An operating model for *AI at scale*

How to maximise the expected business value from AI

Why organisation and ***team structure matters*** when delivering AI at scale

01

Why organisation and team structure matters when delivering AI at scale

Data, analytics and artificial intelligence (AI) provide a competitive advantage across all industries and should be at the forefront of any business transformation agenda. However, existing operating models and resources may not be geared for success, and the need to deliver a robust AI and analytics capability within the wider organisational framework is imperative.

Awareness has never been higher of the benefits of data science and AI, but success is far from simple.

The way in which AI and data science projects and teams are structured plays an essential role in the delivery of business insight. The operating model and ways of working are equally important as the technical tools, algorithms and skills needed to build the solutions themselves. Good organisation is not the magic bullet that will guarantee success but, when combined with other key capabilities, it can greatly improve the likelihood of delivering AI at scale. Ultimately, data science and AI is a team sport.

Businesses must put in place an organisational structure to ensure their data and AI solutions meet the needs of their information consumers, enabling and facilitating business use cases. The ultimate objective of the data science and AI organisation is to deliver business value faster by ensuring solutions optimise processes, improve decision making or free up skilled resources to focus on more value-add activities.

The operating model must take into account the core components of each dimension (shown in Figure 1) for delivering AI at scale. With strong governance as its foundation, it will become a way to create the demand, and then a way to execute and operate at scale. Clearly, the functional dimensions of the operating model must be considered within the context of a wider transformation to deliver AI at scale, and its success will depend heavily on the business's ability to transform processes, people and technology.

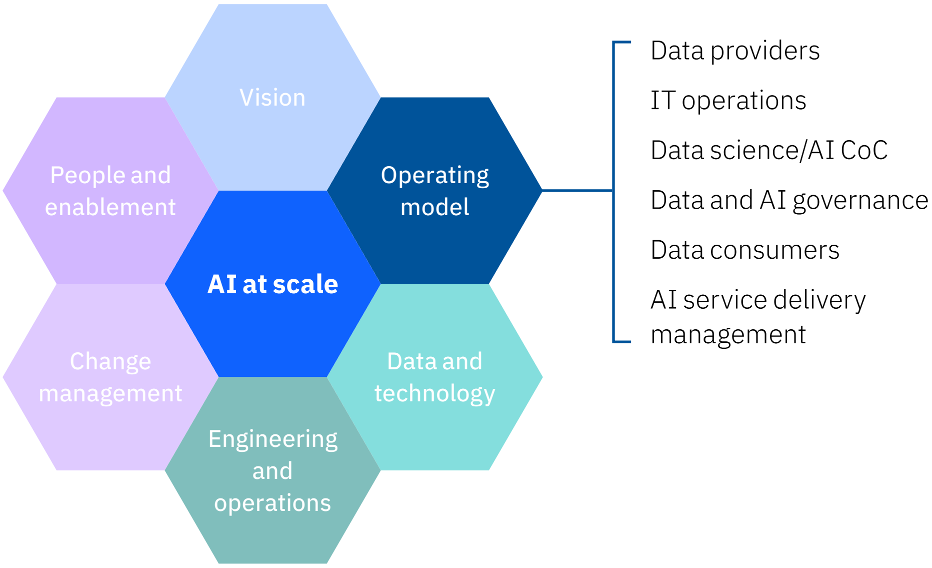

Figure 1: Core dimensions of delivering AI at scale

This paper focuses on the functional capabilities required to define, adopt and establish an operating model for AI. This advice should not be taken in isolation from other critical dimensions of delivering AI at scale, which should be integrated into defining the operating model.

"If it doesn't scale, it is just an interesting project"

Artificial intelligence: The killer app for data, 2019 IBM IBV white paper

Executing a company's vision for AI requires a framework that, at its core, drives innovation through a mechanism to create ideas, prioritise them, test or fail proof of concepts, and execute select minimum viable products (MVP) towards a programme of transformation and industrialisation. Data, analytics and AI organisations need to deliver capabilities that improve business decision making in order to generate value, momentum and credibility. This will ensure that stakeholders maintain their faith and that budgets are not redirected to other pressing priorities.

Target ***operating model for AI:*** functional key dependencies

02

Target operating model for AI: functional dependencies

Organisations are flattening their hierarchical structures in favour of agile and empowered teams. Teams with diverse but complementary skills out-perform and out-innovate more homogeneous ones.

But there are several, often competing functional capabilities that should be taken into account in a data science and AI organisation's operating model.

A. **Data providers**
   Business-critical data will almost certainly be provided by several data providers (source system owners). These may be separated geographically and organisationally, and in some cases may be external.

B. **IT operations**
   Data lakes, data warehouses and cloud platforms are ubiquitous, and AI solutions depend on the IT infrastructures that underpin data flows and processing. This will generally be managed by the overall IT organisation. But according to the project phase, dedicated personnel will be required to ensure the smooth running of the AI solution development, deployment and integration. And when AI solutions become mature enough, dedicated personnel within IT operations will be required to maintain and operate the deployed solutions in production.

C. **Data science and AI skilled resources**
   While data science and AI tools are becoming increasingly easy to use for the non-expert, skilled resources are still in high demand. Not every business will have the luxury of a large team of data scientists, but a dedicated team will be required to undertake the modelling activities and AI solution development.

D. Data and AI governance
Governance of the data stored and processed within the platform is an important function that ensures the quality, integrity and security of the data is maintained in the different data stores and throughout the processing of the data. Careful management, and potentially tracking, of data accessibility and usage will also be required. Data governance should also incorporate capabilities to evaluate the ethics and explainability of the solutions.

E. Data consumers
Business process owners and end users who will benefit must be included in developing the AI solution from the outset. They will be consuming the results of the models or directly using an interface with embedded AI; their ongoing implication in the design and development phase will increase the likelihood of adoption and, consequently, the business benefits. Data availability, timeliness, quality and completeness are key considerations that will underpin the value to the business.

Business process owners and end users who will benefit must be included in developing the AI solution from the outset

F. AI service delivery management
Companies that are investigating or implementing AI solutions are unlikely to only have one initiative in progress. They will require a centralised function to qualify and plan the data science and AI projects in order to drive projects forward, act as gatekeepers throughout the development lifecycle and ensure that change management is embedded in the project design. This function will evolve during the “operate” phase of AI projects, but will continue to provide oversight to ensure that AI solutions and any upgrades bring business value.

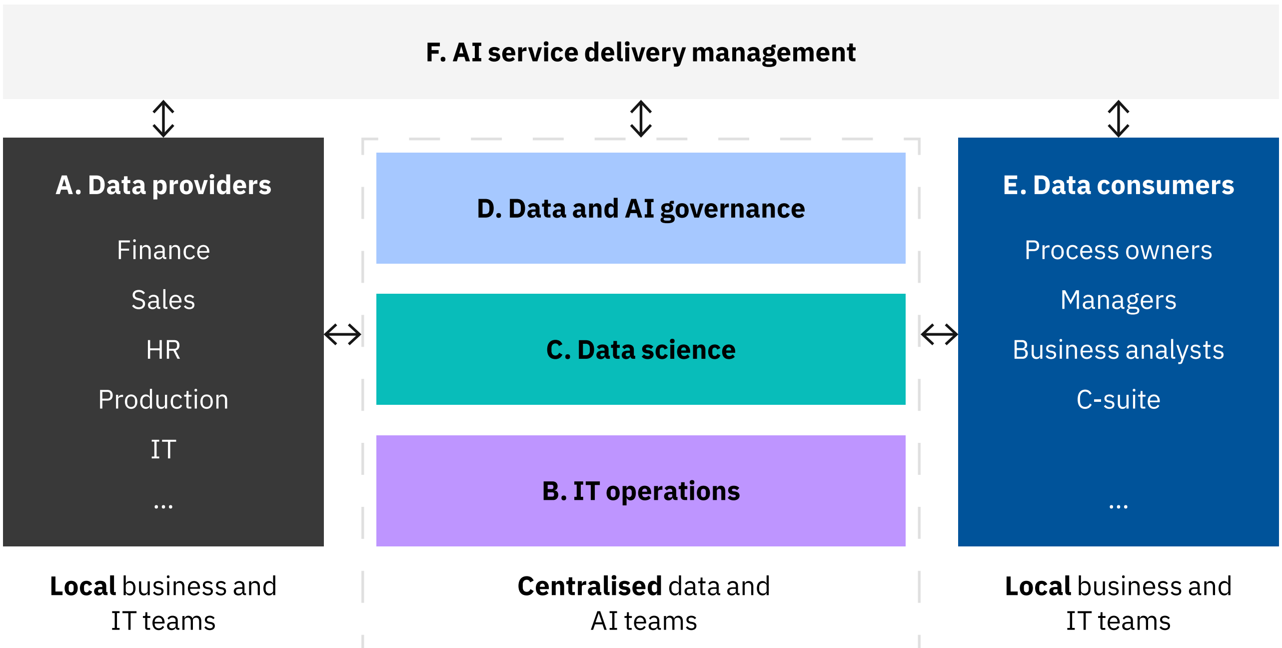

Figure 2: Target operating model for AI at scale

Important factors to consider

When implementing an AI operating model, the following factors will have an impact on the operating cost and the speed with which a return on investment can be expected:

- **The mix of internal versus external resources required to support different operational activities.** This will depend on existing data and AI capabilities, but also on the specialist skills required and the overall workload.
- **The volume of activities that are supported via a centralised team rather than local IT or business units.** Workloads can be spread across a wider community with a centralised team providing the deeper data science and IT skillsets required to optimise and scale solutions.
- **The phasing of analytics use cases.** The speed with which solutions are developed and deployed will typically be a function of the capacity of the data science team, the estimated complexity and the return on investment that can be expected. Clearly the more experienced the team, the easier this is to estimate and plan.
- **Ways of working.** Increased maturity of AI teams should lead to a more pragmatic approach to governance, ensuring that expectations are managed at all levels. Communication and change management should become a fundamental part of the process in order to ensure the adoption of AI solutions and ultimately the realisation of the expected benefits.
- **The principles of insight delivery:**
  - Data is the fuel. There is no AI without information architecture and there are always problems with the data, so adaptability is crucial.
  - Data science is a team sport.
  - Data science is, as Thomas Edison famously said, 99% perspiration. There is no magic.
  - Communication and collaboration underpin any operating model structure.
  - Avoid the hype and concentrate on what is valuable rather than what is interesting.

Communication and collaboration underpin any operating model structure

These factors highlight, once again, the importance of considering the processes, people and technology that make an AI operating model possible.

AI delivery model

The success of AI solutions hinges on the quality of business insights generated by data scientists and business analysts and acted upon by the business to deliver the value.

To achieve this, it is important to distinguish between the operating model and the delivery model.

The operating model provides the framework; the management direction that describes what insight is required and why it will improve something. The delivery model focuses on how insight will be delivered and used.

In IBM's experience, the delivery model has four distinct phases:

1. **Discover** – use case identification, initial analysis and project initiation
2. **Create** – design, build and validation of models and analyses
3. **Execute** – deployment and integration of models into the business process
4. **Operate** – support and maintenance, ongoing monitoring of model performance

The process highlights the importance of a holistic approach with regards to ethics, value and business change factors to deliver sustainable business value.

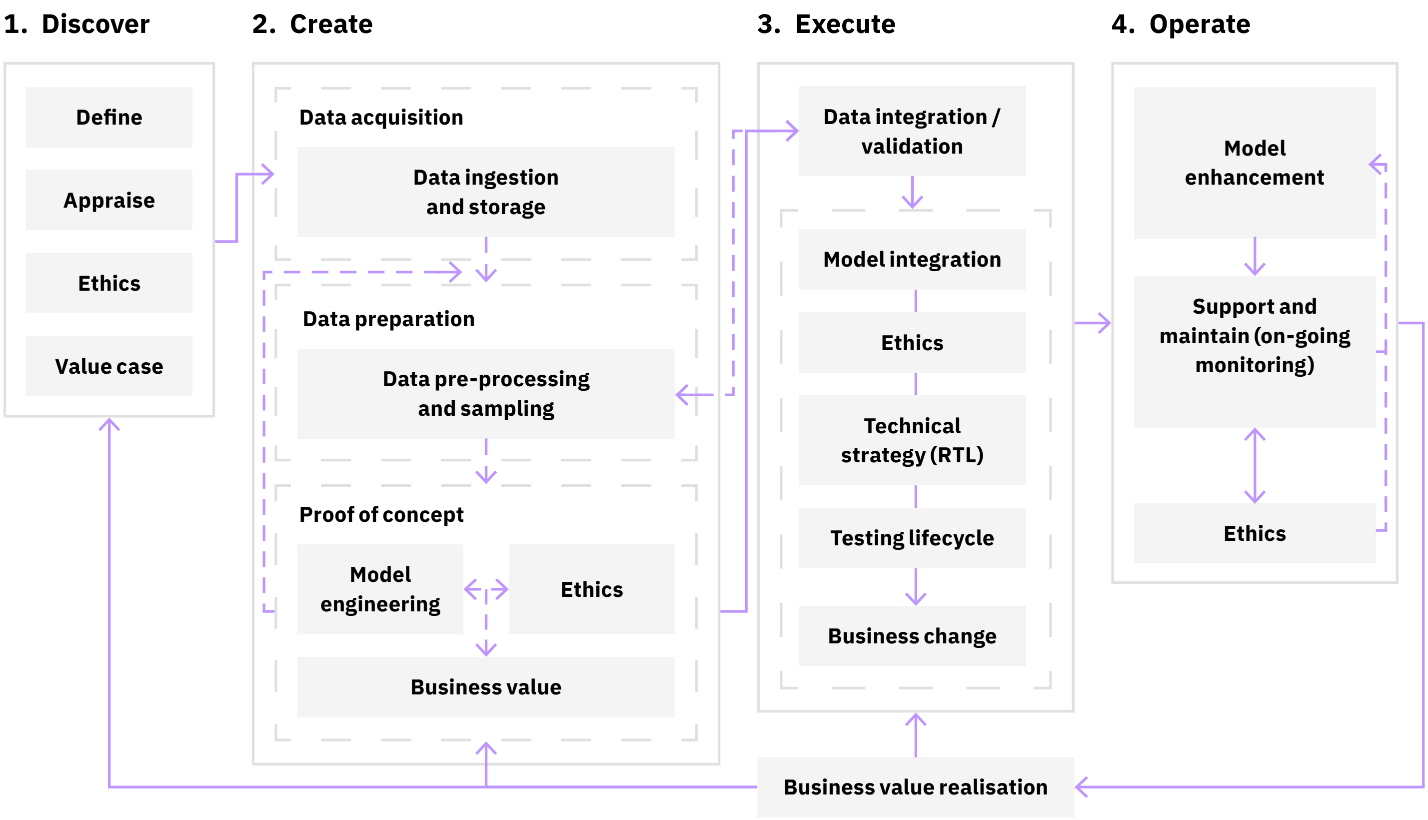

Figure 3: Delivery model for AI at scale

Each team described in the operating model will be involved in the delivery model, but the extent of involvement will depend on:

- **Resource availability**
  Central resources will be limited in terms of overall capacity, so some use cases will continue to be undertaken by a local AI team. Analytics projects will need to be prioritised, taking into account resource availability.
- **Technical or business expertise requirements**
  Certain analytics use cases may require technical know-how or deep expertise in a certain business domain.
- **Data requirements and availability**
  There is always a balancing act between what data the use case requires and what data is available to answer the question; businesses often need to shape the problem according to what is available. Different local teams will be needed to provide support to ensure the data can be ingested into the repository and explained during data preparation.
- **Deployment requirements**
  Usually there will be three types of deployment once a model has been developed:
  1. A "one time" analysis/model: Results are shared with the business group and no deployment is required.
  2. A recurrent analysis/model (daily, weekly or monthly): Deployed either locally or centrally, with results stored in the data lake, to be accessed via a regular report, dashboard or application.
  3. A fully integrated model: Fully integrated within the existing business and IT processes to enable automated or semi-automated decision making that changes the business process flow. This will require the analysis or model to be deployed in a more "industrialised" manner, taking into account the need for training or retraining the algorithm, scoring and feedback mechanisms.

There is always a balancing act between what data the use case requires and what data is available to answer the question

A 2018 *IBM Institute of Business Value* report found that 60% of companies see compliance as a barrier to achieving success in applying AI, in part due to a lack of trust and understanding of the system[1]. Building trust and transparency should form a fundamental part of the AI delivery and operating model and therefore be prevalent at every stage of the solution lifecycle. Trust and transparency are as much an internal requirement for business leaders as an external requirement for customers, clients or external stakeholders using the deployed solution. If business benefits are not clear and realisable, then buy in and focus will wane and wither.

60%

see compliance as a barrier to achieving success in applying AI

An ***adaptable*** approach

03

An adaptable approach

There is no one-size-fits-all operating model for AI and client AI maturity, business size and scope; industry-specific factors such as the geographic spread of data sources will play a part in determining the optimal AI operating model for AI at scale. It should not be set in stone. Based on IBM's experience, it will evolve over time as the need and capacity to scale increases.

In many cases, AI projects will start with small squads and will be driven by local business units or departments – but when the scope increases, a more centralised lab model with increased oversight and IT input will be required. And when scale and industrialisation become the main drivers, the business will need a centralised factory with increased governance of data and solutions.

There are many different types of operating model to consider, and the most appropriate will depend on industry and client-specific factors. Achieving speed-to-value means "doing analytics" in a coordinated, structured way. The table below gives some guidance on where to start.

		Pros	Cons
	**Functional** AI and analytics agenda is not enterprise driven. Teams are dispersed across the organisation with a small central analytics unit.	• Resources are where the current demand of the enterprise is • Resources can more easily specialise in each functional area	• Coordination is challenging • Team is dispersed • Lack of common technical and functional framework • Busiest unit becomes the functional team • Enterprise scaling can be challenging
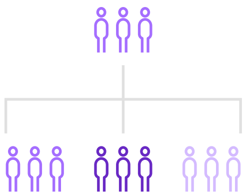	**Centralised** Teams/resources/tools/data are all in a single location and cannot be accessed by units outside the centralised unit.	• Governance, standards, management easier to control • Resources are all in one location • Central management and greater control of team's strategic mission • Collaboration and sharing of assets is easier to facilitate	• Lack of co-create opportunities with business stakeholders • Teams can be a little removed from some business processes and understanding • Adoption of solutions can be more challenging given the distance from the business
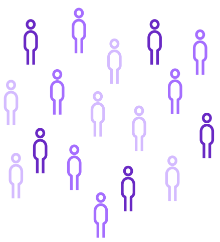	**Decentralised** Resources are dispersed across the organization in different silos with no view of analytics activities outside their respective unit. Multiple teams exist in the organisation but are not joined together strategically.	• Functions have better oversight of resources • Business skills are closer to team • Business understanding for work has greater clarity • Solutions are perceived as "business led", which facilitates adoption	• Team is dispersed • Low visibility of analytics activities • Governance and standards difficult to manage • Low knowledge sharing and asset sharing
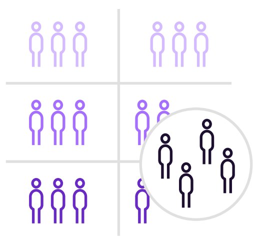	**Factory** Teams are organised in such a way as to prioritise the industrialisation of AI and data science solutions, covering the full lifecycle. The factory could be organised centrally or dispersed according to functional capabilities (e.g. data management, experimentation, testing, industrialisation).	• Facilitates deployment at scale • Synergies across solutions easier to identify, making re-use of assets more effective • Resources are assigned to squads according to technical speciality • Business resources are integrated where required (e.g. as product owners)	• Requires greater organisation and governance overheads • Requires coordination and buy-in across multiple business units (business and IT) • Requires more experienced and higher skilled resources at the start to get the factory up and running
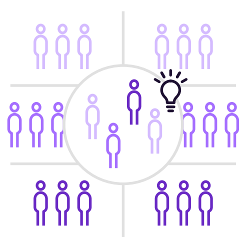	**Centre of excellence (CoE)** Operations and activities are coordinated and tracked from a single CoE, with the majority of the team placed in different areas of the organisation.	• Resources coordinated by the central location but embedded in the business for closer relationships • Better alignment with the enterprise's strategic goals enabled by AI and analytics • More initiatives created across the enterprise	• Governance of resources and tracking of people development needs to carefully be maintained • CoE may turn into a resource pool instead of driving strategic growth • There is a risk that the CoE becomes too focused on technical excellence, neglecting business priorities
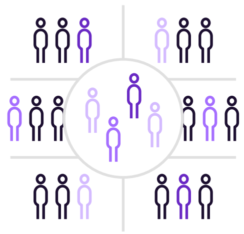	**Consulting** Organisation resources are driven by availability of personnel for projects. The AI and analytics unit charges the organisation for services, primarily for project work.	• Use cases scoped more efficiently • More "skin in the game" • Value and outcomes are measured in a more established framework	• Strategic initiatives can be overlooked for payment of services • Availability of resources can become an issue • No centralised location • Requires an ongoing internal marketing effort to ensure all business units are aware of the service

Figure 4: Comparison of the main options for an AI at scale operating model

Unexpected *challenges*

04

Unexpected challenges

Are data scientists “corporate rock stars” capable of delivering plug-and-play AI with immediate returns, automatically solving business problems with technology that is immune to failure, unbiased and wholly accurate?

Sadly, it is not that simple. From experience, IBM has unearthed a number of challenges that were not immediately apparent during traditional scoping exercises. While there are always tooling, integration and data challenges, success is usually better determined by enterprise-level considerations such as those described below.

Value and delivery speed to value

It can be difficult to translate business problems into AI use cases that data scientists can answer as quickly, effectively and efficiently as realistically possible. There is always a trade-off between accuracy, speed and value. The business, at best, doesn’t care about analytical complexity; it wants something understandable and explainable that will make a difference. Navigating this while delivering progress builds credibility and generates momentum. Credible momentum is important because it builds a virtuous funding cycle.

Governance and communications

Data science is a voyage of discovery, navigating options that may or may not find correlations to prove or disprove the hypothesis. Does finding no fraud in a particular data set you were expecting mean failure? Uncertainty demands flexibility, pragmatic governance with clear objectives, success criteria for acceptance and, above all, regular communication to manage stakeholder expectations. The best approach is to start small and re-enforce iterative progress.

Ownership, funding and commercials

Successful AI projects include multi-discipline teams – combining business and technical skills – that act as a bridge between IT and the business. This three-way relationship is crucial, and clarity of ownership is key to ensure that all stakeholders understand their roles.

Funding the project

It is important to agree upfront who owns and funds the MVP, the scaled solution, deployment and support costs, and who ultimately owns the benefits. This should not be confused with the funding required to deliver business change activities to embed insight and deliver sustainable business value. The operating model needs to demonstrate value to the business as a quantified benefit that is clearly defined and measurable. Challenging as it may be, this must be fully understood before anything is built.

Commercial models and budgeting

There is no guarantee of success with AI and data science projects. Experimentation is key; we do not get value from an AI function by constraining it to things we know will work. But commercial models and budgets are assessed on a case-by-case basis, which limits the ability to experiment. Good relationships and trust can help to secure a budget that works for the delivery team.

A number of additional principles underpin these enterprise-level considerations and can heavily impact project success:

- Start small, avoid hype and embrace value. The most valuable insight is not necessarily the most interesting or fashionable.
- Data is always the challenge. Do we have the right data to solve the right problem? Do we have enough data to be confident, and with the right quality and coverage? Adaptability is a must.
- Business change is difficult, and value cannot be realised if it is not sustainably embedded into business-as-usual processes. Good questions to consider in the early stages are: what decision will be made differently if this (model) was put into production, and how will people accept and use it?
- Productionisation will differ from MVP sprints and mirror deployment of traditional software application lifecycles including non-functional requirements, business usage testing, data integration and support models. These should be tackled as early as possible.

Summary and conclusion

05

Summary and conclusion

The main objective of the AI at scale operating model is to create realisable, explainable and sustainable insight that improves business performance.

The operating model plays an important role in structuring, organising and governing AI solution delivery to the business, and demonstrating momentum and credibility to solve critical business problems.

Businesses must take the following (often competing) functional capabilities into account:

- Data providers
- IT operations
- Data science and AI skilled resources
- Data governance
- Data consumers
- AI service delivery management

The operating model should be considered within the context of a wider transformation to deliver AI at scale, and its success will depend heavily on the ability to transform processes, people and technology.

There is no one-size-fits-all operating model for AI. It will undoubtedly evolve over time as business, IT and industry priorities change.

A clearly defined operating model that takes into account the challenges and complexities of delivering modern AI solutions is key to maximising the expected business value.

Summary and conclusion

Conclusion

Organisations need to exploit the last competitive advantage: data. AI can unlock this, but the operating model is essential in maximising the chances of success. Particular attention should be paid to the importance of business-led collaboration, governance and communication. As Erik Brynjolfsson recently stated: “The bottleneck (for AI) now is in management, implementation, and business imagination, not technology.”[2]

Let’s talk

Want to know how this can work for your business? Get in touch with our experts for a one-to-one consultation.

Damien Holliday
Data Science Lead
for Financial Services,
Cognitive & Analytics
IBM Services
+44-7793-469500 (Office)
damien.p.holliday@uk.ibm.com

Rob Worsley
Data Science Leader, France
Cognitive & Analytics
IBM Services
+33 6 88 38 44 84
rob.worsley@fr.ibm.com

Ciarán Hickey
Analytics Lead, Ireland
Cognitive & Analytics
IBM Services
+353 87 1671907
ciaranhi@ie.ibm.com

1 *Beyond Hype: Shifting Towards Enterprise-grade AI: 2018 Institute for Business Value Report*
https://www.ibm.com/downloads/cas/QQ5KZLEL

2 *The Business of Artificial Intelligence: ERIK BRYNJOLFSSON AND ANDREW MCAFEE*
https://hbr.org/2017/07/the-business-of-artificial-intelligence