Cytora

Achieving digitisation and automation in risk processing - a build versus buy approach

Contents

Executive overview

Our world has been on a journey of inventing, improving and automating many aspects of our lives and work. A significant acceleration of this process happened two centuries ago with the beginning of the industrial revolution. It was then that we discovered how much more can be achieved through creating efficient workflows.

> Over the last 20 years we've been in yet another revolution - this time a digital one. The power of computing and the internet really unlocked what we thought was possible and opened new possibilities to further harness the power of automation.

When it comes to the insurance industry, it was and in many cases still is heavily relying on manual processes. It's for a good reason, after all, most of the work that insurers do rely on multi-step workflows that use a variety of different data sources that are specific to sectors and lines of business. Over time, the growing need to scale and increase the rate of profitable growth pushes insurers to revisit the most time consuming, inefficient and expensive processes and see where the benefits of automation can be gained.

At the heart of unlocking scalability that premium growth is decoupled from the expense growth lies technology. The way insurers can achieve that advantage is through digitising, and automating some of their processes. The technological gains can help minimise the manual, tedious tasks while funnelling the underwriting talent to drive productivity, efficiency and growth. Achieving that scalable growth quickly becomes the core competitive advantage driving market leadership and share price.

The two main paths to reach that state are turning inwards and building their own platforms or engaging a dedicated partner to fill this gap. This whitepaper will explore those two options in great detail to help insurers make the right choice for themselves and their organisations.

Where the inefficiencies truly lie?

One of the most important processes within the insurance business is risk processing (spanning new business submissions, mid-term adjustments or renewals), this is where all the revenue and growth are being made. To outline and design the transformational changes required we must first understand where the challenges are being had currently.

There are four main areas of waste across most insurance businesses:

Time

Underwriters are having to deal with many different forms of intake - through portals, emails, and presentations from different brokers using their own, non-standardised templates. 90% of the intake is in analogue format, and there's no distinction between new business, renewals or mid-term adjustments. All this leads to a time-consuming process of sieving through the data to understand what's what and rekeying information into decision-making systems like CRMs and Underwriting Workbench.

Talent

Underwriting teams are highly skilled professionals; analysing risks and deciding what the business is and isn't going to write. In a situation where they could spend up to a third of their time on risks that fall outside the appetite, their productivity isn't very high. Their expertise is wasted on menial tasks rather than making decisions on the risks that are within the appetite. In addition, risks regularly arrive to underwriters who lack the expertise to underwrite them resulting in capacity being spent on the wrong risks and slow turnaround times.

Data

Typically data being used across a workflow is keyed in manually from the broker submission and from manual lookups performed by underwriters (e.g. underwriters looking up data across multiple websites). Data is not centralised and operationalised into multiple steps in the workflow which creates inefficiency and poor data capture. Different data sources are required depending on the product, line of business and geography.

Control and oversight

Through the nature of incoming risks and manual evaluation there is a lack of control of where risks go (for example, which risks are straight through processed versus sent to an underwriter for review). The vast majority of the data doesn't get captured and often there is no record of the out of appetite risks. That means insurers find it difficult to analyse the quality of risks they are receiving and refine their appetite and products. The lack of standardised data across the workflow means leadership teams are unable to understand team performance levels and how to improve key metrics.

Combining all those factors paints a very good picture why so many insurers select risk processing workflows as perfect candidates for optimisation through automation and digitisation. Knowing where the challenges lie, however, is not enough to know what path of transformation to take. It's important to know what the 'ideal' state would look like. There are a number of aspects that need to be carefully considered.

The "ideal" state

Digitisation

Incoming requests are digitised for quick and efficient processing (no need for rekeying later down the line).

Combine multiple data sources

Multiple data sources are combined to turn risks decision-ready before they reach underwriters. Data sources are integrated centrally and activated in specific steps in the line of business workflows.

This enables insurers to use data cost-effectively and scale as multiple data sources are required and are different based on line of business and country. Insurers are also ambivalent about the types of data sources fulfilling data fields from the combination of the customer (submission) data, internal data and external data. They combine and standardise this so that they can make better use of it.

Digital workflows

Insurers easily combine data sources to create digital workflows. Multi-step workflows enable risks of different types to be separated and routed through different paths, for example, low complexity risks may be straight through processed, higher complexity risks can be sent to experienced underwriters to review. Workflows can be reused allowing insurers to standardise how they process risk submissions across different channels.

Continuous optimisation

Rules are continuously refined to better filter risks, prioritise risks and distinguish between risks of different types to optimise decision-making, helping insurers optimise hit rate in target segments and accelerate their target portfolio.

Productivity

As key workflows like new business processing, renewals, adjustments are fulfilled digitally, underwriter expertise can be applied intelligently - to specific steps in the workflow and types of risks - to maximise decision-quality. For example, underwriters receive decision-ready risks in the workbench and make key decisions on coverage, limits and exclusions to maximise decision quality. Because they are not spending time on rekeying data and out of appetite submissions they are able to quote and bind more risks.

Broker experience

Brokers receive quick responses from insurers across multiple channels enabling brokers to be more efficient in serving their clients.

Data Standardisation

As key workflows are digital and all data sources are integrated into the risk processing workflow, insurers have access to comprehensive, clean and standardised data for analysis and reporting.

Having all this information allows for a very informed starting point for the audit of the current processes understanding how it all applies to the specifics of the given organisation.

What does it taking path of building an in-house risk processing platform look like?

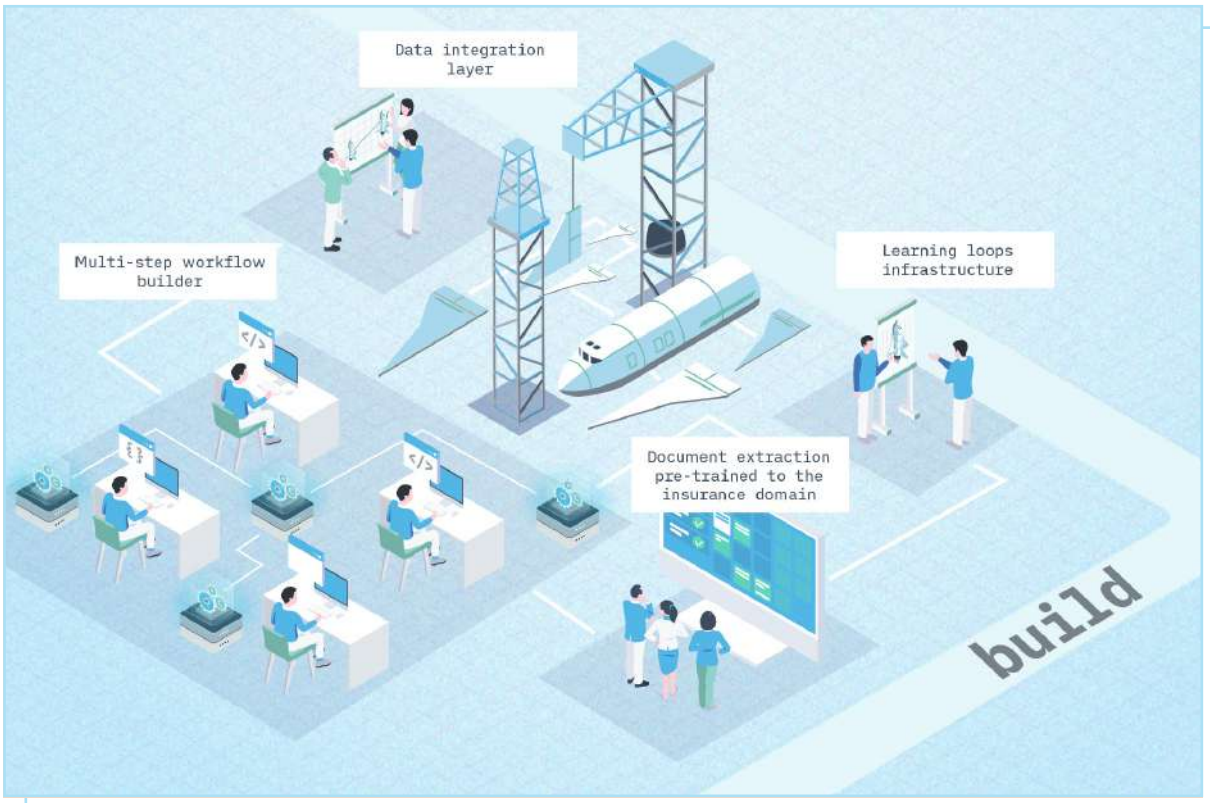

Best software building practices are very clear about the building blocks of success. It is paramount to gather requirements and the exact scope of the work required before the development starts. There are a number of architectural decisions that need to be made early on, and added changes later in the process can throw the delivery dates by a huge margin.

What in this case are the main components of a digital risk processing platform?

The four building blocks of a digital risk processing platform

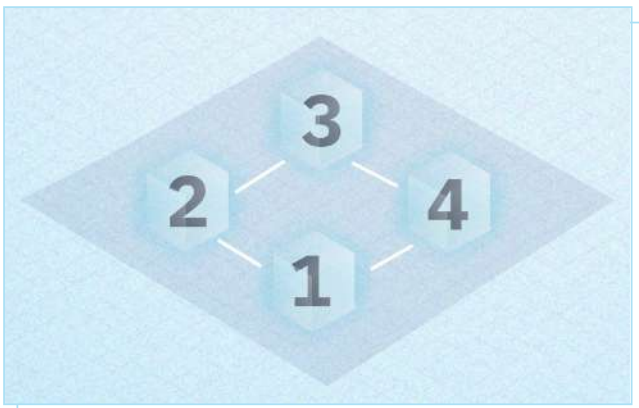

There are at least four technical building blocks of a digital risk processing platform.

When delivered using an agile methodology, each of these needs their own independent product squad to build and maintain.

Here's the scope of what each of the aspects of the product requires:

Multi-step workflow builder

This is a component that can create and execute multi-step digital workflows. The workflow builder needs to support the execution of different steps across digital workflows to process new business submissions, renewals, adjustments and other request types. Some steps are about fetching data and attaching it to work items (e.g. insurance submissions.) Others are about showing the work items to human operators to complete missing data based on confidence thresholds.

For capturing missing data from human operators, the requirement is for a dynamic and visually rich user interface that can render data forms, files and external data (including map data.) Because of the rapid rate of change when it comes to the availability of external data, this risk console should be configuration driven, so that it dynamically shows new data fields as they are integrated into the platform, without requiring a costly IT driven change every time a new data source or extracted target field is added. The interface also should be machine learning optimised, so that every action taken by human operators in the console is captured as training data to continuously improve automation levels. It should include a business rules engine to be able to respond quickly to changing market environments, e.g. enabling management to easily evolve triage rules that are applied for incoming submissions. Usually the workflows finish with output connectors that are responsible for inserting the enriched and extracted data into destination systems such as the CRM, Underwriting Workbench, Policy Admin System and Data Platform.

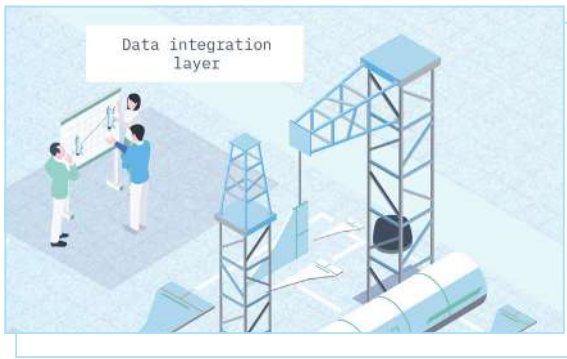

Data integration layer

The data integration layer flexibly connects data flows from upstream and downstream systems (including internal and external data sources, but also policy administration systems, email inboxes, CRMs, data warehouses, information extraction models, and risk scoring models maintained by data science teams). This flexible integration pattern enables neighbouring systems to be used in a digital risk processing flow. Bonus point if it comes with a data library of pre-integrated data sources.

Realistically speaking, this is not something that an insurer is able to easily build in-house. Even though an iPaaS (integration platform as a service such as MuleSoft, SnapLogic etc.) is a good starting point for this, It is not enough. This data platform must include advanced machine learning building blocks to perform entity resolution (e.g. to turn address strings into resolved address IDs, and to turn company names and descriptions into universal company IDs.) These indexing and query resolution problems are sometimes implemented as learning-to-rank tasks in machine learning, similar to the way Google turns a search query into the best search result.

Document extraction pre-trained to the insurance domain

The document extraction building block turns emails, PDFs and Excel sheets and other attachments into standardised digital risk records suitable for automated decision making. The extraction system should be pre-trained, and continuously re-trained.

The initial training to the insurance domain should be based on having already been exposed to a wide range of submissions and other document types from the insurance domain for years. This enables the system to achieve the best possible performance when detecting claims tables, schedules of values, and broker branch identifiers, which can come from the email headers and body, attached PDFs, or large Excel sheets.

Simply extracting content from PDFs is not enough: this component needs to include insurance domain optimized entity resolvers (e.g. to resolve human written sum insured currency amounts into machine readable values that a business rules engine can operate against.)

The system should learn over time: an exception handling user interface should handle missing data, and capture training data for the document extraction learning loops.

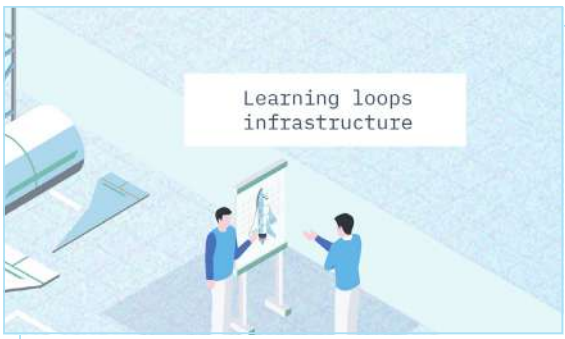

Learning loops infrastructure

Extracting and storing the data is not enough.
To close the learning loop, insurers must also have a way to operationalise insight back into the workflow. Currently a lot of data is lost inside of insurance companies, e.g. whenever a submission is rejected and the data is not key-ed into any systems.

Additionally, when the data does exist, those data points are often not captured in a machine learning ready way, so analytics efforts remain one-off efforts. A learning loop infrastructure should close the loop between:

> collecting the analytics and ML ready data in a data platform,

> deriving insights from it (e.g. suggested rules for the business rules engine, a trained conversion model,)

> serving these model artefacts back into the live flow of risks.

The skills and teams needed to deliver

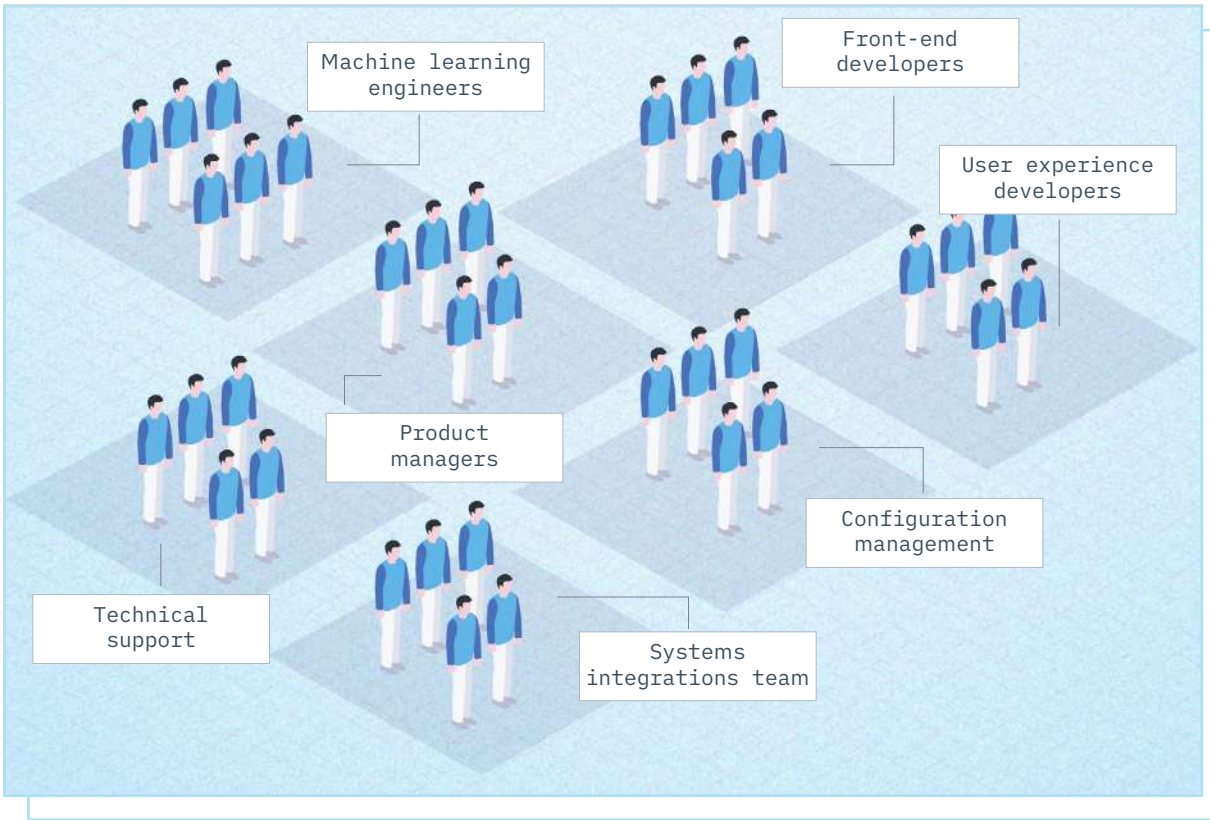

In terms of required skills, the configurable workflow builder needs significant front-end and user experience design resources, whilst the data integration layer, the document extraction system, and the machine learning platform require backend and machine learning engineers. In addition to these software engineers, each of the independent squads will be led by a product manager, and have access to a (possibly shared) cloud platform engineering team.

The initial build phase is even more taxing on resources, and relying on experienced product directors and senior cloud architects with document extraction, machine learning, and insurance experience.

In addition to the four core product development and engineering teams, there are three operational areas that will require staffing:

> Systems integrations (to insert the extracted and enriched data into CRMs, workbench, PAS and data warehouse),

> Configuration management (to encode and continuously update the workflows and business rules, and decide which internal data sources should be used, which requires data research expertise),

> Underwriter adoption and technical support.

The challenges of taking the ‘build’ approach

The above summary gives a sense why this type of platform is complex to build, involving many interconnected parts that only work optimally if they are compatible with each other.

In addition, because of the specialised skills involved in building a digital risk processing platform, growing and maintaining a world class cross disciplinary product team in this area is a significant responsibility and entails ongoing risks. One reason for this is the importance of human-in-the-loop workflows and interfaces. This requires advanced product design and UX capabilities (e.g. see the research of Ge Wang), which is not a core area for insurers so it will be harder to grow and maintain a team of experts in this area.

In-house build - outcomes on the key metrics of success

Time to value

Time to value is slow with an in-house option. This is because even if an insurer has the technical leadership and product vision in place to guide the development of a digital risk processing platform, the time to an MVP would still be at least a year, assuming no adverse unforeseen events like key employee churn, or slow adoption.

We also find that with complex software projects it often takes multiple attempts until a system is developed that truly addresses the use cases, which makes time to value less predictable for insurers that decide to bootstrap something starting from zero.

We have also observed that when insurers take on cutting edge technology projects like this one their key person risk increases significantly. This often happens because of time and budget pressures, which mean that especially complex machine learning and data systems built in-house are often dependent on the early team members who first built them.

Delivery certainty

Control is high with the in-house option, as the direction of the product is fully within the control of the insurer. However internal-know how is likely to be lower. How many full time experts on data, machine-learning and cloud native software development does the organisation have to be able to drive the development of a large scale project like this. Ideally the insurer would trust the project leadership into hands of someone who had built a platform like this before. If there’s lack of a very strong vision for the architecture and the deliverables an insurer is running into the risk of going through a lengthy and costly research and development phase before the actual engineering project can begin.

Run cost

Total cost of ownership is very high for the in-house option, because unlike dedicated software vendors, the insurer will not be able to spread their development costs across 10s to 100s of customers. The dominant expense will be personnel costs, across build and maintenance phases.

This means the insurer will have to carry the cost of the initial build, and of possible failed projects along the way, on their own and mean that the business case is harder to justify than working with a software provider.

Once live with a build option, insurers will need to develop and extend the roadmap themselves, shipping new features continuously to compete and differentiate against other software platforms that over time will develop attractive unit economics. Alternatively, buying a platform enables the insurer to take advantage of the forward looking roadmap and focus on how they steer the roadmap and exploit new features and configure the platform, shaping it to their unique business, to achieve competitive differentiation at lower cost.

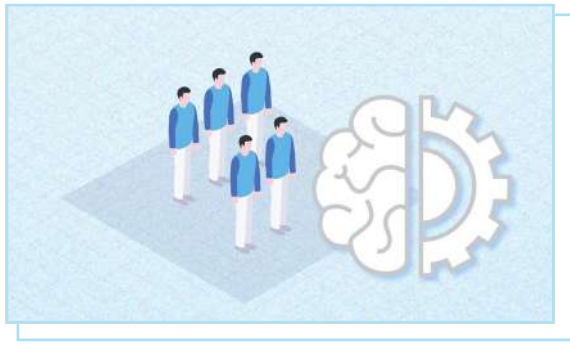

Personnel requirements

Personnel requirements are high, starting with the need for four cross disciplinary product squads to create and maintain the technology building blocks. The need for experienced senior technology leaders is particularly high during the initial build phase, which creates additional project risk due to fierce competition for some of the specialised skills involved.

In addition to the development teams, there is a need for business facing teams that handle systems integrations, configuration of workflows and triage rules, deciding on and configuring data integrations, as well as a team to support underwriters through initial adoption of the new processes.

In-house path - the summary

Building a platform that supports multiple workflows for multiple lines, across potentially multiple geographies is no easy feat. While it may seem easy to begin with ie when the scope is limited and is answering a small use case. However, the scope will naturally grow in size as the requirements are being gathered and other lines of business join in wanting to improve their processes. With that in mind, it's important to analyse the potential entirety of the project and the implications of taking the path of building an in-house solution versus partnering with a dedicated provider.

What does taking the path of working with a provider look like?

When working with a partner there isn't a 'build' stage, however a similar process of gathering requirements and use cases happens at the stage of selecting the right technological partner. In the next part of this white paper we will be focusing on what happens after the decision has been made and will cover implications for an insurer who decides to partner with a software provider to deploy a digital risk processing platform and will provide a flavour of what both parties should do to ensure a successful engagement. For clarity and ease of comparing the options of in-house-build with a ready solution, we will use the same evaluation criteria.

Partner solution - outcomes on the key metrics of success

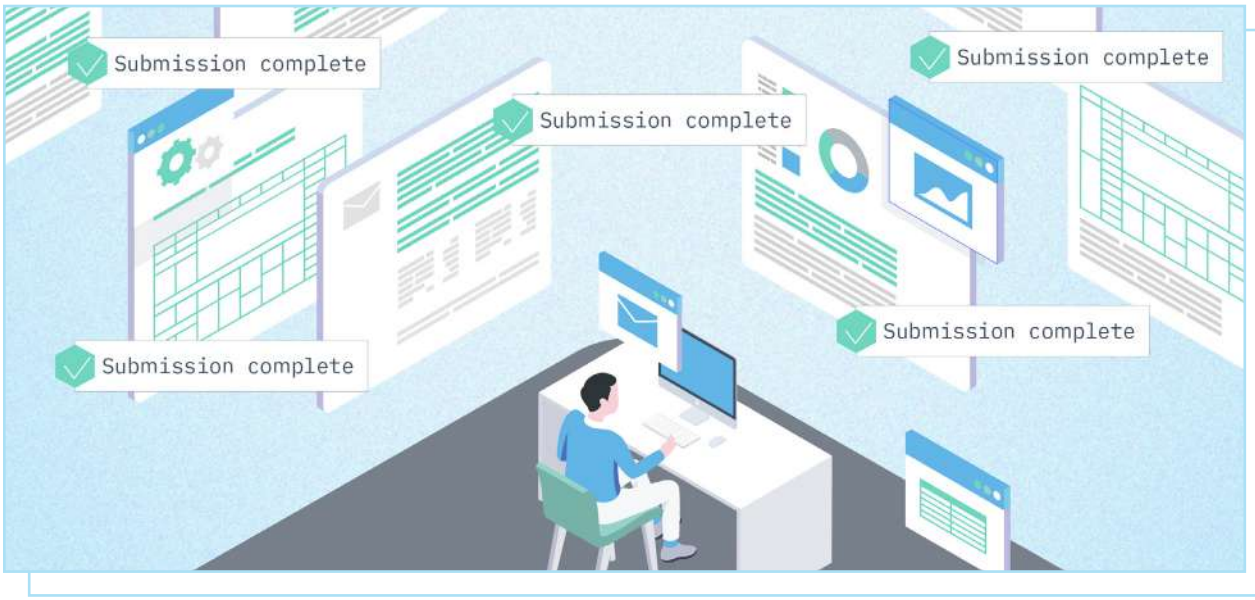

Time to value

When working with a partner time to value depends on the complexity of the business lines the product is deployed to.

In some smaller commercial classes where risks can be priced and underwritten using a small number of data fields ie. turnover or profession/activity will lead to a relatively straightforward codification set of rules to be applied. That makes the job of extracting, validation and augmentation simpler and more standardised.

Larger commercial lines ie 'commercial combined' with insurers who regularly sees commercial risks with multiple properties, specified contents and global turnover splits is more nuanced and therefore configuration takes more time.

Upon starting the partnership, the first period should be spent on clearly articulating the existing process and then working together with the partner to develop the desired future process.

A good software partner will have experience in the insurance sector but will need time to fully get up to speed on the specifics of the individual insurer. If the relationship is going to work and the product successfully deployed this period of time cannot be shortcut and it can be a significant portion of the entire deployment timeline. Once done however, it will be much easier for the partner to create a plan that will be very close to the final timeline. At this point the software partner can also highlight the leading indicators and long term benefit with a high degree of certainty allowing the insurer to quickly feel comfortable committing to the delivery.

Delivery certainty

Delivery certainty is usually quite high when engaging with a software provider. After the agreed onboarding and configuration time as well as training, the insurer should start seeing their underwriting teams being able to use the solution, and see the impact on efficiency and effectiveness.

In order to ensure the successful delivery, it's important to design the deployment process to optimise for the onboarding and ramping up the usage of the solution.

The best method of deploying software that is looking to drive significant long term improvements to business operations (such as digital risk processing) is to start with a number of small deliveries that can kick start the transformation. Once underway the delivery gradually gains traction as smaller successful deployments land and benefits are measured. The impact of trying to do a big bang deployment puts the insurers own resources under unnecessary stress and often causes a longer overall timeline when compared to breaking the project up. The impact of this on the insurer is that they will need to ringfence resource to support the deployment of the delivery.

Another helping factor that impacts the delivery certainty is the software partners' customer success team, who are able to provide operations, underwriting and technical teams with best practices, know how and expertise while guiding them to avoid the most common pitfalls and mitigate any roadblocks that may occur.

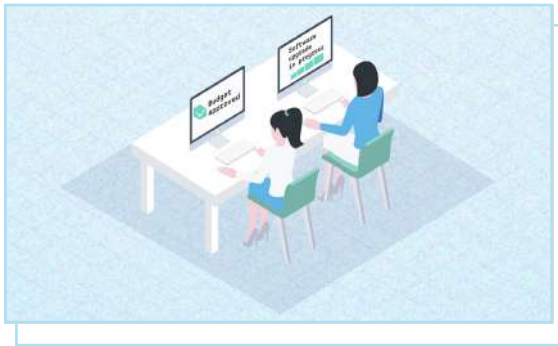

Run cost

When working with a vendor the bulk of the cost is known up front, subject to contract and potential data usage. There's little to no chance of the cost of the project/software spiralling out of control due to changing requirements, development delays and failed attempts of in-house building.

Another benefit is, behind the scenes, continuous software improvement over time. Throughout the duration of the contract the provider will continue developing the product which the insurer can take advantage of. That may be extraction capacity, extra data libraries or advanced routing systems. The insurer should rely on the partner to guarantee up time in an SLA (Service Level Agreement) and provide support.

When selecting a software partner it is important to clearly understand the fixed and variable costs of the delivery. In the specific example of risk digitisation then this could be broken into extraction, augmentation, entity resolution and rules processing. The most significant variable cost will be the volume of records processed and the additional data costs in augmenting the risks with additional data to meet the insurers needs. During deployment it should be expected that these costs can be confirmed and forecast with a high degree of certainty so that the insurer can manage growth effectively and ensure the business case continues to make sense.

It is likely there will be an overall platform cost, which will cover access to the services and possibly a per user charge.

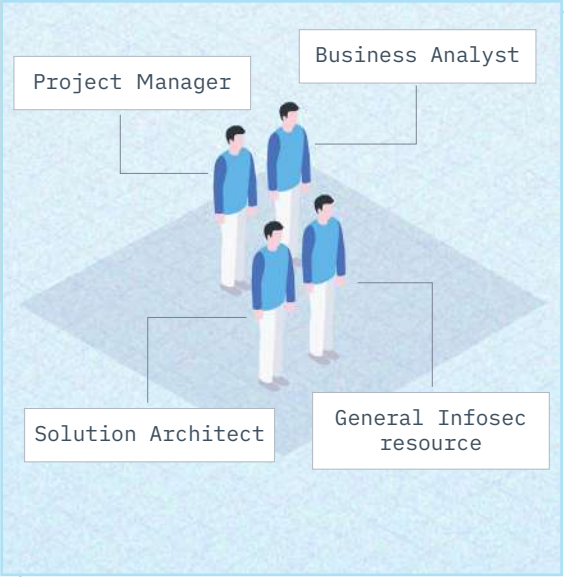

Personnel requirements

There's little cost and implications in terms of personnel when it comes to working with the vendor.

During delivery success can be enhanced by ensuring there is appropriate project delivery resource in place to support the vendor deployment. This would generally consist of:

> Project Manager

> Business Analyst

> Solution Architect

> General Infosec resource

Alongside this delivery resource success can also be improved by ensuring that there is someone with significant experience of the insurers operational processes involved.

Throughout the configuration phase and once deployment is complete, the key resource to achieve success will come from underwriting leadership and the underwriters themselves. The group that should expect to see the benefits and will ultimately be most impacted by the deployment of a digital risk processing system will be underwriting. Early engagement and continued interaction with this group is of paramount importance.

It is important, for the success of the project delivery, to factor in the right amount of time spent on training and onboarding end users.

The partner should make available a few key managers under the customer success function that can support stakeholder management, liaise across all teams and act as a central escalation point during the delivery and then ultimately throughout the lifetime of the relationship. Typically those roles could look like the following:

> Delivery Manager

> Delivery Analyst

> Product Architect

Partner solution - the summary

When engaging with a partner to deliver transformational changes it's important to remember that the success of the project lies both on the software provider as well as on the insurer, although to a lesser extent. In a contrast to in-house solution, where the end-user buying is given or it's been built by the time the delivery, the purchased software requires getting the teams on board of the change to ensure the maximum return on investment.

Conclusion

As digitisation is increasingly a way to achieve genuine competitive advantage insurers are understandably looking for ways to improve their processes, increase efficiency, and drive growth through superior risk selection. It is not surprising that the desire to eliminate waste and focus on profitable risks is high in insurance companies; that in turn enforces a change in risk processing flows.

When investing in any digital risk processing solution, whether in-house built or SaaS, it's important to go have a long and hard look at what are the business goals, by what time they need to be achieved, at what scale across the organisation, at what cost and at what level of delivery certainty. Having clarity on the objectives enables insurers to make clear sighted decisions and ensures concrete progress can be made.

About Cytora

Cytora is the configurable platform that enables commercial insurers to process risks at greater efficiency and accuracy. Cytora digitises every incoming risk, augments them with additional data sources, evaluates them against multiple rules, including appetite and priority rules, and routes them to downstream systems for automated or manual underwriting.