aiimi

Insight Report

AI & Data Governance.

Informing 2026's biggest data and AI opportunities.

DATA & INFORMATION LEADERS
CDOs
DATA PROTECTION OFFICERS
CIOs
INFOSEC LEADERS
INNOVATION LEADERS
HEADS OF DIGITAL, DATA, AND AI

Featuring insights from

- **Global manufacturer**
- **Government agency**
- **Leading GSI**

and more.

Contents.

RECOMMENDED READING FOR:

	CDO & CIO	Head of Digital & Innovation	Data Protection Officer	Information Security Leader	Head of Information or Records Management
CHAPTER ONE Elevating Data Governance to meet the AI agenda.	●	●	●	●	●
CHAPTER TWO Getting your data AI-ready.	●	○	○	●	●
CHAPTER THREE Identifying AI use cases that prove ROI.	●	●	○	○	○
CHAPTER FOUR Compliance in the age of automation and transparency.	●	○	●	●	●
CHAPTER FIVE Deploying AI Agents for smarter data management.	●	○	●	●	●
CHAPTER SIX The evolution of AI Search.	●	●	○	○	●

Introduction

Foreword from Steve Salvin, Aiimi CEO.

During AMP8, water companies face immense pressure to deliver on ambitious capital programmes, respond to climate change, and improve resilience – all while contending with rising costs and maintaining affordability for customers. At the same time, information is exploding in volume and complexity, from decades of infrastructure data to customer interactions and environmental monitoring.

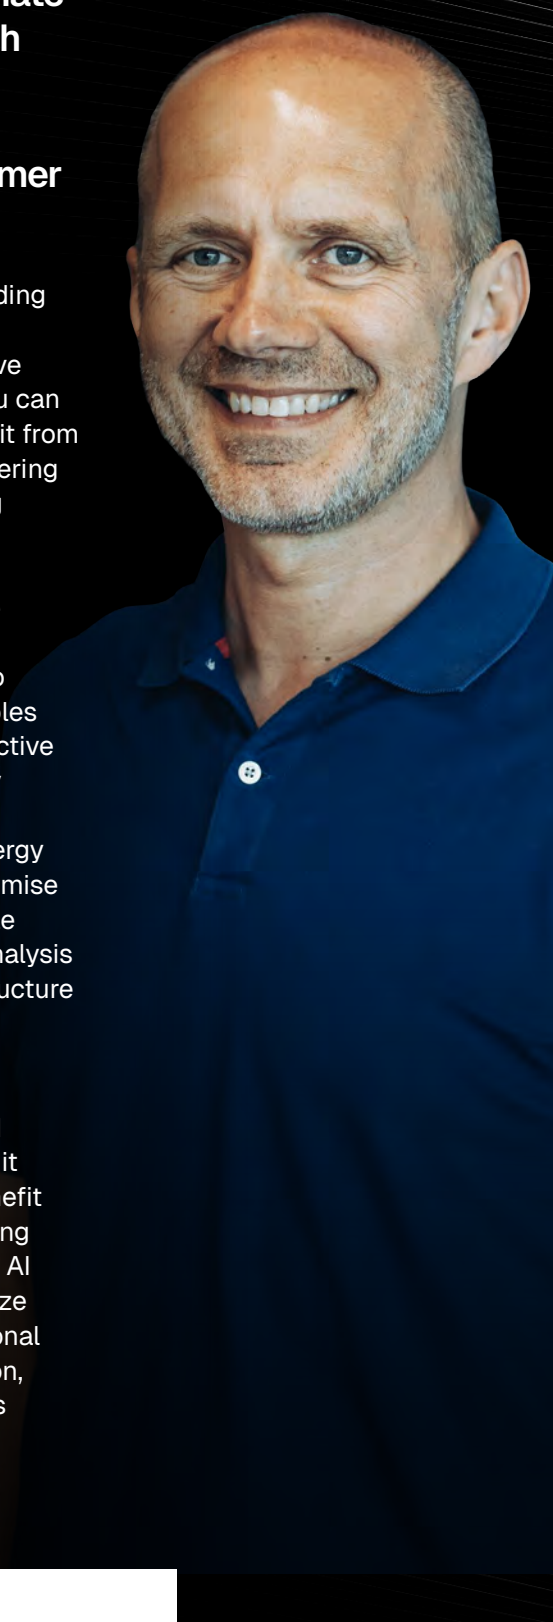

It's not just structured data – it's video, audio, chat, images, and more, across countless platforms, formats, and sources. All this, against a backdrop of unprecedented regulatory scrutiny. So, water companies are grappling with more information, complexity, obligations, and expectations than ever before. Is AI the answer?

In part, yes. AI presents a critical opportunity to help achieve the goals of AMP8 and beyond. But while the sector is rich in data, siloed systems mean much of its potential value remains locked away. Simply pointing AI at the problem doesn't solve it – in fact, it exposes the underlying problems, like weak data foundations, fragmented governance, and missing context. AI doesn't paper over the cracks; it magnifies them. It's why disillusionment in AI is rising, and why organisations are spending more than ever on platforms, people, and compute power – just to stay standing, not to get ahead.

Right now, most organisations are still chasing their data, not controlling it. We need to change that – to get control over data and make it work for you.

This means visibility and governance across your entire data estate: knowing where your data is, what it means, and what it's worth. It means building an environment where always-on compliance becomes a competitive advantage, not a burden. Then you can really start pulling levers to benefit from your data: joining the dots, uncovering insight, building trust, and turning information into action.

From here, you can operationalise AI with confidence – applying the right data, with the right model, to solve the right problem. This enables practical use cases such as predictive maintenance tools, which identify asset risks in real time to reduce downtime and leakage; smart energy management systems, which optimise pump usage in line with renewable energy availability; and climate analysis platforms, which pinpoint infrastructure vulnerable to future flood risk.

The water sector's future will be centred on resilience, control, and understanding. At the same time, it is uniquely well-positioned to benefit from AI transformation. By elevating data governance to align with the AI agenda, water companies can seize opportunities to enhance operational resilience, support decarbonisation, and create better outcomes for its customers and the environment.

aiimi

Chapter One

Elevating Data Governance to meet the AI agenda.

The real value in data governance has always been in improving the quality, completeness, and relevance of the data that feeds your most critical decisions. Now, AI promises to accelerate workplace decisions and actions like never before – but AI needs the right data to deliver on these big promises.

For leaders responsible for data governance at scale, now is the time to prove the value of getting governance right – and lead the way with AI adoption.

AI has added fuel to the data governance fire; in fact, it's pretty much burned the house down. The rules of data governance – how to start, where to start, and what you can achieve with good data – have totally changed.

Leaders have the chance to create a whole new set of founding principles for modern data governance, where AI is at the centre. And, importantly, the business case for change is strong.

⟶ cont.

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

Chapter One

Targeted governance for AI success.

If AI success is critical in your organisation right now, it might be time to press pause on extensive data governance initiatives and start getting focused. You don't have to boil the ocean to run a successful and valuable governance initiative; instead, start small and stay focused.

Improve data governance, data classification, and data quality in the areas you need it most. Importantly, align these focus areas to AI use cases that are critical to your business. Find the areas where there will be a clear return on your investment in managing and improving data quality. Most often, these areas tend to be aligned with AI priorities.

After years of trying to address the entirety of an organisation's data governance needs simultaneously, there's often no clear end in sight – and, most importantly, there's no obvious return on investment. Targeting data governance and management work to the areas where you're strategically investing in AI makes sound business sense.

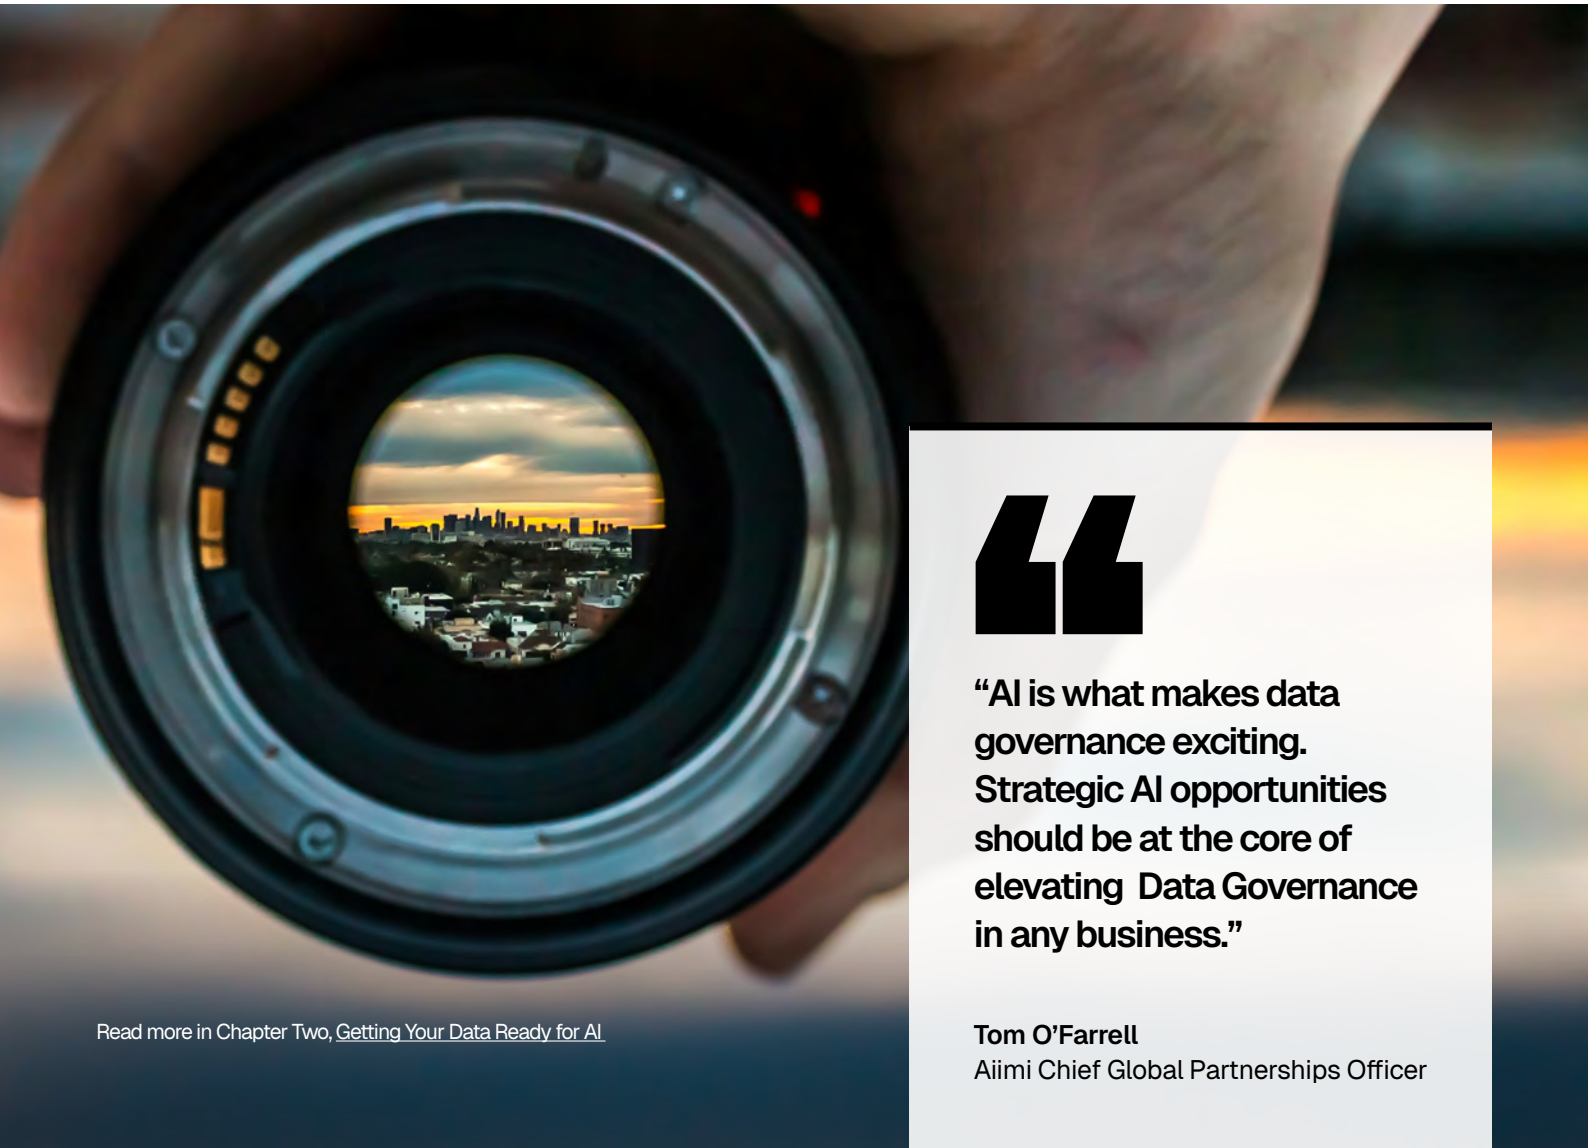

Read more in Chapter Two, Getting Your Data Ready for AI

> **"AI is what makes data governance exciting. Strategic AI opportunities should be at the core of elevating Data Governance in any business."**

Tom O'Farrell
Aiimi Chief Global Partnerships Officer

Chapter One

Automating Governance: AI's opportunity

Whether it's to support AI initiatives or fuel migration to the Cloud or out of a legacy system, AI can change how we govern and manage data. AI and automation have a huge role to play; making complex data governance manageable, achievable, and far more effective.

AI should be a tool that helps achieve better data governance, faster. Data governance has to address how we set permissions and where we store data. Every enterprise is full of hundreds of systems and sources, each with their own complex permission structure, and AI can navigate this.

Where previously you might have asked users to take on the task of classifying and labelling their data to build understanding of what the data relates to and who should be able to see it, now there's a practical route to automating this. Instead of messy dropdowns and checkboxes that create inevitable inaccuracies and inconsistencies, combining natural language technologies with unsupervised Machine Learning techniques means data can be classified automatically. Starting with clustering to establish course-grain groupings for data and validating this with users, we can then train an AI model to run classification at a larger scale. AI can now add labels – for example, if a document has been classified as an invoice, it can capture the payee details, invoice number, payment date etc., all automatically. And by using feedback from users to reinforce the model's training, AI gets better the more data it classifies and enriches.

Once the rules of the game are set, Agentic AI can even go ahead and file documents, using automated classification and labelling to keep data safe and compliant. AI can move files according to complex permission structures, flag the latest versions or 'golden records' for a single source of truth, and more – and it never gets tired or sidetracked.

When you're focused on providing data to power a business-critical AI initiative, quality and completeness are top priorities. AI is good at finding and filling gaps in data too, reaching much further than humans could. If there are gaps in structured data, Large Language Models can seek out the relevant data in unstructured content, joining the dots to build a more complete, high-quality picture of your data.

When AI is this good at managing and classifying data, realising value from data governance no longer has to take years and involve every data owner and user. A lack of buy-in and participation from users doesn't matter in the same way – AI only needs us to oversee and validate. AI means scale is no longer a problem; data is growing faster than ever, but AI keeps up. It doesn't lose interest after a few months; it stays alert to changes at any point and updates classifications or metadata accordingly.

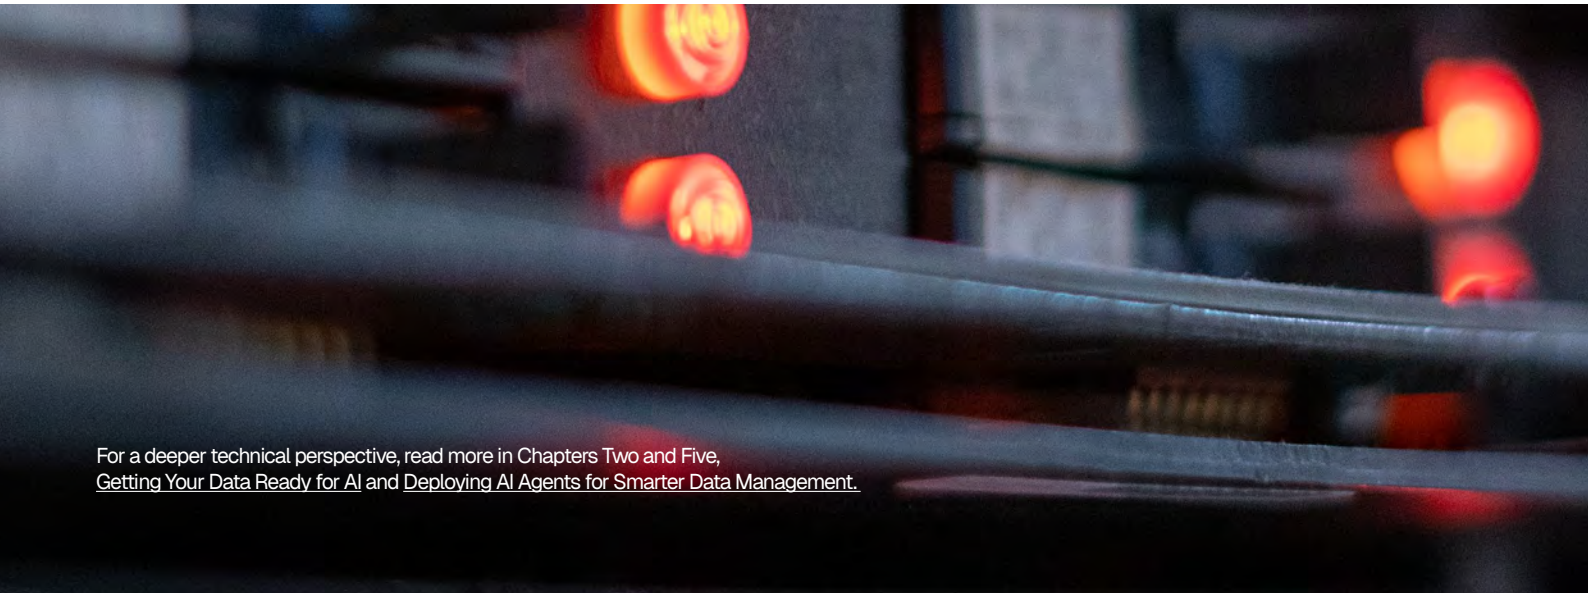

For a deeper technical perspective, read more in Chapters Two and Five, Getting Your Data Ready for AI and Deploying AI Agents for Smarter Data Management.

Chapter One

Best next steps for AI-driven data governance:

Shift your data governance approach to focus on the areas that matter most for your AI priorities – start small, target high-value use cases, and align governance efforts with your most important business questions. Use AI and automation to classify, label, and manage data efficiently, reducing manual effort and improving accuracy. As you see results, expand your efforts in manageable steps, building a strong business case for modern, AI-driven governance, laying the groundwork for successful AI adoption across your organisation.

↘

- → **Identify your data governance priorities, guided by high-value AI opportunities**
- → **Evaluate tools that use AI to power data classification**
- → **Stay close to ROI from data governance to build the case for expansion**

aiimi

Chapter One

CASE STUDY

Saving on data storage – and setting the stage for AI success.

For one global professional services firm Aiimi works with, automated data governance saved £1m in storage costs within just six months by removing obsolete data. Through automated data governance, they're now laying the foundations for their internal Data Service, curating gold-standard data for almost 30,000 employees.

Read more about this case study:

aiimi

Chapter Two

Getting your data AI-ready.

Speaking to executive teams across different sectors, it's clear that enterprise organisations are still early in their AI journey. While the conversation used to be about 'operationalising' AI, what leaders now want is to empower everyone in their business to use AI tools – giving people the superpowers they need to be more productive.

However, there's a growing realisation that the data landscape within most organisations is in disarray. The quality of information is often poor, and visibility into what data exists and where it lives is limited. For the first time, we're starting to see widespread recognition that enabling the workforce with AI is intrinsically linked to building a solid data foundation.

→ cont.

2

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

Chapter Two

AI enablement and data foundations.

The challenge, then, is achieving both these objectives and doing so without falling behind. This demands a shift away from embarking on decade-long projects to build the perfect data foundations. Instead, it's about taking a more pragmatic, value-driven approach: identifying high-value business challenges where AI can supercharge the solution, then focusing on improving the data associated with those specific areas. It means tackling AI enablement and data quality in a series of manageable 'slices', rather than trying to do it all at once and missing the boat.

"Data Governance is vital to protecting an organisation's (and its customers') information when deploying AI. Information might need to be withheld from AI services because it contains sensitive information such as personal data, or because it is protected by client contractual conditions, for example. Conversely, effective data classification can drive the right information to an AI process to improve the accuracy and relevance of LLM outputs."

Matt Eustace
Aiimi Data Protection Officer and
Head of Solutions Engineering

Chapter Two

The central role of classification.

A critical part of this journey is the ability to find, understand, and classify information. Most organisations have between 500 and 1,000 defined information classifications. Fortunately, AI is really good at clustering, pattern-matching, and automatically labelling documents – and can be enhanced further by combining machine learning with rules-based techniques. Once information is classified, AI can identify data that hasn't been classified and therefore likely holds limited business value. It can find and remove duplicates and near-duplicates, retaining only the most recent or authoritative versions. AI can flag data that sits outside of retention policies, prompting archival or deletion to maintain compliance, and identify data that's stored in the wrong place – such as personal OneDrives, email, or DocuSign – then move it to an appropriate location.

Part of this is about permissions and access. We know the average employee can access only around 1% of an organisation's information. This means that no matter how advanced your AI capabilities, if your people don't have access to the information they need, they still won't be able to find it and make use of it with AI. It's therefore essential to tackle overly restricted or misplaced information up-front.

Classification is the foundation that enables these crucial steps forward. With effective classification, information retrieval becomes much more efficient and effective, with irrelevant material excluded from search results. Removing unclassified or low-value data can immediately eliminate 30–40% of the information estate, reducing noise and improving both user experience and cost efficiency. One of Aiimi's customers, a leading professional services firm, saved £1m in data storage costs within a year, simply by removing duplicates and obsolete data – that's before even considering the additional benefits that AI-powered search brings.

Another significant advantage of classification is the ability to build collections of information relevant to specific use cases. For example, contracts may be scattered across DocuSign, SharePoint, email, and personal drives. By classifying all contracts, regardless of where they're stored, we can then use agentic AI to assemble a comprehensive collection, deduplicate it, verify signatures, and ultimately whittle it down to the golden records – the fully executed, authoritative contracts we want people to query against and interact with.

Beyond the immediate value for the user, this also provides the data and business process owners responsible for these contracts with a clear view of their information landscape – enabling them to understand how data is being used by the business and where it lives, and helping them fix governance policies and processes around information creation and storage.

Chapter Two

An evergreen information landscape.

What's particularly exciting is the potential for classification to become a dynamic, evergreen process. With published business classification schemes, AI can cluster and summarise new data, map it to existing classifications, and train models to label information automatically across the enterprise. It can also detect new clusters of information that don't fit into existing classifications, alerting humans to the emergence of new document types or business needs. This ensures the classification scheme remains up-to-date and relevant, able to keep up with evolving requirements such as new regulations.

By leveraging industry-specific classification frameworks, organisations can accelerate this process, reducing manual effort and improving accuracy. The vision is for AI-powered classification agents to autonomously cluster, label, and manage all enterprise data, surfacing only the small percentage that requires human attention. This delivers instant business benefits and cost savings, while also laying the groundwork for a continuously improving, AI-enabled data landscape.

Chapter Two

Making sense of data quality.

The CARTIE framework – Completeness, Accuracy, Reliability, Timeliness, Integrity, and Existence – offers a structured way to categorise data quality issues and consider appropriate mitigations.

Completeness.

The data covers the full scope required, without gaps or missing values.

Accuracy.

The data is valid and properly formatted for use.

Reliability.

The data is consistent and dependable across time and processes.

Timeliness.

The data is refreshed and updated at the right frequency, avoiding staleness.

Integrity.

The data remains consistent and aligned as it moves between systems.

Existence.

The data actually exists, rather than being absent or uncollected.

Given that the level of data quality required depends on the specific use case and the problem at hand, it's essential to consider the use case first before making any assessments about data quality requirements. Pre-trained models open up more opportunities to use AI with lower data-quality demands. Value-driven experimentation enables organisations to determine which data quality issues truly need attention (and which don't), in order to focus their time and money where it will have the greatest impact.

High data quality sensitivity	Training a model on your company's own data
	Searching internal documents for critical answers (RAG)
	Automated decisions based off your data
	Customising models with your organisation's data
Low data quality sensitivity	Co-Pilots to draft content (Emails etc.)
	Sentiment analysis
	Grouping similar topics from feedback or documents
	Automating routine tasks / workflows

Chapter Two

Best next steps for AI-ready data:

Take a pragmatic, step-by-step approach – don't wait for a perfect data foundation before enabling AI. Use AI to help classify, clean, and organise your information, making it easier for people to find what they need and for AI tools to deliver value. Over time, aim to make classification and data management an ongoing, automated process, so your information landscape stays healthy and ready for whatever comes next.

↘

- → **Begin integrating AI by taking practical, incremental steps**
- → **Leverage AI to improve data and its discoverability and usefulness**
- → **Progress towards automated information governance**

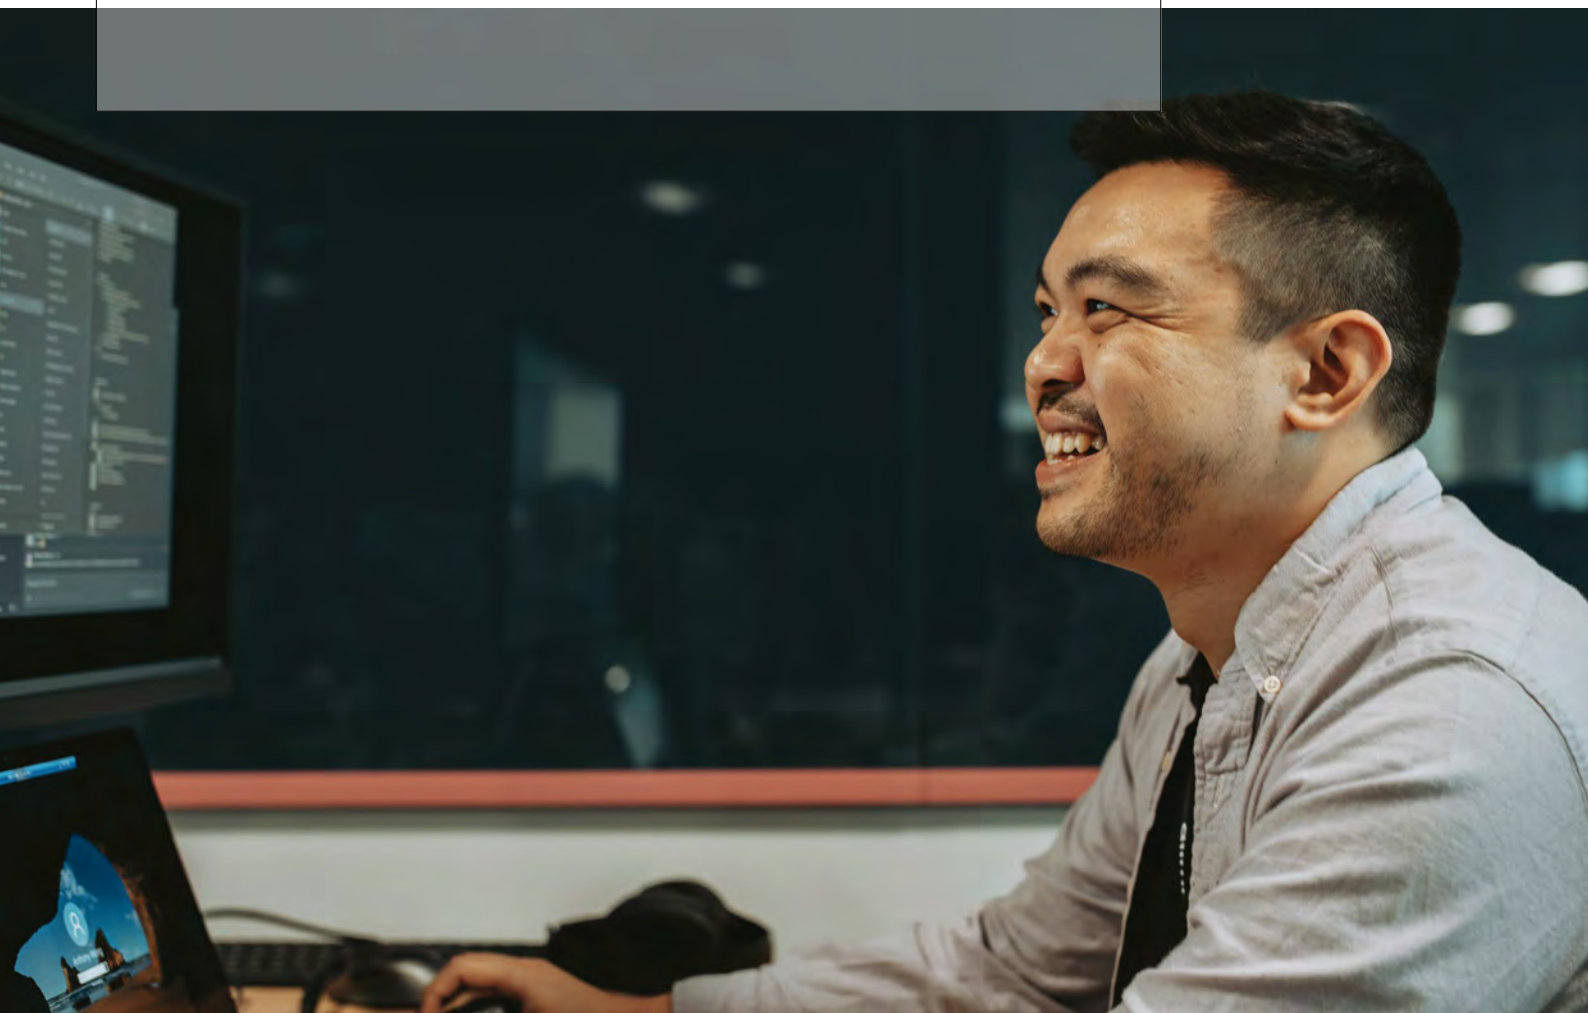

aiimi

Chapter Three

Identifying AI use cases that prove ROI.

Business leaders want to use AI strategically, but it can be difficult to know where to start or how to guarantee value.

The first hurdle is demonstrating exactly where to invest in AI and why. Data quality and internal skills gaps are common barriers, and traditional approaches to strategy simply aren't cut out for AI – the technology is evolving too rapidly.

To unlock the value of AI in the enterprise, the most effective starting point isn't technology or data – it's questions. When you start with, "What questions does my organisation want to answer with AI?", you can focus your efforts on the data that truly matters, narrowing your field of vision to what will drive the most impact.

→ cont.

> **"Our biggest barrier is choosing the right things to adopt and where."**
>
> **Risk & Compliance Leader**
> Insurance firm

> **"Senior leadership buy-in is our biggest barrier."**
>
> **Risk & Compliance Leader**
> Retailer

Quotations from Aiimi's 2025-2026 AI Now Leadership Survey.

3

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

Chapter Three

Here are just a few questions brought to us by leading organisations:

Water & Utilities
For every live burst or pollution event, can we identify the root cause and best next actions as quickly as possible, using insights from previous incidents? Better yet, can we predict the next event and recommend preventative actions, or optimise resource deployment?

Manufacturing
How do we fix a faulty machine for a customer, leveraging knowledge from every previous fault and fix? Even better, can we predict failures, adjust servicing schedules, and recommend maintenance to avoid breakdowns altogether?

Local Authorities
When a resident or employee initiates a Data Subject Access Request, can we automatically collect all relevant data, redact sensitive information, and route it for review?

Armed with valuable questions like these, you can begin to identify the data needed to answer them. This is where due diligence comes in: does the necessary data exist, and is this a worthwhile experiment? If the answer is 'yes', you've found your first opportunity for a manageable data governance experiment that aligns directly to value. This approach also challenges the notion that data governance means lengthy, costly transformation projects with uncertain returns.

Engaging people, not just adding technology.
The only way to get real value from AI right now is to experiment safely and learn quickly; the best approaches are agile, iterative, and action focused. By identifying, evaluating, and prioritising the best AI use cases that will release value quickly, you're able to win over stakeholders for greater investment in AI. For example, when an alternative fuel provider worked with us to create their foundational AI strategy, focusing on value and safety enabled them to secure the trust and buy-in needed from senior leaders to pursue further AI opportunities. For some organisations, the focus of AI strategy might be on ethics, especially for those in the public sector or dealing with sensitive data. For others, it might be more about scaling internal AI capability or getting a grip on fragmented AI initiatives that are already underway. Whatever the desired outcomes, the process should start with engaging your users. Proven Discovery techniques, such as workshops, interviews, and exploration, enable you to surface the all-important questions your people need answers to and the decisions they need to make. This helps build an inventory of AI use cases clearly aligned to strategic value, alongside early insights into technical, ethical, and operational considerations.

"Before weighing up the value of any AI solution or initiative, the critical question is: what needs to change? From pilot through to scale, the focus should stay on delivering measurable, business-relevant results. This creates a golden thread of value, connecting AI initiatives directly to meaningful impact."

Paul Sliwinski
Aiimi Customer Success Director

Chapter Three

The importance of data quality.
No AI use case will deliver value without complete, high-quality data, and poor data is often the biggest stumbling block in the adoption of AI and automation. The promise of AI can be a powerful motivator for tackling data governance, to improve data quality in preparation for strategic AI initiatives. Fortunately, AI is also a tool that helps us to improve data quality, faster. It accelerates data discovery, classification, and governance at scale, to give a clear picture of AI project feasibility upfront and deliver better results in implementation.

It's crucial to remember that responsible AI must underpin any AI strategy or project, to ensure its use is ethical and compliant. AI outputs must be transparent and traceable, enabled by citation, explainability, and the capability to continuously monitor and audit model use, sentiment, accuracy, and toxicity.

Humans – the next big thing in AI?
Despite predictions of AI replacing humans, the reality is that successful enterprise AI depends heavily on human expertise – not just in managing implementation, but in pinpointing the right use cases where AI can deliver real value. This shift towards 'humans-as-a-service' is evident in OpenAI's $10m investment in a new AI consulting arm and Synechron's $1bn valuation. These moves point to a growing recognition that AI is not plug-and-play; it requires strategic guidance, contextual understanding, and human judgement to align its use with complex business needs. In practice, it's not about AI replacing humans, but about humans enabling AI – and AI enabling humans.

> "For our organisation, the biggest barrier to AI adoption is the underlying quality of our data, and reluctance to invest in improving foundational data maturity."
>
> **Senior Data Leader**
> Utilities industry

> "We need to be sure that our data is properly managed before letting AI loose on it."
>
> **IT Manager**
> Professional Services firm

> "The complex, siloed nature of our business is the biggest barrier to adopting AI at scale."
>
> **Senior IT Leader**
> Professional Services firm

Chapter Three

Hype does not equal ROI.

This approach is about shifting the focus from shiny new AI use cases to outcome-led use cases. AI should be contextualised so that it is directly relevant to the business, ensuring every innovation or AI idea is firmly rooted in the organisation's specific context and needs.
The process begins with a clear understanding of the problem – asking why it exists and how it can be resolved – before jumping to solutions. Being KPI-led is essential, with a focus on core business metrics and quantifiable impact. Solutions should be framed in terms of the value they can deliver; rather than setting out to "build a chatbot," the real goal might be to "reduce 20% of inbound service calls through intelligent triage" – with a chatbot being just one possible option. Ultimately, it's about choosing the right tool for the job, rather than defaulting to AI just for the sake of it.

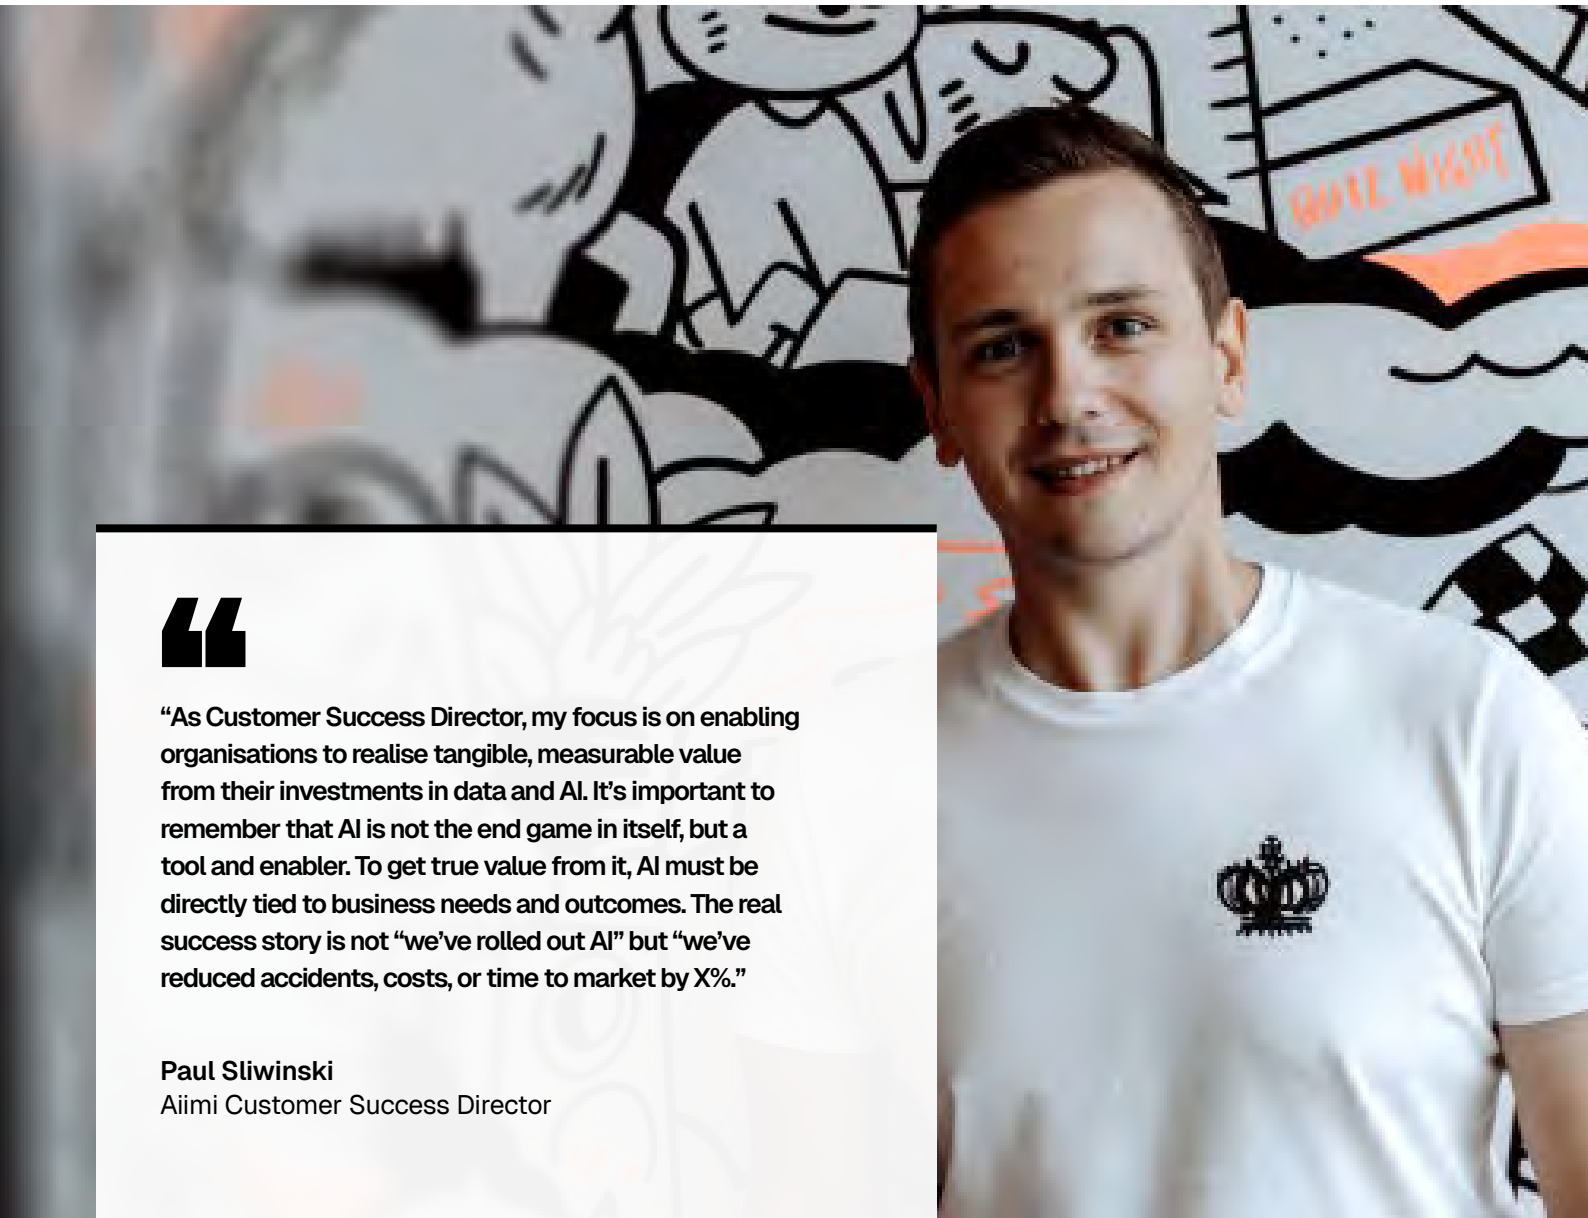

"As Customer Success Director, my focus is on enabling organisations to realise tangible, measurable value from their investments in data and AI. It's important to remember that AI is not the end game in itself, but a tool and enabler. To get true value from it, AI must be directly tied to business needs and outcomes. The real success story is not "we've rolled out AI" but "we've reduced accidents, costs, or time to market by X%."

Paul Sliwinski
Aiimi Customer Success Director

Chapter Three

Defining value for AI.

AI initiatives should always be referenced against business problems and value drivers; identifying the value you're targeting, understanding how the business realises that value, and establishing clear success measures to ensure you're on the right track.

Key categories of value driver:

- **Commercial** (baseline, business case, contractual obligations)
- **Productivity** (output improvements)
- **Efficiency** (speed, time saved, automation)
- **Performance** (system and human impact)
- **Social** (safety, confidence, wellbeing)
- **Risk** (reduction of accident or operational risks)
- **Economic** (cost reduction, revenue generation)
- **Strategic** (alignment with outcomes and goals)

Benefit realisation must be considered from the outset: how will the benefits be achieved, and is AI truly the best solution, or might something else be more appropriate? It's also vital to consider who is impacted, who influences success, and to recognise that different business areas will prioritise different factors – such as risk versus growth. Alongside this, organisations must weigh up considerations like data quality, risk mitigation, and AI readiness – remembering that AI will be a cost centre until it starts enabling better decisions and delivering measurable impact.

Chapter Three

Metrics for AI impact.

Tracking the right success metrics is critical.

Non-functional metrics, such as the number of data sources integrated, accuracy percentages, or data throughput, are typically the focus of technical or data teams and should be designed to support functional outcomes.

Functional metrics, on the other hand, directly reflect business impact – such as percentage reduction in safety events, increases in productivity, or hours saved.

Opportunity cost is also important: what new value is unlocked by freeing up resources?

Ultimately, benefit realisation should underpin your business case – whether that's revenue increase, cost reduction, risk mitigation, safety improvements, innovation, or strategic alignment. The key is to choose metrics that best relate to your business drivers and value realisation.

Chapter Three

Evolution of value.

Value from AI evolves over time. In the short term, the focus should be on evidencing that the proof of concept (PoC) works. Over the longer term, it's about operationalising the solution, driving adoption, and realising sustained benefits. Continuous improvement is essential: define what needs to be included in the benefits case, establish how you'll track it, and be ready to adjust your approach to maximise value as you learn and grow.

It's easy to get stuck in a cycle of building endless pilots to test the art of the possible, celebrating demos over outcomes, and burning budget without a clear path to ROI. In these scenarios, hype often drowns out strategy. The most common gaps include a lack of focus on real business challenges, an absence of scaling strategies, and no clear ownership of ROI or accountability for impact.

Winning with AI means solving real business problems at scale, with measurable outcomes. The golden thread of value must run through every initiative: start with a real business problem, map the impact on KPIs and ROI, quantify value early and often, assign ownership for outcomes, and plan for success and scale from day one.

Key principles for success:

- → Measure value in ways that matter to the business (not vanity metrics).
- → Prioritise initiatives based on their value and start measuring early.
- → Plan for long-term benefit realisation (e.g., ROI ratio over multiple years).
- → Drive economies of scale to ensure AI capability is sustainable.

Chapter Three

Best next steps for AI use cases:

The best enterprise AI use cases start with the right questions, align with business value, and are delivered through agile, experiment-led approaches. By engaging users, fixing data quality, and embedding responsible AI, organisations can quickly start releasing demonstrable value from AI, fuelling the appetite for further AI projects.

↘

- → **Identify impactful business challenges and frame clear questions to guide your AI initiatives**
- → **Involve users early, improve data quality, and prioritise responsible practices to accelerate value**
- → **Use agile, experiment-based methods to deliver quick wins and build momentum for more AI projects**

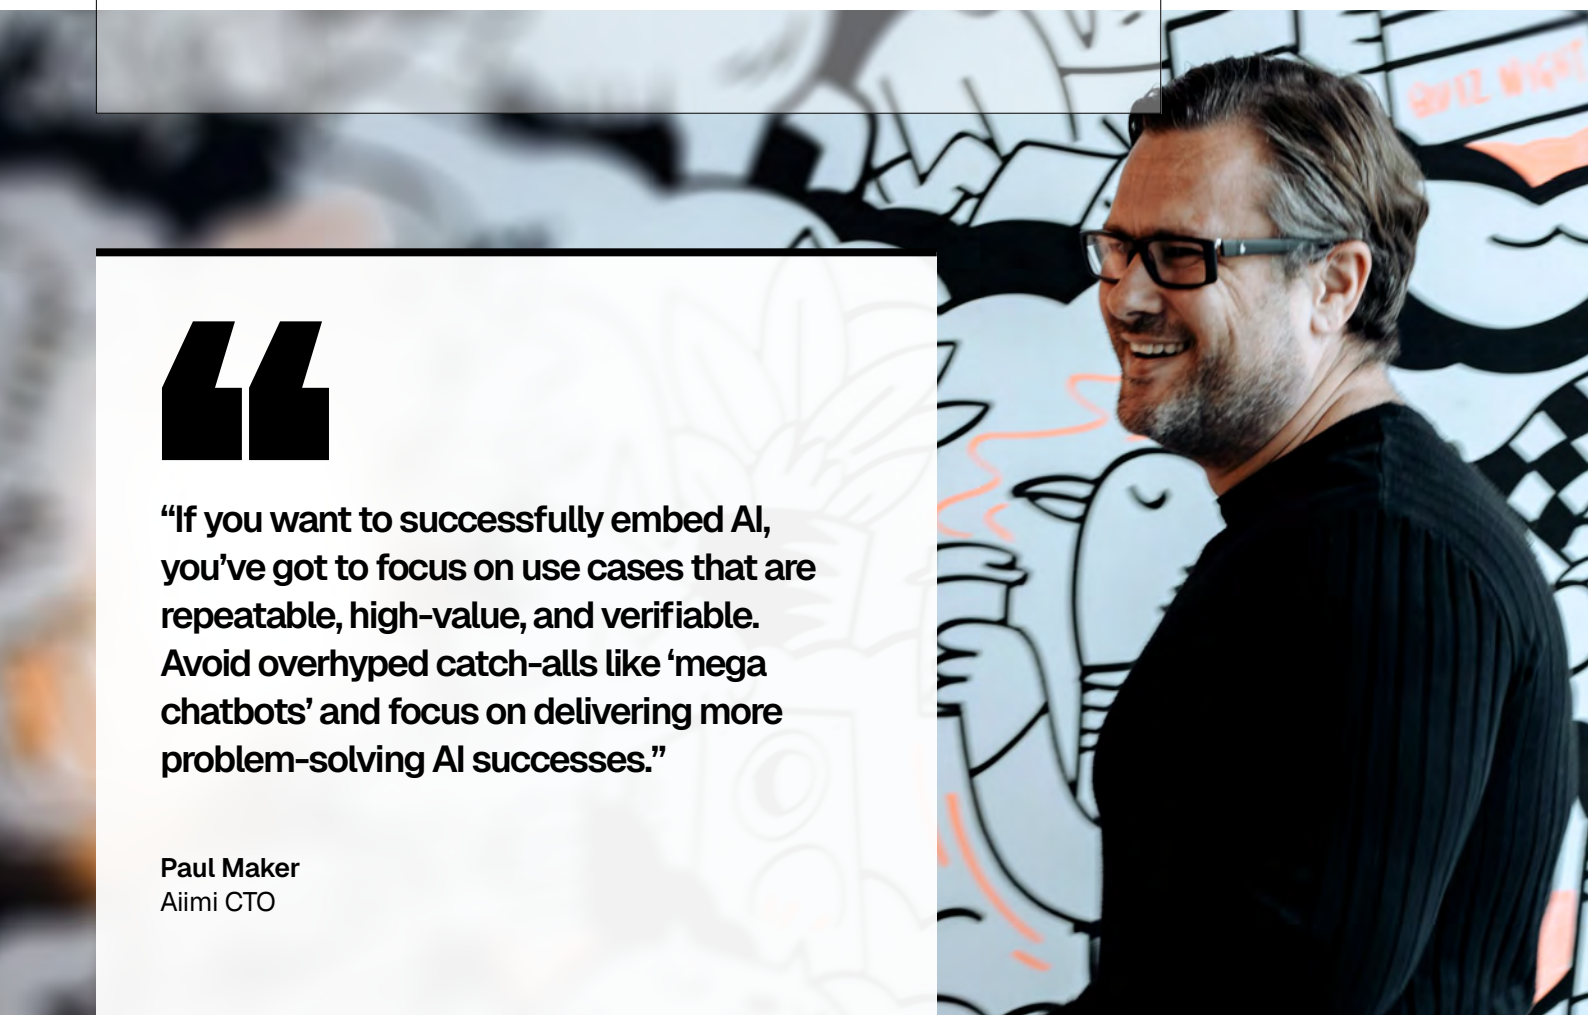

> **"If you want to successfully embed AI, you've got to focus on use cases that are repeatable, high-value, and verifiable. Avoid overhyped catch-alls like 'mega chatbots' and focus on delivering more problem-solving AI successes."**

Paul Maker
Aiimi CTO

aiimi

Chapter Three

CASE STUDY

Creating a three-year AI roadmap for the Government's Insolvency Service.

The Insolvency Service knew AI could supercharge its work supporting the UK economy. But with so many potential applications, it was difficult to know where to start.

Through extensive research, we worked with their users to identify the best AI use cases, then applied our four-part framework (right) to evaluate and prioritise this list and develop a three-year AI roadmap.

Value.

AI should unlock better decisions, productivity gains, cost reduction, safety improvements, or something new. We plan how to measure impact and benchmark the current state, so you can see the difference AI makes.

Data.

AI needs good data to succeed. We work out whether the right data exists to make this use case fly - and what kind of data governance work might be needed.

Risk.

All AI comes with considerations - regulatory, compliance, or ethical. Risk could mean physical safety, IP exposure, or brand damage. We consider what guardrails you'd need to progress an AI use case.

Feasibility & Cost.

What would it actually take to go from AI prototype to AI that thousands of users can adopt safely and securely? We stay grounded in AI that's going to make an impact.

Read more about this case study:

aiimi

Chapter Four

Compliance in the age of automation and transparency.

In most organisations, personal storage systems like email and OneDrive present significant compliance challenges. These platforms are typically locked down to individual users or those they explicitly invite, unlike document libraries or network drives that are governed at a much higher level. As a result, every person's OneDrive becomes a complex web of permissions, often used as a private workspace or a 'safe space' for work in progress.

Over time, however, people forget to delete old files, and these spaces can accumulate sensitive or customer information that should never have been stored there in the first place. This data can linger indefinitely, creating hidden risks.

The problem is compounded by the way information is shared. In the moment, a contract or sensitive file might be shared with a few colleagues for a specific business process. Years later, when that information is needed – perhaps in the context of legal hold or export control – it can be impossible to locate, as permissions may have changed or access may have been revoked. These scenarios create powerful business cases for compliance and Data Security Posture Management (DSPM) initiatives that focus on cleaning up personal storage and ensuring sensitive data is properly managed. And this challenge extends to everything within the enterprise.

→ cont.

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

4

Chapter Four

THROWBACK TO THE FUTURE

"Over a decade ago, I worked with the CEO of a major global bank who'd been hired to lead a transformation towards a more customer-centric retail experience. His vision was to create a welcoming space that broke away from the cold, transactional feel of traditional banking – replacing reinforced glass with cosy sofas and staff bending over backwards to look after you.

To deliver consistently excellent customer service, the CEO believed that every piece of information held about a customer, and every conversation with or about them, should be seen as the customer's property. He encouraged his employees to imagine the customer sitting on their shoulder, listening and watching, whenever they were resolving an issue or completing a task on their behalf. Without this mindset, he said, customer service would always fall short.

This vision is now closer than ever to being a widespread reality. Today, if a business holds information on a customer or employee, then why shouldn't that individual have real-time access to it? This will be the foundation for maturity in data privacy and compliance – a dynamic, live 'case file' that is always up to date and accessible to the data subject at any time. This isn't just a technical aspiration; as regulatory demands and customer expectations outgrow manual processes and patchwork solutions, it becomes an absolute necessity. And it's not just about people; we should look for these live, gold-standard collections of data wherever they add value – physical assets, for instance.

Privacy and data management will be key market differentiators, and organisations that lead in these areas will set new standards. Over time, these data-first security practices will become the norm, and it will seem inconceivable that we ever operated in any other way. To get there, organisations will depend on robust information classification and automated governance for visibility into where sensitive data resides and who can access it – not just as enablers of compliance, but as the building blocks of trust and competitive advantage."

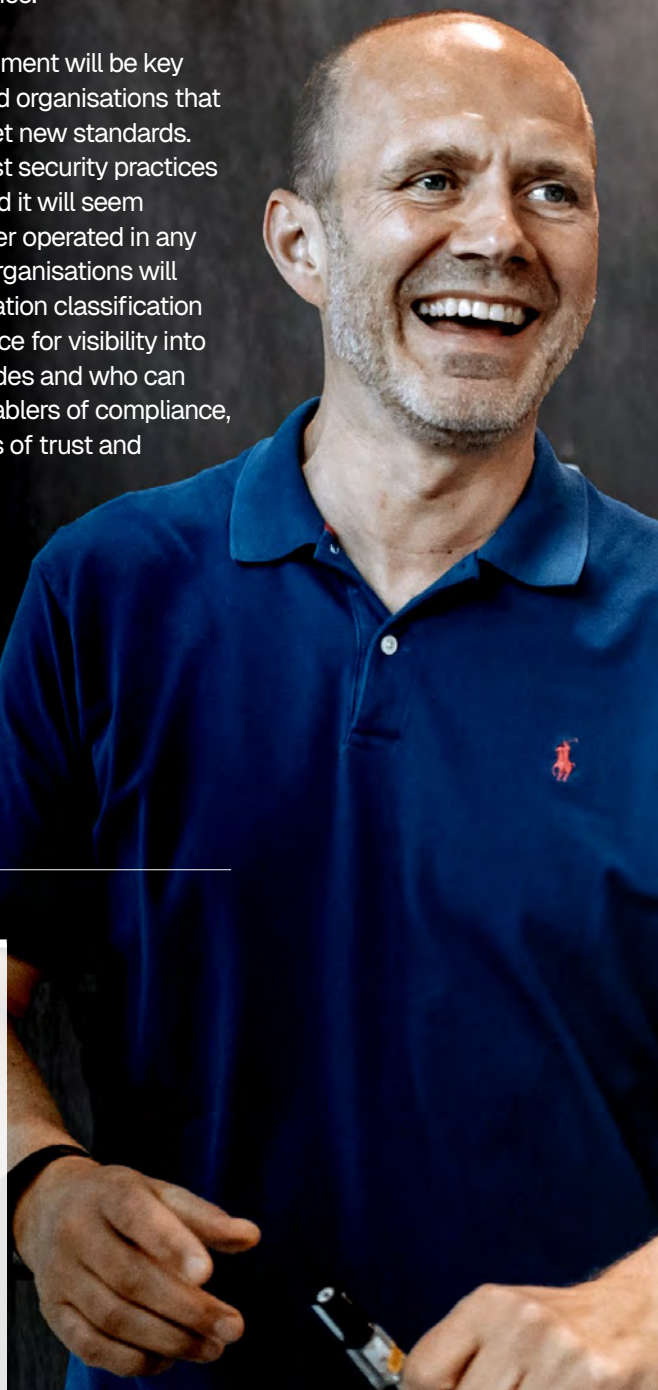

> **"Privacy and data management will be key market differentiators, and organisations that lead in these areas will set new standards."**
>
> **Steve Salvin**
> Aiimi CEO

Chapter Four

From manual to agentic.

Today, we already have the capability to classify complex data, create collections, and audit and manage risk within source data - for example, migrating information to a secure location during a clean-up operation, or redacting sensitive information that should not be disclosed, while still allowing relevant information to be shared. Some parts of those processes already harness AI. But the future vision is for intelligent AI Agents to handle the entire process end-to-end, delivering fully automated data management.

Currently, automated redaction is supplemented by human review to ensure compliance with legal and staffing requirements. The goal is to further systematise this process, so that over time, more of the review can be handled by intelligent agents with a high degree of confidence. Eventually, this could lead to a scenario where, for example, a Data Subject Access Request (DSAR) is fulfilled almost instantly: the system assembles and redacts the relevant information, a human reviewer gives final approval, and the data is released – all within minutes.

The key steps in this journey are:

- → Moving from semi-manual data discovery processes to automated, agent-driven processes
- → Transitioning from manual redaction to automated data protection systems, with human oversight
- → Gradually increasing the confidence and capability of automated systems, reducing the need for manual intervention
- → Exploring innovative techniques to ensure accuracy and security, such as using one AI to redact and another to review or challenge the redactions.

Open data and changing expectations.

A natural extension of these compliance and data security capabilities is opening up the data you can share – like responding to Freedom of Information (FOI) and Environmental Information Requests (EIR). Rather than treating these as one-off processes, there is an opportunity to make information more openly available by default. For example, in the water industry, there is a push to publish as much information as possible via Stream, the Ofwat-funded open data initiative, in order to derive greater value from it as a collective. If this approach becomes standard, it would make sense for organisations to publish responses to FOIs or EIRs on open data sites, making them accessible to all rather than just the requester.

This shift reflects a broader change in expectations. As more organisations embrace openness and transparency around data, we're moving towards a world where much of our data will either be open or individually owned, with less existing in between.

Chapter Four

Best next steps for compliance:

To move forward, organisations should prioritise auditing and cleaning up data, ensuring sensitive data can be found and is properly managed. Automated tools enable accurate classification, redaction, and governance at scale, and transparent data practices will make real-time information accessible to data subjects by default. Fostering a culture of data ownership and transparency prepares teams for the shift to real-time, agent-driven compliance processes.

↘

- → **Prioritise data audits and clean-up, with proper management of sensitive information**
- → **Use automated tools for scalable classification, risk-scoring, redaction, and governance**
- → **Encourage data ownership and transparency in readiness for agent-driven, real-time compliance and accessibility**

aiimi

Chapter Four

CASE STUDY

Saving an international retailer 3,000 hours with 4x faster Disclosure Requests.

International retailer Holland & Barrett faced a surge in complex DSARs, with manual processes taking up to 90 days and straining resources. Aiimi's Workplace AI Platform transformed their approach, automating setup, redaction, and data discovery across systems. DSAR setup now takes 15 minutes, average delivery times are down to a week, and the retailer is saving 3,000 hours annually – all while returning accurate, trusted data.

Read more about this case study:

"I don't know where we'd be without the Aiimi platform – it's been a lifesaver. The difference it's made to our delivery timeline is incredible. And, crucially, we're pulling data we can trust and which enables a data subject to receive their data more quickly."

Hollie Mela
DPO Holland & Barrett

aiimi

Chapter Five

Deploying AI Agents for smarter data management.

Looking ahead, we see data governance and management being driven by a suite of configurable AI Agents, working 24/365 to organise and improve enterprise data stored across thousands of corporate systems.

The central use case for deploying these AI Agents is to enhance security and compliance, and to ensure enterprise data is primed and ready for AI. At the same time, these agents will dramatically improve information retrieval, enabling businesses to shrink the world down to only the information needed to perform a given task or make an informed decision.

AI-powered data management will provide organisations with unprecedented visibility into their data landscape. Businesses will be able to see exactly how much data they have, where it's stored, and filter this information based on infinite criteria. From this central vantage point, data management and remediation steps can be taken to ensure data is always under control.

→ cont.

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

5

Chapter Five

Classification.

A core capability of these AI Agents will be advanced data classification – creating and executing classification models for enrichment, correction, reinforcement training, monitoring, and auditing. Whether operating autonomously or with human oversight, these agents will apply a variety of classifications to content and data records, such as document type, security level, export control, personal information, and risk level. Organisations can tailor their own classification schemes or adopt industry-standard/off-the-shelf schemes. As part of this process, AI Agents will automatically identify and label data that is redundant, obsolete, or trivial (ROT), ensuring it can be excluded from enterprise search and retrieval processes to improve results, and empowering data governance teams to manage ROT appropriately.

Migration.

In the future, we see migration as an integrated part of the always-on Data Governance & Management experience. This capability will be consumed by AI Agents and configured by data governance teams, with background jobs capable of safely and reliably moving data from any source to any source. This will include the ability to manipulate metadata and permissions, handle exceptions and errors, and provide audit and analytics on how migration jobs have performed. This capability will be used for cloud migrations, legacy decommissioning, archiving, disposition, and for moving data into the right location.

Permissions and access.

Another key function of AI Agents will be to monitor and manage data permissions – identifying data that is either too widely accessible or too tightly restricted, and giving data governance teams the ability to adjust permissions or relocate data as needed. The primary use cases will be to improve security and compliance by ensuring sensitive information is protected, and to improve information retrieval and knowledge management by ensuring valuable data is accessible to those who need it.

Data quality.

AI Agents will play a crucial role in improving and maintaining data quality, by profiling data records to identify incomplete or incorrect attributes, then automatically finding and correcting these issues using information stored in related unstructured data. This systematic approach will help data governance teams improve data quality automatically, and even generate more granular data attributions to support deeper analytical reach.
As part of this continuous data quality work, AI Agents will generate and maintain real-time collections of ‘golden records’ – kite-marked as the most up-to-date, authoritative version that the business wants people to use and query against. These collections can draw from any source, including document libraries, OneDrive, email attachments, or corporate data stores. Golden records can then be used by chatbots or RAG systems to find accurate, relevant information and perform tasks. They will also help data governance teams identify and remediate any inappropriate data storage practices.

Chapter Five

Best next steps for smarter data management:

By deploying AI Agents and automation for smarter data management, organisations can transform the way they handle information – making it more secure, compliant, accessible, and ready for further AI use cases. This approach also drives greater efficiency, reduces manual effort, and delivers significant cost savings by automating routine data management tasks and streamlining workflows across the enterprise.

↘

- → **Use AI Agents to strengthen information security, compliance, and accessibility, for further AI projects**
- → **Replace repetitive manual data management tasks with automation to boost productivity**
- → **Achieve cost efficiencies by streamlining how data flows across the organisation**

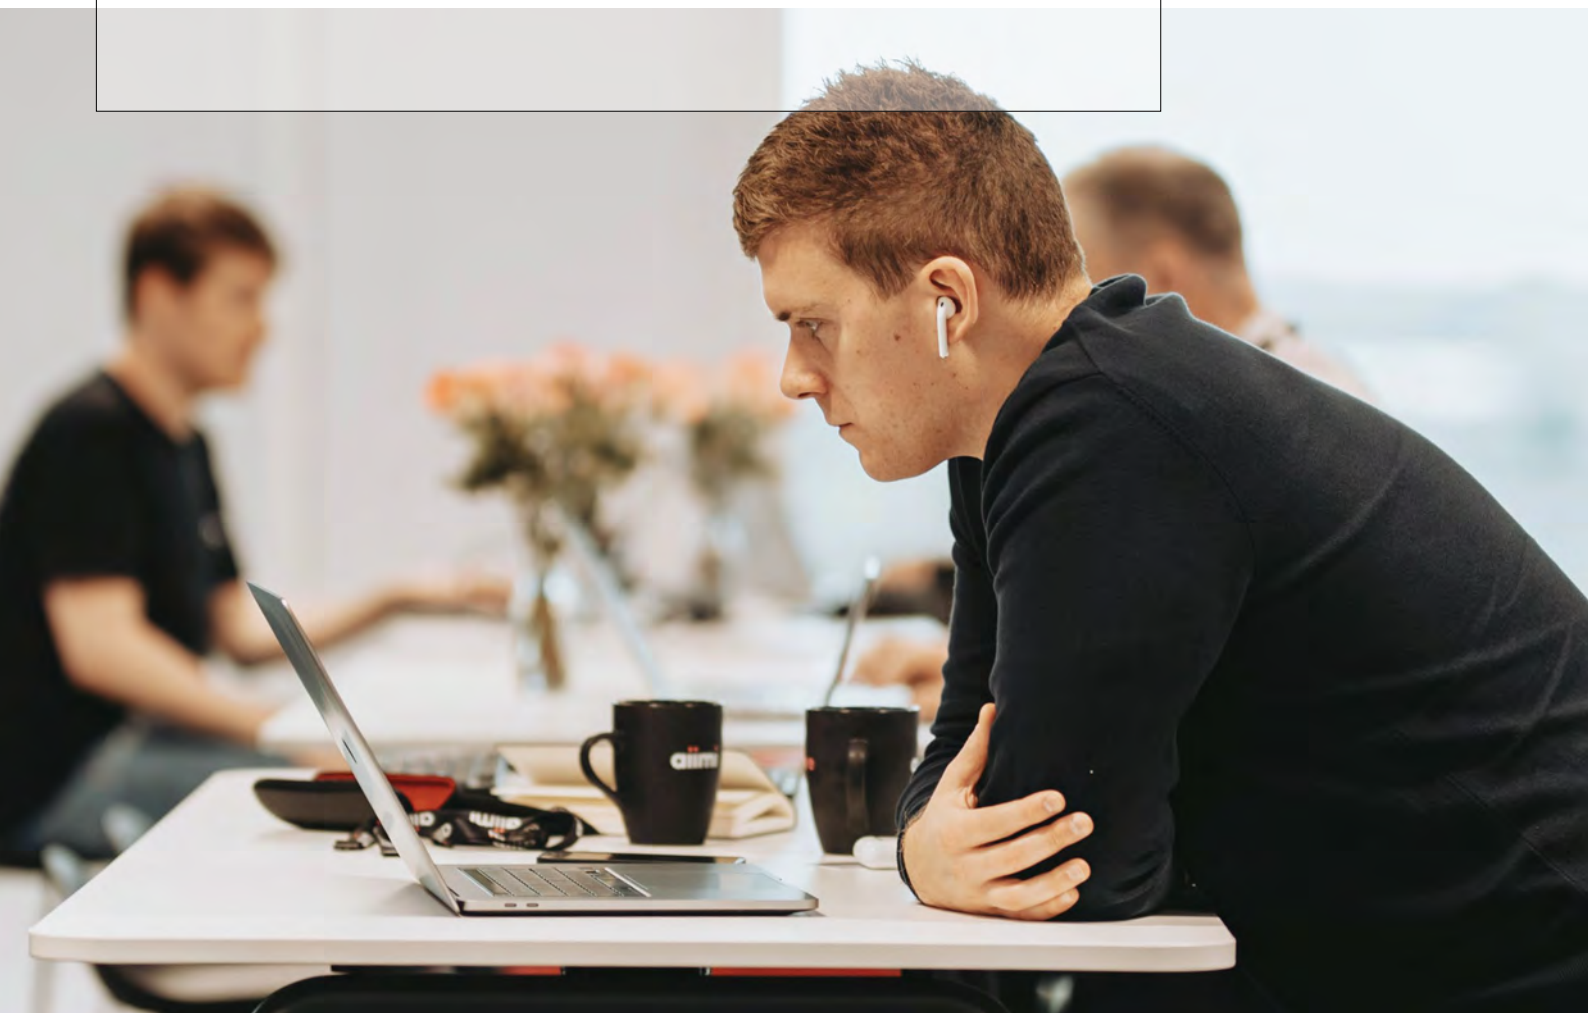

aiimi

Chapter Six

The evolution of AI search.

A key challenge in enterprise search has always been relevance. Traditional search often returns thousands of matches, overwhelming users with irrelevant results and making it difficult to find the two or three items that really matter. **Several strategies can be combined to address this:**

- **Better classification:** By classifying information more effectively, we can narrow search results to only what's relevant to the query or prompt.
- **Semantic understanding:** By developing capabilities around semantic layers, we can interpret the meaning behind a search, rather than relying solely on keywords.
- **Contextual awareness:** By factoring in the user's role, intent, and current task, we can surface the most relevant information for each individual.

The result is a connected enterprise where, despite the vastness of data, users can quickly get to what they need. For example, three users searching the same term would see different results – a finance user searching for 'Databricks' might see invoices and purchase orders, while a commercial user would see statements of work, and a sales user would find proposals. This personalised approach is the direction of travel: classifying and connecting everything, then shrinking the world down to the smallest possible, most relevant packet of information.

6

ESSENTIAL READING FOR:

CDO & CIO | HEAD OF DIGITAL & INNOVATION | DATA PROTECTION OFFICER | INFORMATION SECURITY LEADER | HEAD OF INFORMATION OR RECORDS MANAGEMENT

Chapter Six

The universal data layer.

We've always envisioned the Aiimi search app as a single, go-to place for enterprise customers with many systems and search options. While some scenarios call for a dedicated app – such as mobile access where corporate systems aren't available – the broader vision is to augment existing search capabilities by connecting our virtual data layer via APIs or Model Context Protocols. This allows systems like ServiceNow, Workday, SAP, or Microsoft to tap into our index and search across multiple repositories, returning unified results.

In the last couple of years, the proliferation of AI assistants in every application has only reinforced the need for this approach. These assistants are limited to the data within their own app and can't connect to information across other systems. Aiimi's platform enables search and AI assistants to access enterprise-wide information, breaking down those silos to deliver far greater value to the user.

A common question is whether the rise of APIs and MCP clients in every repository makes our platform redundant. The answer is no. **There are two main approaches to enterprise search:**

- **Federated search:** Each system is searched individually, and results are aggregated. This approach is often slow and struggles with permissions and result prioritisation.
- **Enterprise view:** All systems are indexed in one place, enabling fast, consistent, permission-aware search across the entire enterprise.

Our experience shows that the enterprise index approach delivers a better user experience and more reliable results. While federated search may work for connecting a few systems, as organisations bolt on more and more connections, complexity and resulting inefficiency quickly escalate.

Aiimi's journey began with enterprise search, and our customers and founding partners helped drive that. We saw that the need to replace Google Search Appliance was a common challenge – and while many organisations have some form of enterprise search (often within the Microsoft 365 landscape), critical business information is scattered across hundreds of corporate systems, including outside of Microsoft.

From the outset, our goal was to create a single platform that could search across all these systems – documents, data, media, web content, messages – so users could find what they needed, no matter where it was stored.

Chapter Six

Beyond search: conversations and actions.

Another significant change in recent years is the shift in user expectations. As consumers, we've become accustomed to conversational interfaces – tools like ChatGPT and also chat-based e-commerce journeys. This shift is transforming enterprise search. Users now want to bypass the step of landing on a document or website and instead receive the answer or action they need, along with recommendations for other relevant content that they hadn't even considered. As we become more comfortable with conversational prompts, it becomes easier to capture user intent and deliver not just search results, but actionable insights and generative responses – enabling users to not just find something but do something.

For AI assistants and agents to be effective, they must be anchored in high-quality, trustworthy information. To achieve this, enterprise search must connect users and systems with authoritative sources and deliver nuanced, context-aware insights at scale. Chat applications such as Microsoft Copilot can use classification labels to determine which documents should be included in the information retrieval process, enabling more relevant and context-aware responses. And by classifying data outside of Microsoft systems and making this available to tools like Copilot, organisations can extend these benefits outside the systems these apps would otherwise be limited to, unlocking even greater value. Ultimately, the credibility and impact of AI-powered automation will depend on how well enterprise search platforms can surface accurate, relevant knowledge.

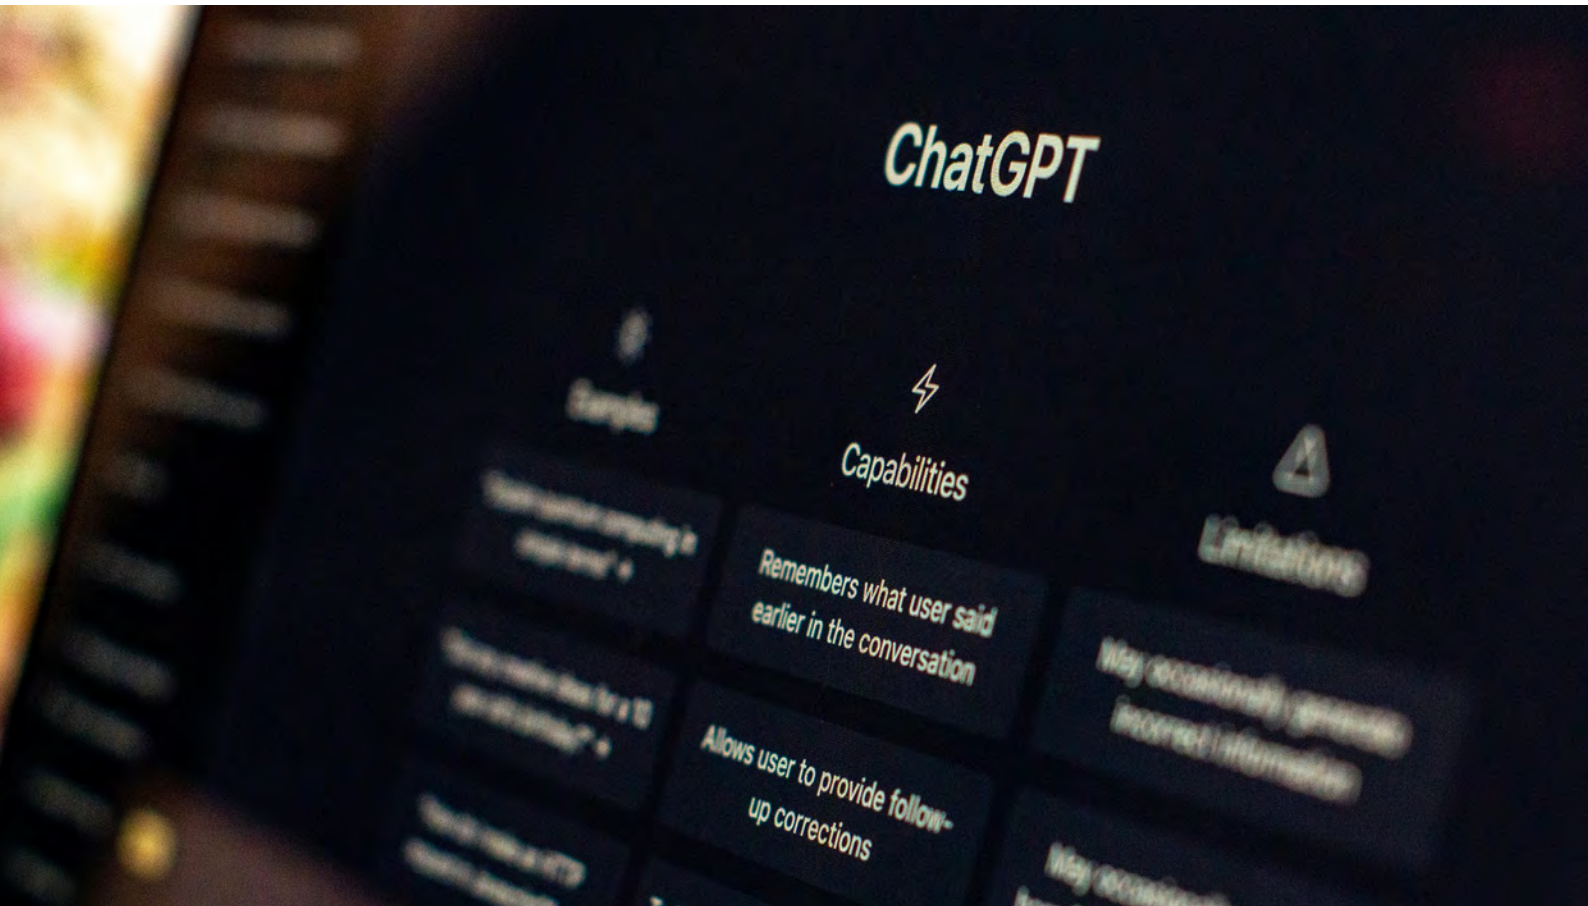

Chapter Six

Best next steps for AI search:

Organisations should focus on AI capabilities for advanced classification, semantic understanding, and contextual awareness to deliver more relevant, personalised search experiences. As user expectations shift towards conversational, action-oriented interfaces, investing in AI-driven assistants that leverage enterprise-wide data will be key. Ultimately, the goal is to provide users with not just search results, but actions and answers, wherever and however they work.

↘

- → **Focus on AI capabilities that deliver more relevant search experiences**
- → **Invest in conversational AI assistants that can tap into enterprise-wide data**
- → **Move beyond search results to provide actionable insights and answers**

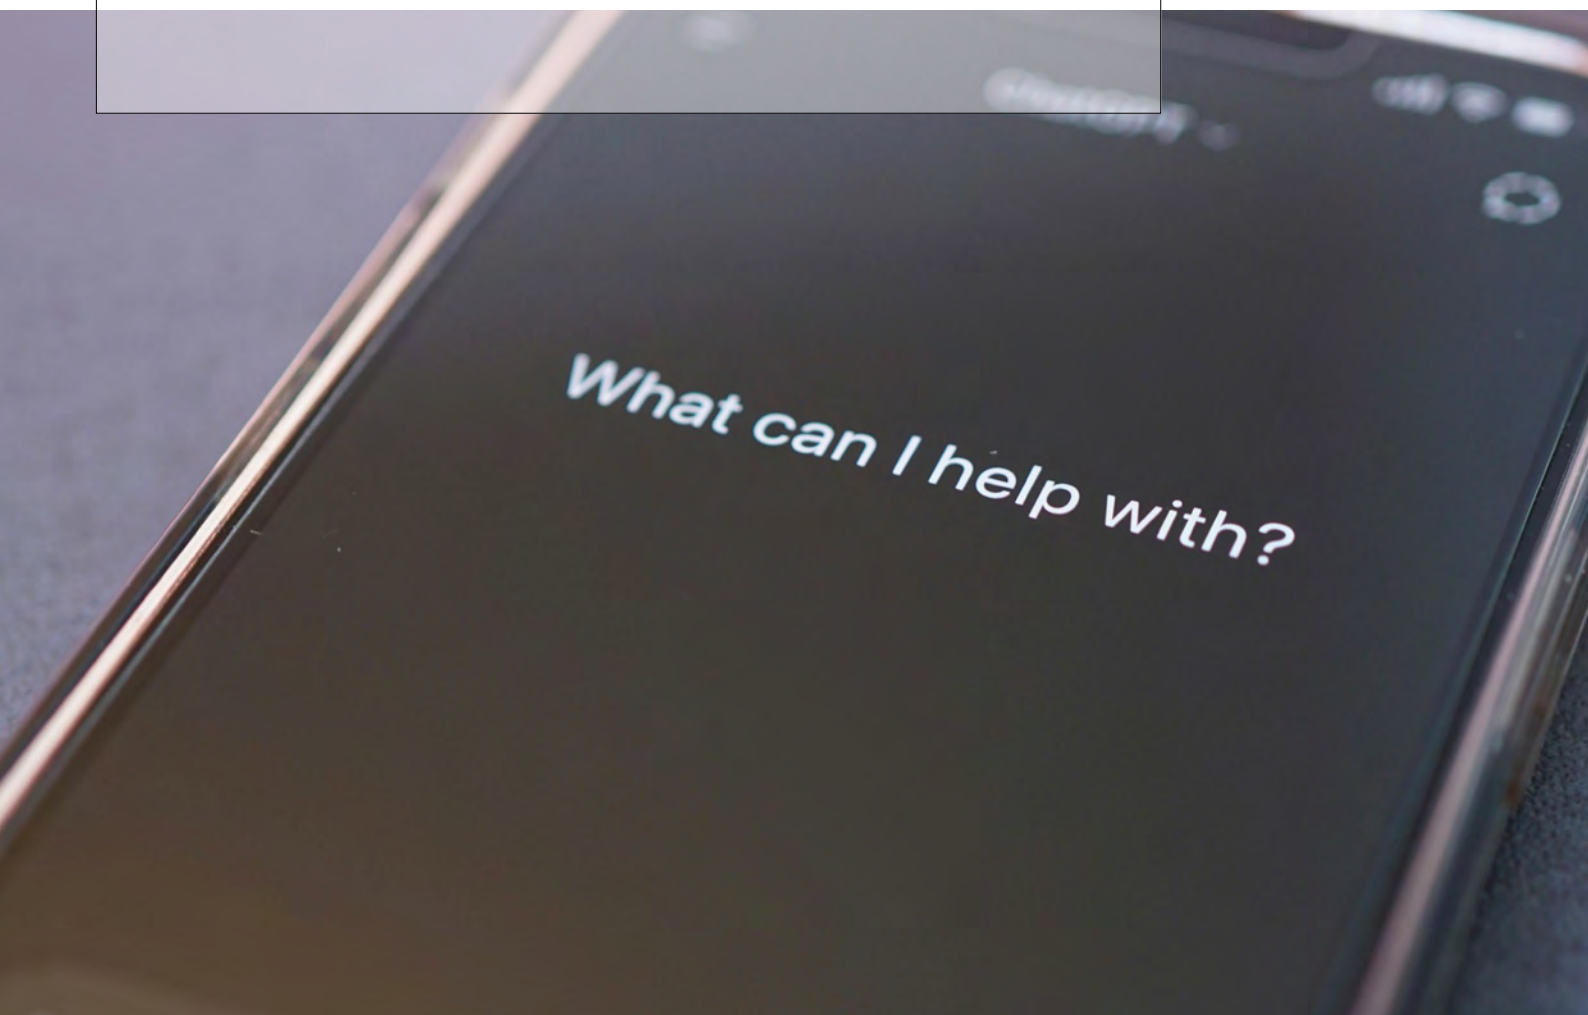

Chapter Six

CASE STUDY

Enabling AI search with 330% ROI for a major manufacturer.

This global manufacturer wanted the ability to navigate hard-to-reach historical data, in order to optimise processes, unify its information, and, ultimately, save time and money.

We implemented Aiimi Workplace AI to index 400 million documents and drawings across multiple sites, creating a comprehensive knowledge network. This has enabled users to quickly find relevant data, identify subject matter experts, and evidence design decisions, significantly improving efficiency. Searches that previously took up to 90 minutes now take as little as five – saving 30 minutes to four hours per search and delivering up to 80% efficiency gains. With a 75% first-search success rate and over 330% annual ROI, the solution has become the go-to platform for Product Engineering, driving operational, quality, and manufacturing improvements, reducing risk, and unlocking value from historic data.

Read more about this case study:

Aiimi Events

Aiimi expert events: aligning AI ambitions with value.

AI has become a strategic priority, reshaping how organisations create value, build resilience, and stay competitive. Our expert-led events give business leaders the practical tools and insights needed to drive meaningful progress with AI.

Aiimi's recent NavigateAI event, hosted in partnership with Iron Mountain, brought together over 40 leaders and practitioners from sectors including retail, manufacturing, government, engineering, asset management, education, insurance, and healthcare. Through a packed agenda of roundtables, presentations, and workshops, we explored the practical realities of AI adoption.

These events offer a unique opportunity to gain transferable insights from real-world AI journeys, while engaging in open, collaborative discussions about challenges, lessons learned, and future ambitions. Structured around the pillars of Value, Data, Risk, and Feasibility, our sessions give organisations the tools to move beyond isolated pilots to enterprise-wide AI transformation.

Join us at our next event to connect with peers, learn from industry experts, and leave with actionable strategies to accelerate your AI journey.

"The conversations very quickly turned to value; where to find it, how to measure it, and when to walk away. In the end, technology is just a tool – what matters to leaders is whether it creates something worth having."

Joshua Swords
Aiimi Head of Data & AI Engineering and speaker at Navigate AI

Hear about upcoming events or explore opportunities to join a session.

→

"A great event - my first Aiimi Navigate event and the topic of the Operationalisation of AI could not be more relevant! For me, the key takeaway was the continued criticality of Data Quality to ensure AI deployment success."

Head of Digital at global engineering group

"Such a high-quality event – I left the room far more confident about getting started, and inspired about what might be possible."

Head of Economic Development at Milton Keynes City Council

"A genuinely insightful experience. Each workshop provided valuable takeaways – it was great to hear practical insights from industry leaders and connect with others passionate about the future of AI and data."

Senior Data Engineer at Pluxee


aiimi
WE ARE AIIMI.
POWERING AI
EVERYWHERE
AT WORK.
AIIMI.COM