Innodata

Q2 2024 Investor Presentation

AUGUST 2024 | NASDAQ: INOD

Disclaimers

Forward-Looking Statements

This presentation contains forward-looking statements within the meaning of Section 21E of the Securities Exchange Act of 1934, as amended, and Section 27A of the Securities Act of 1933, as amended. These forward-looking statements include, without limitation, statements concerning our operations, economic performance, financial condition, developmental program expansion and position in the generative AI services market. Words such as “project,” “believe,” “expect,” “can,” “continue,” “could,” “intend,” “may,” “should,” “will,” “anticipate,” “indicate,” “predict,” “likely,” “estimate,” “plan,” “potential,” “possible,” “promises,” or the negatives thereof, and other similar expressions generally identify forward-looking statements..

These forward-looking statements are based on management’s current expectations, assumptions and estimates and are subject to a number of risks and uncertainties, including, without limitation, impacts resulting from ongoing geopolitical conflicts, including between Russia and the Ukraine, Hamas’ attack against Israel and the ensuing conflict and increased hostilities between Iran and Israel; investments in large language models; that contracts may be terminated by customers; projected or committed volumes of work may not materialize; pipeline opportunities and customer discussions which may not materialize into work or expected volumes of work; the likelihood of continued development of the markets, particularly new and emerging markets, that our services support; the ability and willingness of our customers and prospective customers to execute business plans that give rise to requirements for our services; continuing reliance on project-based work in the Digital Data Solutions (DDS) segment and the primarily at-will nature of such contracts and the ability of these customers to reduce, delay or cancel projects; potential inability to replace projects that are completed, canceled or reduced; our DDS segment’s revenue concentration in a limited number of customers; our dependency on content providers in our Agility segment; the Company’s ability to achieve revenue and growth targets; difficulty in integrating and deriving synergies from acquisitions, joint ventures and strategic investments; potential undiscovered liabilities of companies and businesses that we may acquire; potential impairment of the carrying value of goodwill and other acquired intangible assets of companies and businesses that we acquire; a continued downturn in or depressed market conditions; changes in external market factors; the potential effects of U.S. monetary policy, including the interest rate policies of the Federal Reserve; changes in our business or growth strategy; the emergence of new, or growth in existing competitors; various other competitive and technological factors; our use of and reliance on information technology systems, including potential security breaches, cyber-attacks, privacy breaches or data breaches that result in the unauthorized disclosure of consumer, customer, employee or Company information, or service interruptions; and other risks and uncertainties indicated from time to time in our filings with the Securities and Exchange Commission.

Our actual results could differ materially from the results referred to in any forward-looking statements. Factors that could cause or contribute to such differences include, but are not limited to, the risks discussed in Part I, Item 1A. “Risk Factors,” Part II, Item 7. “Management’s Discussion and Analysis of Financial Condition and Results of Operations,” and other parts of our Annual Report on Form 10-K, filed with the Securities and Exchange Commission on March 4, 2024, as updated or amended by our other filings that we may make with the Securities and Exchange Commission. In light of these risks and uncertainties, there can be no assurance that the results referred to in the forward-looking statements will occur, and you should not place undue reliance on these forward-looking statements. These forward-looking statements speak only as of the date hereof.

We undertake no obligation to update or review any guidance or other forward-looking statements, whether as a result of new information, future developments or otherwise, except as may be required by the U.S. federal securities laws.

Non-GAAP financial measures

Financial information contained in this presentation includes certain financial measures that are calculated and presented on the basis of methodologies other than in accordance with generally accepted accounting policies in the United States of America (GAAP), such as Adjusted EBITDA, which include or exclude certain items from the most directly comparable GAAP financial measure. These non-GAAP measures differ from reported GAAP measures and are intended to illustrate what management believes are relevant period-over-period comparisons and are helpful to investors as an additional tool for further understanding and assessing Innodata’s expected ongoing operating performance. A reconciliation of historical non-GAAP financial measures to the most directly comparable GAAP financial measure is included in this presentation. Any non-GAAP financial measure used in this presentation is in addition to, and not meant to be considered superior to, or a substitute for, measures prepared in accordance with GAAP.

About Innodata

Why Invest in Innodata

The technology partner-of-choice for data quality, scale, and agility

Enormous total addressable market (TAM) opportunity for generative AI IT services of $200 billion by 2029[1]

Supporting Big Tech builders and enterprise adopters of language models (LLMs)[2] and other AI for enterprises at scale

Proven, transferrable model that provides opportunity to expand beyond Big Tech

Strong financial position: accelerated revenue growth, strong operating leverage, healthy balance sheet

Experienced leadership team and 35+ year legacy of developing large-scale, high-quality data

Innodata is uniquely positioned as a "picks and shovels" beneficiary of the Generative AI era to deliver long-term value and drive growth

1 McKinsey & Co. "Tech services and generative AI: Plotting the necessary reinvention." June 12, 2024. https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/tech-services-and-generative-ai-plotting-the-necessary-reinvention?stcr=E53D1C8EC1264C69BC3242D9A1CC1543&cid=other-eml-alt-mip-mck&hlkid=32250fd913c1495cbd0c4bfa9bbdb1b1&hctky=12981724&hdpid=4b848883-5426-4f0e-a8df-fed78d178c58

2 A large language model (LLM) is a large-scale artificial intelligence system trained on extensive textual data that leverages deep learning to comprehend, generate, and work with human language effectively. Generative AI refers to a category of AI technologies, including LLMs, designed to create original content, such as text, images, music or other forms of media.

Innodata provides technology, tools, and platforms supported by domain data specialists, technologists, data scientists, and testers across the globe

20+ Delivery Locations

Global infrastructure supporting customers across all time zones

5,000+ Global Experts

Wide range of subject matter expertise across finance, healthcare, legal, science, etc.

85+ Languages

Experts supporting domains for specialty generative AI

"Magnificent 7"

- Contracts with five of the Mag 7
- Contracts with two more Big Tech YTD 2024

SME Domain Specialization

- Data Science
- Computational Linguistics
- Linguistics
- Programming
- Law and Regulation
- Medicine and Allied Health
- Life Sciences
- Physical Sciences
- Business & Finance

Accreditations

Our competitive advantage

Innodata supports multiple AI services subsegments

BUILDERS (LARGE TECH)

AI Scaled Data Solutions

- Data Curation
- Data Transformation
- LLM Instruction Datasets / Demonstrations (Multilingual) (Supervised Fine-Tuning, RLHF, DPO)
- Data Annotation (Across Modalities)
- LLM Trust & Safety

ADOPTERS (ENTERPRISES)*

AI Services and Platforms

- Consulting
- AI Services (AIOps, CoPilots, Model Trust, Safety & Evaluation)
- Digital Services (Cloud, Legacy App and Product Design and Modernization)
- AI Stack Services (Foundation Models, Tooling, Data Architecture, Training Data Preparation, Data Governance, Data Migration)
- Business Process Management / Managed Services
- Vertical Platforms

SCALED DATA SOLUTIONS COMPETITORS

accenture, ALEGION, appen, cloudfactory, clarifai, Dataloop, DEFINED.AI, iMerit, sama, scale, shaip, TELUS International

AI SERVICES AND PLATFORMS COMPETITORS

cognizant, EigenTechnologies, HIVE, (h[s]) HYPERSCIENCE, IBM, Infrrd, Infosys, LeewayHertz, Leverton, QuantumBlack AI by McKinsey

VERTICAL SOLUTIONS COMPETITORS

CISION, COMPLIANCE.AI, eNoah, HUMAN API, Hummingbird, Intrado, KANTAR MEDIA, Meltwater, MUCK RACK, RiskRight EHS, Verisk

*Includes both current and prospective (roadmap) capabilities.

Our technology platforms support both Big Tech and enterprise AI use cases

Vertical Platforms — Niche Industry Workflows

Innodata

Regulatory Change Management Platform
Integrated workflow application enabling legal and compliance teams to make informed decisions. One of world's largest banks is charter customer.

Agility PR Solutions

Media Intelligence Platform
Connect with media influencers, amplify messages, monitor coverage, and measure public relations impact. ~1,500 customers.

Synodex

Medical Data Intelligence Platform
Extracts detailed medical data from patient records for AI analytics, automated processing, and expert review. ~12 customers, including several of largest U.S. and global life insurers.

Horizontal Platforms — AI Data Pipelines

Innodata

LLM Annotation Platform
ML Annotation Platform
Customized platforms for accelerating development of LLM (supervised fine tuning, RHLF) and ML (labeling, annotation) training data at scale

Innodata

Document Intelligence Platform
Extracts intelligence from documents in accordance with a company's proprietary taxonomy or industry taxonomy

Innodata

LLM Safety and Evaluation Platform[1]
Establishes LLM model compliance, bias, and security

Foundational Goldengate Platform[2]

Proprietary, State-of-the-Art Low-Code AI Platform is our Core AI Technology Stack

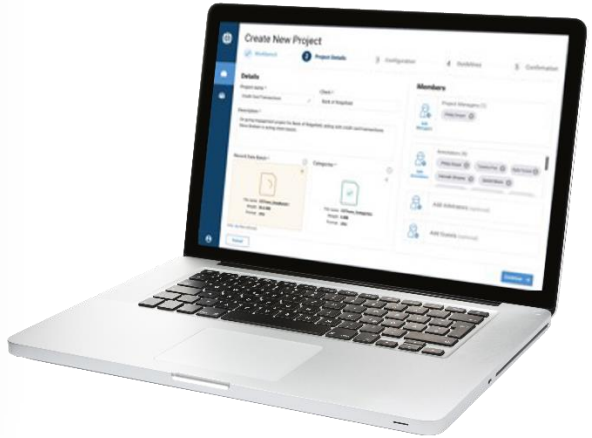

1 Presently in development.
2 Our Goldengate Platform is a proprietary AI foundation model we deeply trained over past six years. We believe it is state-of-the-art for over 50 knowledge tasks. It is at the core of many of our engagements, tools, and enterprise deployments.


Well Positioned in
Growing Market

We have evolved our model to expertly service the world's leading enterprises on generative AI

Applied AI and Machine Learning (ML)

Launched Innodata Labs

Developed proprietary Goldengate foundational AI platform

Used real world insights to expand addressable market and support organizations in their AI journeys

Classical AI for Leading Tech Companies

Rapidly scaled two Big Tech companies

Delivered high-quality data for training cutting-edge AI models

Created new type of training data critical for large language models (LLM)

Generative AI for Leading Tech Companies

Signed seven Big Tech for LLM data engineering, driving significant growth

Includes five of "Mag Seven"

Harnessing our data capabilities and platforms to meet emerging requirements of mega-cap technology companies

Expanding Generative AI Capabilities for Enterprise + Public Sector Opportunities

2017 onwards → 2021 – 2022 → 2022 – 2023 INFLECTION POINT → 2024 – Beyond

WELL-POSITIONED IN GROWING MARKET

We believe we have opportunities for continued growth & expansion

Magnificent Seven Global Tech Companies & Other Large Technology Companies Building Foundation Models

Billions of dollars committed to LLM strategies.

Enterprise AI Services and Platforms (Direct Distribution and via White Label)

A decade from now, all successful companies will be AI companies.

AI/LLM-Enabled Vertical Platforms

Successfully re-inventing the way work is done through LLM-integration/augmentation.

Federal Government

Nascent opportunity.

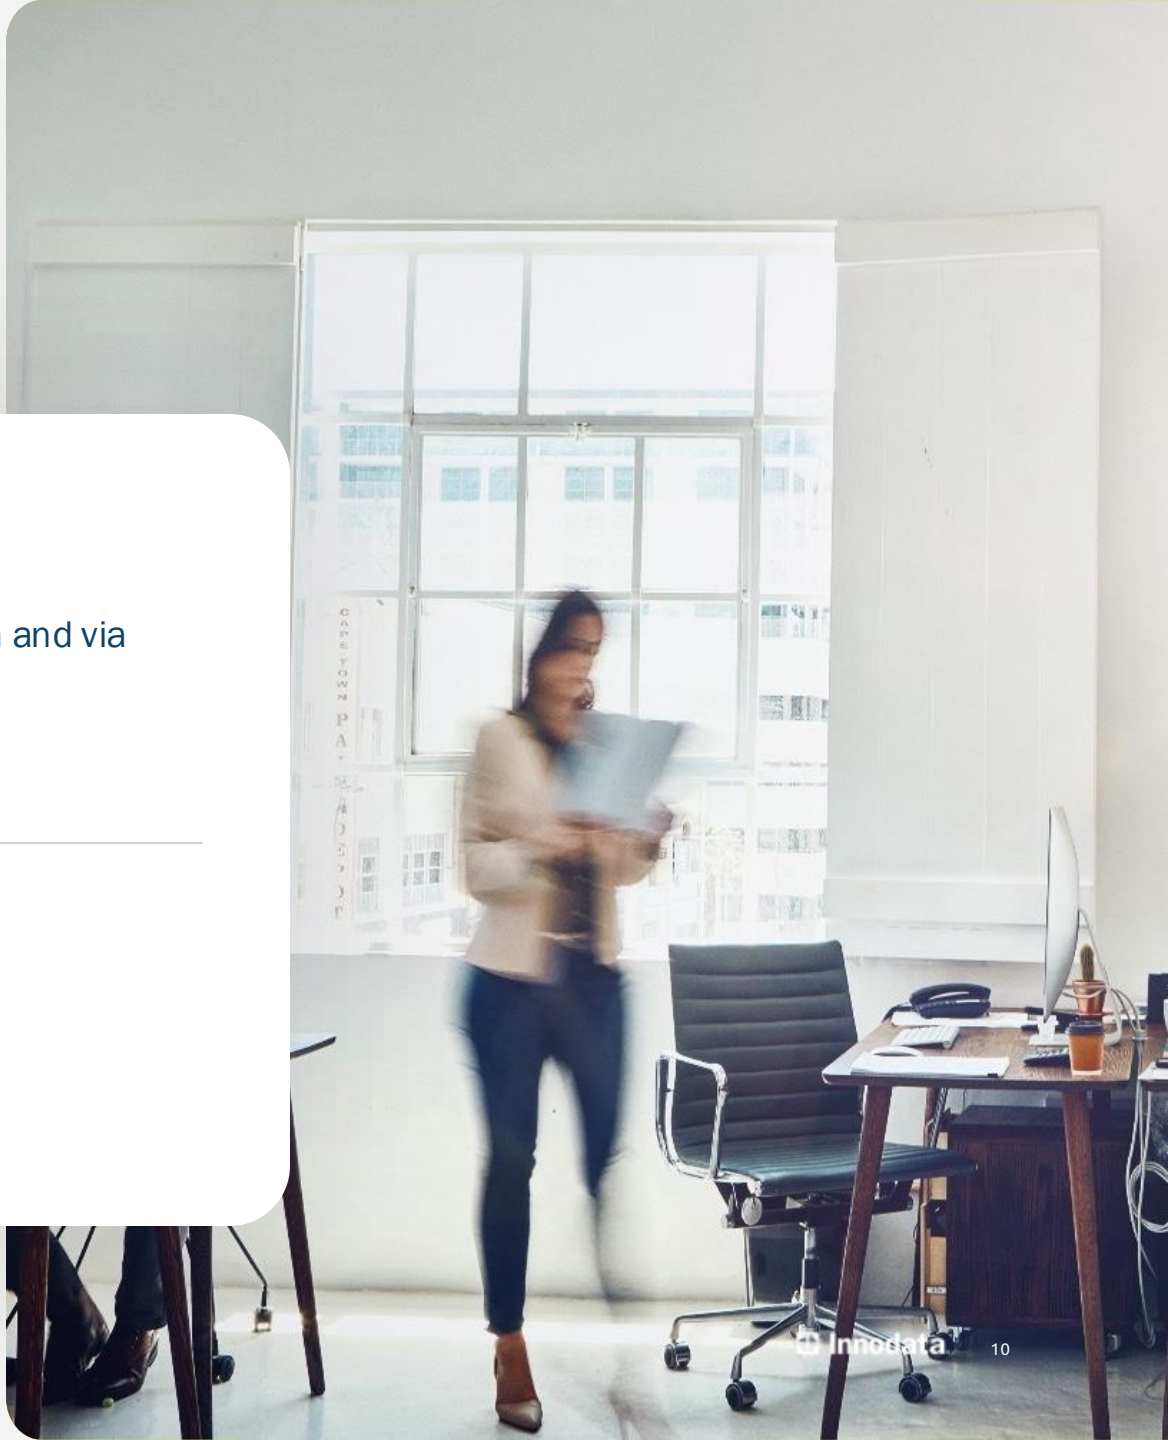

Generative AI is anticipated to be the next-generation computation platform

We believe the generative AI market opportunity will rival the internet and mobile in significance

- "Magnificent Seven" AI arms race now underway — billions allocated[1]
- McKinsey & Co. believes generative AI will potentially:
  - Automate work activities that absorb 60-70% of employee time today by midpoint 2045
  - Add $2.6 trillion - $4.4 trillion to global economy across
    - 63 immediate use cases alone[2]
- Enterprises rank gen AI/LLM as #1 software spending priority[3]

Positioned for next technology wave

- There's limitless use cases for regenerative AI
- We've recreated our skill data in a diversity of languages and use cases
- New technology is fundamentally built on the data we provide

Enormous generative AI addressable market opportunity[4]

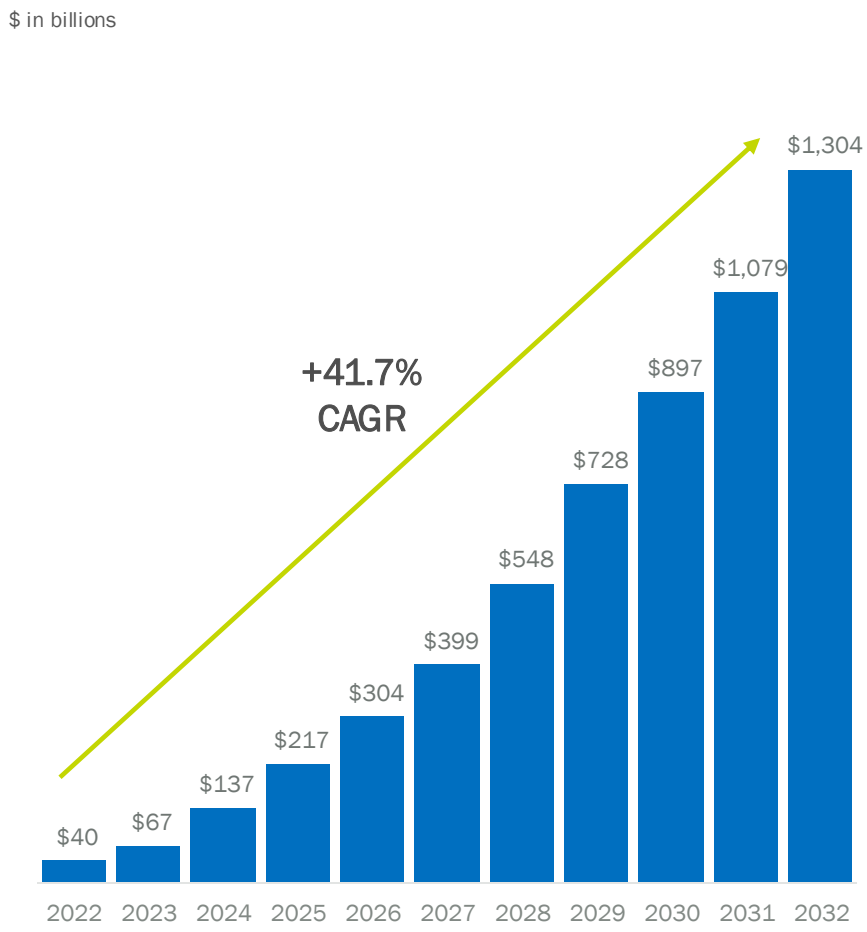

1 Magnificent Seven includes Apple, Amazon, Alphabet, Meta Platforms, Microsoft, Nvidia, and Tesla. 2023 earnings call transcripts for GOOG, AMZN, and META.
See also: https://www.nytimes.com/2023/08/05/technology/tech-nvidia-chips.html.
2 "The Economic Potential of Generative AI: The Next Productivity Frontier", McKinsey & Company, June 2023.
3 "Generative AI: Fuel for Growth Across Tech", Bloomberg Intelligence, August 2023.
4 Bloomberg Intelligence – Generative AI Market Size (Jun-23).

Strategy to Lead Global AI Technology Company Serving Both AI Builders and AI Adopters

Our Vision: To Be the Leading Global AI Technology Company Serving Both AI Builders and AI Adopters

Land

New Customers & Grow Revenue

Drive growth in targeted segments.

Focus Sales and Marketing efforts around priority customers.

Expand

Engagement with Existing Customers

Listen to the needs of our customers and be the highest quality provider.

Utilize our deep domain expertise when engaging with our customers and ensure that our teams are well-versed in our solutions.

Innovate

New Capabilities & Platforms

Expand addressable market by developing new offerings and capabilities.

Continue to invest in our AI technology stack and in our dev/engineering capabilities.

Optimize

Operations & Platforms

Harness the power of AI to drive continuous improvement and long-term value creation.

Regularly evaluate cost structure and investments to simplify and streamline.

Q2 2024 Highlights and Financial Overview

Q2 2024 Key Highlights

66%
YoY organic revenue growth

+2
New Big Tech customers won YTD 2024 (Five of "Magnificent 7" won in 2023)

$87.5M[3]
Program expansions with Big Tech customer (bringing value of customer relationship to approximately $110.5 million in annualized run-rate revenue)

60% or more
organic revenue growth anticipated for 2024 (up from 40% guidance prior quarter)

33%
Adjusted Gross Margin[1] (Excluding unusually high recruiting costs[6], Adjusted Gross margin is 44% for Q2'24)

9%
Adjusted EBITDA[1] (Excluding unusually high recruiting costs[6], Adjusted EBITDA margin was 20% for Q2'24)

$16.5M[2]
Cash on balance sheet

67[4]
Net Promotor Score (2023 survey)

$107.3M
June 2024 LTM Revenue

~84% / 16%[5]
Services / SaaS

1 Non-GAAP measure. Please see "Reconciliation of Non-GAAP Measures" on Slide 21.
2 Cash and short-term investments.
3 Anticipated annual revenue run-rate after ramp-up.
4 Represents the weighted average (based on FY'23 revenue) of the Net Promotor scores for each of our three business segments.
5 SaaS revenue represents revenue from our Agility Media Intelligence business segment.
6 Recruiting costs were $3.6 million in Q2'24.

Q2 2024

In millions US$

Q2-2024 YoY

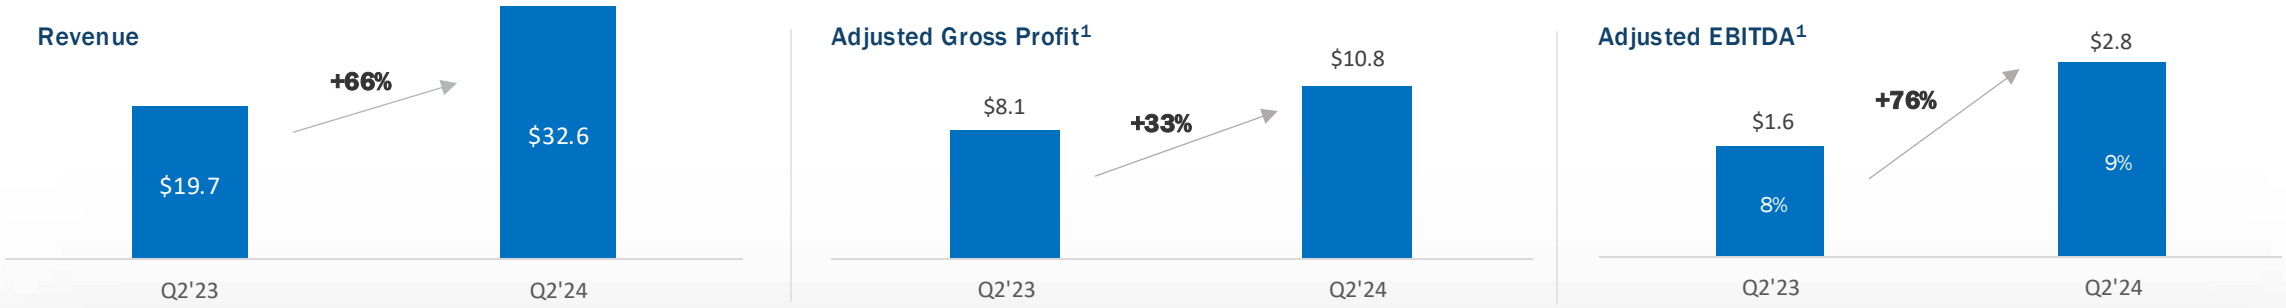

In millions US$

Historical 2022 through Q2-2024

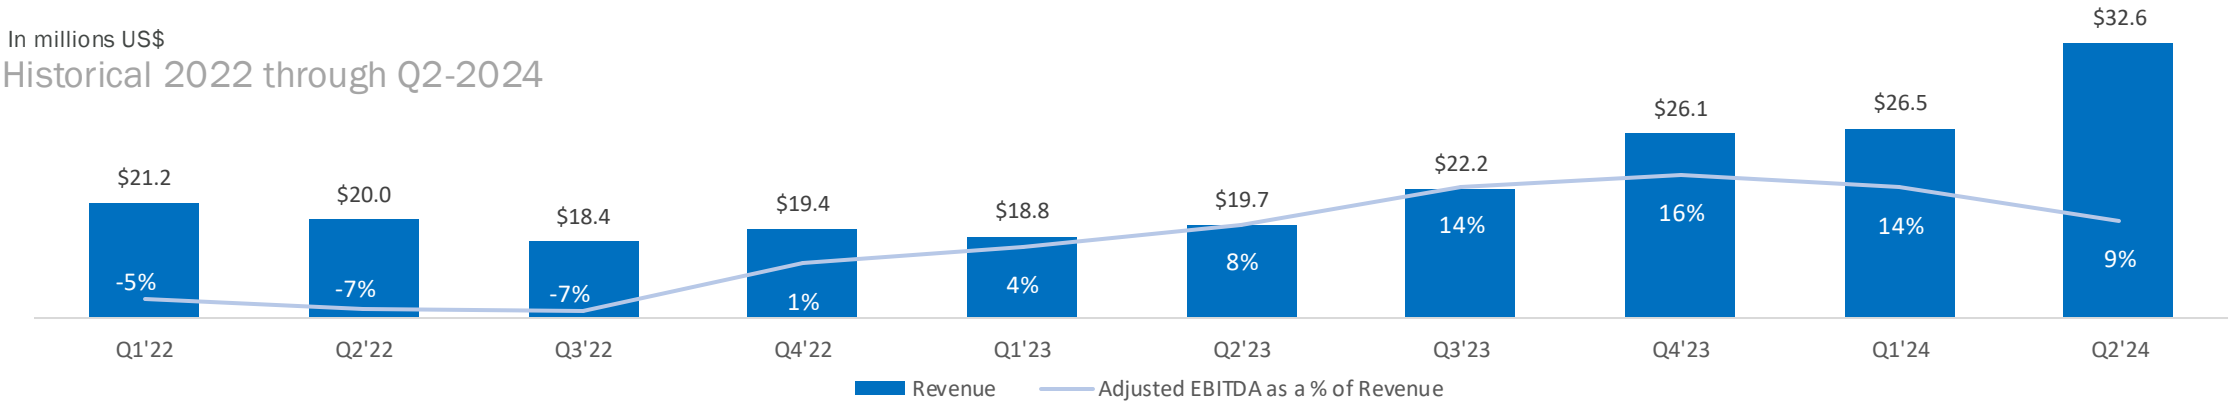

$3.6m of atypically high recruiting costs reduced profitability for the quarter. Backing out recruiting costs, our Adjusted Gross Margin would be 44% (vs. 33%), and Adjusted EBITDA would be 20% (vs. 9%), demonstrating strong operating leverage.

1 Non-GAAP measure. Please see "Reconciliation of Non-GAAP Measures" on Slide 21.

1H 2024

In millions US$

June 2024 YTD

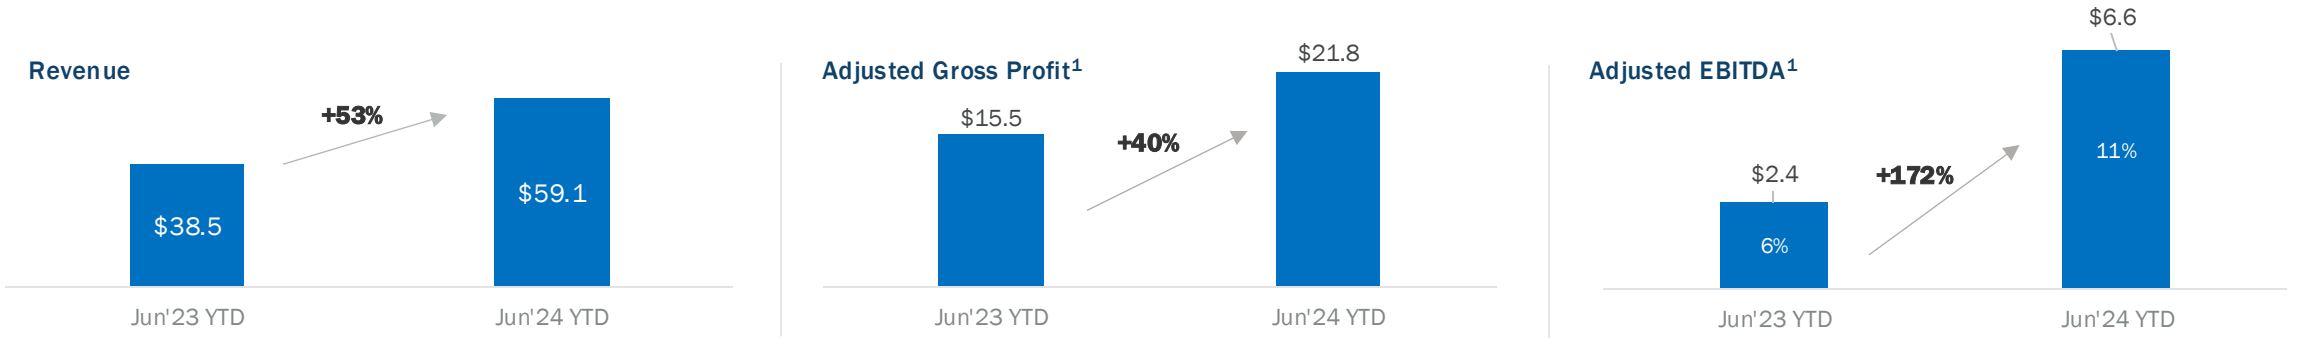

$3.6m of atypically high recruiting costs reduced profitability in Q2. Backing out recruiting costs, our 1H Adjusted Gross Margin would be 43% (vs. 37%), and Adjusted EBITDA would be 17% (vs. 11%), demonstrating strong operating leverage.

1 Non-GAAP measure. Please see "Reconciliation of Non-GAAP Measures" on Slide 21.

Revenue Growth

Q2-2024 was a record revenue quarter, capping five consecutive quarters of growth

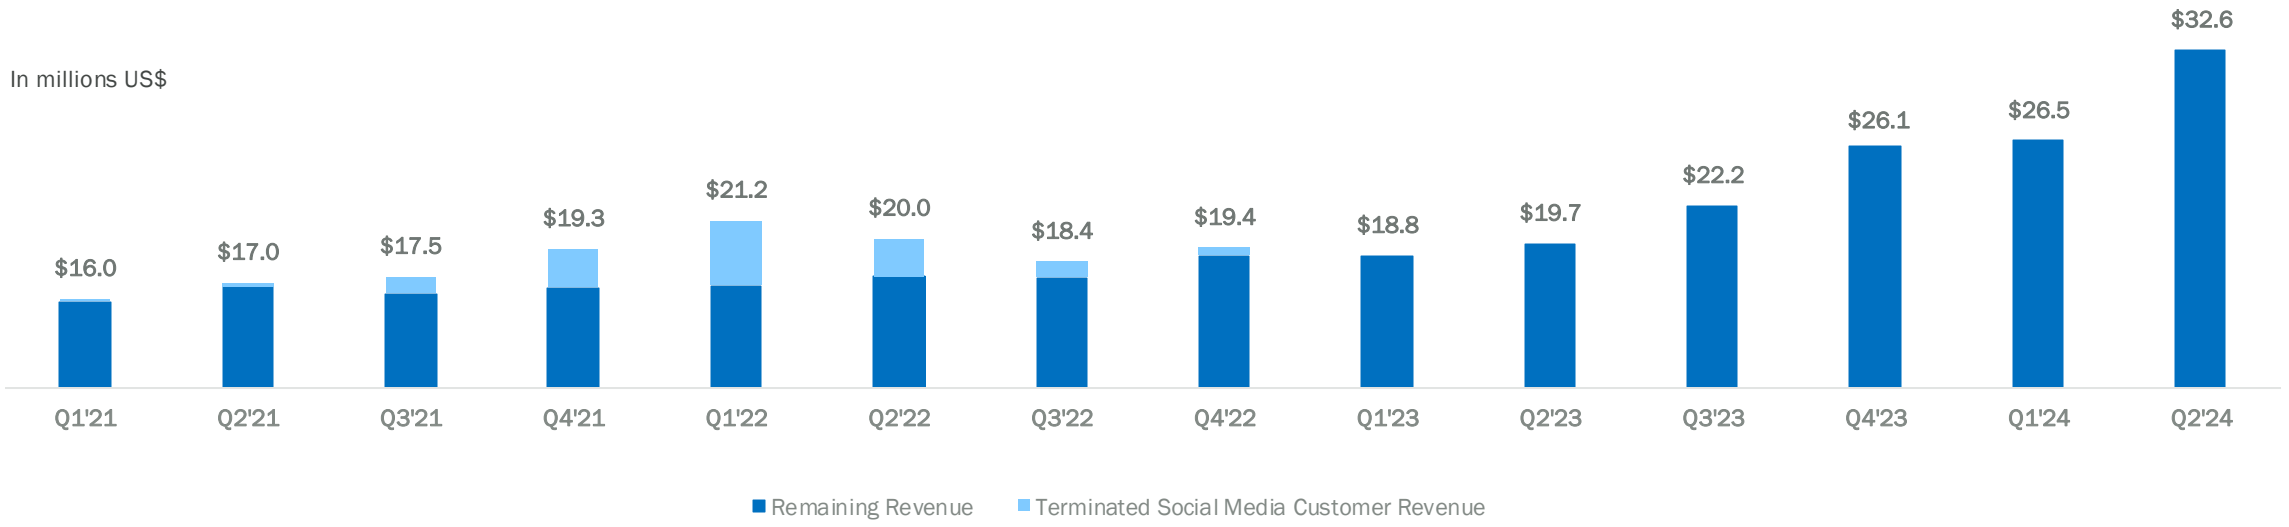

Our relationship with a large social media company, for which we provided AI data training services, terminated as the result of a highly-publicized take-private. Revenue from this relationship is depicted in light blue, above.

We increased our FY24 revenue guidance to 60% or more as a result of strong momentum in our business

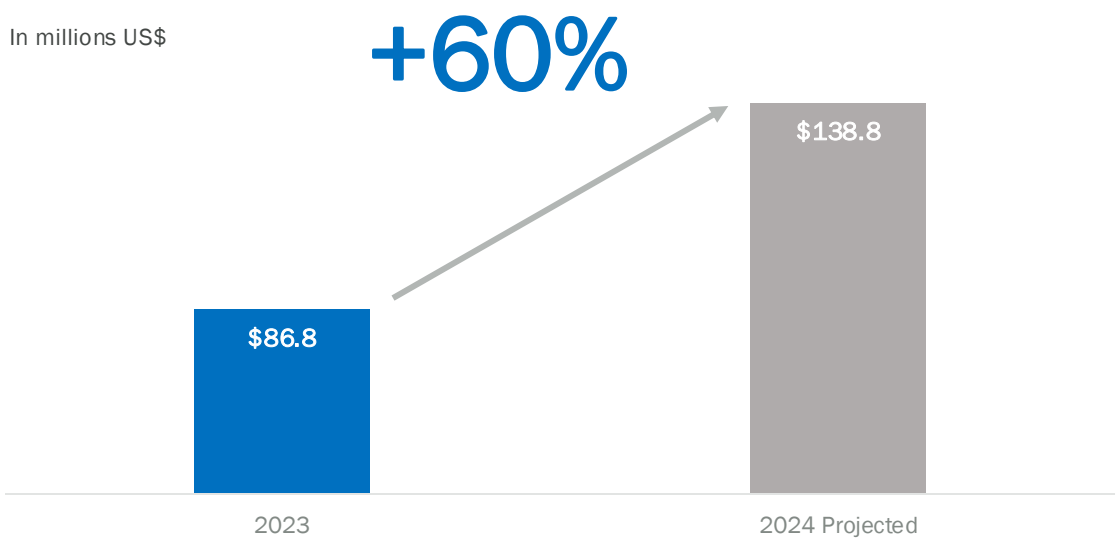

We are forecasting to triple our Adjusted EBITDA[1] from $2.8m in Q2'24 to $8.4m in Q3'24, demonstrating strong operating leverage in our business

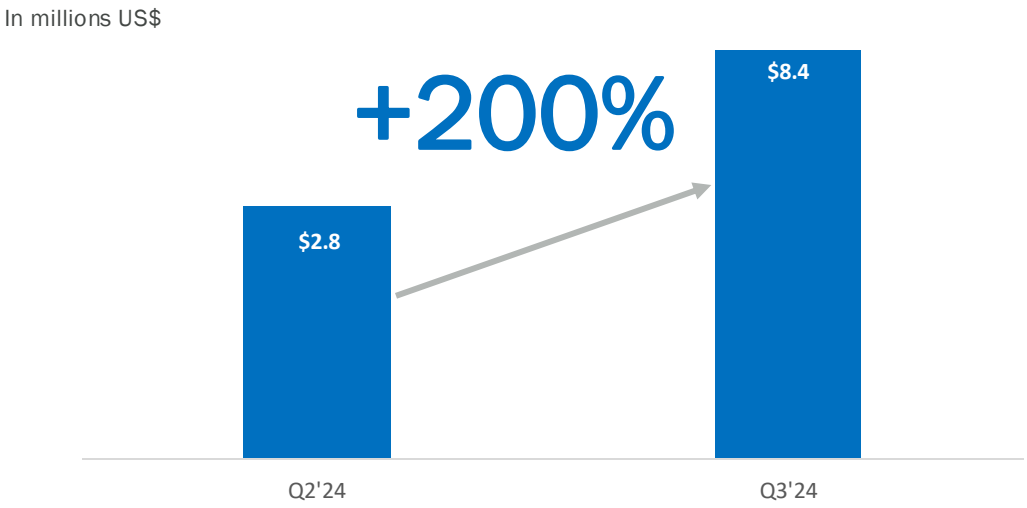

1 Non-GAAP measure. Please see "Reconciliation of Non-GAAP Measures" on Slide 21.

We have a healthy balance sheet with high net cash

Cash and short-term investments

$16.5M

of cash and short-term investments as of Jun'24

Balance sheet

Clean

balance sheet with no external debt/borrowings

Available credit facility

$30M

line of credit with Wells Fargo, no draw-downs to date[1]

1 Subject to a borrowing base limitation, with an accordion feature that enables it to expand to up to $50 million subject to the approval of Wells Fargo.

Reconciliation of Non-GAAP Measures

$ in Millions	FY'22	FY'23	H1'24	Q2'24
Net Income (loss) attributable to Innodata Inc. and Subsidiaries	**-$11.9**	**-$0.9**	**$1.0**	**$0.0**
Tax Provision	$1.5	$1.0	$0.7	$0.3
Interest Expense	$0.0	$0.4	$0.2	$0.1
Depreciation and Amortization	$3.9	$4.7	$2.7	$1.4
Severance	$0.0	$0.6	$0.0	$0.0
Stock-based compensation	$3.3	$4.0	$2.0	$1.0
NCI	-$0.1	$0.0	$0.0	$0.0
Adjusted EBITDA	**-$3.3**	**$9.9**	**$6.6**	**$2.8**

$ in Millions	FY'22	FY'23	H1'24	Q2'24
Gross Profit attributable to Innodata Inc. and Subsidiaries	**$27.5**	**$31.3**	**$19.0**	**$9.4**
Depreciation and amortization	$3.8	$4.6	$2.6	$1.4
Severance	$0.0	$0.3	$0.0	$0.0
Stock-based compensation	$0.2	$0.3	$0.2	$0.1
Adjusted Gross Profit	**$31.5**	**$36.5**	**$21.8**	**$10.8**
Gross Margin %	**35%**	**36%**	**32%**	**29%**
Adjusted Gross Margin %	**40%**	**42%**	**37%**	**33%**

Appendix

Data engineering services are at the core of generative AI and LLMs

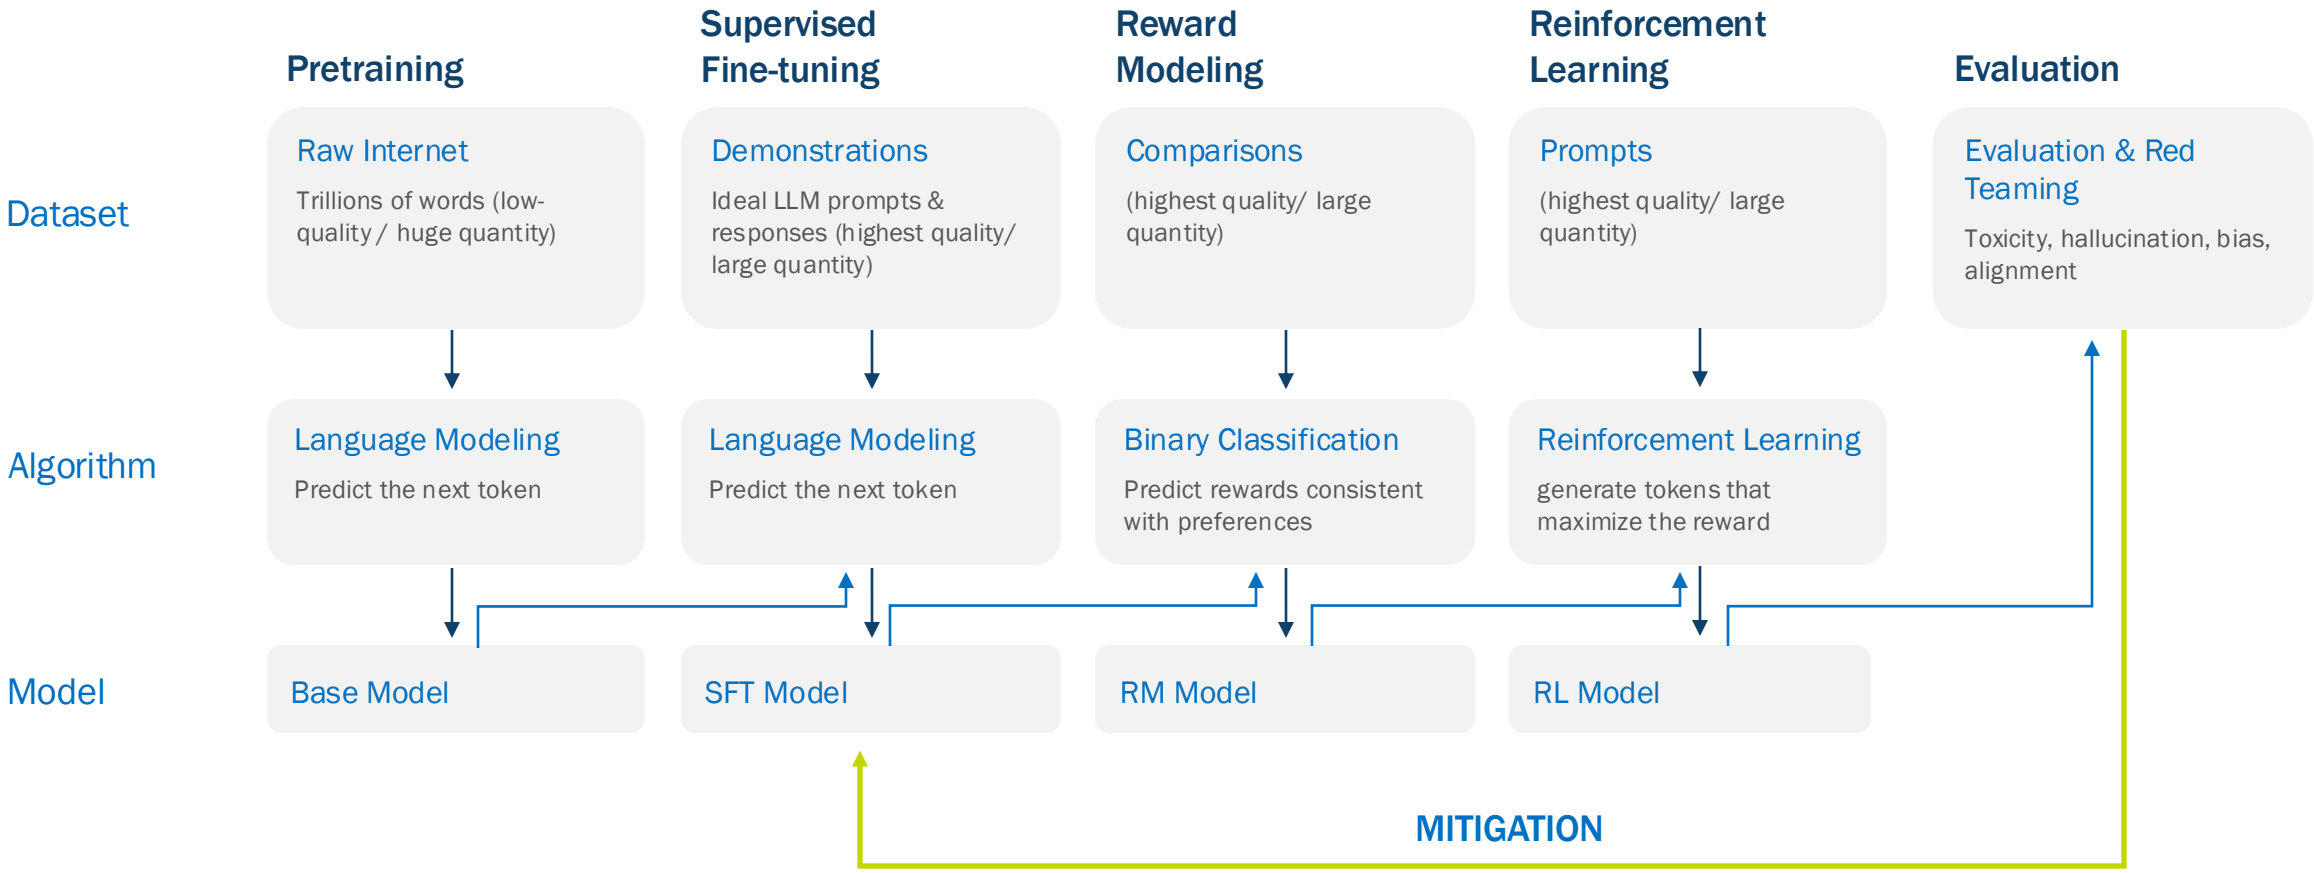

We're helping align LLM foundation models with human values

The Challenge

One of the world's leading hyperscaler cloud providers and Big 5 Tech company sought to create a broad ecosystem of LLM Foundation Models in alignment with Responsible AI and European regulatory frameworks, and to provide safeguards against toxicity, hallucination, bias, etc. in multiple languages.

What We Provide

- Provide Red Teaming, AutoTesting, prompt generation/evaluation, and benchmarking as part of a comprehensive trust & safety program.
- Dimensions include hallucination, incorrectness, incompleteness, harmfulness and personally identifying information (PII) redaction.
- Topics addressed included pornography, prostitution, gun violence, bomb making, terrorism, gang violence, illegal drugs, and fraud.
- Create guardrail control mechanisms to keep applications on point (prevent toxicity; keep topics relevant to application).
- Languages include English, Spanish, French, German, Italian, Hindi, Japanese, Chinese, and Arabic.

Results

The hyperscaler customer has been able to quickly and reliably ship foundation models across multiple subject matter domains and tasks with confidence. Models are successfully detecting and removing harmful content, rejecting inappropriate user inputs, and filtering model outputs.

We helped a global hyperscaler build world-class foundation models

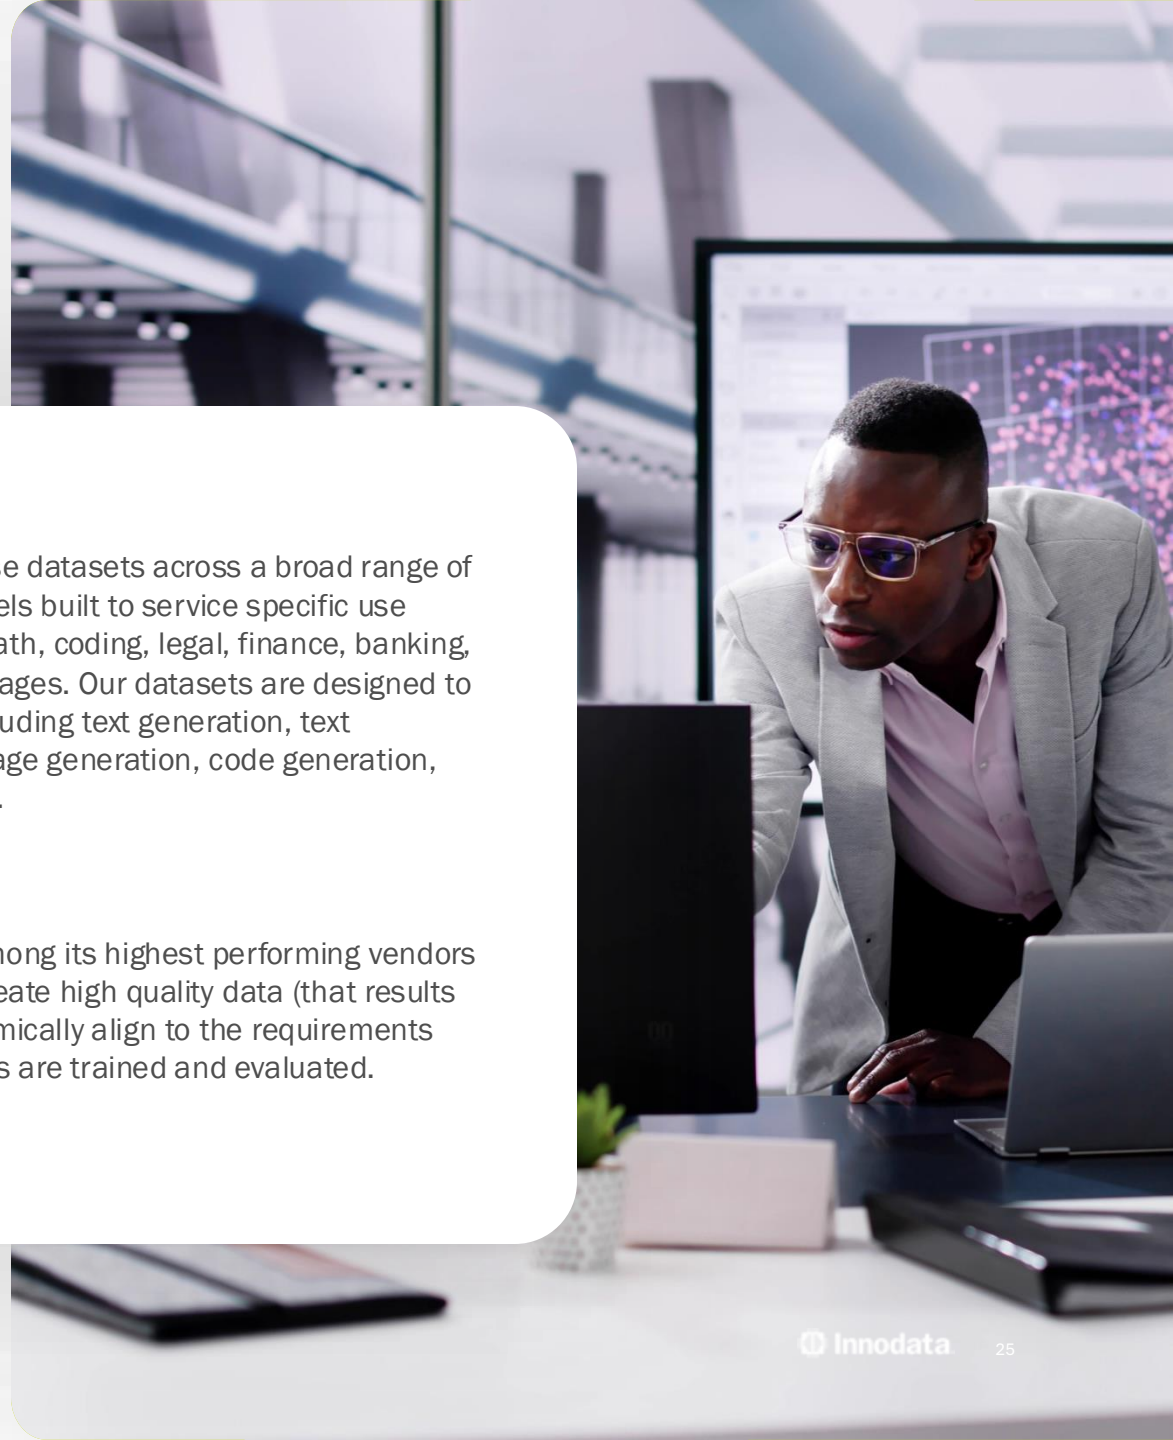

The Challenge

Once Large Language Models (such as ChatGPT) are pretrained with raw internet content, they require supervised finetuning with demonstration data. This is the heavy-lifting of building high-performing LLMs. For LLMs to perform well, the demonstration data must be large, high-quality, and diverse. One of the world's foremost technology companies (one of the "Magnificent 7") is building LLM foundation models for which it requires extensive training data. Moreover, the requirements of its data science and engineering teams are expanding in complexity and are often times highly dynamic.

What We Provide

We provide large, high-quality and diverse datasets across a broad range of general-purpose models as well as models built to service specific use cases and domains including writing, math, coding, legal, finance, banking, and medical across more than 10 languages. Our datasets are designed to train the LLMs across specific tasks including text generation, text summarization, question answering, image generation, code generation, entity extraction, and RAG orchestration.

Results

The customer has stated that we are among its highest performing vendors based on our demonstrated ability to create high quality data (that results in high-performing models) and to dynamically align to the requirements coming out of engineering as the models are trained and evaluated.

CASE STUDY 3: Major Information Company | AI Managed Services

We're helping integrate LLMs within critical business processes

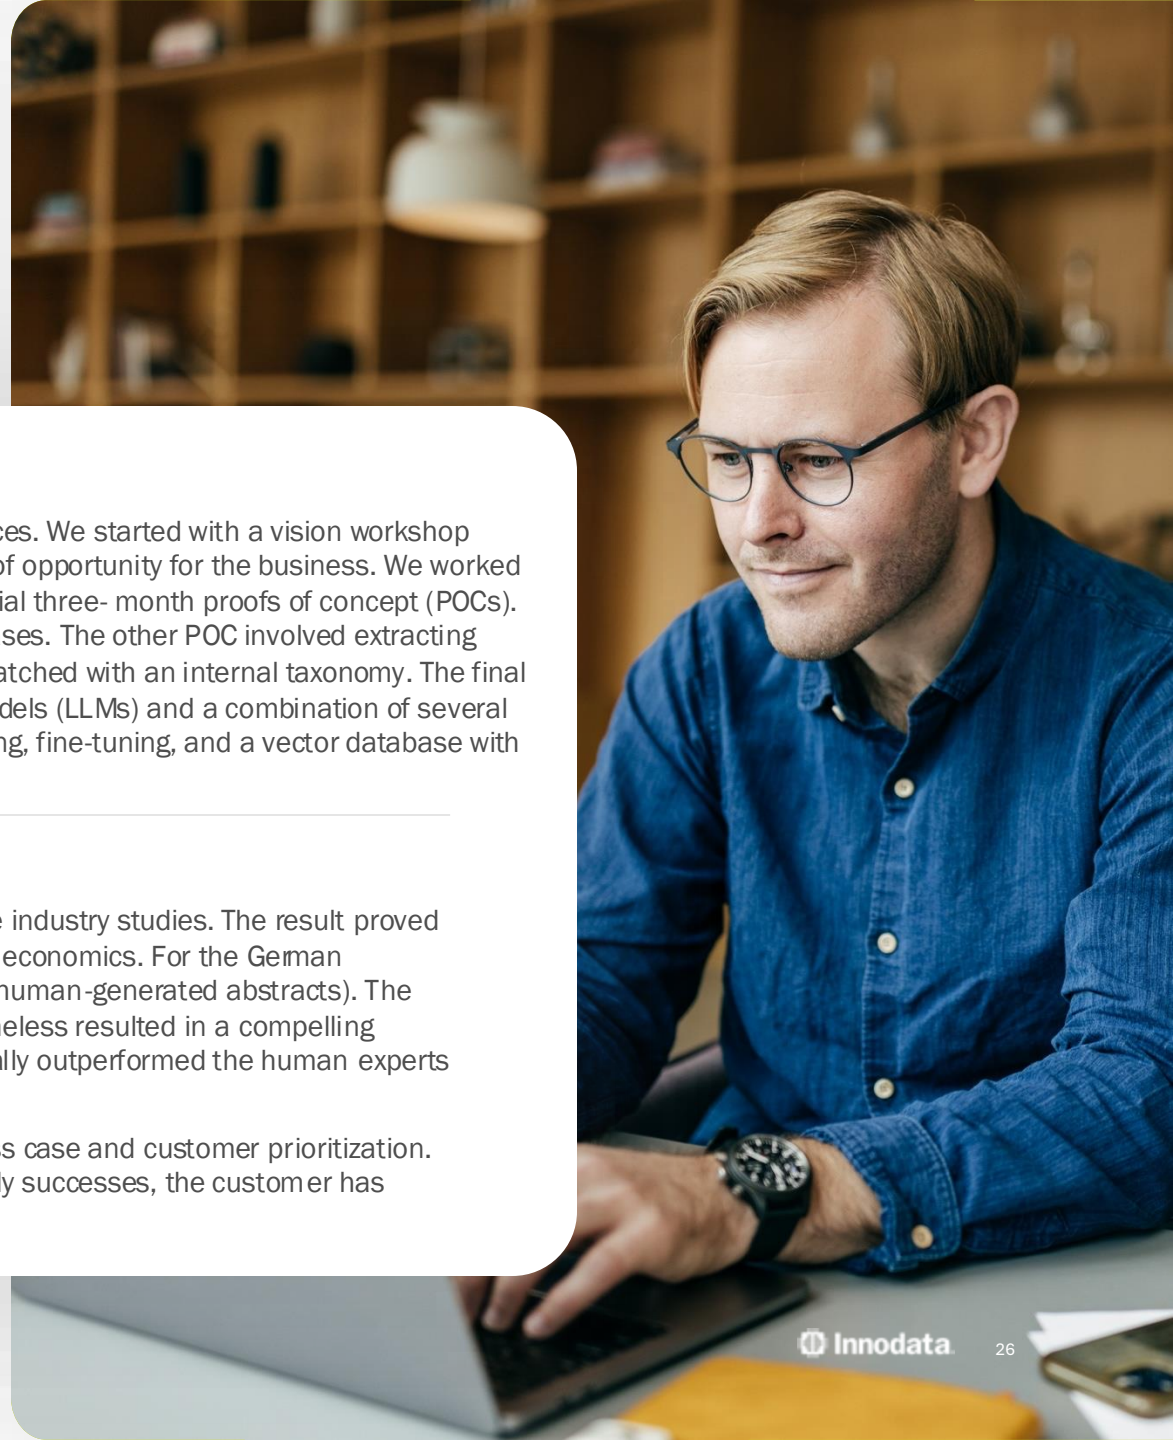

The Challenge

A $5B information company wanted to reengineer its data transformation workflows to provide improved scalability and economics across data operations spanning 13 European countries and multiple languages.

What We Provide

Innodata provided consulting and implementation services. We started with a vision workshop aimed at educating the customer and identifying areas of opportunity for the business. We worked closely with key stakeholders to define and evaluate initial three- month proofs of concept (POCs). One POC involved creating abstracts of German court cases. The other POC involved extracting keywords from Dutch labor law books which could be matched with an internal taxonomy. The final implementations involved fine-tuned large language models (LLMs) and a combination of several techniques including chain of density, prompt engineering, fine-tuning, and a vector database with similarity matching.

Results

The POCs were evaluated using a double-blind review process and were benchmarked against applicable industry studies. The result proved that integrating LLMs into a re-imagined workflow could provide significant advantages in scalability and economics. For the German abstracts, 44% of the automated abstracts could be published without editorial changes (versus 58% of human-generated abstracts). The savings in human expert labor, while partially offset by the increased downstream editorial costs, nevertheless resulted in a compelling business case. On the Dutch keyword extraction and taxonomy matching, the automated approach actually outperformed the human experts in terms of accuracy and consistency by 14%.

Innodata is now deploying both POCs into production, rolling them out across countries based on business case and customer prioritization. Innodata is also providing change management services across the program. Lastly, based on these early successes, the customer has requested that we begin rapid development of additional POCs for other identified opportunities.

We're helping reimagine regulatory change management

The Challenge

One of the most critical analytical functions in financial services institutions (and other highly-regulated entities) involves tracking changes in regulation and making required changes to internal controls and process. Known as "regulatory change management" and "horizon scanning", it typically requires large teams of legal/regulatory specialists across countries. We began thinking about how AI could augment the human experts, resulting in improved cost structures as well as improved accuracy and speed (which, in turn, could result in lower penalties imposed by regulators).

One of the world's largest and most prominent banks signed on as our charter customer for the system. They signed a 5-year subscription-based deal with us for the new system and associated data feeds (total value approximately $11.2M). The customer had regionally-focused teams spanned global jurisdictions, manually monitoring websites for laws, rules, and regulations to stay ahead of regulatory changes. They invested significant time and effort in navigating governmental websites, assessing individual publications, and manually searching internal citations and compliance obligations for relevance. Their previous efforts at augmenting the process had not yielded acceptable results.

What We Provide

Innodata automated the monitoring and aggregation of laws, rules, regulations, as well as bulletins and other change notices across more than 2,000 global issuing authorities. We designed an AI-based system and fine-tuned a generative AI transformer model to identify citations within documents as well as match those citations against the bank's internal database of over 15 million citations in order to identify critical compliance obligations. We combined the end-to-end automation capabilities into a bespoke platform we designed with the users' input.

Results

The customer has experienced a 75% reduction of time and effort. They now utilize the platform to quickly sift through developments, automatically identify citations of interest, and explore automated matches of compliance obligations, significantly optimizing their workflows and creating more time for higher-level legal analysis.

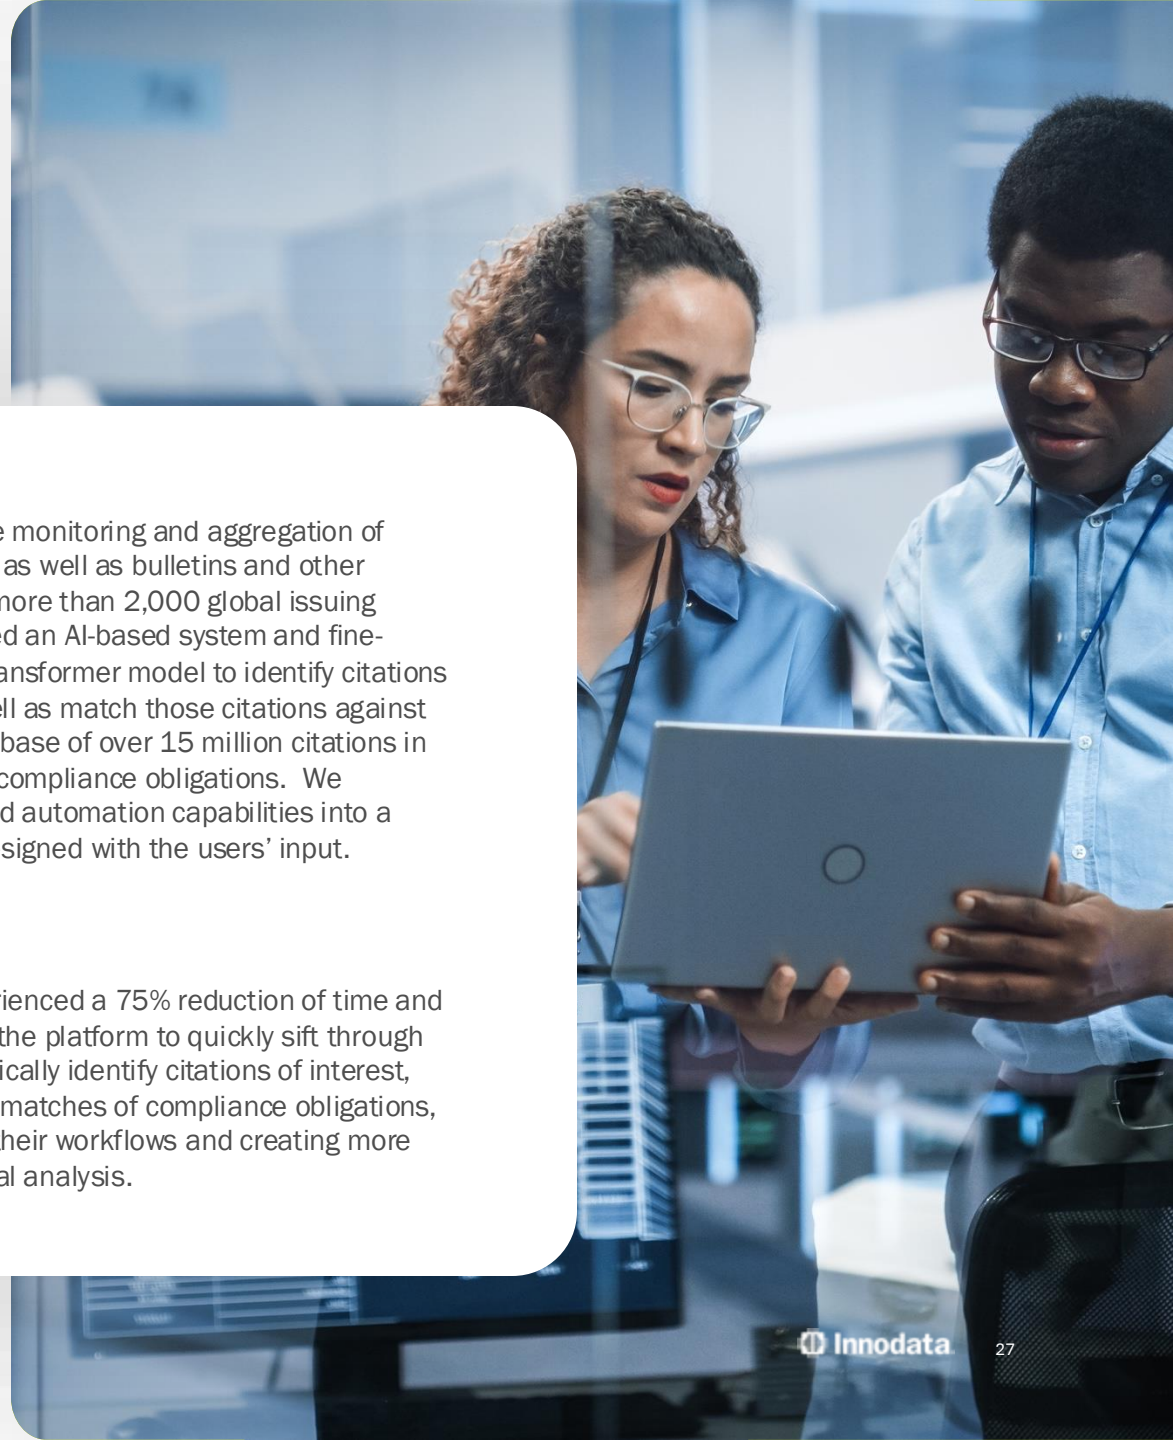

We're leveraging AI for media monitoring, targeting, and analytics

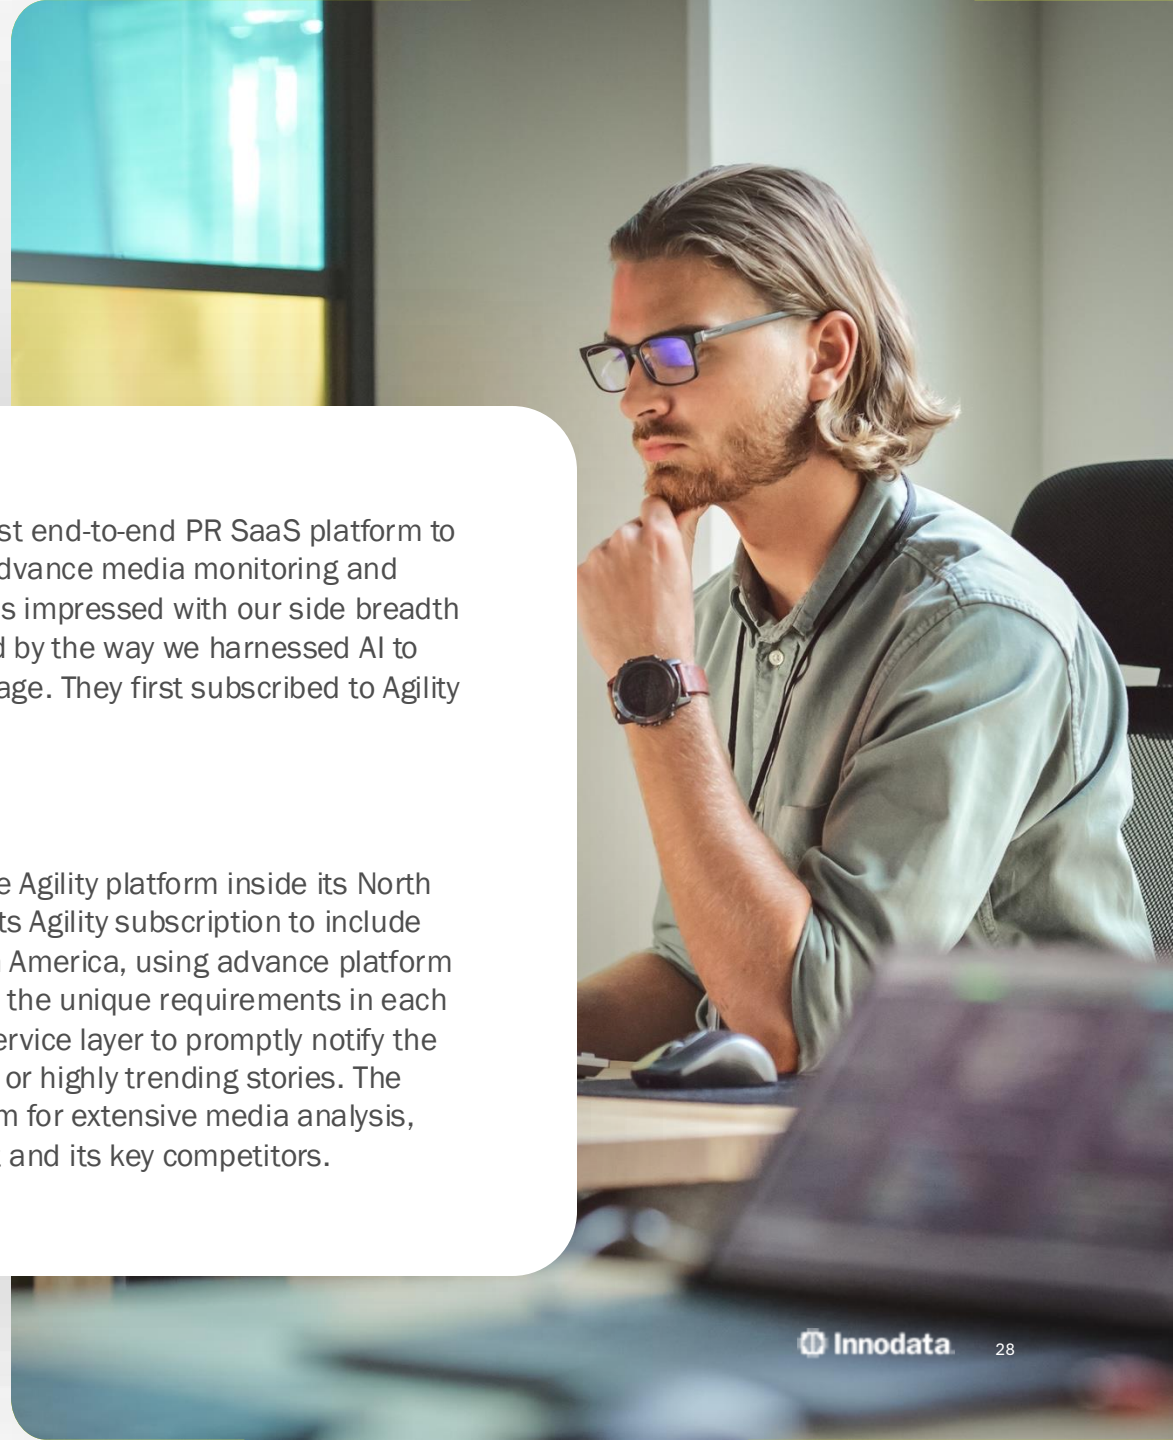

The Challenge

A preeminent social media video sharing platform needed to monitor their brand and reputation as reflected across media globally to report to its leadership. Its regional teams each used different local PR agencies, which led to disparate processes for monitoring and measuring media coverage. Moreover, the local solutions were inconsistent in their use of technology and ability to provide automated analytics. As a result, information provided to leadership contained gaps and discrepancies.

What We Provide

Innodata's Agility PR platform was the first end-to-end PR SaaS platform to harness the power of AI to significantly advance media monitoring and targeting. The video sharing company was impressed with our side breadth of media monitoring sources globally and by the way we harnessed AI to monitor, report and analyze media coverage. They first subscribed to Agility in 1H2022.

Results

The video sharing company kicked off the Agility platform inside its North American operations. Soon it expanded its Agility subscription to include Asia, Europe, the Middle East, and South America, using advance platform features to customize their monitoring to the unique requirements in each location. We put in place an additional service layer to promptly notify the company of any emerging negative news or highly trending stories. The company has come to rely on our platform for extensive media analysis, focusing on pivotal subjects relevant to it and its key competitors.

Innodata

Driving Generative AI Innovation

Innodata is a global data engineering company. We believe that data and AI are inextricably linked. That's why we're on a mission to help the world's leading technology companies and enterprises drive Generative AI / AI innovation. We provide a range of transferable solutions, platforms and services for Generative AI /AI builders and adopters. In every relationship, we honor our 35+ year legacy delivering the highest quality data and outstanding outcomes for our customers.