

A Buyer's Guide to Conversational Data Intelligence

Scaling services, enhancing efficiency and the customer experience with AI

Contents

What is Conversational Data Intelligence?	4
Why does Conversational Data Intelligence matter?	6
The digital transformation landscape	7
Cloud NLP APIs	8
Chatbot Tools	10
Conversational Data Intelligence	12
Conversation Intelligence	14
Customer Experience Management	16
Intelligent Document Processing	18
Robotic Process Automation	20
Data Science and Machine Learning platforms	22
Surveillance	24
Process Mining	26
Workflow / Case Management / Business Process Management	28
The requirements of a Conversational Data Intelligence platform	30
Natural Language Understanding	32
Model choice	33

Accuracy	34
Data efficiency	35
Model Operations	36
Training User Interface	38
User Support	40
Monitoring & Alerts	41
Availability	42
Security	43
Integrations	44
Implementing Conversational Data Intelligence	45
Step 1: Discover	46
Step 2: Discuss	48
Step 3: Deploy	50
Build or buy?	52
Re:infer Conversational Data Intelligence	53
Key takeaways	54

What is Conversational Data Intelligence?

Conversational Data Intelligence is a category of enterprise software that combines machine learning, natural language processing (NLP) and the specialist knowledge of employees to drive understanding and automation in business communications and conversations.

Through machine-human collaboration, Conversational Data Intelligence creates structured data from masses of unstructured enterprise communications - such as email, chats, calls and tickets. Through a zero-code platform, it gives users the ability to analyse previously hidden business processes and automate low-skill, manual communications processes that used to depend on human reading comprehension.

Why does Conversational Data Intelligence matter?

Conversational Data intelligence fills a major gap in business understanding and process efficiency. Business runs on communication. Almost every operational process and customer workflow involves human interaction at key points of the process. Customers reach out when they want something, employees collaborate to get work done.

However, the scale of digital communications is growing too fast for human workers to process efficiently. The average employee now sends or receives 126 emails a day, and up to 40% of their time is spent in Outlook alone. It's a massive drain on productivity and staff morale - **30% of workers see their inbox as the biggest distraction from actual work, and 22% say they want to quit their current role due to exploding email volumes.**

It isn't just employees who suffer. Service levels are impacted, handling times increase, business process efficiency slumps, and customer dissatisfaction grows. The problem is only getting worse and no business can afford to let it continue. However, the communications challenge isn't a problem that they can solve with humans alone.

Conversational Data Intelligence provides a solution. It gives organisations easy access to advanced NLP and machine learning capabilities that anyone in the organisation can use. These zero-code platforms analyse masses of unstructured communications at speed and at scale, pulling out the most valuable information and creating structured data that can then be used for analysis and automation.

This is game-changing for process insight and efficiency. For the first time, businesses have total visibility and the ability to understand everything that is happening in their comms channels. They also have the raw data needed to create intelligent automated workflows from manual comms-based processes. This frees employees from the chore of manual comms processing, letting them focus on only the most valuable work.

The digital transformation landscape

The enterprise has started to look beyond traditional automation solutions. They have new imperatives and priorities - to boost operational efficiency, improve service delivery and customer experiences, and create fresh business value - alongside the familiar need to cut costs. Old rules-based systems aren't up to the task - expensive to run and update, inflexible and limited to automating only simplistic manual tasks.

Businesses need new ways to scale automation to every corner of the enterprise, even to its most complex processes. But to achieve this level of hyperautomation, business leaders must leverage the latest AI and machine learning tools for intelligent automation. A new technology stack is emerging, comprising the best-of-breed solutions in AI and automation.

Conversational Data Intelligence sits within this wide ecosystem of technologies. It is one of many tools used to drive efficiency, competitive advantage and digital transformation. However, it plays a unique role within the solution stack being the only tool designed to understand the unstructured natural language used in digital communications. This allows the category to serve as a crucial intelligence and integration layer, linking the intelligent automation stack together and allowing each technology to be scaled across the enterprise.

An overview of the major technologies within the digital transformation landscape is provided:

Cloud NLP APIs

OVERVIEW

Most major cloud vendors offer cloud-hosted NLP models and APIs that are production-ready for commercial use. These most often form the building blocks of NLP systems built in-house.

Cloud NLP APIs save enterprises the cost of building basic NLP functionalities such as Named Entity Recognition (NER), sentiment-analysis, text classification, summarisation, question answering, translation, language detection, tokenisation, and Part-Of-Speech (POS) tagging. However, these provide little value on their own, and must be integrated with a separate NLP solution before they can be properly used by the business. Large teams of data scientists, engineers and developers are needed for the creation of these systems.



PRIMARY USE CASES

- Diverse NLP use cases developed in-house



PREFERRED VENDORS

- AWS Comprehend
- Azure Cognitive Services
- Google AutoML
- IBM Watson
- Salesforce Einstein



ADVANTAGES

- Easy access to high-performance NLP algorithms
- Reduce the DevOps side of machine learning in production
- A low-risk method for testing new cutting-edge models without installing anything locally
- Reduces ML model development time



DISADVANTAGES

- High initial investment due to heavy data engineering requirements, with need for data engineers to build pipeline and engineer into processes
- Slow time to value as many features (including a functional UI and connectors to other systems) have to be built before the solution can be deployed
- Heavy dependence on data science and AI expertise to create a workable solution
- Follow a 'take it or leave it' model, with no way to optimise data

Chatbot Tools

OVERVIEW

Chatbots, also known as conversational AI, are a category of products used to deflect simple requests and workflows from their service teams. Chatbots can automate simple workflows, such as answering frequently asked questions and processing basic requests like password changes.

Chatbots are typically rules-based and cannot deviate from their fixed activities. Some, but not all, are designed to hand off the interaction to a real person when they cannot complete the request.

→ PRIMARY USE CASES

- Support request automation
- FAQs and query automation
- Identity and verification

→ PREFERRED VENDORS

- Cognigy
- Druid AI
- Google Dialog Flow
- IBM Watson Assistant
- Microsoft LUIS + bot framework
- Amazon LEX
- Rasa

→ ADVANTAGES

- Increased productivity from agents due to contact deflection
- FTE savings
- Average handling time reduction

→ DISADVANTAGES

- Can disrupt and devalue customer journeys, especially when gating off human interaction from the customer
- Only supports chats and is unable to understand longer form messages such as emails and tickets
- Performant across only the most basic requests and workflows
- Expensive to set up, maintain and scale
- Limited or no analytical capabilities
- Requires an additional channel - often implementers see an increase in volumes as it's now easier to contact them

Conversational Data Intelligence

OVERVIEW

Conversational Data Intelligence, also known as NLP for Mass Conversations Intelligence and Conversational Intelligence as a Service, is a new class of enterprise software. It combines machine learning, NLP and human employees to drive understanding and automation in business communications.

Conversational Data Intelligence creates structured data from masses of unstructured communications. This enables users to discover and analyse previously opaque business processes, and automate conversation-based manual processes that depend on natural language understanding.

→ PRIMARY USE CASES

- Communications Mining - extracting insight from communications data
- Automation discovery
- Automation of conversation-based workflows
- Auto-triage
- Auto-routing

→ PREFERRED VENDORS

- Re:Infer

→ ADVANTAGES

- No machine learning expertise required
- Enhanced business and customer insight - better customer and employee experiences
- Improved operational efficiency
- Integration with RPA permits the automation of complex conversation-based workflows - like responding to an email or auto-routing a request to the correct team
- No code
- Fast and easy integration with all other intelligent automation solutions
- Rapid implementation and model training time

→ DISADVANTAGES

- Does not support document understanding - but this can be solved by integration with a suitable OCR solution
- Conversational Data Intelligence is an enterprise-scale solution and works best with a large volume of communications, but you don't need a large amount of training data

Conversation Intelligence

OVERVIEW

Conversation Intelligence solutions use a form of NLP to record and understand the spoken words of humans, picking out the important parts of the conversation for later review and analysis. It is most commonly used in sales and customer service functions to streamline the manual reviewing of calls and the quality assessment process.

→ PRIMARY USE CASES

- Streamlining manual call analysis

→ PREFERRED VENDORS

- Gong
- Chorus

→ ADVANTAGES

- Improved sales outcomes
- Fast and easy implementation

→ DISADVANTAGES

- Direct ROI is difficult to measure
- Limited ability to customise

Customer Experience Management

OVERVIEW

Customer Experience Management (CEM) solutions provide end-to-end tools for managing the customer experience. They help users understand customer feedback and behaviours across every stage of the customer journey, and implement strategic plans to improve customer satisfaction, loyalty and advocacy.

CEM most often provides analytics on surveys and customer feedback channels. Some solutions also provide analytical support for other channels such as chat, tickets and social media.

→ PRIMARY USE CASES

- Feedback analysis

→ PREFERRED VENDORS

- Clarabridge
- Medallia
- Qualtrics

→ ADVANTAGES

- Improved customer intelligence and customer outcomes
- Improved customer loyalty and retention

→ DISADVANTAGES

- Simplistic, rules-based AI models provide no ability to tailor analysis around your specific customer group
- Limited ability to understand long-form customer conversations
- No automation capability
- Expensive and slow to implement across your different customer channels

Intelligent Document Processing

OVERVIEW

Intelligent Document Processing (IDP) solutions extract information from unstructured and semi-structured document formats. They use AI technologies such as computer vision, optical character recognition and machine learning to transform unstructured and semi-structured information into usable data for the purposes of classification, categorisation and automation.



PRIMARY USE CASES

- Unstructured data processing
- End-to-end automation



PREFERRED VENDORS

- ABBYY FlexiCapture
- Automation Anywhere IQ Bot
- Ephesoft Transact
- Kofax TotalAgility
- OpenText Intelligent Capture
- Rossum
- UiPath Document Understanding



ADVANTAGES

- Lower processing costs
- Faster document processing
- Easy integration with RPA for end-to-end automation
- Great customisation capabilities



DISADVANTAGES

- Limited success when used for unstructured natural language found in communications data
- Performance is impacted by varied inputs
- Limited intent recognition

Robotic Process Automation

OVERVIEW

Robotic Process Automation (RPA) software boosts the speed and efficiency of business processes by automating existing repetitive tasks in a workflow. Business users build and deploy software 'bots' that mimic the actions of human workers - interacting with multiple software applications to execute processes and tasks.

RPA enables reproducible and error-free execution of structured business processes, freeing your workers from monotonous, repetitive tasks and enhancing their productivity by enabling them to focus on higher-value activities.

→ PRIMARY USE CASES

- Order processing
- Data entry
- Onboarding

→ PREFERRED VENDORS

- UiPath
- Blue Prism
- AutomationAnywhere
- Microsoft Power Automate

→ ADVANTAGES

- Increased cost savings and efficiency due to straight-through processing
- Improved service quality due to reduced handling times
- Easy integration with other categories (incl. Conversational Data Intelligence and IDP) opens up more processes for automation

→ DISADVANTAGES

- Rarely provides real-time capability
- Lack of - or only rudimentary - NLP capabilities means RPA has little scope to automate business processes based in human conversation
- The cost of RPA bot development can rapidly spiral when their use is scaled across the organisation.

Data Science and Machine Learning platforms

OVERVIEW

Data science and machine learning platforms are enterprise machine learning tools that operationalise model development and lifecycle management. They aim to reduce the dependence on skilled data scientists to build machine learning and deep learning models. Solutions like AutoML systems are fed manually-labelled training data as an input and produce an optimised machine learning model as an output.

→ PRIMARY USE CASES

- Model governance

→ PREFERRED VENDORS

- AWS SageMaker
- Azure Machine Learning Studio
- Dataiku
- Datarobot
- Google AutoML

→ ADVANTAGES

- Cost savings due to accelerated model development by data scientists

→ DISADVANTAGES

- Usually single-purpose and lack tooling for annotating data and collecting training data
- Lack wider automation capabilities
- Manual data labelling means long training times and slow ROI
- Requirement for costly data science expertise
- General purpose means reduced performance in specialised use cases

Surveillance

OVERVIEW

Surveillance describes a broad category of monitoring tools used in heavily regulated industries for compliance purposes. A surveillance solution will scan employee conversations on digital channels, typically employing keyword matching technology to surface potential security or compliance risks. Once a potential breach has been identified, a compliance professional will usually be alerted for deeper investigation.



PRIMARY USE CASES

- Compliance monitoring



PREFERRED VENDORS

- Behavox
- Digital Reasoning
- Nice Actimize



ADVANTAGES

- Faster and more accurate compliance monitoring than manual investigation alone
- Improved regulatory adherence
- Many solutions provide reporting tools and functions for compliance reporting



DISADVANTAGES

- Lack of sophisticated AI means high false positive rates that compliance pros must manually filter out
- Single-purpose tool which usually can't be adapted for other use cases

Process Mining

OVERVIEW

Process Mining is a class of tools and solutions that help users monitor, improve and extract value and insight from structured processes in an enterprise.

Process Mining constructs and visualises business processes using data extracted from business IT systems. This presents users with an event map that helps them understand and improve known processes, and identify new ones.

→ PRIMARY USE CASES

- Automation opportunity discovery
- Process discovery
- Process monitoring
- Process enhancement

→ PREFERRED VENDORS

- Celonis
- UiPath Process Gold
- Minit
- Signavio (acquired by SAP)

→ ADVANTAGES

- Cost reduction due to greater automation opportunities
- Process understanding based on data rather than opinion
- Improved customer and employee experiences

→ DISADVANTAGES

- Process Mining struggles to analyse unstructured processes
- Limited insight into root causes (which can be solved with Conversational Data Intelligence)

Workflow / Case Management / Business Process Management

OVERVIEW

Workflow or workflow management tools help businesses to manage the work of their employees. They provide capabilities for many common processes and tasks, including case management, business process management (BPM) and customer relationship management. They typically help users to keep track of different workflows, efficiency and service levels under a single platform. Workflow, case management and BPM tools are most often used by business leaders to increase visibility, improve productivity and reduce operational risk.



PRIMARY USE CASES

- Managing work across business services functions, e.g. tech support, customer support etc.



PREFERRED VENDORS

- Salesforce Service Cloud
- Microsoft Dynamics
- ServiceNow
- Pega
- Appian



ADVANTAGES

- Increased process efficiency
- Enhanced visibility and management information



DISADVANTAGES

- Limited NLP capabilities mean workflow tools will struggle to analyse or automate certain processes
- Full implementation can be costly and slow, leading to only partial deployments and low coverage.

The requirements of a Conversational Data Intelligence platform

Once an enterprise has identified its need for Conversational Data Intelligence, it should decide whether it will build its own custom solution or purchase one that's already available. Before a final solution is made, it's important to consider the necessary components and requirements of a successful Conversational Data Intelligent platform.

Natural Language Understanding

Natural Language Understanding model choices are fundamental to the success of an implementation. A poor decision at the start of the design phase can cripple a system from the outset, drastically limiting its chances of a positive deployment.

To improve your chances of success, models should be considered across a number of dimensions:

Model choice

For the majority of use cases, Analytics and Automation models should be supervised. Purely unsupervised models can't be directed by human expertise, nor can their performance be tested or measured. This makes them unsuitable for enterprise analytics or automation workloads. Supervised models should typically satisfy the following requirements:

- Theme recognition: The ability to recognise high-level themes
- Sentiment recognition: Understanding of emotions - such as frustration and confusion - from text
- Intent recognition: The ability to extract specific requests and concepts chosen by the user
- Entity extraction: The ability to extract crucial, business-defined entities - these could include references such as customer names, case numbers, monetary sums etc.

Accuracy

The more accurate the model, the greater its ROI. The chosen model must be accurate in its predictions, otherwise analytics can't be trusted nor automations permitted in case of bad decisions and outcomes. There's also an important requirement for models to self-improve as data volumes grow. A model that can't remain accurate while it scales can only accrue negative ROI for the enterprise.

Model accuracy requires:

- Downstream application feedback loops to enable continuous model improvement and accountability
- Quick access for business process owners and analytics consumers to interrogate model performance, accuracy and false predictions
- Modern GPUs, and Transformer-based architectures that outperform humans on the latest NLU benchmarks
- A constant R&D investment to keep parity with current advances.

Data efficiency

Models must maintain accuracy and performance with minimal human supervision. If a lot of data is required to train a model, the cost of model training will quickly spiral and positive ROI will be hard to achieve.

Data-efficient supervised learning requires:

- Active learning - where human SMEs are queried for supervised labels through intelligent prioritisation
- Rapid retraining to ensure active learning uses all available SME annotations
- Unsupervised pre-training on large datasets representing the primary languages for your intended application
- Semi-supervised learning using unlabelled data and labelled data to minimise supervised data requirements
- Transfer learning to reduce task-specific training sets
- An easy-to-use UI that SMEs can use to train a model quickly

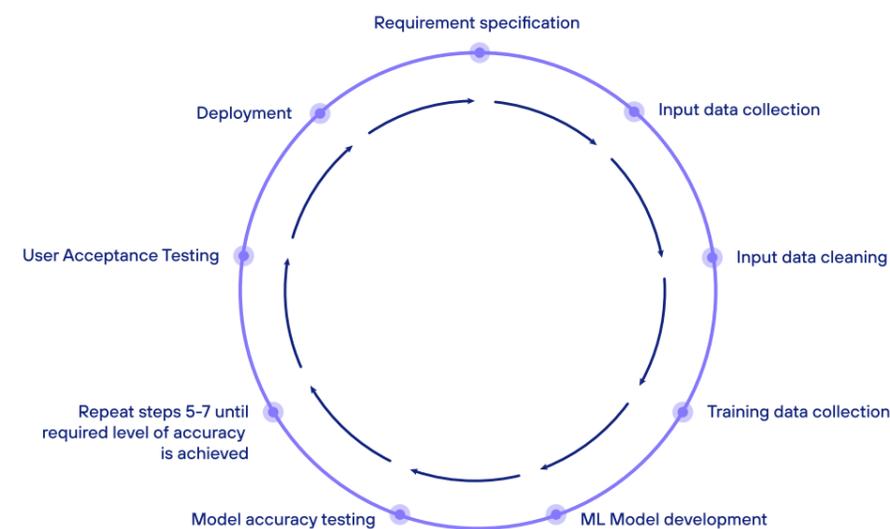
Model Operations

On selecting a model, a Conversational Data Intelligence solution needs the resources to maintain it over its full lifecycle - from specification to deployment and production.

Machine learning software development has an associated lifecycle development sequence. Ensuring control, compliance and oversight throughout this process is critical in the enterprise. However, because Machine Learning software is defined by data not just code, the machine learning operations (MLOps) sequence differs from traditional lifecycle development sequences.

The MLOps development lifecycle:

- Requirement specification - lead by the business
- Input data collection - lead by data engineers
- Input data cleaning - lead by data engineers and/or data scientists
- Training data collection - data annotation led by business, but tested by data scientists
- ML Model development - lead by machine learning engineers
- Model accuracy testing - lead by machine learning engineers and or data scientists
- Repeat steps 5-7 until required level of accuracy is achieved. This needs to be rapid - almost instantaneous for model development to be efficient and effective
- User Acceptance Testing
- Deployment



Enterprise machine learning development pipelines also require tooling for:

- Model versioning
- Model forking
- Model sharing and re-use
- Model determinism
- Model testing
- Model back-up
- Model scaling
- Permissioning against Develop, test and deploy actions

Training User Interface

The efficient collection of training data requires an intuitive, easy-to-use UI that business people can use to generate supervised data quickly and efficiently.

Data scientists and data engineers have the technical knowledge to build machine learning models, which they do in most business contexts. However, they often don't know the business, the customer, or the business process well enough to then train or test a model.

Therefore, training data annotation and model testing must be done by subject matter experts (SMEs) in the business who live and work in these processes every day. Rapid model development requires tooling and an easy-to-use UI that allows technical experts and non-technical SMEs to work together to rapidly create business-specific Conversational Data Intelligence models.

Many machine learning projects that require SME knowledge for annotation fail if the SMEs aren't given tools to make data annotation fast and efficient - and also if the ML models themselves aren't sufficiently data-efficient.

A Conversational Data Intelligence UI must:

- Represent model predictions and confidences - users can visually inspect what the model understands, correct and improve it
- Represent model accuracy statistics - the next best training actions are obvious
- Train in real time - SMEs don't waste time on training concepts that there's enough data for the model to already understand

- Show data in its context - metadata such as sender, file type and size should be visible in the thread
- Be adaptable - new concepts, themes, and entities are creatable on the fly in the UI

User Support

To start generating ROI, a Conversational Data Intelligence platform needs a fast adoption cycle. Users, whether they have technical ability or not, require training materials and tutorials to become proficient in the platform's use.

If adoption isn't encouraged during the first few weeks of deployment, it is unlikely that the platform will 'stick' and become a lasting part of employee workflows. Furthermore, if users do not receive help and advice from platform advocates or experts, it will take them more time to learn how to extract the most value from the platform.

Users and SMEs will need:

- A full training programme to understand the key concepts and requirements of the platform
- A centralised knowledge centre for self-service and reference materials
- Access to platform experts and success personnel for best practice and advice

Monitoring & Alerts

To provide long-term value, a Conversational Data Intelligence platform must be able to change and adapt rapidly. Communications, whether between employees and customers, are never static. Changing market dynamics and customer requests demand real-time monitoring and analytics, to identify and respond to emerging trends and themes in their communications.

The platform should have full alert capabilities. User teams should be notified of the most important events and request types, changes in volumes, as well as surges in activity and root-cause drivers. These are critical for Communications Mining, where the value of the solution comes from drawing insights from masses of data over an extended period of time, breaking it down to make it understandable and actionable.

Availability

A key non-functional requirement is the availability of the solution and its ability to recover from disaster events. By its nature, a Conversational Data Intelligence system needs to run GPU-enabled workloads on a horizontally scalable platform. Very high-level guarantees must be provided, delivering effective zero downtime both scheduled and unscheduled. Anything less can result in platform failures where data and training progress are lost, creating delays and wasting precious resource.

Security

Conversational data from employees and customers is highly sensitive. You must be able to train, test and deploy natural language understanding models with SME annotations against collections of different channels of data. Each of these channels may have different permissioning considerations - such as legal entities, geographies, user-groups and customer categories - that have to be taken into consideration.

A Conversational Data Intelligence platform, therefore, needs a permissioning model that supports all combinations of the following throughout the entire data pipeline:

- Users
- Data-sources / channels
- Datasets and - combinations of data-sources
- Organisations / teams
- Actions - train, validate, deploy and trigger

Security questions to ask any vendor:

- Where will our data be stored?
- Who can access our data?
- How often are penetration tests conducted?
- Will our data be backed up?
- Is our data encrypted at rest?
- Is our data encrypted in transit?
- How can I get data in and out of the platform securely to use in downstream systems and applications?

Integrations

Scalability is very important for a Conversational Data Intelligence solution. Many implementations begin with a smaller, confined use case, but the impetus to scale the platform across the business will grow as the returns become apparent. That's why it's crucial for the platform to have easily integratable APIs which enable it to be accessed across different teams and departments.

Quick and easy integration is also critical for extending the capabilities of the platform. There must of course be strong pipelines and connectors for ingestion and extraction of data from all desired communications channels. Communications Mining - the extraction of useful data from unstructured communications - is achievable with few integrations. However, automation use cases require a Conversational Data Intelligence solution with pre-built connectors to other RPA and automation systems.

Ultimately, a platform that can easily integrate with the rest of the intelligent automation ecosystem offers far more value, and has much greater potential to scale within the organisation.

The most important systems for integration:

- CEM
- CRM
- IDP
- OCR
- Process Mining
- RPA
- Workflow
- Mail and calendar server

Implementing Conversational Data Intelligence

Implementing a Conversational Data Intelligence solution has to be a carefully considered process. A poor decision at any stage can lead to problems down the line, jeopardising the success of the project as a whole.

At Re:infer, we have found success in implementing a three-stage implementation process - centred around opportunity discovery, collaboration between problem owners and Re:infer experts, and the actual deployment of the solution.

Whether you have identified a vendor platform, or have decided to start building your own Conversational Data Intelligence solution, this is our recommended implementation approach:

Step 1: Discover

- Detailed requirements review
- Technical validation

1. Detailed requirements review

Understand the opportunity for Conversational Data Intelligence in your enterprise. A vendor should work with you to understand the pain point at the centre, mapping out the related processes and identifying the main stakeholders affected by it. Often the opportunity for Conversational Data Intelligence will be much larger than it seemed before this requirements review.

Common questions

- What are your key strategic drivers?
- What are your short-, medium- and long-term objectives to deliver your strategic drivers?
- How does Data & Analytics support the delivery of your objectives, and what analytics capabilities do you have within your organisation?
- If you could listen to every conversation associated with these processes, what would you want to know and how could that support your objectives?
- Do you have an Automation Centre of Excellence within your organisation, and if so, how mature is it?

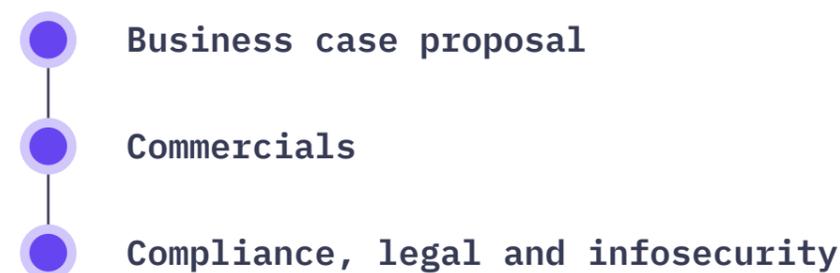
2. Technical validation

After scoping the business need for Conversational Data Intelligence, you need to understand the technical requirements. A vendor should work with you to build a detailed understanding of the technical environment and infrastructure the solution will be integrated into - as well as its future state. A key part of this is validating the non-functional technical requirements of the solution.

Common questions

- Do you have a cloud or on-premise deployment model preference?
- What core systems / applications are in your ecosystem ?
- Which of these core systems / applications support your communication channels?
- Do you have Natural Language Processing capability, and if so where is it deployed?

Step 2: Discuss



3. Business case proposal

Information and data collated during the Discover phase will inform a detailed and costed business case that's presented to the rest of the business, demonstrating the value of implementing a Conversational Data Intelligence solution. When working with a vendor, they should take the lead in constructing the proposal in collaboration with the business leader.

A breakdown of costs to implement the solution should be provided, covering potential services and annual license fees. If required, automation software and implementation costs should also be included to provide full visibility.

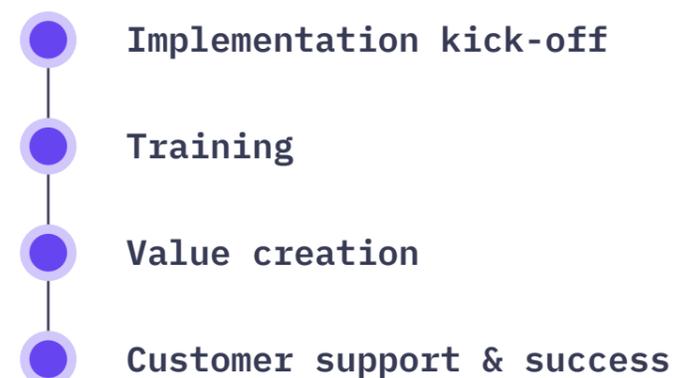
4. Commercials

If purchasing from a Conversational Data Intelligence vendor, this is the stage where both sides agree and sign a contract once the business case has been approved. Common starting points for the enterprise include Proof of Concept and land deals. Land deals often include a contract break helping the buyer minimise risk through a smaller initial use case, while being able to access the full capabilities of a Conversational Data Intelligence platform.

5. Compliance, legal and infosecurity

On completion of a contract, the customer and vendor will exchange all necessary legal and security information required. In highly regulated industries, this stage can last considerably longer.

Step 3: Deploy



6. Implementation kick-off

The Conversational Data Intelligence platform is deployed to your system. All connectors with data sources and complementary solutions are set up, and any technical implementation challenges are worked through. Working with an experienced vendor or skilled implementer will greatly expedite this process.

7. Training

The intended business users will be introduced to the platform and trained in its use. Platform advocates and experts should provide numerous tutorials and training sessions, while users are encouraged to practice model training in-platform during the working day. Before using the solution actively in their work, users should have successfully trained at least one production-ready practice model.

8. Value creation

The solution has been deployed and employees are using it in their everyday workflows to extract insights and/or build intelligent, automated workflows. The priority here is to iron out any remaining deployment issues, whether hard or soft, and actively seek out new ways to extend the value of the solution. A vendor should be well-placed at this stage to develop added features and functionalities, tailoring your solution to maximise its potential.

9. Customer support & success

In a vendor-customer relationship, a Conversational Data Intelligence provider should work to ensure you can extract the optimal value from its platform. Customer support should be available from a dedicated team of engineers, while a customer success team should be on hand to assist platform adoption through on-demand training and provide the best strategies on its use.

Build or buy?

An organisation may decide that it has the in-house capability needed to build its own Conversational Data Intelligence solution. However this is, at least, a multiple-year undertaking and there is no guarantee that the solution will make it into production.

Model development requires large teams of expensive data science talent. Poor decisions at any stage can seriously compromise a platform, limiting the value it can bring should it reach production.

A Conversational Data Intelligence solution is difficult and expensive to build, but it's even more costly to run. MLOps has strenuous and expensive talent requirements that only increase as the solution is scaled up and integrated with more systems in the enterprise.

Furthermore, proprietary, self-built Conversational Data Intelligence platforms often lack intuitive training tools due to time constraints in the development phase. Without such tools, data annotation becomes a much longer and more resource-intensive process. SMEs are more likely to become unmotivated and disengaged from the process, which can contribute to the overall failure of the platform to achieve adoption.

By contrast, the benefit of purchasing from a vendor is that it allows businesses to unload the initial time and investment of developing their capability in-house, while outsourcing the cost of model governance.

A cloud-based solution will give users access to a platform that is managed, maintained and updated by an external team. After a short implementation period, they have access to a platform that is tried, tested and fully production-ready. This frees business users to instead focus on the task of training and utilising the models they create.

Re:infer Conversational Data Intelligence

The Re:infer Conversational Data Intelligence platform provides a centralised NLP capability for the enterprise. It gives customers all the tools they need to train powerful AI models quickly and easily, enabling them to mine, monitor and automate their business communications. And this is without any need for expensive data science and AI development expertise.

Re:infer is the most mature and sophisticated Conversational Data Intelligence solution on the market. It is underpinned by the latest transformer-based machine learning models, delivering the highest accuracy and data efficiency available. Our research team tracks emerging trends in machine learning and NLP, ensuring useful innovations are implemented across all models. Model development and governance are fully managed by Re:infer, freeing our customers to focus on extracting value from their communications data.

Re:infer also provides a fully no-code, user-friendly platform. No technical experience is required to build machine learning or language models in Re:infer. Users receive a fully guided experience from discovery to deployment, with prompts and alerts to advise the next best training actions. Re:infer offers unique features and capabilities that enable faster training times, including a model health score that helps them see when a model is ready for deployment.

Key takeaways

- **Businesses need to deliver NLP benefits quickly and sustainably**

You should focus on reducing time-to-value while ensuring the solution can be scaled and maintained across the enterprise. Only then can you reliably deliver and scale enhanced analytics and automation capabilities.

- **Self-building a Conversational Data Intelligence solution is a multi-year effort, with no guarantee of success**

Model development is difficult, time-intensive and expensive given the need for data scientist talent. It also requires careful thought around automation, reporting, and user experience. As well as this, the solution needs to meet enterprise and regulatory requirements and be integrated with source systems and control frameworks.

- **A vendor solution reduces risk and ensures fast time-to-value**

A purchased platform, such as Re:infer, provides you with the latest most up-to-date algorithms, short training cycles, automatic model validation, out-of-the-box controls, and pre-built integration with enterprise systems and security frameworks. It provides you with the fastest time-to-value and the most reliable, value generating solution in the long term.

To see Conversational Data Intelligence in action, and learn how it enables cost savings, improved efficiency and superior customer service, [watch our Re:infer Platform Tour](#).

To see how Re:infer drives greater process understanding and end-to-end automation, [start your free trial today.](#)

Re:infer