retailers – are often the catalysts for change. In pharmaceuticals and life sciences, the challenges of monitoring for unanticipated adverse drug reactions in the open market, analysing the incoming data, and rapidly responding with corrective actions are often behind the desire to drive more automation. In manufacturing, diminishing returns from outsourcing is driving organisations to look again at automating support processes.

End-to-end automation of work is not always possible or desirable. But even where end-to-end automation is not the right answer, there are ways to use digital platforms to improve the work environment.

...and it's cheaper and more applicable than ever

From about 2000 until 2015, organisations had three main technology choices if they wanted to automate aspects of administrative work: they could implement packaged business software applications (automating aspects of practices like ERP and CRM); they could build their own applications using traditional custom software development tools (probably based on Java or .NET technologies); or they could use more business-friendly, model-driven platforms to build systems (primarily focused on workflow automation and business rules automation).

Now, those three choices remain – but they're joined by more: not least, Robotic Process Automation (RPA) tools, low-code application development tools, low-code application integration platforms and various AI-related technologies.

Intelligent Automation in leading organisations: some examples

Intelligent Automation strategies are already maturing in leading organisations around the world, with particular hotspots in back-office processing – typically, in finance and accounting activities like accounts payable and accounts receivable, as well as in servicing simple customer service requests (for example statement or bill requests, change of address requests, and so on). Leading organisations are also using Intelligent Automation strategies to boost effectiveness in front-office activities, particularly where there are issues with legacy systems and challenging customer expectations.

UK-based general insurer Aviva has pursued a clear Intelligent Automation strategy in its customer services division, using a combination of workflow, case management and RPA technology to increase agents' effectiveness in dealing with requests – and seeing a 40% efficiency increase in customer request processing in its pensions business alone. Scandinavian bank Nordea has used RPA extensively, as a way of avoiding having to hire hundreds of workers to deal with access requests under the recently-introduced GDPR legislation. Online retailer Shop Direct has used RPA technology to automate around 40 back-office processes; while telecoms multinational Vodafone uses Intelligent Automation across its global shared services operation to drive efficiencies in the accounting, HR, and other business services it offers to Vodafone's operating businesses.

So how and where should you think about using these different technologies? That's what the next section covers.

Your automation strategy

Given the wide variety of automation-related technologies being pitched to organisations, and the speedy pace of automation technology evolution, it's vital that you develop an automation strategy that clearly and deliberately sets out how and where you will look to apply different kinds of automation technology capability – and why you are looking to apply different technologies in different places. The key thing to remember is that – despite what some vendors will try and tell you – different kinds of automation technology are specialised to deal with different aspects of work automation.

Not all work is created equal (for automation)

We all know intuitively that not all work is equally amenable to automation: some work is highly suited to extensive automation, whereas the opportunity to automate other work is much less obvious. What is it that makes the difference? And if you can distinguish between different types of work, can you use that distinction to guide your teams in applying the right kinds of automation technology to the right projects?

A good starting point is to think about categorising work according to how much expert discretion is required to perform it. The more expert discretion is required, the less amenable the work will be to automation. But an obvious question then springs to mind: what feature(s) of the work in question are we talking about when we talk about 'expert discretion'? And conversely, what features are able to be prescribed in advance? It's by answering these questions that we start to map out our framework for automation strategy.

Three types of work: programmatic, transactional, exploratory

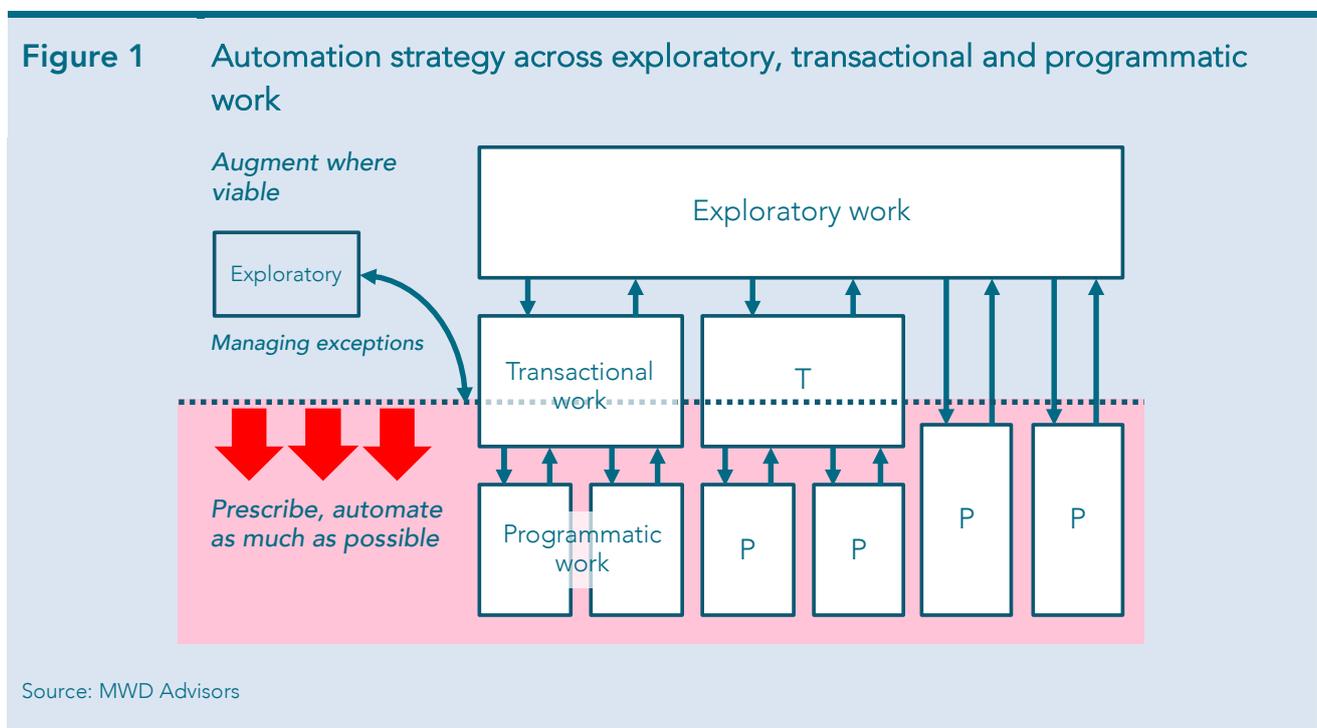
In practice we find that there are three types of work commonly found 'in the wild' in organisations, with expert discretion tending to be required for different blends of features across the three types:

- **Programmatic work.** With programmatic work, almost all features of the work in question – from tasks and decisions to the definition of work flow, work inputs and triggers (that signal work needs to be done), and resources required to complete the work – can be prescribed and designed in detail; expert discretion is only very rarely needed. This doesn't mean that the work is already automated, but it does mean that in theory it's possible for the work to be almost completely automated (with people only handling any exceptions that occur). Many types of clerical back-office administration work fall into this category.
- **Transactional work.** With transactional work, many of the tasks and decisions involved, and most of the flow of work, can be prescribed and designed in advance. The goal of the work is often also prescribed in quite a structured way. However human discretion and expertise is required to carry out some tasks, as well as in deciding when, how and where work needs to be progressed. A good example of a process revolving around transactional work is mortgage application processing – where the steps and decisions required are almost always definable ahead of time, but human discretion is often required to review decisions and minimise risks of fraud.
- **Exploratory work.** With exploratory work, the goal of the work will be defined ahead of time, but the set and sequence of tasks and decisions needing to be performed, and the people or roles needing to perform them, are very unlikely to be decidable ahead of time. There may be some high-level waypoints or milestones that are common to a particular type of exploratory work (perhaps to ensure regulatory compliance, or management approval) but they provide a very loose, rather than tight, structure. In exploratory work, as the label suggests, the overall experience for both the work participants and the 'customer' of the work is that of a set of possibilities being explored, rather than a recipe being followed. Common examples of processes revolving around exploratory work include complaint resolution and fraud investigation.

Automating and augmenting work across the three work types

When we look at organisations and their end-to-end business processes, we see a common big-picture pattern: the work within business processes tends not to be of the same type all the way through. Instead, processes tend to be composed of our three types of work. Specifically:

- The work carried out at the 'front end' of business processes – the work that's driven from external (customer, partner, supplier) requests and interactions – tends to be somewhat exploratory and needs some degree of expert discretion to deliver the best result.
- In these situations, the work that flows 'downstream' as a result of initial interactions with customers, partners and suppliers – the work that's involved in taking actions and making changes in various business systems – tends to be more transactional or programmatic in nature.
- Where work not driven directly by external parties is exploratory in nature (fraud investigation or compliance checking are good examples), the work that flows downstream – the work involved in recording outcomes, updating reference systems and so on – is still transactional or programmatic.



Considering again how our three different types of work are automatable to differing extents, a clear strategy presents itself (as shown in figure 1):

- Analyse the work that gets done today in your key business processes to identify which parts of the work are exploratory, which are transactional and which are programmatic.
- Where you find programmatic work (and you will typically find it 'downstream' from work that requires more human discretion to complete), explore opportunities to drive automation into this work.
- Where segments of your business processes are more exploratory or transactional in nature, look to use automation technologies to *augment* the work of human knowledge workers – providing smart assists for those people, and so putting the 'human in the middle' of work, rather than trying to automate what humans do.

Automation gets intelligent: specific roles for specific technologies

Now we've looked at a high level at how you should look to use technology to augment or automate work across our three work types (programmable, transactional and exploratory), in this section we dig more deeply to look at the roles that specific technologies can play.

Robotic Process Automation (RPA)

Robotic Process Automation (RPA) is a business and technology practice that deals with the configuration and management of software 'robots' that act as synthetic application users, automating highly repeatable, highly structured tasks across business software systems. It provides a non-invasive alternative to the creation and use of specialised integration APIs, or programmed integration by other means (for example hooking into underlying databases via triggers, or hooking into application code directly).

There are two common ways in which the core automation technology is delivered: 'attended' and 'unattended' RPA. Both these implementations of RPA have value, but it is slightly different in each case. In 'attended' RPA, the robots that automate interactions with users' desktop applications execute within each user's desktop environment itself – working 'side by side' with each user. In 'unattended' RPA, robots execute on a separate server or virtual machine (or multiple of these), automating interactions with applications 'behind the scenes' and carrying out their work when triggered by other software systems.

Unattended RPA

Unattended RPA technology is particularly useful today as a way to automate and orchestrate frequent, highly routine and rules-based interactions with IT systems within programmable work, where those systems don't have handy integration APIs. However, RPA isn't the best option in every situation. For one thing: although RPA bots can work faster than humans as they 'type' and 'read' information through an application's user interface, their overall performance is limited by the speed with which those applications can render and update through their user interfaces. For this reason, unattended typically gets used to automate clerical administrative work in the delivery of finance and accounting, HR and IT services, rather than automating aspects of core processes where minimising latency is crucial (such as the execution of digital products and services).

Attended RPA

In attended mode, the goal is not to have RPA bots automate complete tasks end-to-end. Instead, users invoke RPA bots themselves at specific points in their work, using them as 'automated assistants'. At the end of a process to respond to a customer enquiry, for example, a human user might invoke an RPA bot to update a handful of legacy reference systems to close out the process; in a call centre environment, an agent might invoke a bot to automate logging-on tasks across multiple systems and applications at the start of a shift.

Attended RPA technology is a useful augmentation for workflow deployments in automating aspects of transactional work, and for case management deployments in automating aspects of exploratory work: it enables organisations to increase the 'coverage' of automation in processes that are fundamentally still driven by knowledge workers, to include 'long tail' automation opportunities where implementing more sophisticated (but complicated to implement) integration technologies is difficult to justify.

AI at the edge: Intelligent Capture and related AI services

In situations where work is being done to fulfil requests from customers, partners or suppliers (for example, fulfilling a shipment delivery address change request) it's often the case that such requests arrive in a variety of formats over a variety of channels – for example requests might arrive via email, fax, PDF or via a web form.

In situations where work inputs arrive in the form of documents, you should consider using intelligent capture technology to digitise the data in those documents and identify and extract key data fields. Intelligent capture technology, for example, can automate the digitisation of contracts and invoices, and automatically identify extract key terms and items (like contract signatories, subjects and key clauses; or like PO numbers, invoicing addresses, totals and products, payment terms and so on).

Third-party commercial or open-source AI libraries can further enhance the flexibility of this technology by improving the accuracy of image or document classification, as well as improving the accuracy of detecting important features in the data that will be used to drive subsequent decisions (and improving the later discoverability of assets, by enriching document metadata).

AI in the flow of work: recommendation models

Most decision management capabilities available today are designed to support teams implementing decisions using prescribed business rules that always produce a definitive result (for example: if the customer is a premium customer, have an agent call the customer back; otherwise, send an email response to the customer's request). However, beyond automating decisions with prescriptive rules, there's a lot of potential value in guiding decisions through automated *recommendations*, based on predictive statistical models. For example, a customer retention team might build a model that predicts how likely customers are to switch providers (or 'churn'), based on perhaps dozens of data features related to a customer's record and their service usage history with the organisation. These models don't provide definitive results: they provide confidence levels. So, for example, a predictive customer churn model might highlight that for a given customer with a given history, that customer is currently 81% likely to churn. A trained call centre agent can then use that information to make an informed judgement.

It's now possible to apply predictive models in real time, in the flow of work, as decisions need to be taken.

Historically, the cost of computing meant that predictive model generation and maintenance would typically be carried out by a specialised team using dedicated infrastructure, and that models were only applied to business data periodically. However in recent years, the technology (and the cost of technology) has improved drastically – so it's now possible to apply predictive models in real time, in the flow of work, as decisions need to be taken. This means that where predictive models used to be used primarily to do things like shape email marketing campaigns, they can now be used interactively as part of conversations with customers.

In the context of transactional and exploratory work, where work features can't necessarily be prescribed through up-front design, there are opportunities to use predictive models to make recommendations that guide the flow and results of work. We've seen examples of predictive models being used to:

- Recommend best-next actions for workers to take as they progress cases – matching case features against those of historical cases to recommend courses of action.
- Recommend colleagues who can help with progressing a case, working on a task or making a decision – based on the features of the case and the skills of all available colleagues provisioned on the system.
- Recommend guidance documents and procedures that people working on a particular case stage or task should use, based on case features and the use of reference documents across historical cases.

Case study: psHealth

psHEALTH, a technology provider specialising in process automation solutions for healthcare institutions, has been working with Appian since 2009. Co-Founder Ingolv Urnes shared details of his team's work exploring opportunities to drive intelligent automation for clients.

psHEALTH is finding more and more clients with a general interest in automation – starting with an interest in the digitisation of paper forms and correspondence; then broadening out to the use of case management approaches to distribute, co-ordinate and track casework; and then to automate steps within the progress of cases.

One way that psHEALTH is responding to this interest is with its new Advanced Referral and Triage (ART) offering. Designed to help automate the clinical referrals process, ART takes incoming referral requests from a variety of channels, then uses a combination of off-the-shelf OCR, custom-built machine-learning prediction models and business rules technology (using Appian rules) to analyse incoming requests and make automated referral recommendations to staff. The system is now also able to make fully automated referral decisions in some cases (where clinical risk is low), should a client want to do that. In addition, in implementing ART with clients, psHealth found that staff were spending significant time copying and pasting information from booking systems, looking that information up on other systems, then using the result to input data into further systems. With a very mixed systems landscape and incumbent systems suppliers without incentives to modernise, a middleware-based approach to automating this work was infeasible – so psHEALTH explored the option of using RPA technology. Where clients' staff were spending perhaps 2-3 minutes performing this cutting-and-pasting for each situation, an RPA bot would do the same work in around one second.

Now, with ART, psHealth clients have processed around 500,000 referrals, and psHealth estimates that ART can save between 5-10 minutes on each interaction. With a client typically processing between 100-200,000 referrals each year, that adds up to a huge time and cost saving. However, as well as savings, there are other benefits: not least, the ability for clients to have access to much better operational intelligence. In a market where performance metrics are difficult to come by in a timely fashion (because of the widely distributed and federated nature of operations) the system, based on the Appian platform, enables clients to quickly get better operational insights, much closer to the point of referral.

psHEALTH hasn't experienced skills development challenges in its pursuit of intelligent automation with its ART product. As Urnes says: "this isn't a complex technical job - it's not about hard-core engineering, but more about hard-core curiosity." In other words: getting long-term value out of using intelligent automation strategies is about being prepared to really look at the data that an implementation provides, understanding what the data is telling you and acting on that in creative ways to improve over time.

psHEALTH is now exploring more opportunities to drive automation improvement through machine-learning models – specifically, exploring whether machine-learning models can reliably recommend interventions based on past knowledge of interventions and outcomes. To date ART has been improved over time using what Urnes calls a "slow learning" approach – updating triage rules manually based on expert observation, rather than driving recommendations from unsupervised learning algorithms.

Based on psHEALTH's experience, Urnes offers two key recommendations to others. First, know the problem you're trying to solve. Make sure you start with a clear measurement of the end-to-end processing time for manual processes, with breakdowns for individual steps and stages, so you can prioritise which work to automate. Second, start simple. Before you get involved with complicated AI technology, start with encoding prescriptive rules to drive recommendations, and seeing how this delivers value. Then, where this falls short, work with simple AI options (like machine-learning enabled OCR) to learn how these things work before diving into unsupervised learning algorithms and full end-to-end decision automation.

Case study: IQVIA

IQVIA is a multinational life sciences company that operates clinical trials and other research and development programs on behalf of pharmaceutical companies, and has been working with Appian since 2015. The Appian platform is deployed in 60 countries to orchestrate cross-industry processes as well as drive internal workflows.

As IQVIA's CTO Malcolm Postings explains: "We are moving towards an operation that can combine assets and capabilities on demand into collaborative, autonomous, real-time services which can link dynamically to future loosely-coupled, outcome-based, dynamic pathways which are more customer driven. To do this, we're having to create a new mindset for designing the workflows of the future. Specifically, as services become more customer focused, 'real-time' service discovery, adoption and aggregation will be key. We also need to break down large, tightly-coupled complex processes into separate plug-and-play, best of breed, autonomous loosely-linked services – and we need a platform that will support us in doing that."

Postings and his team have worked to find ways to combine use of the Appian platform with complementary automation-focused technologies, and Postings cites two examples of significant initiatives.

The first example is in 'Regulatory Startup' – the set of regulatory activities required to start a clinical trial or other pharmaceutical research program. It's a highly document-intensive process that can take many months. Using Appian's platform IQVIA managed to reduce clinical trial setup time by 25%, but the company felt it could do better. Because IQVIA operates in multiple countries, when it operates regulatory startup activities on a global basis, very often key documents will be written in local languages. Now, IQVIA uses machine-learning models to auto-translate documents, invoking these models as services from Appian workflows.

The second example is in the processing of large volumes of paper documents into clinical trial 'master files' (digital collections of all essential documentation relating to a trial for regulatory scrutiny). These documents are of varying quality and formats aren't always known ahead of time. In order to automate as much document processing as possible, IQVIA has built a 'processing pipeline' using machine-learning prediction models, off-the-shelf Optical Character Recognition (OCR) technology, XML processing, and natural language translation services. Together, this pipeline ingests digital documents; analyses their structure and quality (pagination, orientation, and so on); extracts the text; transforms the text into structured XML; translates text into English; and classifies them according to type. The pipeline supports 10 languages today, and the IQVIA team is working on further broadening this support. Google's Natural Language API is the core of this part of the pipeline, but IQVIA has augmented it with home-grown pre- and post-processing, which improves accuracy from around 50% to around 85% (which is good enough to drive the document classification algorithm).

IQVIA's found Appian skills development straightforward – teams in the US and India have been trained with no challenges – but team development in AI and machine learning has been more difficult. Nevertheless IQVIA built a team of data scientists in the US and India over 9 months. The bigger challenge for IQVIA has been in the area of change management: working initially without the benefit of a formal change management stream, the team found that its systems were poorly received. Strengthening this is now a major priority.

Looking to the future, IQVIA's main initiative is to take the AI and machine-learning technologies it's used so far and 'harden' them into a container-based platform it can deploy at scale – potentially making the platform available to other commercial entities, as well as to its own internal projects.

Postings offers two lessons learned from his time driving intelligent automation. First: it's easy to ignore 'soft skills' – but they're crucial. As well as the right technical people, you need people who can understand the future state you're planning and deliver it in a way that people will embrace. Second: it's vital that you think beyond initial project delivery, and consider how you'll support systems in the field over time. With technology capabilities maturing and expanding quickly, intelligent automation can't be a 'once and done' project.

Case study: University of South Florida (USF)

The University of South Florida (USF) is a large public university serving over 50,000 students across 14 colleges. It's used the Appian platform widely on projects across the university, since the platform's introduction to revitalise the institution's in-house custom software development capability around 4 years ago.

The Appian platform is already leveraged in the context of an agile development practice; development projects are highly iterative, delivering working code in two-week sprints. This cultural change has been imperative as the institution has worked to eliminate significant technical debt and demand backlog that built up over preceding years. Still, though, USF wants to find ways to get things done faster – both in software development terms, and also in terms of the operations that drive day-to-day experiences at the university.

Sidney Fernandes, the USF's CIO, explains that USF's business model places students and faculty at the core. Everything USF does to change its operations is designed to make students' and faculty's lives as easy as possible. Automation is seen as a lever that can increase operational efficiency, and so improve the experiences of students and faculty; but as well as that, it's seen as a way to improve auditability. With around 50,000 students, there are multiple opportunities to make impacts on their experience by improving efficiency; there's no grand ambition to "change the world" but rather to continually work on relieving operational pressure points.

Fernandes' team is now actively exploring opportunities to employ RPA to deal with what he calls the "last mile" of business process automation. RPA is a great tool to substitute for more heavyweight, complex custom integration where the conditions are right.

USF has conducted a RPA proof-of-concept in the HR onboarding process, integrating an Appian-based RP onboarding application with PeopleSoft: it took just a couple of days for the work to be done with Blue Prism. The alternative – using MuleSoft to integrate the system-of-record with the Appian-hosted workflow – would have taken much longer.

USF is also exploring the potential benefits of predictive analytics. It's already working with a technology partner, Civitas Learning, to predict in advance which students might drop out, based on their patterns of work submission and attendance; student advisors can use these predictions to nudge students 'at risk' and encourage them to step up. It's now looking at using predictive analytics in more areas – as well as integrating the Appian platform with the Twilio telephony gateway, enabling Appian applications to send SMS messages to students highlighted by Civitas' models as needing attention.

For Fernandes, the broader set of technologies being investigated by USF to drive digital operations is a "digital transformation ecosystem" rather than a seamless platform. Fernandes sees that the individual technologies in play (workflow, analytics, RPA, and so on) are maturing fast, but independently – so it's vital to continually evaluate each technology category on its own merits and look for robust integration points.

From a skills and practices point of view, the biggest challenge has not been the technology itself – "the technology will take care of itself". Rather, it's been in embedding the Appian platform and other complementary technologies effectively in an agile software delivery culture and process. USF has had to work hard to create cross-functional teams with real business stakeholder and client representation, and to make sure that all of the participants are properly educated about what agile delivery really means – such as the implications of delivering minimum viable products (MVPs) and subsequent rapid iterations.

USF's CIO offers two lessons to others looking to embark on similar initiatives. First, make sure business stakeholders are truly comfortable with the idea of incremental development, where value is delivered over time as requirements are discovered and elaborated – not all at once, in a (high-risk) 'big bang' fashion. Second, always focus on solving business problems, and make sure the problem that technology teams see is the same problem that the business sees – it's easy for technology teams to make false assumptions about the kinds of solutions business teams need for their problems.

Your next steps

The three case studies shown in this report help to highlight not only the kinds of ways that automation technologies can be applied; they also highlight important next steps and key considerations you need to take into account if you're serious about Intelligent Automation.

Foster an automation technology ecosystem

Don't fixate on finding one supplier that will provide an end-to-end portfolio of Intelligent Automation technologies. Instead, look for a core platform that provides a 'backbone' for Intelligent Automation and provides a variety of integration options and technology partnerships to enable you to quickly bring new technologies into the mix.

Remember that across the variety of automation technologies currently available and ready for enterprise use, different technologies play different automation roles across the features of different types of work. Be careful not to overstretch any particular technology, trying to use it to automate work features (or work types) where it's not suited. Be aware of the problems that each technology is best at solving, and be aware of the risks of overstepping those boundaries.

Don't forget: Intelligent Automation is about *business change*

It's perhaps tempting to think that Intelligent Automation is all about the application of technology – after all, it's about using technology to take aspects of work out of human hands and brains. This is a dangerous way to think, though. Intelligent Automation is very much about business change, for four main reasons:

- Unless your teams understand the business context for the work you're considering automating, you can't effectively prioritise or design work automation or augmentation projects.
- Unless your teams work in an agile, collaborative way that involves business stakeholders, you risk solving the wrong problems or delivering tools that present themselves in unintuitive ways.
- Unless you work right from the start to engage those people whose roles and tasks will be affected by work automation or augmentation, you risk those people rejecting the solutions you deliver.
- Unless you first explore your wider organisation for existing staff with relevant skills (for example, people with analytics or data science skills) you risk paying for new staff you don't need or hiring expensive consultants you don't need.

Review your assumptions and analyses regularly

Be aware that over time, the boundaries that you draw between exploratory, transactional and programmatic work are likely to shift. As this happens, it pays to be aware that more automation may be possible and profitable.

As automation technologies mature and are able to more economically automate more and more human-discretion tasks (and this will surely happen), you may find that work you would initially characterise as transactional (for example) could be reclassified as programmatic and automated straight through.

A good start is an annual architecture review process that explores how work is classified, reappraises technology state-of-the-art and investment cases, and determines whether reclassification of work (and therefore the application of new automation patterns) is necessary.