HEXAWARE

CLIENT IMPACT SERIES

Artificial Intelligence Led, Human Intelligence Perfected.

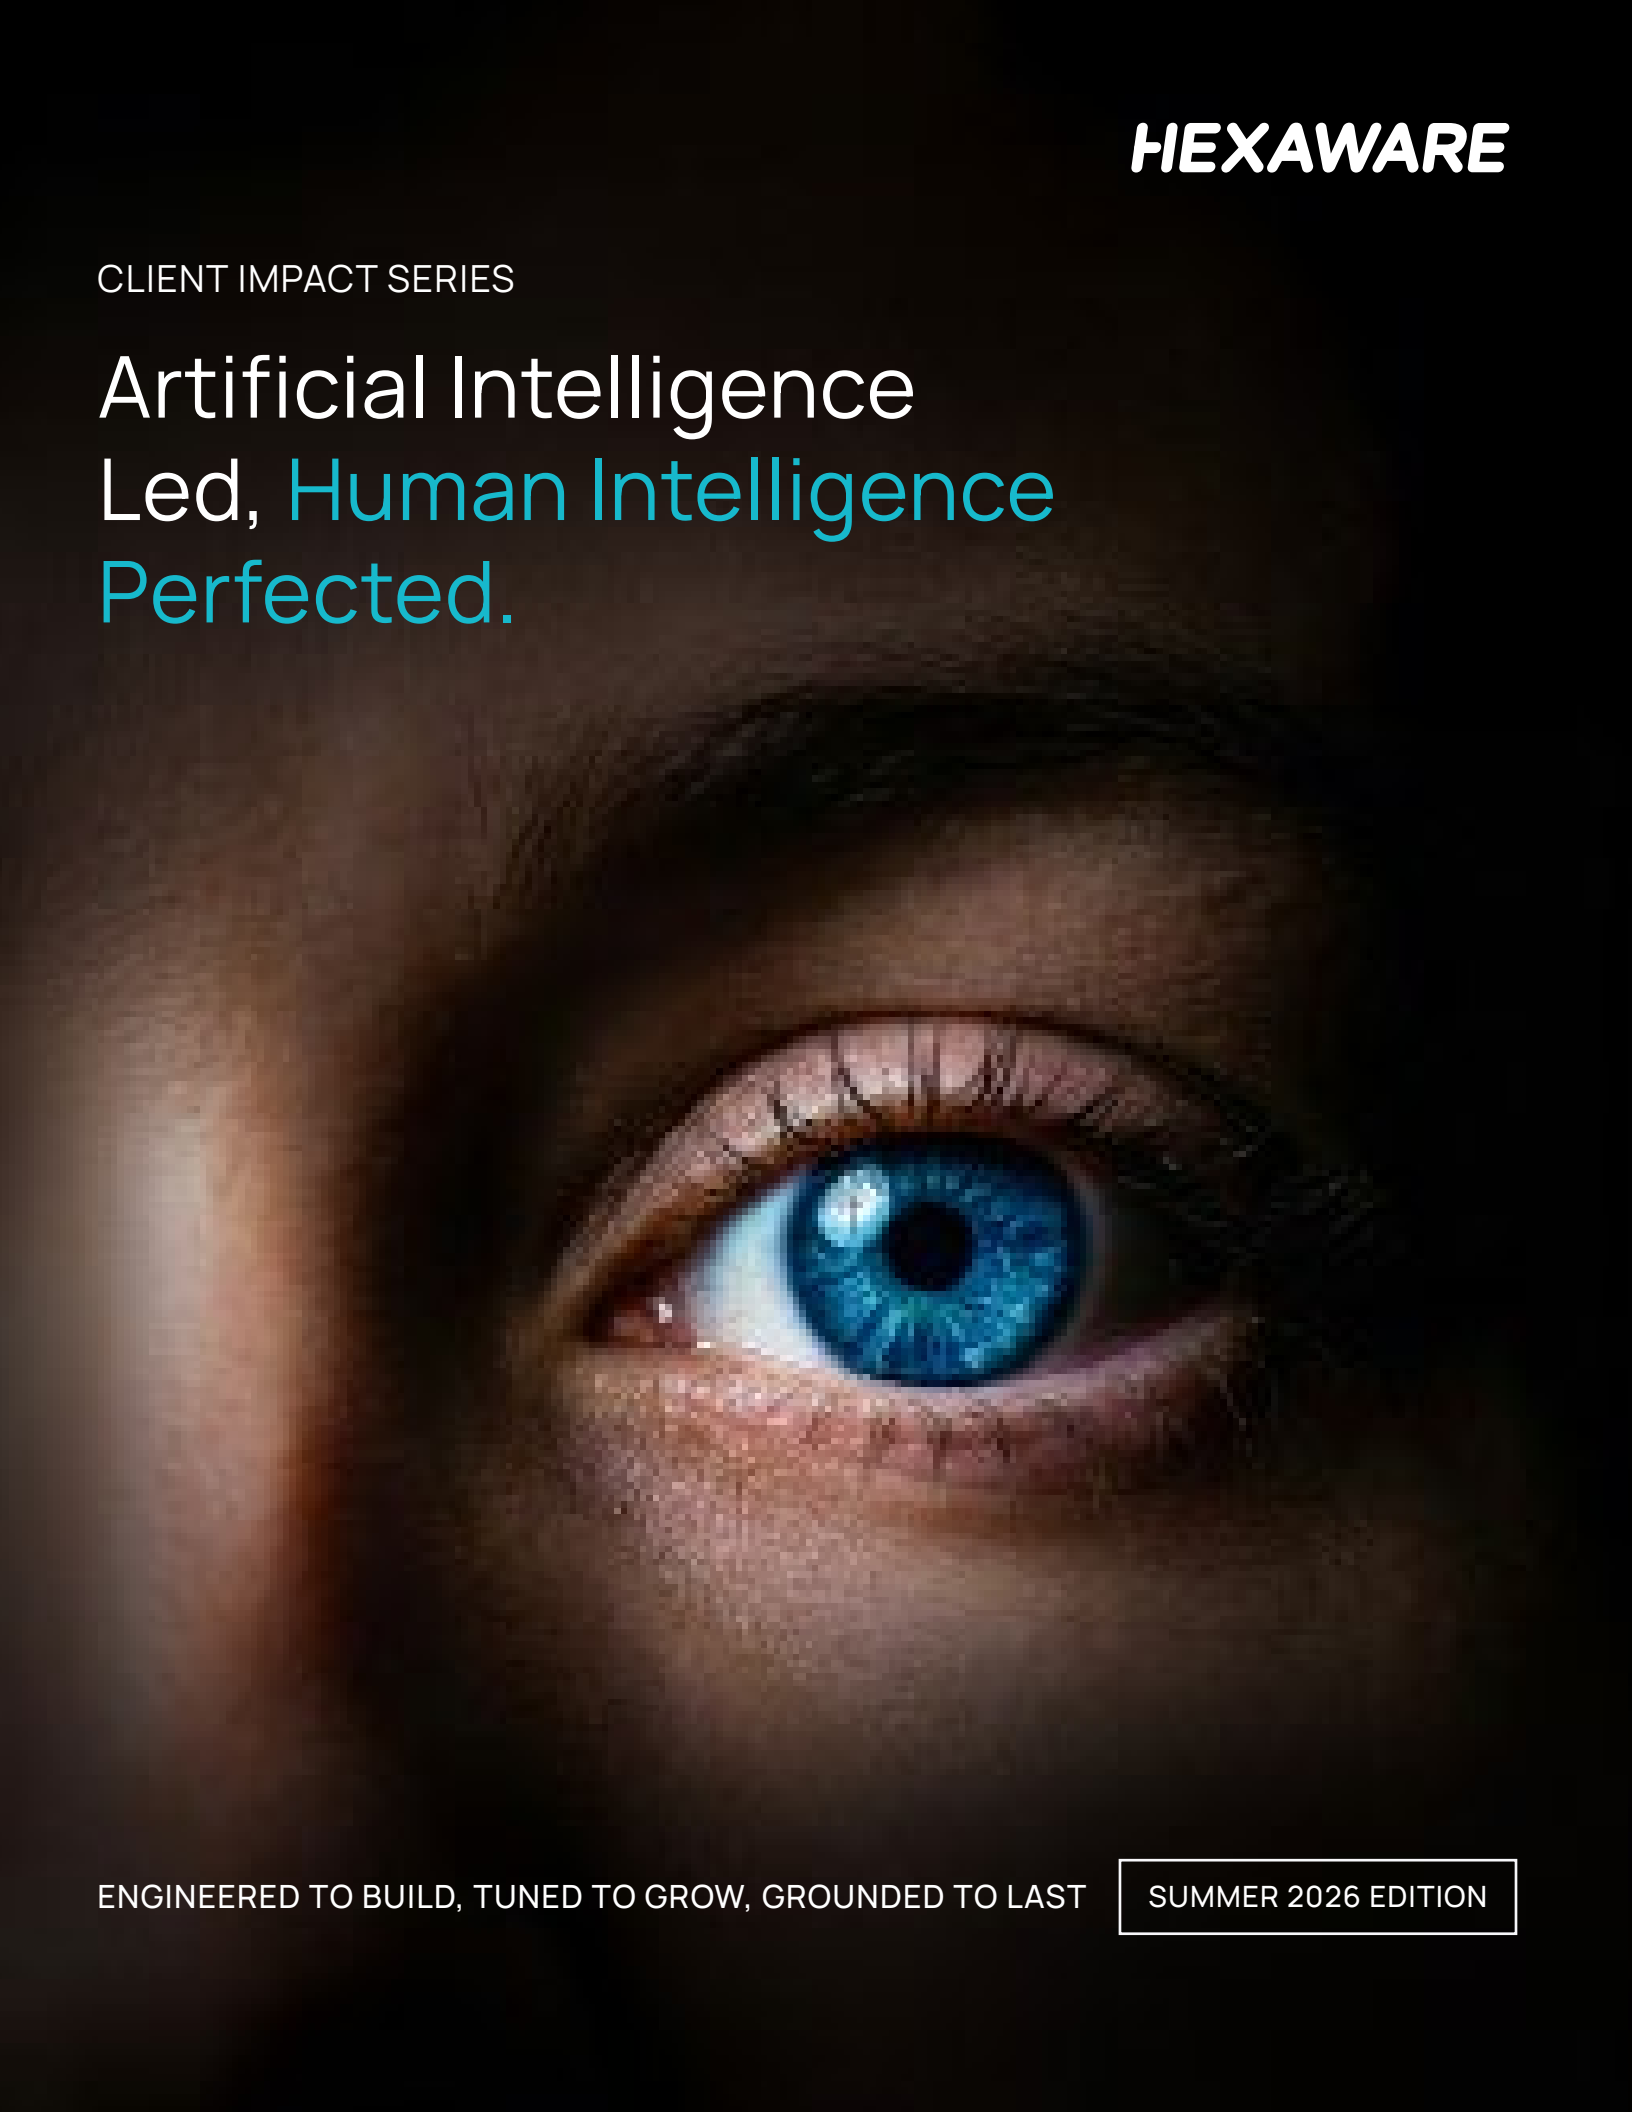

ENGINEERED TO BUILD, TUNED TO GROW, GROUNDED TO LAST

SUMMER 2026 EDITION

R. Srikrishna (Keech)

CEO & Executive Director,
Hexaware Technologies

Dear All,

Three months ago, we published the first volume of this series—a collection of real customer stories showing what AI actually does when it is put to work inside enterprise technology and business operations. We are back with Volume II, and the story it tells is a more urgent one.

AI is no longer a conversation about potential.

It is real, it is deployed, and it is being used today by organizations across every industry we serve. What began as pilots and proofs of concept has become production reality.

The velocity of adoption is exponential—and this edition proves it.

Volume I opened with case studies spanning SDLC transformation, IT outsourcing, data modernization, and our Zero License model. In the three months since, the number of material and verifiable customer case studies has more than doubled. This is not a curated handful of flagship stories—it is a body of evidence that now spans every vertical we operate in, from financial services to healthcare, from manufacturing to travel, from technology to professional services. The pace at which our customers are moving from experimentation to measurable impact is accelerating faster than any technology shift we have tracked before.

This series exists to inspire what comes next.

Every story in these pages is a demonstration of the art of the possible—proof that AI can be applied meaningfully to nearly every dimension of technology and business, in every industry, when it is grounded in real engineering discipline and real customer outcomes rather than hype. Our hope is that as you read Volume II, you see not just what our customers have achieved, but a preview of what is now within reach for your own organization.

Thank you for being part of this journey with us—as customers who trust us to deliver, and as investors who believe in where this is headed. We look forward to Volume III and to the stories you help us write between now and then.

Hexaware AI: Built to Scale Responsibly

Move AI from Pilot to Production in Record Time

AI pilots often stall in proof-of-concept: fragmented experiments, ungoverned costs, compliance risk, and the inefficiency of rebuilding agents for every new use case. Agentverse™ is Hexaware's next-generation enterprise AI agent platform, designed to unify the entire AI agent lifecycle on a single, governed foundation. It delivers that through three layers:

Agentverse™ Core
A secure, scalable AI landing zone for hybrid and multi-cloud environments (Azure, AWS, NVIDIA, GCP), deployed in 4 weeks with centralized audit trails, FinOps guardrails, and policy-aware connectors.

Agentverse™ Studios
Developer, Testing, and Ops Studios to build, validate, and operate agents. Drag-and-drop creation with no SDK expertise, AI-grade QA with RAI validation and drift monitoring, and unified deployment with real-time observability.

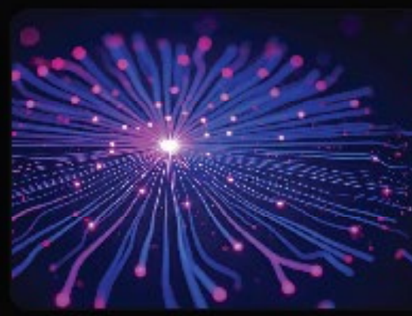

Agentverse™ Agents
600+ ready-to-deploy agents spanning six categories (Technology, Domain, Functional, System, Process, and Role), reusable across business units to reduce development time.

Together, these three layers move organizations from AI experimentation to enterprise execution, with governance integrated as a horizontal control plane across RBAC, audit trails, RAI checks, and FinOps.

This Client Impact Series shows what that looks like in practice, through a curated set of use cases designed to help you identify where Agentverse™ can create the greatest value for your organization.

600+

Ready-to-deploy AI agents built for real business and technology workflows.

Success Stories

Driving impact for our clients – the world can't miss!

Table of Contents

Banking

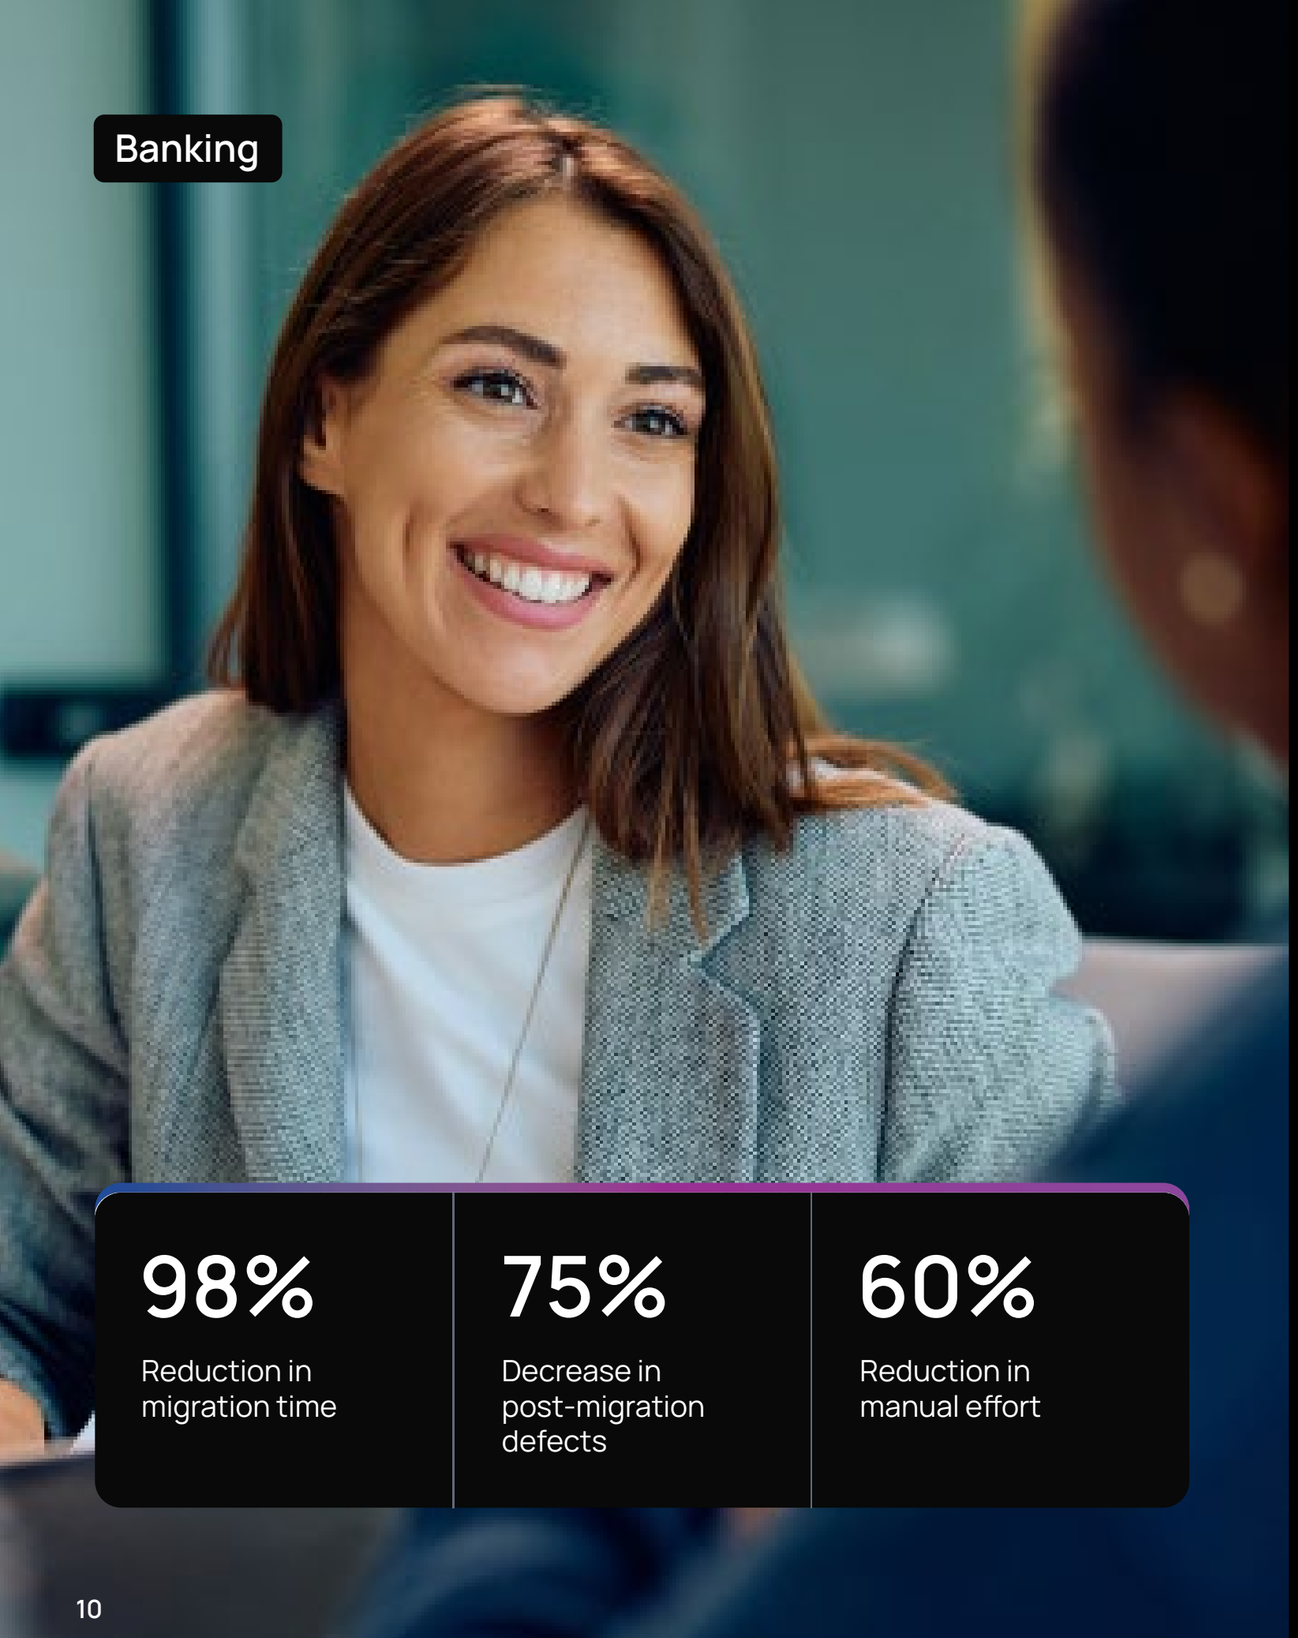

98%	75%	60%
Reduction in migration time	Decrease in post-migration defects	Reduction in manual effort

Client

A North American Retail and Commercial Bank

Modernizing legacy Customer Information Control Systems (CICS) without the usual disruption

The client needed to modernize legacy CICS and aging .NET systems while protecting quality, compliance, and business continuity. The priority was to move faster without increasing risk in a tightly regulated environment.

Solution

- Implemented a multi-agent software development lifecycle model with planner, executor, and reviewer agents.
- Automated COBOL copybook parsing, business logic extraction, and field mapping generation.
- Generated Angular user interface and .NET APIs through AI-driven code automation.
- Embedded human-in-the-loop checkpoints, validation gates, and domain-specific context engineering.

Benefits

- Reduced migration time by 98%, dramatically accelerating modernization.
- Cut post-migration defects by 75%.
- Reduced manual effort by 60%, freeing teams to focus on higher-value decisions.
- Enabled a more incremental, lower-risk migration path for compliance-sensitive operations.

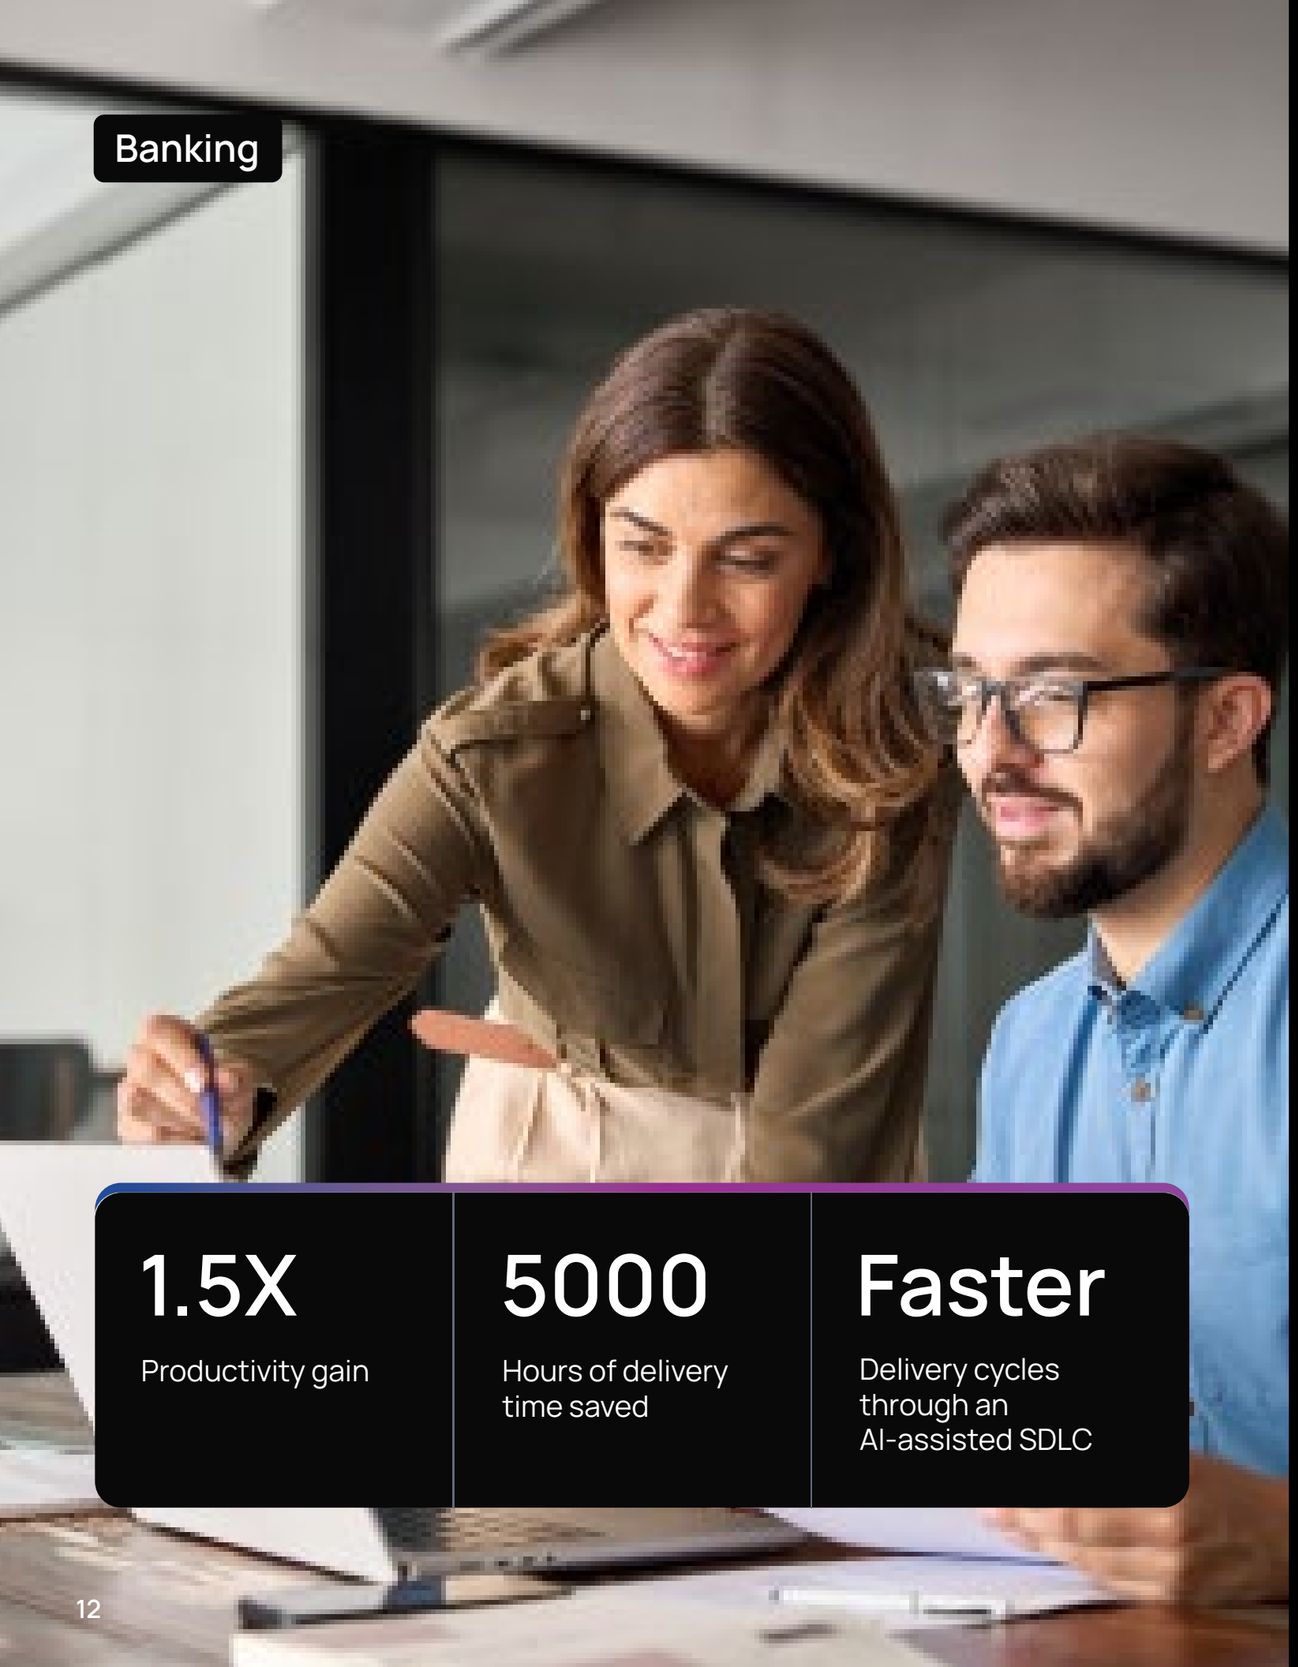

Client

An American Bank Offering Personal, Business, and Commercial Banking Services.

Giving delivery teams AI support across the entire software development lifecycle (SDLC)

The client wanted to remove friction across requirements, architecture, development, and QA so teams could move with more confidence and less rework. The result was a more connected SDLC with AI support where teams needed it most.

Solution

- Implemented AI-driven support across requirements analysis, architecture, coding, and quality assurance (QA).
- Used BA agents to turn raw requirements into structured delivery artifacts.
- Automated test case creation, QA scripts, and regression execution.
- Deployed six AI agents to support delivery teams end-to-end.

Benefits

- Delivered a 1.5x productivity gain across the SDLC.
- Saved up to 5,000 hours of development time.
- Improved code quality and reduced rework through AI-assisted validation and reviews.
- Accelerated delivery cycles with stronger test coverage and faster pull-request turnaround.

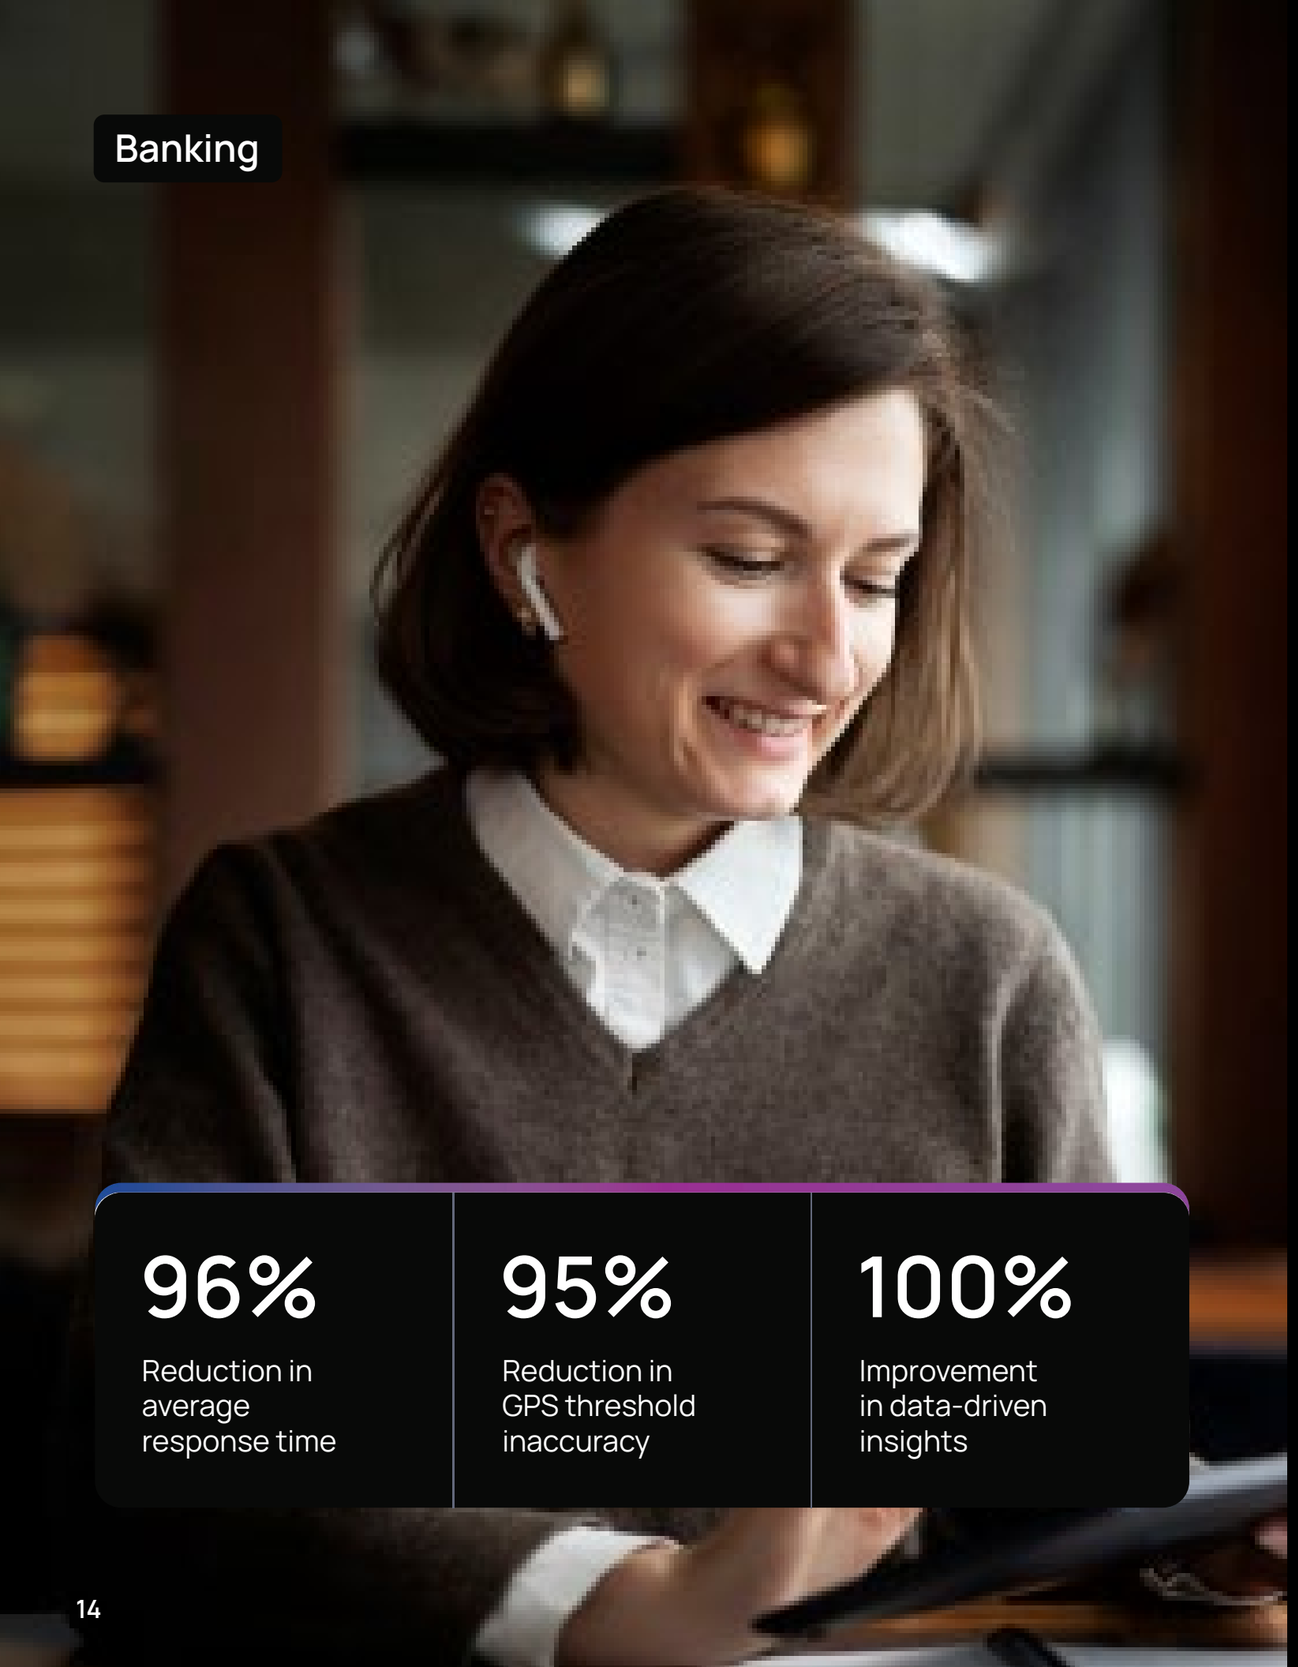

Client

A European Multinational Bank

Helping farmers protect herds—even where connectivity is limited

With a strong agriculture customer base, the bank wanted to support farmers managing large cattle operations in low connectivity regions. The opportunity was to make herd monitoring more proactive, practical, and data driven so issues could be spotted before they became costly problems.

Solution

- Deployed a GPS- and AI-enabled livestock monitoring solution built for large-scale farms, including low-connectivity environments.
- Created clear visualizations so farmers could manage thousands of bovines with far less manual tracking.
- Used AI to detect abnormal behavior and early health signals, including sickness and calving difficulties.
- Enabled visibility into location, grazing patterns, behavior, water availability, and other operational factors.

Benefits

- Reduced average response time by 96%, helping farmers act faster when animals needed attention.
- Cut GPS threshold inaccuracy by 95%, improving confidence in location-based monitoring.
- Drove a 100% improvement in data-driven insights for day-to-day herd management.
- Made large-scale cattle operations easier to oversee in regions where traditional monitoring is harder to sustain.

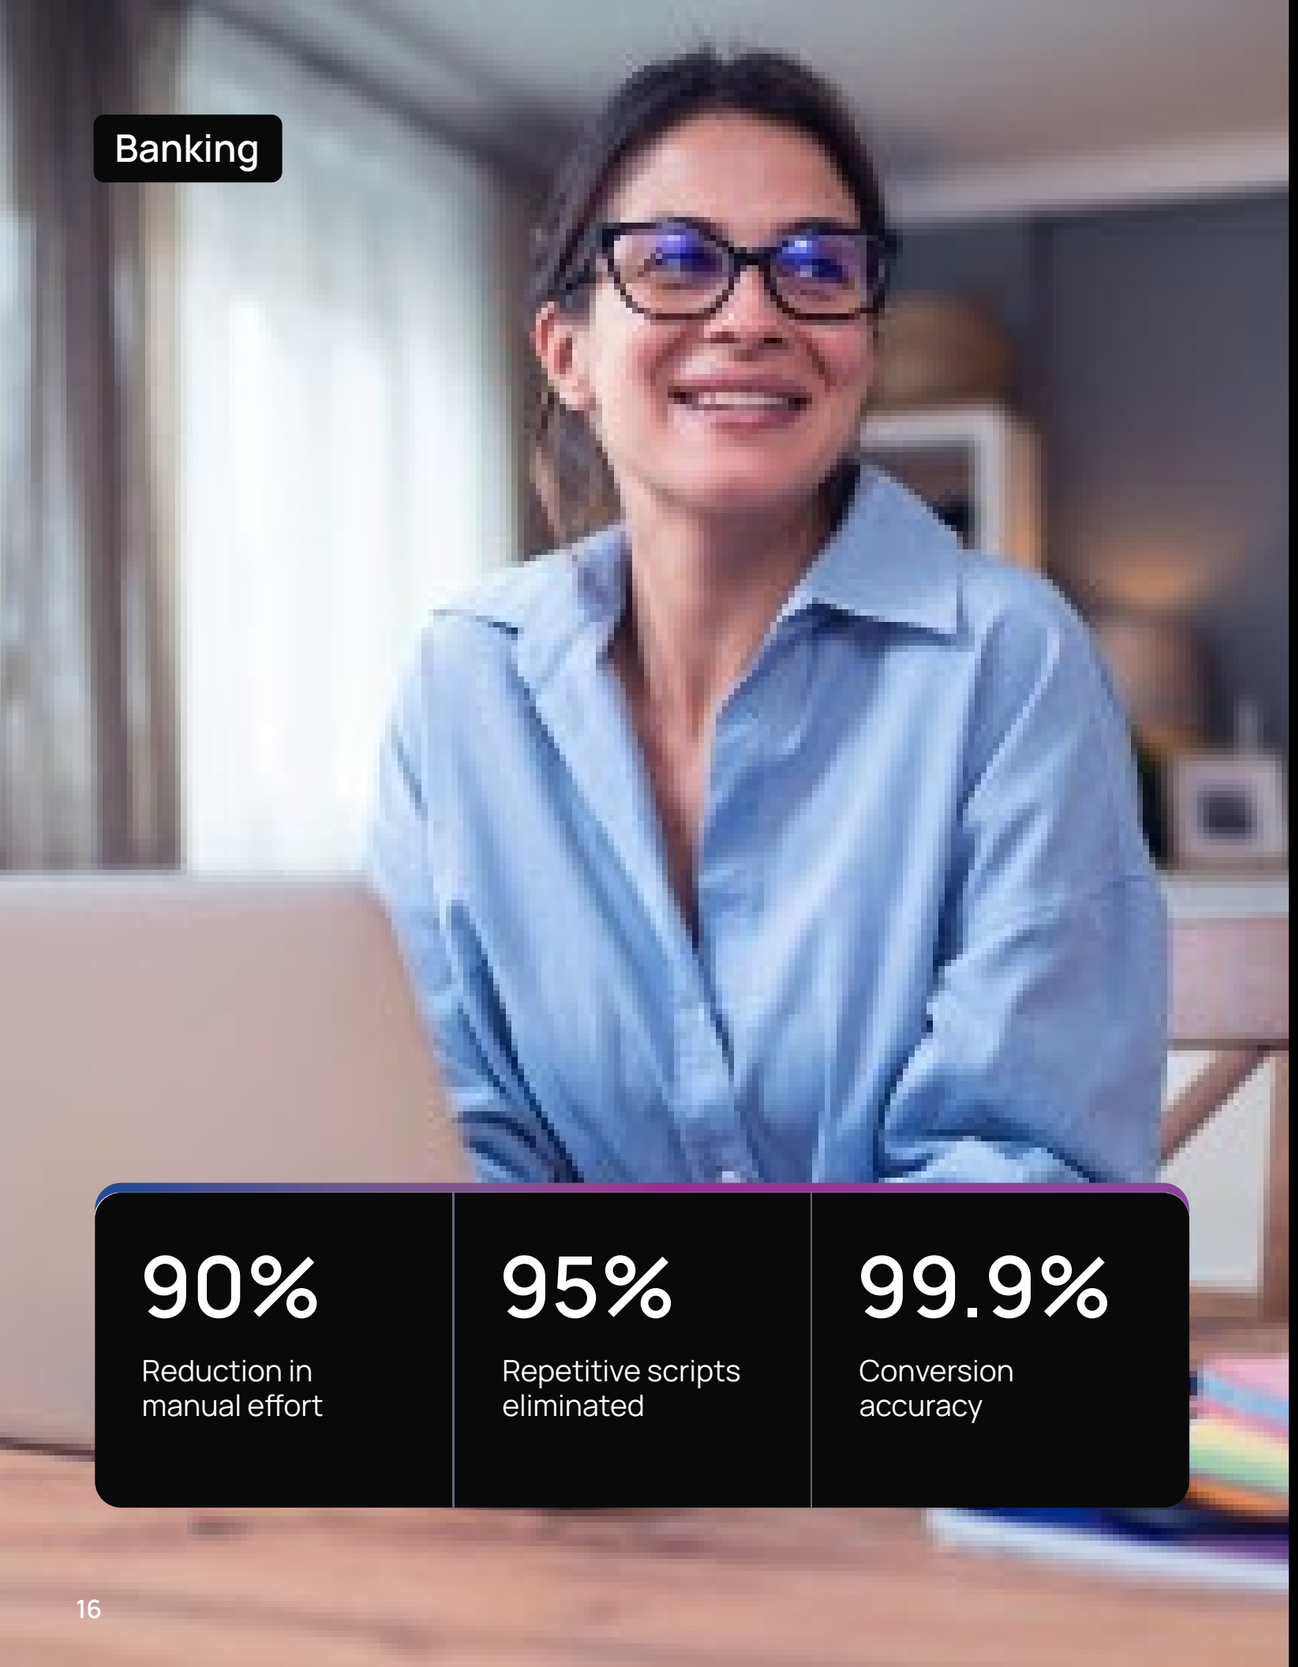

90%	95%	99.9%
Reduction in manual effort	Repetitive scripts eliminated	Conversion accuracy

Client

One of the Oldest and Largest Banks in Australia

AI-led SAS modernization: 25,000+ scripts converted to open-source PySpark with a 90% cut in manual effort

A large share of the bank's analytics ran on SAS, a licensed proprietary platform that had become a strategic liability. Licensing and operational costs climbed with data volume rather than flexing with it, the platform scaled poorly, and proprietary lock-in blocked cloud migration and modern AI/ML adoption. Compounding this, retiring SAS skills were hard to replace as new engineers favored Python and Spark.

Solution

- Delivered the Zero License Factory for Data and AI, replacing the licensed SAS platform with an AI-led modernization service on open technologies.
- Ran deterministic reverse engineering with pre-built Python parsers to fully document the existing SAS estate, removing about 80% of uncertainty before conversion.
- Generated target-state blueprints, architecture documents, and execution playbooks autonomously through agentic blueprinting, eliminating weeks of manual analysis.
- Rebuilt the estate into deployment-ready PySpark notebooks, with ModelContext Protocol binding AI agents to the bank's security and compliance guardrails.

Benefits

- Eliminated recurring SAS license and renewal costs through an open-source ecosystem.
- Handed the bank's analytics logic back as code it owns outright, free of proprietary lock-in.
- Delivered scalable processing and faster insights on a cloud-native foundation.
- Reduced dependency on scarce, retiring SAS skills by moving to Python and PySpark.

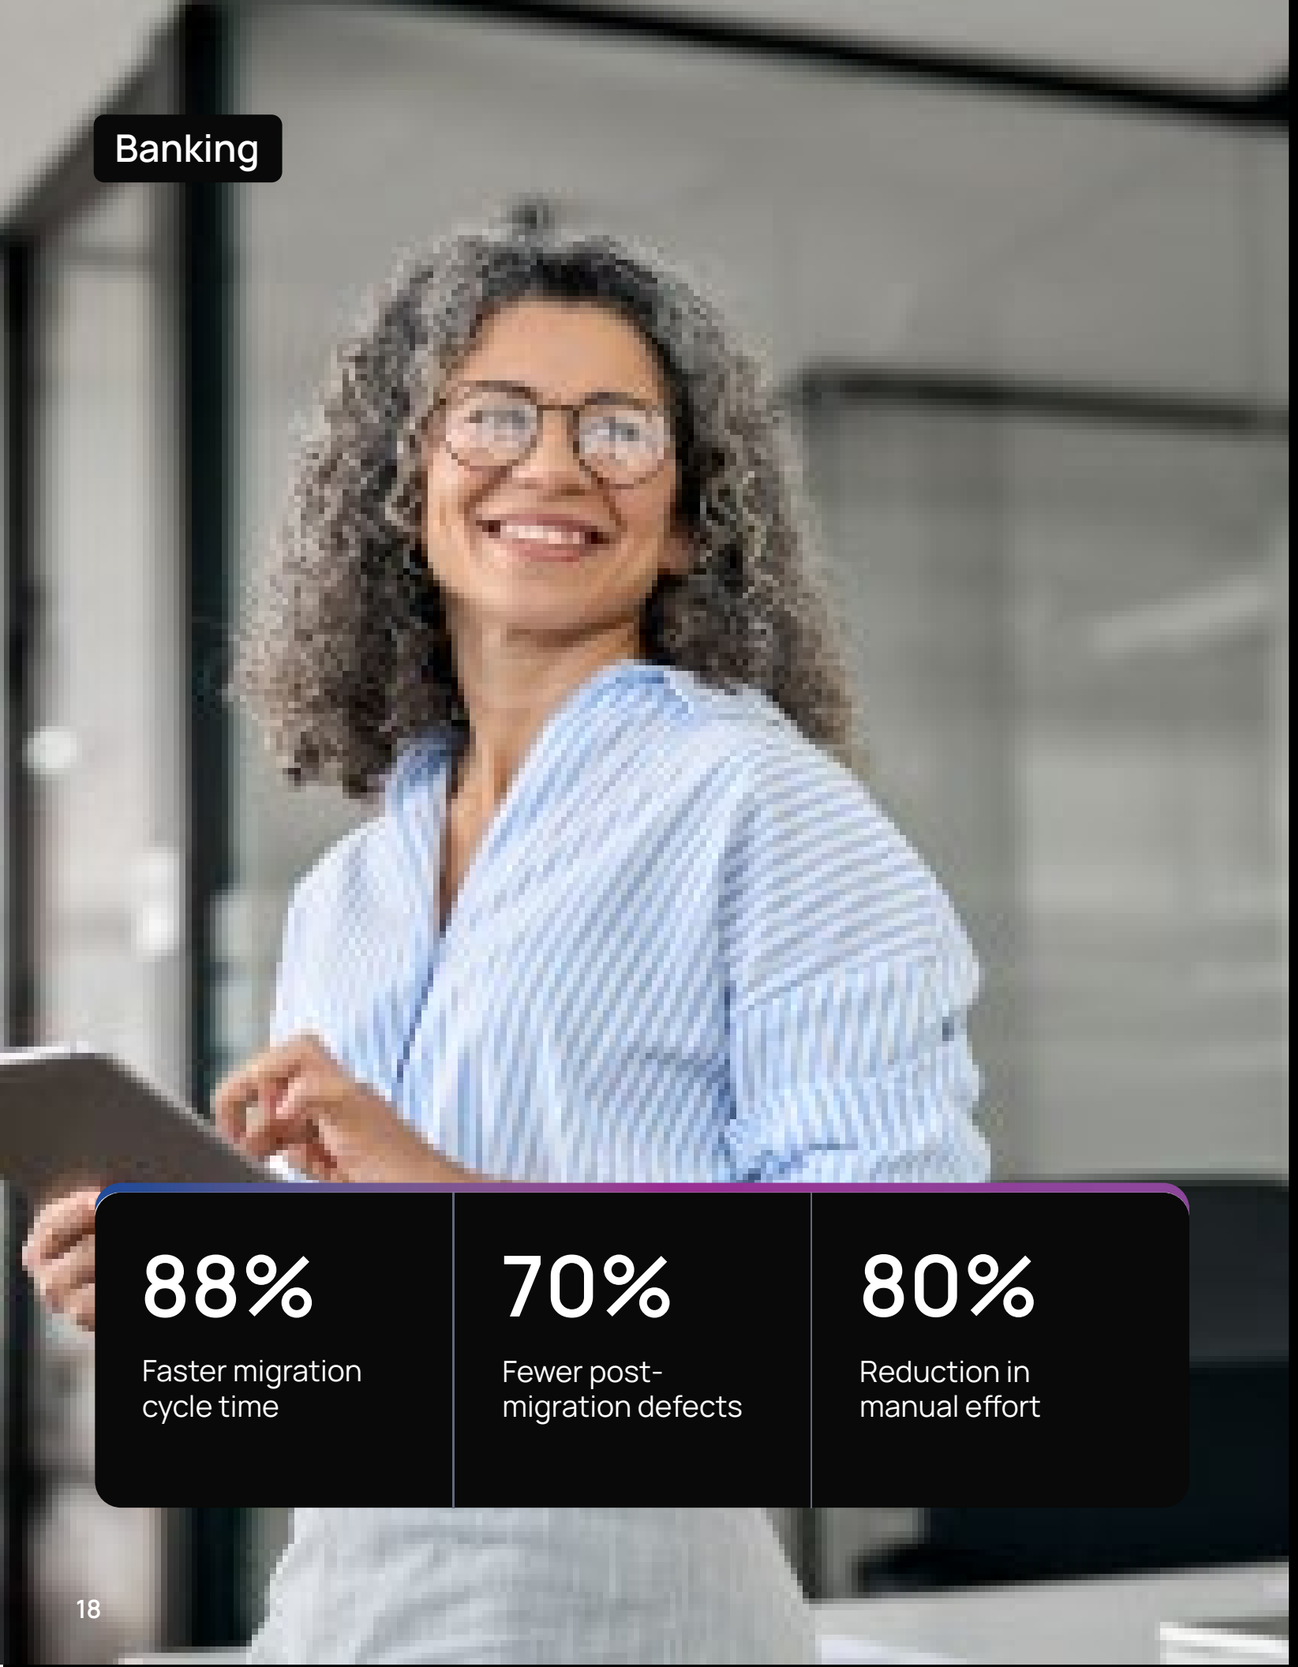

Client

A Major US Financial Institution

Modernizing teller APIs without the middleware license burden

The Teller API layer ran on aging IBM IIB and DataPower middleware with .NET-dependent integrations, creating delivery drag, reliability risk, and high run cost. Manual modernization was too slow and too risky for a regulated banking environment, with teams spending days per migration unit on scarce legacy knowledge. The client needed an API-first, cloud-native model to ship faster without trading off compliance.

Solution

- Delivered an AI-engineered Zero License factory that replaced IBM IIB and DataPower with cloud-native, API-first components the client owns outright.
- Implemented a three-agent planner-executor-verifier pattern with human-in-the-loop checkpoints at critical decision gates.
- Used a deterministic forensics and blueprinting platform to extract legacy logic into a verified Knowledge Index with no LLM hallucination risk.
- Established a governed plan-build-review-ship workflow with traceable approvals, audit logs, and policy-driven quality gates.

Benefits

- Cut migration cycle time by 88%, from 4 days to 4 hours per comparable unit.
- Reduced post-migration defects by 70% through verifier-gated automated validation.
- Lowered manual effort by 80% via AI-assisted planning, generation, and documentation.
- Removed licensed IBM IIB and DataPower dependency, giving the client an integration layer it controls.

Banking

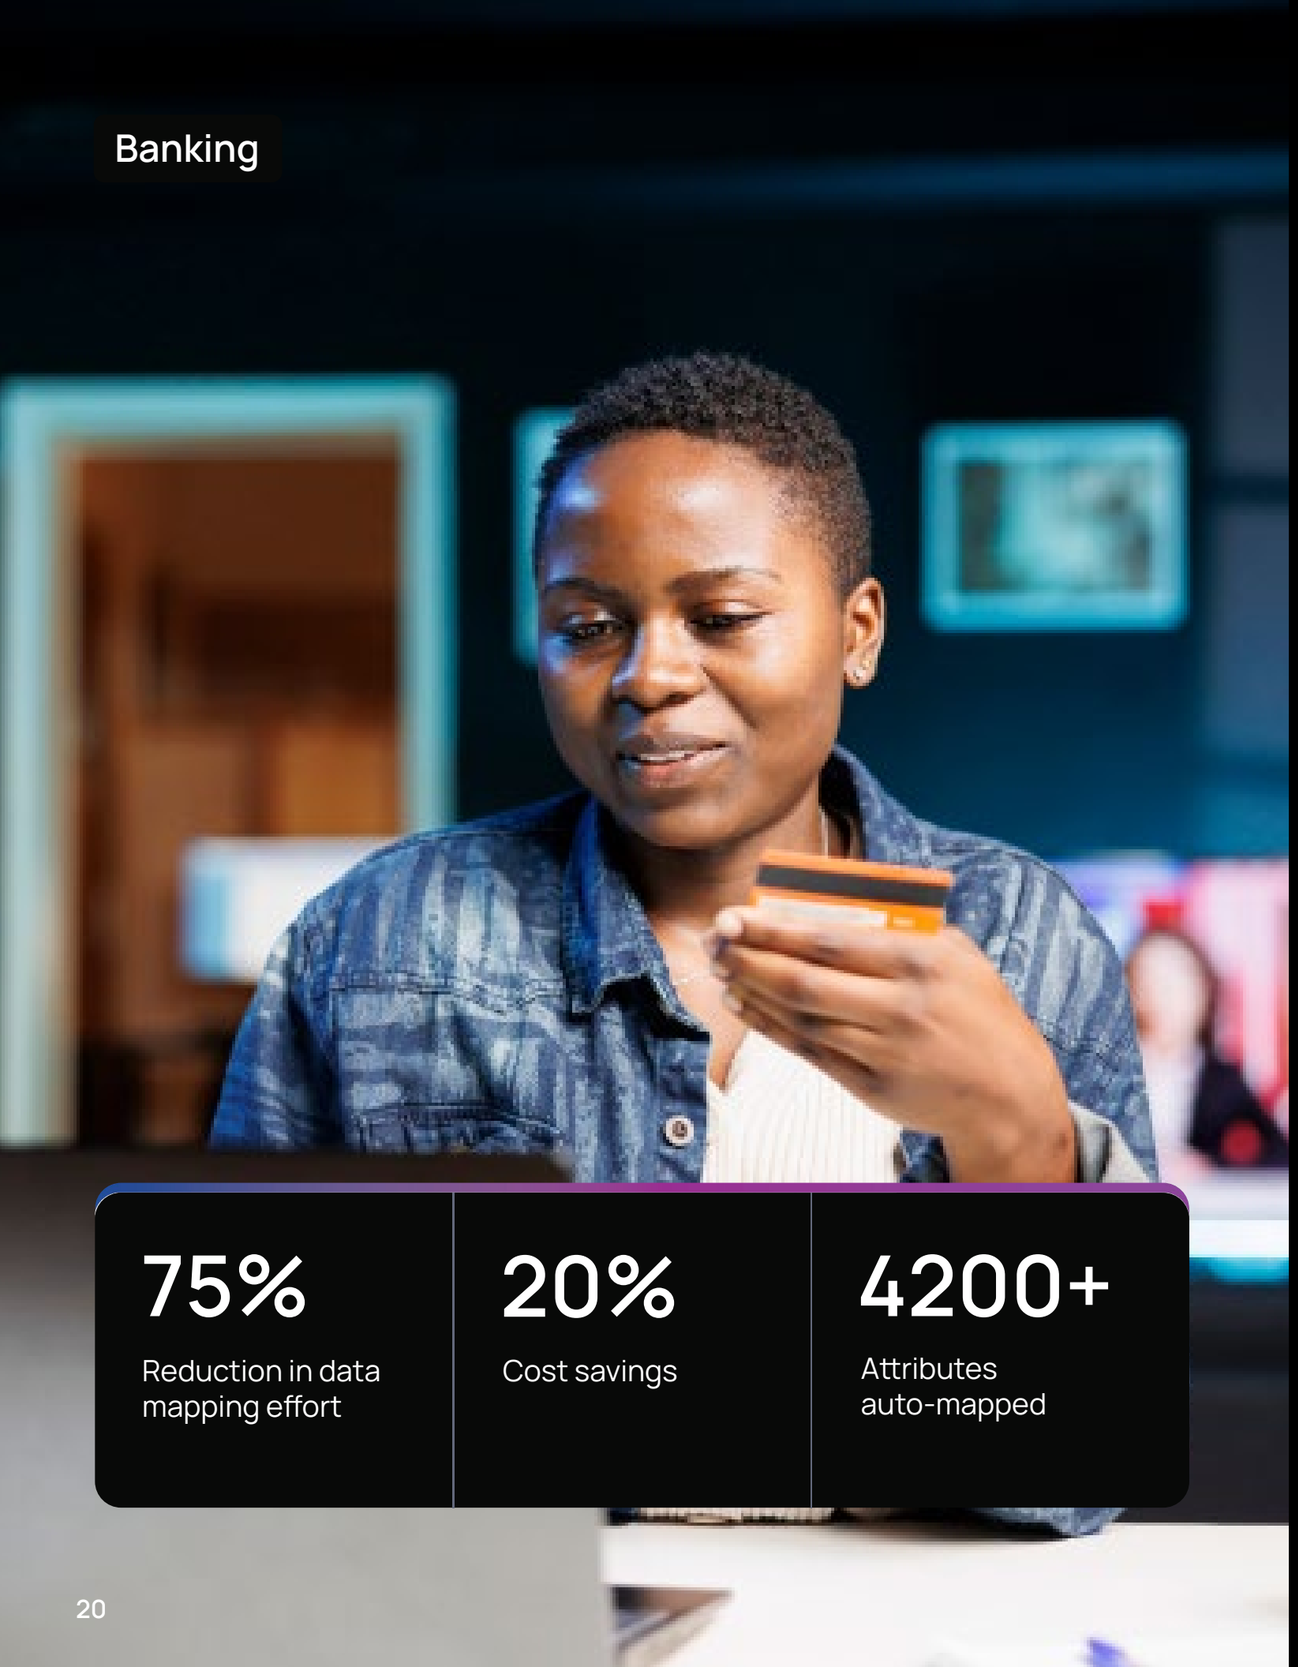

75%	20%	4200+
Reduction in data mapping effort	Cost savings	Attributes auto-mapped

Client

A North-American Multinational Banking and Financial Services Corporation

Taking the risk and rework out of card-platform migration

The client needed to move quickly on a card processor migration without letting thousands of complex attribute mappings slow progress or introduce errors. The goal was a faster, safer path to modernization with less manual effort at every step.

Solution

- Built an AI-powered mapping utility to automate attribute mapping between TSYS and Fiserv systems.
- Supported CSV, JSON, and XML ingestion to handle both structured and semi structured inputs.
- Applied data cleaning, masking, preprocessing, and post-processing to improve transformation quality and security.
- Used semantic similarity scoring to generate ETL-ready, column-level mappings with confidence scores.

Benefits

- Reduced data mapping effort by 75%, easing pressure on migration teams.
- Delivered 20% cost savings across the migration effort.
- Automatically generated mappings for 4,200+ attributes, accelerating delivery at scale.
- Lowered the risk of data loss, poor quality, and operational disruption during migration.

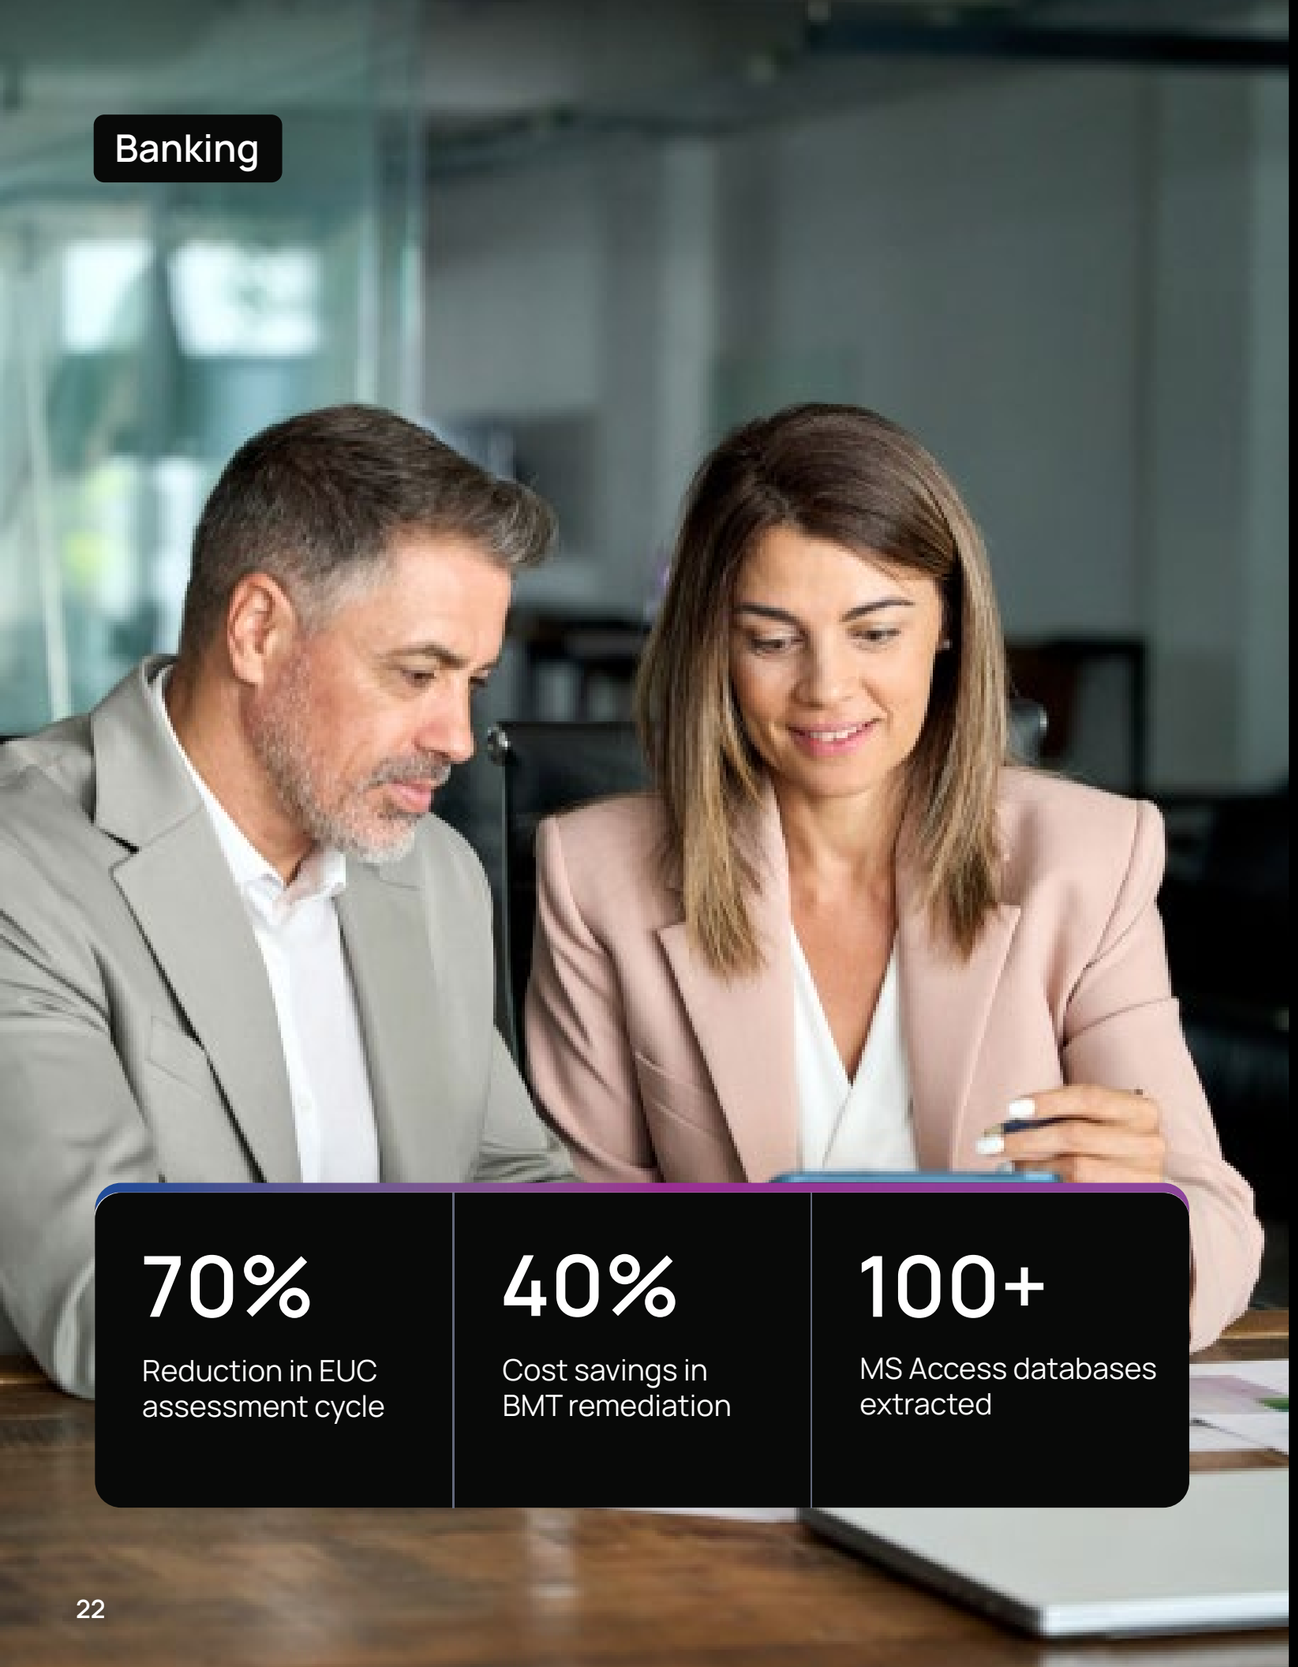

Client

A North-American Multinational Banking and Financial Services Corporation

Bringing decades of end user computing (EUC) sprawl under control

Years of uncontrolled EUC growth had left the bank with fragmented tools, limited documentation, and rising governance risk. It needed a practical way to assess, rationalize, and modernize complex assets without slowing critical business operations.

Solution

- Applied a phased modernization approach covering accumulation, remediation, acceleration, and assimilation.
- Automated EUC discovery and deconstruction to speed identification of high-risk assets.
- Built a tailored accelerator for MS Access and VBA-heavy EUCs.
- Used GenAI and adaptive large language models (LLMs) for classification, documentation, linkage mapping, and business logic extraction.

Benefits

- Reduced the EUC assessment cycle by 70%, shrinking timelines from six months to eight weeks.
- Cut business managed technology (BMT) remediation costs by 40%.
- Extracted and analyzed 100+ MS Access databases to support modernization planning.
- Gave the bank a clearer roadmap to reduce operational, security, and regulatory risk.

Client

One of Mexico's Four Largest Banking and Financial Services Companies

Turning regulatory packs from a transformation bottleneck into a faster, reusable workflow

Manual compliance processes were slowing every cloud migration and consuming thousands of SME hours. The client needed a way to move faster, stay audit-ready, and make regulatory work reusable instead of starting from scratch each time.

Solution

- Re-architected the platform on Google Cloud with an event-driven microservices design.
- Embedded Vertex AI and retrieval augmented generation (RAG) to automate drafting, reviews, and citation-backed responses.
- Automated creation of 300-page regulatory packs through an AI document engine.
- Integrated enterprise systems and smart workflows to reduce manual data gathering and approval delays.

Benefits

- Improved regulatory pack turnaround time by 60%.
- Saved 1,225 SME hours per pack.
- Delivered 99.9% performance reliability through scalable cloud operations.
- Made cloud transformation easier to scale with stronger compliance guardrails and reusable workflows.

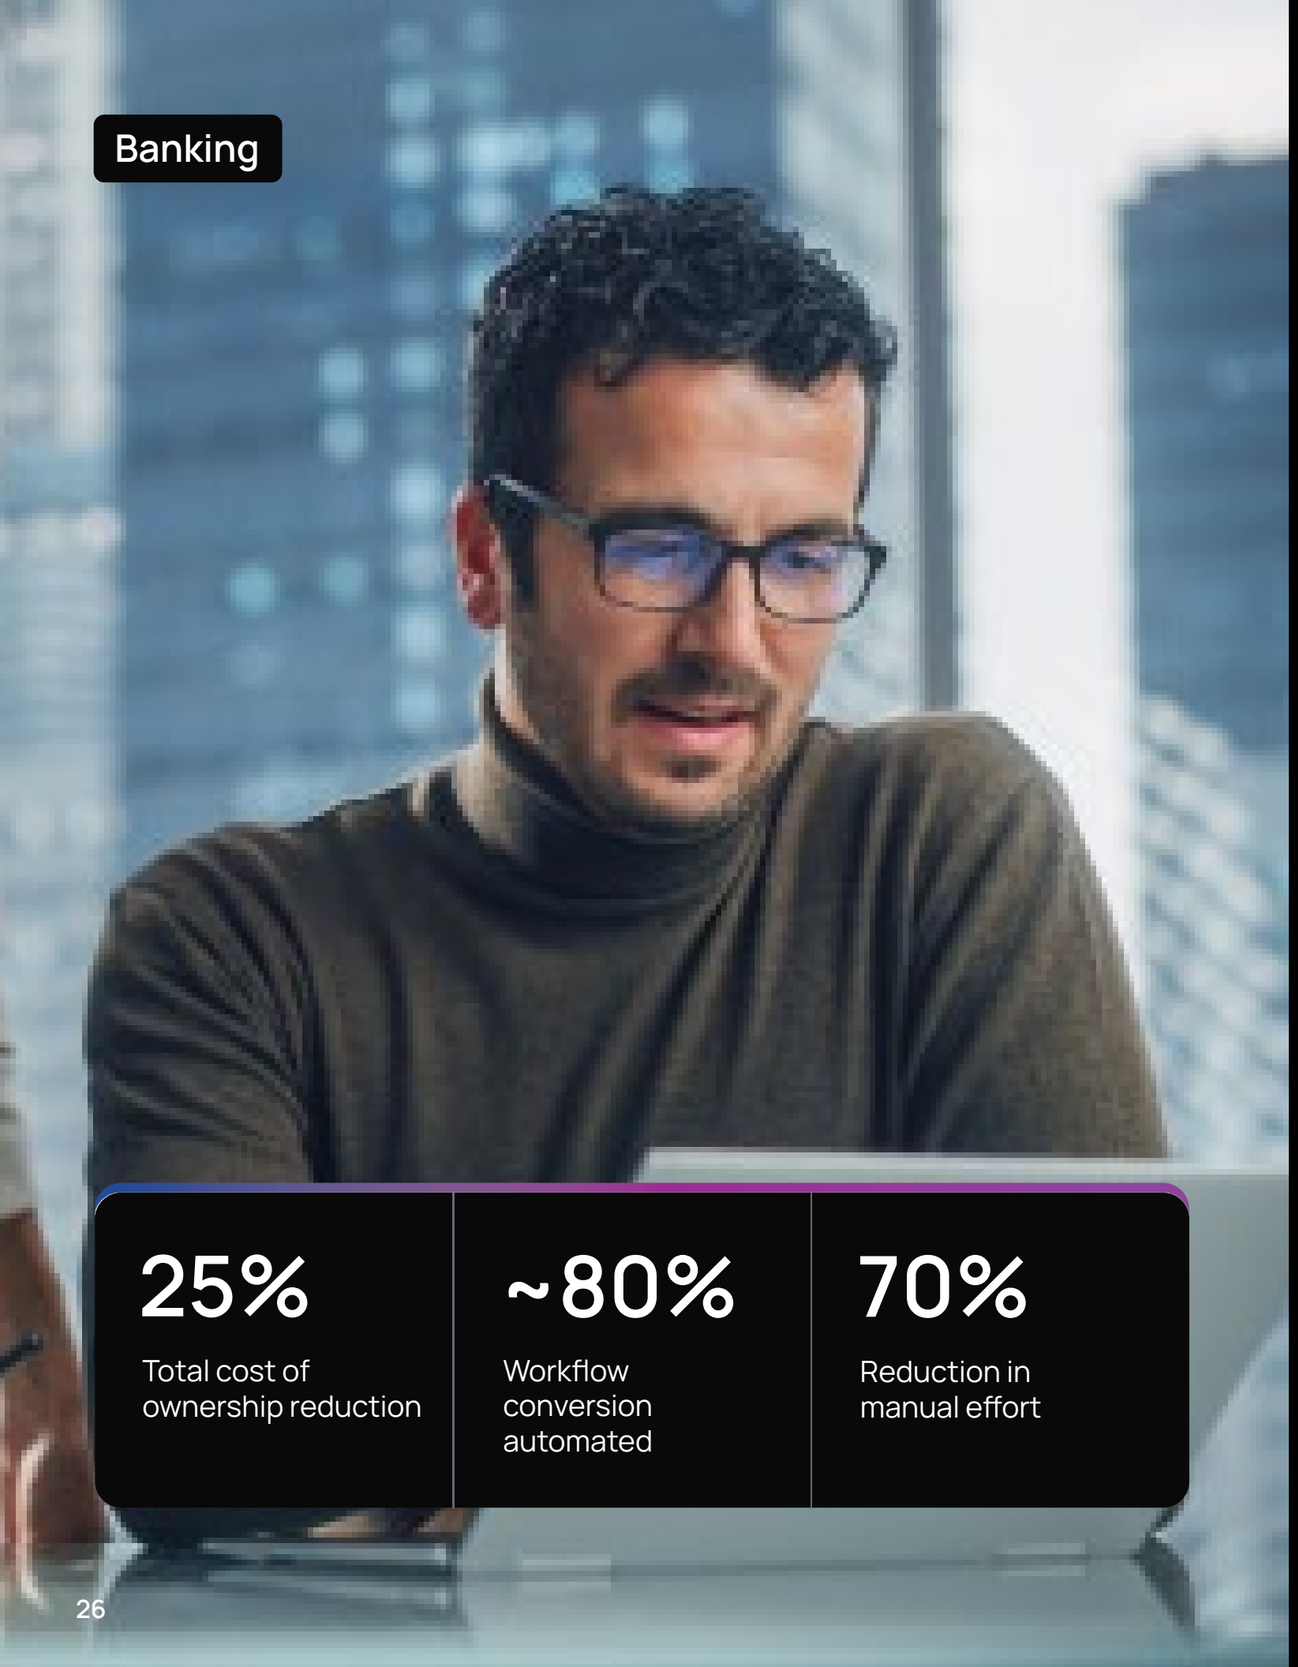

Client

A Major North American Financial Institution

Replacing licensed ETL with AI-led PySpark

The bank ran its data pipelines on Alteryx, a licensed ETL platform whose cost kept rising with use. Complex transformation logic and inter-workflow dependencies drove up development effort, and manual conversion led to inconsistent code quality. The legacy ETL landscape was hard to scale and maintain, holding the bank back from current data engineering technology.

Solution

- Delivered a Zero License model that converted Alteryx workflows into open-source Python and PySpark code the client owns.
- Implemented a three-phase approach combining deterministic reverse engineering, agentic blueprinting, and high-velocity forward engineering.
- Used Model Context Protocol to bind AI agents to the client's architecture patterns, security standards, and compliance requirements.
- Established a validation-first workflow in which every line of generated code was tested for functional equivalence before deployment.

Benefits

- Cut total cost of ownership by 25% by eliminating the Alteryx license outright.
- Automated around 80% of the workflow conversion, reducing manual effort by 70%.
- Gave the bank an open, scalable foundation it controls and can extend, free of legacy-technology constraints.

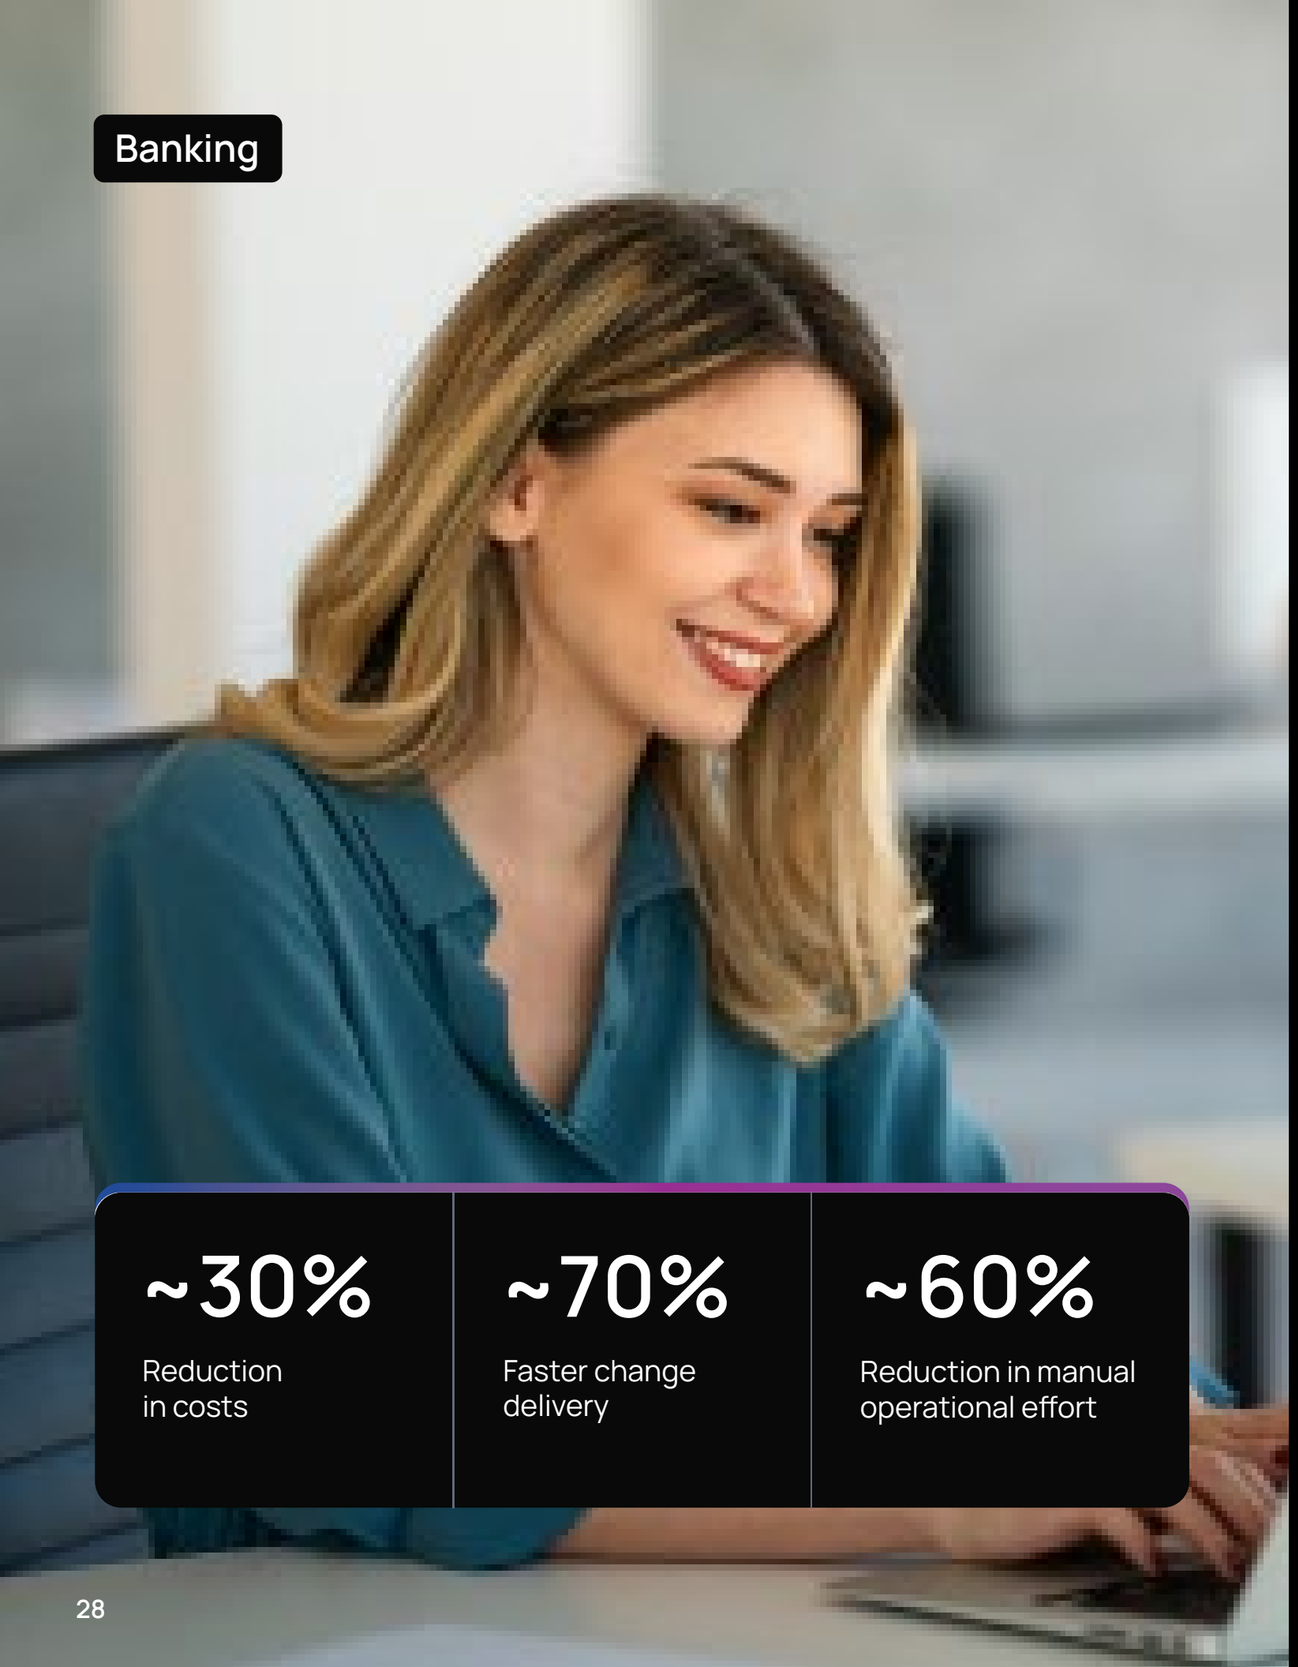

Client

A UK-Regulated Specialist B2B Transaction Bank

Reimagining a global B2B bank from the ground up with AI-embedded architecture

Legacy systems and rising operational complexity were driving costs faster than revenue growth. Limited platform scalability, manual processes, and fragmented technology stacks made it difficult to expand into new markets, maintain compliance, and deliver change at speed. The client partnered with Hexaware to define a “Bank of the Future” vision and initiate a greenfield reimagining of its platform, with AI embedded across the full software delivery lifecycle and operations.

Solution

- Built a capability-led global banking platform on reusable, composable building blocks using a plug-and-play factory model for rapid market adaptation.
- Embedded agentic AI across delivery and operations through a spec-driven, semi-autonomous multi-agent workflow with built-in governance.
- Applied a strangler fig implementation to deliver incremental value, maximize reuse of existing assets, and minimize operational disruption.
- Translated client standards into production-ready UI, APIs, integrations, and data assets, with governance, security, and compliance built in by design.

Benefits

- Up to 40% gains in developer productivity and increased test automation through AI-enabled delivery.
- Up to 50% reduction in production defects and up to 60% faster incident resolution, supporting 99.9%+ platform availability.
- 90%+ automated compliance checks, delivering consistent governance across a regulated banking environment.
- A scalable, cloud-native platform that decouples growth from operational complexity, improving the cost-to-income ratio over time.

Healthcare, Life Sciences, & Insurance

Healthcare, Life Sciences, & Insurance

8X	40%	$10M
Engineering velocity delivered	Efficiency gain across SDLC phases	Savings unlocked through AI-led optimization

Client

An American Health Information Technology and Clinical Research Company

Turning AI into velocity with 8× faster, smarter, and leaner processes

The client faced legacy-heavy systems, siloed SDLC processes, and manual-intensive operations across eight critical programs. These challenges resulted in soaring delivery timelines, high operational costs, and limited scalability—necessitating an urgent AI-led transformation.

Solution

- AI-led Hospital Pricing Data Consolidation: Scaled team capacity from 200 to 6,000 websites, delivering unmatched 30x efficiency.
- End-to-End SDLC Transformation: A tool-agnostic design approach saved 40% effort with reusable artifacts and compliance gap remediation, 10x faster tech and functional specifications generation.
- AI-Powered Code Modernization: Automated Perl-to-Python migration, mainframe code analysis, achieving up to 8× effort reduction. SOQL to Snowflake SQL Converter with 90% effort reduction.
- Live Reporting Automation: JIRA-Confluence MCP integration cut release reporting time by 98% (12 hours to 10 minutes).

Benefits

- Productivity Boost: Story points per release surged from 60 to 100 with 50% higher velocity.
- 3× Faster Delivery: Clinical data services moved from concept to production in just six weeks.
- $10M+ Cost Savings: Unlocking significant savings across SDLC phases.
- Improved Scalability: Enabled cross-functional collaboration across 10+ business units and 50+ AI solutions.

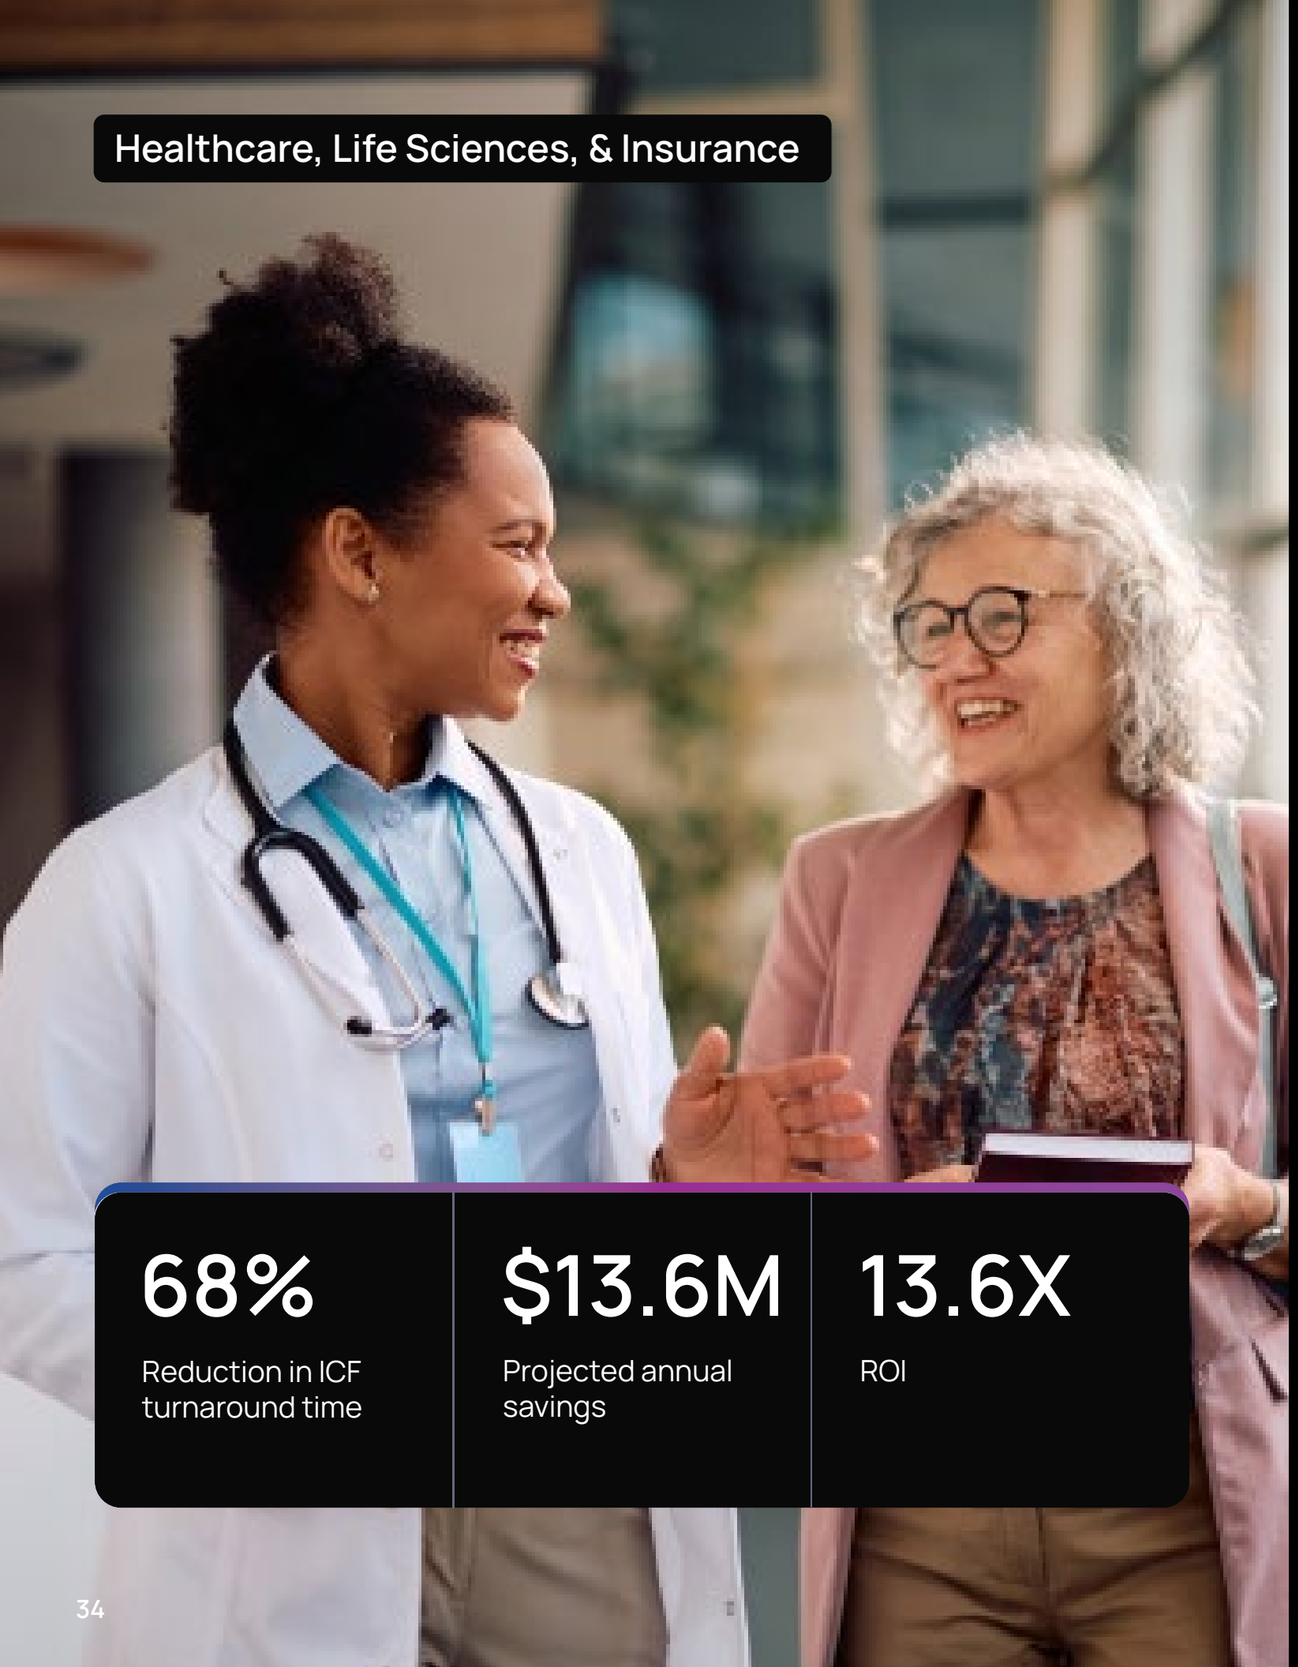

Client

A Major US-Based CRO Supporting Rapid Clinical Research and Drug Development

Cutting informed consent form (ICF) turnaround from days to hours without adding compliance risk

Manual ICF customization was slow, fragmented, and hard to scale, creating delays and rework across trial readiness. The client needed a faster, more dependable process that could improve quality while lowering the burden on teams.

Solution

- Built a multi-agent AI platform to automate country-specific ICF customization.
- Created a role-based portal with secure access, track changes, and a full audit trail.
- Applied AI-driven document detection, modification, and validation aligned to HIPAA and customization plans.
- Designed the Azure-based architecture for adoption, future learning, and integration with clinical systems.

Benefits

- Reduced ICF turnaround time by 68%.
- Projected annual savings of $13.6 million.
- Delivered a 13.6x ROI outlook.
- Improved consistency, compliance, and scalability for a document-intensive process.

Healthcare, Life Sciences, & Insurance

94%
Reduction in HCP engagement cost

45%
Faster MLR cycle

40%
Faster APAC brand launch

Client

A Swiss Multinational Pharmaceutical Corporation

Rebuilding commercial operations around precision, speed, and lower cost

Fragmented data, slow content cycles, and uneven CRM adoption were limiting commercial impact across APAC markets. The client needed a smarter operating model that could improve coverage, accelerate launches, and do more at a lower cost to serve.

Solution

- Designed an AI-native commercial intelligence ecosystem tailored by brand and market.
- Built a cardiovascular patient operating model with six AI agents across eight pipeline stages.
- Enhanced Salesforce and Veeva workflows with next-best-action and territory optimization.
- Established a compliance-first model with unified data governance and human approval loops.

Benefits

- Reduced healthcare professional (HCP) engagement cost by 94%, lowering per-contact cost from $180–220 to $8–12.
- Cut MLR review cycles by 45%, reducing timelines from 14–18 weeks to 7–9 weeks.
- Improved APAC brand launch speed by 40%.
- Increased patient program enrollment by 38% while improving HCP targeting precision by 30–40%.

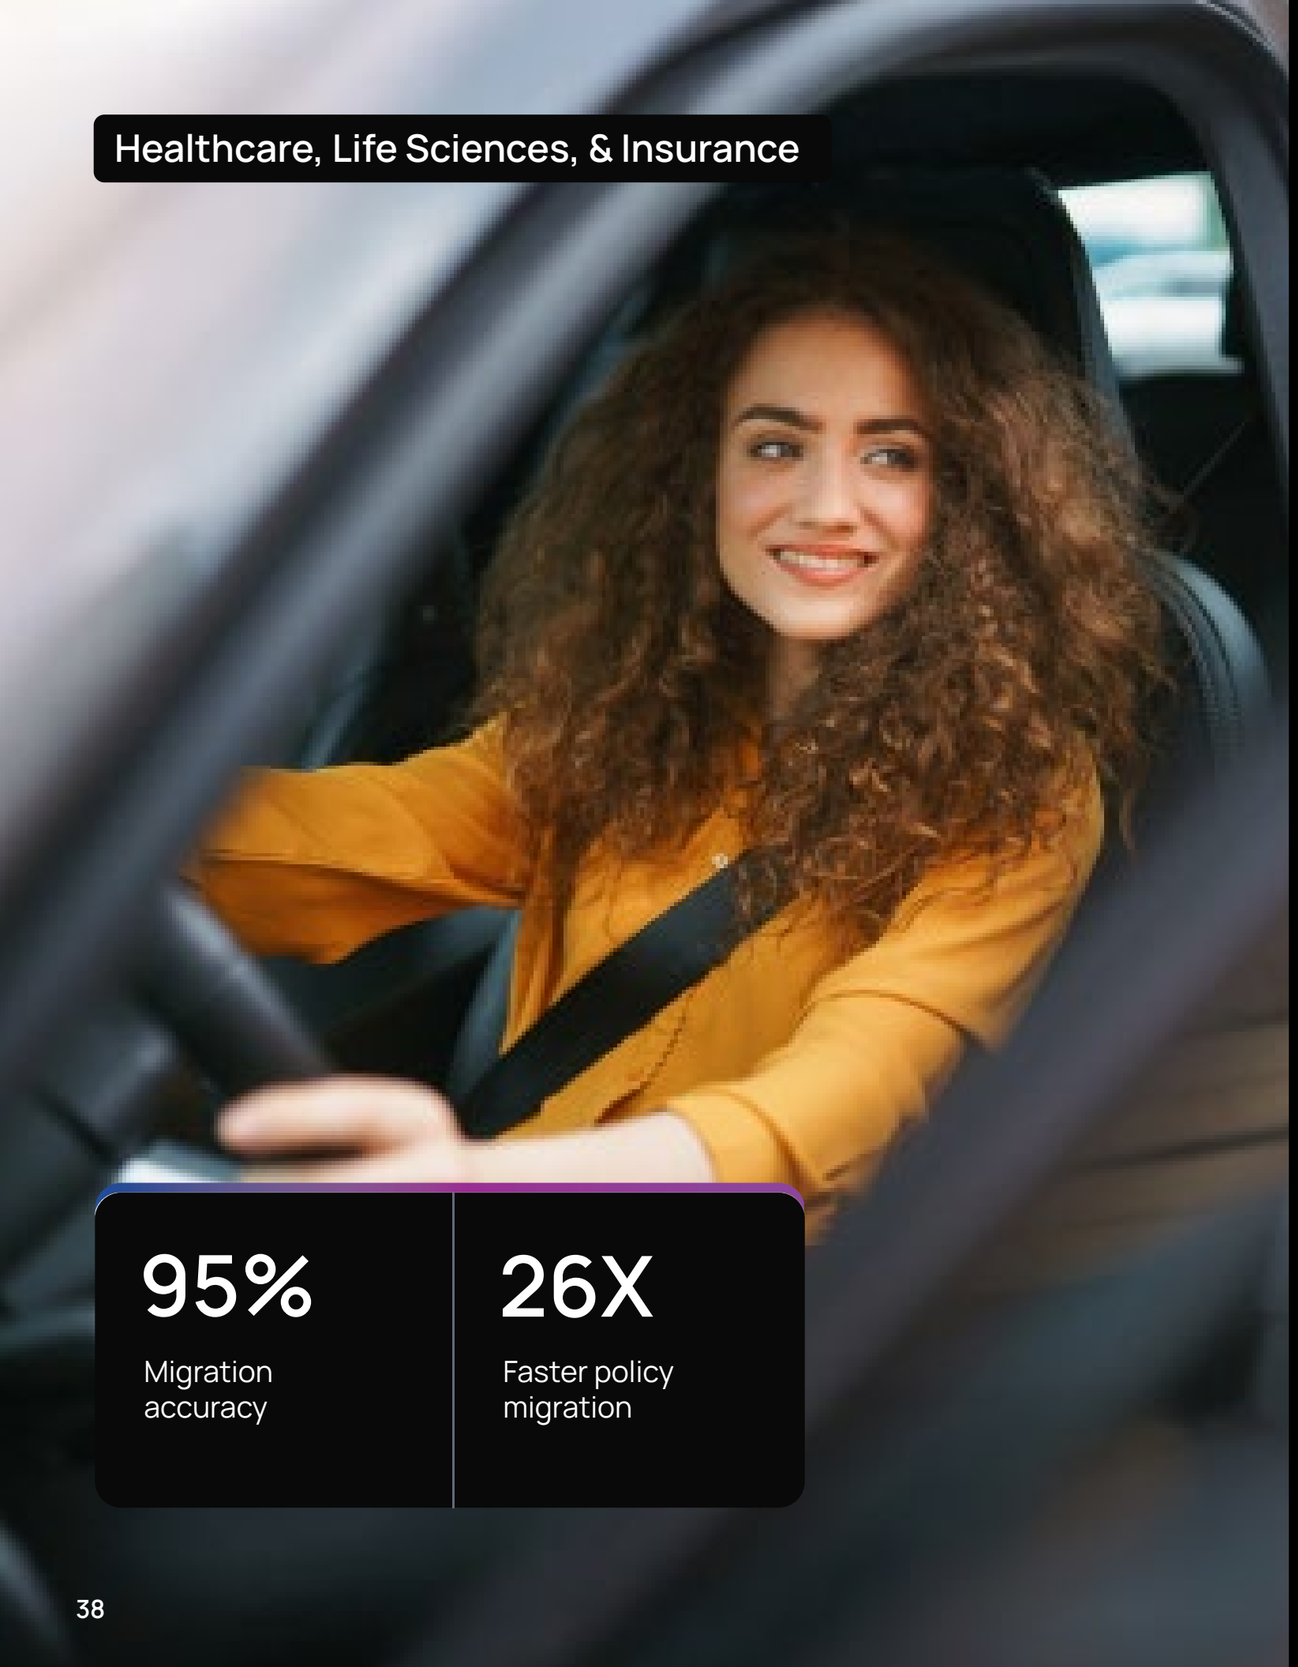

Healthcare, Life Sciences, & Insurance

95%
Migration accuracy

26X
Faster policy migration

Client

A Leading Belgian Non-Life Insurer

Migrating 100,000+ policies with agentic AI

With source data unavailable and only policy documents to work from, the client faced migrating 100,000+ insurance policies manually — a slow, error-prone effort estimated at roughly 12,500 working days that risked major operational delays.

Solution

- Deployed an agent-driven migration approach using Azure AI services and OpenAI.
- Extracted and mapped policy data directly from source documents.
- Completed end-to-end migration in three weeks using five consultants.
- Automated validation to preserve accuracy at scale.

Benefits

- Reduced coding cycle time by 96% versus the OCR approach.
- Increased migration velocity by 26X—from 1200 working days to 45 working days.

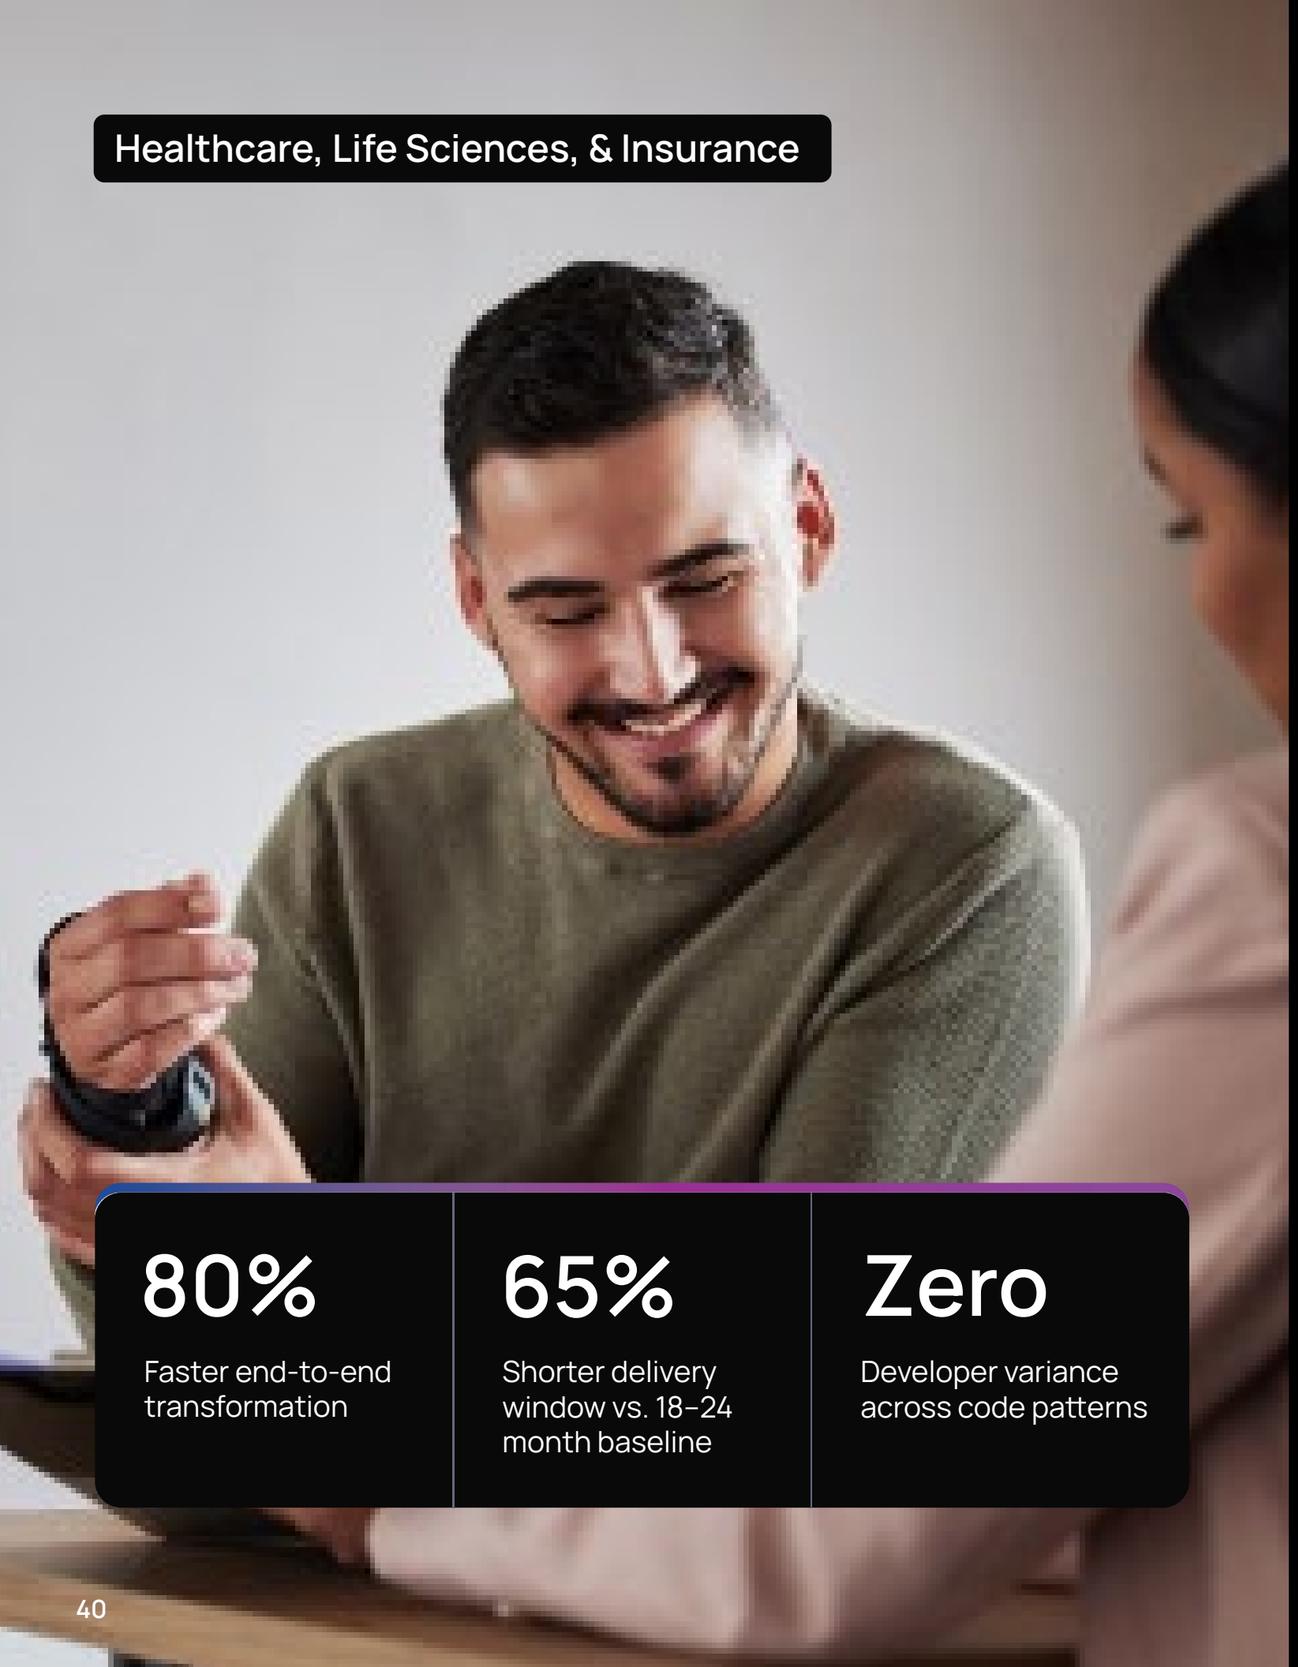

Client

A Global Insurance and Asset Management Group

Accelerating Windows EC2-to-EKS modernization with an AI-native factory model

Addressed a complex modernization mandate where 200 custom-built .NET applications on AWS EC2 Windows had to move to AWS EKS with Linux containers, without creating delivery variance or audit risk.

Solution

- Built a two-engine transformation model combining Amaze® Replatform Engine for deterministic assessment and a Claude Code Skill Library for agentic, pattern-based code transformation.
- Automated repeatable modernization patterns across .NET upgrades, React UI translation, MSMQ-to-AWS messaging migration, Kerberos removal, secrets management, containerization, and Kubernetes deployment readiness.
- Preserved governance by ensuring every change was traceable, reversible, mapped to an approved transformation pattern, validated through build and test checks, and opened as an auditable pull request.

Benefits

- Achieved 80% faster end-to-end transformation.
- Maintained zero developer variance across code patterns.
- Delivered a reusable transformation playbook for the broader application portfolio and future operating entities.

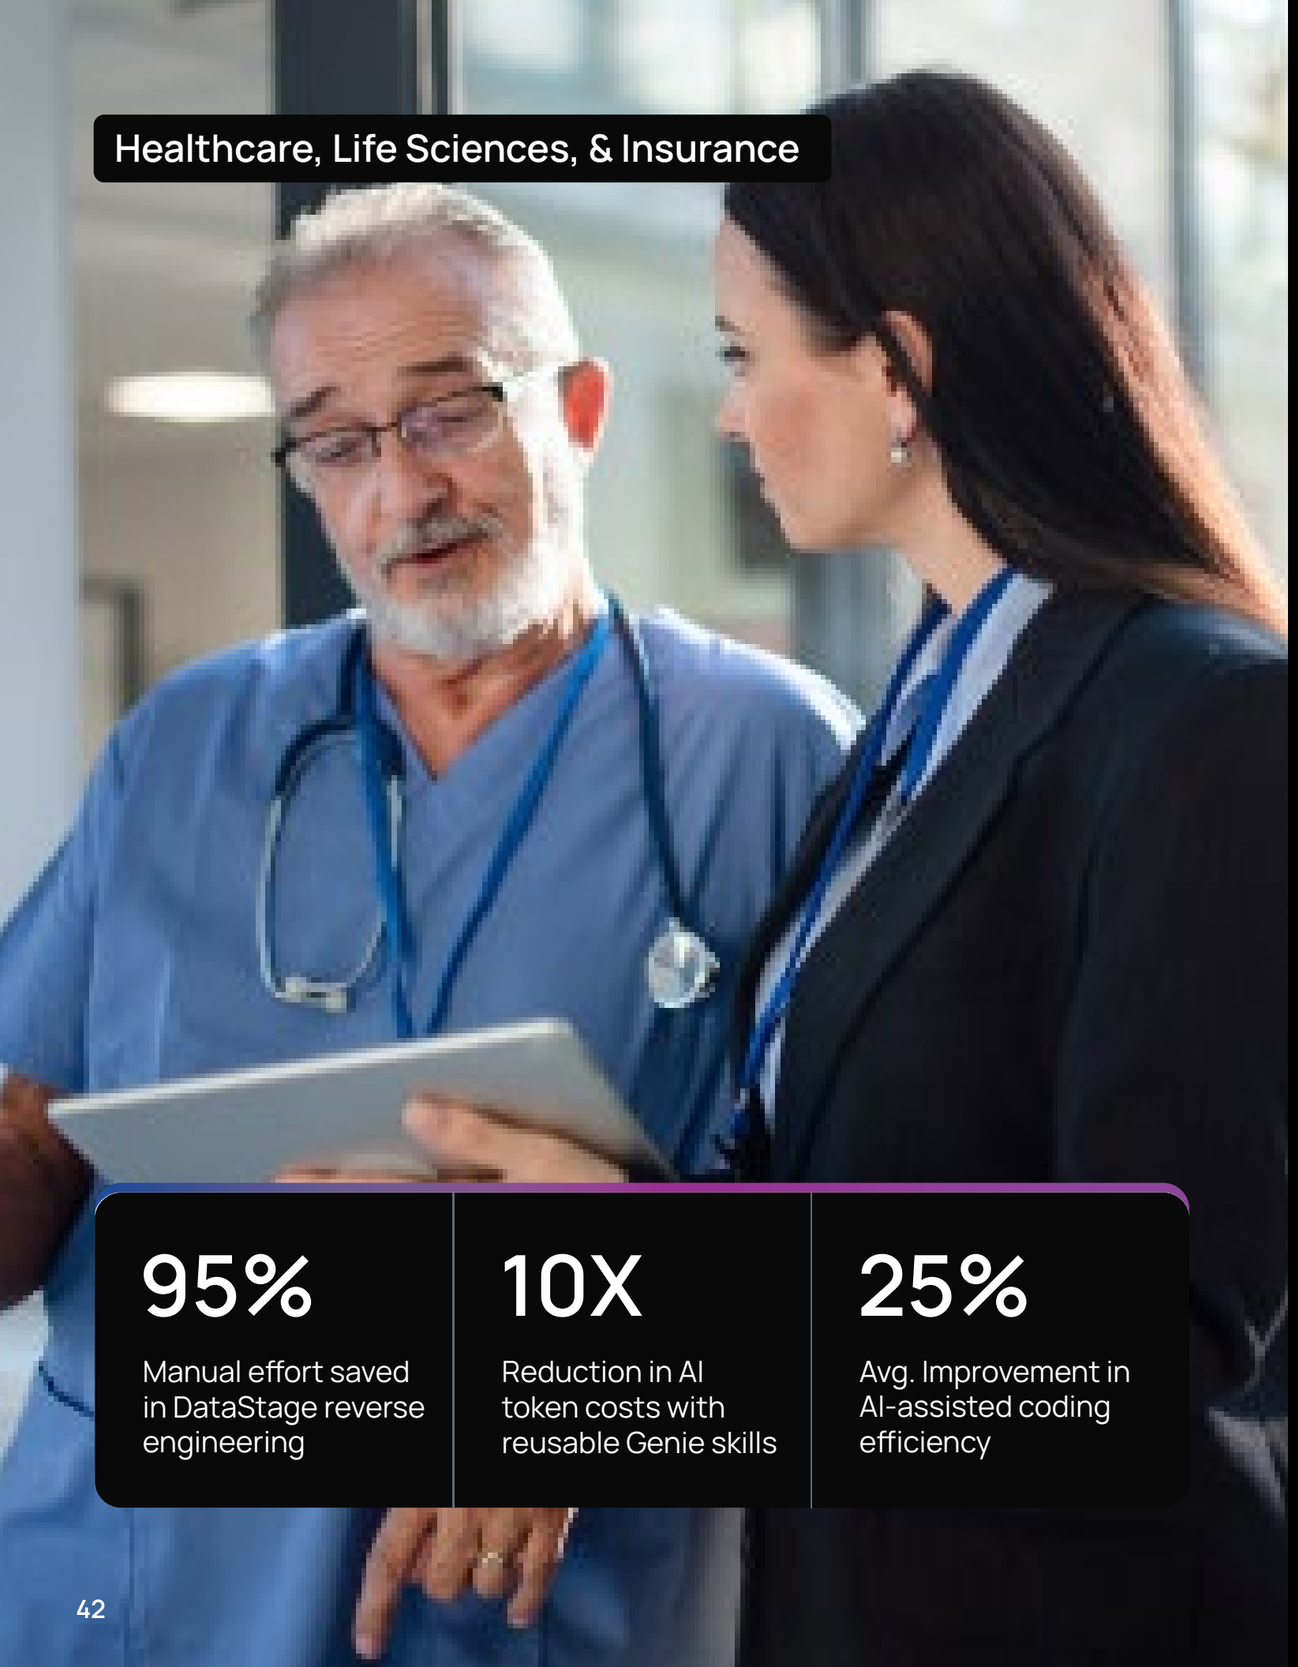

95%
Manual effort saved in DataStage reverse engineering

10X
Reduction in AI token costs with reusable Genie skills

25%
Avg. Improvement in AI-assisted coding efficiency

Client

A Leading U.S. Managed Care Organization Focused on Government-sponsored Healthcare Programs

Accelerating healthcare data modernization with AI-assisted engineering

The client was executing multiple strategic data initiatives simultaneously, including CMS interoperability, legacy EDWH decommissioning, HEDIS engineering, and data lake modernization. The teams needed to accelerate development while ensuring data quality across complex healthcare data platforms.

Solution

- Leveraged Databricks Genie across four active data engineering programs.
- Supported AI-assisted development for CMS interoperability, HEDIS, and Data Lake modernization initiatives.
- Used Genie for data analysis and validation during migration activities from legacy data warehouse workloads to a Databricks-based data lake architecture.
- Created a reusable skill to reduce token usage.
- Enabled engineering teams to accelerate coding, test case generation, and validation activities.

Benefits

- Reduced manual effort by 95% through AI-assisted DataStage reverse engineering, cutting analysis time from 3 days to 1 hour across 91 jobs.
- Lowered AI operating costs by 10x with reusable Genie skills for token optimization.
- Improved developer productivity by an average of 25% through AI-assisted coding across multiple programs.
- Accelerated testing and delivery by reducing manual FHIR validation effort by 25%+ and cutting metadata versioning time by 25%.

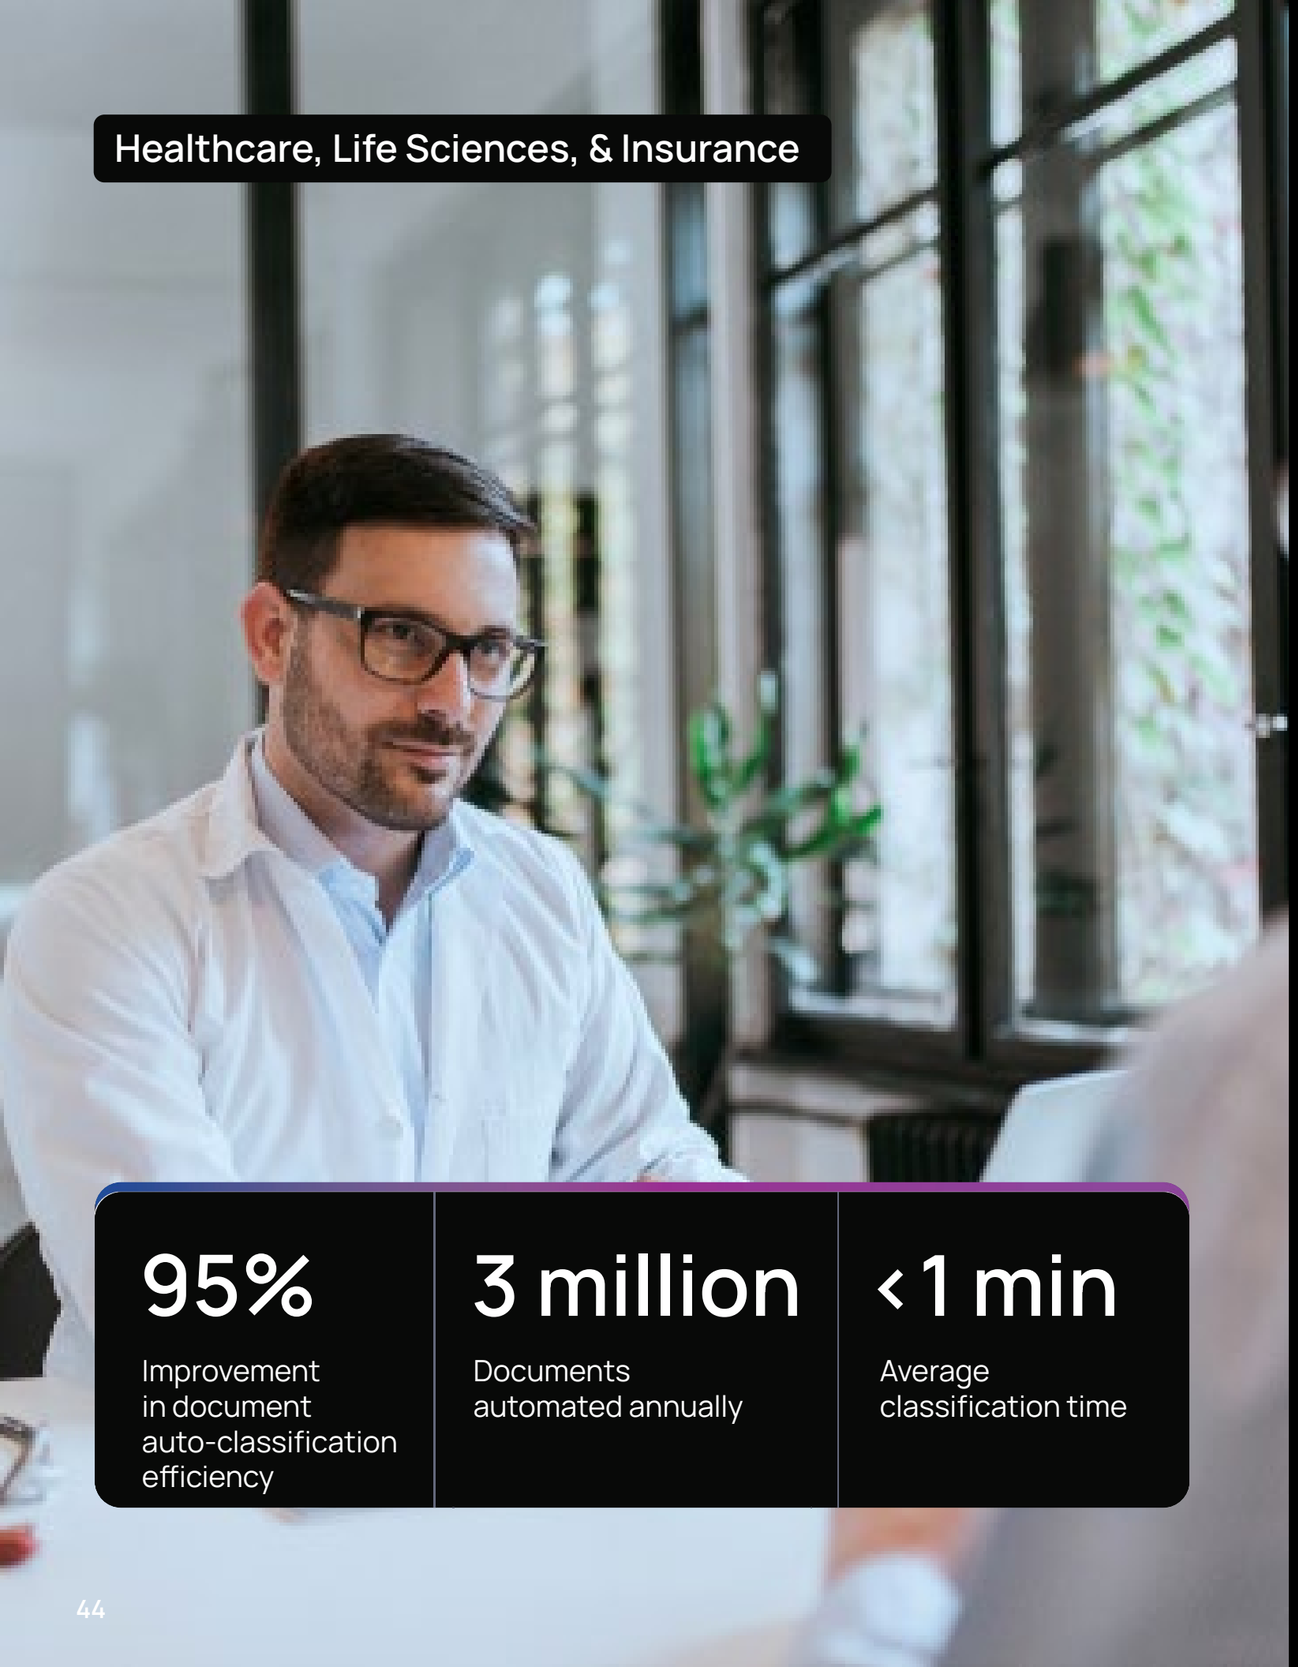

Client

A Global Clinical and Preclinical Trials and Research Enterprise

Classifying millions of clinical documents without slowing the business

Manual document classification was too slow for the volume and complexity of the client's trial documentation. The business needed a scalable, compliant way to organize records faster without compromising accuracy.

Solution

- Applied AI automation and GenAI to records management, classification, sorting, and audit-trail integrity.
- Built a deep-learning library and text analytics framework to create taxonomies and models.
- Automated classification at scale to support an annual volume of roughly three million documents.
- Designed the workflow to support regulatory compliance and growing document volumes.

Benefits

- Improved auto-classification efficiency by 95%.
- Automated classification of about three million complex medical documents annually.
- Reduced average classification time to under one minute from roughly 20 minutes manually.
- Improved scalability and access to critical records without sacrificing compliance.

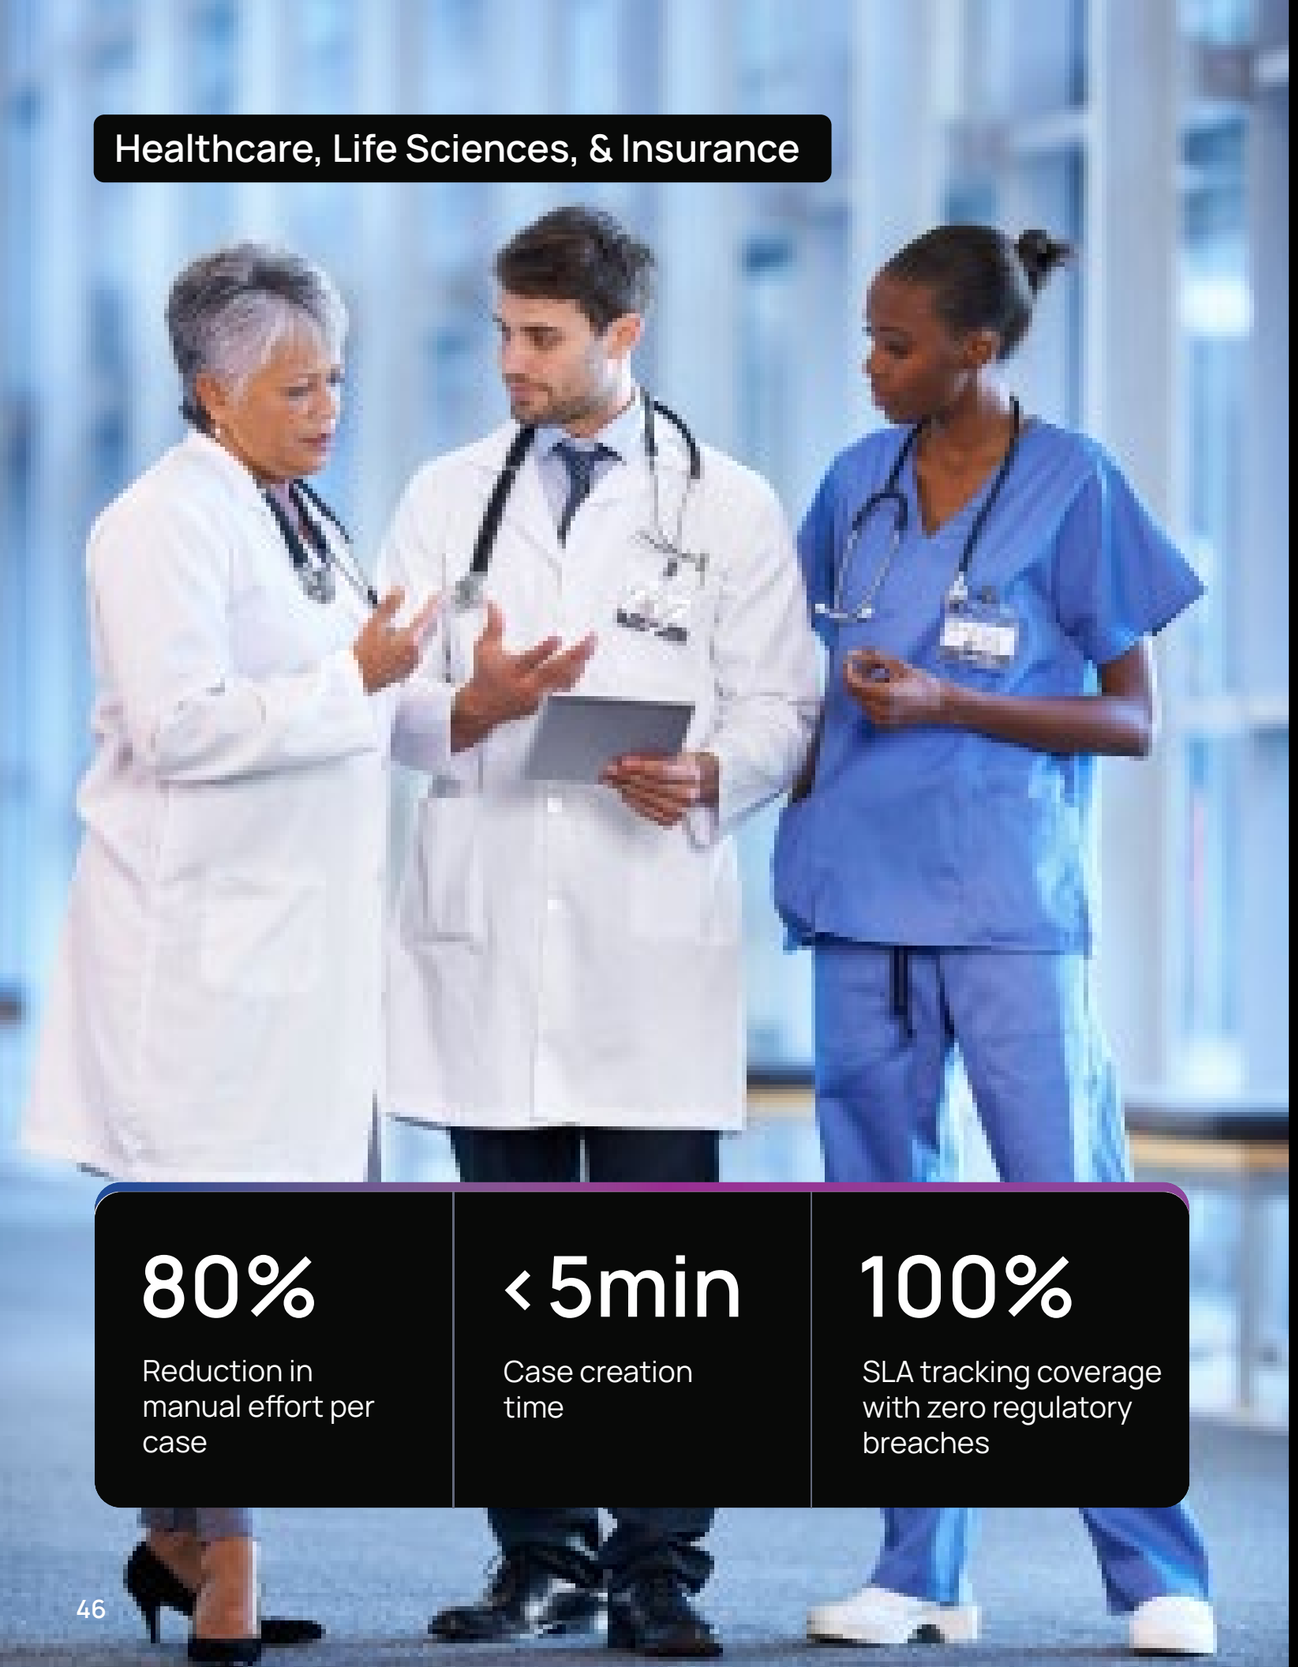

80%
Reduction in manual effort per case

<5min
Case creation time

100%
SLA tracking coverage with zero regulatory breaches

Client

A Leading Global Contract Research Organization (CRO)

Taking pharmacovigilance (PV) intake from inbox overload to audit-ready speed

An email-led PV intake process was creating operational drag, inconsistent tracking, and unnecessary regulatory exposure. The client needed a more dependable way to process adverse event cases at scale without adding headcount linearly.

Solution

- Deployed an end-to-end AI agent system covering the full PV intake lifecycle.
- Automated mailbox monitoring, metadata extraction, and attachment handling.
- Used AI to classify cases, flag edge cases, and route work to the right teams.
- Monitored SLA deadlines in real time and automated quality checks and escalation workflows.

Benefits

- Reduced manual effort per case by 80%.
- Cut case creation time to under five minutes from a roughly 40-minute baseline.
- Achieved 100% SLA tracking coverage with zero regulatory breaches.
- Reduced QC processing time by about 60% while improving audit readiness.

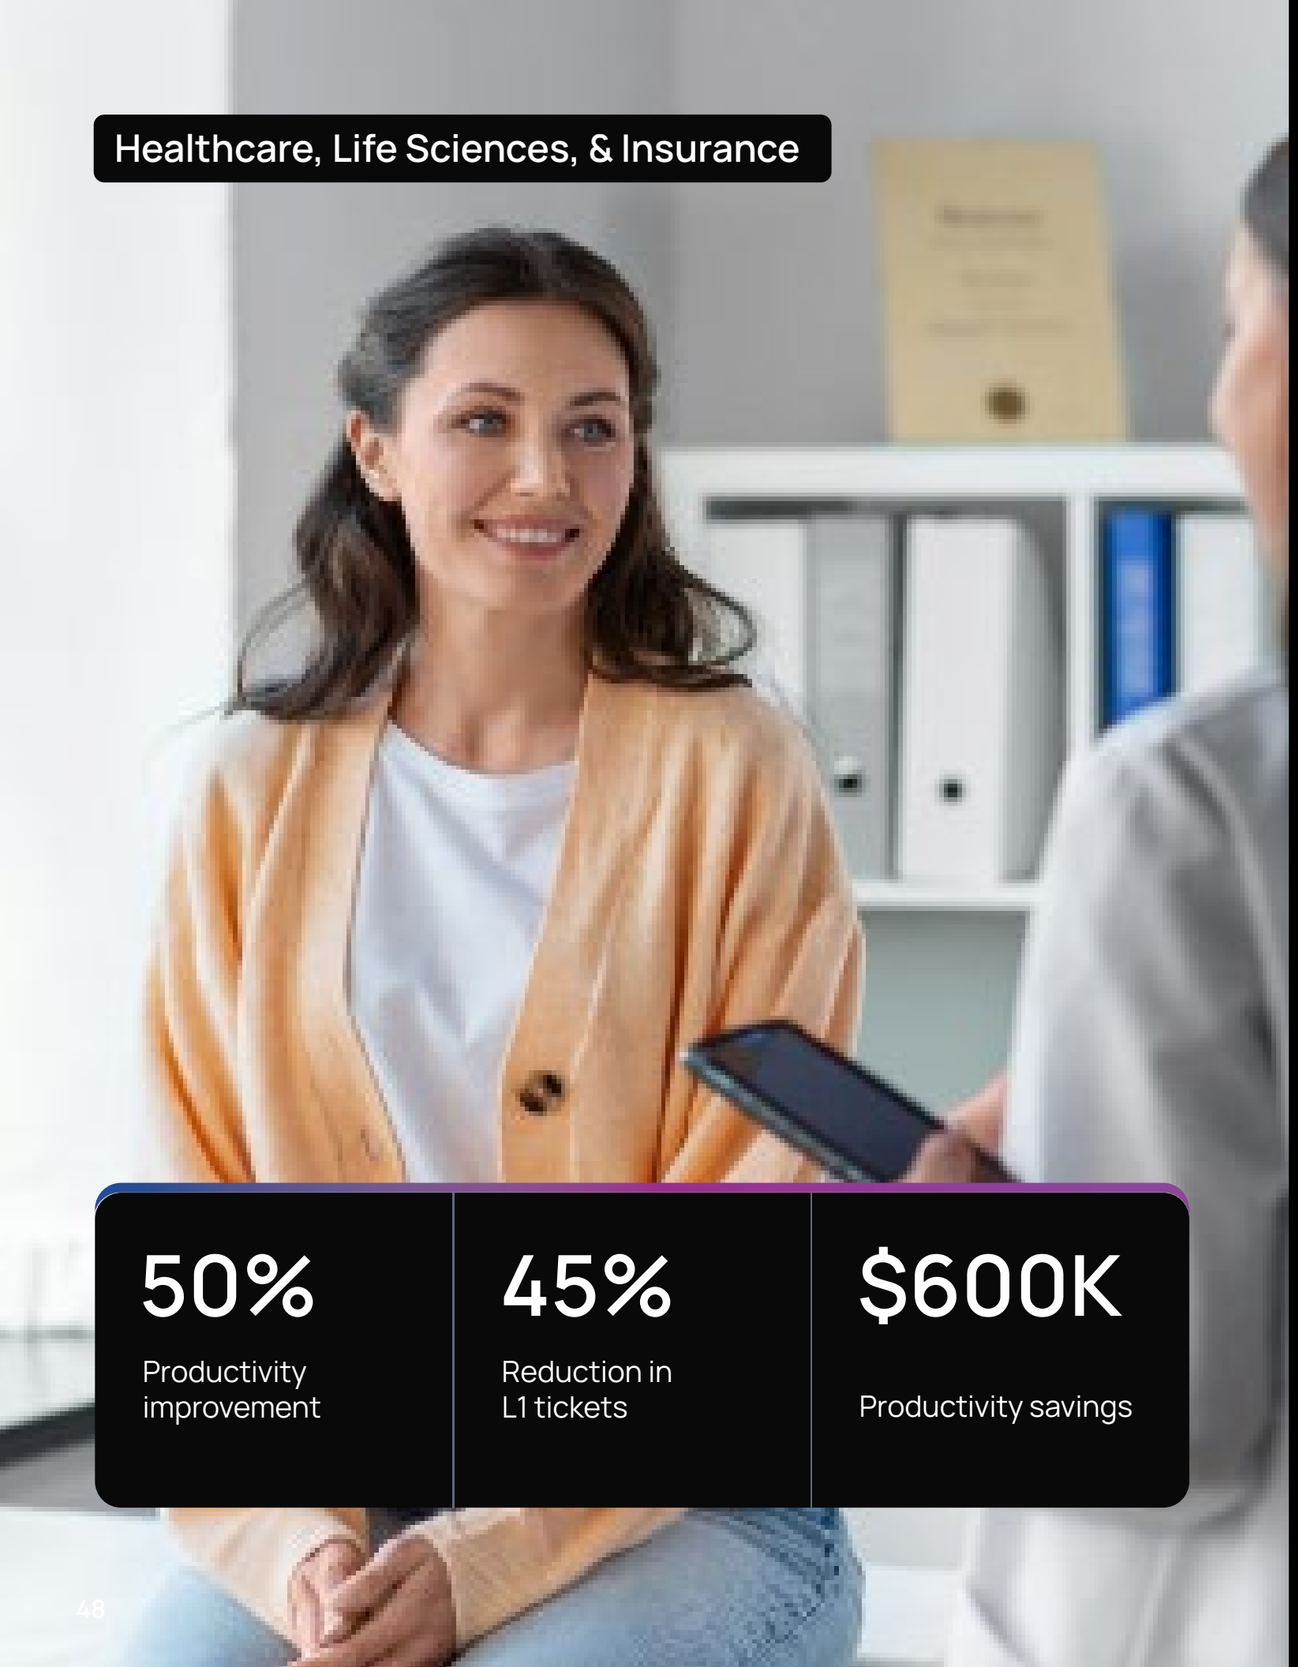

Healthcare, Life Sciences, & Insurance

50%	45%	$600K
Productivity improvement	Reduction in L1 tickets	Productivity savings

Client

A Global MedTech & Women's Health Technology Company

Scaling Agentic AI across ServiceNow-driven IT operations

The client's global ServiceNow platform spans 14 custom applications on one of its global instances, supporting incident management and enhancement work under a formal AI governance model. As the client scaled agentic AI across ITSM, HRDS, ITOM and SecOps, and testing workflows, it needed a delivery partner to keep pace responsibly.

Solution

- Automated ITSM workflows to summarize incident details, resolution notes, and chat threads.
- Customized AI agent using ServiceNow AI/LLM to be deployed based on the existing workflow.
- Streamlined onboarding and offboarding through end-to-end employee provisioning and access management.
- Automated test case creation and execution using NowAssist.

Benefits

- 50% productivity improvement.
- Delivered approximately $600K of productivity savings.
- Enabled agentic AI to resolve incidents and tasks without L1 intervention.

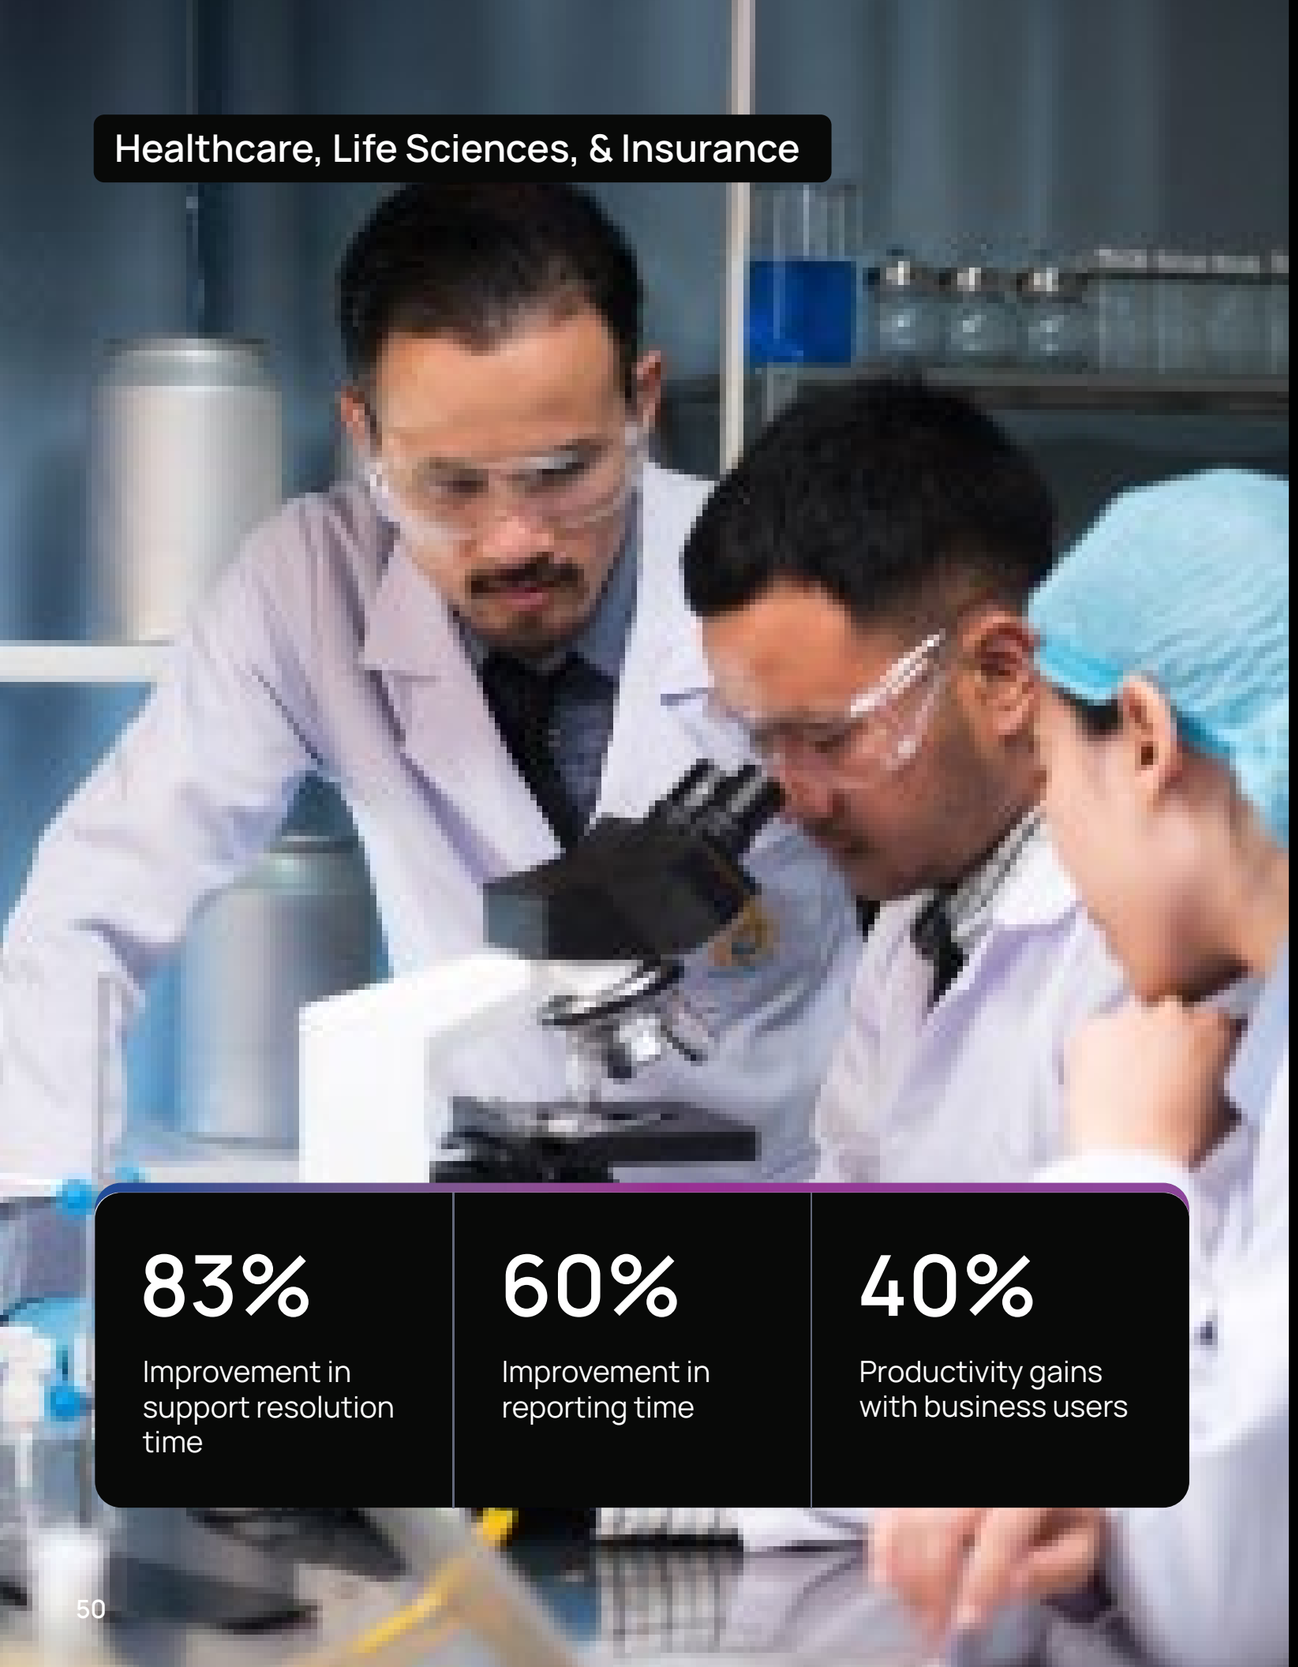

Client

A Leading US Based Life Sciences Tools Company

Scaling agentic AI across customer support, treasury, and product operations

The client manufactures lab equipment, laboratory chemicals, and bio-processing materials for life sciences customers worldwide. With several agentic AI initiatives running concurrently across support, treasury, and product teams, they needed a coordinated approach to scale AI safely.

Solution

- Deployed a multi-agent tool, powered by OpenAI, to resolve tickets, access knowledge articles, and fetch required supporting data and workflows. This will speed up support resolution.
- Built an AI/ML pipeline for treasury credit worthiness and credit risk forecasting to speed up business decisions.
- Built Agentic workflows for product/chemical content enrichment and lab experiment data workflows.

Benefits

- Reduced customer support resolution time by 83%, cutting average ticket resolution from 30 minutes to 5 minutes.
- Reduced reporting time by 60%, from 750 hours to 300 hours.
- 40% productivity gains across AI-enabled workflows.

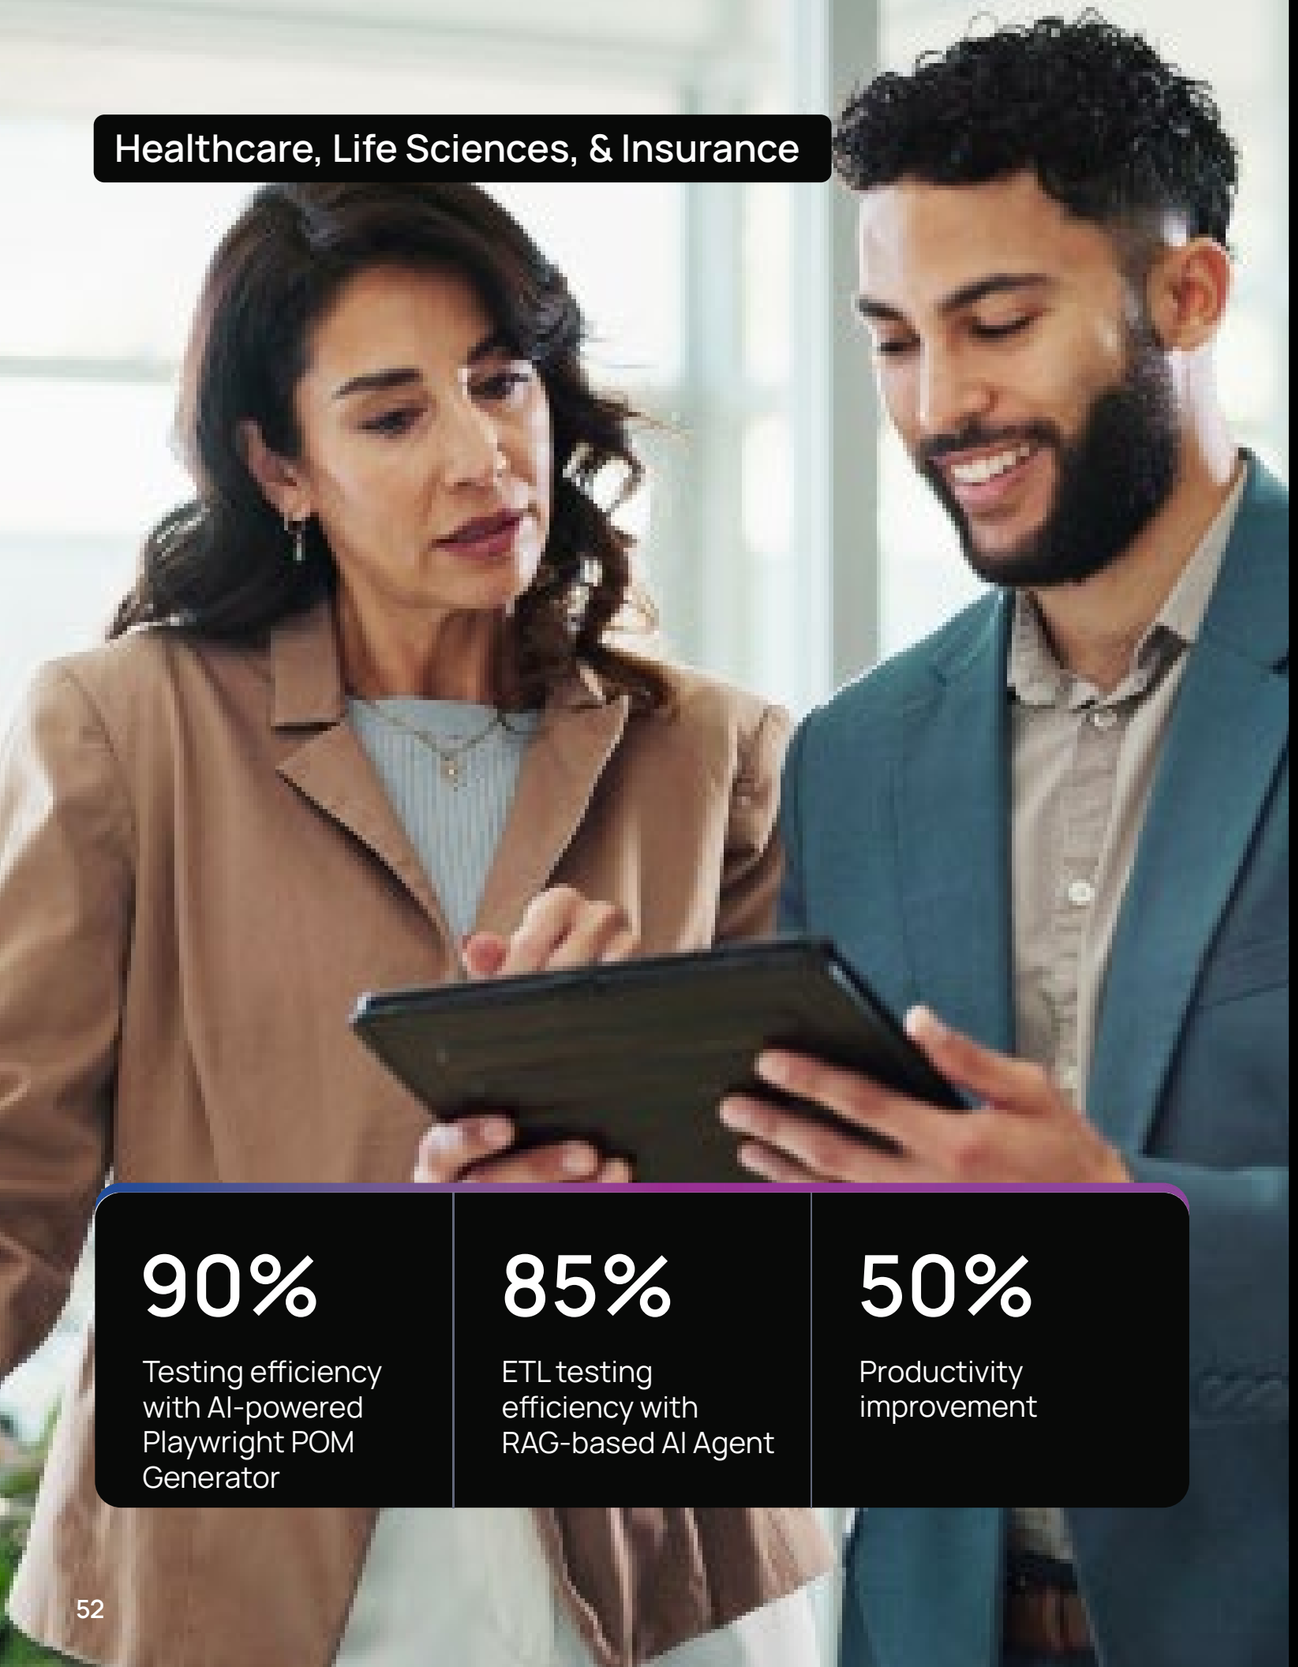

90%
Testing efficiency with AI-powered Playwright POM Generator

85%
ETL testing efficiency with RAG-based AI Agent

50%
Productivity improvement

Client

A Leading Global Insurance Provider Operating Across Europe, North America, and Asia

Accelerating application modernization and testing with AI-powered engineering

The client sought to improve software engineering productivity across application modernization, testing, and infrastructure initiatives while managing a large portfolio of support and transformation programs. The objective was to accelerate delivery, enhance quality, and meet contractual productivity commitments through AI-driven automation.

Solution

- Deployed a Retrieval-Augmented Generation (RAG)-based ETL Testing Agent to automate semantic data validation and improve testing accuracy.
- Implemented an AI-powered Playwright Page Object Model (POM) Generator to automate the creation of reusable test assets.
- Playwright is integrated with Github Copilot tool and with Human-in-the-loop to enhance productivity and savings.
- Automated Node.js environment upgrades using an AI coding assistant, reducing manual effort and improving dependency management.

Benefits

- **Improved testing efficiency by 90%** with the AI-powered Playwright Page Object Model Generator.
- **Increased ETL testing efficiency by 85% using the RAG-based AI Agent.**
- **Improved productivity by 50% through AI-assisted testing.**
- **Accelerated AI adoption with reusable testing automation.**

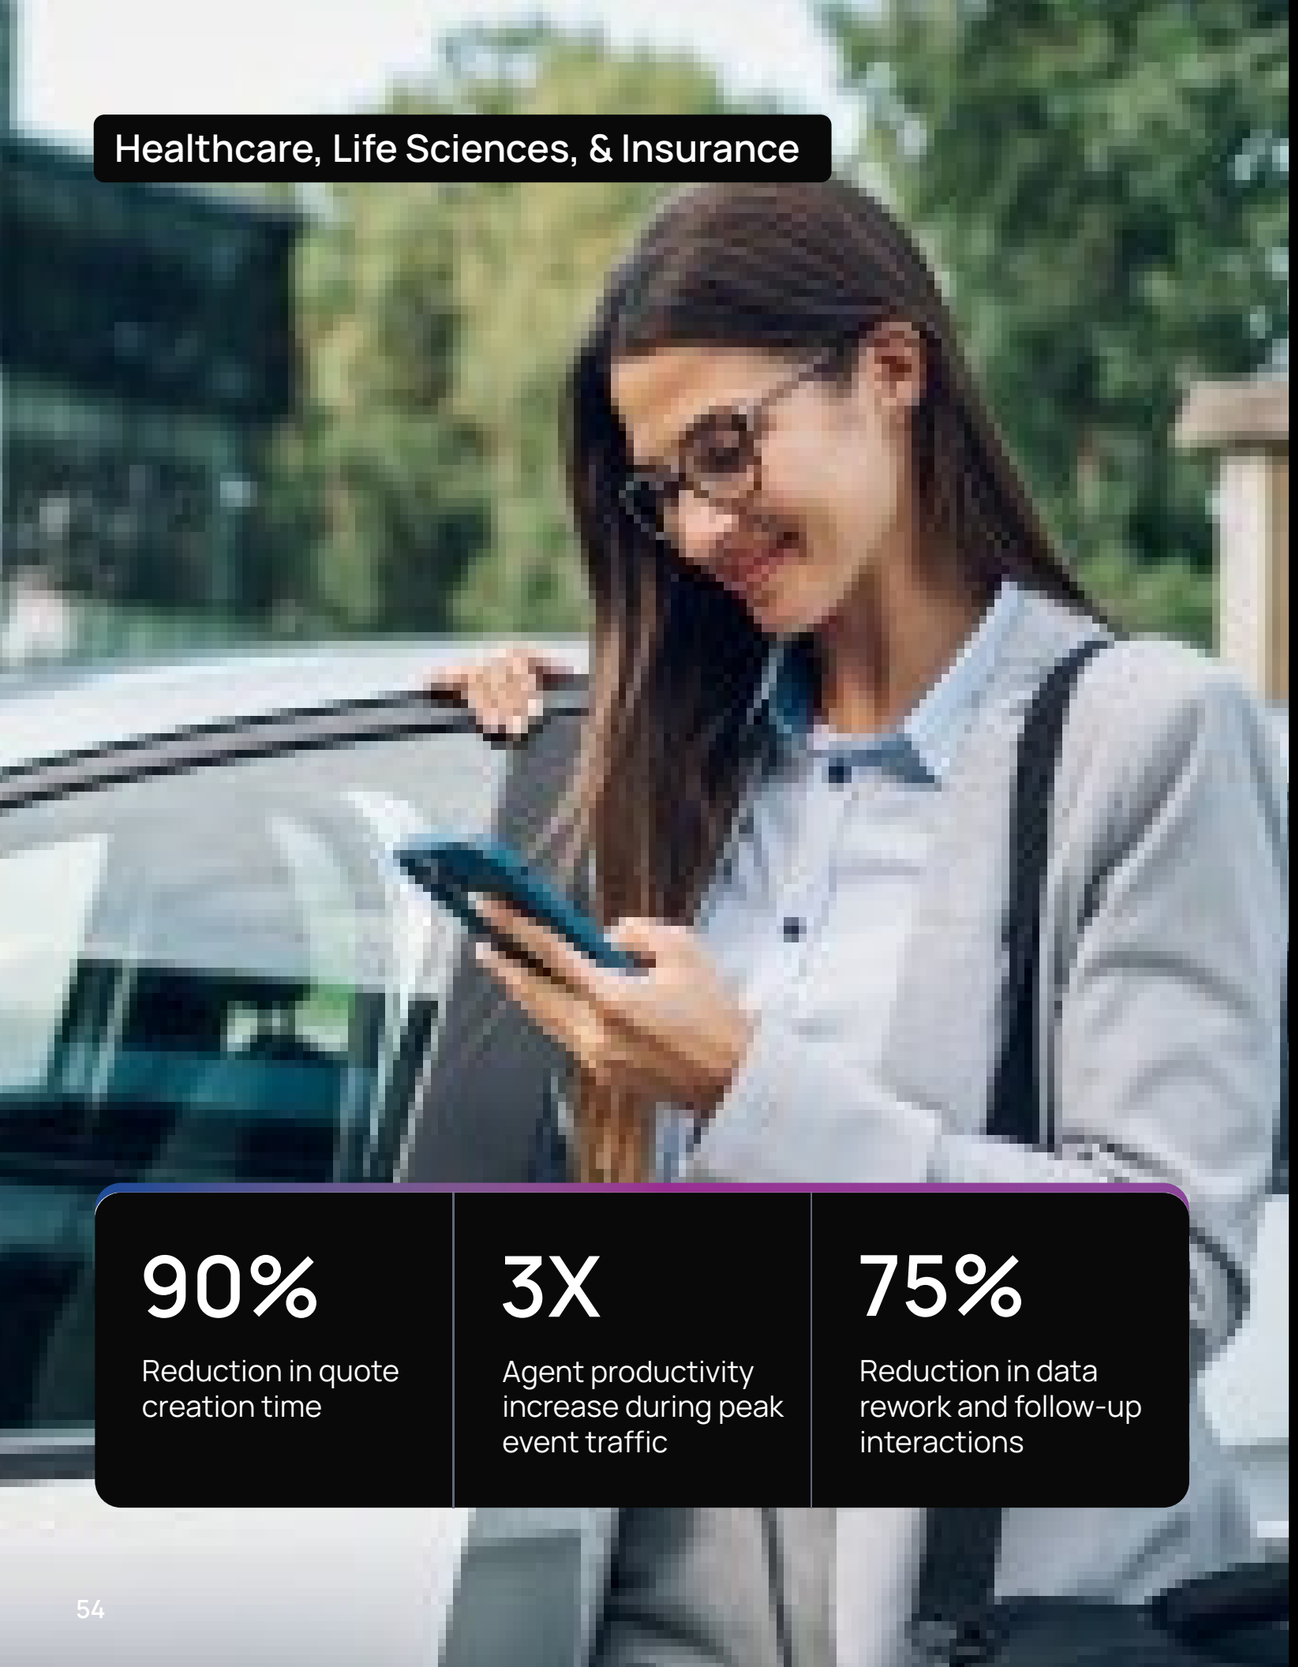

Client

A Leading European Non-Life Insurer

From photo to premium in under 60 seconds: delivering instant car insurance quotes with GenAI

Customers purchasing vehicles expected instant insurance cost insights, but traditional quotation journeys required manual data entry that created friction, errors, and abandonment. Rule-based estimation approaches could not scale to the growing diversity of vehicle models, limiting both speed and differentiation at a critical decision moment.

Solution

- Launched a GenAI-powered solution generating instant insurance and loan estimates from a single vehicle photo.
- Applied computer vision to detect vehicle make, model, fuel type, and variant automatically from the image.
- Enriched and validated vehicle attributes using Generative AI before passing structured data to pricing engines via secure APIs.
- Generated real-time insurance premiums and loan estimates, enabling customers to proceed to digital onboarding without re-entering any data.

Benefits

- Raised quote completion rates from 45% to 78%, a 33 percentage point increase driven by reduced journey friction.
- Increased agent throughput from 5-6 to 18-20 quotes per hour during peak event traffic.
- Reduced AI-extracted data correction rate to under 3%, compared to 12% with manual entry.
- Established a scalable AI foundation adaptable to new vehicle models and additional insurance products.

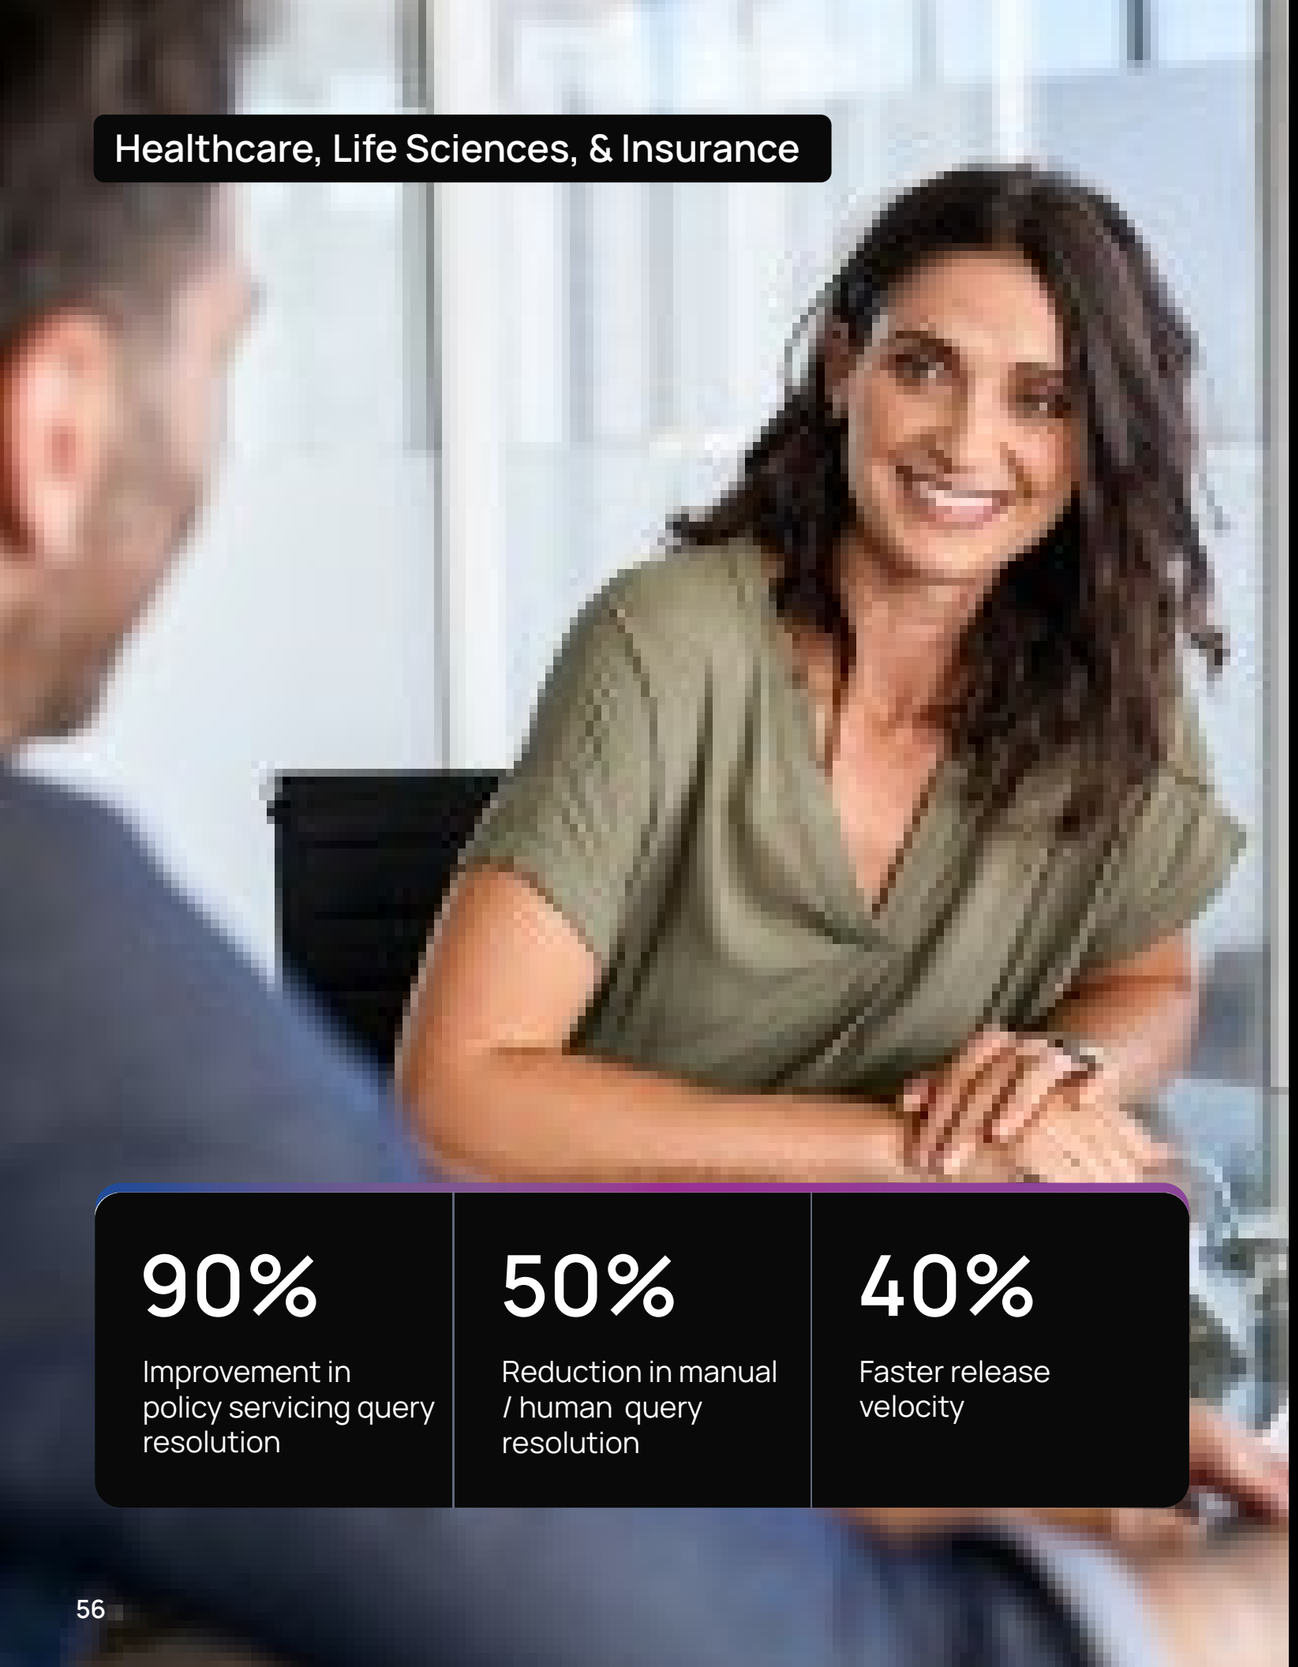

90%
Improvement in policy servicing query resolution

50%
Reduction in manual / human query resolution

40%
Faster release velocity

Client

A Leading UK Insurance Provider

Driving innovation with AI-powered solutions for operational & delivery excellence

Burdened by legacy systems, increasing technical debt, and infrastructure constraints, the client faced slower development cycles and operational inefficiencies that negatively affected productivity, revenue growth, and time-to-market.

Solution

AI-Assisted Application Modernization:

- Utilized RapidX® and AI forward engineering to modernize the client's policy servicing system.

GenAI-Powered Software Engineering:

- Leveraged Cursor IDE and Generative AI tools to reduce lead times for development and debugging.

Azure OpenAI Knowledge Assistant:

- Deployed a knowledge chatbot using Azure OpenAI for improved customer service.

Intelligent Workflow Automation:

- Delivered scalable and automation-ready solutions for advisor registration and administration processes.

Benefits

- Accelerated software delivery.
- Higher operational accuracy and compliance.
- Scalable and modernized systems for long-term growth and innovation.
- Increased workforce productivity.

Healthcare, Life Sciences, & Insurance

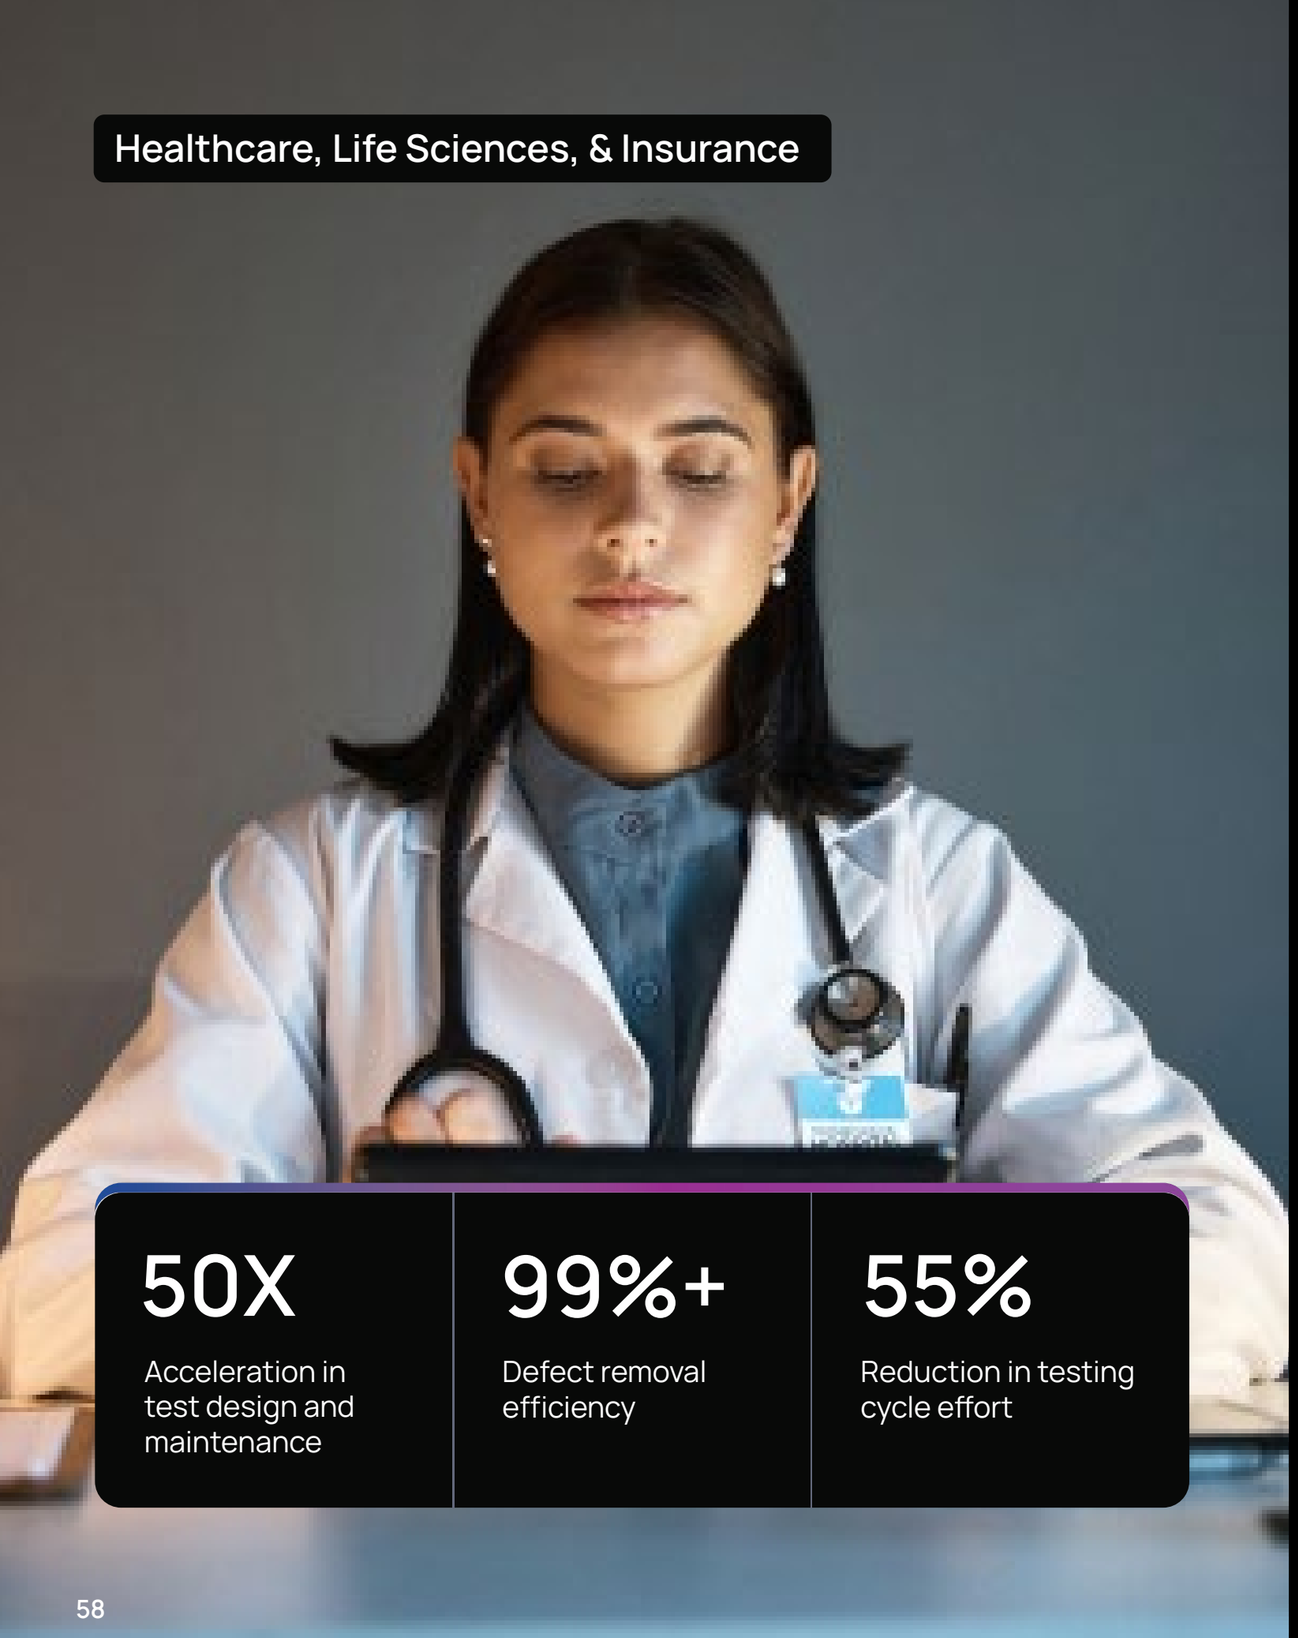

50X	99%+	55%
Acceleration in test design and maintenance	Defect removal efficiency	Reduction in testing cycle effort

Client

A Premier Healthcare Provider In The UK

Accelerating healthcare releases without compromising quality

Manual testing and fragmented workflows were slowing feature releases and putting service reliability at risk across patient-facing systems. The client needed a more autonomous testing model that could keep pace with change while protecting user experience.

Solution

- Delivered a GenAI-powered digital assurance transformation tailored to healthcare workflows.
- Implemented autonomous orchestration for functional and regression testing across API and UI layers.
- Used Tensai® for Autonomous Testing to generate domain-aware test cases from acceptance criteria.
- Established a continuous, end-to-end workflow for high-quality test design and execution.

Benefits

- Achieved up to 50x acceleration in test design and maintenance productivity.
- Raised test coverage to 70–100% across critical healthcare and CRM applications.
- Reduced overall testing lifecycle effort by 55%.
- Improved reliability with 99%+ defect removal efficiency and 100% SLA adherence.

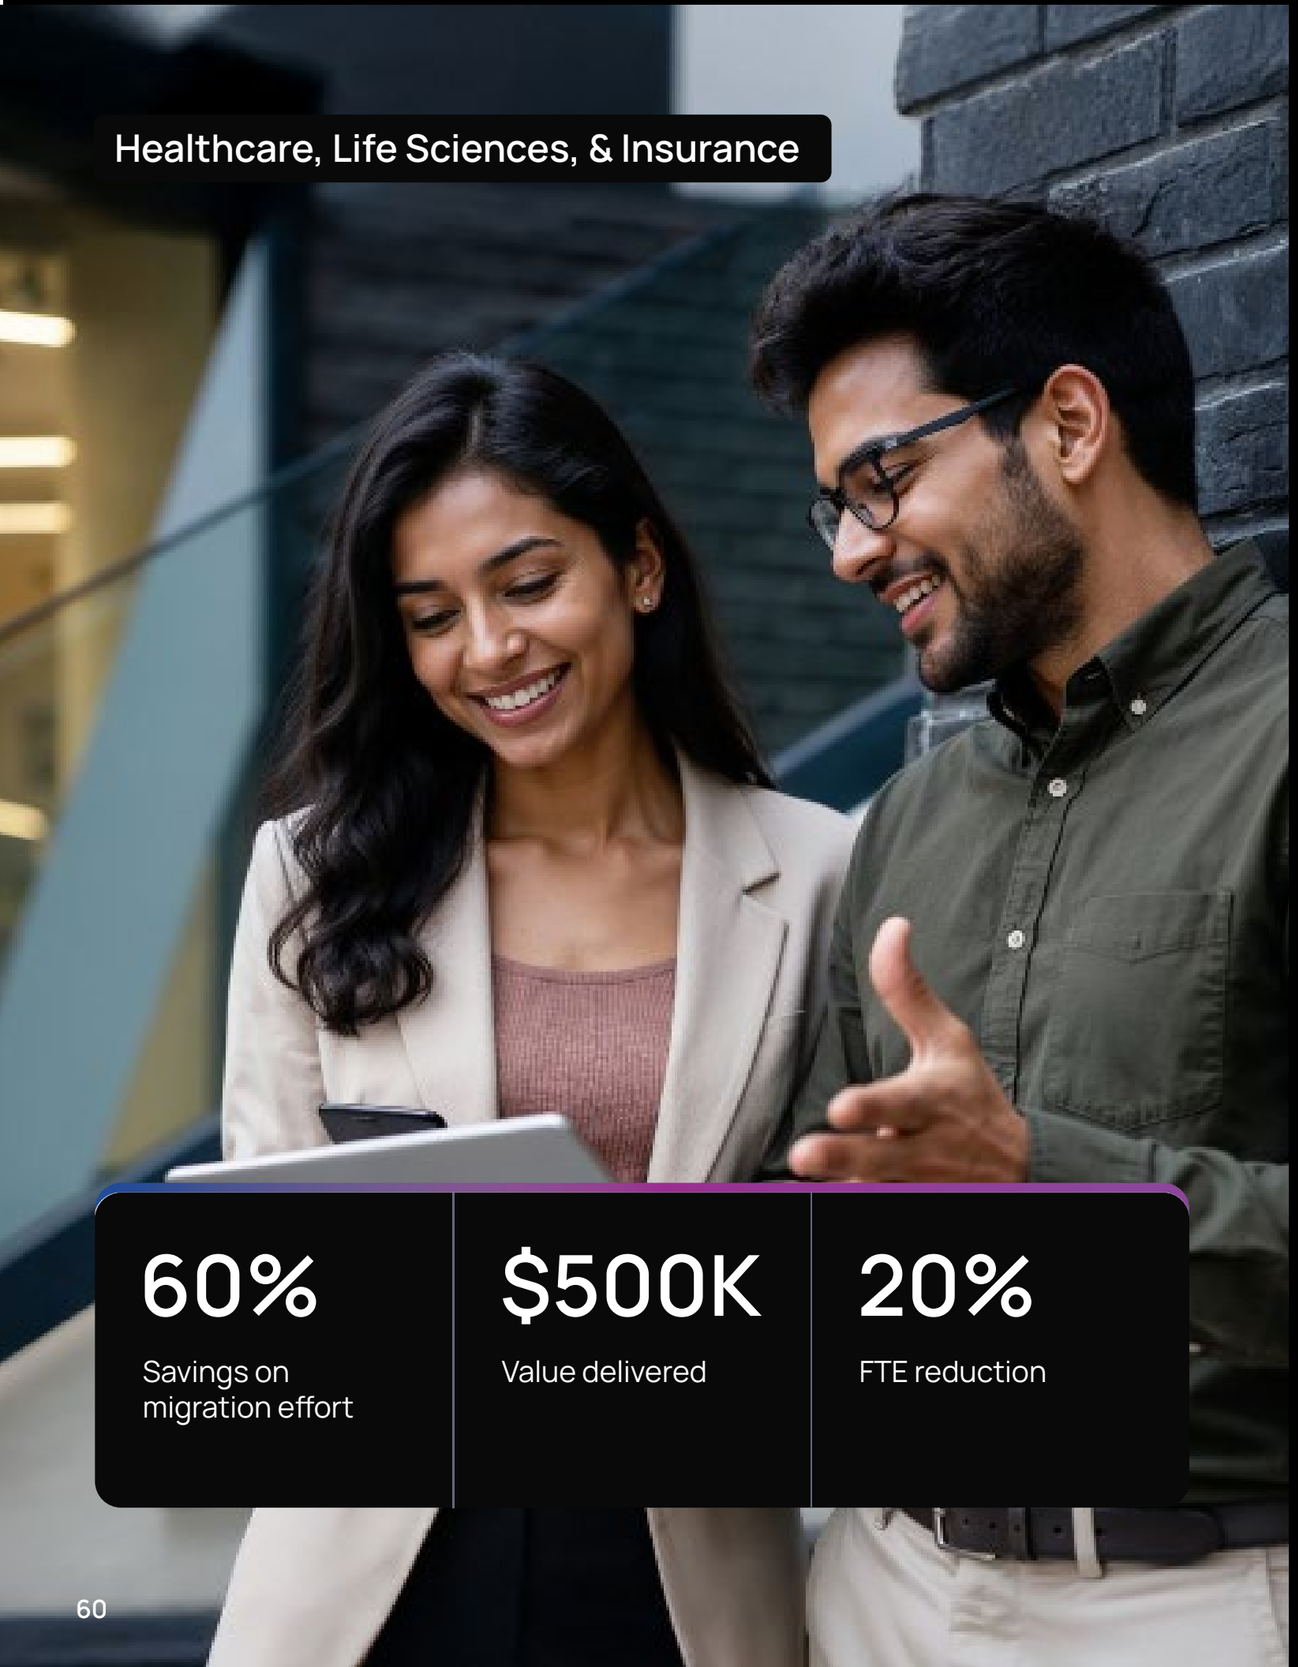

Client

A Leading Belgian Banking & Insurance Group

Modernizing SAS analytics onto Azure with GenAI

The client needed to modernize its legacy SAS analytics onto a scalable cloud platform to cut infrastructure costs and enable big data processing. Migrating to PySpark and Spark SQL on Azure required converting a large SAS program estate without disrupting operations.

Solution

- Migrated SAS programs to PySpark and SparkSQL on Azure.
- Used Mistral AI to accelerate program conversion at scale.
- Focused AI acceleration on simple and medium-complexity programs.
- Re-platformed analytics onto Azure for scalable big-data processing.

Benefits

- Achieved ~60% savings on simple and medium-complexity migration.
- Delivered approximately $500K in measurable value.
- Enabled scalable cloud analytics through generative AI-assisted conversion.

Healthcare, Life Sciences, & Insurance

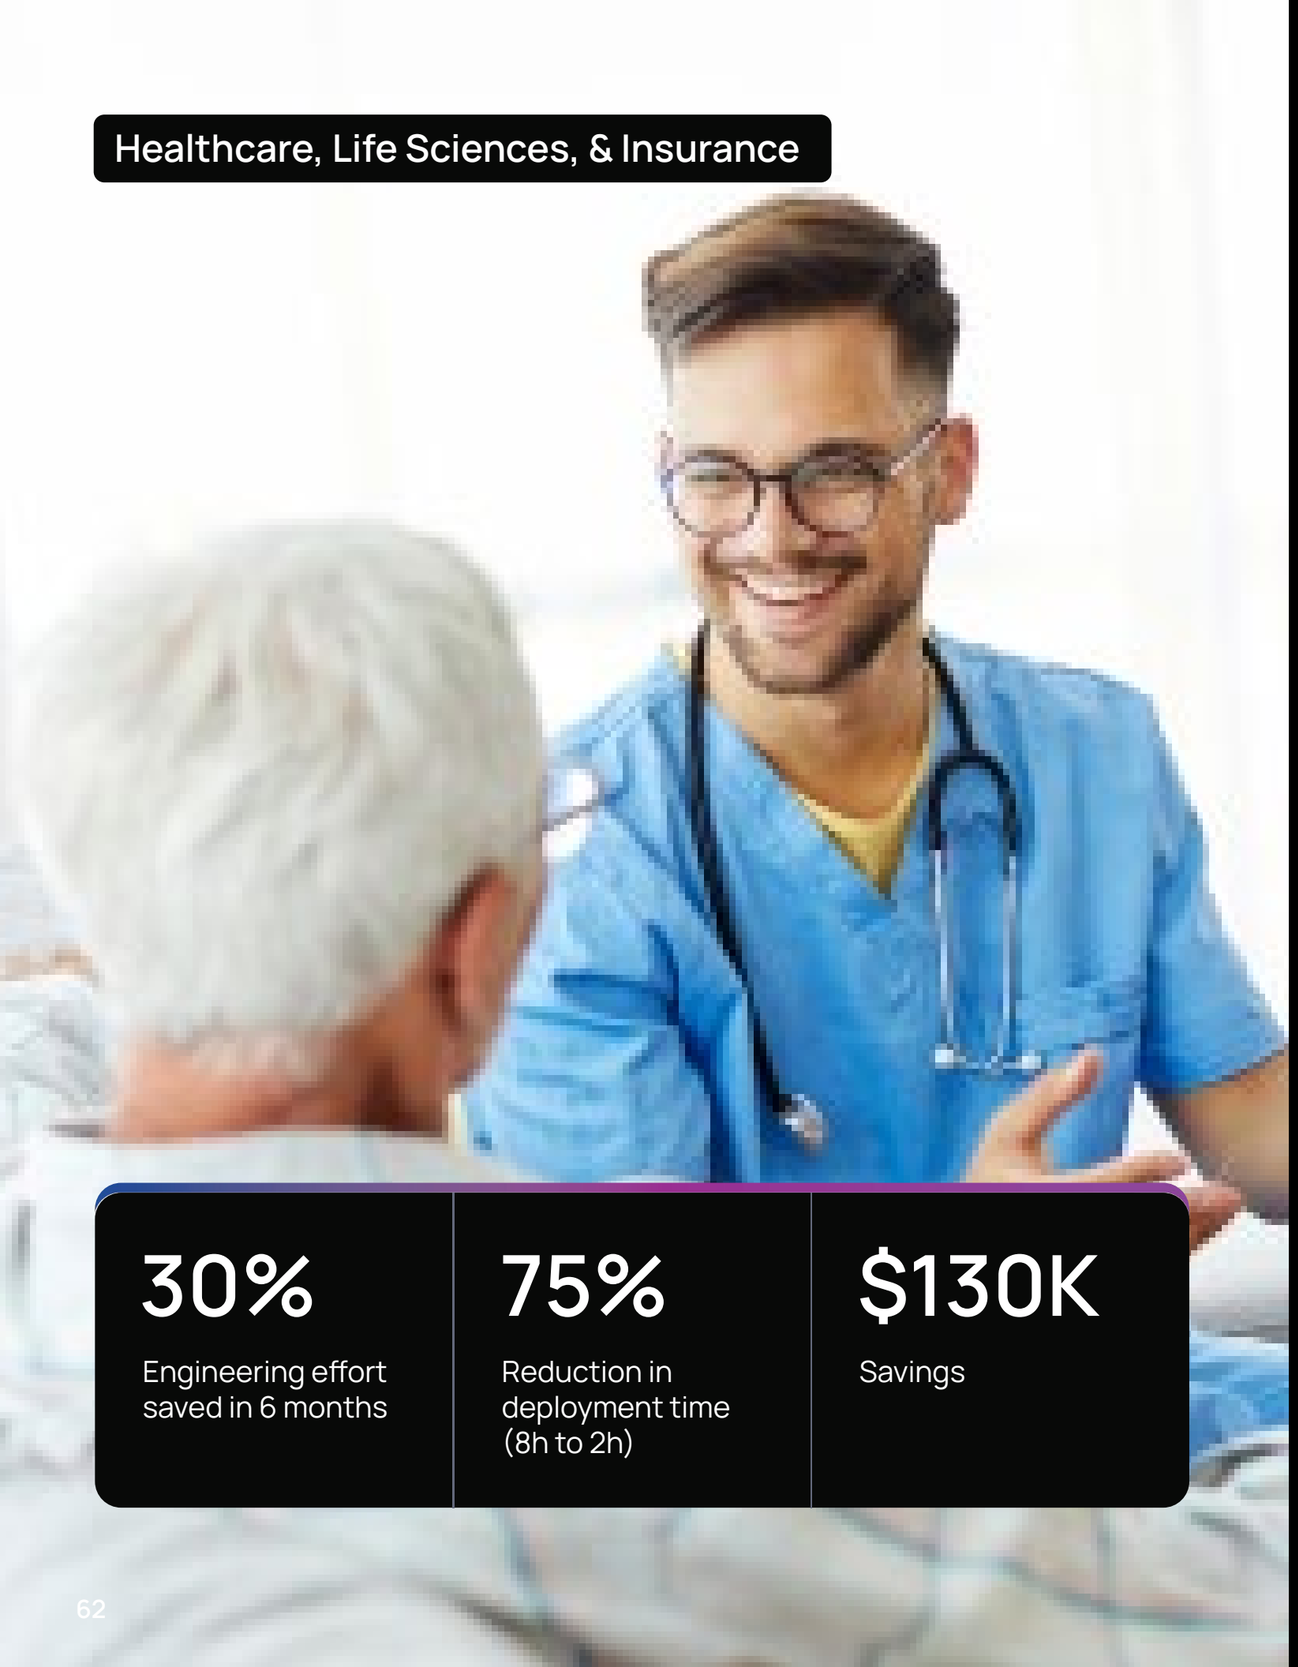

30%	75%	$130K
Engineering effort saved in 6 months	Reduction in deployment time (8h to 2h)	Savings

Client

A Leading UK Healthcare Provider Delivering Integrated Fitness and Private Healthcare Services

Accelerating quality engineering and enterprise application support with AI-driven automation

The client needed to improve software quality and operational efficiency across its QA automation and enterprise application support landscape. Manual testing and lengthy deployment cycles were slowing delivery, while the organization sought greater productivity through AI-enabled automation.

Solution

- Implemented Tensai® for Autonomous Testing, our AI-powered automation platform to streamline quality engineering.
- Enabled Workday AI Agent to accelerate Finance and HR platform deployments.
- Automated testing across identified business applications and critical functionalities.
- Improved engineering productivity by reducing manual testing and deployment effort.

Benefits

- Saved 3,591 engineering hours over six months through AI-assisted QA automation.
- Increased productivity by automating testing and deployment workflows.
- Improved delivery speed and operational efficiency across QA and enterprise applications.

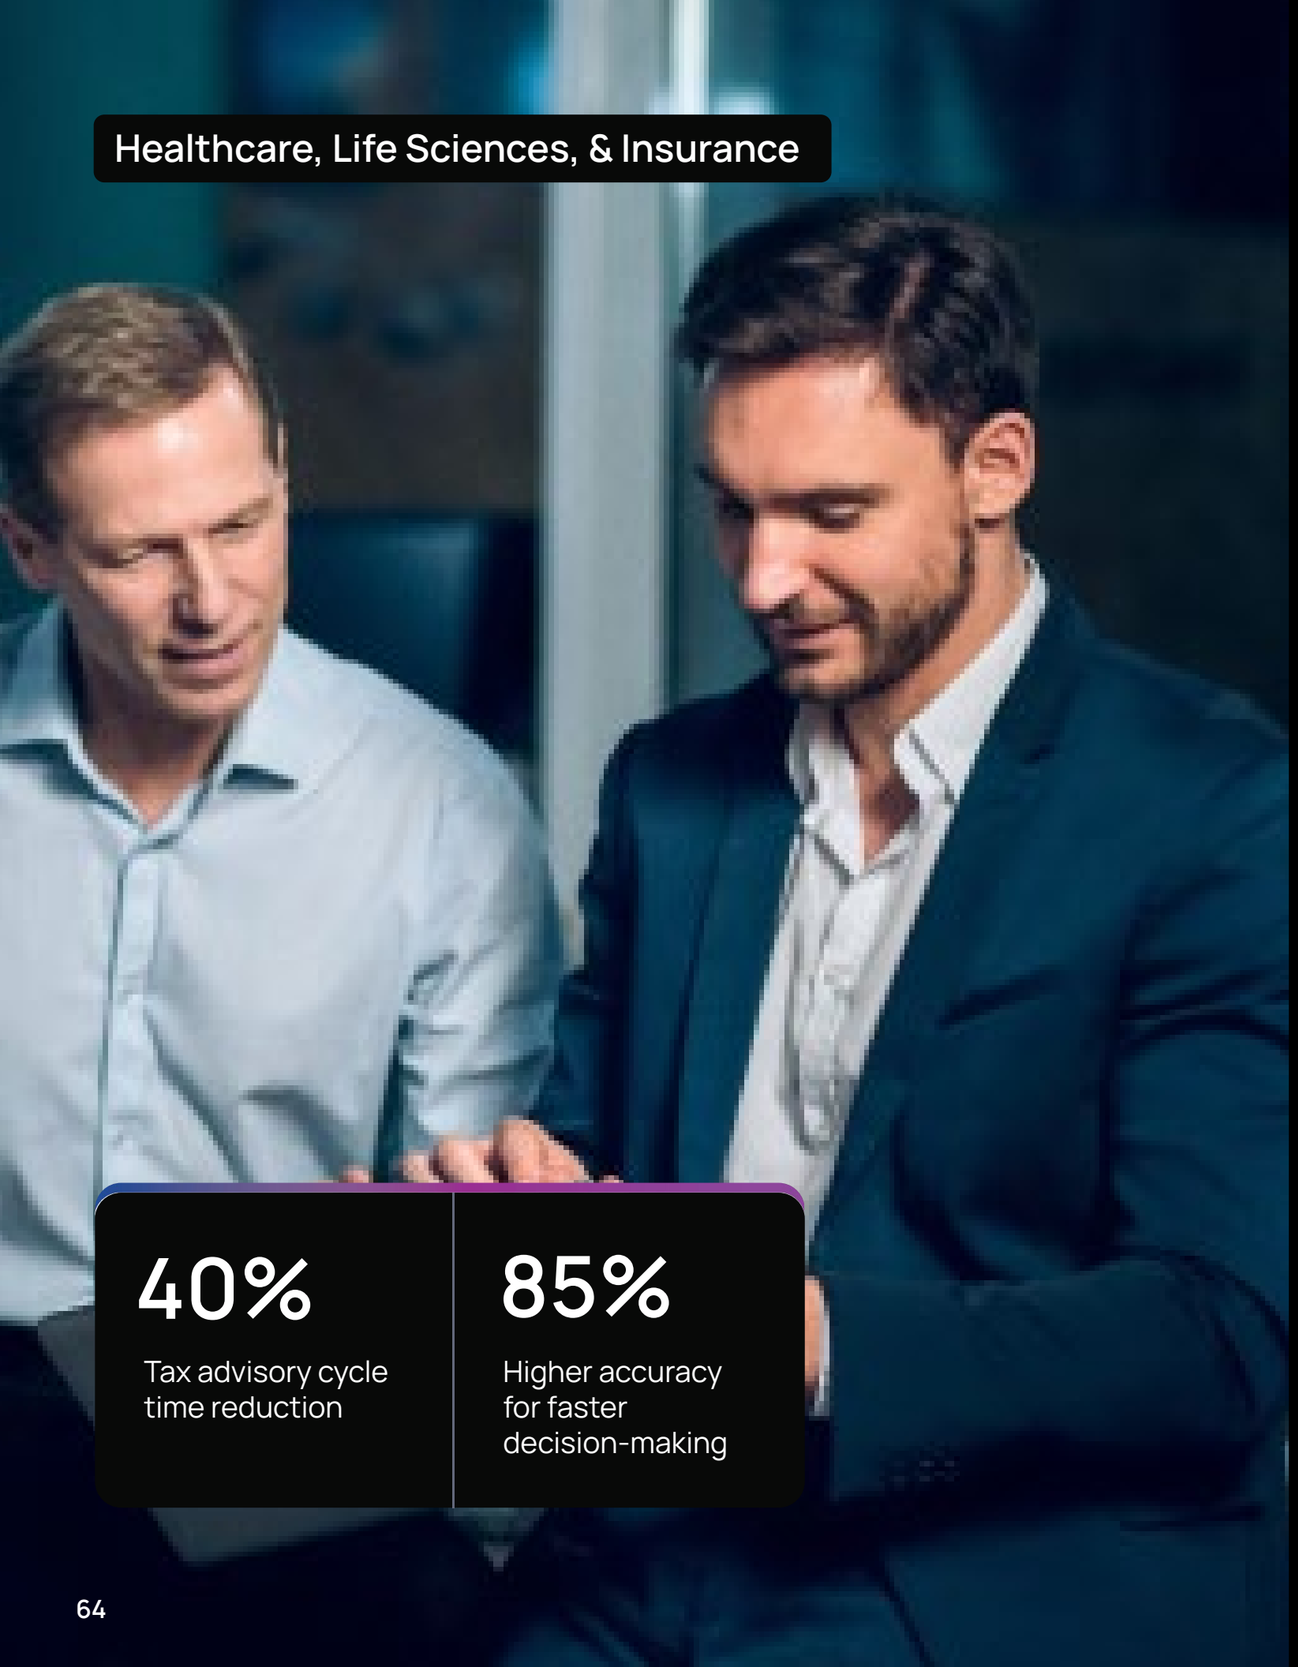

Client

A Global Tax Innovator

Empowering tax advisory with AI to improve accuracy and efficiency

The client faced inefficiencies in manual tax advisory processes, which required extensive human effort as tax knowledge is distributed across multiple sources and advisors. They sought an AI-driven solution to reduce cycle times and improve accuracy.

Solution

- Developed AI components by leveraging ServiceNow AI assistant for tax assessment and advisory processes.
- Leveraged Databricks AI/ML and GenAI capabilities to drive intelligent automation, predictive operations, and enhanced tax consultant experience within ServiceNow.
- Enabled agentic AI to automate tax advisory risk-based change approvals with reduced human involvement.
- Implemented a scalable, efficient system to ensure high alignment and accuracy with consultant decisions.
- Enhanced agent productivity and self-service adoption.

Benefits

- 40% improved process efficiency in overall tax advisory.
- AI solutions aligned 85% with tax consultant decisions.
- Reduced tax advisory cycle time by 40%
- Enhanced knowledge reuse and decision-making consistency.

Client

A Global Care Innovator

Productivity boost with Claude-powered solutions by Hexaware

The client, a global healthcare technology leader, faced inefficiencies in code generation, SQL queries, and test automation, and lacked AI tools to drive productivity.

Solution

- Integrated the Claude generative AI tool for coding, test automation, and documentation.
- Trained consultants on Claude to maximize AI-driven productivity.
- Implemented AI-powered solutions for app development, DevOps, and QA support.
- Delivered measurable productivity improvements through generative AI integration.

Benefits

- Reduced dependency on the domain expert to better understand medical codes through the use of an AI model.
- Enhanced team efficiency with AI tools, reducing manual workload.

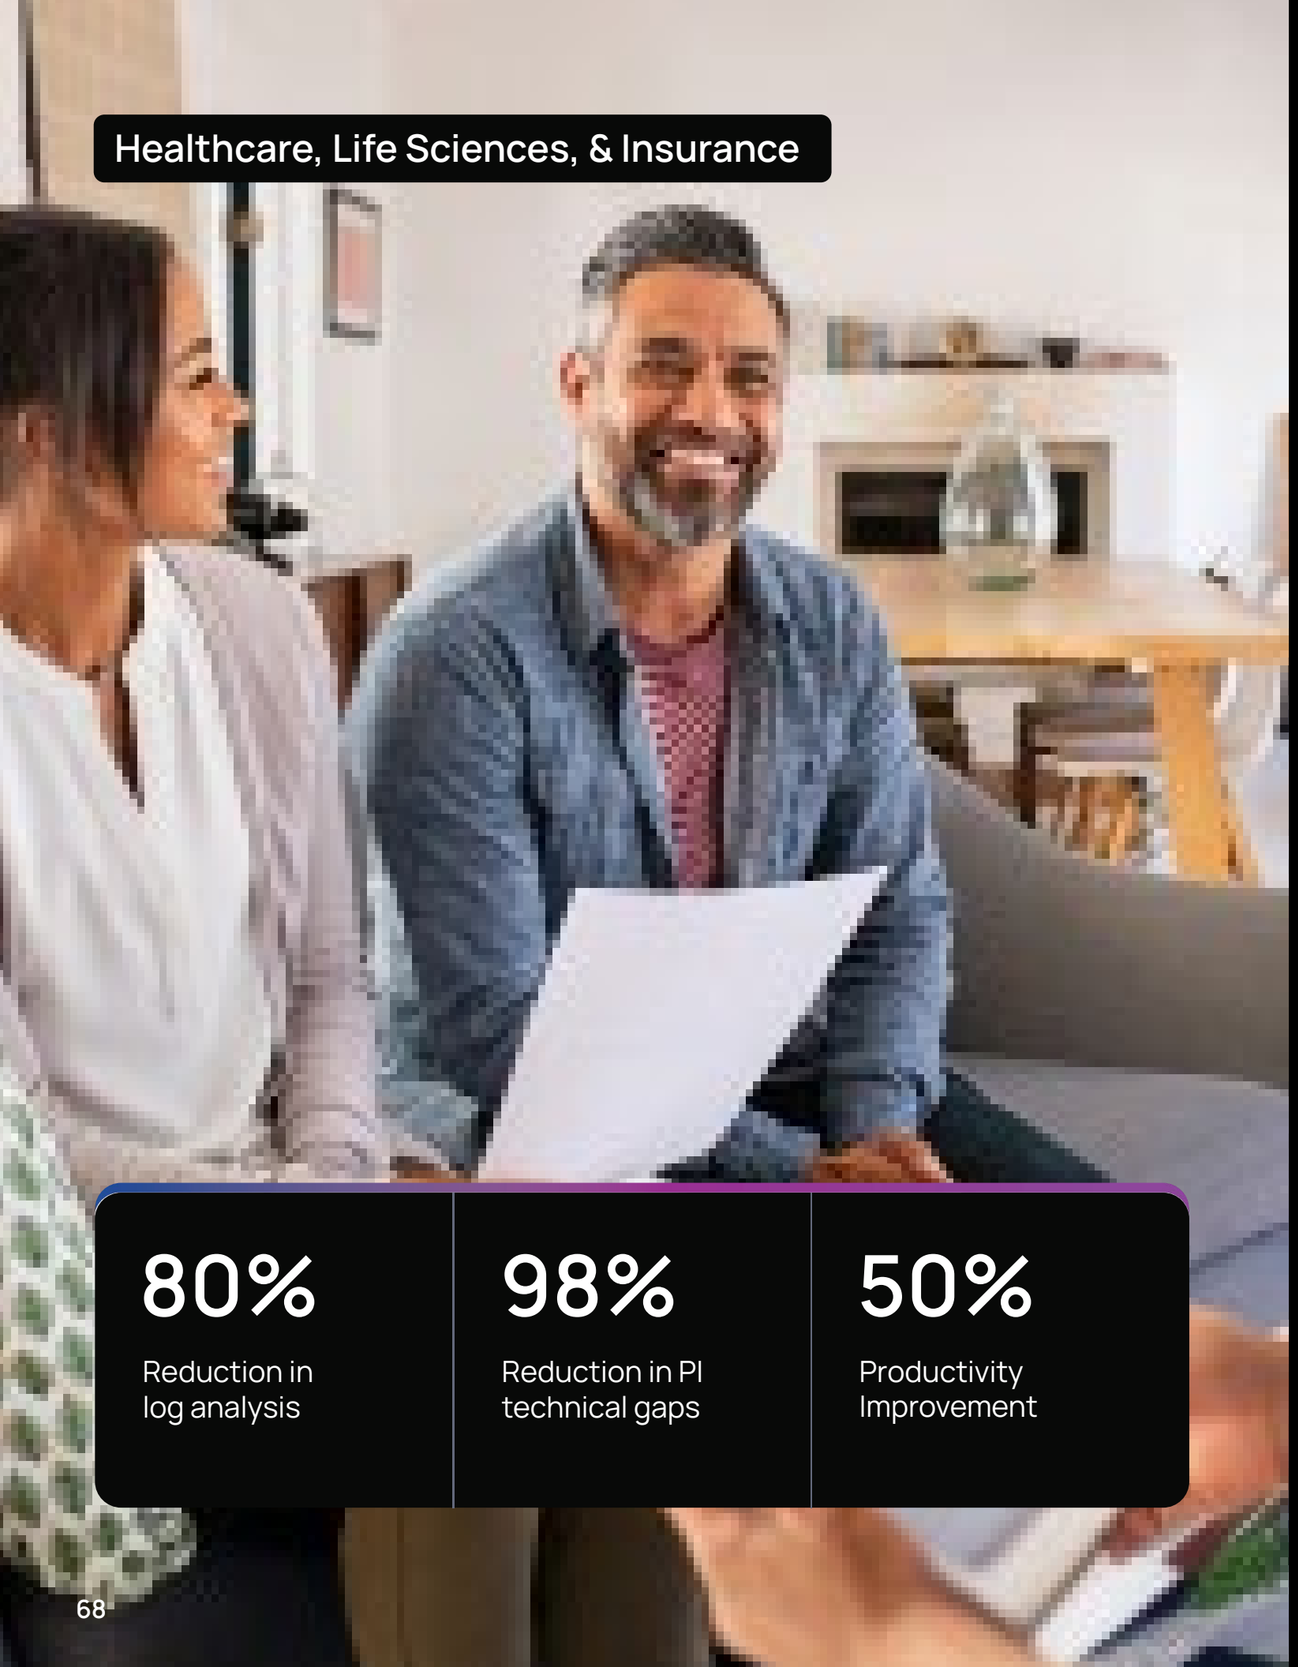

Client

A Leading Life and Non-Life Belgium Insurer

Revolutionizing efficiency with AI-driven solutions by Hexaware

The client faced inefficiencies in log analysis and PI gap analysis, with manual processes consuming extensive time (30 minutes to 4 hours). They aimed to reduce turnaround times and improve operational efficiency.

Solution

- Developed an AI-driven log analysis agent that reduced analysis time from 30 minutes to 5 minutes.
- Designed a PI gap analyzer agent that automated technical gap analysis, cutting time from 4 hours to 5 minutes.
- Leveraged GitHub Copilot for code assistance across the CI/CD development lifecycle.
- Created Proof-of-Concept (POC) AI agents to demonstrate productivity enhancements for future scalability.

Benefits

- Accelerated analysis processes by over 80%, ensuring faster insights and decision-making.
- Enhanced productivity and efficiency by 50% for the CI/CD enhancement processes.
- 98% reduction in technical gaps during the PI process.

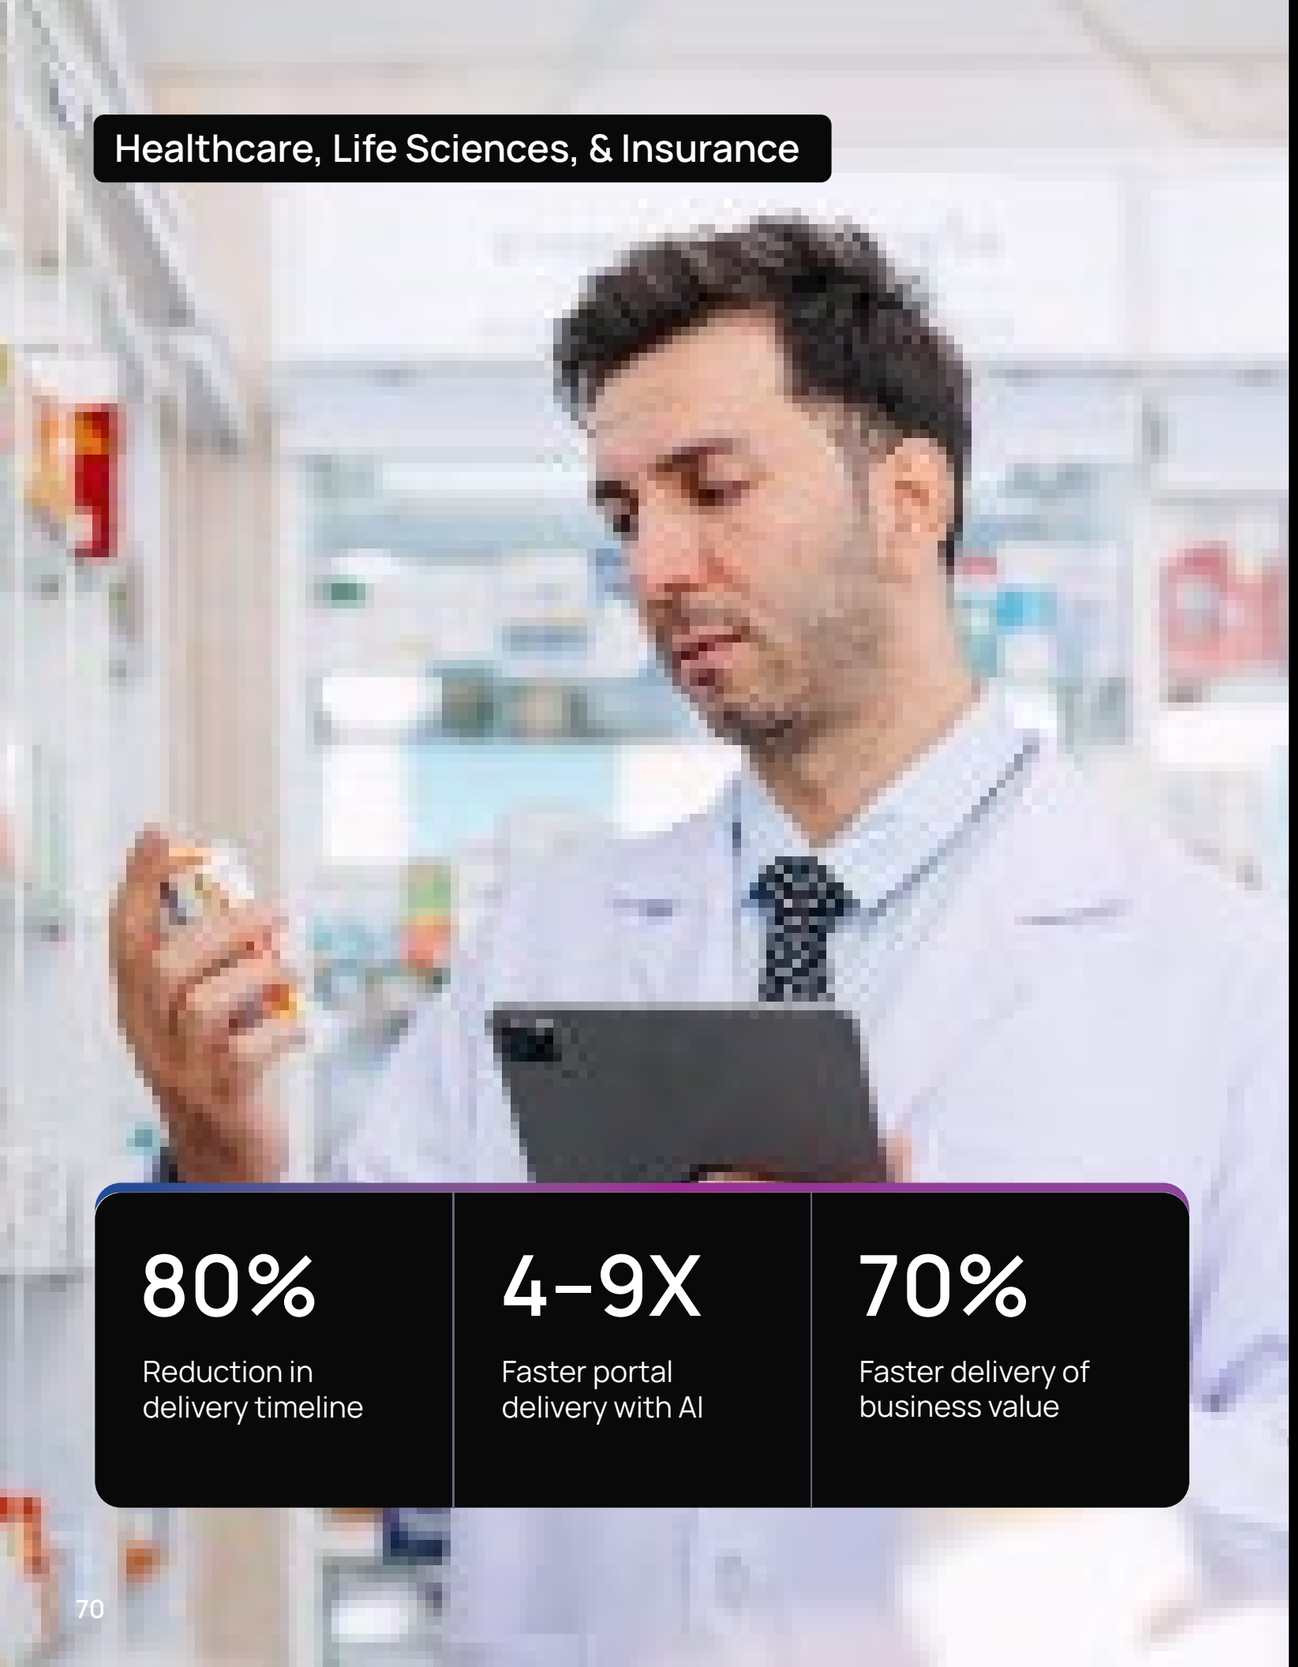

80%	4–9X	70%
Reduction in delivery timeline	Faster portal delivery with AI	Faster delivery of business value

Client

A Leading US Pharmacy Benefit Manager (PBM)

Delivering five production portals with AI-first engineering

The client needed to replace fragmented legacy applications with a modern digital platform while rapidly delivering five production-grade customer portals, digitizing patient enrollment, eligibility verification, and medication fulfillment within aggressive timelines.

Solution

- Used OpenAI Codex and Replit to build five production portals in under three months.
- Applied AI across the SDLC, from requirement definition to development and deployment.
- Digitized the patient journey with registration, eligibility verification, and pharmacy routing.
- Delivered AI-accelerated development aligned to the client's workflows through collaboration.

Benefits

- Delivered five production-ready portals in under three months.
- Reduced delivery timeline by approximately 80%, from 11 months to 3.
- Achieved 4x–9x faster delivery across individual portals with AI-assisted engineering.

Client

A Global Biotech Leader

Shortening study onboarding while strengthening GxP confidence

Manual onboarding, fragmented data systems, and limited test data were slowing trials and putting compliance pressure on the business. The client needed a governed, AI-enabled data foundation that could move faster without cutting corners.

Solution

- Showcased eight live AI agents and a synthetic test data generator during evaluation.
- Consolidated fragmented systems into a unified, GxP-validated Databricks platform.
- Automated study onboarding and SDTM preparation with built-in quality checks.
- Enabled 24/7 intelligent monitoring and AI-first incident triage.

Benefits

- Accelerated study onboarding by 63%, reducing timelines from 6–8 weeks to 2.5 weeks.
- Achieved 100% pipeline test coverage without relying on real patient data.
- Reduced study data tabulation model (SDTM) preparation time by four weeks.
- Optimized cloud costs by more than 30% while maintaining 99%+ support SLA compliance.

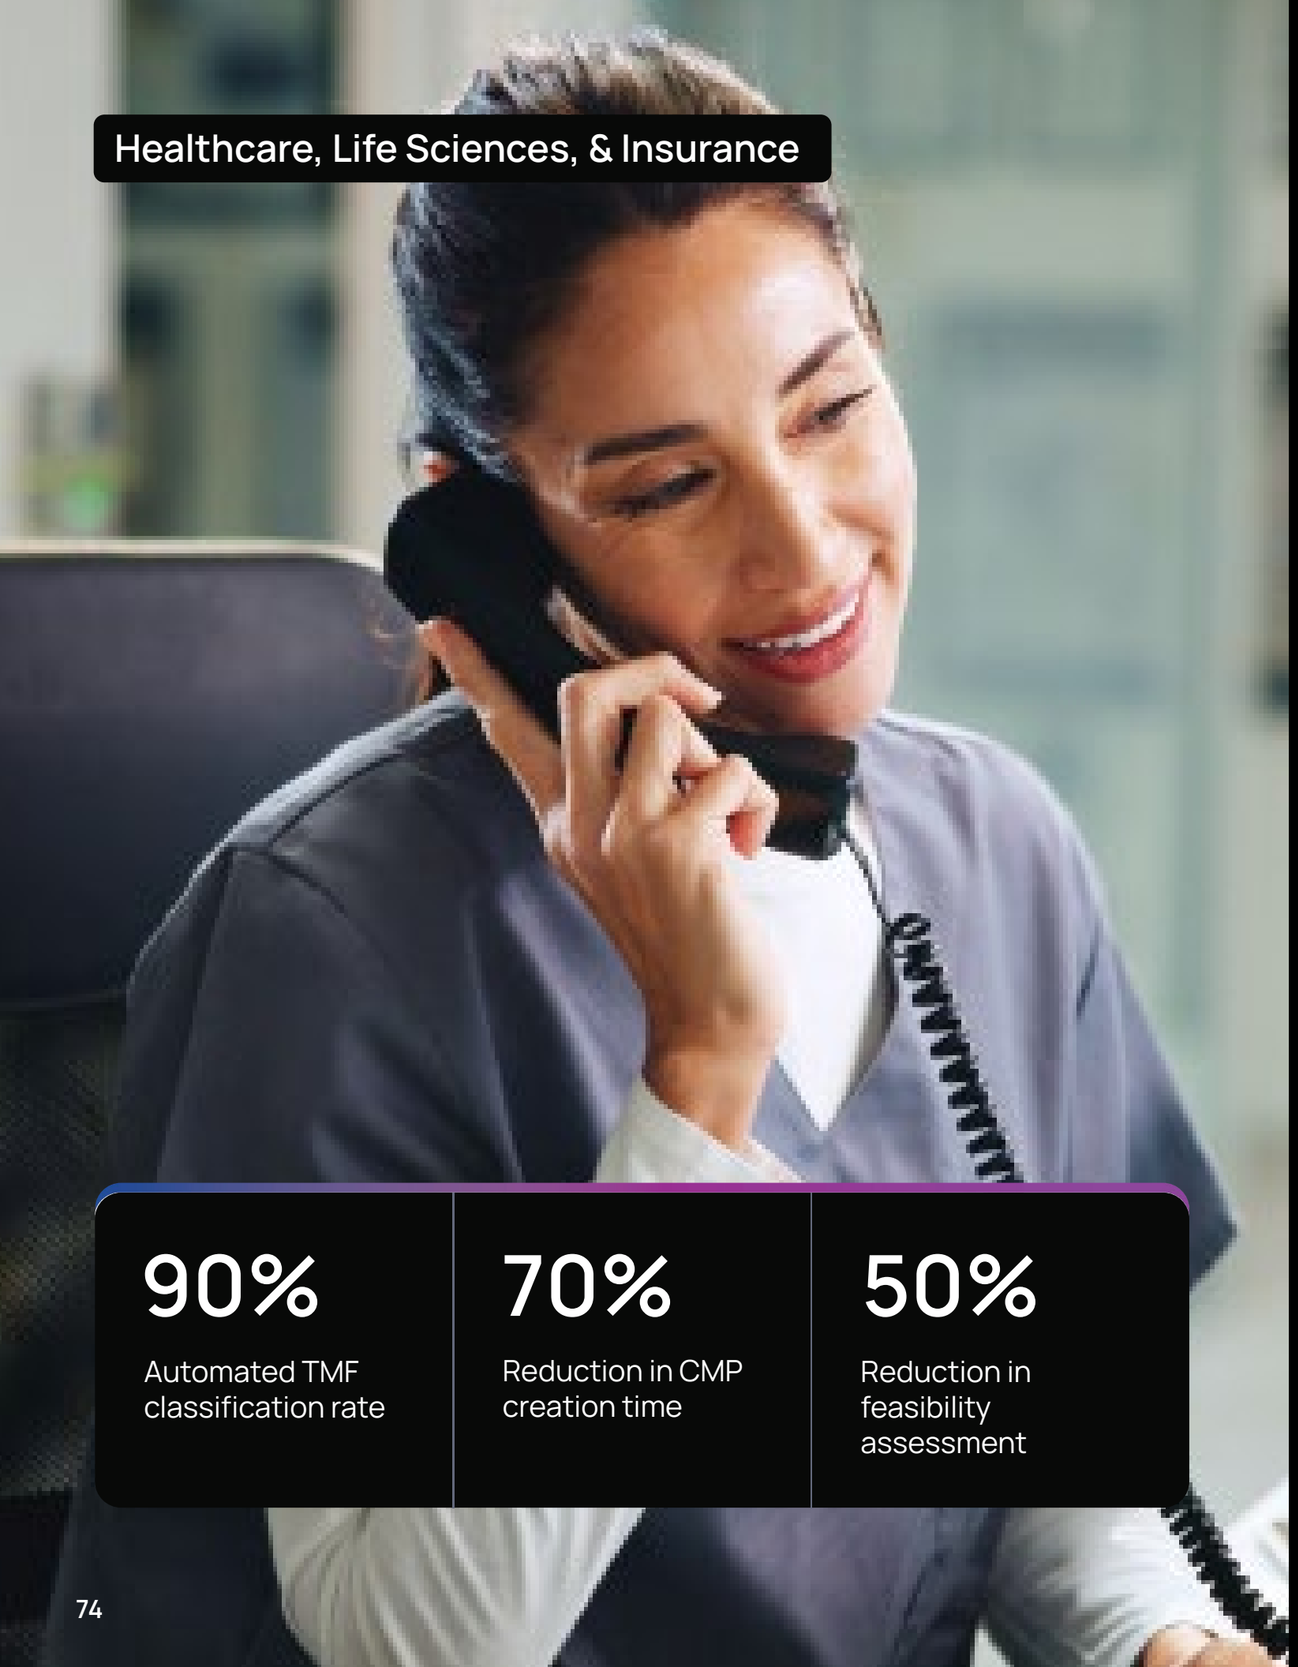

Client

A Full-Service Contract Research Organization (CRO) Specializing In Oncology, Rare Disease, and Central Nervous System (CNS) Trials

Building an AI-native CRO model clients can see and trust

The client wanted to improve operational efficiency while strengthening its market position in an increasingly AI-driven CRO landscape. The need was not just internal productivity, but sponsor-facing differentiation backed by measurable operational gains.

Solution

- Established a three-horizon AI CoE to guide phased transformation.
- Delivered early wins with AI agents focused on high-impact workflows in the first six months.
- Built on Microsoft Azure for ecosystem fit and easier scaling.
- Implemented an AI governance framework and embedded clinical AI expertise within operations.

Benefits

- Automated 90% of trial master file (TMF) classification work.
- Reduced clinical monitoring plan (CMP) creation time by 70%.
- Cut feasibility assessment effort by 50%.
- Strengthened sponsor confidence with a more efficient, AI-native delivery model.

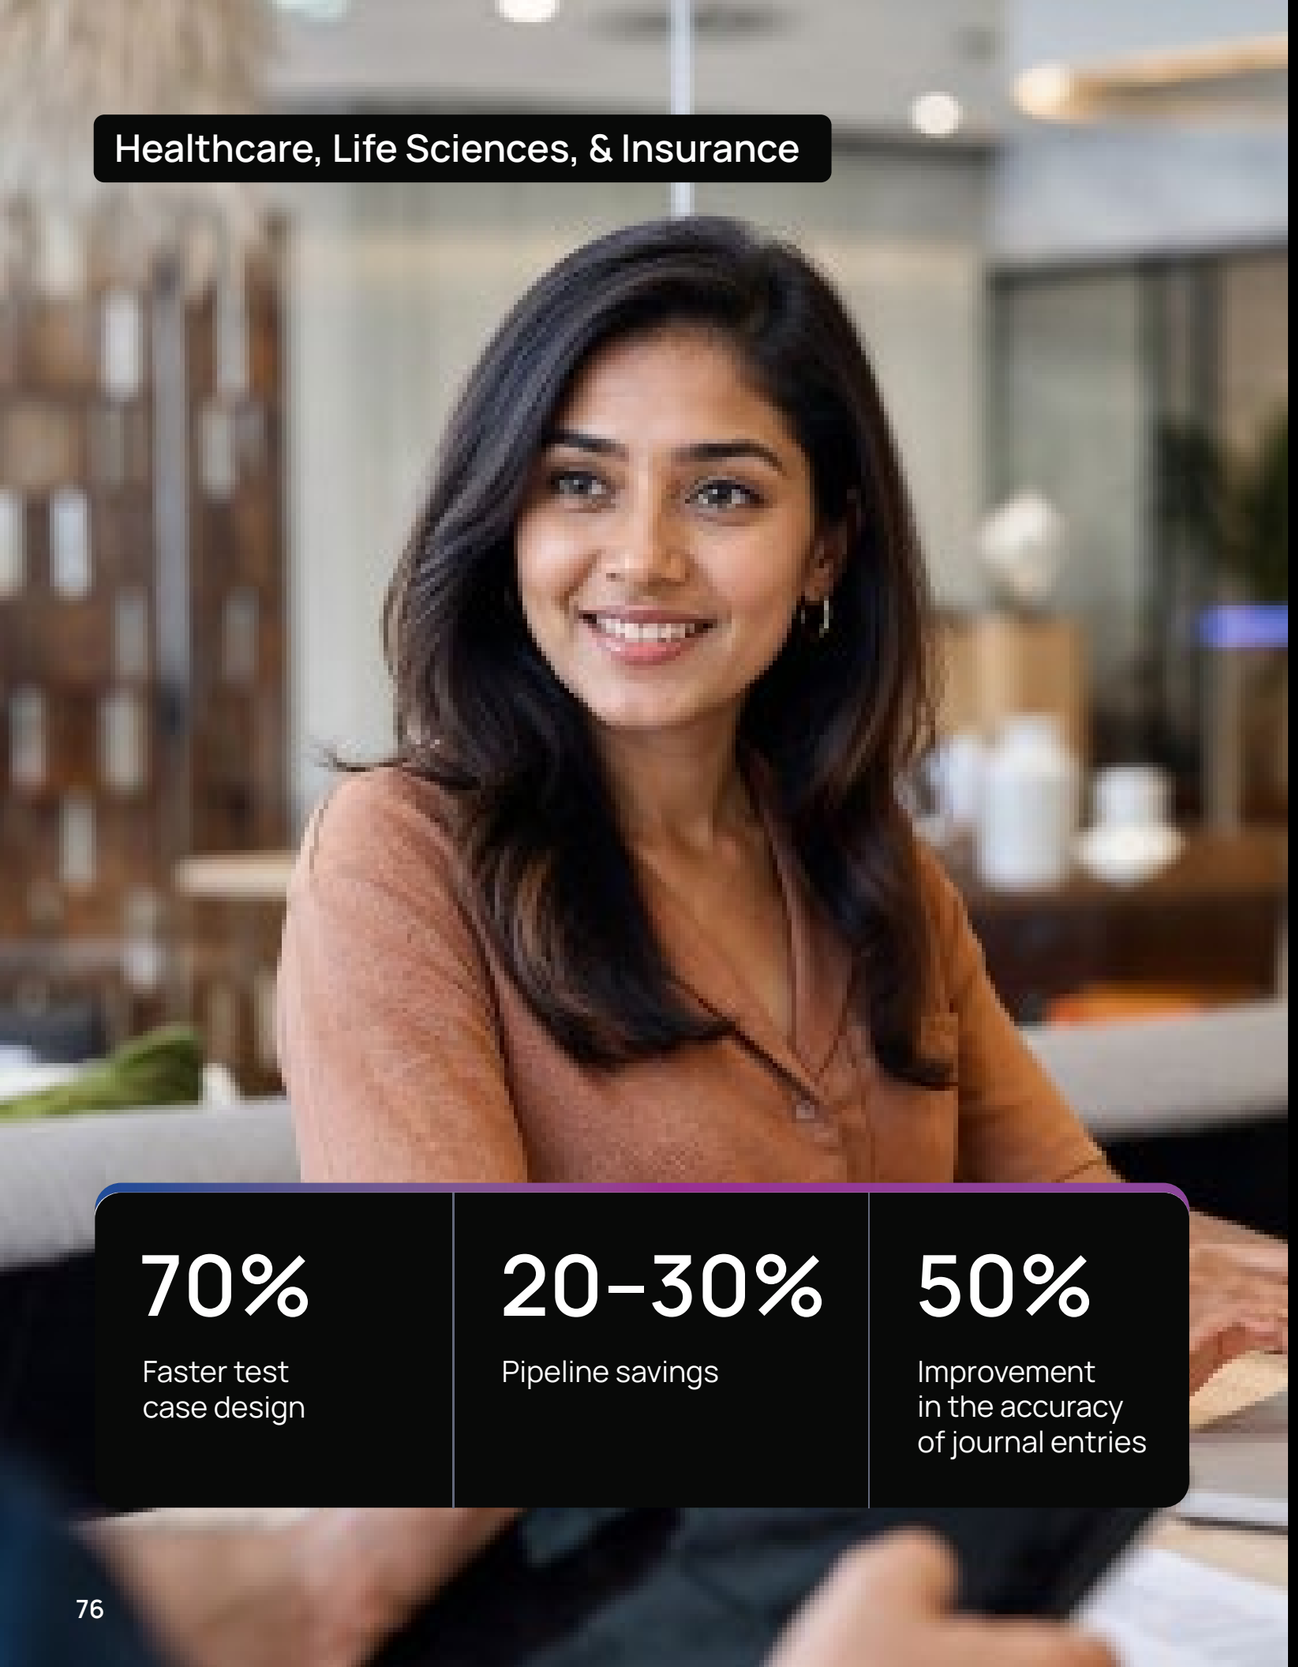

Client

A Leading UK Warranty & Insurance Provider

Scaling AI across infrastructure, testing, digital, and automation

The client, a UK provider of appliance and electronics warranties, insurance, and repair services, ran parallel workstreams — infrastructure support, testing, digital/eCommerce, and process automation — and needed AI embedded across all four to lift productivity.

Solution

- Deployed agentic AI (GitHub Copilot, MS Copilot) to modernize infrastructure support.
- Applied GitHub Copilot to accelerate system test case design across product lines.
- Built a generative AI chatbot on Amazon Bedrock and Claude Sonnet to digitize customer journeys.
- Used Workday GenAI for Journal Insights and Spend categories recommendations.

Benefits

- Accelerated test case design by up to 70% with AI-assisted testing.
- 20–30% pipeline savings across delivery workstreams.
- Reduced manual data entry by 90% for customers & support through AI automation.
- 50% improvement in journal entry accuracy.

Healthcare, Life Sciences, & Insurance

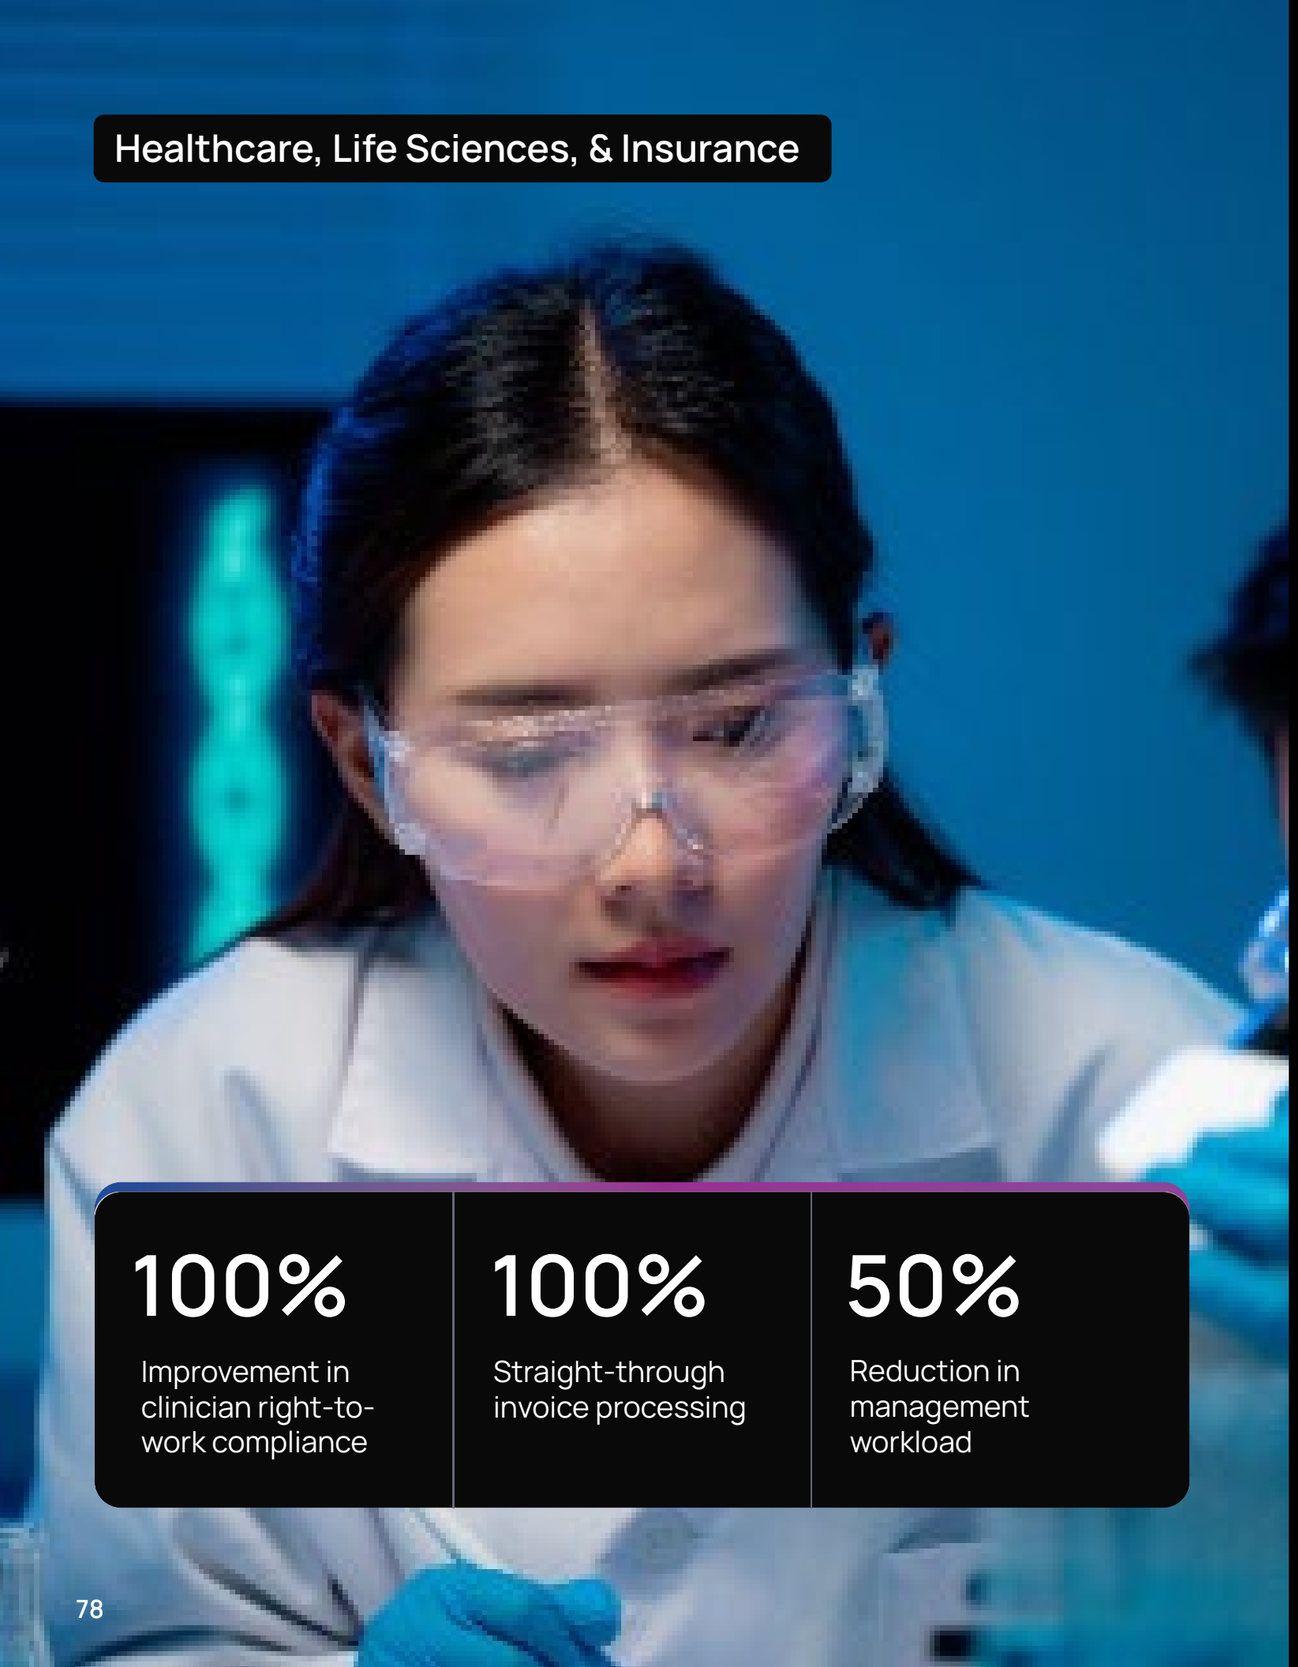

100%	100%	50%
Improvement in clinician right-to-work compliance	Straight-through invoice processing	Reduction in management workload

Client

A UK Healthcare Provider Delivering a Digital-First Blended Care Model

Helping healthcare operations scale when clinician capacity is under pressure

Clinician shortages, patient backlogs, and fragmented systems were making it harder to meet demand and maintain compliance. The client needed workforce operations that could move faster, stay governed, and reduce administrative strain.

Solution

- Deployed an AI-powered onboarding system to accelerate clinician integration.
- Automated compliance checks, timesheets, and rota scheduling.
- Added AI-driven dashboards for real-time governance and operational visibility.
- Integrated cloud-based platforms to support smoother NHS collaboration.

Benefits

- Improved clinician right-to-work compliance by 100%.
- Reduced management workload by 50%.
- Achieved 100% straight-through invoice processing.
- Accelerated onboarding and reduced administrative burden, so more effort could stay focused on care delivery.

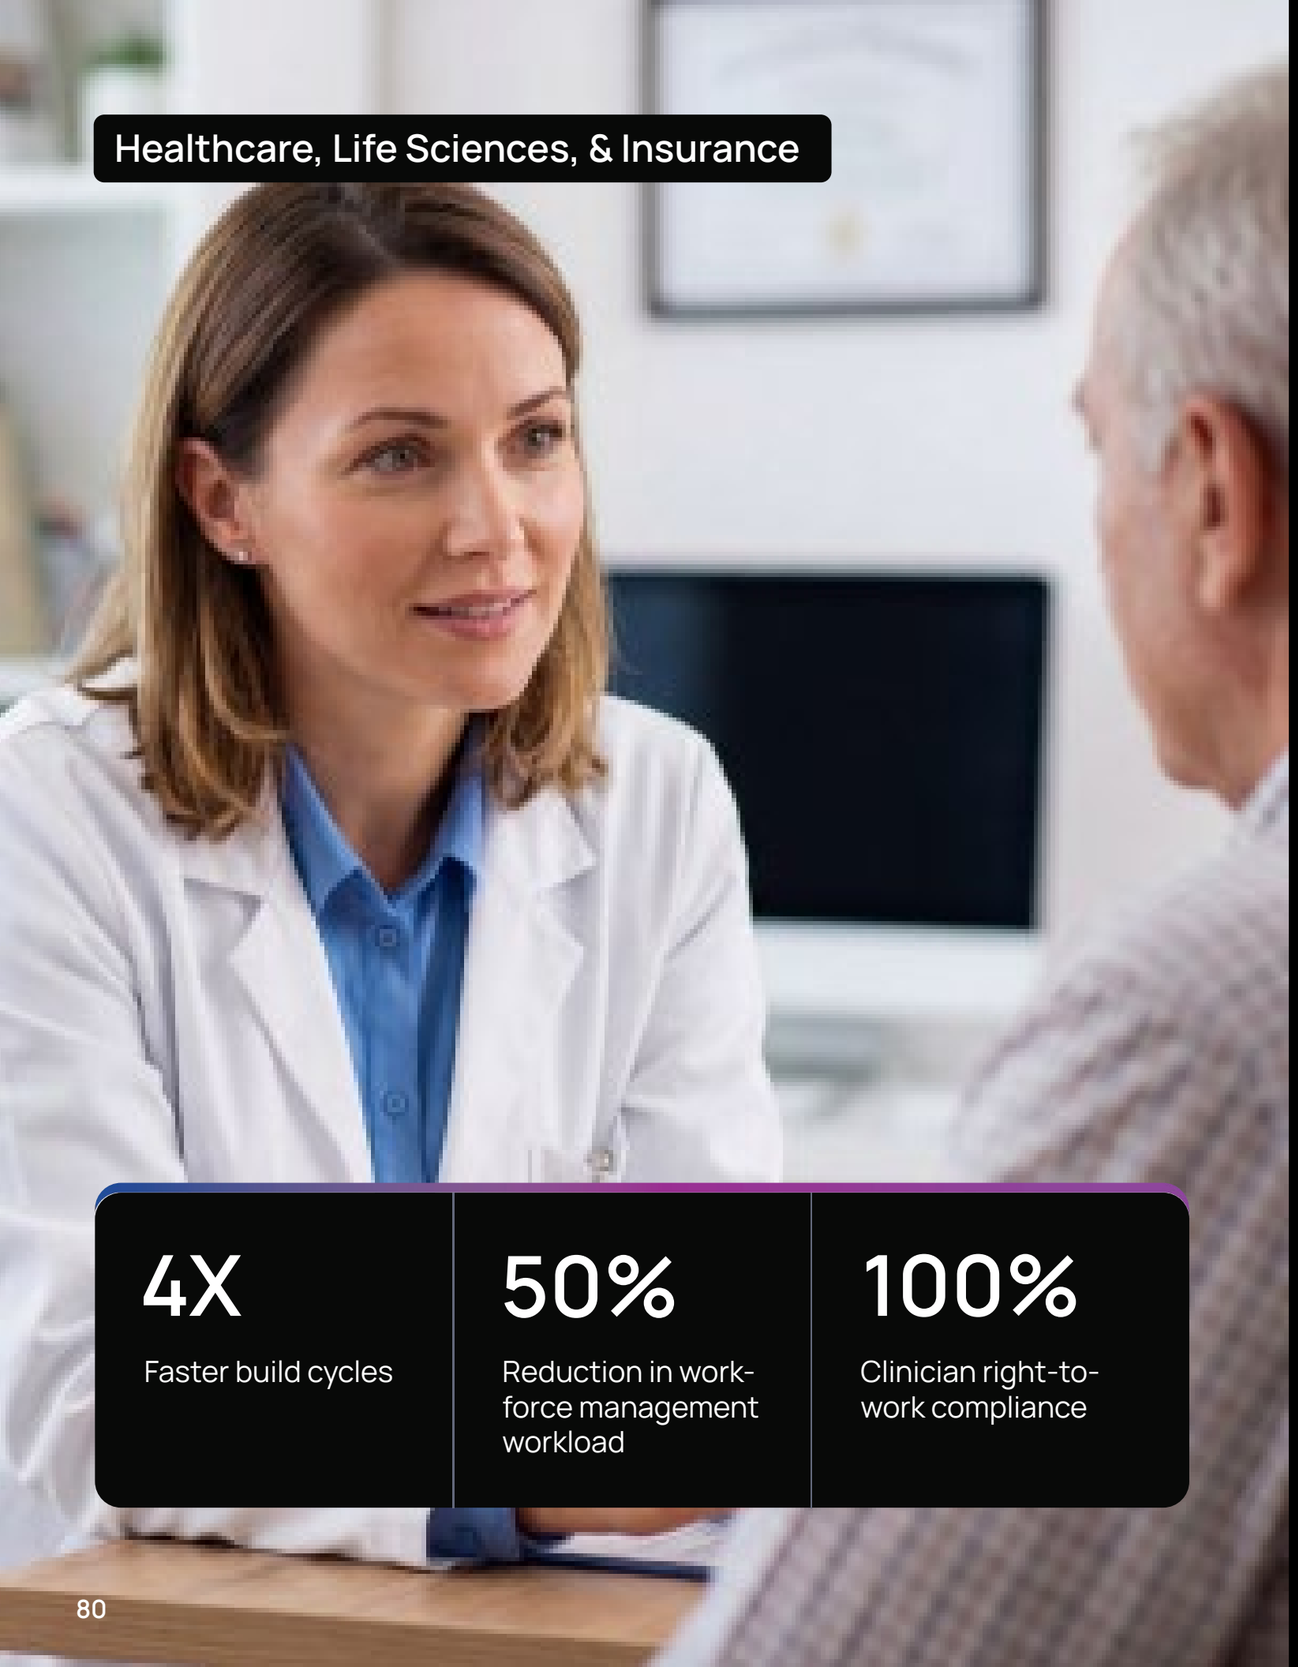

Client

A Healthcare Provider Supporting NHS Capacity

Accelerating clinician onboarding with an AI-powered low-code platform

Addressed manual-heavy processes across clinician recruitment, onboarding, rota management, contact management, compliance tracking, and patient-facing engagement, where fragmented data limited operational visibility.

Solution

- Conducted design-led discovery to map current journeys, identify pain points, define future-state service blueprints, and create an implementation roadmap for AI-led transformation.
- Built and validated functional prototypes using AI models, historic data, and rapid proof-of-value cycles before scaling the solution into a secure, production-ready platform.
- Leveraged Journey36's cloud-based App36 platform, Axonyx.ai governance capabilities, and Hexaware's delivery model to accelerate development while supporting responsible AI use in healthcare operations.
- Delivered a low-code/no-code digital platform that gives the client a single source of truth for clinician onboarding, compliance, right-to-work checks, and operational visibility.

Benefits

- Achieved 4x faster build cycles.
- Reduced workforce management workload by 50%.
- Reached 100% clinician right-to-work compliance.

8%
Improvement in loss ratio

100+
New causes of loss discovered

33%
Reduction in underwriting effort

Client

A Global Marine Insurance Company

Sharper underwriting decisions start with better risk signals

Manual underwriting judgment and fragmented historical data were limiting loss prediction accuracy. The client needed deeper pattern recognition to price risk more effectively and improve portfolio performance.

Solution

- Analyzed 16 years of text-based claims and loss data within an AI analytics framework.
- Identified 68 leading causes of loss and hundreds of previously hidden categories.
- Built six marine-specific taxonomies to improve natural language understanding.
- Deployed an AI model that predicts claim patterns for new vessels to support underwriting decisions.

Benefits

- Improved loss ratio by 8%.
- Reduced underwriting effort by 33% after one year.
- Discovered 100+ new causes of loss to improve risk understanding.
- Enabled more informed pricing and underwriting decisions for new business.

90%
KT sessions documented automatically

70%
Reduction in effort to locate approved clinical trial specifications

25%
Effort saved generating synthetic clinical trial test data with AI

Client

A Leading Global Biotechnology Company Developing Innovative Immunotherapies

Standardizing knowledge transfer documentation for clinical data operations

The client needed to improve the consistency and efficiency of knowledge transfer across its clinical data engineering teams while operating in a highly regulated GxP environment. Manual KT documentation was time-consuming, and compliance restrictions limited the adoption of AI tools.

Solution

- Implemented Hexaware's Tensai® GPT to automate knowledge transfer documentation.
- Leveraged Databricks on AWS with Terraform to support scalable clinical data engineering and operations.
- Standardized KT outputs into a consistent, reusable format.
- Reduced manual effort while enabling AI-assisted documentation within the client's GxP compliance requirements.
- Developed clinical data platform knowledge estate from fragmented documents into a governed, searchable enterprise asset.

Benefits

- Automated 90% of knowledge transfer session documentation.
- Generated synthetic clinical trials test data, delivering 25% effort savings for business users.
- Reduced search effort by 70% for locating approved specifications and clinical study artifacts.

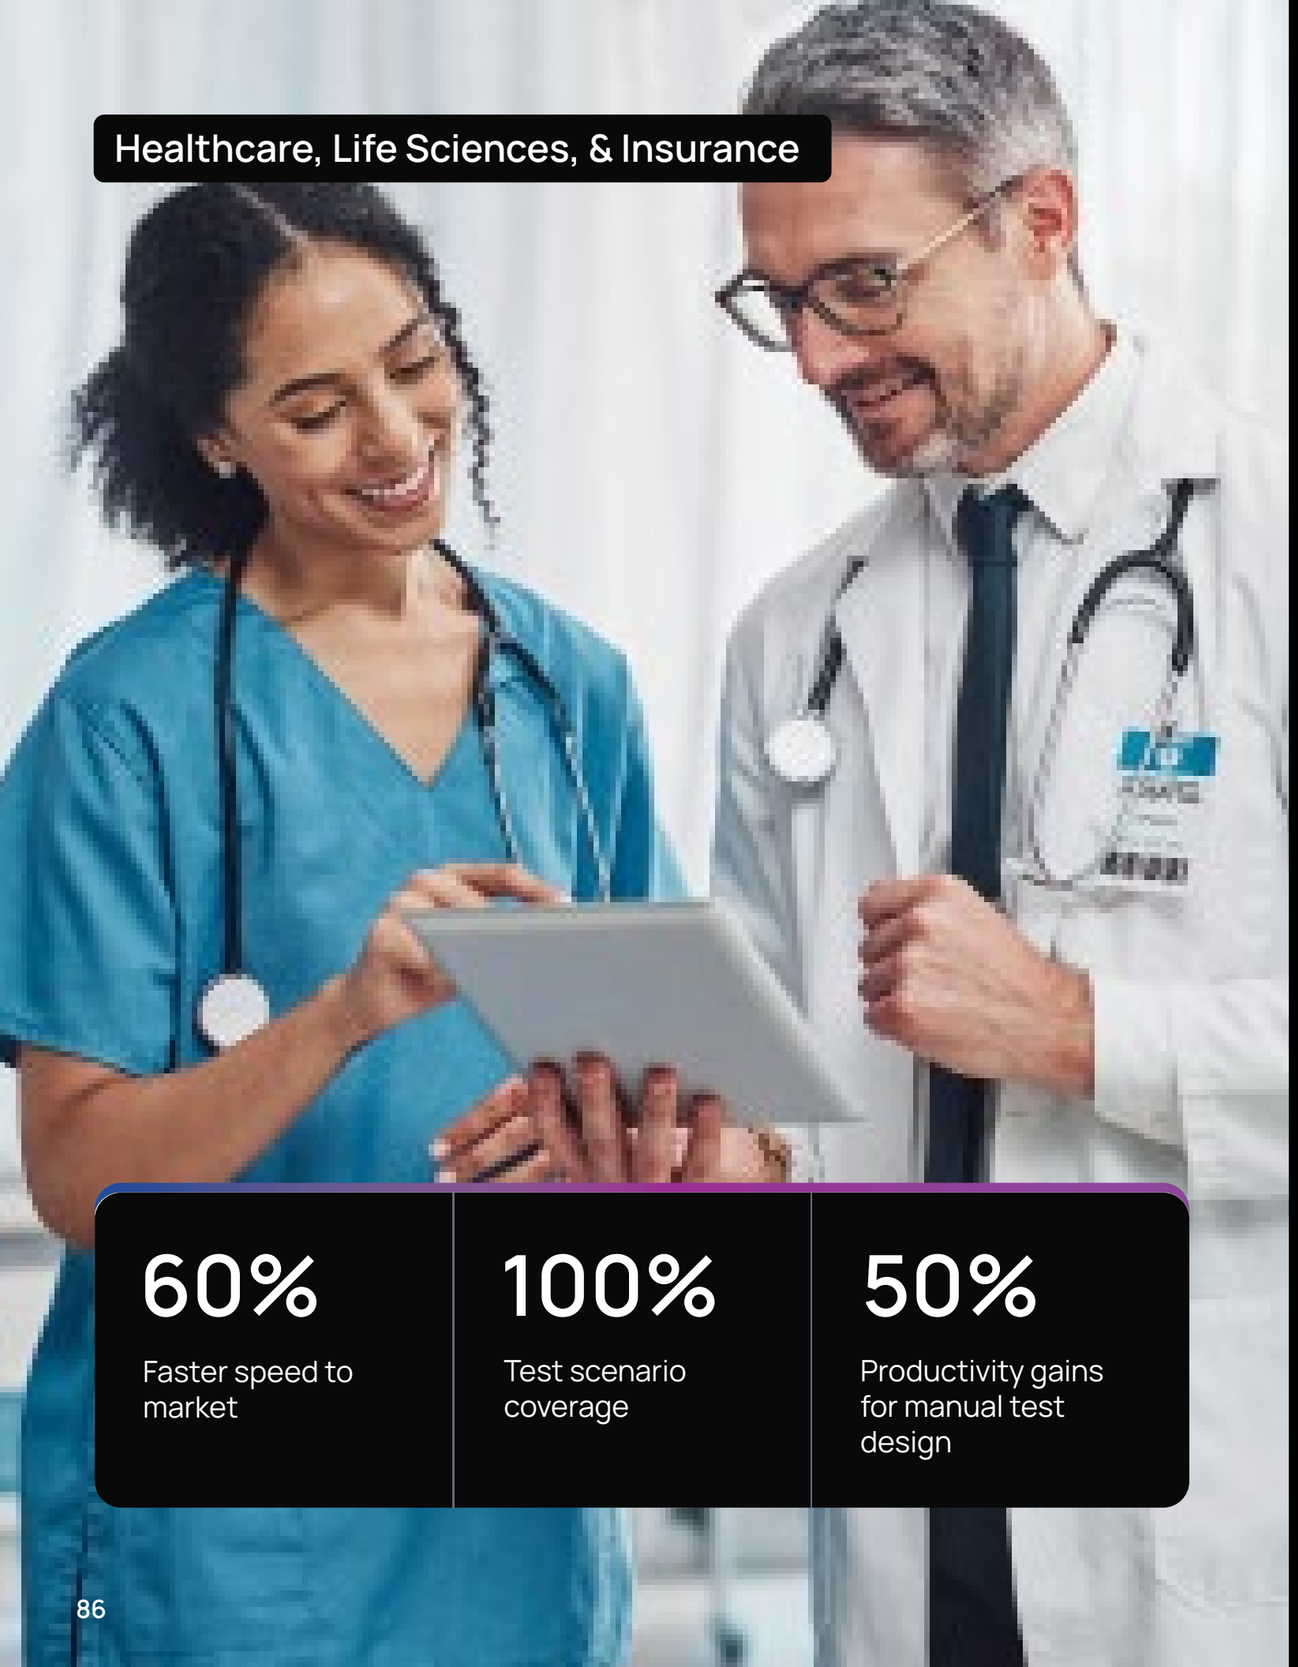

Client

American Health Information and Clinical Research Company

Using GenAI to clear the path from support to faster release readiness

The client wanted a self-service support experience that would reduce friction around troubleshooting and accelerate adoption of GenAI in enterprise testing. The result was a support model that helped teams move faster while protecting data privacy.

Solution

- Developed a GenAI-powered self-service support solution that used the client's existing knowledge base.
- Assessed feasibility and designed a scalable model aligned to the client's enterprise architecture.
- Used Tensai® to accelerate adoption and address data privacy concerns in testing workflows.
- Built a pipeline of 12 validated use cases in six weeks.

Benefits

- Improved speed to market by 60%.
- Achieved 100% coverage for test scenarios.
- Delivered 50% productivity gains for manual test design.
- Reduced ticket volume by 15%, allowing teams to focus on more complex issues.

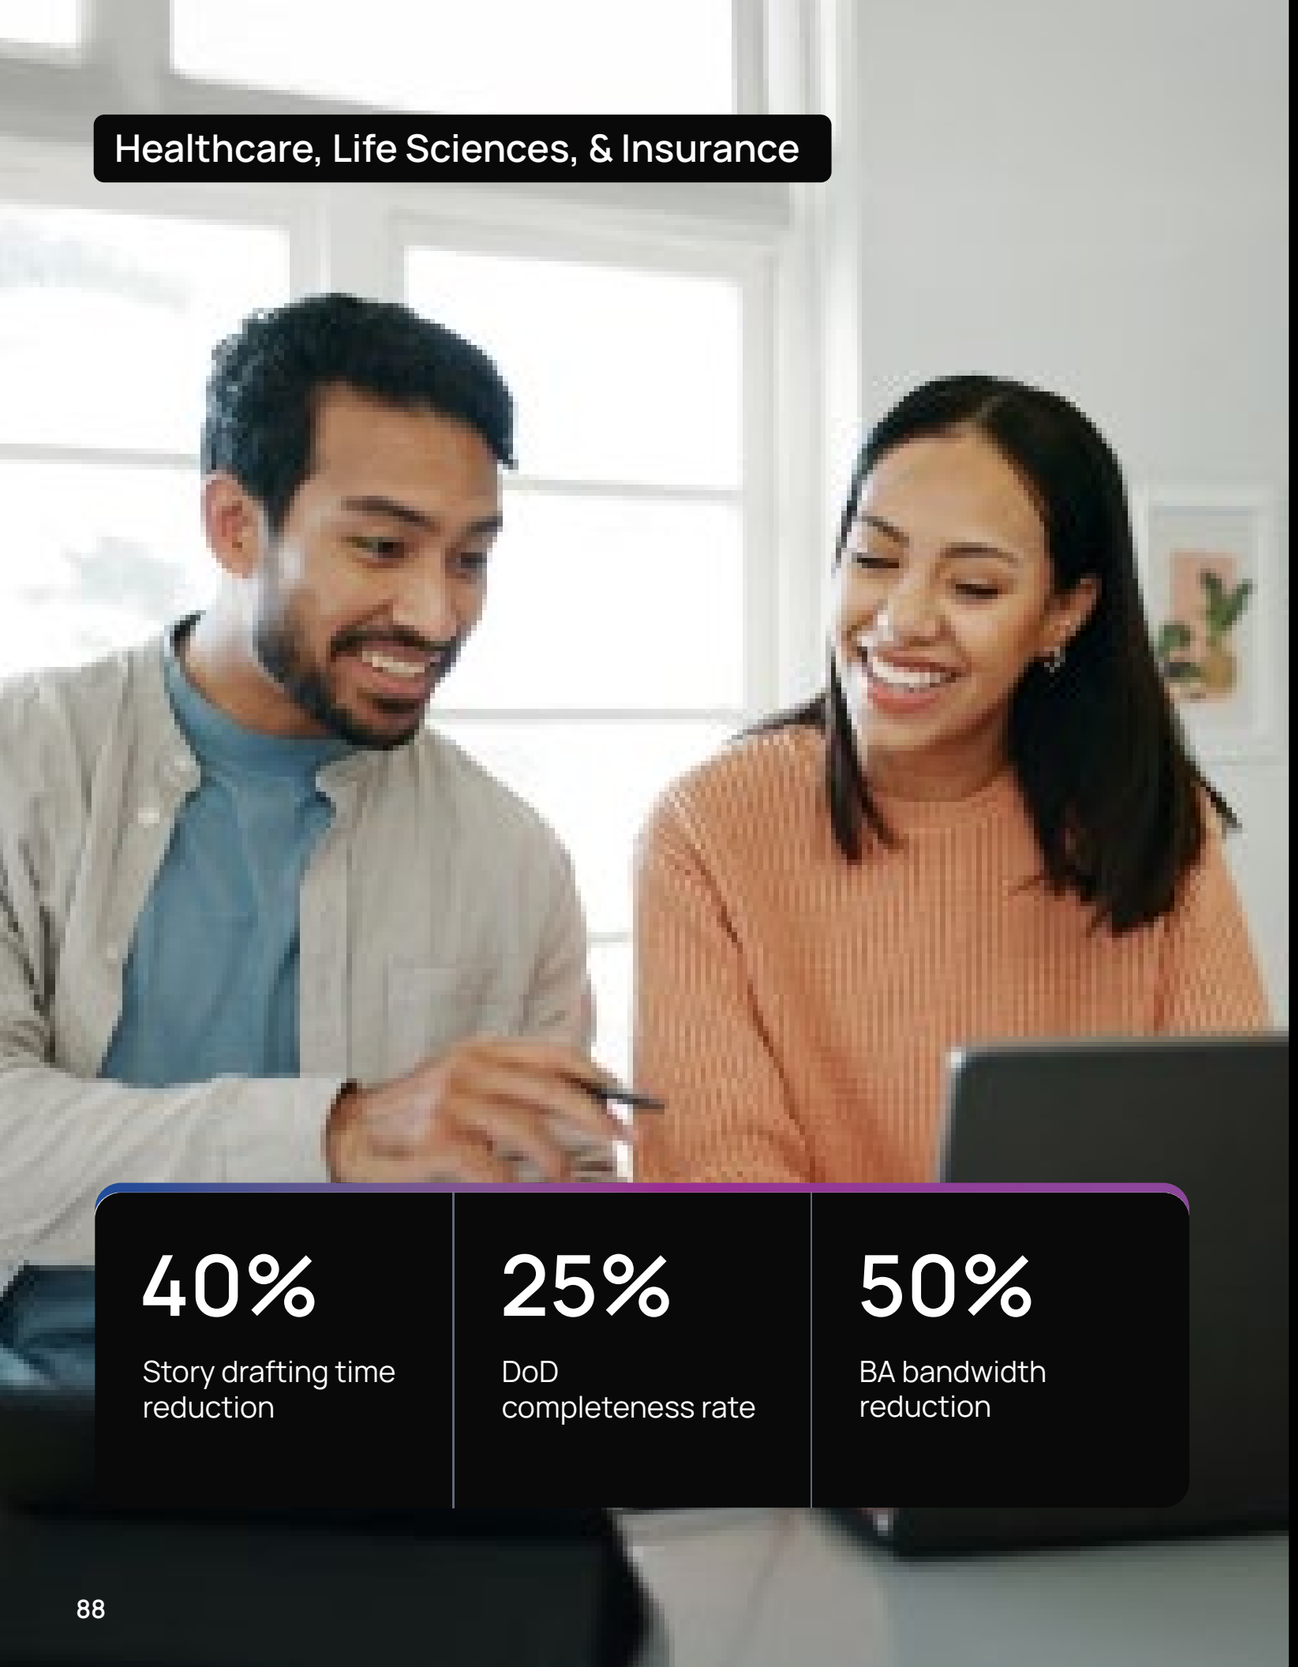

40%
Story drafting time reduction

25%
DoD completeness rate

50%
BA bandwidth reduction

Client

A Leading Insurance Repository for Life and Non-Life Insurance in India

Accelerating application development with AI-assisted requirements engineering and code quality

The client sought to accelerate product development while accommodating evolving business requirements. The engagement focused on using AI to automate requirements analysis, user story creation, and code quality validation to improve development speed and consistency.

Solution

- Leveraged Microsoft 365 Copilot to generate user stories, Definition of Done (DoD), acceptance criteria, and functional test cases from business requirements.
- Used GitHub Copilot to automate code quality checks, coding standards validation, and code coverage reporting.
- Enabled AI-assisted development across a Java, ReactJS, Oracle PL/SQL, and Azure technology stack.
- Established a foundation for AI-driven software engineering to support on-time project delivery.

Benefits

- **Automated code quality checks and coverage reporting.**
- **Enabled 24×7 story creation,** eliminating dependency on senior Product Owner availability.
- **Reduced effort for backlog refinement meetings.**

Healthcare, Life Sciences, & Insurance

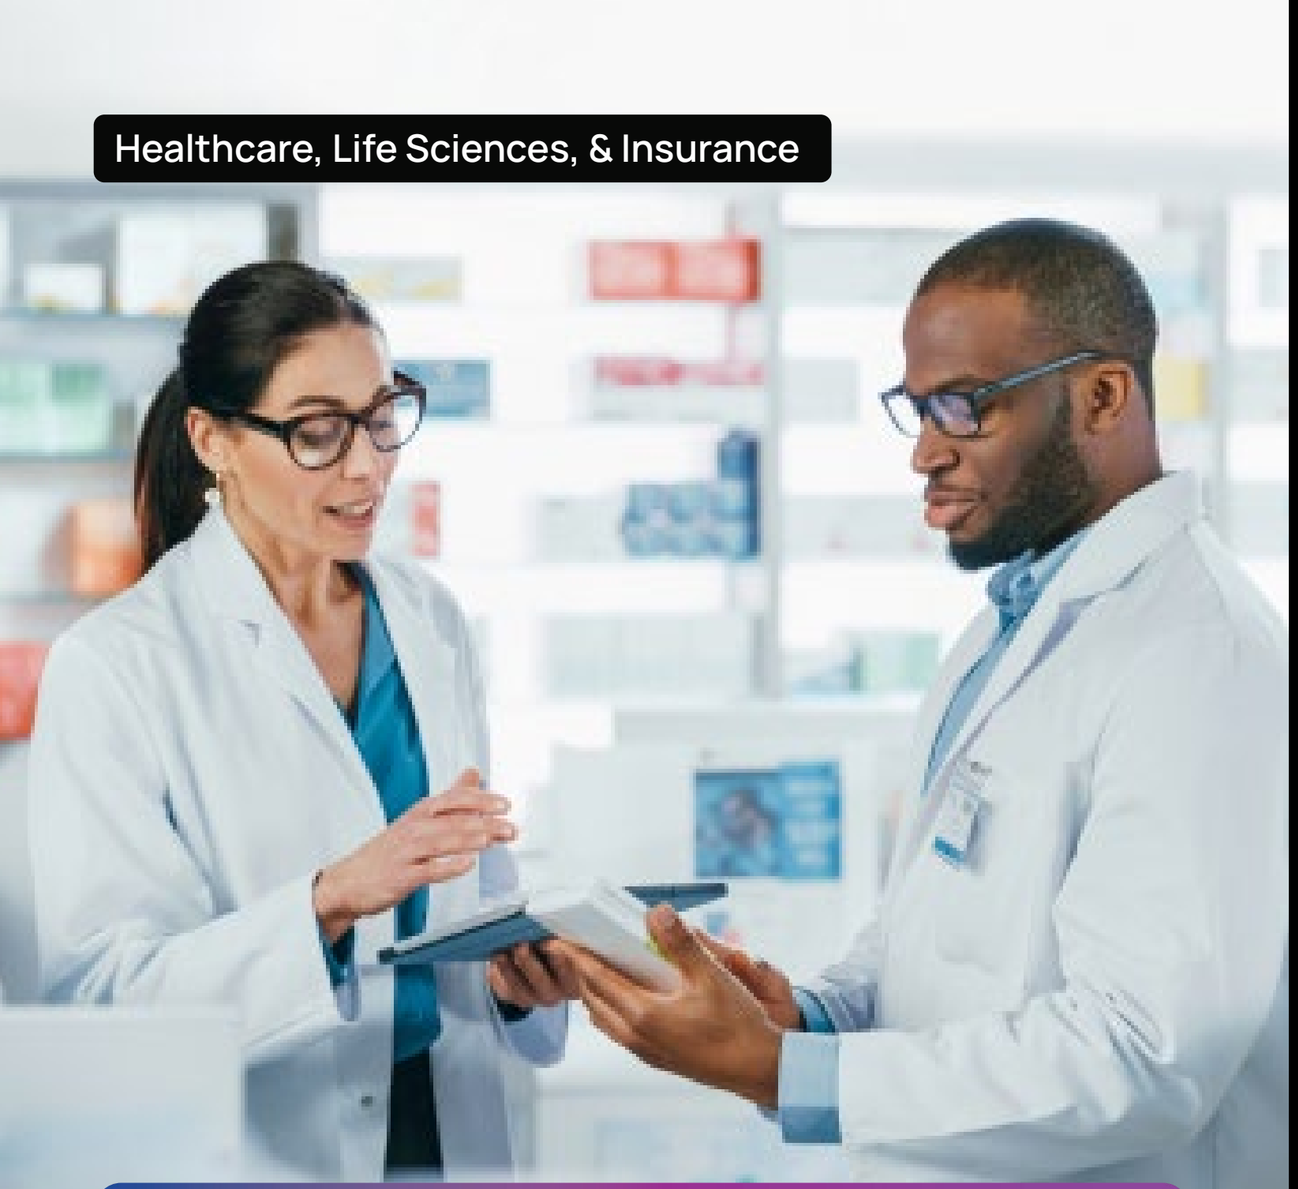

20%	60%	40–60%
Increase in brownfield development velocity	Reduction in cycle time	Faster delivery of business value

Client

A Leading US Pharmacy Benefit Manager (PBM)

Embedding agentic AI across the engineering lifecycle

The client's pharmacy benefit and fulfillment applications run on a broad stack, including Java, Node, ReactJS, .NET, SSRS/SSIS, and Python. Manual coding and review cycles were slowing releases, creating pressure to embed AI directly into the engineering lifecycle.

Solution

- Delivered engineering services covering PBM and pharmacy fulfillment application development and enhancement.
- Deployed Codex, an agentic AI solution, directly into core SDLC workflows.
- Augmented day-to-day development with Codex to lift individual developer velocity.

Benefits

- Brownfield development velocity increased by 20%.
- Cycle Time reduced by up to 60%.
- Business value delivered up to 60% faster.

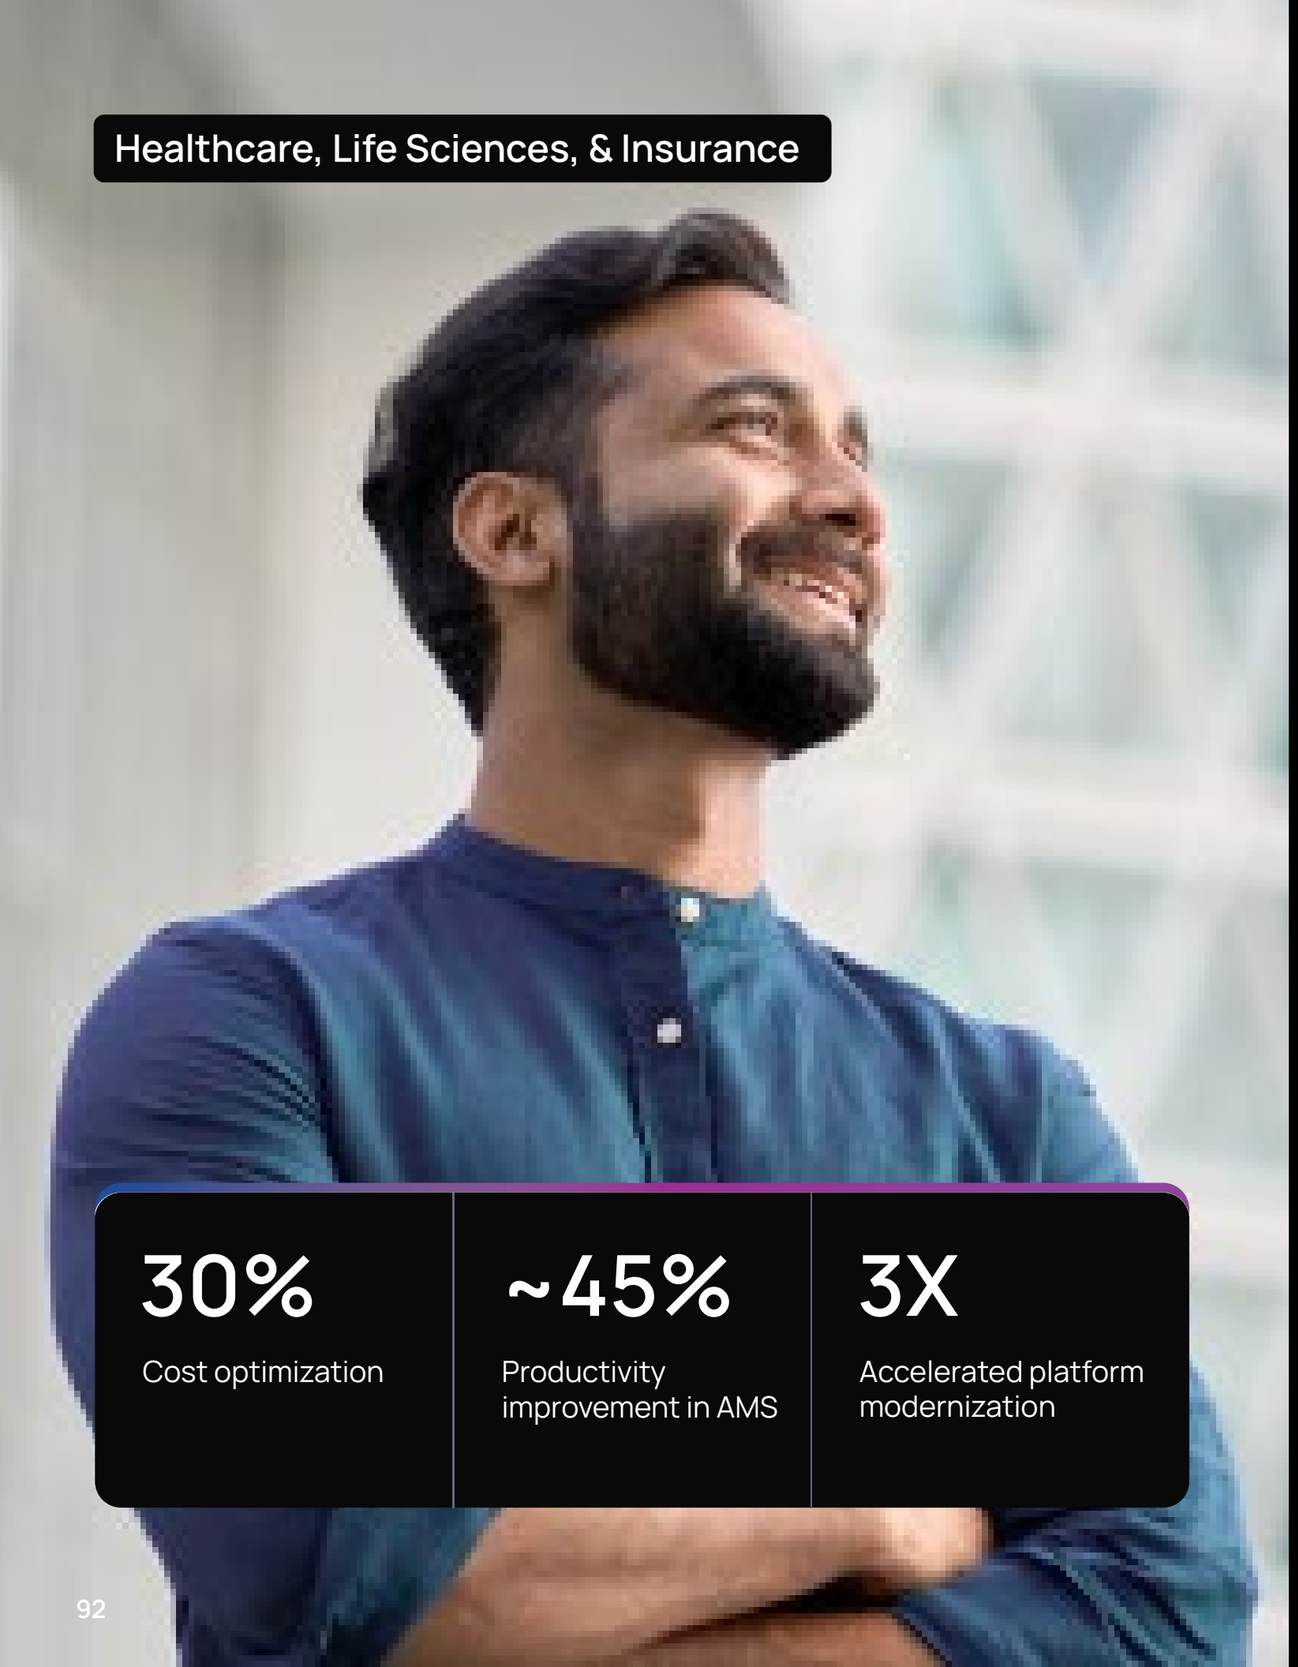

Client

A Global Commercial & Specialty Insurer

Empowering AI-driven operational excellence for enhanced efficiency

The client struggled with limited AI adoption, lengthy approval decision-making cycles, outdated technology platforms, and productivity constraints stemming from reliance on a single AI tool (GitHub Copilot).

Solution

- Leveraged Agentic AI, complemented by Generative AI, for AI-assisted development and structured execution.
- Integrated GitHub Copilot for code remediation, script generation, and test case creation.
- Accelerated migration to modern platforms, including ASP.NET Core.
- Enhanced operational efficiency with SDLC framework and AI-powered tools.

Benefits

- Improved operational efficiency by 30% and reduced decision-making delays.
- Accelerated platform modernization by 3x, meeting compliance requirements.
- Enhanced productivity 45% via AI-assisted development.

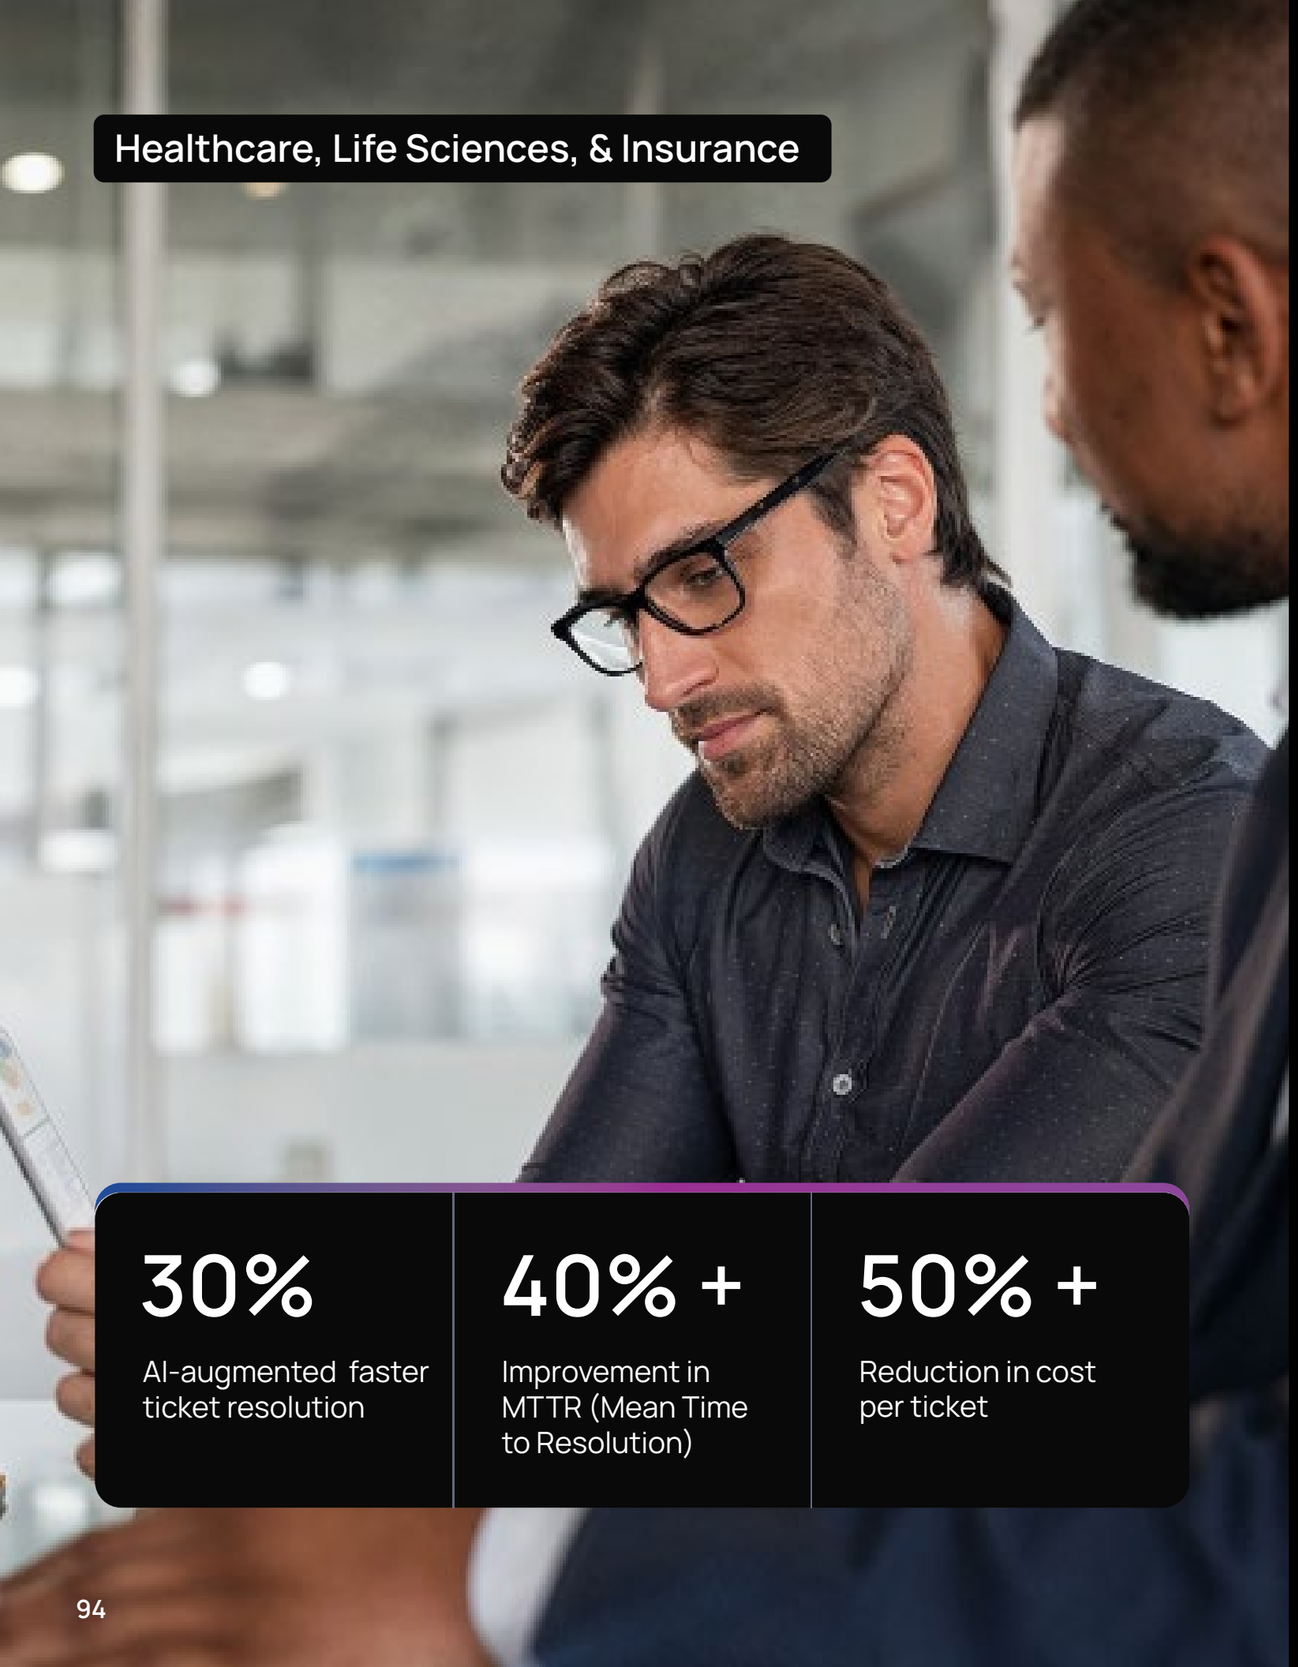

Client

A Leading Indian General Insurance Provider

Transforming application support with AI-assisted service operations

The client needed to improve productivity across L1, L1.5, and L2 application support for six critical business applications. To accelerate service delivery and reduce ticket volumes, Hexaware proactively invested in AI-driven automation to establish a foundation for intelligent support operations.

Solution

- Established a Hexaware-funded AI team to drive AI adoption across application support services.
- Leveraged Gemini 2.5 Flash-Lite and Claude Opus 4.5 to evaluate AI-assisted service operations.
- Applied AI to automate L1 and L2 support activities and improve service efficiency.
- Enabled AI adoption across an ITSM, SQL, and Excel-based support environment.

Benefits

- **Established a scalable foundation** for AI-led service transformation.
- 100% Automated SOP creation through a Watch and Learn Agent.
- End-to-end automated ticket resolution with integration to an ITSM tool.

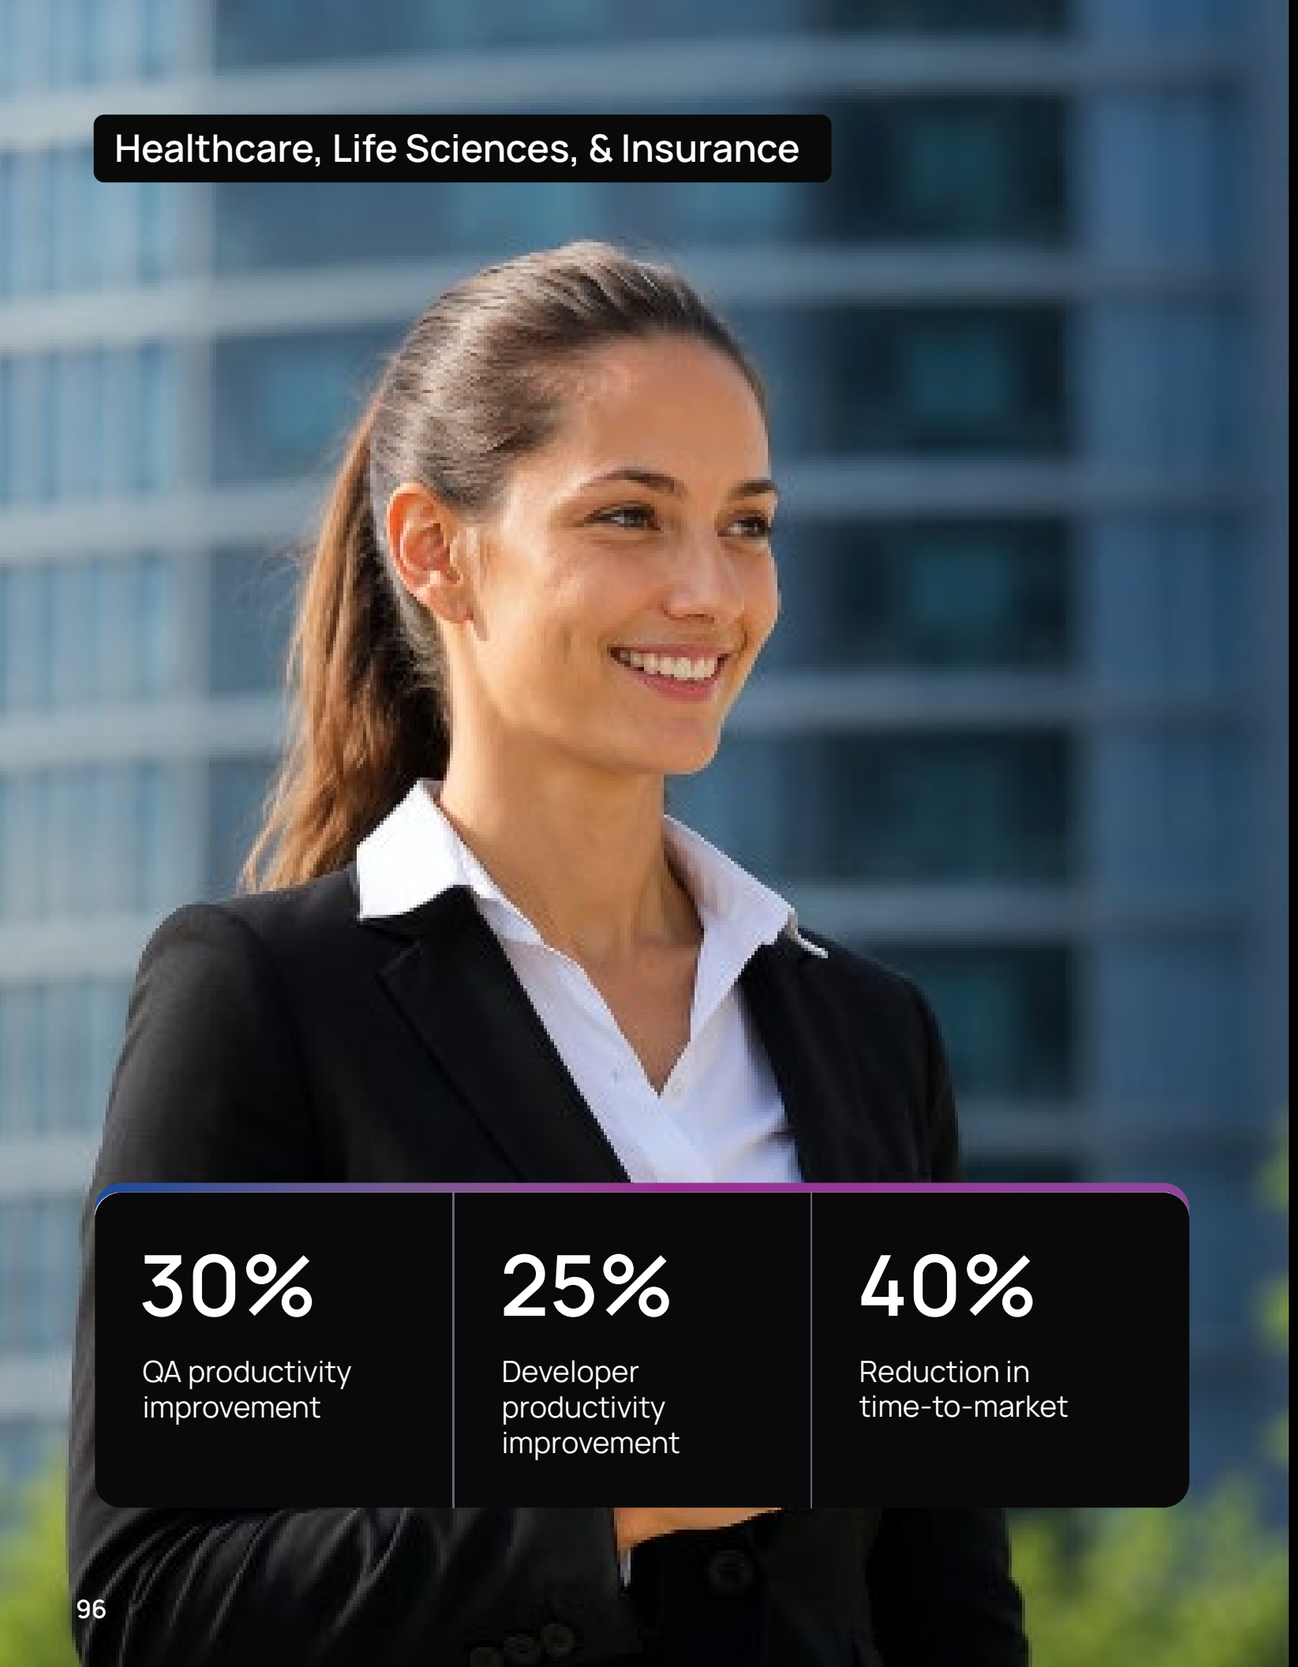

Client

A Leading Global Commercial and Specialty Insurance Provider

Accelerating low-code application delivery with AI-assisted software engineering

The client sought to accelerate software development across ServiceNow App Engine initiatives while improving engineering productivity. The engagement focused on introducing AI-powered development and testing within the client's enterprise platform, with licensing approvals and pilot activities underway.

Solution

- Implemented **ServiceNow Now Assist** to accelerate development on the **ServiceNow App Engine** platform.
- Deployed Hexaware's **Tensai® for Autonomous Testing**, our agentic AI platform, to improve testing productivity.
- Implemented **Agentverse™** for AI-powered document extraction.
- Enabled AI-driven quality engineering and software delivery across the ServiceNow ecosystem.

Benefits

- **Improved QA productivity** by 30% through agentic QA.
- **Increased developer productivity** by 25% with AI-assisted development.
- **Reduced time-to-market by 40%** through accelerated SDLC.
- Accelerated SDLC across ServiceNow App Engine initiatives.
- Built a scalable foundation for AI-powered software delivery.

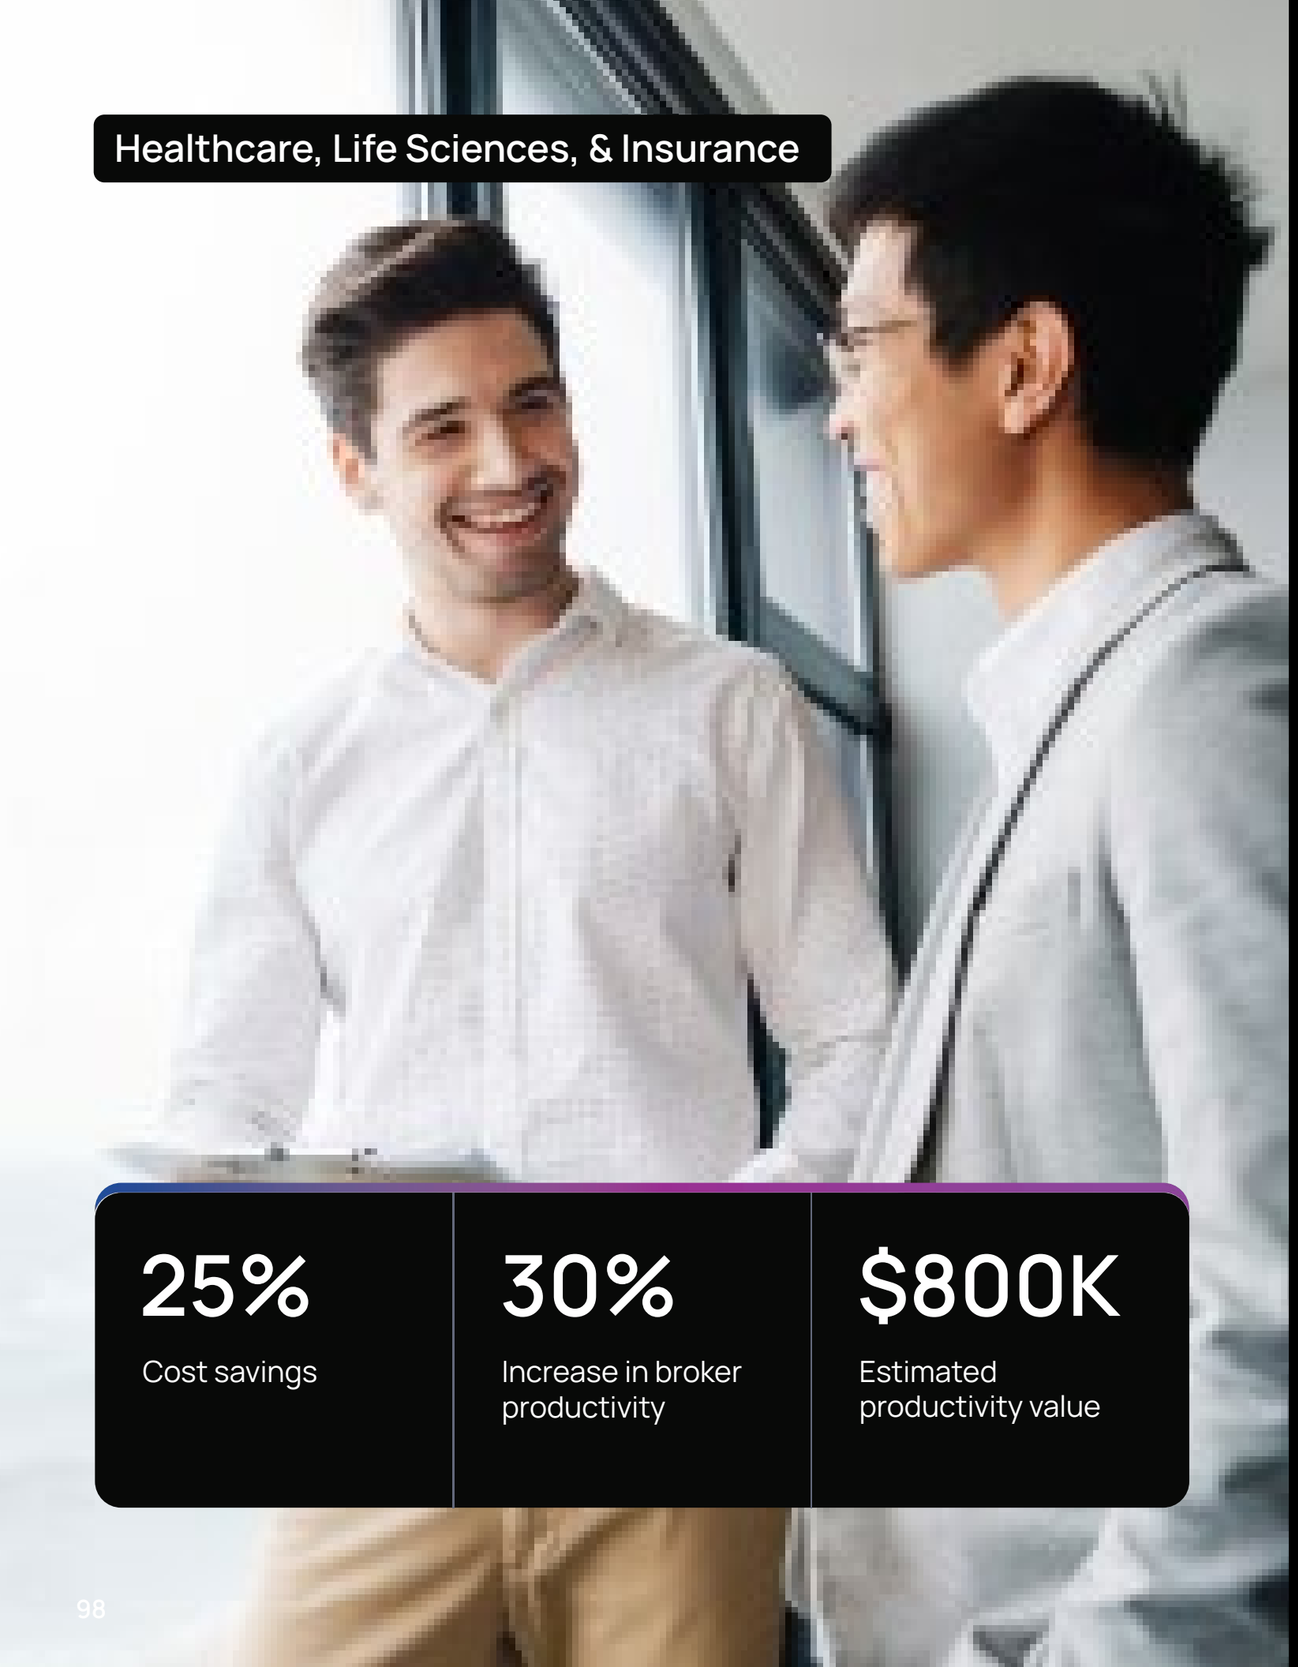

25%
Cost savings

30%
Increase in broker productivity

$800K
Estimated productivity value

Client

One of the Largest Insurance Brokers and Insurance Intermediaries

Building a secure foundation for enterprise-scale AI-assisted software engineering

The client sought to improve engineering productivity across multiple application development and support engagements while adhering to strict security and governance requirements. The objective was to improve engineering productivity and deliver contractual cost savings.

Solution

- Adopted GitHub Copilot and AON GPT across application development and support engagements.
- Enabled AI-assisted development using approved enterprise AI tools.
- Trained 120 consultants to accelerate AI adoption across delivery teams.
- Supported engineering across 9 statements of work (SOWs) spanning insurance platforms and business applications.

Benefits

- Delivered 25% cost savings through AI-assisted software engineering.
- Improved broker productivity by 30% with AI-enabled workflows.
- Accelerated AI adoption with 120 trained consultants.
- Improved engineering productivity using secure, enterprise-approved AI tools.
- Established a scalable foundation for AI-enabled software delivery.

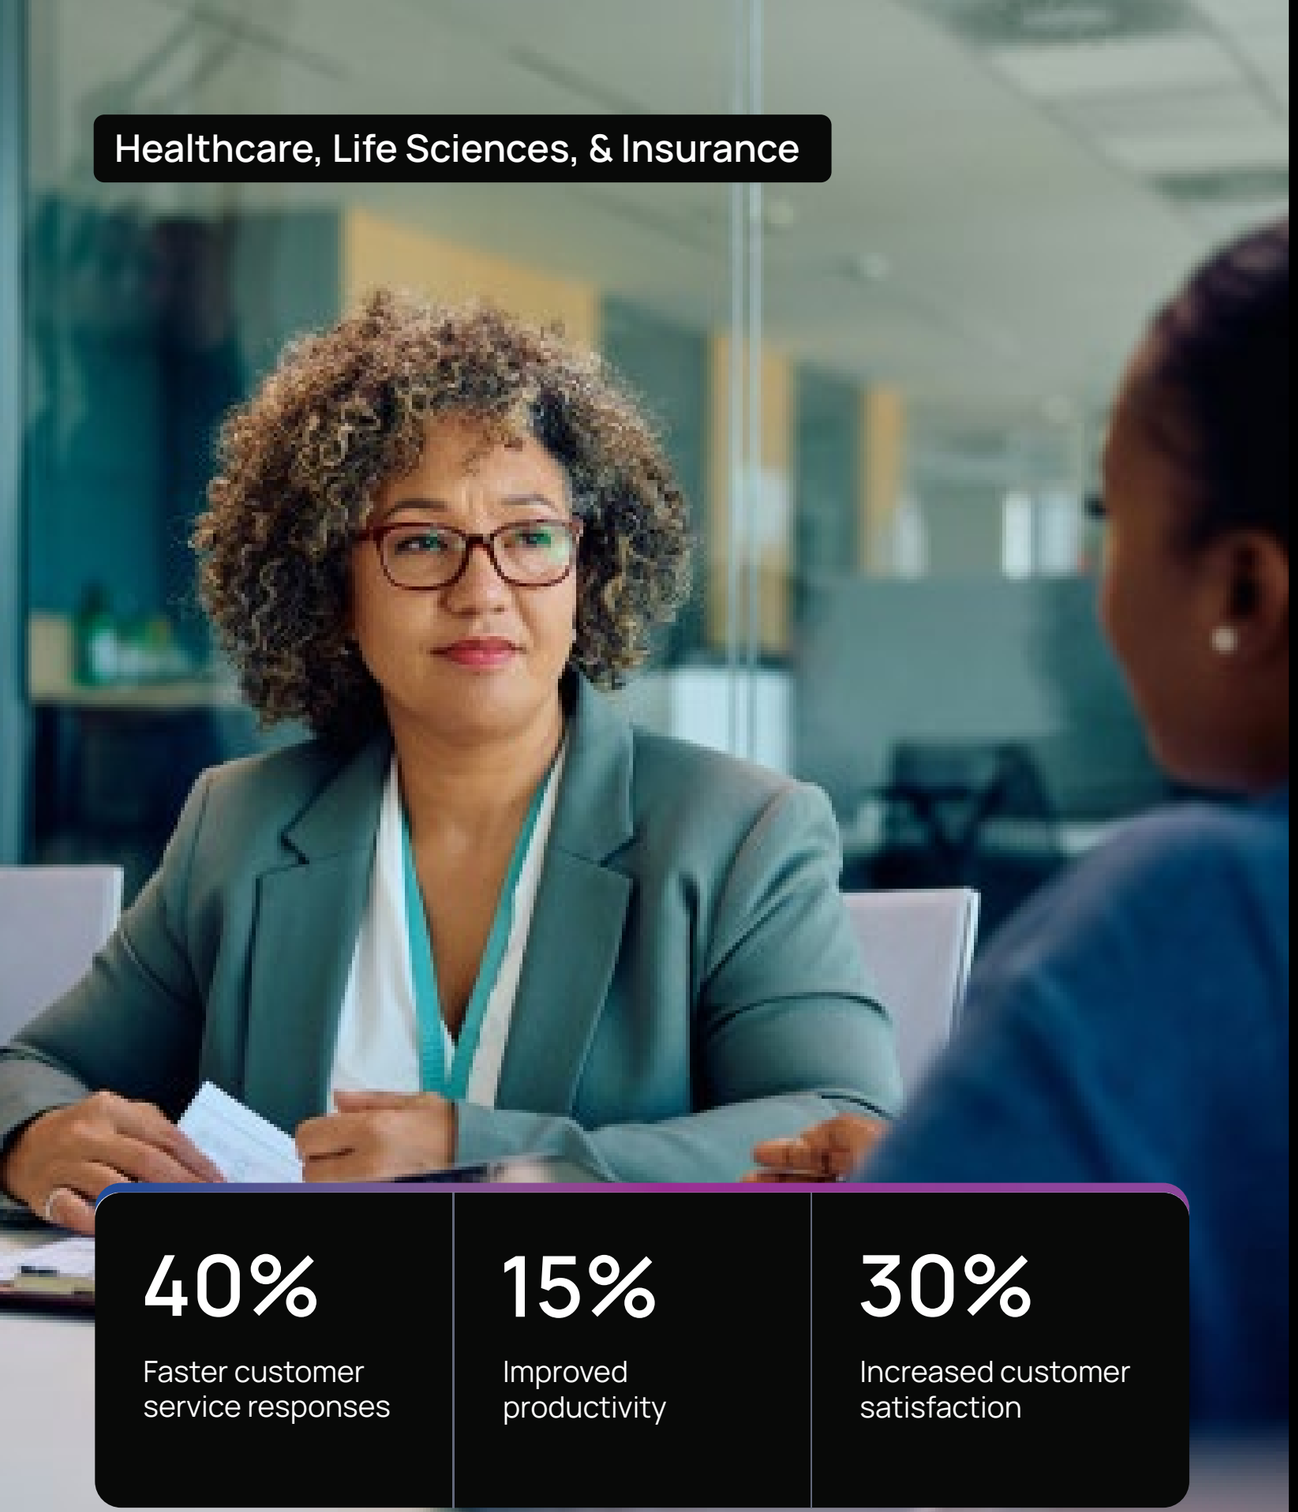

Client

An Insurance and Asset Management Company

Equipping service teams with answers in the moment

Customer service executives needed faster access to product knowledge and customer insights so every interaction could feel more informed and more personal. The goal was better responsiveness without increasing complexity for frontline teams.

Solution

- Created a GenAI assistant for product and policy inquiries in close partnership with the client's product team.
- Used a sophisticated retrieval augmented generation (RAG) architecture across public and private data sources.
- Designed the solution to support GDPR requirements, including session-level data purging.
- Surfaced real-time answers and insights during live customer interactions.

Benefits

- Improved customer service response times by 40%.
- Increased team productivity by 15%.
- Raised customer satisfaction by 30%.
- Reduced agent ramp time while creating more opportunities for personalized recommendations and cross-sell.

Client

A Global Biopharma Leader

Empowering AI-driven efficiency in biopharma operations

The client faced limitations in autonomous decision-making across systems, lacked context awareness in code generation, and required extensions to automation for multi-step processes, impacting efficiency and scalability.

Solution

- **Agentic AI Implementation:** Integrated advanced AI tools like GitHub Copilot for agile development and decision-making.
- **Proactive Issue Prevention:** Transitioned from reactive ticket handling to proactive issue resolution and self-service.
- **Custom Extensions:** Developed custom AI extensions to address gaps in multi-step automation using Azure services.
- **Accelerated Solution Development:** Introduced automation in workflows, achieving faster development cycles.

Benefits

- Achieved 30–50% faster ticket resolution.
- Delivered 2–5x accelerated solution development.
- Reduced 10% support costs by eliminating repetitive requests.
- 100% elimination of routine weekend health check work.

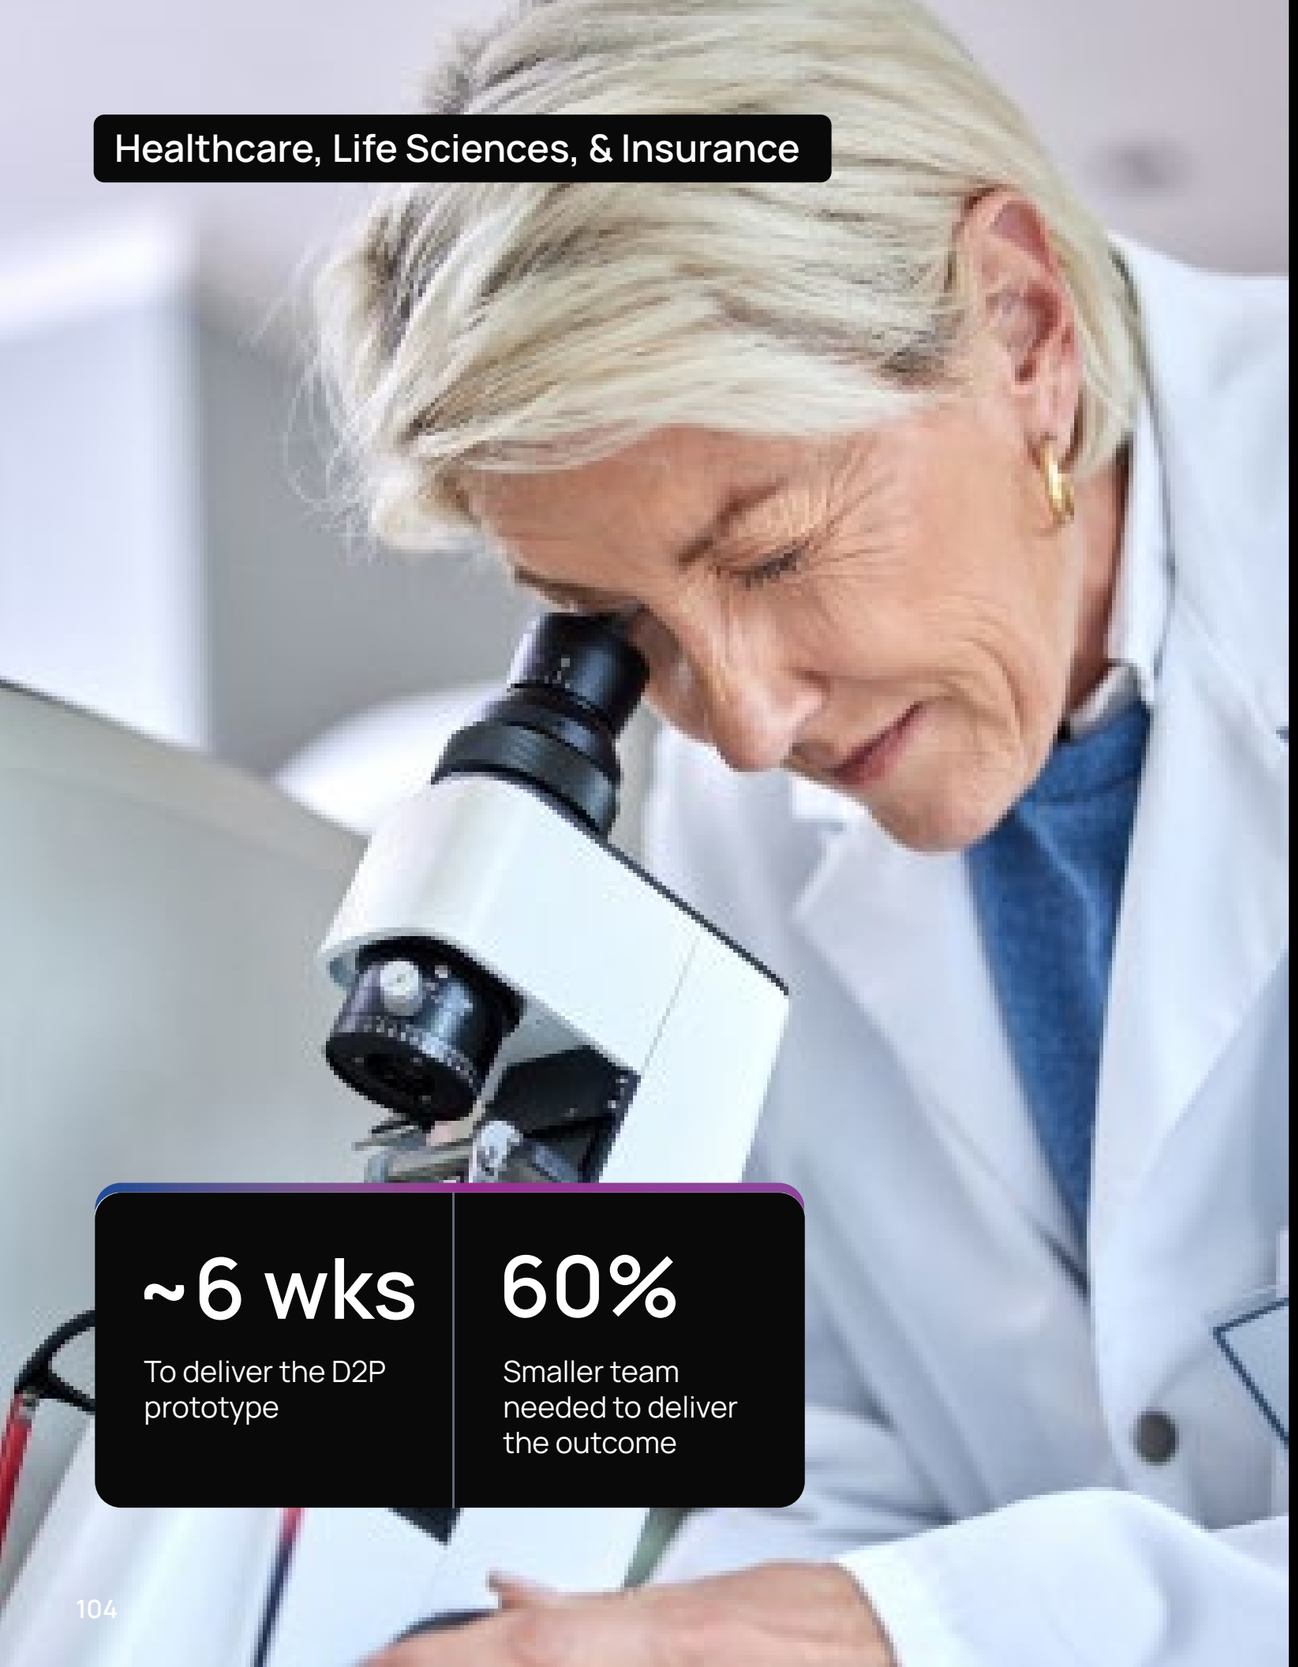

Client

A Leading Pharmacy Benefit Manager (PBM) That Uses Data Analytics, Clinical Programs, and Personalized Engagement to Improve Patient Outcomes

Launching a high-fidelity D2P prototype before the market moved on

Regulatory shifts were reshaping the market, and manufacturers needed stronger direct-to-patient experiences. The client wanted a fast, credible prototype that could help it compete quickly without sacrificing realism or compliance.

Solution

- Delivered a fully functional direct-to-patient prototype using Hexaware's Vibe Coding framework.
- Mapped the end-to-end patient and administrative journey from physician visit through fulfillment.
- Built a high-fidelity patient experience plus a dynamic admin portal for white-labeled demos.
- Separated marketing and clinical content to support regulatory safety and demo readiness.

Benefits

- Delivered the prototype in about six weeks, helping the client respond quickly to market change.
- Achieved the outcome with a team roughly 60% smaller than a traditional approach might require.
- Created a credible source of truth for future product development and manufacturer conversations.
- Enabled faster go-to-market discussions with a more realistic, demo-ready experience.

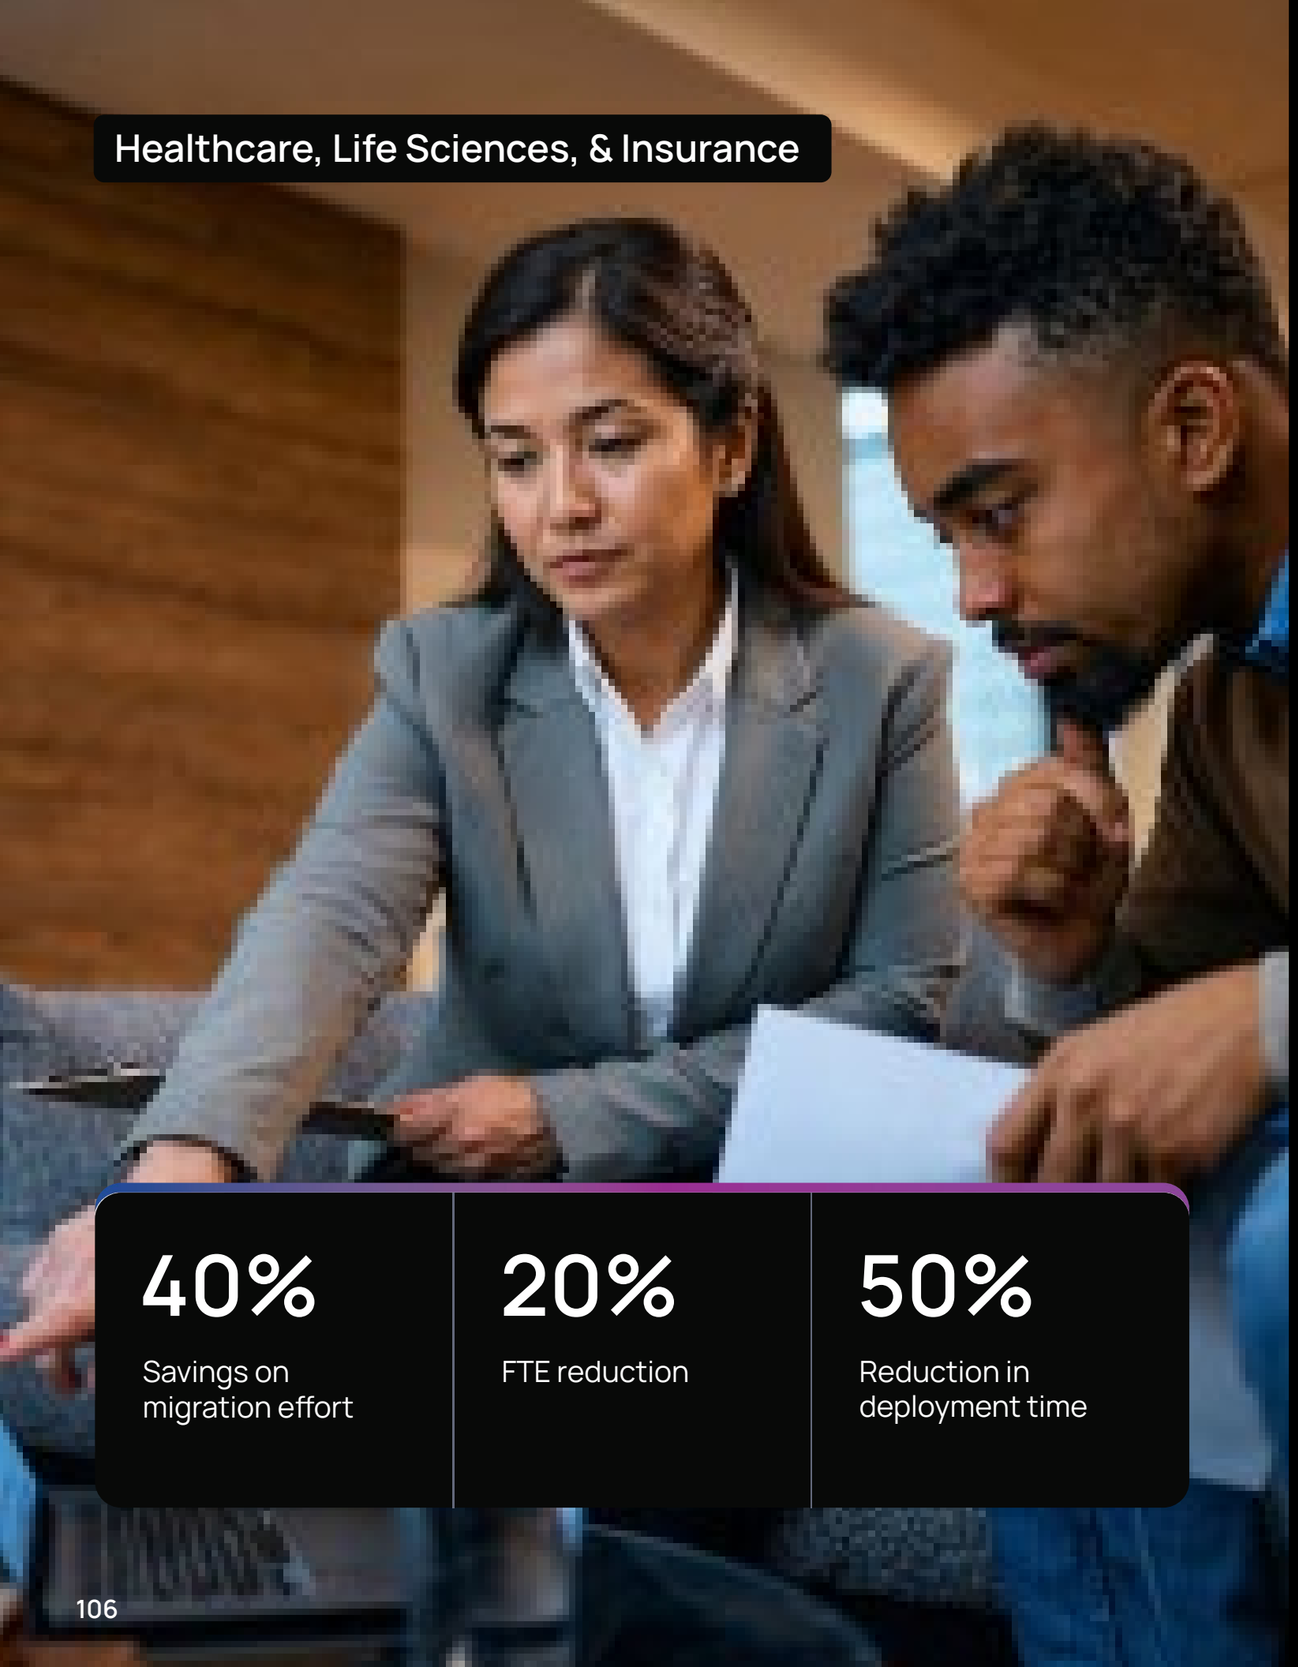

Client

A Leading Belgian Banking & Insurance Group

Migrating legacy SAS analytics to Oracle PL/SQL with GenAI

The client's business-critical analytics ran on legacy SAS programs that were costly to license and maintain, with a shrinking talent pool to support them. They needed to migrate to Oracle PL/SQL without disrupting operations or introducing risk.

Solution

- Migrated SAS programs to Oracle PL/SQL across the analytics estate.
- Used Amaze® for Data to automate program conversion.
- Focused AI acceleration on simple and medium-complexity programs.
- Validated migrated logic to preserve accuracy and business continuity.

Benefits

- Achieved ~40% savings on simple and medium-complexity migration.
- Delivered 20% FTE reduction.
- Accelerated migration through generative AI-assisted conversion.

Client

A Leading Belgian Non-Life Insurer

Automating policy mutations, claims servicing, and customer queries with agentic AI

Every customer email requesting a policy update, claims service, or related inquiry required manual review and data entry by back-office teams before processing could begin. This repetitive, time-consuming workflow constrained operational efficiency, increased processing times, and limited scalability.

Solution

- Built a standalone API triggered automatically when a request email arrives.
- Used OpenAI to extract the policy number, mutation type, claims detail from each email.
- Generated a structured JSON payload from the extracted details.
- Fed the payload into the existing mutation, claims API to process changes automatically.

Benefits

- Reduced back-office effort by ~25% annually.
- Achieved 95% accuracy in customer request resolution.
- Accelerated policy and claims changes through end-to-end AI automation.

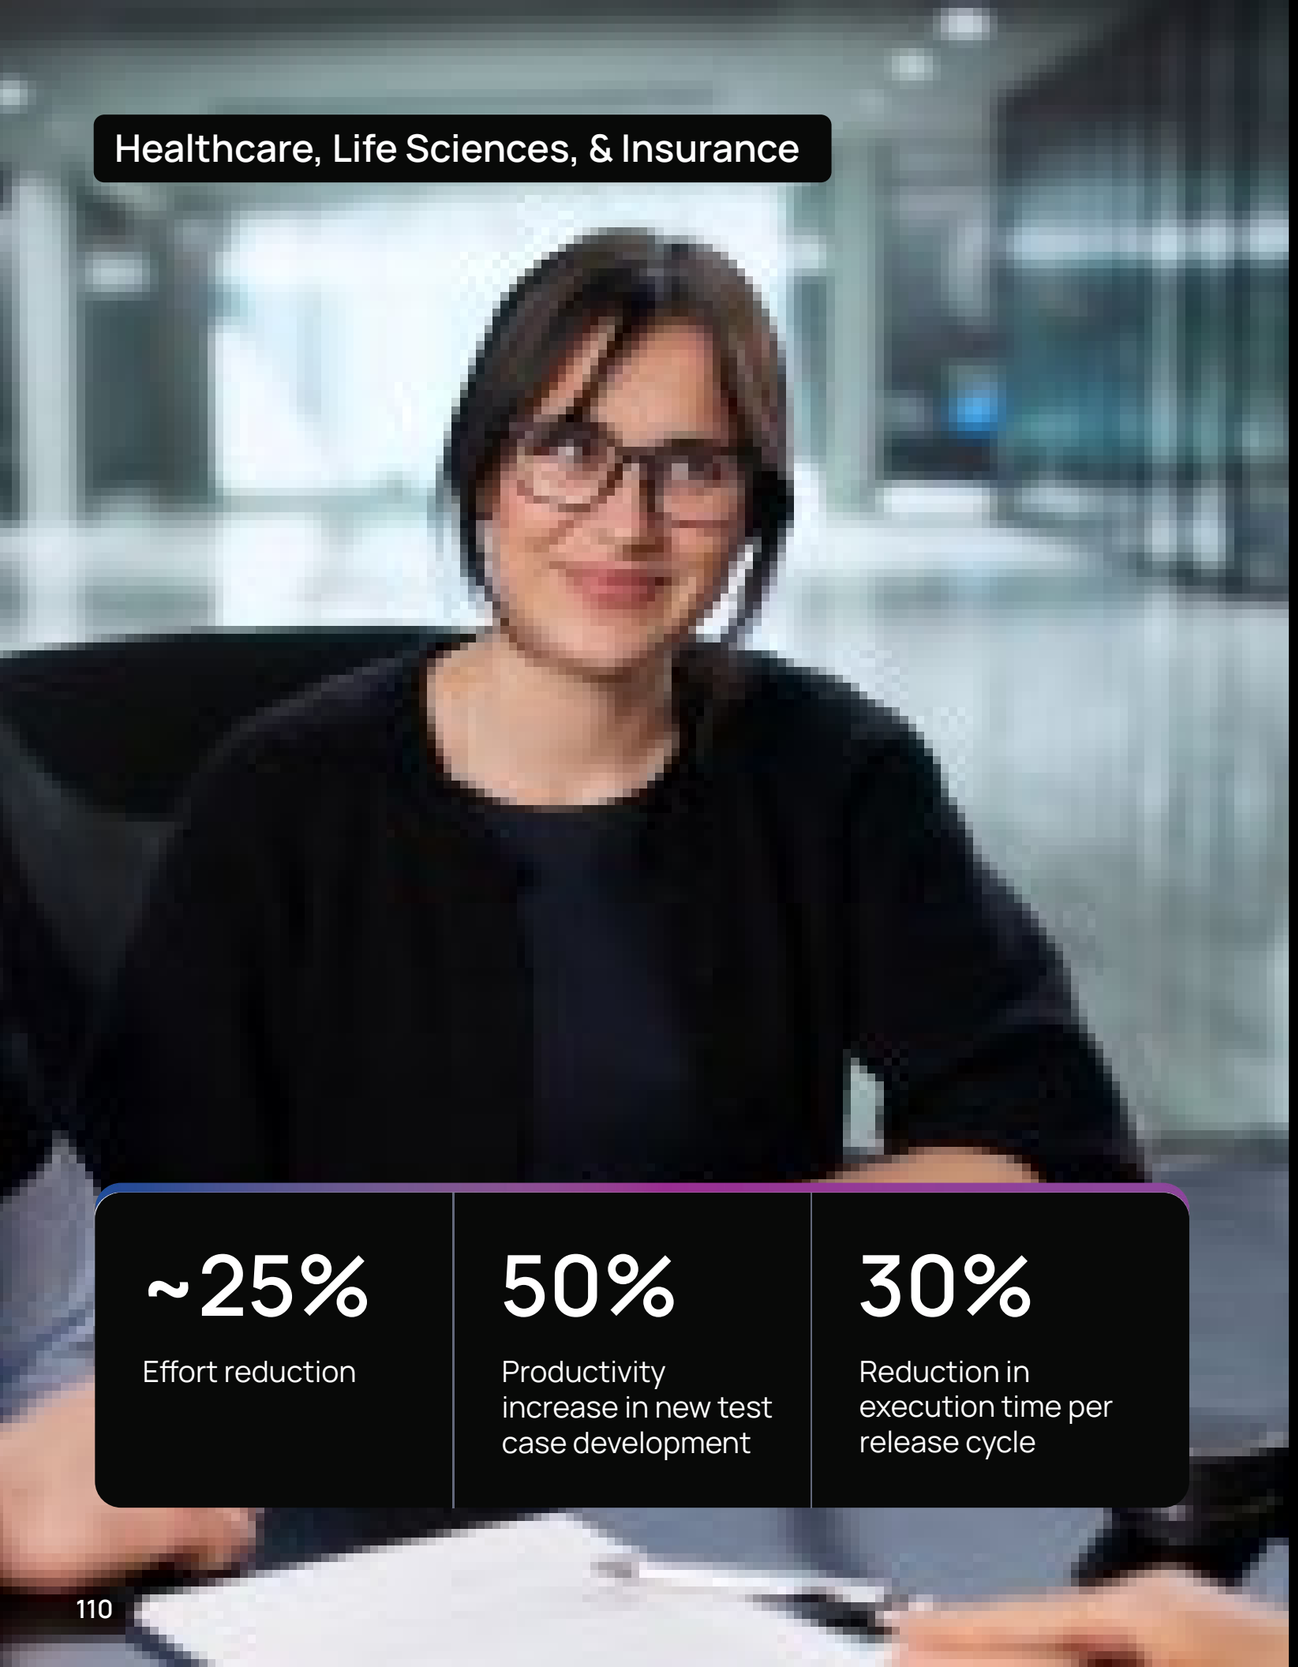

Client

A Leading UK Insurer

Driving improved efficiency through generative AI-powered testing

The client faced testing inefficiencies due to manual processes across claims handling and reconciliation, limited AI license usage (restricted to 4 users), and constrained monthly AI usage.

Solution

- **Accelerated AI-Driven Testing:** Enabled seamless adoption of GitHub Copilot to enhance test automation through generative AI capabilities.
- **Modernized Test Automation Frameworks:** Leveraged Playwright, TypeScript, and BDD methodologies to streamline testing for claims processing and reconciliation applications.
- **Quantified Efficiency Gains:** Established standardized metrics to measure productivity improvements, effort reduction, and testing effectiveness across projects.
- **Optimized Regression Testing:** Applied AI-driven test selection to identify the most relevant test cases for execution, enabling risk-based testing and reducing regression effort while maintaining coverage.

Benefits

- Realized 20–25% productivity gains across testing initiatives through AI-assisted test automation and streamlined workflows.
- Improved new test case development efficiency by 50%, enabling faster coverage expansion and reduced manual effort.
- Maintained high levels of quality and operational continuity across critical claims, reconciliation, and FNOL processes.
- Cut regression testing execution time by 30% per release, accelerating deployments while preserving risk coverage and application stability.

Client

A Global Commercial Insurer

Accelerating insurance testing with AI-powered agentic automation

Manual test case design and slow automation script development created delivery bottlenecks, limiting measurable value and hindering efforts to reduce cycle time.

Solution

- Deployed an agentic AI framework to automate test case design and script generation.
- Integrated GitHub Copilot to boost developer productivity across the testing team.
- Built Playwright-based agents for end-to-end functional test automation.
- Established structured AI adoption within the client's SDLC for scalable impact.

Benefits

- Improved test design and script generation productivity by 30%.
- Reduced manual effort by building a test automation framework, resulting in 15% cost reduction.

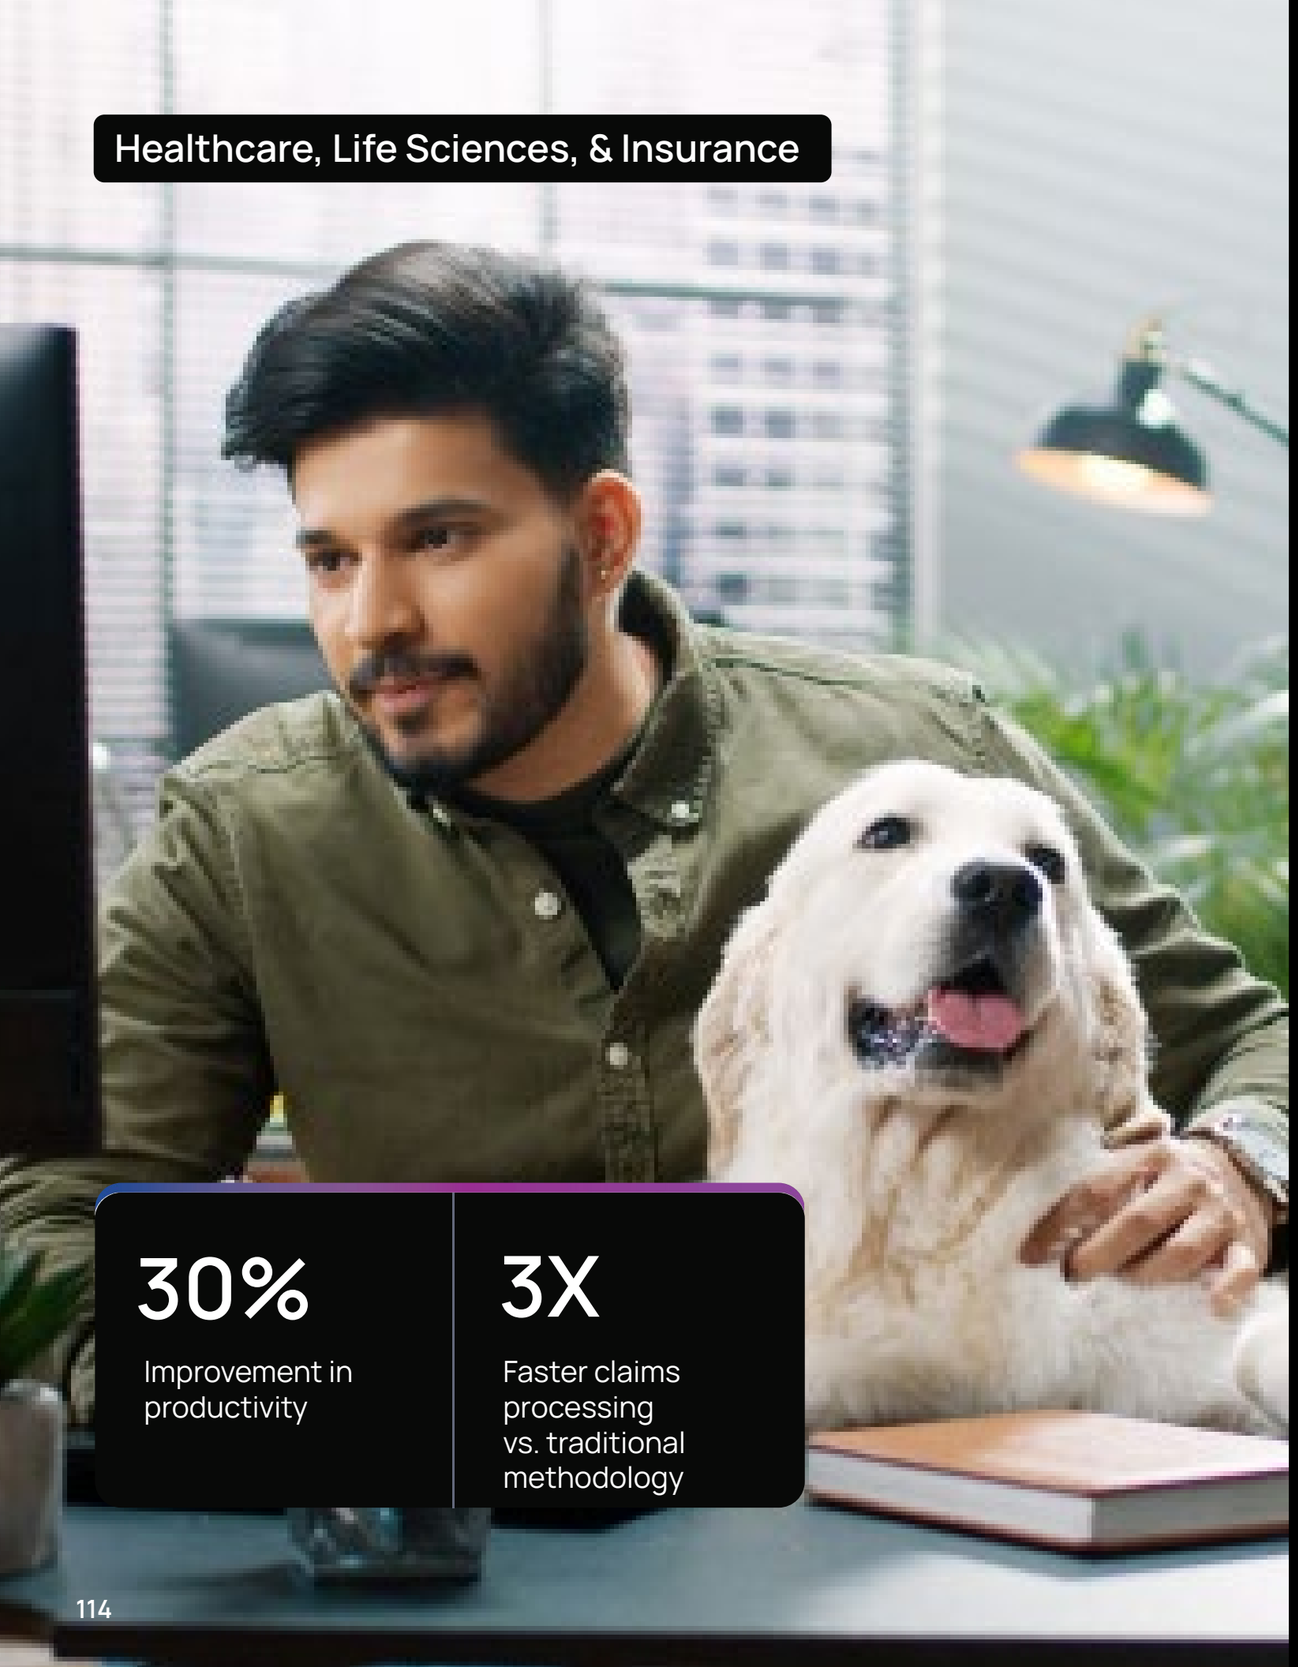

Client

A Leading Pet Insurance Provider in North America

AI-powered claims adjudication

The client processes thousands of claims documents daily. It needed to improve the accuracy and efficiency of claims adjudication at that scale without adding manual effort.

Solution

- Implemented an agentic AI document classifier using GPT 5.5 to read, classify, extract, analyze, and route medical records, invoices, and insurance forms.
- Automatically extracted invoice details and loaded them into the database.
- Analyzed and distinguished valid from invalid claims through estimation detection.
- Performed fraud detection by validating document structure, consistency, and validity.

Benefits

- Automated processing and classification of 1,000+ documents in the claims life cycle.
- Increased business-user productivity by 30% through faster claims turnaround.
- Delivered up to 3x faster claims processing compared to traditional methods.

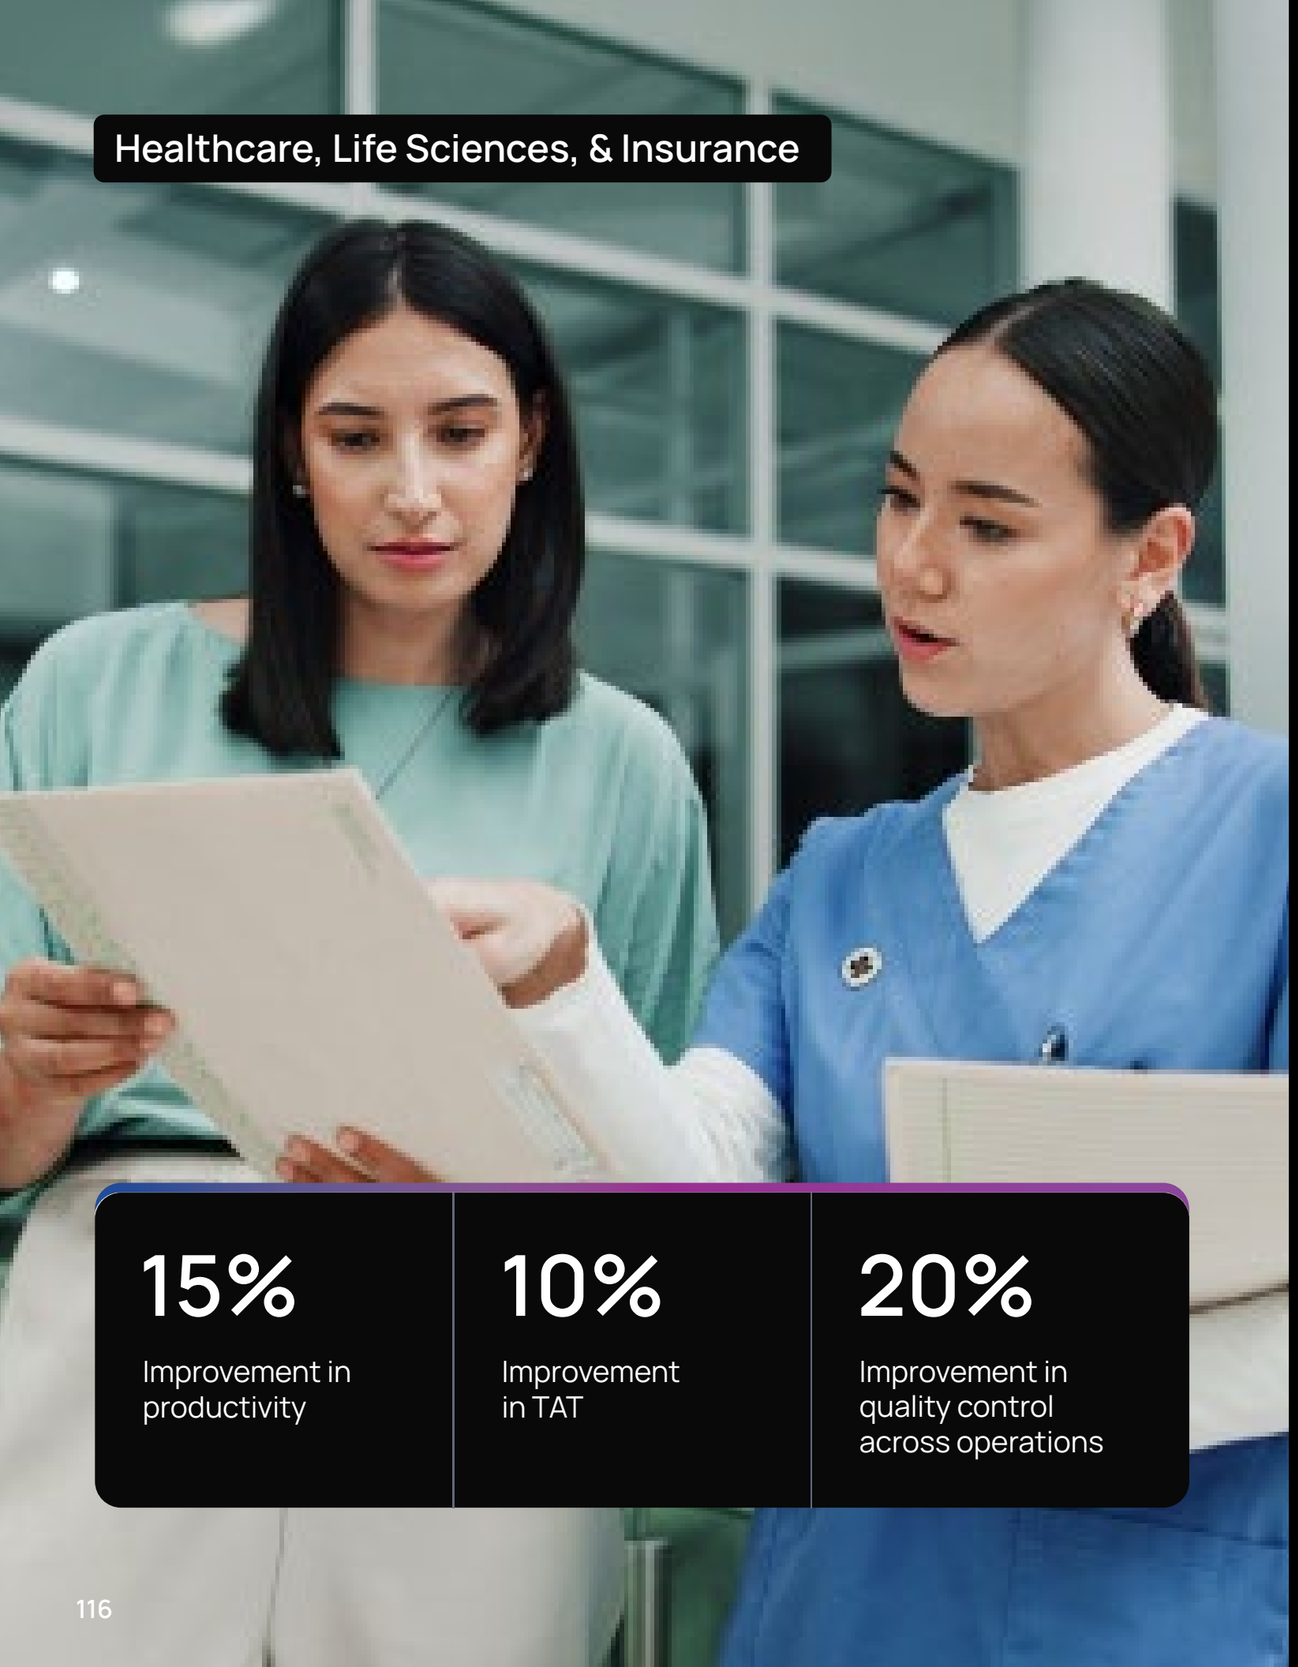

Client

A Leading North American Third-Party Administrator for Health Plans

Laying the foundation for AI-driven operations reporting and AI assisted SDLC across managed services

As part of its managed services modernization strategy, the client introduced AI to establish a unified framework for data insights and analytics, enabling smarter operations and driving innovation across application, data, infrastructure, and business service functions.

Solution

- Deployed Hexaware's Tensai® to automate infrastructure managed services with ServiceNow, Amazon Connect, and ELK, enabling sentiment analysis for calls and near real-time insights into infrastructure operations.
- Established a framework to measure productivity improvements through monthly reporting.
- Enabled AI-assisted development using approved enterprise AI tools.
- Trained all developers on GitHub Copilot to accelerate AI adoption across delivery tasks and activities.

Benefits

- Automated operations reporting with Hexaware's Tensai®, improving productivity by 15%.
- Reduced reporting turnaround time by 10% through AI-driven automation.
- Improved operational quality by 20% with AI-enabled reporting.
- Built a scalable foundation for enterprise AI adoption and future AI initiatives.

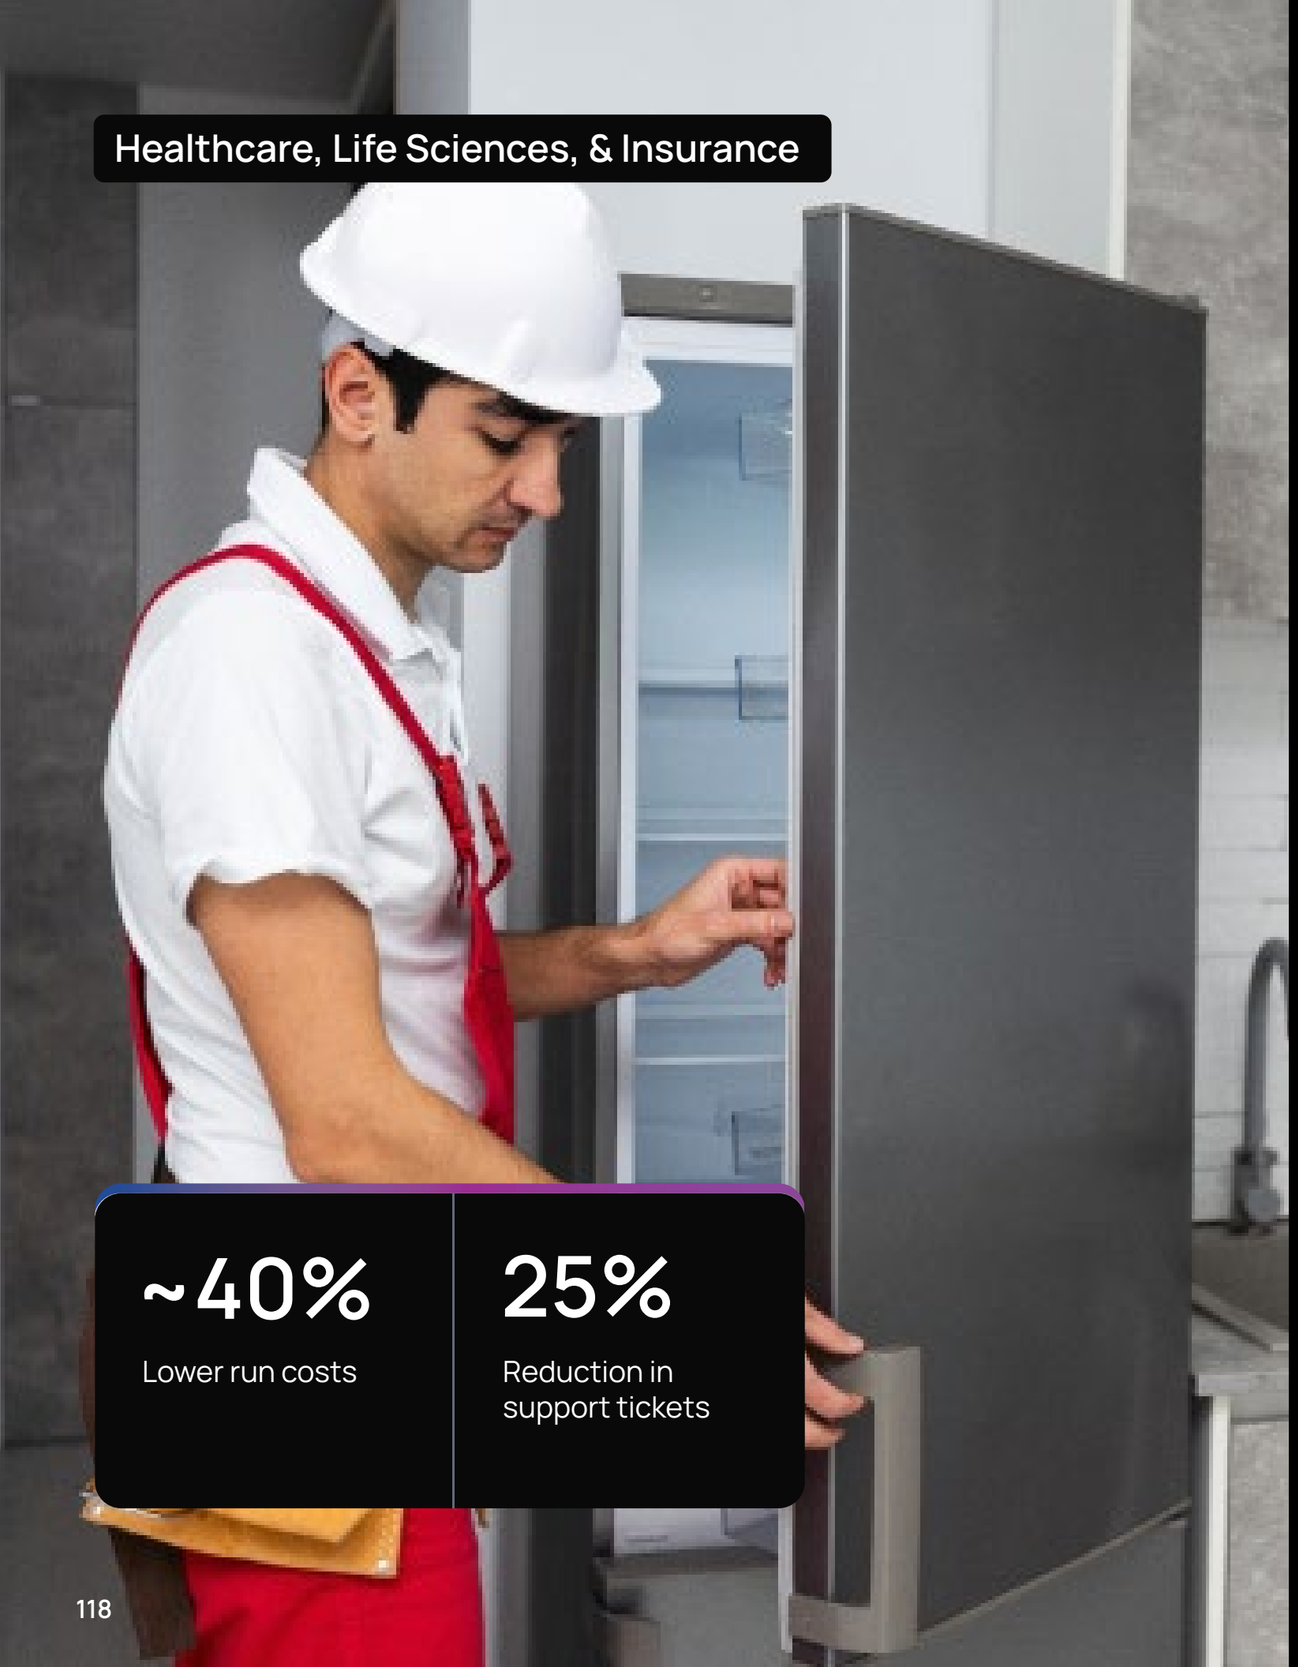

Client

A Century-Old, UK-Headquartered Appliance Protection and Repair Insurer

Modernizing Workday finance operations so growth does not create more friction

As the business grew, manual journal reviews, fragmented reporting, and integration issues were increasing audit risk and operational drag. The client needed finance operations that could scale with more control and less effort.

Solution

- Implemented Workday Illuminate and Journal Insights to automate anomaly detection.
- Standardized reporting through real-time KPIs and unified dashboards.
- Stabilized Finance and HCM AMS operations with proactive support and preventive maintenance.
- Improved integrations, security roles, and governance across finance processes.

Benefits

- Reduced run costs by about 30–40%.
- Cut support tickets by 25% through proactive fixes and automation.
- Improved month-end close speed and financial accuracy with stronger anomaly visibility.
- Strengthened audit readiness through better reporting and more consistent controls.

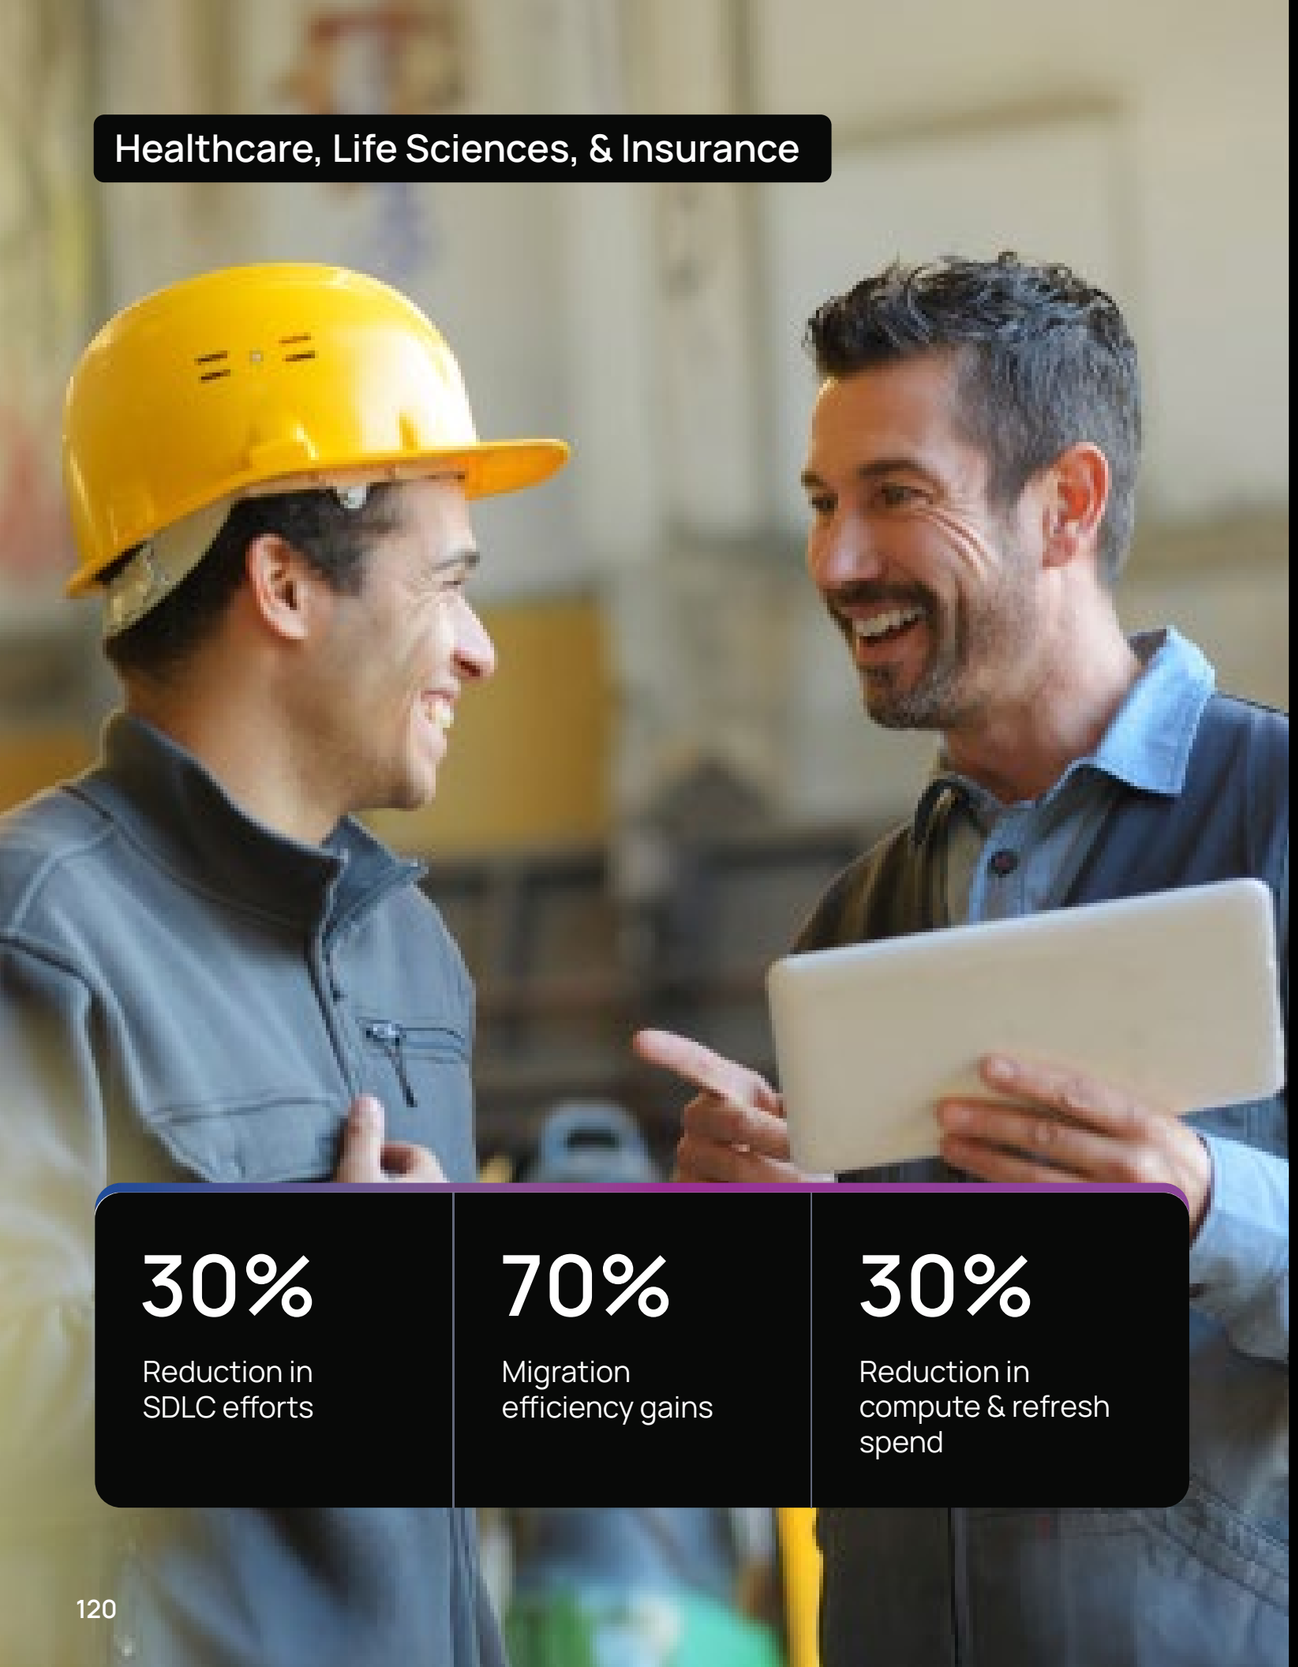

Client

A US Workers' Compensation Care Management Company

Modernizing the data platform with Tensai®-accelerated migration

The client manages physical medicine, rehabilitation, and medical bill reviews for the workers' compensation industry. Its reporting was fragmented across Sisense, Qlik Sense, and MicroStrategy, creating pressure to modernize onto a governed, AI-ready data platform.

Solution

- Assessed the current reporting ecosystem across Sisense, Qlik Sense, and MicroStrategy.
- Built a centralized modern data platform with medallion architecture, CDC-based ingestion, and real-time/batch reporting capabilities.
- AI-assisted (Cortex-driven) optimization of transformation logic and SQL workloads.
- Applied Tensai® and Claude-assisted development to speed up migration analysis and coding.

Benefits

- Reduced analysis & testing effort by 30% using Tensai®.
- Faster time-to-market with rapid onboarding of new sources/reports.
- Delivered a governed, medallion-based architecture to modernize enterprise reporting.
- Separation of compute & storage enabling independent, parallel ETL/BI/analytics development.
- Achieved scalability for ~120+ TB and high concurrency.

Healthcare, Life Sciences, & Insurance

25%

Savings in the annual contact center expenditure

Client

A Firm Specializing in Investments, Pensions, and Retirement Financial Services

Creating a contact center experience customers actually want to use

The client wanted to modernize engagement with digital channels, conversational AI, and analytics while still improving service quality. The aim was to lower cost without making the experience feel less human or less helpful.

Solution

- Progressively introduced digital channels, conversational AI, and speech and text analytics.
- Took end-to-end ownership of contact center modernization to simplify rollout.
- Used AI insights to support better agent decisions and more consistent service.
- Designed the model to future-proof the contact center as customer needs evolved.

Benefits

- Reduced annual contact center expenditure by 25%.
- Improved service quality by giving agents better insight during customer interactions.
- Created a more scalable engagement model with stronger digital adoption.
- Laid the groundwork for a more resilient, analytics-led customer experience.

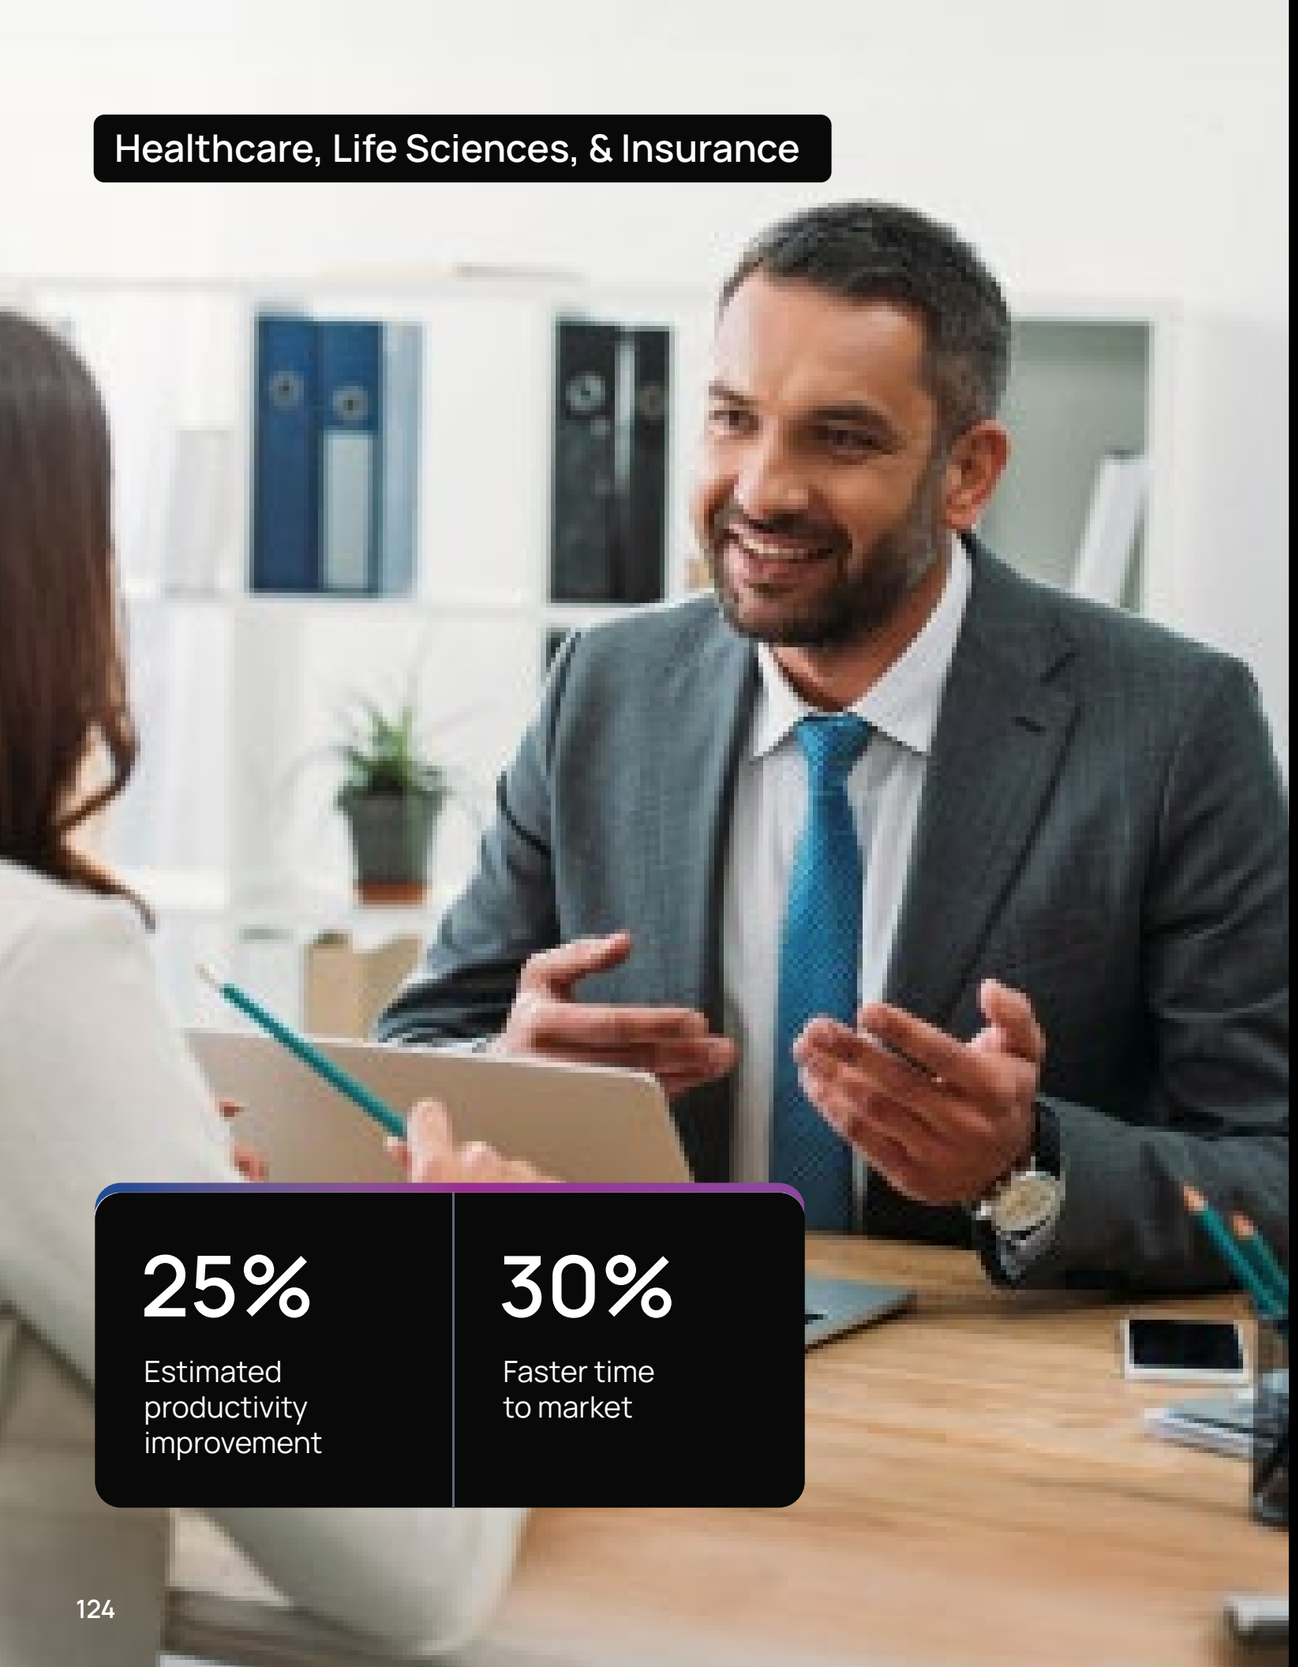

Client

A Leading UK Specialty Insurance & Reinsurance Provider

Building AI-assisted delivery velocity into digital underwriting

The client's pricing and underwriting relied on manual, legacy workflows that limited pricing agility and slowed policy issuance. To modernize underwriting on Python and .NET, Hexaware brought Claude and GitHub Copilot into its own development workflow to accelerate pricing modernization delivery.

Solution

- Migrated Excel-based pricing models to a Python-based digital model, using GitHub Copilot to:
  - Reverse-engineer accelerators to decode the pricing algorithms
  - Develop Python code, run unit tests, and conduct code reviews
  - Generate documentation
- Used Claude and GitHub Copilot within Hexaware's development workflow to accelerate delivery.
- Established velocity and AI-authored code as core productivity metrics.

Benefits

- Reusable boilerplate created for model migration framework.
- High accuracy in digital pricing model conversion.
- Model maintenance turnaround time was reduced due to standardization in model code.

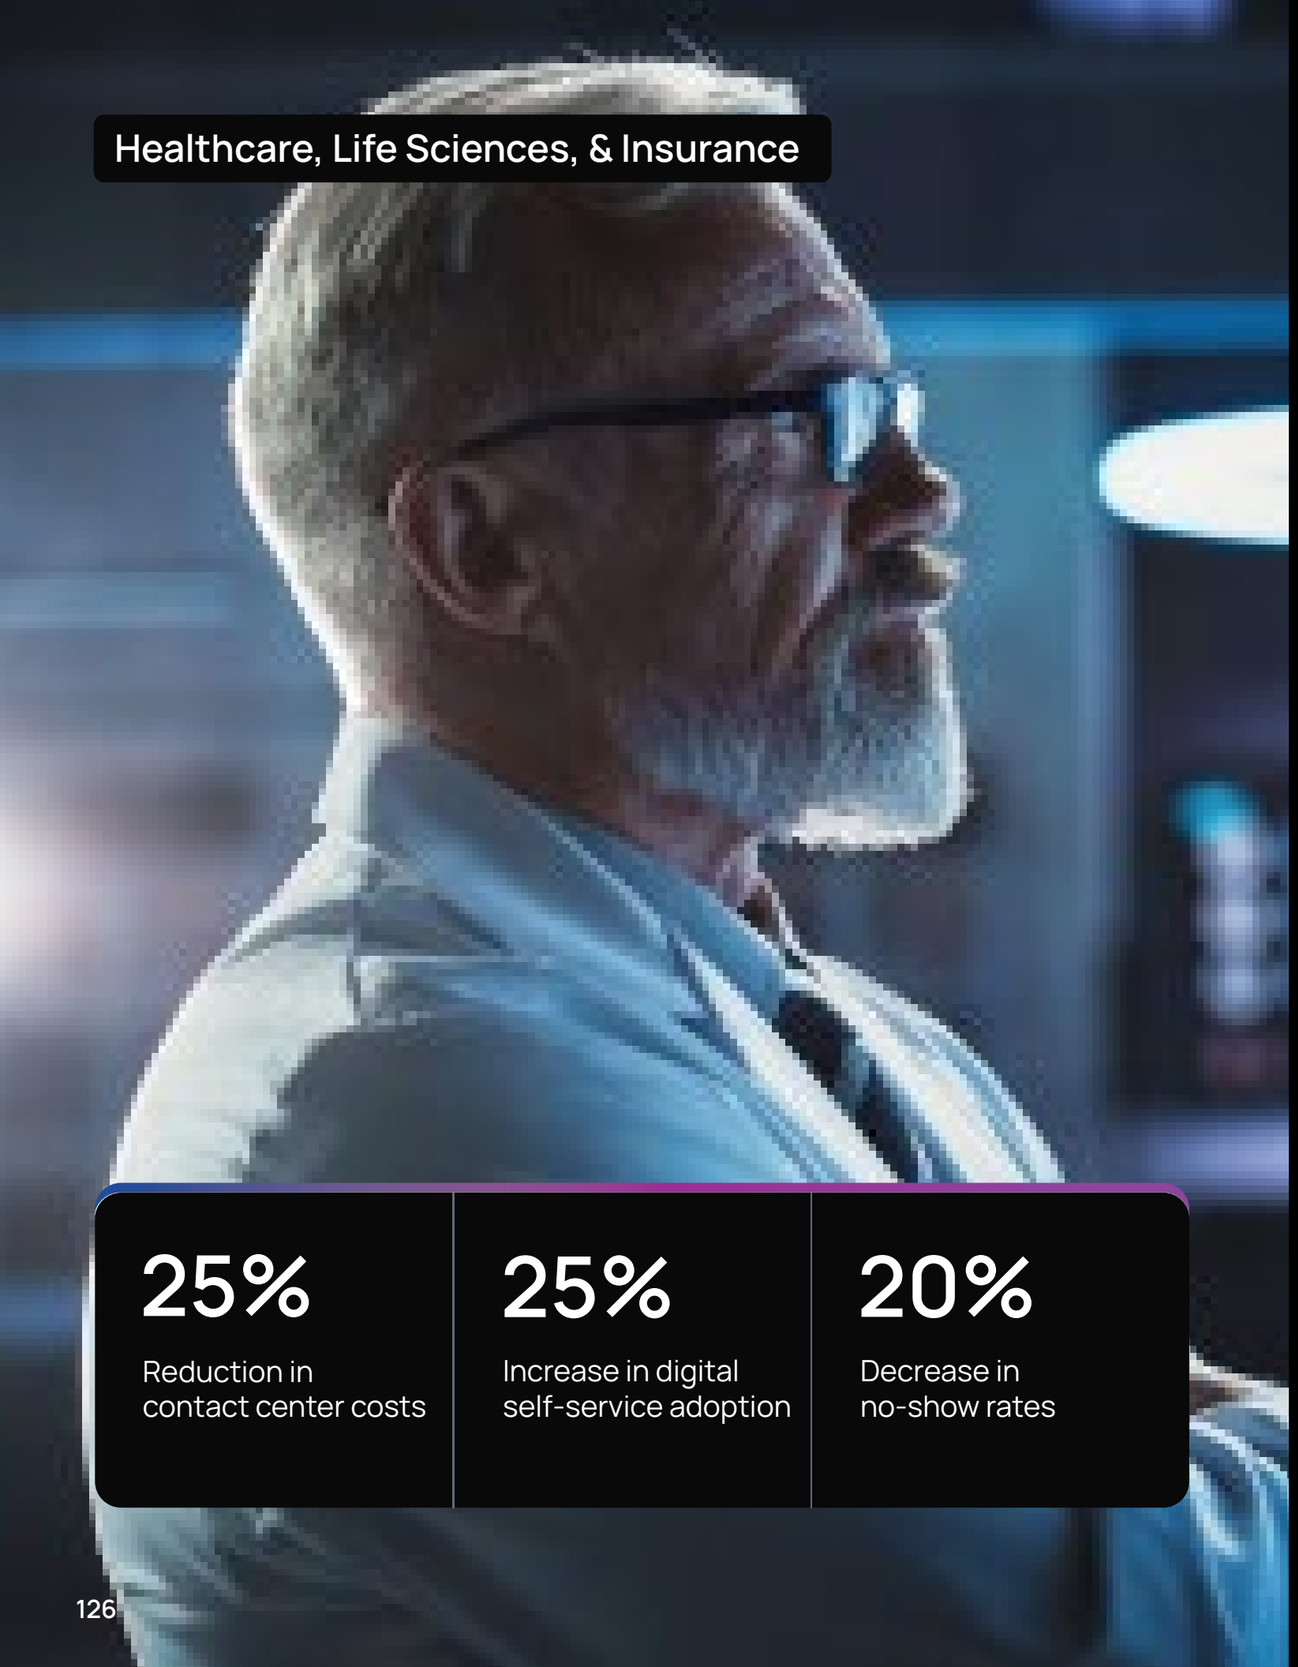

Client

An American Technology-Driven Healthcare Services Company Focused on Non-Emergency Medical Transportation

Making care access feel simpler across every channel

Siloed data, heavy voice dependence, and inconsistent member experiences were making it harder to scale support and proactive care. The client needed one connected engagement model that was easier for members to use and easier for teams to run.

Solution

- Unified voice, chat, and digital self-service in one omnichannel platform.
- Introduced multilingual AI virtual assistants for instant member support.
- Created a centralized data hub with real-time visibility into member journeys.
- Added intelligent routing and agent-assist capabilities integrated with existing dispatch and care systems.

Benefits

- Reduced contact center operating costs by 25% in the first year.
- Lowered no-show rates for transportation and home care appointments by 20%.
- Increased digital self-service adoption by 25%, reducing inbound call volume.
- Improved responsiveness and member satisfaction through more proactive, personalized engagement.

Financial Services

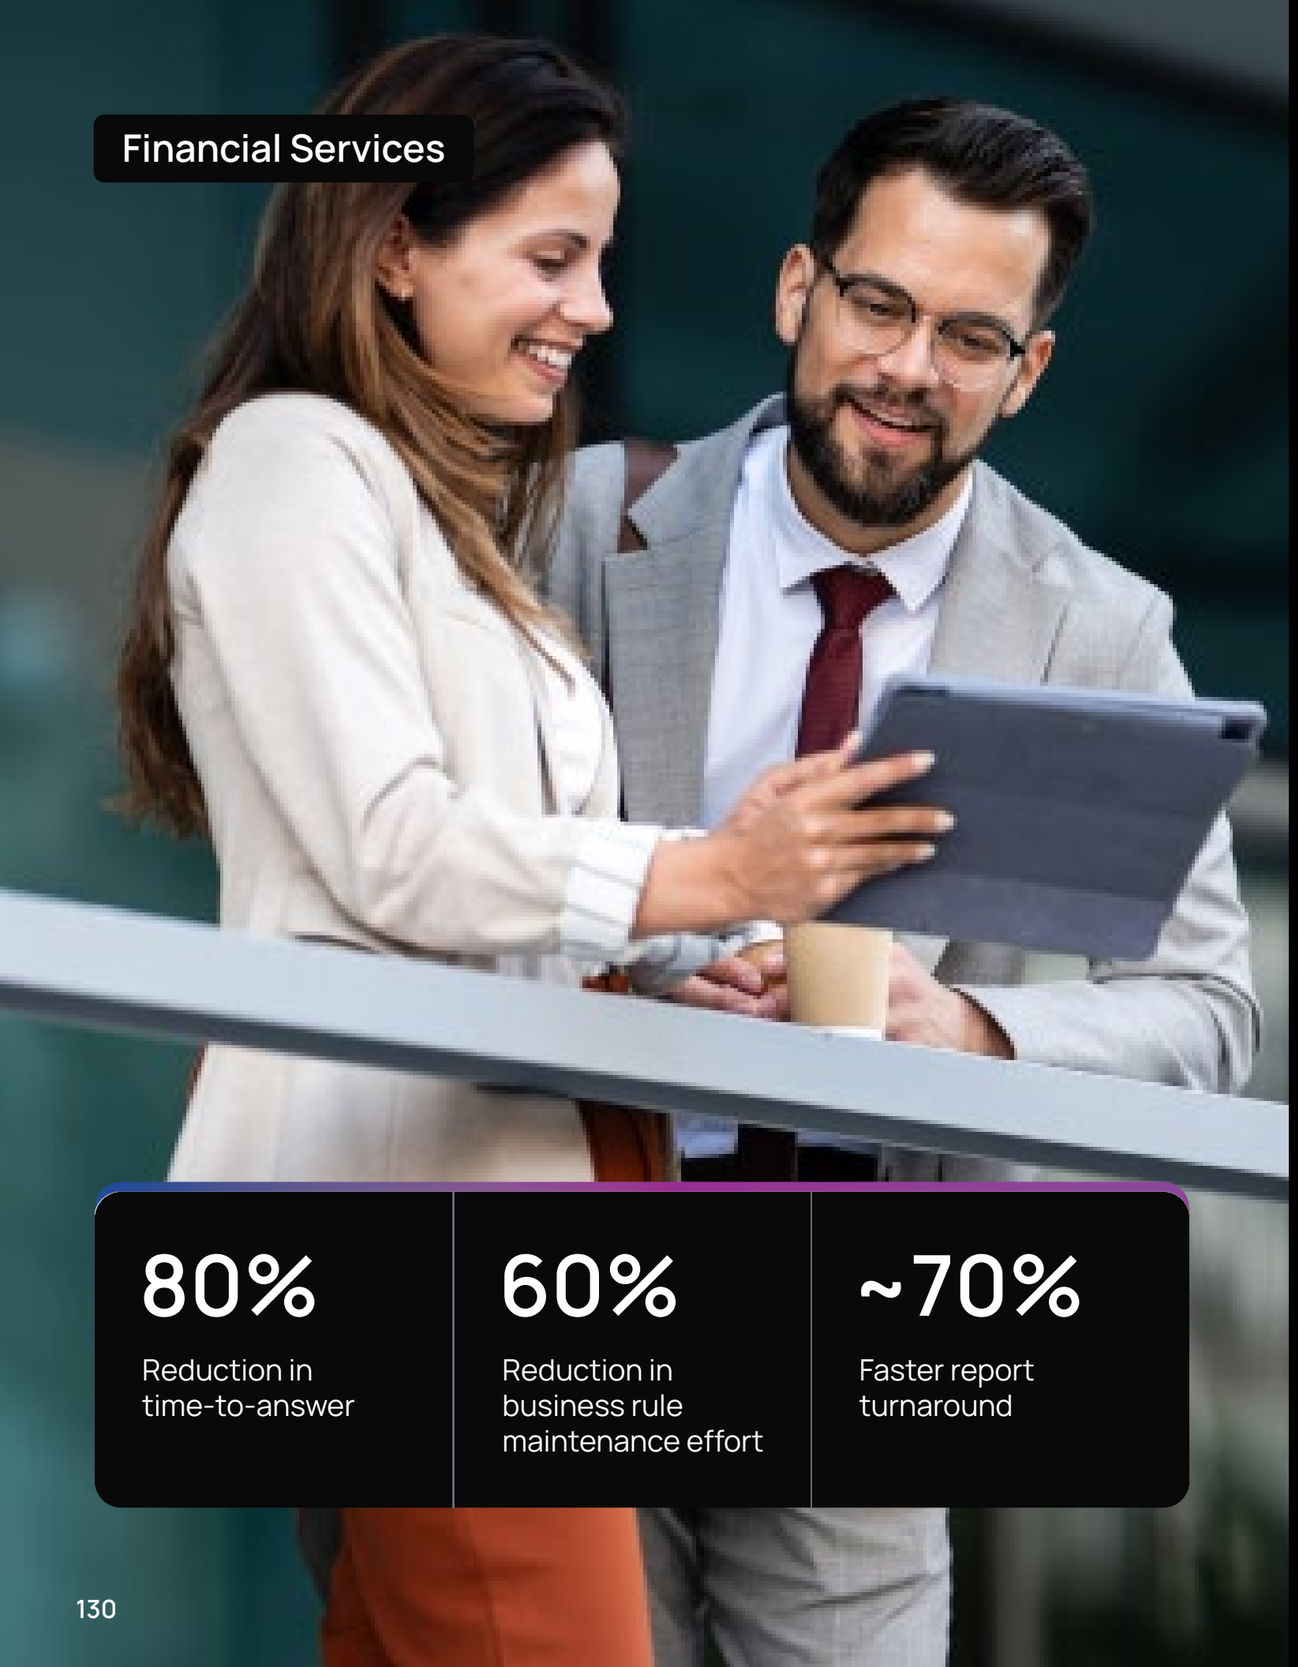

Financial Services

80%	**60%**	**~70%**
Reduction in time-to-answer	Reduction in business rule maintenance effort	Faster report turnaround

Client

The Largest Source of Mortgage Financing in the United States

Replacing BI complexity with answers people can actually get to

A sprawling BI environment made risk analytics harder than it needed to be, with duplicated logic, heavy maintenance, and limited self-service. The client needed faster access to trusted answers without rebuilding everything from scratch.

Solution

- Replaced fragmented dashboards with an agentic AI analytics platform using a multi-agent React framework.
- Enabled natural-language queries with governed SQL generation and validation.
- Centralized business logic in a metadata and glossary layer to improve consistency.
- Added semantic retrieval, full logging, and secure deployment within the client's cloud perimeter.

Benefits

- Reduced time to answer by 80%.
- Cut business rule maintenance effort by more than 60%.
- Improved report turnaround by about 70%.
- Made self-service analytics more practical without sacrificing auditability or control.

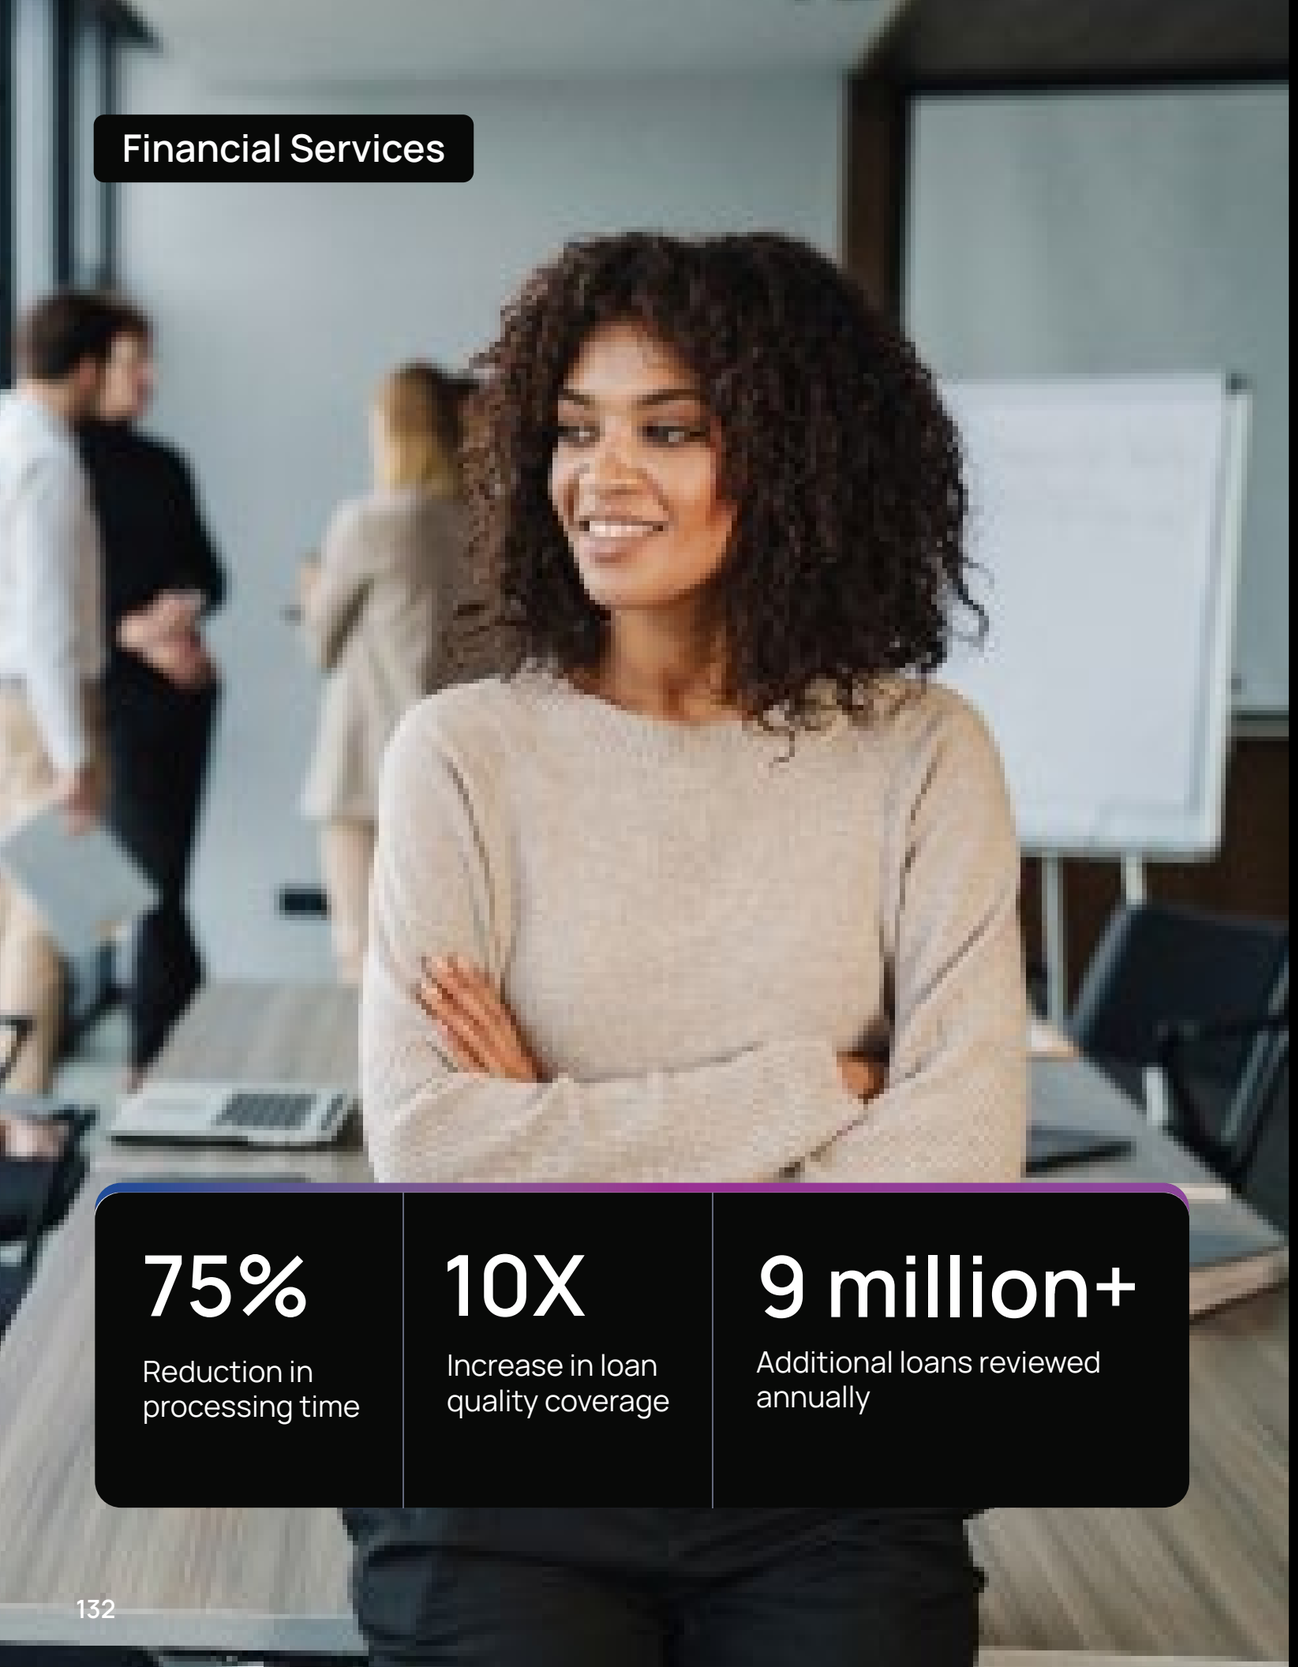

Client

The Largest Source of Mortgage Financing in the United States

Reviewing more loans, faster, before quality issues become risk

Manual review of lengthy loan files could not keep pace with volume or the need for consistent quality checks. The client needed a more scalable way to validate more loans without adding linear effort.

Solution

- Built an end-to-end AI platform for ingestion, extraction, validation, and discrepancy generation.
- Used advanced optical character recognition (OCR) to process large (up to 800 pages), complex multi-format loan packages.
- Applied an agentic AI pipeline to extract entities, run business rules, and classify discrepancies.
- Added a human-audit portal and secure AWS VPC deployment for oversight and control.

Benefits

- Reduced processing time by 75%.
- Increased loan quality coverage by 10x.
- Enabled review of 9 million+ additional loans annually.
- Improved consistency and speed in a mission-critical quality assurance workflow.

Client

A Firm Specializing in Wealth Management, Commercial Banking, and Investor Services

Taking client onboarding from fragmented to frictionless

The client wanted onboarding operations that were more accurate, compliant, and efficient without forcing a rip-and-replace of existing systems. The priority was to simplify the process for clients while reducing the operational burden behind the scenes.

Solution

- Implemented a custom AI-driven KYC and AML solution.
- Automated front-to-back onboarding processes, including claims-related NLP workflows.
- Enhanced fee, billing, and invoicing operations through broader automation.
- Integrated the solution with existing ERP and CRM environments to support adoption.

Benefits

- Reduced software licensing costs by 50%.
- Accelerated customer onboarding by 25%.
- Lowered client onboarding process costs by 70%.
- Created a more consistent, compliant onboarding experience across shared services operations.

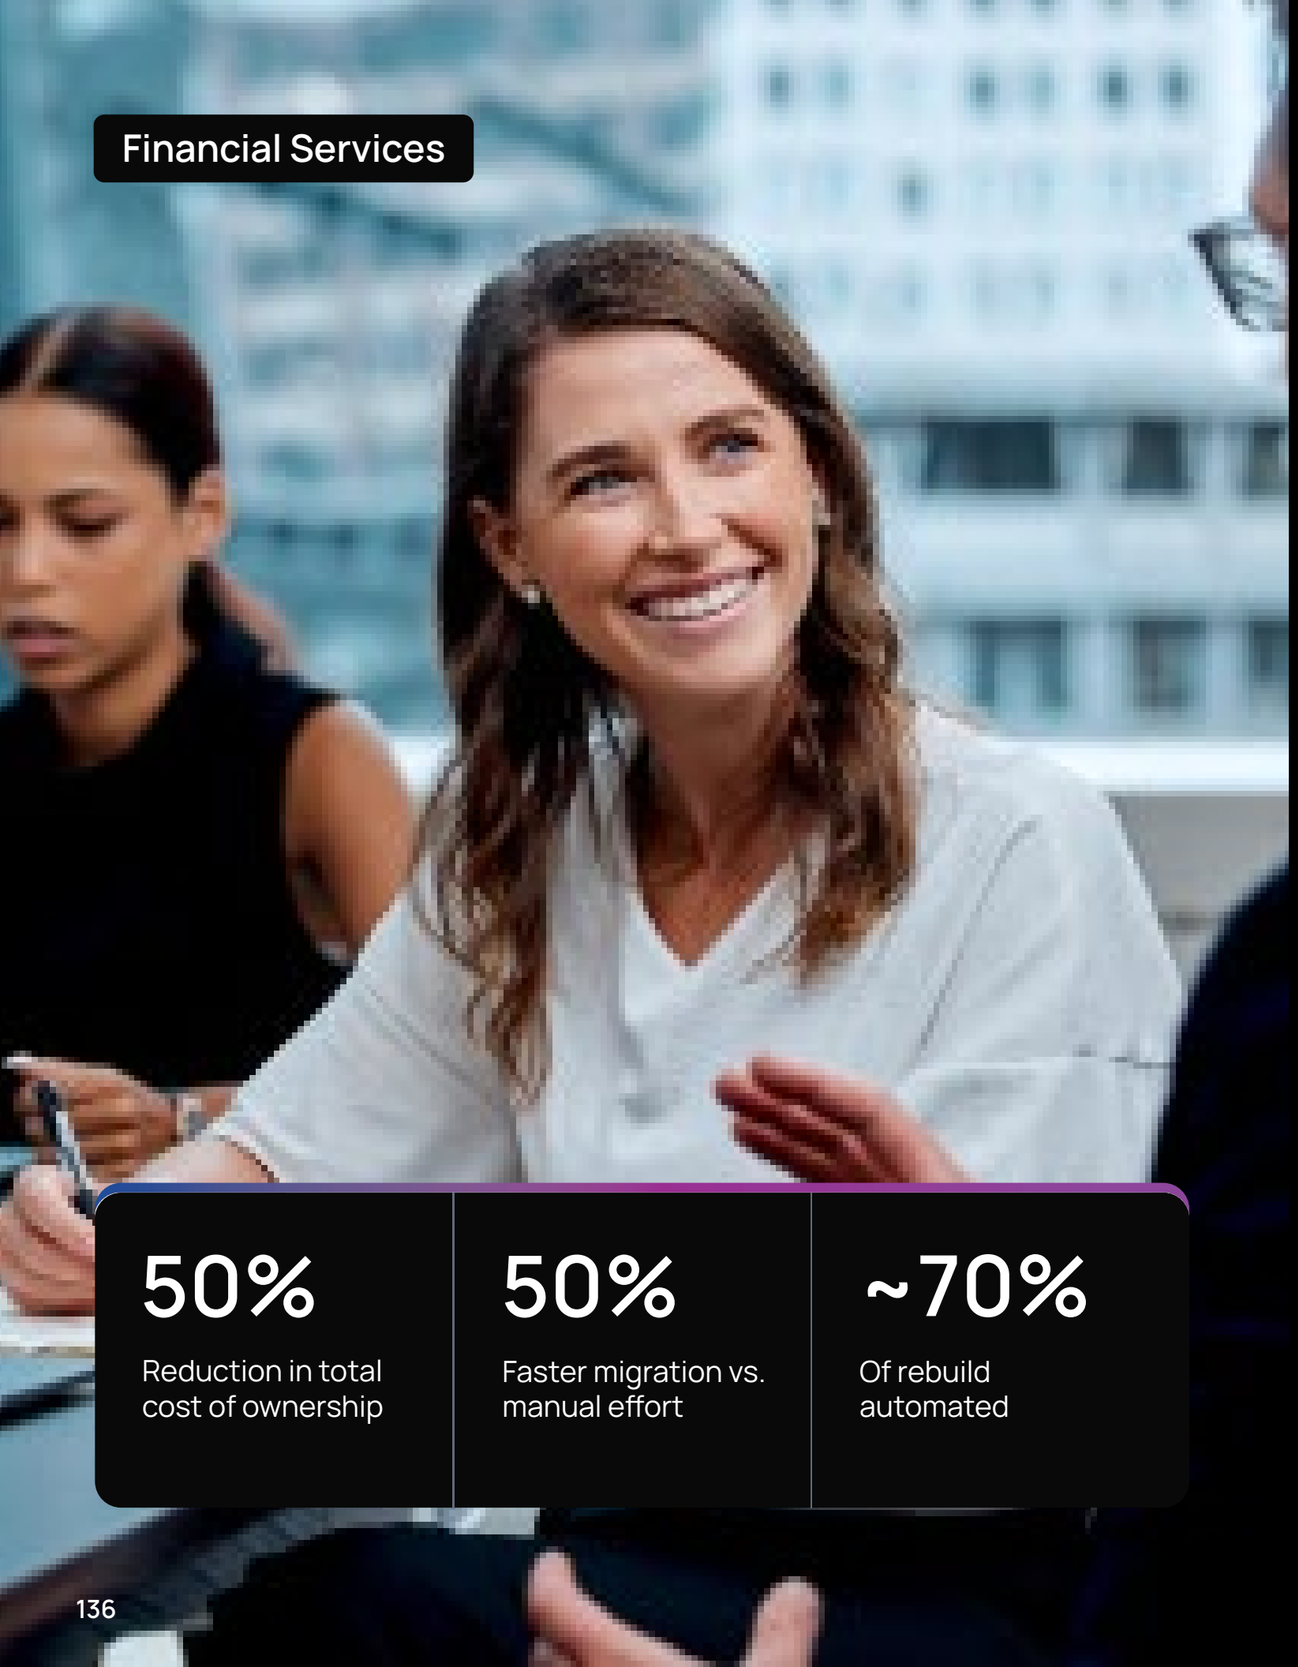

Client

A Leading US Mortgage Firm

TIBCO modernization cuts total cost of ownership by 50% with cloud-native microservices

A business-critical TIBCO BusinessWorks 5.0 application, running for a decade and integrating multiple upstream and downstream systems, had become a costly liability. High licensing and maintenance costs, vendor lock-in, limited cloud readiness, and integration complexity slowed change and made the stack hard to scale. A conventional rewrite would have meant higher manual effort, longer timelines, and greater delivery risk.

Solution

- Replaced licensed TIBCO middleware with open-source Java Spring Boot microservices via the Zero License model.
- Ran deterministic reverse engineering to fully map processes, dependencies, and logic before conversion.
- Generated target architecture, blueprints, and playbooks autonomously through agentic blueprinting.
- Rebuilt processes as reusable Spring Boot microservices with REST APIs secured by OAuth 2.0.
- Validated functional equivalence with hash-based row validation and generated automated tests.

Benefits

- Eliminated proprietary middleware dependency for a client-owned stack.
- Preserved business behavior and downstream flows via retained JMS messaging.
- Created a reusable blueprint extensible to other TIBCO processes.

50%	50%	~70%
Reduction in total cost of ownership	Faster migration vs. manual effort	Of rebuild automated

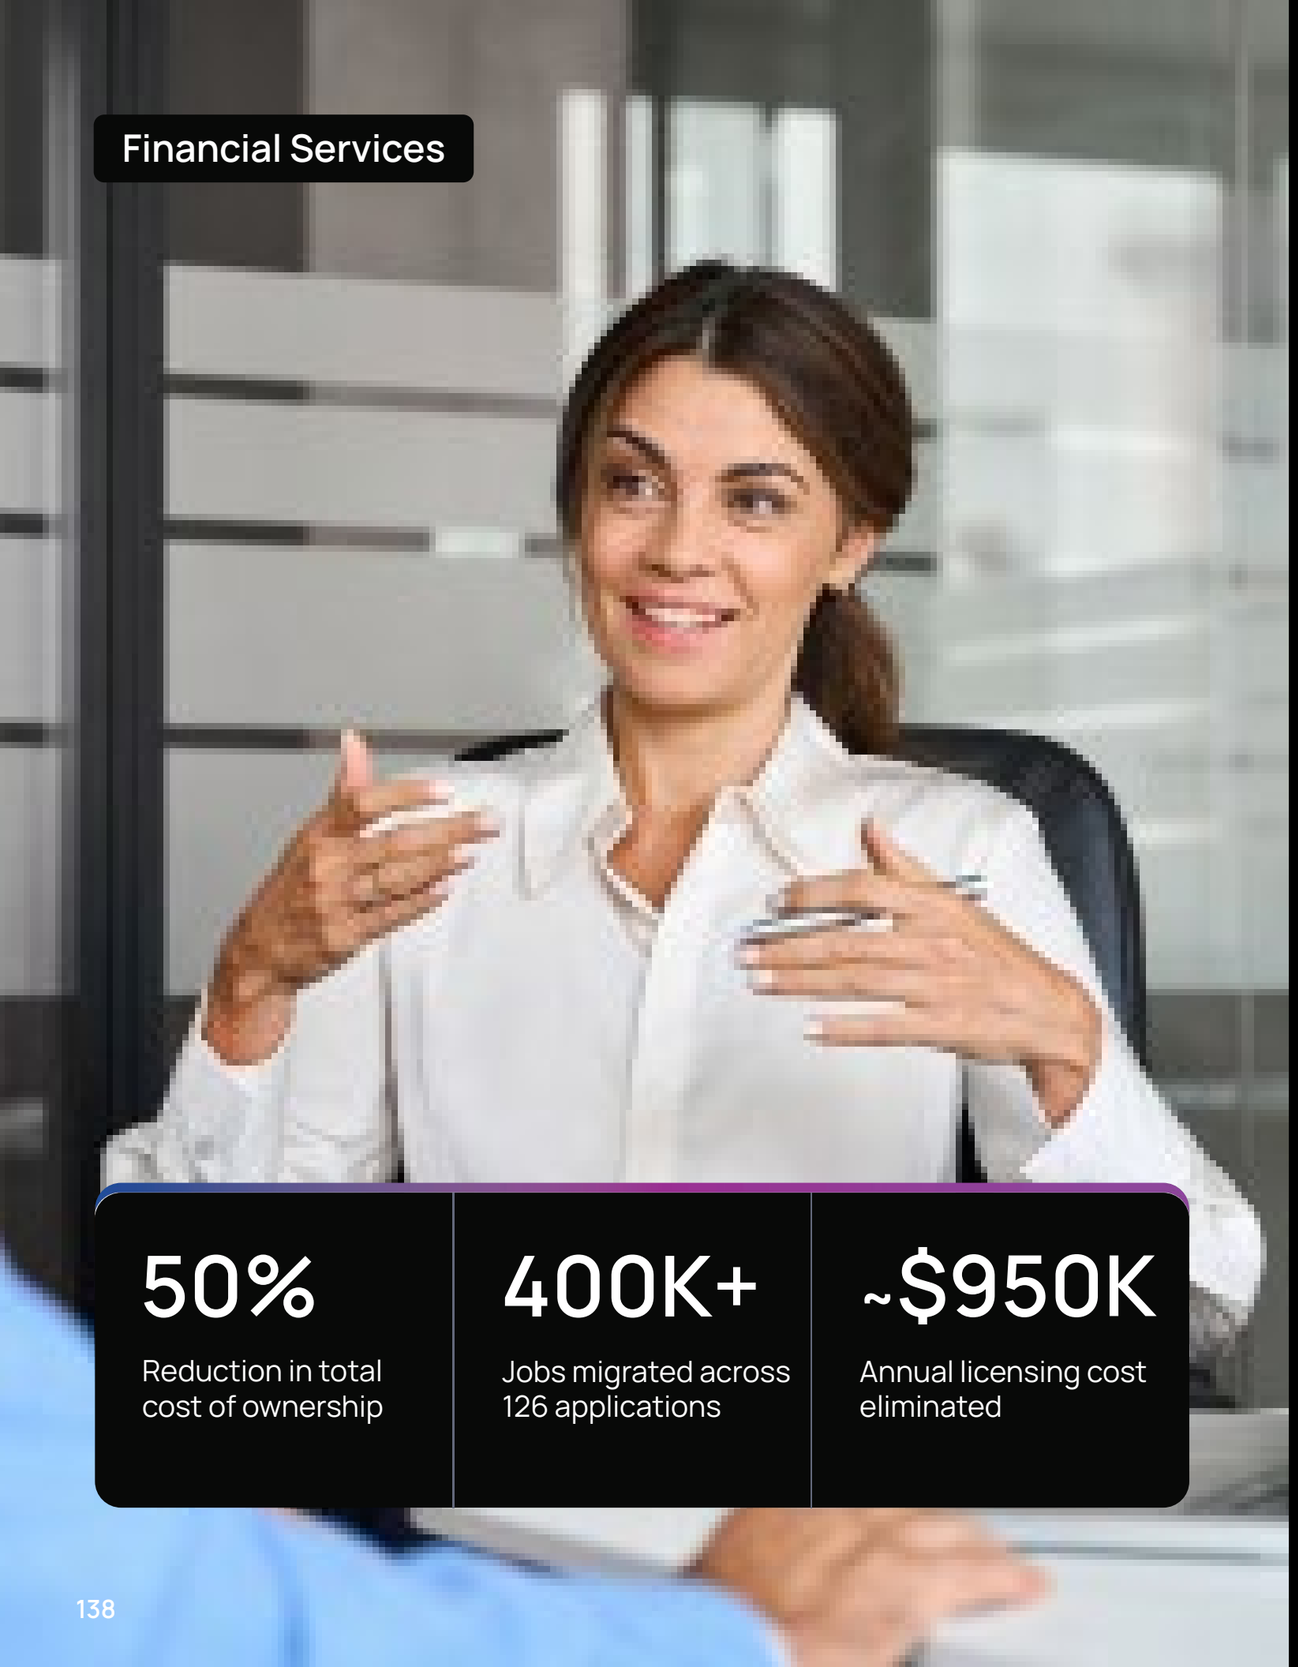

Client

A Large US Mortgage Enterprise

Replacing a licensed job scheduler with open-source Airflow to cut TCO by 50%

The enterprise ran its workload scheduling on Autosys, a licensed legacy platform spanning over 400,000 automated jobs across 126 applications, including critical loan management and data operations. Licensing, infrastructure, and maintenance ran to roughly $1 million a year with steep renewal escalation risk, while the platform's lack of native AWS support, rigid job chains, and dependency on scarce specialist skills held back the cloud roadmap.

Solution

- Replaced licensed Autosys with open-source Apache Airflow on AWS (MWAA) via the Zero License model.
- Ran deterministic reverse engineering to map all jobs, JILs, scripts, calendars, and dependencies before conversion.
- Generated target architecture, blueprints, and Airflow DAGs autonomously through agentic blueprinting.
- Rebuilt the estate as cloud-native Airflow, with MCP binding AI agents to enterprise guardrails and Terraform provisioning infrastructure.
- Validated functional equivalence through parallel Autosys runs and smoke testing before phased cutover.

Benefits

- Removed proprietary licensing, infrastructure cost, and reseller dependency.
- Enabled faster change and release cycles on code-driven CI/CD.
- Delivered horizontal scalability and stronger reliability through managed Airflow.
- Reduced dependency on scarce specialist Autosys skills.

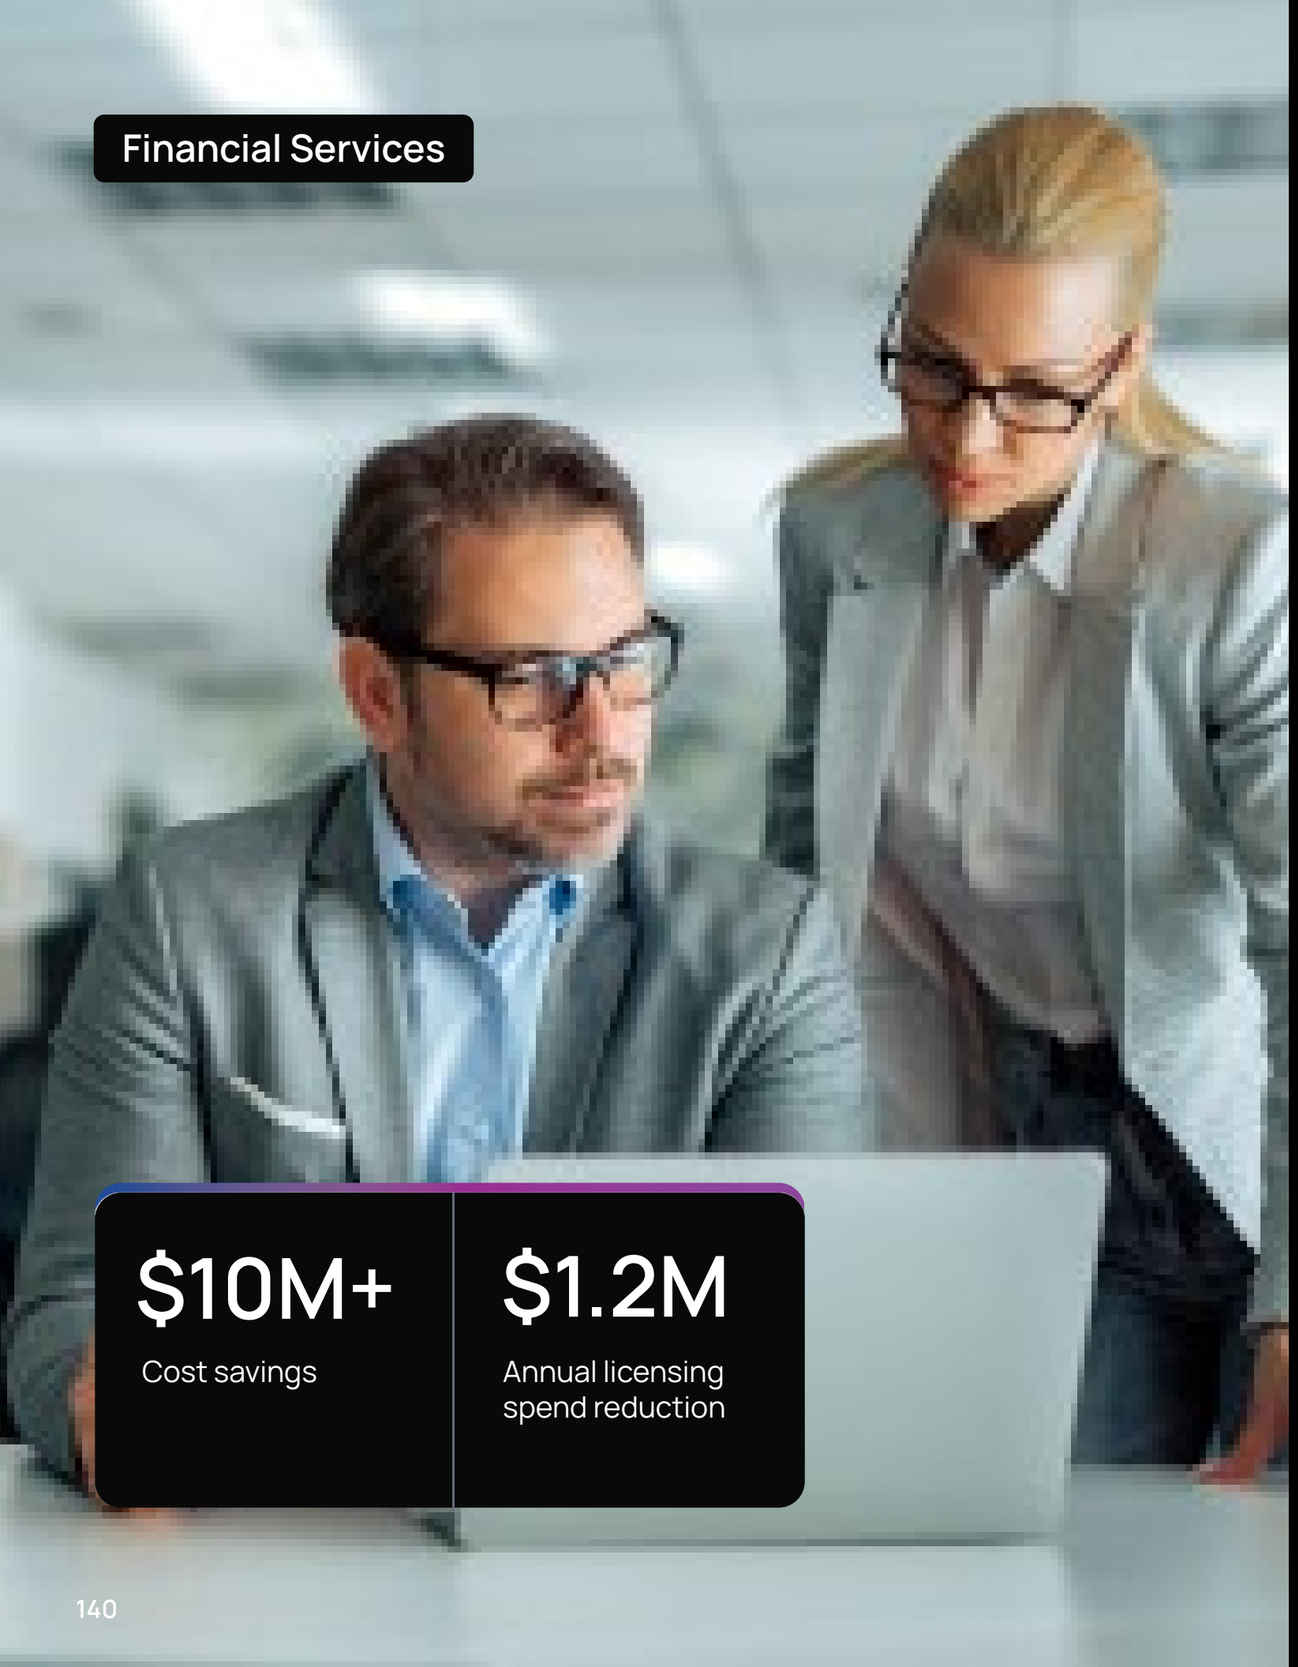

Client

A Leading Secondary Mortgage Enterprise

Aligning a Tier-0 application to a cloud-native strategy

A mission-critical, Tier-0 external-facing loan data correction application had been successfully built on a low-code platform. As the enterprise matured its cloud-native, microservices modernization strategy, it sought to bring this application in-house onto an architecture it fully owned and controlled, gaining the flexibility to patch independently within its 30-day vulnerability SLA, reduce recurring licensing spend, and align the application with its broader engineering roadmap.

Solution

- Delivered a complete cloud-native rewrite via the Zero License model, using Java, Spring Boot, and Angular microservices on AWS.
- Ran deterministic reverse engineering to fully document the existing application and build a verified knowledge index before conversion.
- Generated target architecture, blueprints, and execution playbooks autonomously through agentic blueprinting.
- Rebuilt the entire application in 30 weeks, with MCP binding AI agents to the client's architecture, security, and compliance guardrails.
- Mapped every component to a cloud-native equivalent across UI, data model, workflows, and database layers.

Benefits

- Reduced recurring licensing spend by moving to a fully owned open-source stack.
- Gained full compliance control through independent patching inside the 30-day SLA.
- Streamlined operations into a leaner cloud-native model across fewer environments.

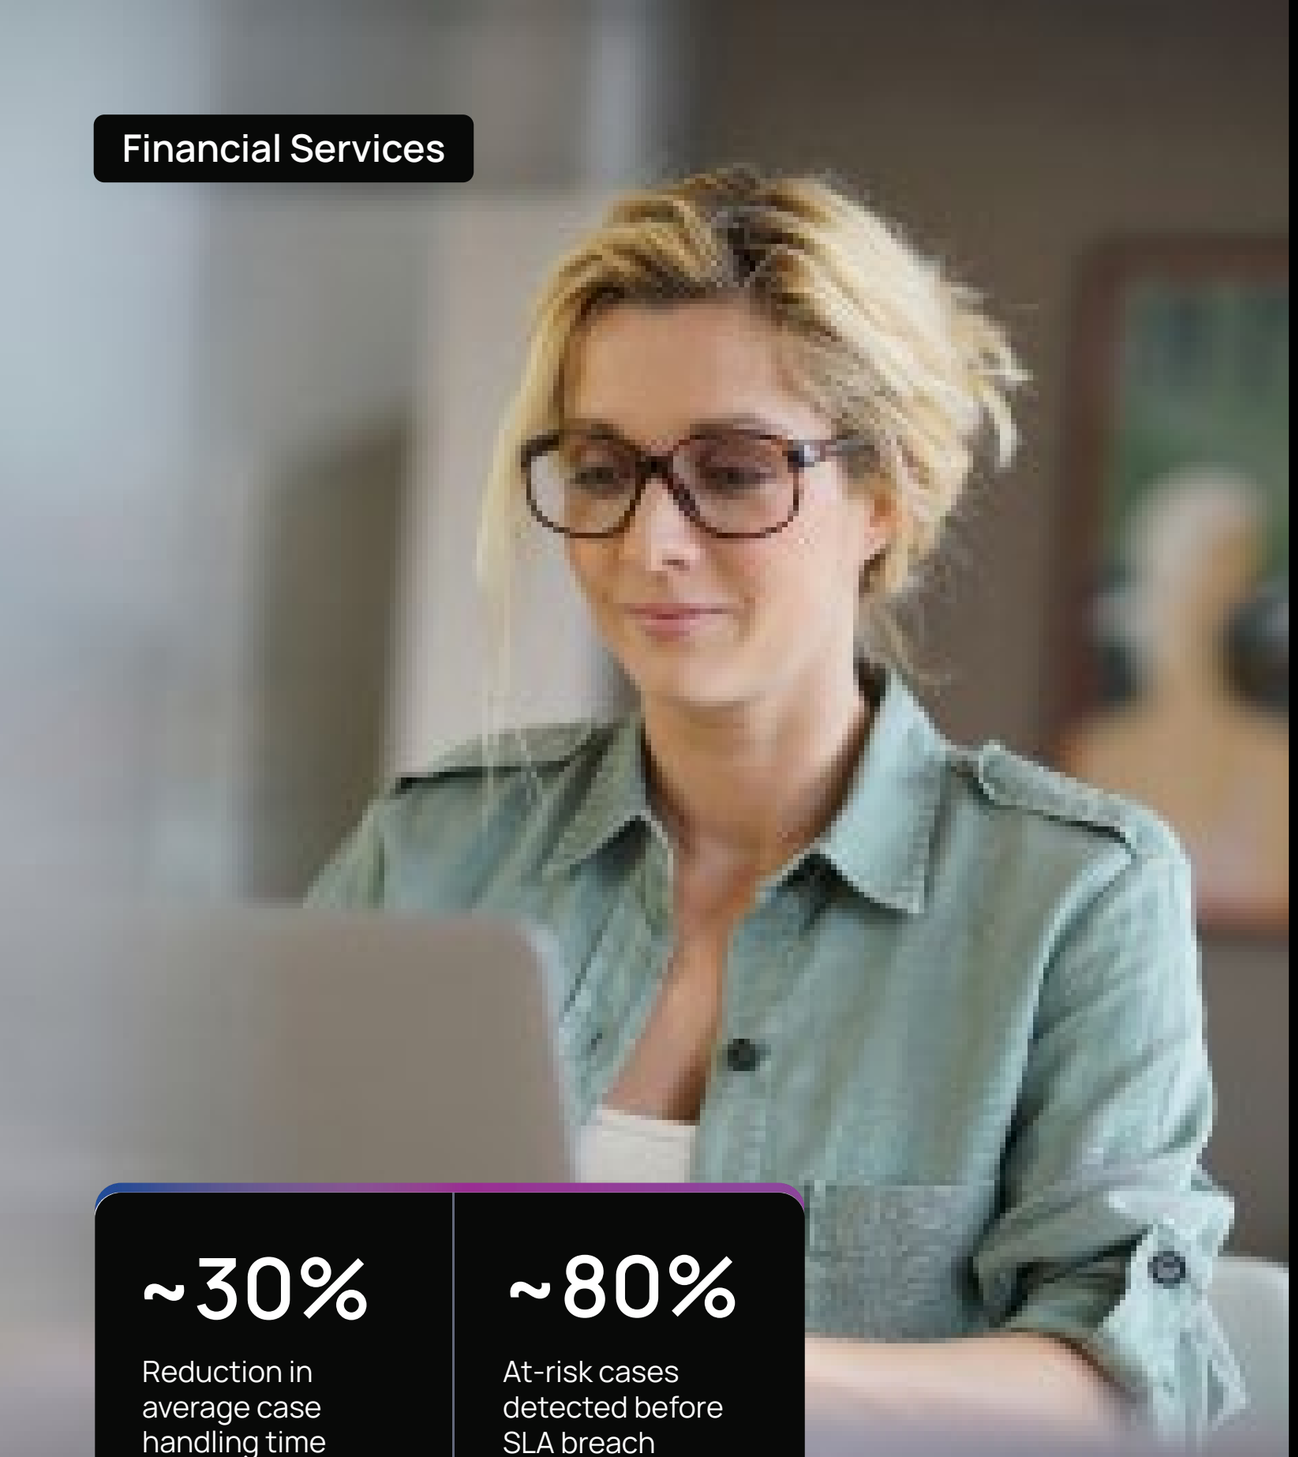

Financial Services

~30%

Reduction in average case handling time

~80%

At-risk cases detected before SLA breach

Client

A Leading Global Financial Markets Infrastructure Data & Analytics Provider

Preventing escalations, closing context gaps, and unlocking revenue with four Salesforce Agentforce AI agents

Across a high-volume, round-the-clock support operation spanning multiple time zones and asset classes, the client faced four compounding problems: context loss at shift handovers, reactive escalation management, inconsistent agent coaching, and revenue signals going undetected in service interactions.

Solution

- Four AI agents built natively on Salesforce Agentforce.
- Shift Handover Agent auto-summarizes case history, actions, and sentiment on every handover.
- Escalation Prediction Agent scores risk on every inbound message before escalation occurs.
- Opportunity Identification Agent detects upsell signals and auto-creates linked Opportunity records.
- Case Quality Agent surfaces next best actions for agents and coaching notes for managers.

Benefits

- Zero context loss across shift changes, reducing handle time and repeat customer effort.
- At-risk cases identified and acted on before SLA breach or customer escalation.
- Support converted into a revenue channel through automated opportunity surfacing.
- Structured, data-driven coaching on every closed case without added management overhead.

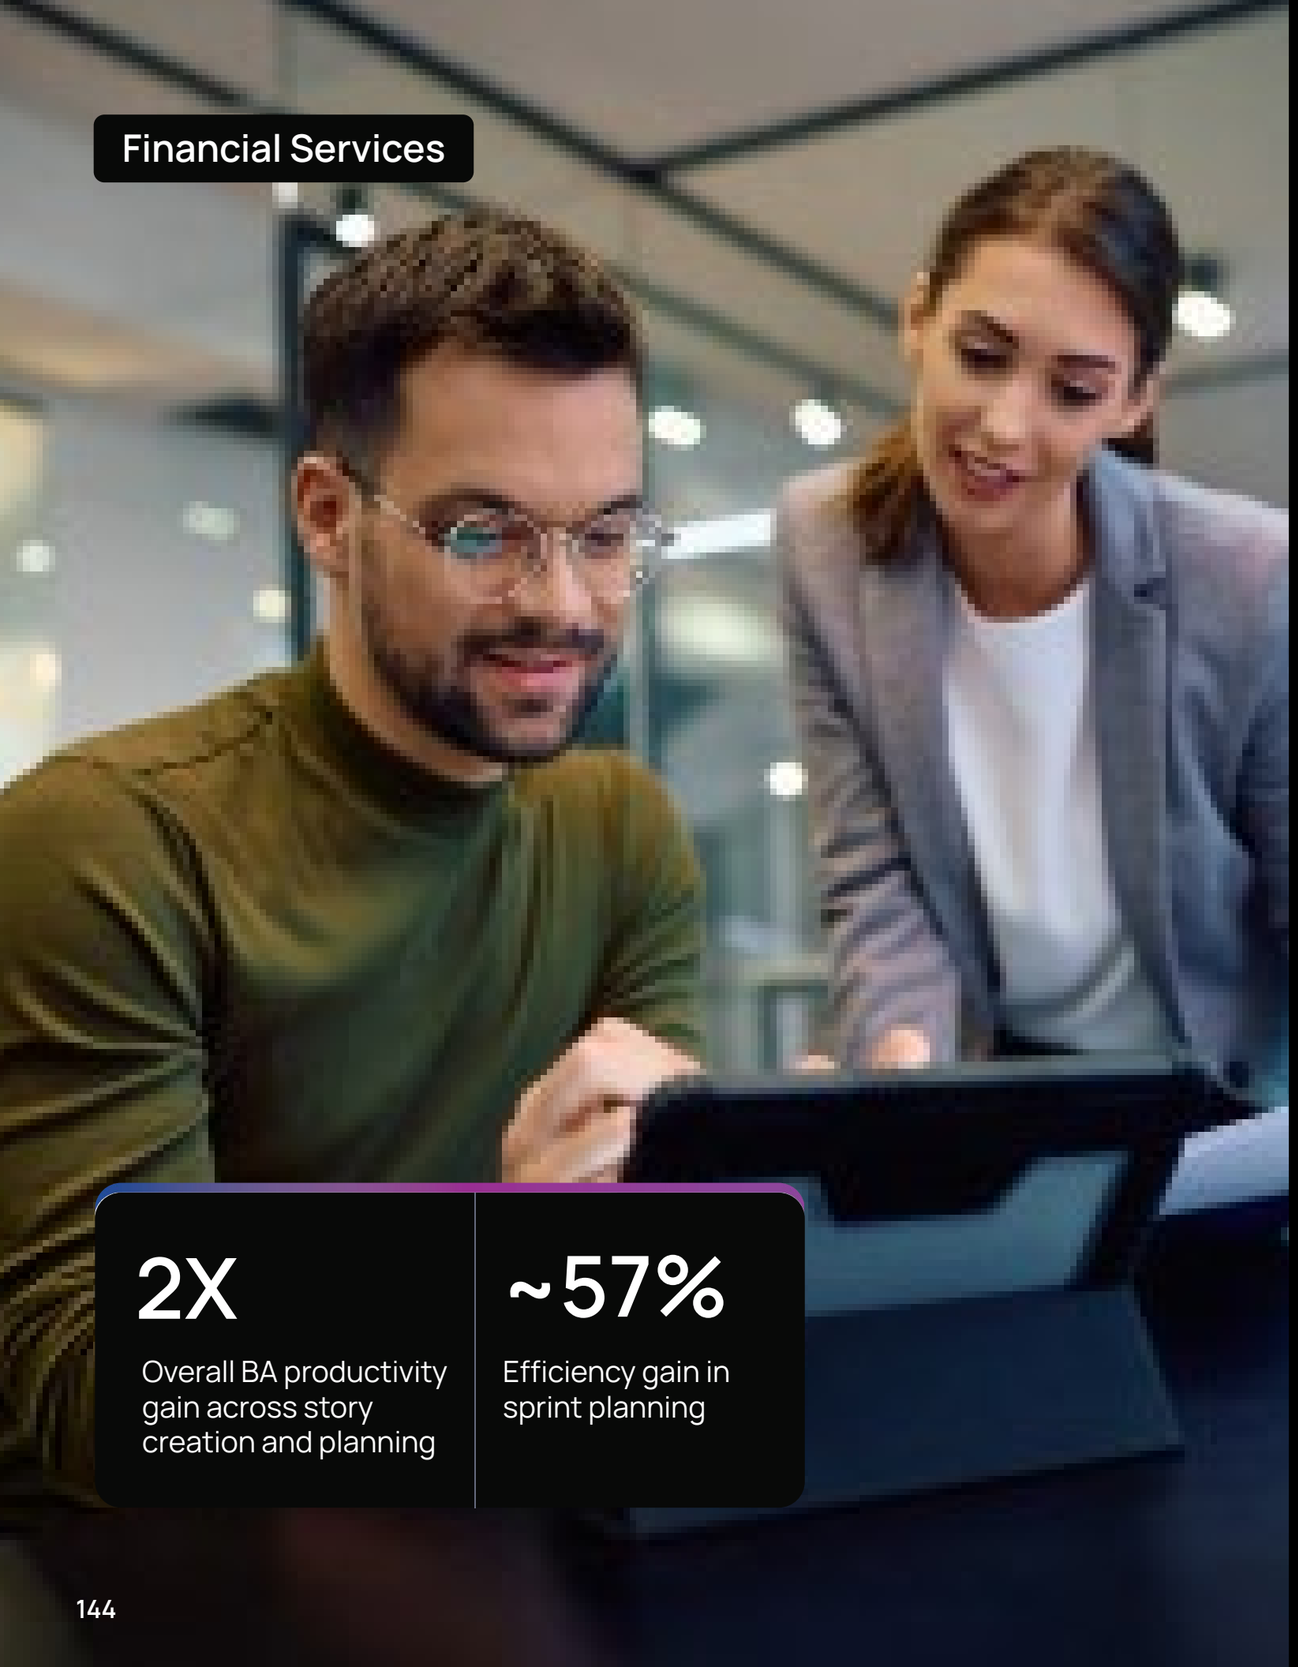

2X

Overall BA productivity gain across story creation and planning

~57%

Efficiency gain in sprint planning

Client

A Leading Global Financial Index & Data Provider

Embedding AI into the SDLC to double business analysis productivity

The client's application delivery team ran a functional but manually intensive process. Story enrichment, Jira hierarchy creation, sprint sizing, and planning all depended on the business analyst's individual recall and repeated mid-sprint involvement, creating overhead that worked for one team but could not scale across the wider organization.

Solution

- Built a self-proposed suite of AI-powered BA tools on GitHub Copilot within VS Code, using only org-licensed platforms with no new infrastructure.
- Delivered a guided story enrichment workflow taking raw requirements through enrichment, validation, and direct Jira creation in a single agent session.
- Automated Jira hierarchy creation with parent keys, dependency links, and bottom-up sizing injected across Portfolio Epics, Epics, and Stories.
- Enforced 6 quality dimensions at authoring via the BA Output Validator, catching gaps before Jira submission and developer sprint testing.

Benefits

- Reduced BA mid-sprint clarifications and post-creation Jira corrections significantly.
- Preserved QA capacity for release regression by identifying more story gaps at authoring.
- Enabled rollout to other delivery teams via scalable team profile configuration.
- Codified BA quality standards in agent instructions, applying uniformly across all stories.

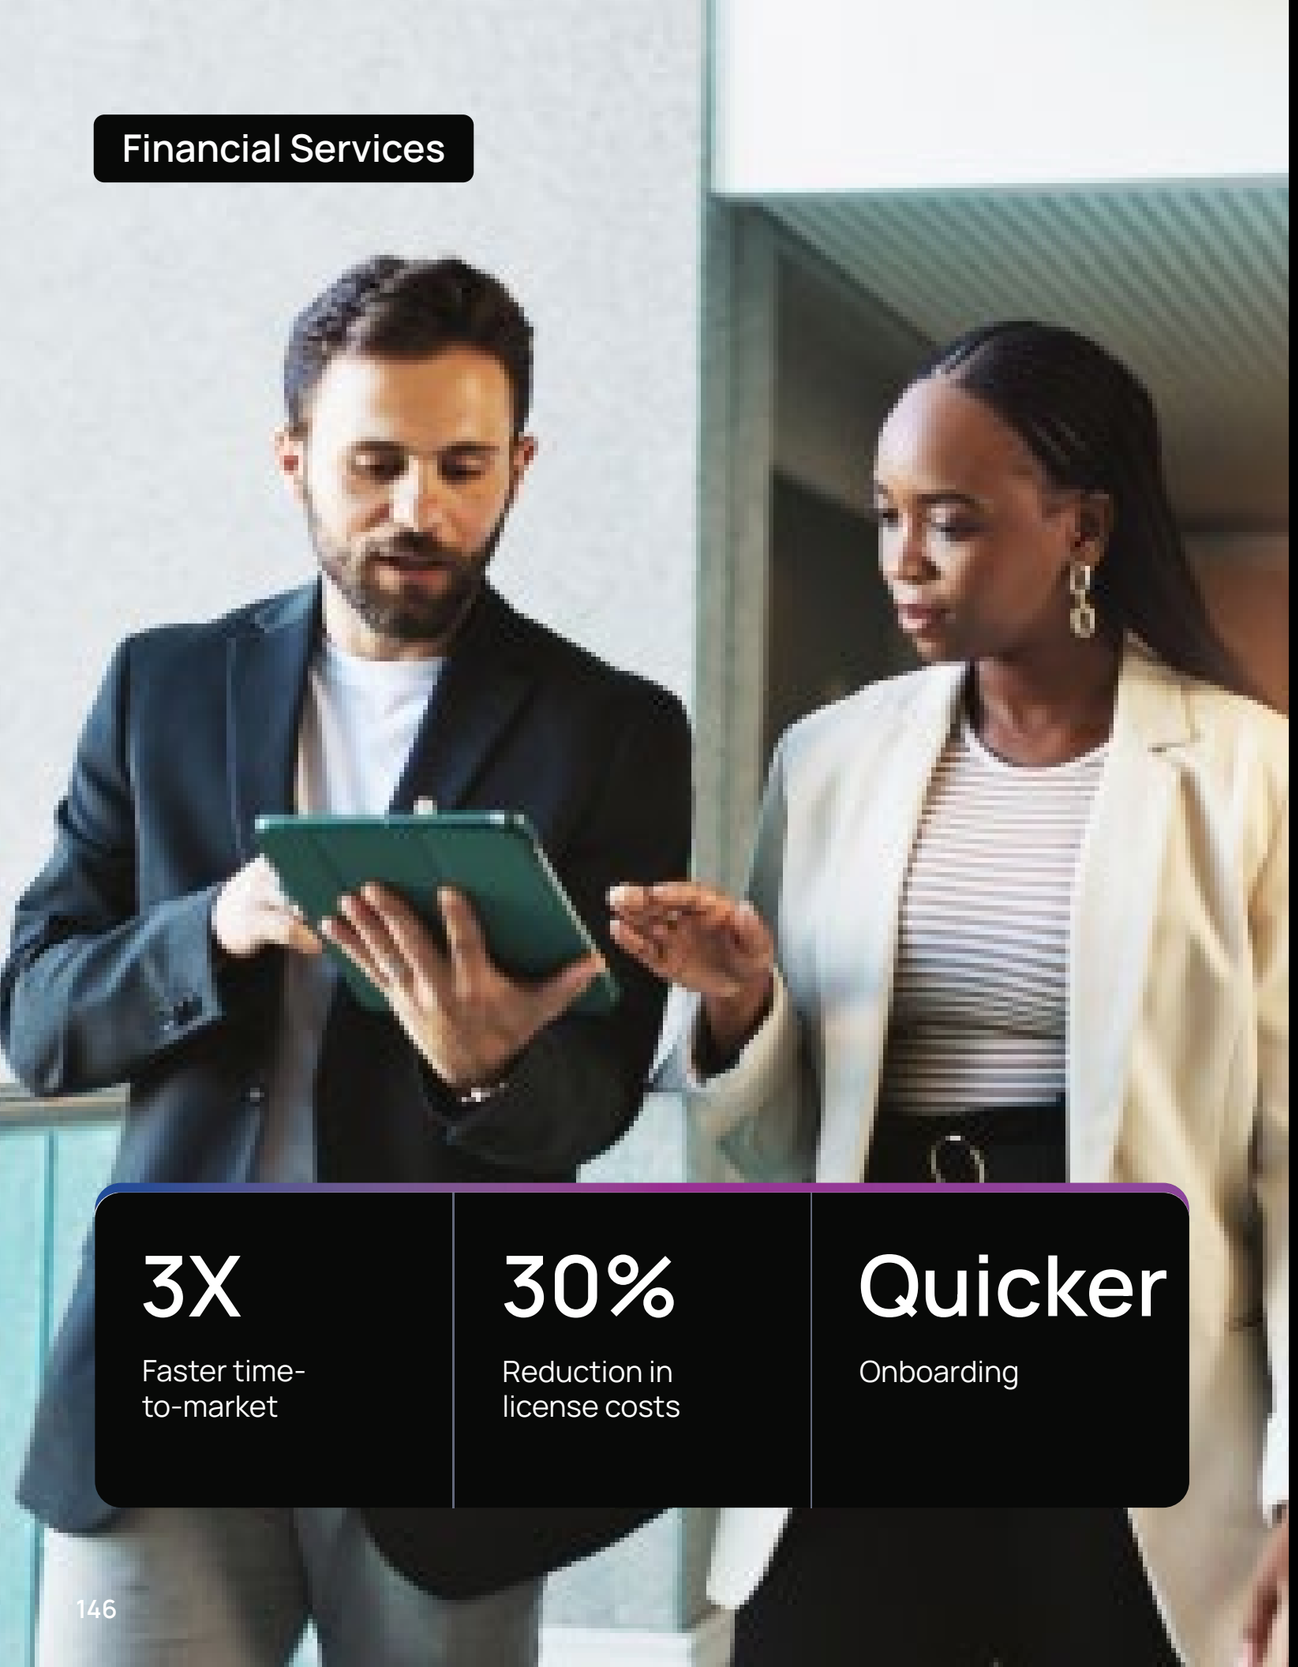

Client

A Leading Global Asset Manager

Empowering agility with AI-powered, Zero License transformation

The client faced complex, rigid customizations in their internal wealth management platform, leading to inefficiencies, high costs, vendor dependency, delayed feature launches, and suboptimal user experience for advisors and clients.

Solution

- Eliminated SaaS licensing costs with an AI-driven, cloud-native Advisor Portal.
- Used deterministic forensics to capture business logic, dependencies, and a validated knowledge index without large language model (LLM) hallucination risks.
- Automated generation of architecture documents, business requirement documents (BRDs), and delivery roadmaps using agentic blueprinting.
- Rebuilt the portal using GitHub Copilot and model context protocol (MCP), ensuring rapid delivery, scalability, compliance, and architectural alignment.
- Enabled a multi-tier access model with date-driven permissions and audited admin impersonation.

Benefits

- **Cost Reduction:** Eliminated recurring SaaS licensing fees and reduced future enhancement costs.
- **Faster Time-to-Market:** Leveraged AI-driven automation and containerized architecture for accelerated deployment.
- **Enhanced User Experience:** Delivered a clean, intuitive portal enabling advisors to access reports and documents effortlessly.
- **Quicker Onboarding:** Simplified understanding of code structures and patterns for new developers.

Retail

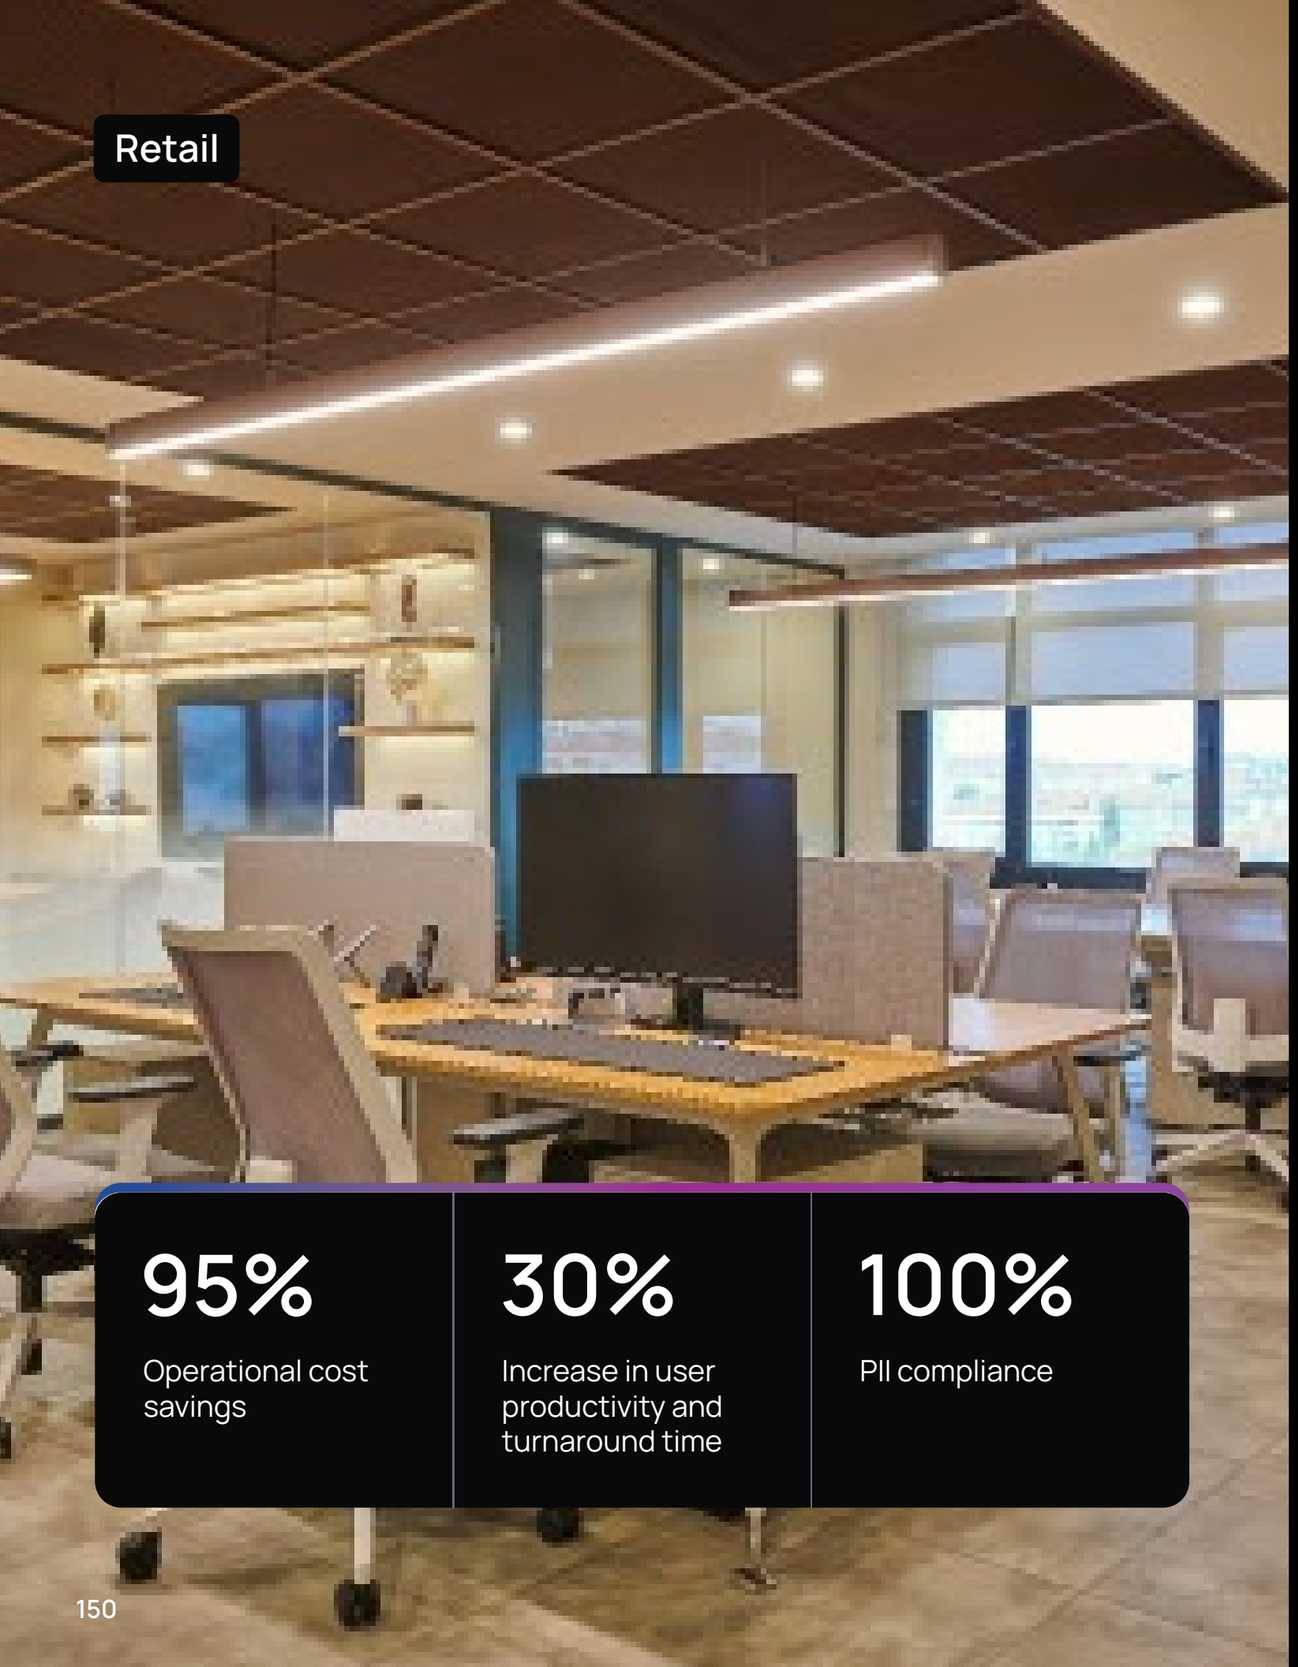

95%	30%	100%
Operational cost savings	Increase in user productivity and turnaround time	PII compliance

Client

A Workplace and Residential Furnishings Company

Giving teams a secure enterprise GPT they will actually use

Employees needed a safer, simpler way to access enterprise knowledge without creating new data security concerns. The opportunity was to combine usability, governance, and insight in one internal AI experience.

Solution

- Built a user-friendly enterprise GPT experience on Azure OpenAI.
- Added data retention, conversation tracking, and secure handling within the company network.
- Integrated Azure Active Directory for secure identity and access management.
- Enabled model selection and analytics to support performance monitoring and continuous improvement.

Benefits

- Reduced operational costs by 95%.
- Improved user productivity and turnaround time by 30%.
- Achieved 100% compliance for PII data handling.
- Created secure, self-serve access to enterprise knowledge for marketing and related teams.

75%	37%	136%
Increase in revisits for engaged customers	Increase in yearly spend	Increase in click-through rate

Client

A US-Based Convenience Store and Gas Station Chain

Turning scattered customer data into better-timed engagement

Customer data was spread across 250+ applications and siloed systems, making it hard to personalize effectively or move quickly on insight. The client needed a more unified view of behavior to increase engagement and spend.

Solution

- Implemented a personalized engagement solution triggered by micro-moments in customer behavior.
- Created an AI-generated fingerprint for each customer to predict propensity across those moments.
- Recommended engagement strategies designed to maximize spend and revisit behavior.
- Analyzed data across offline purchases, mobile apps, fuel stations, and loyalty systems.

Benefits

- Increased revisits among engaged customers by 75%.
- Raised yearly spend by 37%.
- Lifted click-through rates by 136%.
- Reduced marketing waste by targeting the right segments at the right time.

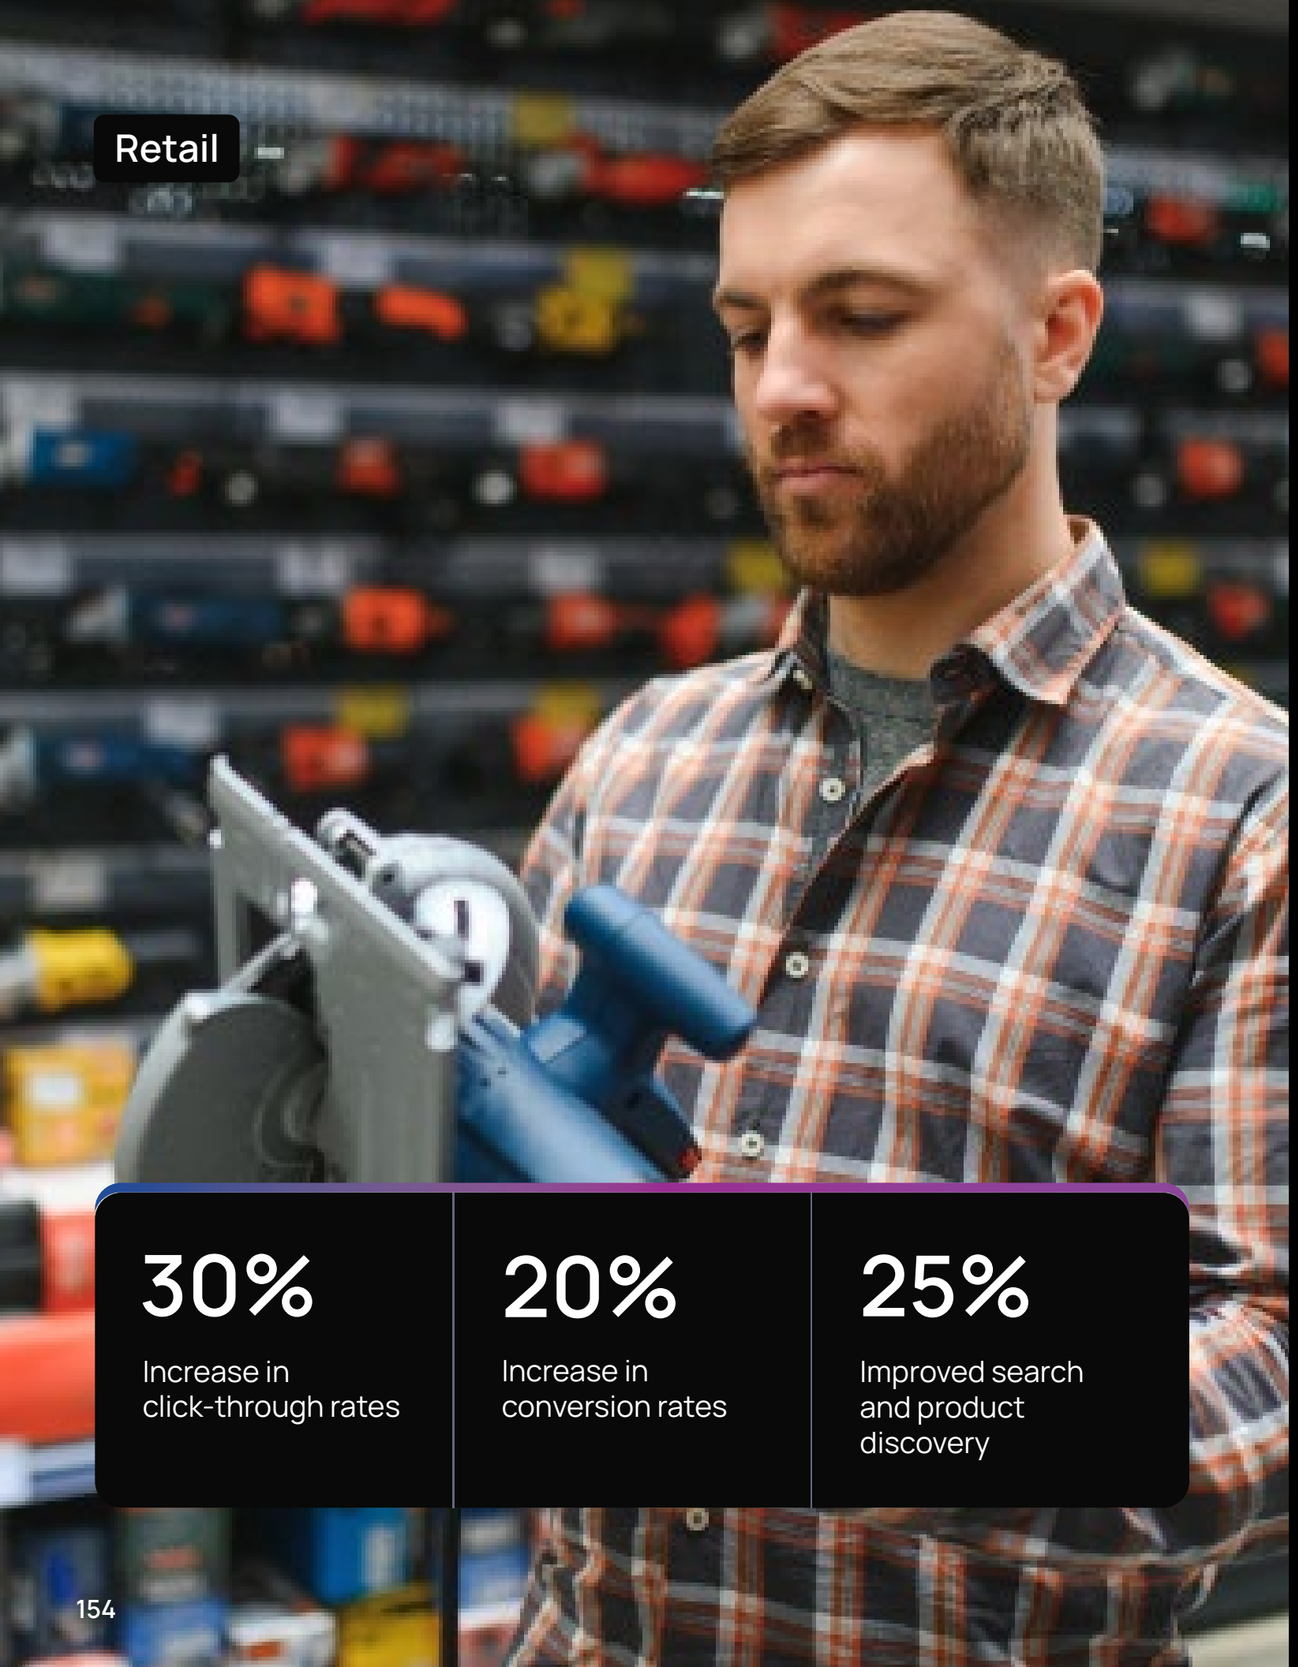

Client

A US-Based Hardware Wholesaler and Retailer

Making product discovery and conversion more personal at scale

Product data quality, audience segmentation, and pricing signals were all affecting how easily customers could find and buy the right products. The client wanted GenAI to make merchandising, search, and support more relevant in real time.

Solution

- Used GenAI to enrich product descriptions for different target audiences.
- Derived product-level attributes to improve search, discovery, and affinity baskets.
- Built a GenAI assistant for retailer merchandising, stock, substitution, and discontinued-item queries.
- Created a vendor-service assistant to handle payment and short-payment inquiries.

Benefits

- Increased click-through rates by 30%.
- Improved conversion rates by 20%.
- Improved search and product discovery by 25%.
- Created a more responsive digital shopping experience for both retailers and vendors.

Energy & Utilities

Energy & Utilities

95%
Operational cost savings

20%
Increased user productivity

20%
Improvement in turn around time

Client

A Mobility and Alternative Energy Solutions Company

Making enterprise knowledge easier to access and easier to act on

Users needed a more intuitive and secure way to interact with enterprise knowledge so they could move faster and serve customers more effectively. The goal was a conversational experience that felt simple without losing governance.

Solution

- Built an interactive web experience for registration, login, and chatbot use.
- Enabled personalized knowledge bases with cited answers for enterprise users.
- Improved document management and retrieval through a GenAI-powered conversational layer.
- Designed the solution to keep enterprise interactions secure and easy to scale.

Benefits

- Reduced operational costs by 95%.
- Increased user productivity by 20%.
- Improved turnaround time by 20%.
- Made customer-facing and internal interactions more responsive and intuitive.

Client

A Leading UK Water Utility

Automating regulatory compliance reporting with AI to eliminate manual effort and penalty risk

Environment Agency Compliance Assessment Report (CAR) forms were unstructured, image-based, or handwritten, requiring significant manual effort to interpret and re-enter. Inconsistent data capture created incomplete records, audit gaps, and a persistent risk of missed regulatory submissions and associated financial penalties.

Solution

- Advanced OCR digitizes scanned and handwritten CAR forms with high accuracy.
- NLP models extract, classify, and structure compliance data automatically.
- Automated validation flags missing or inconsistent data before submission.
- Real-time dashboards give business users visibility into compliance status, risks, and trends.

Benefits

- Near-elimination of missed EA reporting deadlines.
- Significantly improved data accuracy, completeness, and audit traceability.
- Teams freed from repetitive tasks to focus on analytical and governance work.
- Scalable foundation for AI-led automation across additional compliance processes.
- Strengthened regulator relationships through consistent, high-quality reporting.

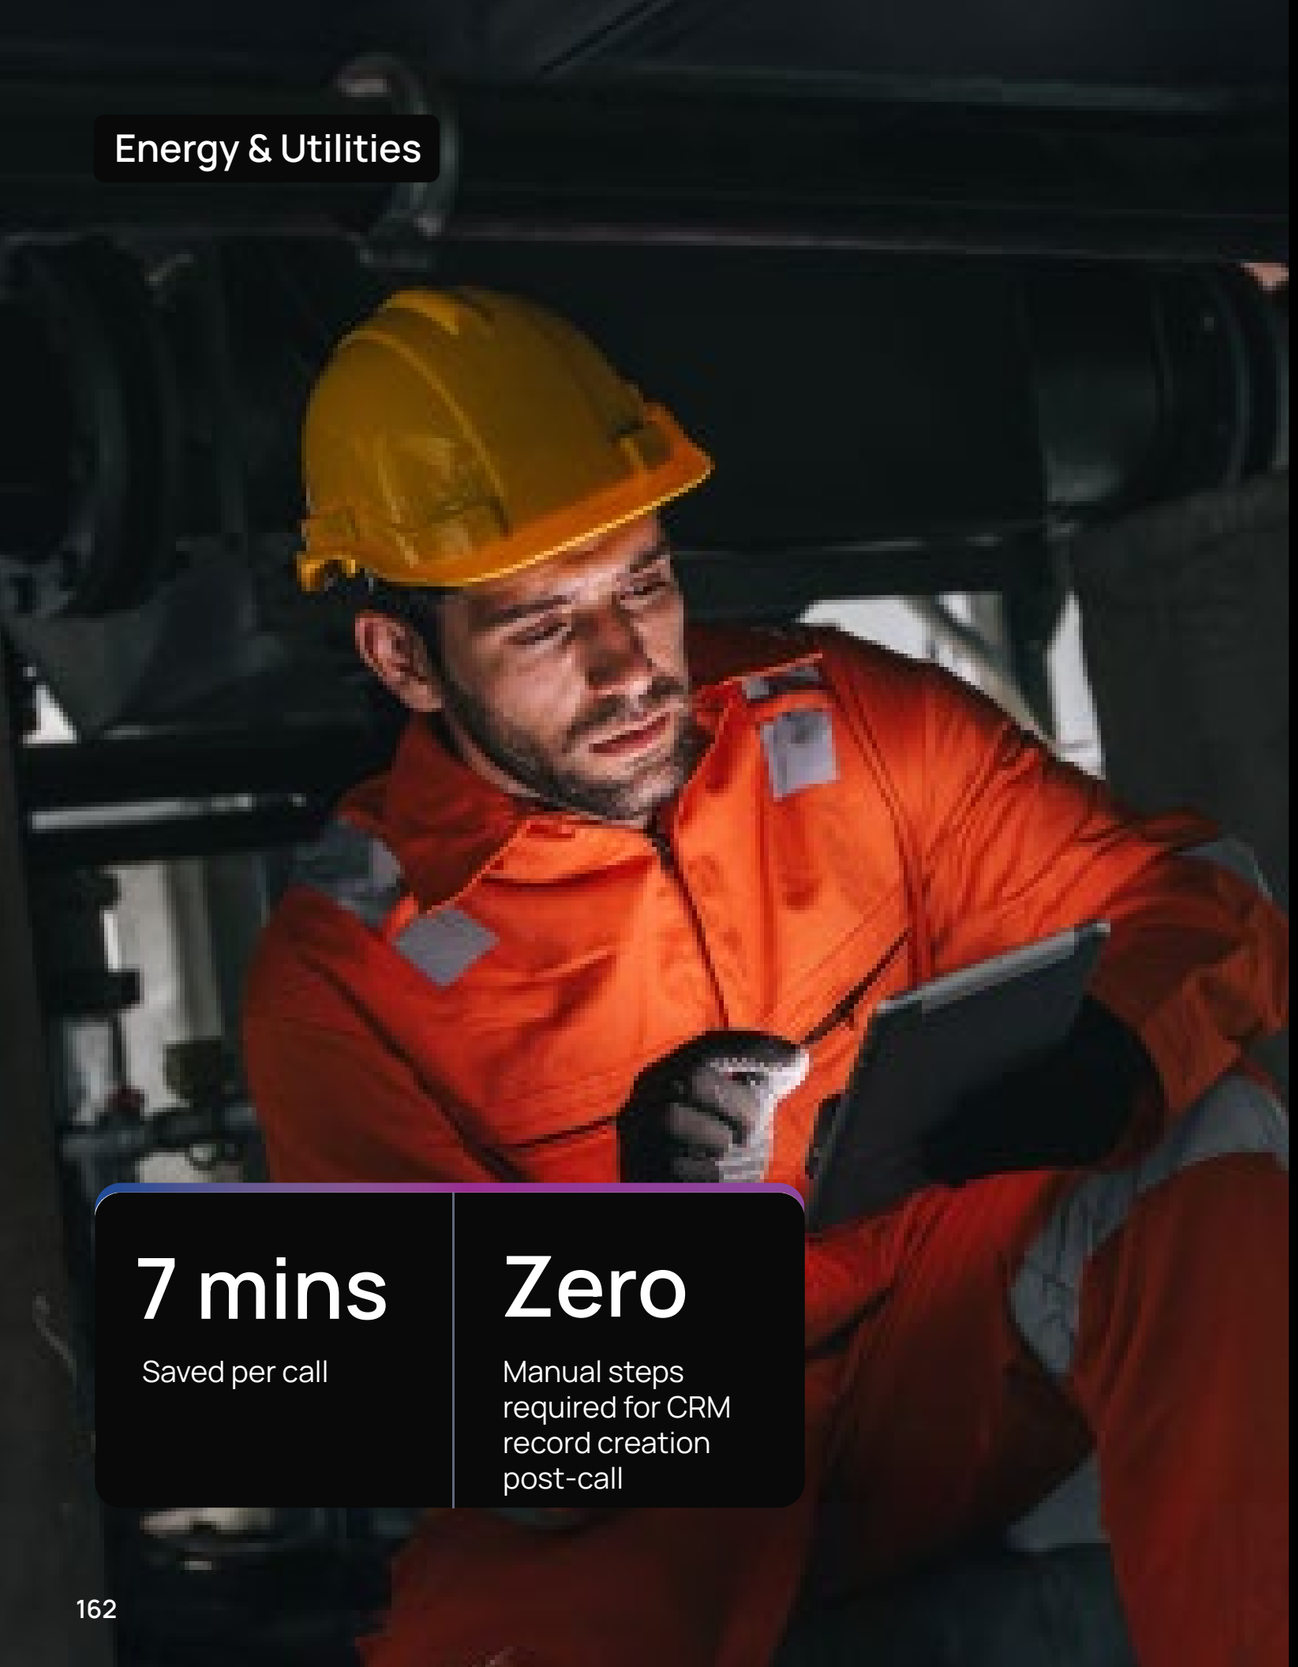

Client

A UK Water Utility Undergoing Large-Scale Digital Transformation

Transforming call center operations with an AI agent to cut handling time and eliminate manual data entry

Sub-optimal phone-based customer service was driving long call-handling times, inaccurate regulatory transcript data, and manually created cases with incorrect next steps. With Ofwat customer satisfaction scores directly influencing the client's ability to raise water charges, improving service quality was both an operational and a financial imperative.

Solution

- Deployed a custom AI Agent built in Copilot Studio and Power Platform, natively integrated with Dynamics 365 CRM Customer Service.
- Automated post-call creation of cases, phone calls, and call transcripts in CRM, eliminating all manual data entry after each call.
- Configured business-user-managed script questions and prompts, removing IT dependency for future script updates.
- Built on the client's existing Microsoft licensing, avoiding additional infrastructure or platform costs.

Benefits

- Reduced average call-handling time, freeing agents to service more customers per shift.
- Eliminated regulatory transcript risk through accurate, automated data capture on every call.
- Improved case accuracy, ensuring the right solutions and next steps are applied consistently.
- Enabled continuous process improvement through a structured audit trail of every interaction.

Transportation & Logistics

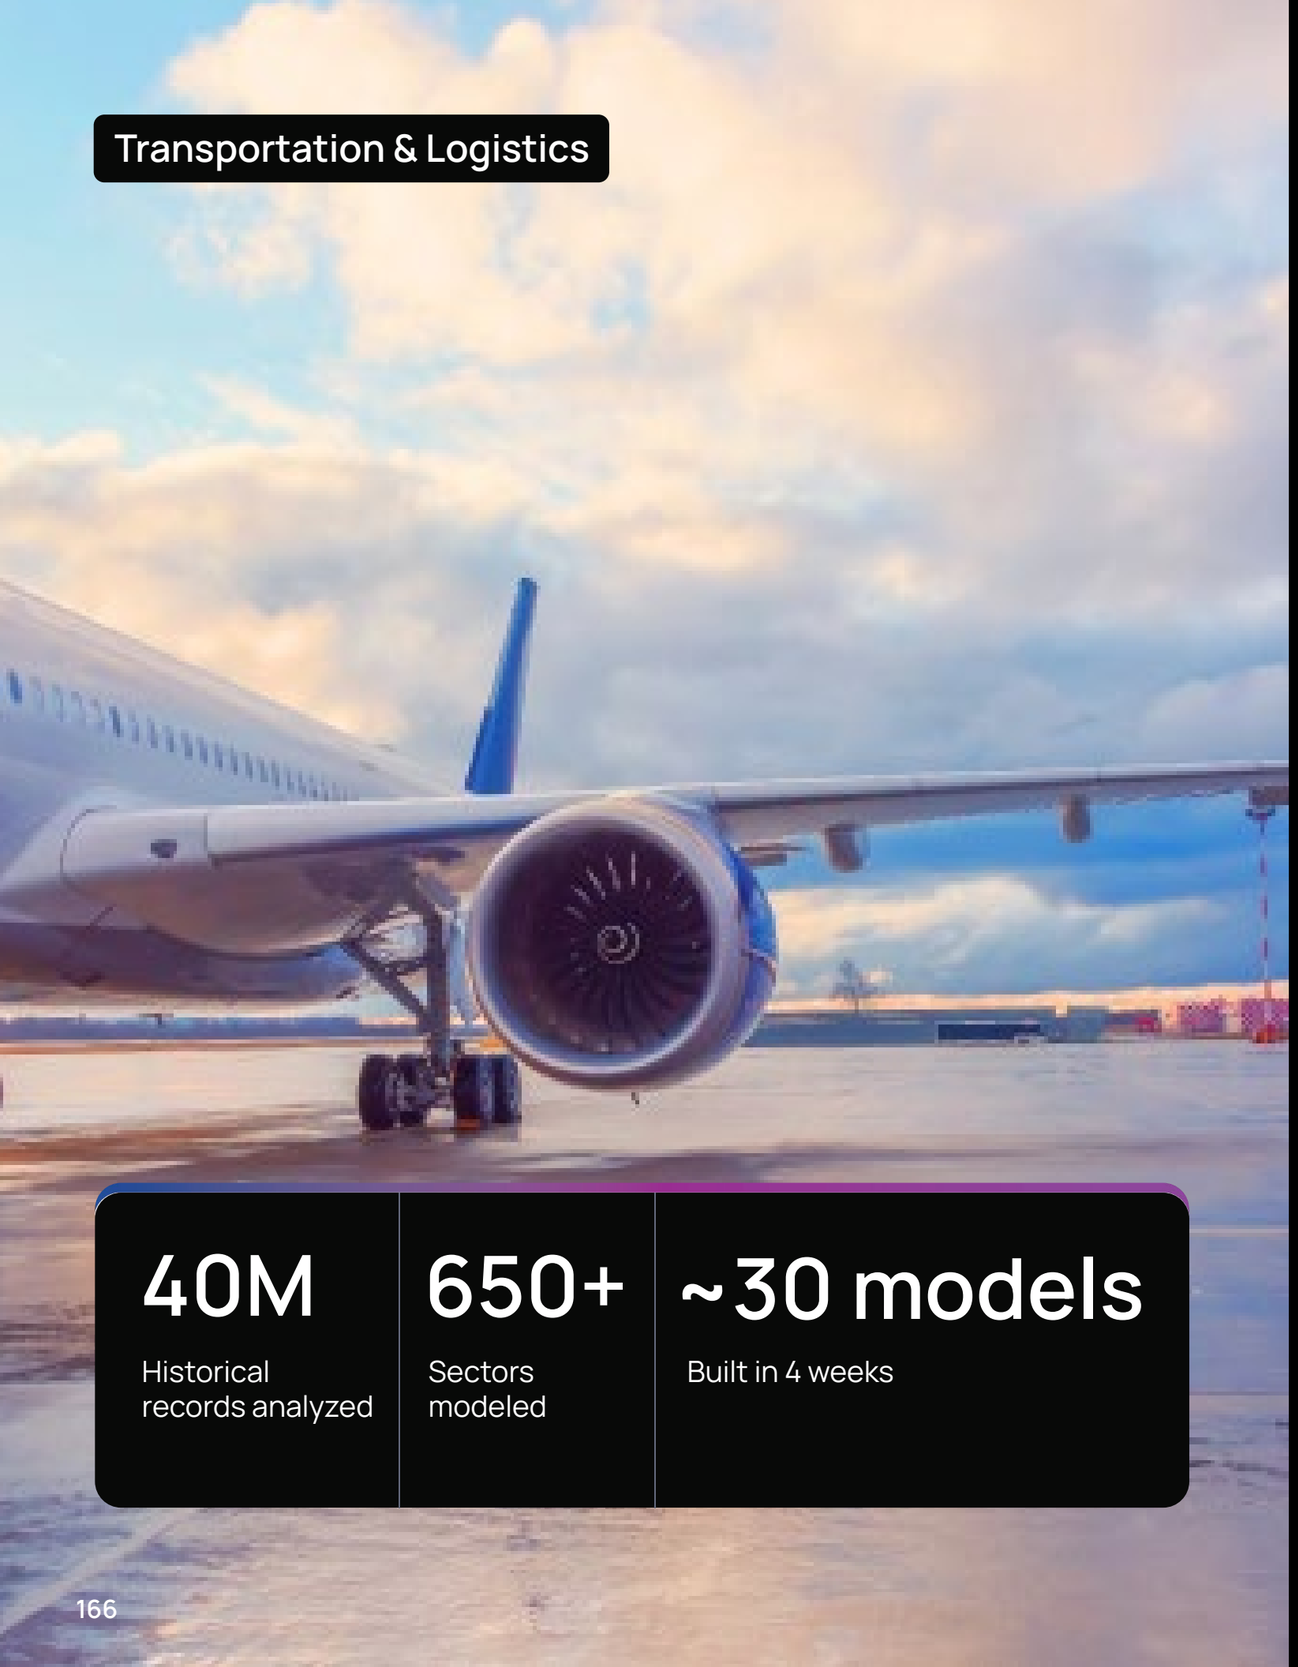

Transportation & Logistics

40M
Historical records analyzed

650+
Sectors modeled

~30 models
Built in 4 weeks

Client

An International Airline

Scheduling flights with a clearer view of near-term demand

The airline needed more accurate short-term passenger forecasts to plan routes and aircraft capacity with greater confidence. Better forecasting meant better decisions in the timeframes that matter most to network planning.

Solution

- Analyzed more than 40 million historical records across 650+ sectors and 20+ forecasting features.
- Built 20–30 segmented models within four weeks to forecast passenger counts for the next 30 days.
- Designed the solution to support both existing and newly introduced sectors.
- Used iterative modeling to improve precision across a complex route network.

Benefits

- Helped the network planning team schedule more effectively in high-performing sectors and airports.
- Improved short-term visibility into dynamic demand patterns not captured in older weekly forecasting methods.
- Supported better aircraft capacity decisions with fresher forecasting signals.
- Created a repeatable forecasting approach that can scale as route conditions change.

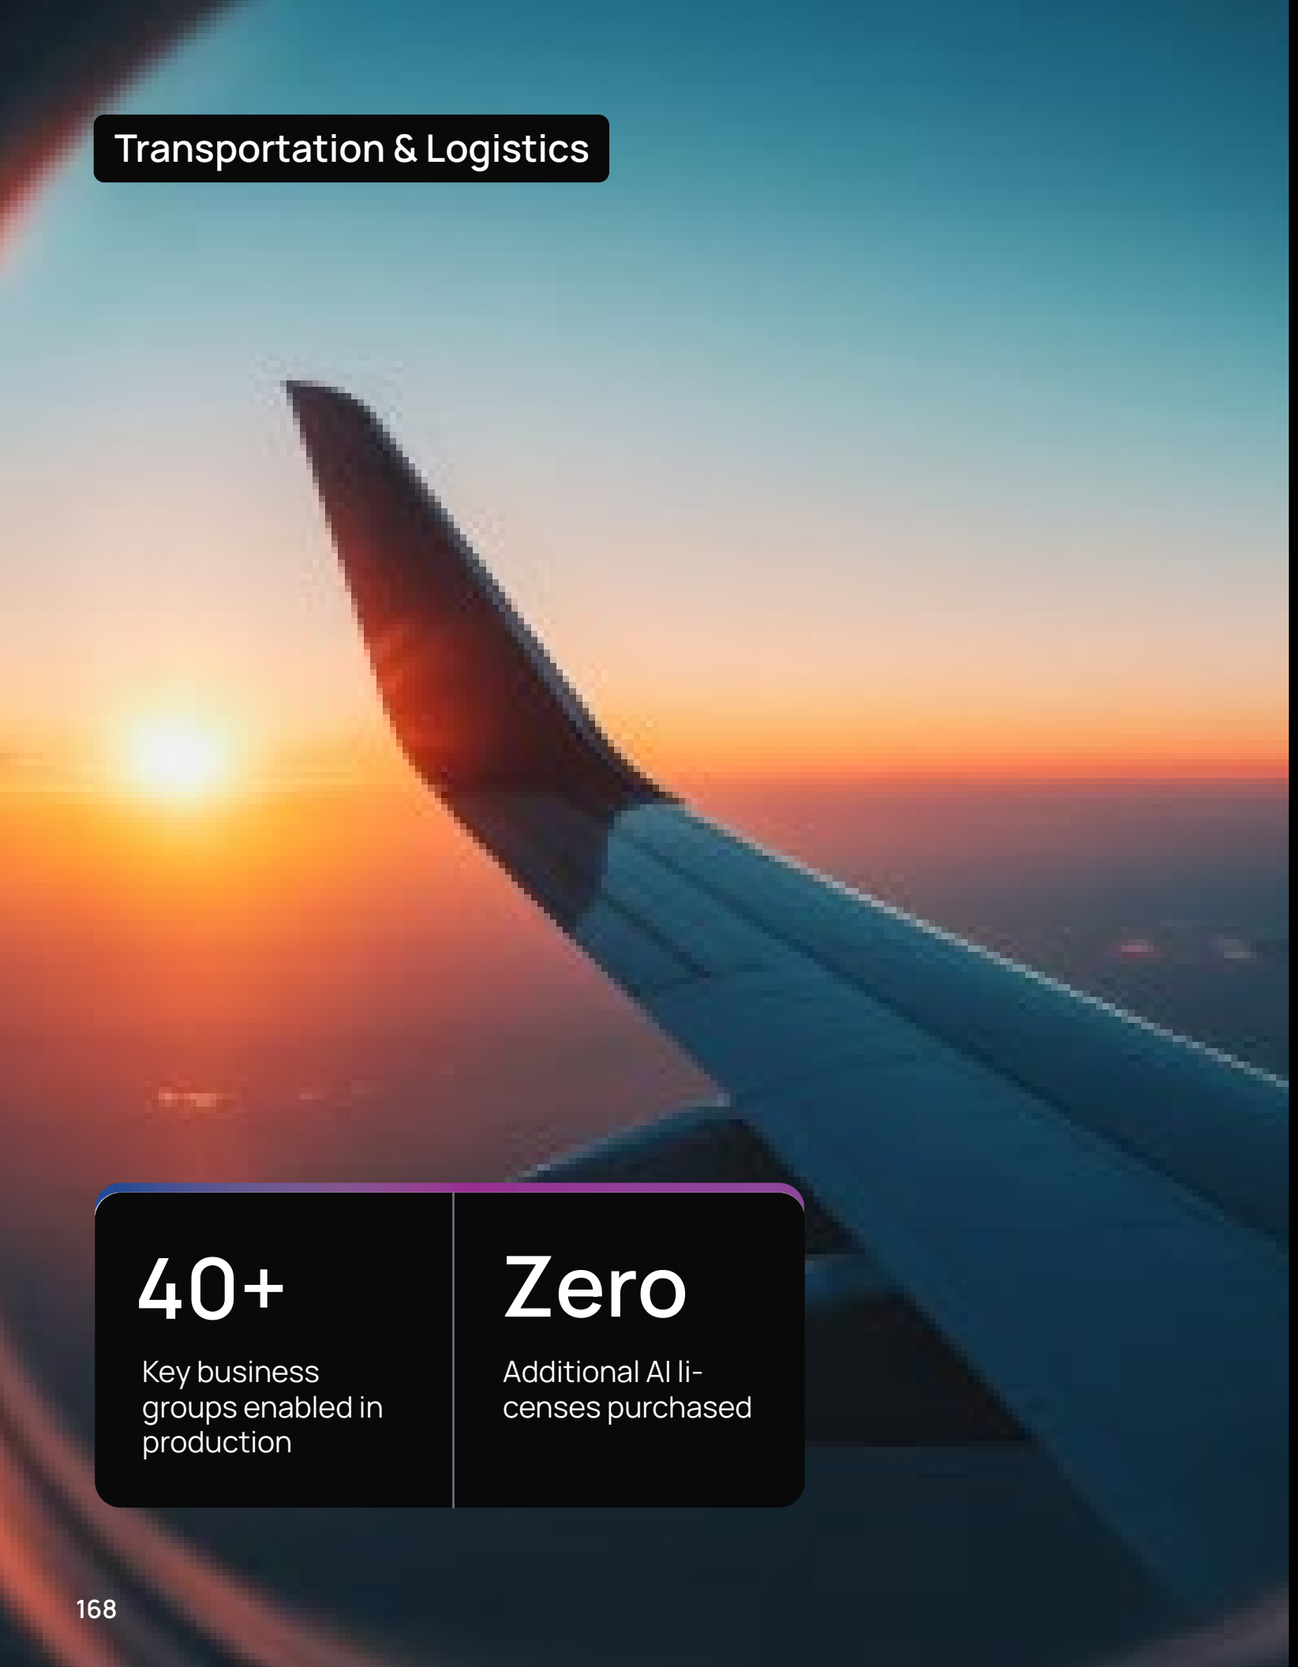

Client

A Leading Global Air Transport Communications & IT Specialist

Accelerating incident resolution through AI-Powered summarization on ServiceNow

Agents working on ServiceNow incidents had to manually review all related communications and transaction histories before they could act, with no quick-summary functionality available. This added significant time to every resolution, created inconsistency in how incidents were handled, and led to delays from a lack of automation. The client wanted an AI-based solution that worked within their existing ServiceNow Now Assist licensing, without purchasing additional AI agent licenses.

Solution

- Implemented ServiceNow Now Assist's Incident Summarization feature, giving agents a one-click AI-generated summary for faster resolution.
- Enabled the feature for a defined set of groups and users to keep AI agent consumption within the existing licensed allocation, avoiding additional licensing costs.
- Deployed a usage monitoring tool to track consumption of the AI feature across the organization.

Benefits

- Agents can now get an instant AI-generated incident summary, reducing the time and effort required to begin resolution.
- Faster, more consistent incident resolution with reduced manual dependencies across key business groups.
- Success of the solution has increased the client's confidence in adopting more AI-based capabilities in the future.

Travel & Hospitality

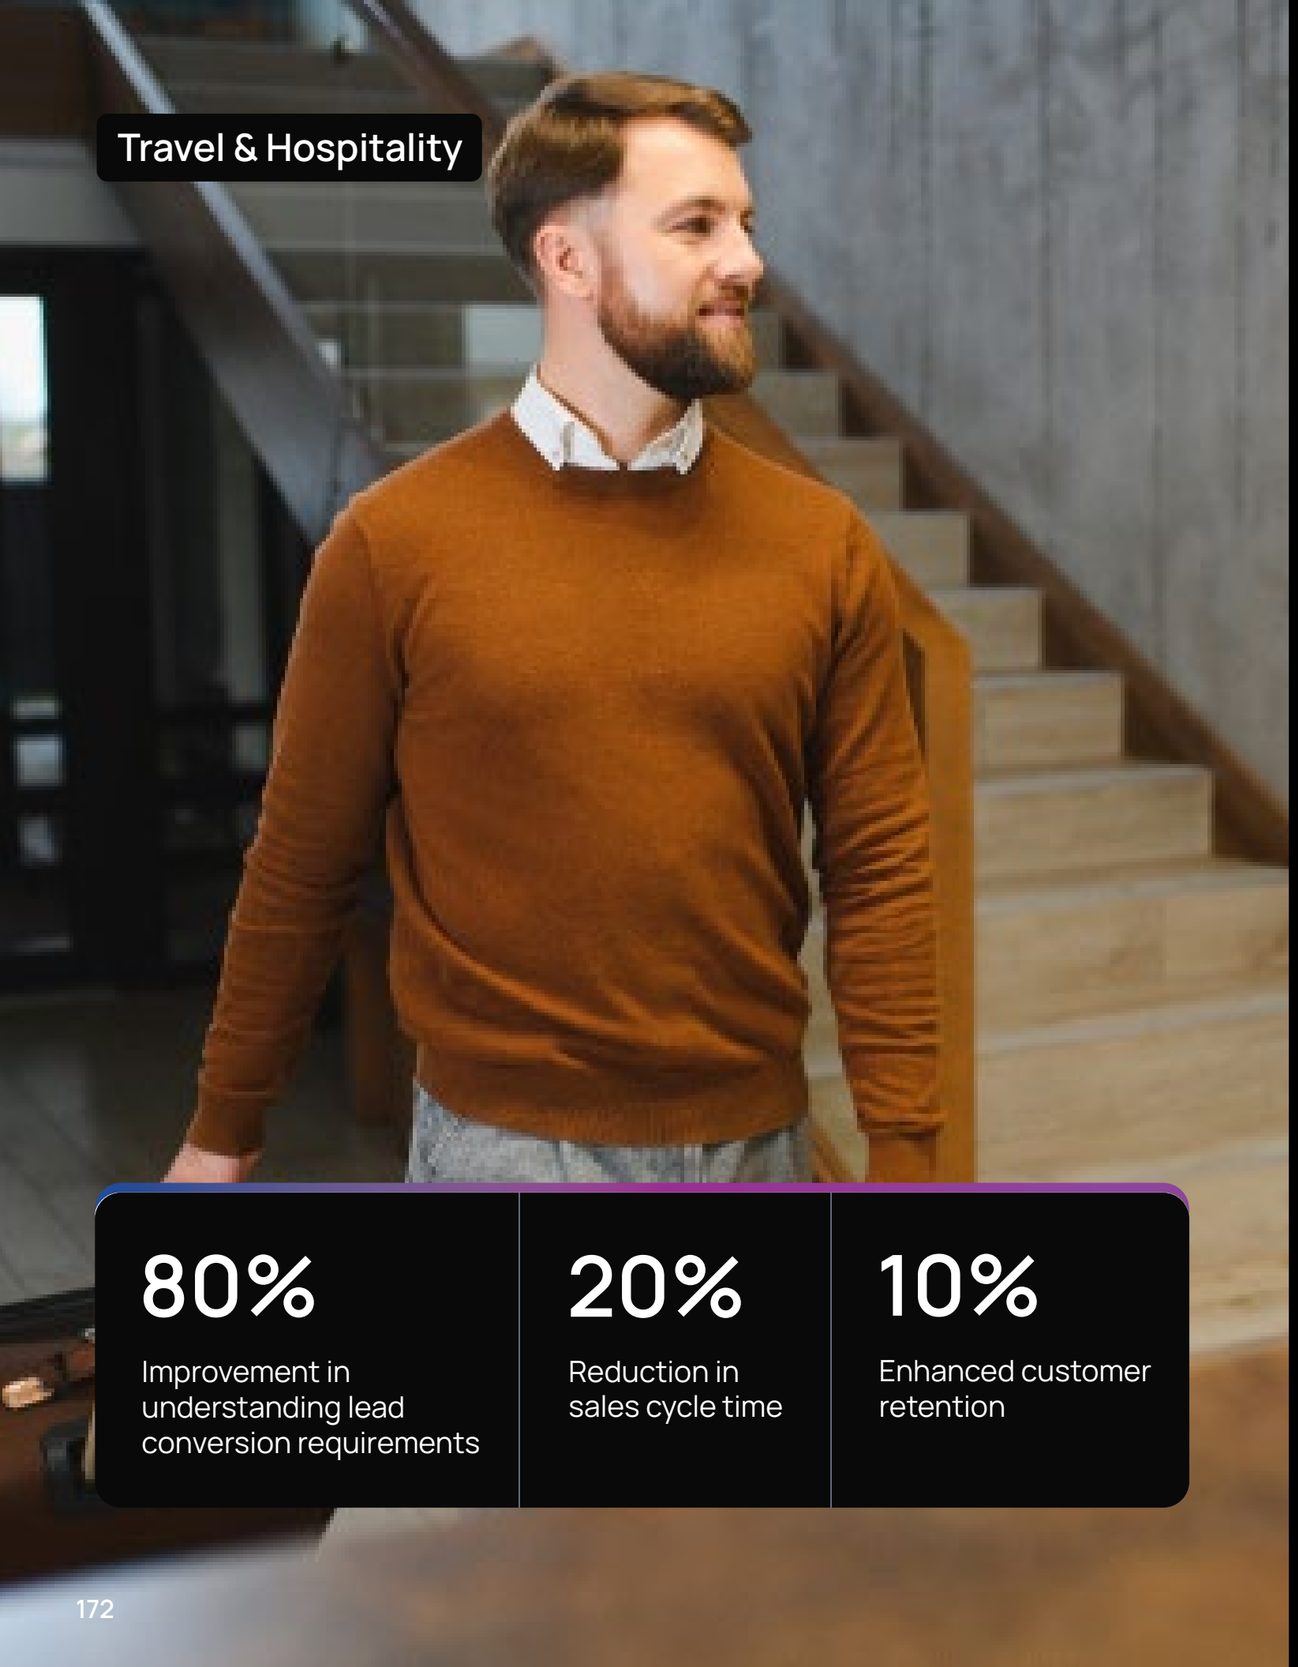

Client

A Renowned Hotel and Hospitality Chain

Capturing every opportunity before it slips through the cracks

Lead capture was fragmented across channels and formats, making it too easy for opportunities to be delayed, missed, or routed slowly. The client needed faster intake and cleaner CRM data so sales teams could respond with more confidence.

Solution

- Implemented AI-driven opportunity capture integrated directly with the client's CRM.
- Automated intake from manual channels, such as email, to reduce human dependency.
- Extracted event agenda data from Word, PDF, and XLSX RFP formats.
- Converted and presented captured information directly in CRM for rapid action.

Benefits

- Improved understanding of lead conversion requirements by 80%.
- Reduced sales cycle time by 20%.
- Improved customer retention by 10%.
- Ensured opportunities were captured accurately and routed to the right team faster.

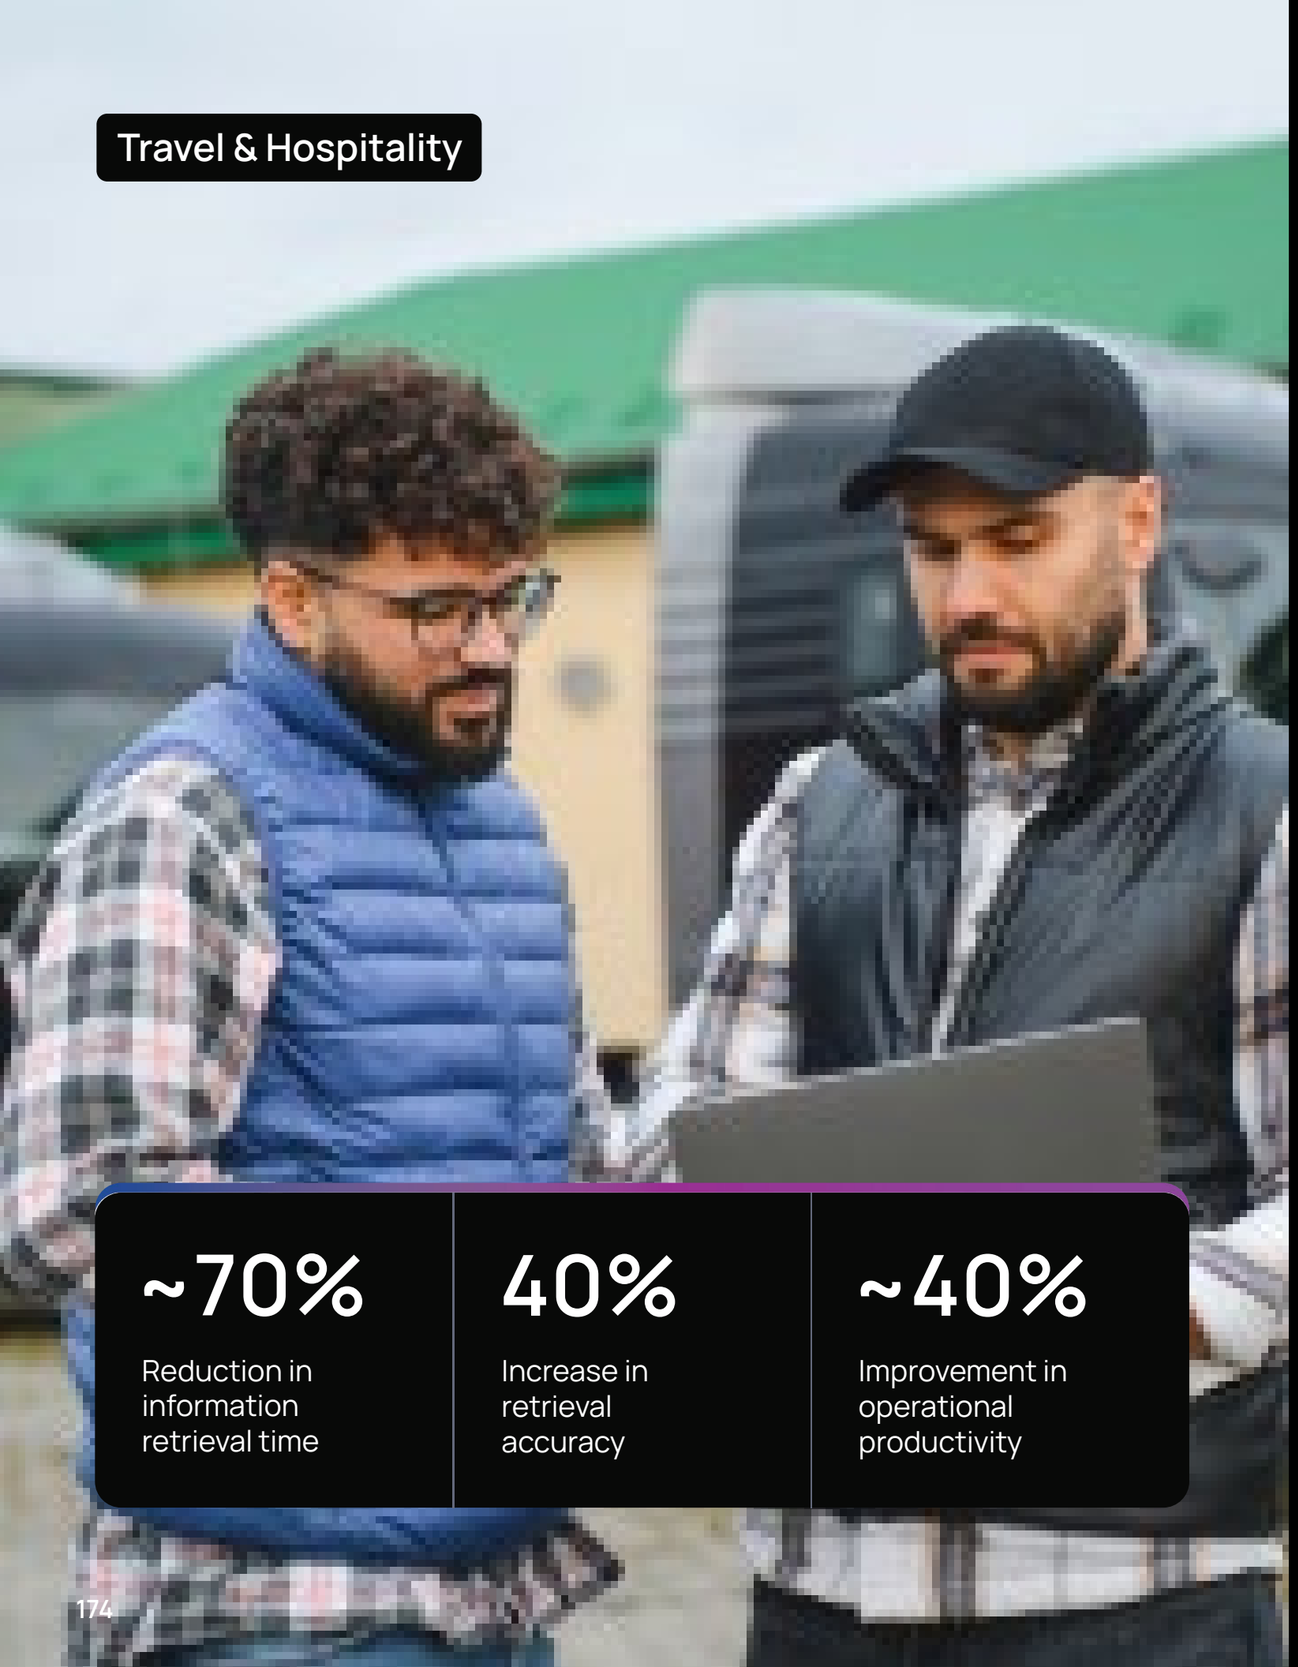

Client

A Leading US Logistics Provider Specializing in Truck Leasing, Fleet Solutions, and End-to-End SCM

Turning enterprise knowledge into a faster decision advantage

Important knowledge was trapped across systems, documents, and manual processes, making it too slow to find and too inconsistent to trust. The client needed a more intelligent way to retrieve information and act on it at operational speed.

Solution

- Implemented an AI-powered platform unifying structured and unstructured enterprise data.
- Added intelligent document processing to automate classification and extraction.
- Enabled hybrid semantic and keyword search plus RAG for grounded responses.
- Built governed ingestion, indexing, observability, and agent-orchestrated workflows.

Benefits

- Reduced information retrieval time by about 70%.
- Improved operational productivity by about 40%.
- Increased knowledge retrieval accuracy by 40%.
- Made enterprise knowledge easier to scale, govern, and use across the business.

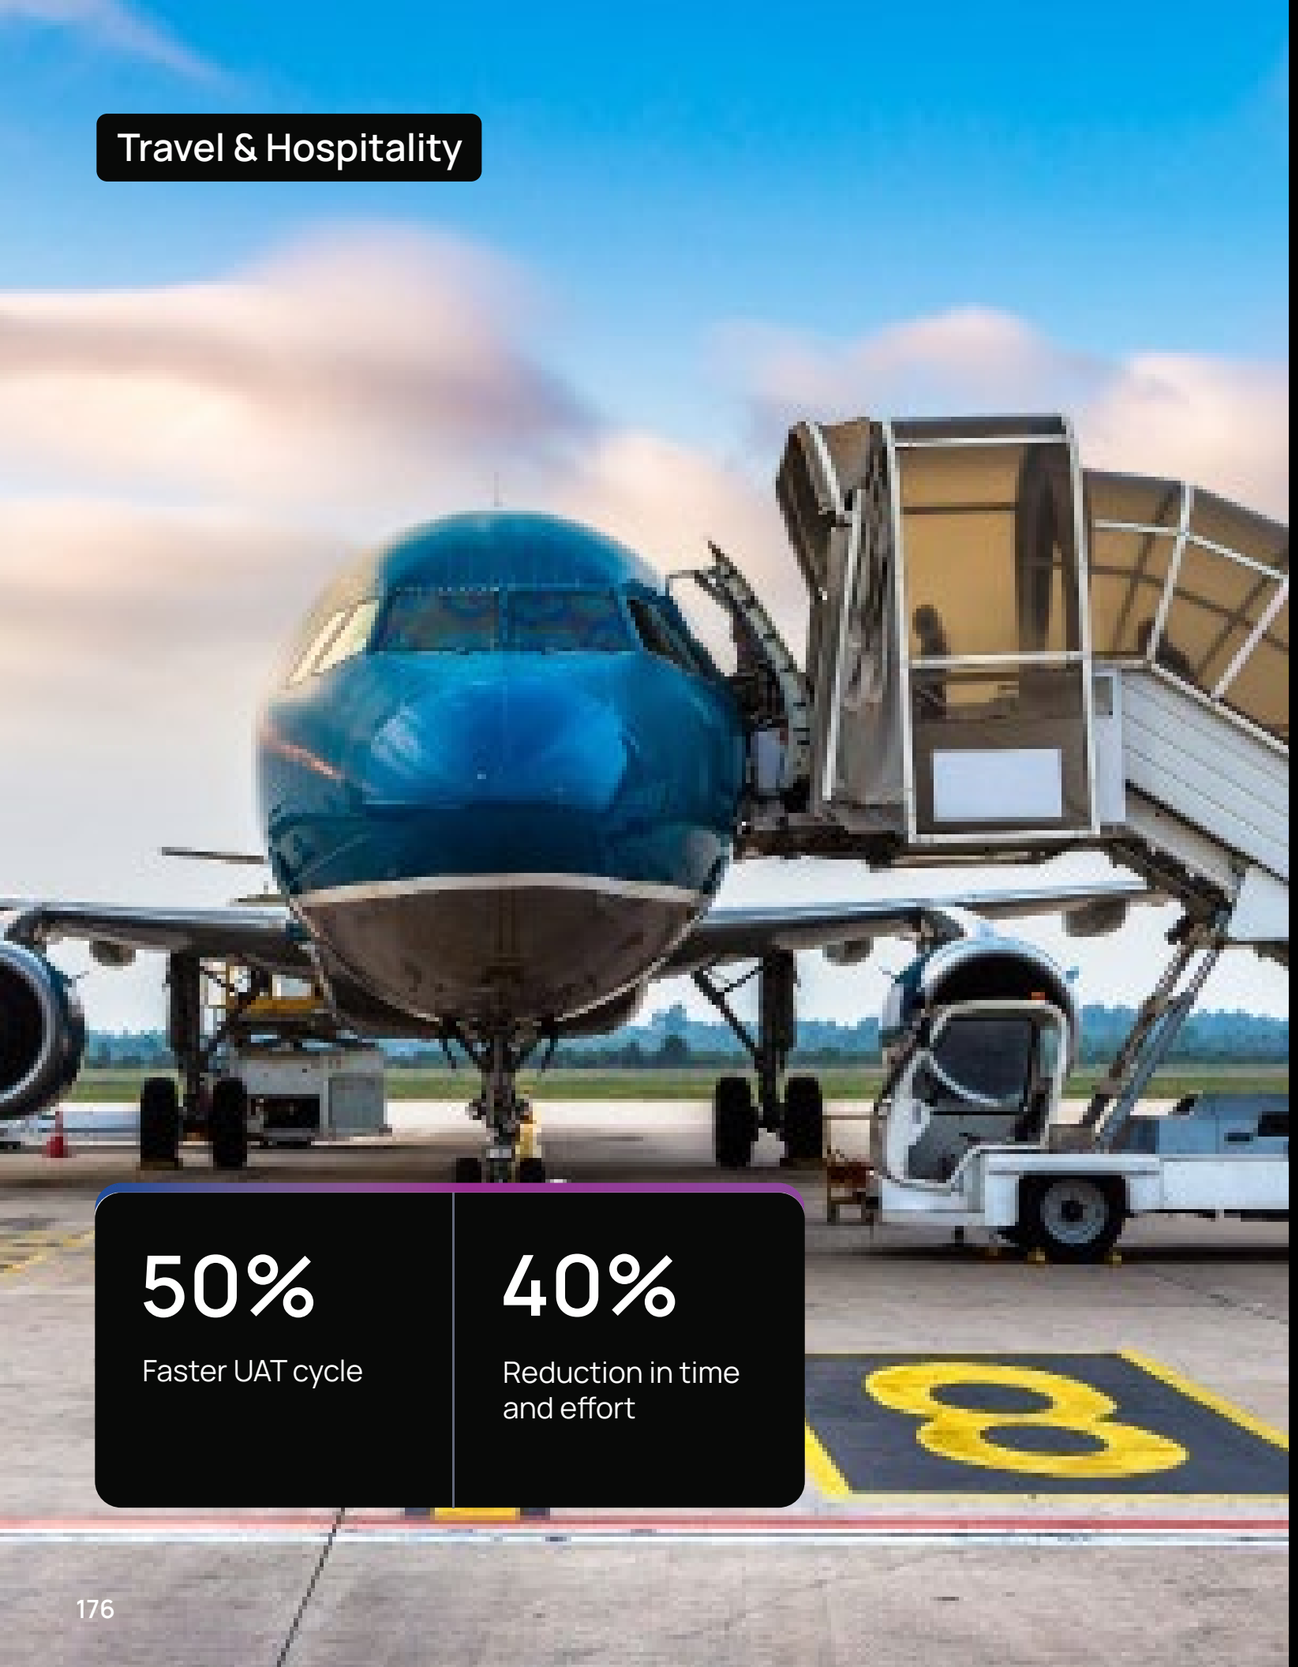

Client

A Major US Airline Recognized For Technology-Driven Innovation

Helping business teams move from ideas to working AI experiences faster

The client wanted to infuse AI into business-led innovation without forcing every idea through a slow, IT-dependent path. The result was a collaboration model that helped teams see possibilities faster and validate them in working product flows.

Solution

- Reimagined crew experiences by treating pilots and flight attendants as customers.
- Built an end-to-end companion app journey for flight attendants.
- Used walkthroughs and live prompting sessions to demonstrate rapid AI-driven feature creation.
- Aligned prompts and workflows to business goals, partner ecosystems, and real-world operations.

Benefits

- Reduced user acceptance testing (UAT) cycle time by 50%.
- Cut time and effort by 40%.
- Delivered key features two quarters early.
- Expanded stakeholder appetite for broader AI-led transformations across pilot payroll and crew scheduling.

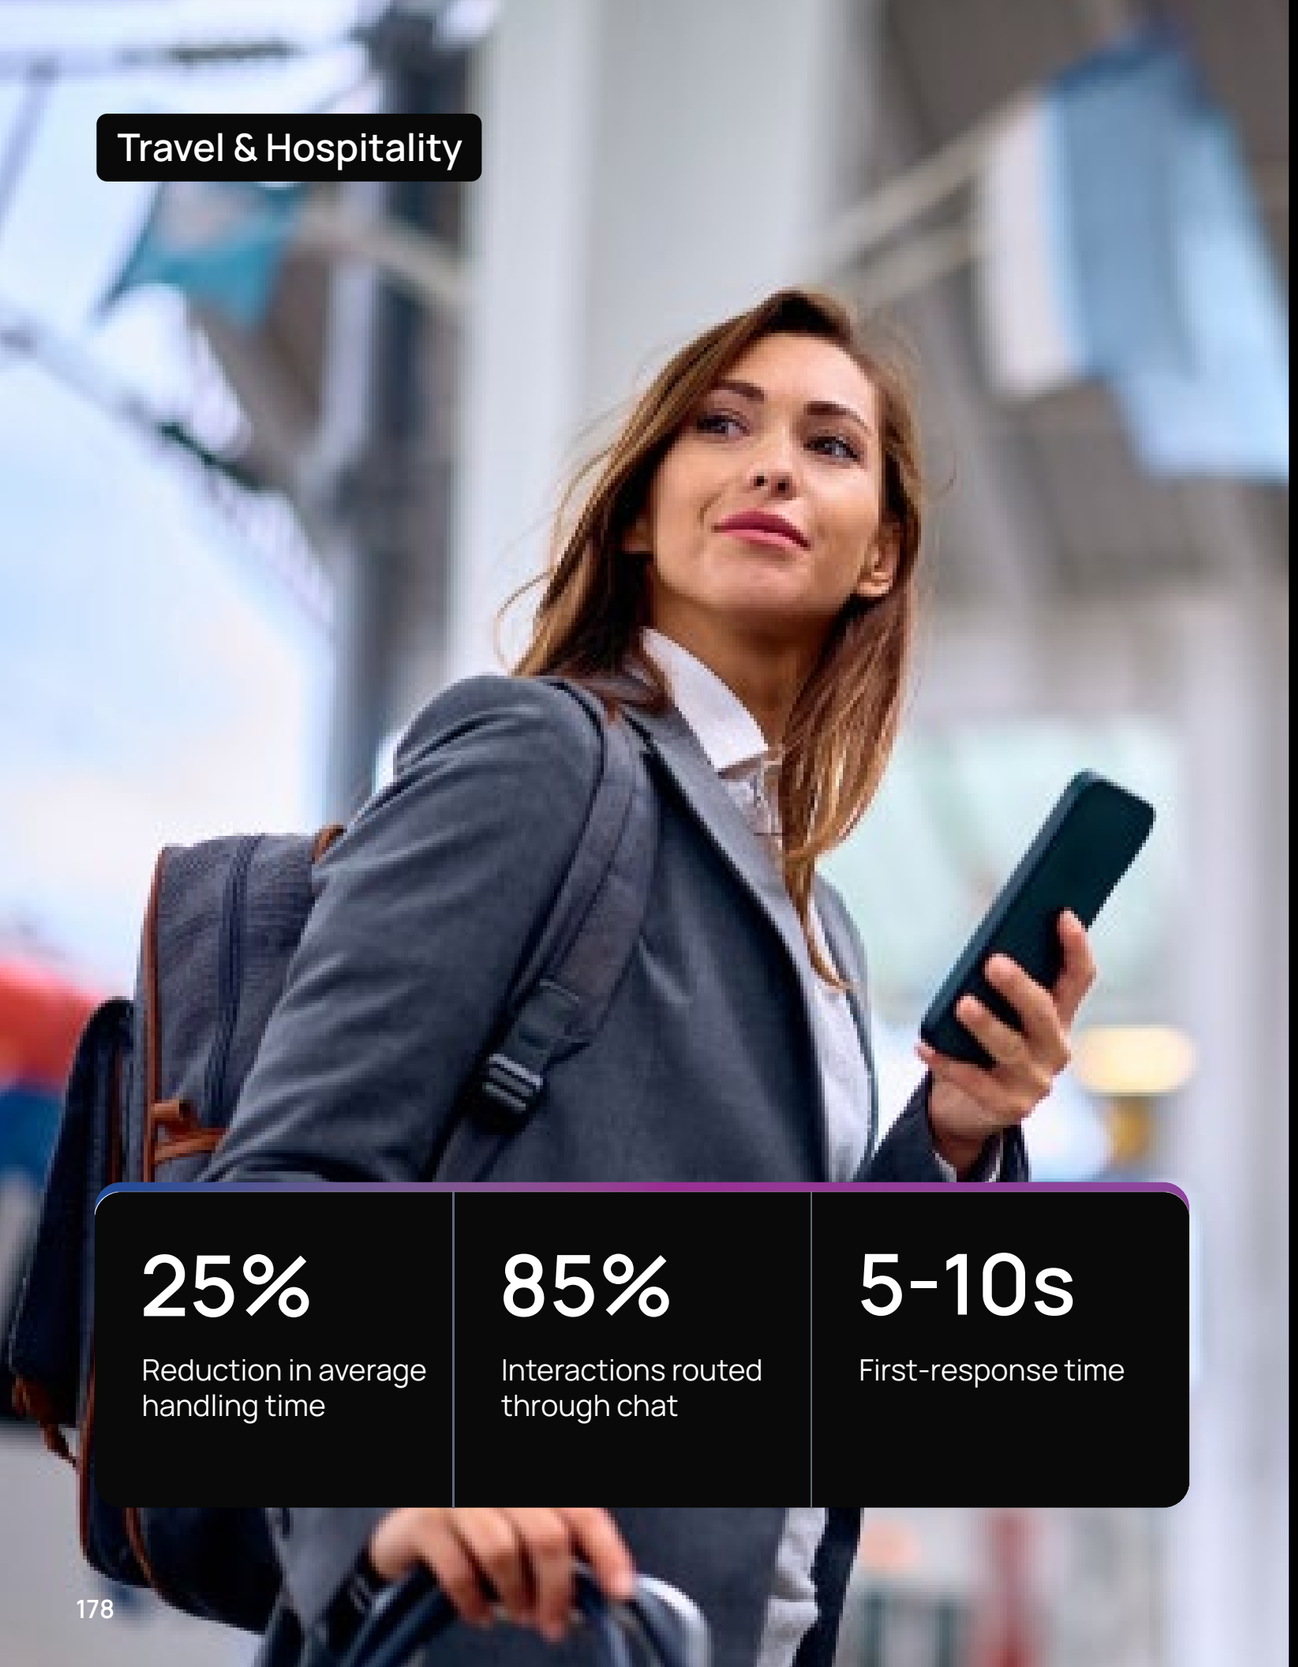

Client

A High-Performance Concierge Services Provider

Delivering faster customer support across every channel

Manual, channel-dependent support was driving longer handling times and inconsistent experiences. The client wanted a more responsive service model that could shift routine interactions to self-service without lowering quality.

Solution

- Implemented a multilingual self-service chatbot powered by NLP and LivePerson GenAI capabilities.
- Automated validation, FAQ handling, and out-of-hours support.
- Enabled intelligent routing and prioritization for seamless escalation when needed.
- Extended 24/7 support across messaging channels customers already use.

Benefits

- Reduced average handling time by 25%.
- Routed 85% of interactions through chat, increasing self-service adoption.
- Achieved first-response times of just 5–10 seconds.
- Made support easier to access through whichever channel the end customer preferred.

Client

A Leading Global Airline Serving 100+ Destinations Across Europe, Asia, and the US

Replacing command memorization with AI-guided crew operations support

Legacy flight operations system (FOS) commands were hard to learn, easy to misuse, and slow to retrieve during live operations. The client needed faster, more reliable access to command guidance to improve training, quality assurance (QA), and operational response.

Solution

- Built an agentic AI solution with advanced retrieval augmented generation (RAG) for accurate FOS command retrieval.
- Created a secure model context protocol (MCP) server with OAuth 2.0 authentication for controlled tool access.
- Indexed legacy manuals through a scalable ingestion pipeline.
- Used GPT-4.1 reasoning and session-aware interactions to deliver traceable, cited guidance.

Benefits

- Reduced integrated operations center (IOC) training effort by 10x.
- Improved QA deliverables by 5x.
- Reduced command-entry errors and inconsistencies through natural-language guidance.
- Improved response speed during irregular operations with instant lookup and clearer guidance.

Technology, Products, & Platforms

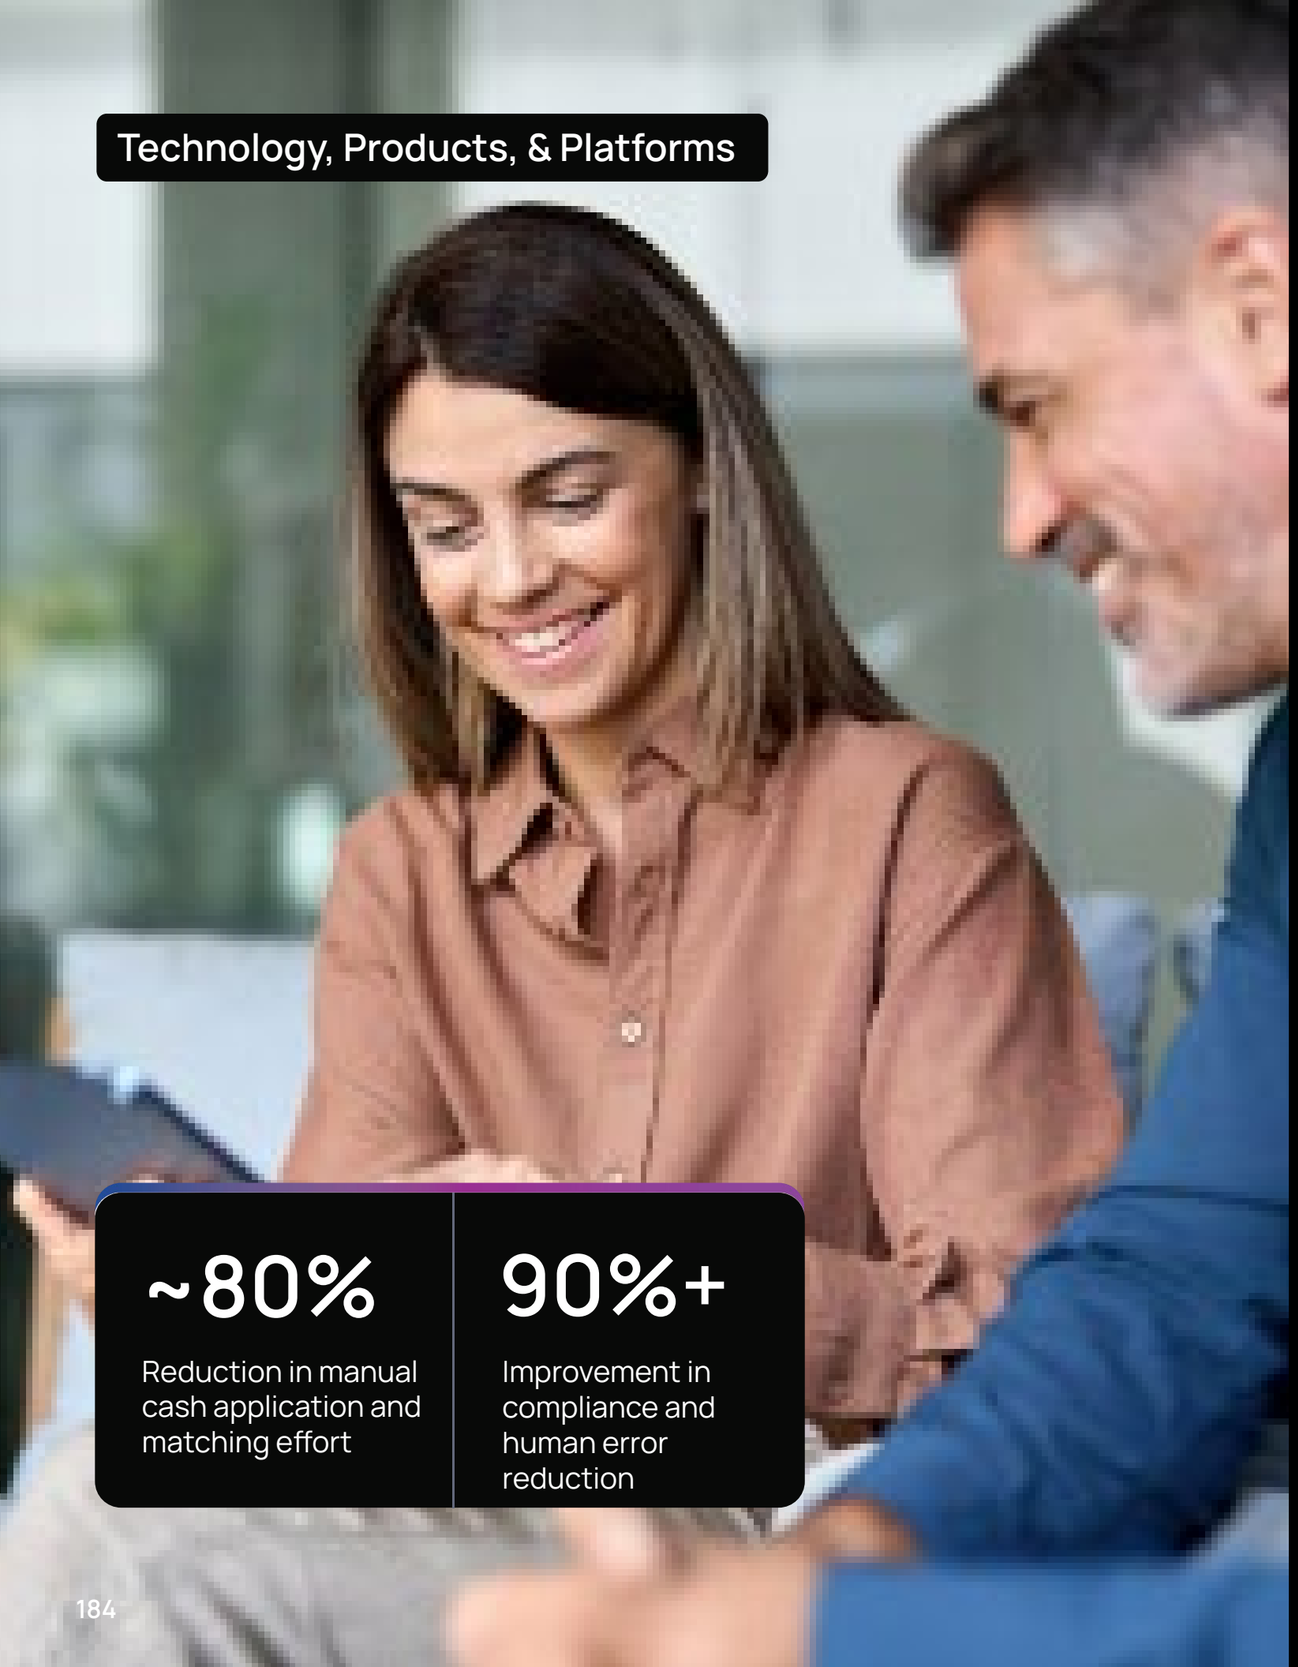

~80%

Reduction in manual cash application and matching effort

90%+

Improvement in compliance and human error reduction

Client

A Leading Smart Building Technology Solutions Provider

Transforming accounts receivable with an end-to-end agentic AI cash collection platform

The client's accounts receivable operation was largely manual and reactive, with bank statements matched by hand, no real-time prioritization, and analysts spending most of their time on data cleansing rather than collections.

Solution

- Deployed a coordinated Agentic AI platform covering the full cash collection cycle, from bank statement ingestion to dunning, across the client's existing ERP, bank feeds, and email systems.
- Automated cash application and AI-driven 3-way matching across bank, invoice, and remittance, with intelligent remittance email responses generated after successful matching.
- AI-powered risk scoring and worklist prioritization give analysts a single, ranked worklist based on aging, P2P status, and payment behavior each day.
- Automated P2P tracking and consistent customer follow-ups via email, replacing manual, analyst-driven dunning entirely.

Benefits

- Analysts freed from data cleansing to focus on collections.
- Faster cash conversion through proactive, risk-prioritized dunning.
- Deteriorating accounts surfaced earlier via continuous AI scoring.
- Full audit trail across matching, P2P, and customer communications.
- Platform scales with transaction volume without adding headcount.

90%

Reduction in test execution time

Client

An American Multinational Technology Company

Making localization testing fast enough for global release speed

Manual localization testing across 70 languages was slowing release cycles and introducing too much inconsistency into validation and documentation. The client needed a more reliable way to scale quality assurance (QA) without adding more manual overhead.

Solution

- Built an AI-powered localization testing platform for accuracy, speed, and scale.
- Enabled one-click QA automation for end-to-end execution.
- Automated screenshot-based issue detection, annotation, and side-by-side documentation.
- Added AI-driven linguistic, visual, and translation-quality validation across platforms.

Benefits

- Reduced test execution time by 90%.
- Improved localization QA consistency across a large global language footprint.
- Accelerated stakeholder review with automated visual documentation and issue descriptions.
- Reduced manual error in a high-volume, high-variation testing workflow.

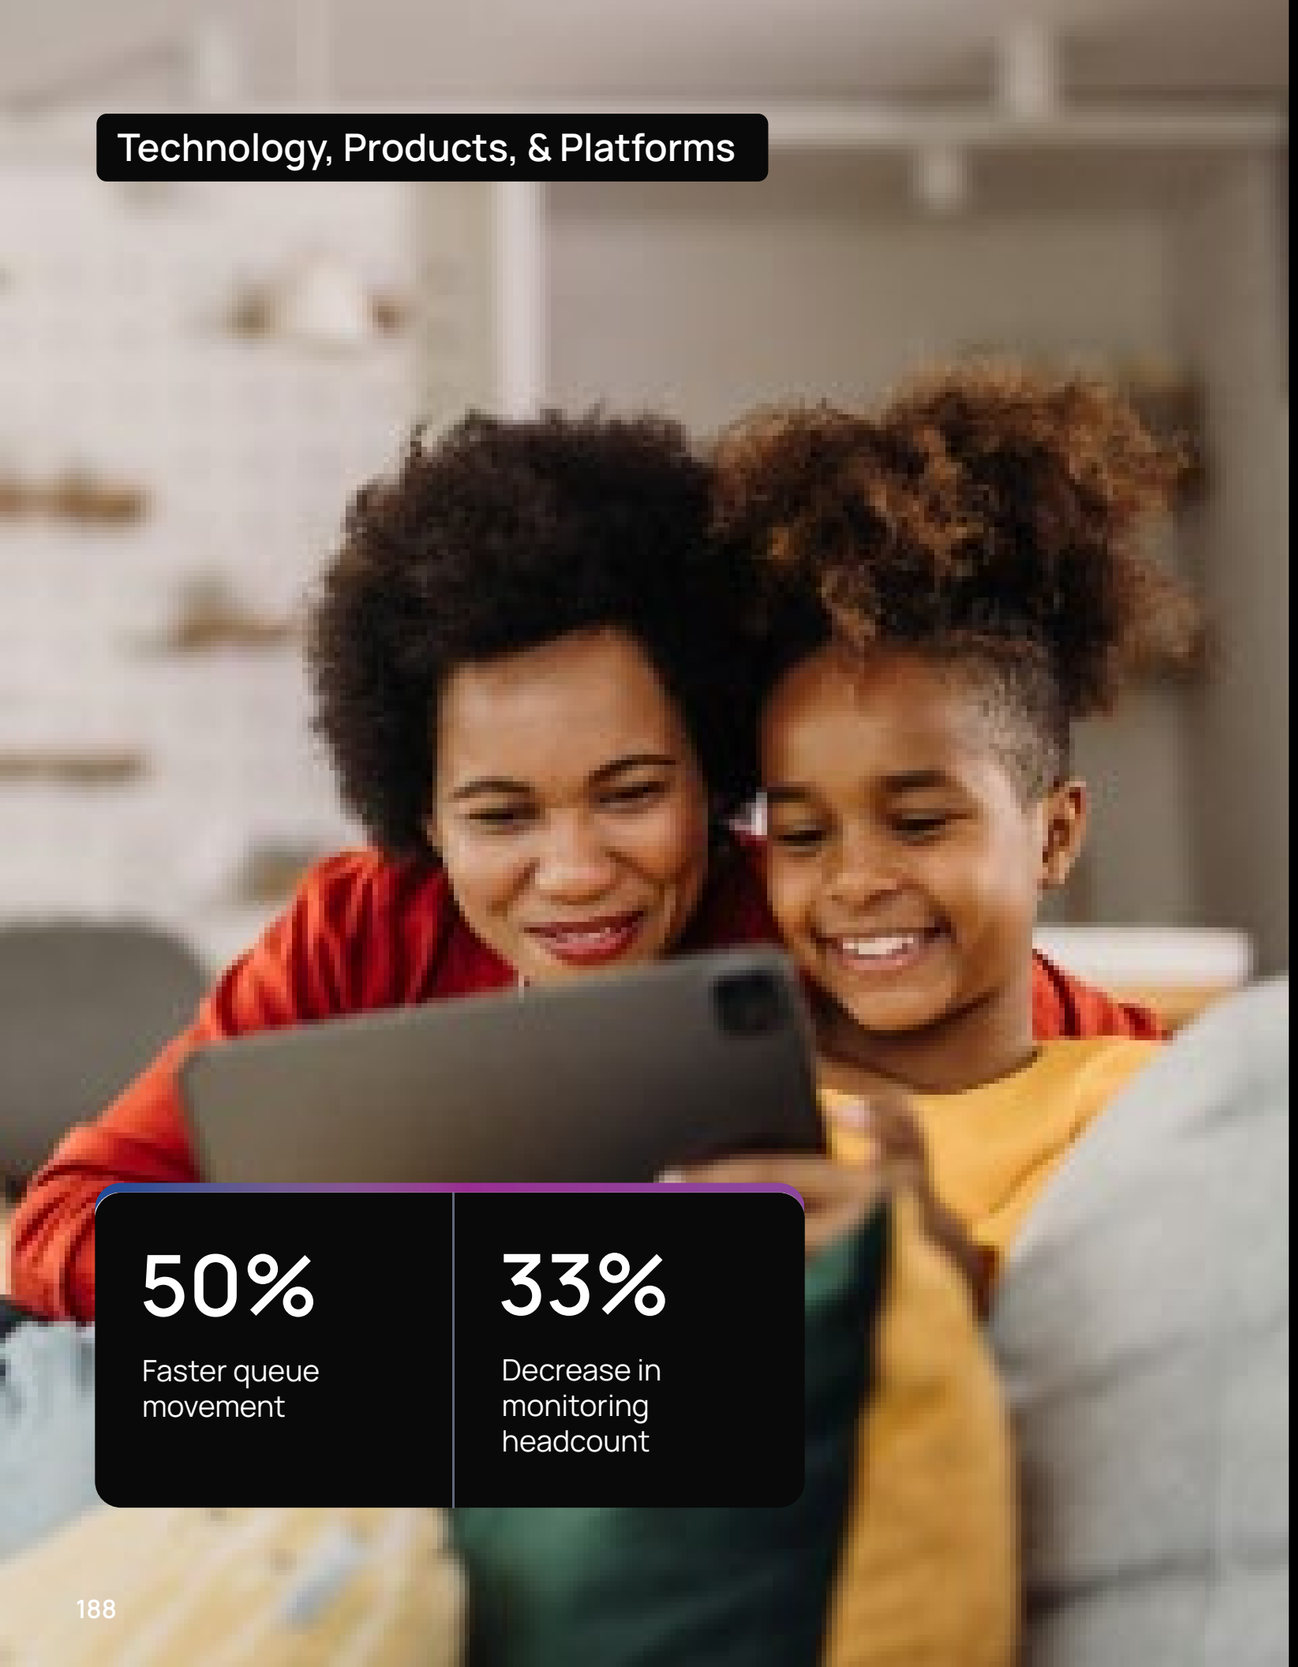

50%
Faster queue movement

33%
Decrease in monitoring headcount

Client

A Global Cloud Computing Leader

Seeing queue pressure in real time before the experience breaks down

Event teams needed privacy-safe visibility into crowding and queues so they could respond faster and guide people more effectively. The client wanted less manual monitoring and a better attendee experience at scale.

Solution

- Automated crowd counting in queues to free event staff for higher-value attendee support.
- Ran models locally so only non-PII data flowed to backend notification systems.
- Triggered alerts for crowded areas such as registration and high-traffic event zones.
- Designed the solution for real-time visibility without sacrificing privacy.

Benefits

- Improved queue movement by 50%.
- Reduced human resources needed for monitoring by 33%.
- Increased participant satisfaction through faster, more informed guidance.
- Made large-event operations more responsive without relying on more staff.

30%
Reduction in service costs

50,000
Installations supported

18
Countries covered

Client

A Global Leader in Materials Engineering Solutions

Moving from reactive installs to smarter, AI prioritized service

As chip-installation volumes increased across countries and sites, the client needed a more intelligent way to reduce recovery time, improve first-time resolution, and spot issues before they escalated. The goal was a more proactive field operation at global scale.

Solution

- Built a phased AI solution to improve first-time-right and mean-time-to-recovery in installations.
- Automated response to known symptoms across 50,000 installations in 18 countries.
- Expanded the model with text analysis to identify emerging issues earlier.
- Combined text and IoT data for real time monitoring and next-best-action recommendations.

Benefits

- Reduced service costs by 30%.
- Improved support coverage across 50,000 installations in 18 countries.
- Enabled more proactive issue prevention through text and IoT-driven insights.
- Helped operations teams prioritize the next best action faster and more consistently.

40%
Lower per-language engineering overhead

25%
Faster regional rollout

18%
Improvement in ASR and NLU accuracy

Client

A Global Wearables Leader

Embedded linguistic engineering for multimodal AI at global scale

The client integrated advanced LLMs to deliver multimodal experiences combining voice, images, and natural language for premium consumers worldwide. Off-the-shelf AI software could not provide the phonological, morphological, and semantic modeling required, and third-party vendor annotations delivered inconsistent data that degraded model accuracy. The client needed a strategic services partner to embed linguistic engineering, quality assurance, and domain expertise directly into their AI operating model.

Solution

- Delivered ownership of linguistic assets, including lexicons and ontologies, across Automatic Speech Recognition (ASR), Text-to-Speech (TTS), and Natural Language Understanding (NLU) workflows for 15+ target languages.
- Implemented phonological, morphological, and semantic models capturing regional variants.
- Used strategic sampling, golden datasets, and linguistic red-teaming to govern third-party annotation quality.
- Established dynamic playbooks and KPI monitoring to drive continuous improvement and linguistic asset evolution.

Benefits

- Lowered per-language engineering overhead by 40% via modular, reusable linguistic assets and playbooks.
- Compressed regional rollout timelines by 25%, enabling confident launches across Asia and MENA within planned schedules.

Professional Services

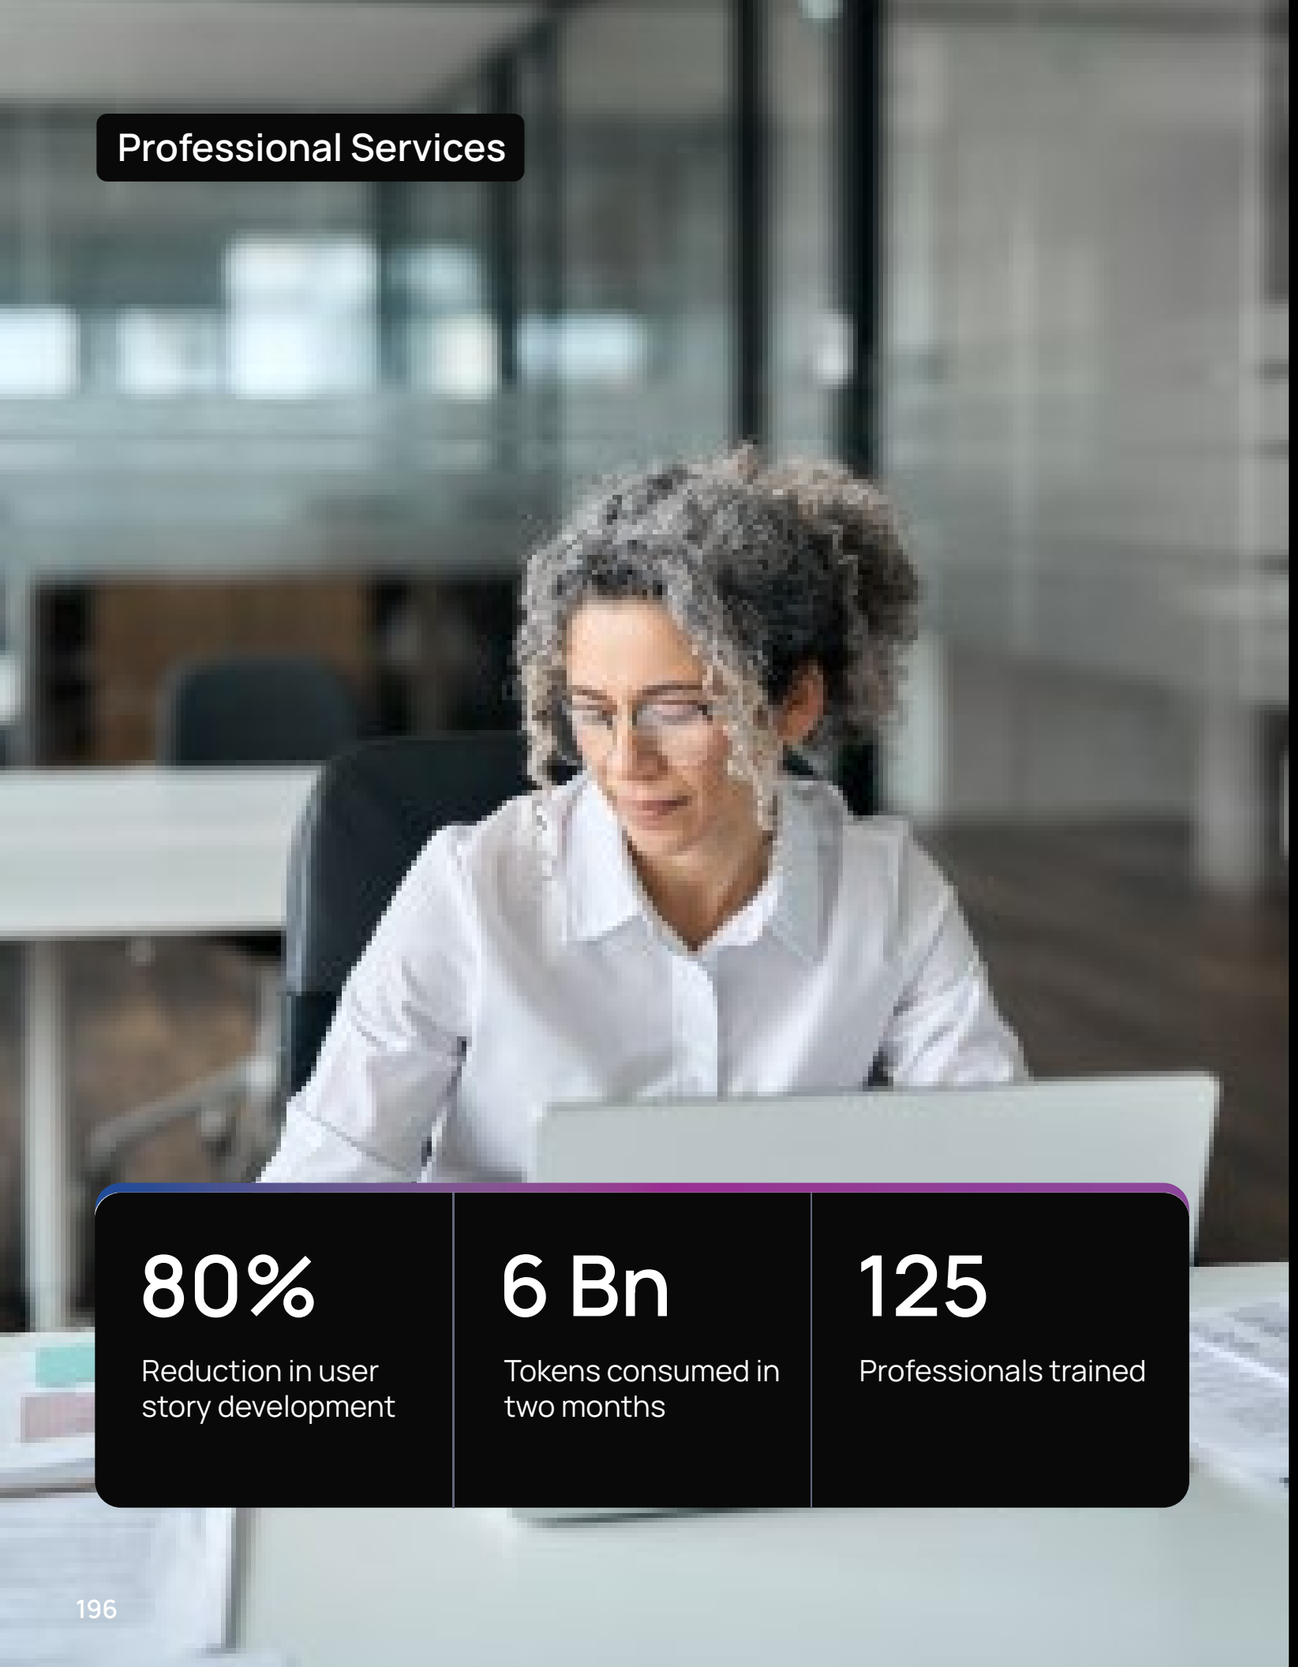

Client

A Leading Global Professional Services Organization From the Big 4

Turning audit innovation into something teams can use in weeks, not quarters

The client wanted to embed AI across audit workflows to improve quality, accelerate insights, and build trust in responsible adoption. The opportunity was as much about changing how work gets done as it was about building new tools.

Solution

- Used markdown-based UIs to create fast, consistent interfaces and speed deployment.
- Shifted key interactions from manual inputs to AI-generated outputs for human review and approval.
- Enabled agent-to-agent orchestration across audit workstreams.
- Built on a scalable AI tool chain designed for secure development and review.

Benefits

- Reduced user story development effort by 80%.
- Reached measurable results within just 60 days.
- Consumed six billion tokens in two months to power scaled experimentation and delivery.
- Trained 125 professionals on AI platforms, helping the capability spread beyond a single team.

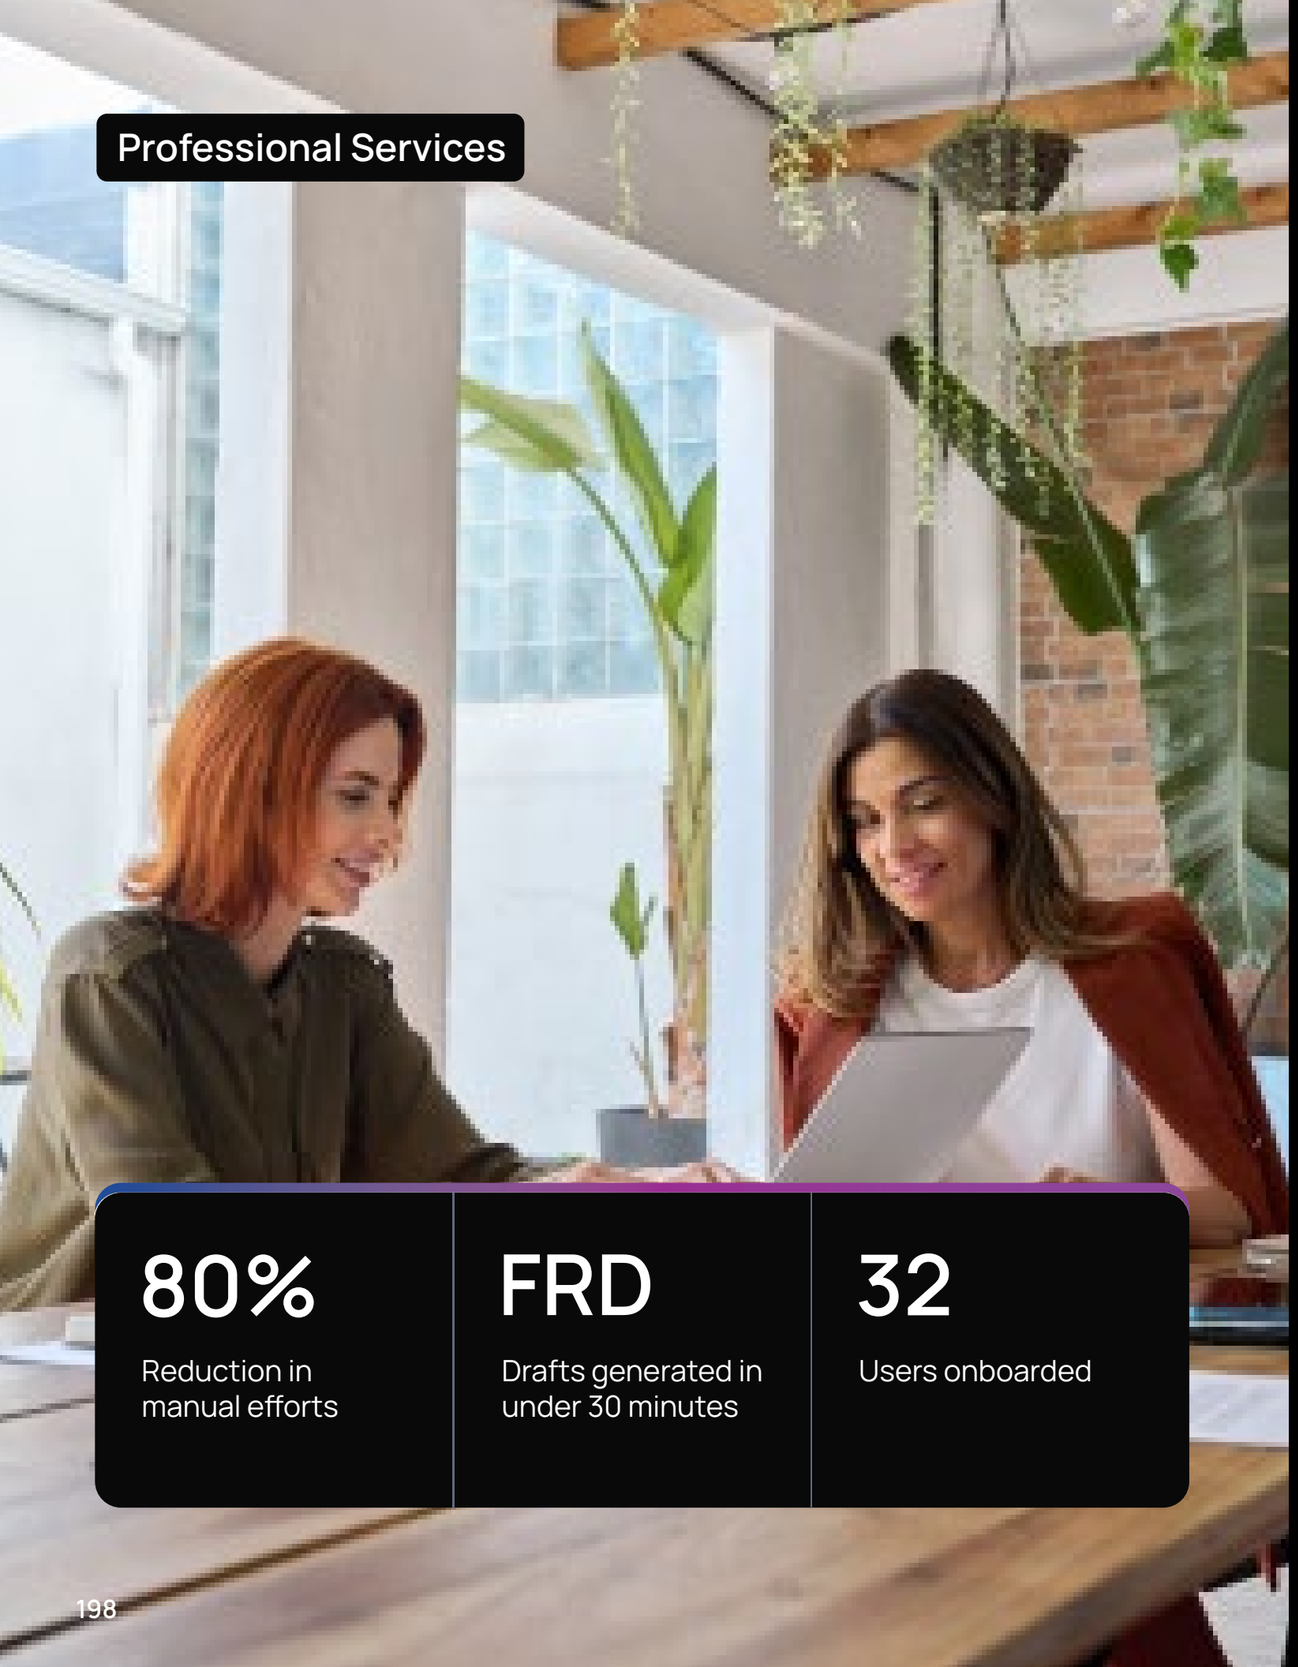

Client

An American Public Accounting Firm

Standardizing financial reporting direction (FRD) creation so quality no longer depends on who writes it

Inconsistent FRD creation across enterprise resource platform (ERP) teams was driving variable quality, heavy manual effort, and scaling challenges. The client needed a more repeatable documentation model that would improve both speed and consistency.

Solution

- Designed a GenAI-driven workflow to modernize the FRD documentation lifecycle.
- Used an MVP approach to validate value quickly and support broader rollout.
- Automated conversion of raw inputs into standardized blueprint documents.
- Built a reusable master FRD template to support consistency across teams.

Benefits

- Reduced manual effort by 80%.
- Generated FRD drafts in under 30 minutes, with manual effort cut to roughly 20–30%.
- Onboarded 32 users and completed 28 projects, including 21 final deliverables using AI-generated FRDs.
- Created the firm's first enterprise-grade AI application for ongoing operational use.

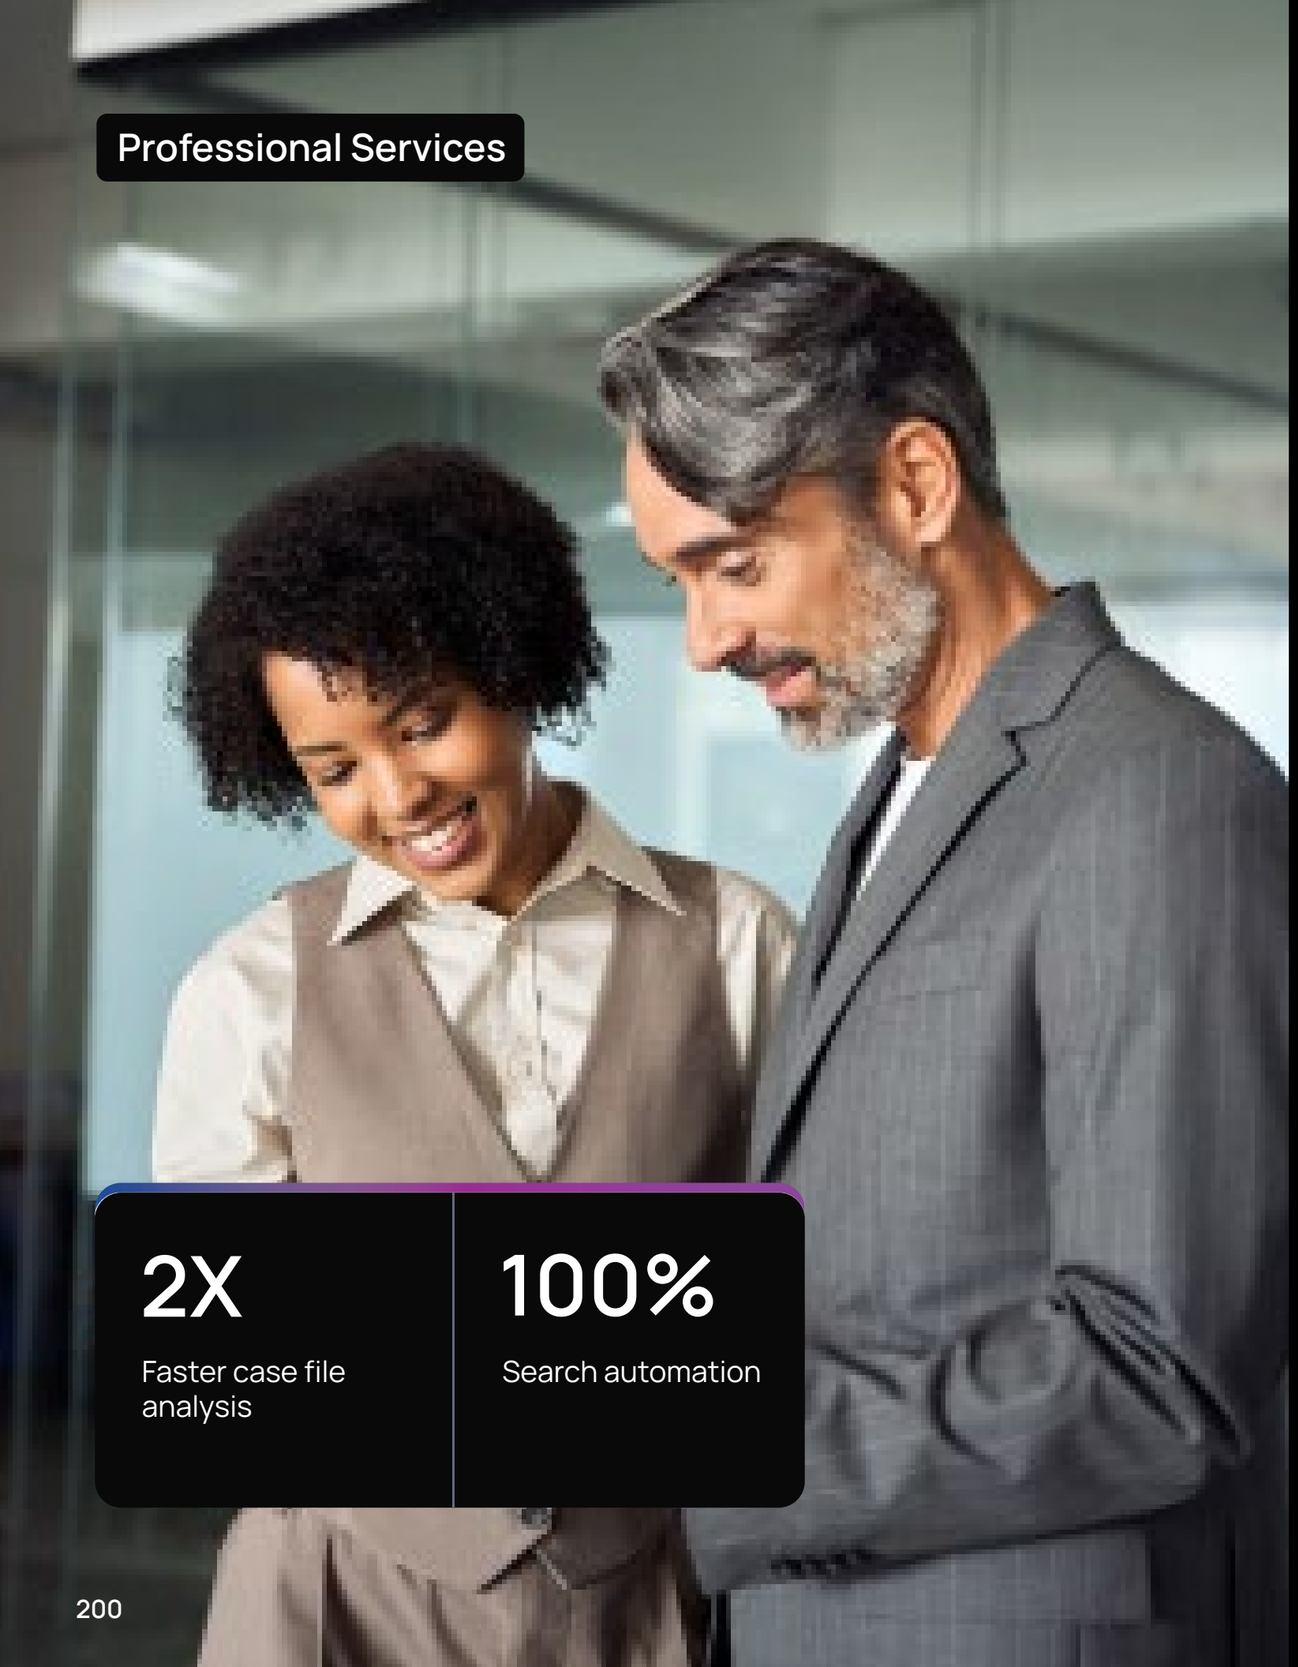

Client

A Corporate Legal Services Company

Helping paralegal consultants get to critical facts faster

Medical consultants were spending too much time digging through hundreds of pages of case files to find the facts that mattered most. The client needed a faster, more consistent way to surface key information without missing important context.

Solution

- Developed an AI-powered search solution for PDFs up to 500MB in size.
- Extracted simple and complex case information, including names and root causes of injury.
- Designed the system for very large files, often 700–1,000+ pages long.
- Automated deep document analysis to reduce manual searching and review effort.

Benefits

- Doubled case file analysis speed.
- Achieved 100% search automation across the targeted workflow.
- Improved consistency and accuracy in extracting critical information.
- Reduced manual effort on large, high-stakes medical case reviews.

Manufacturing & Consumer

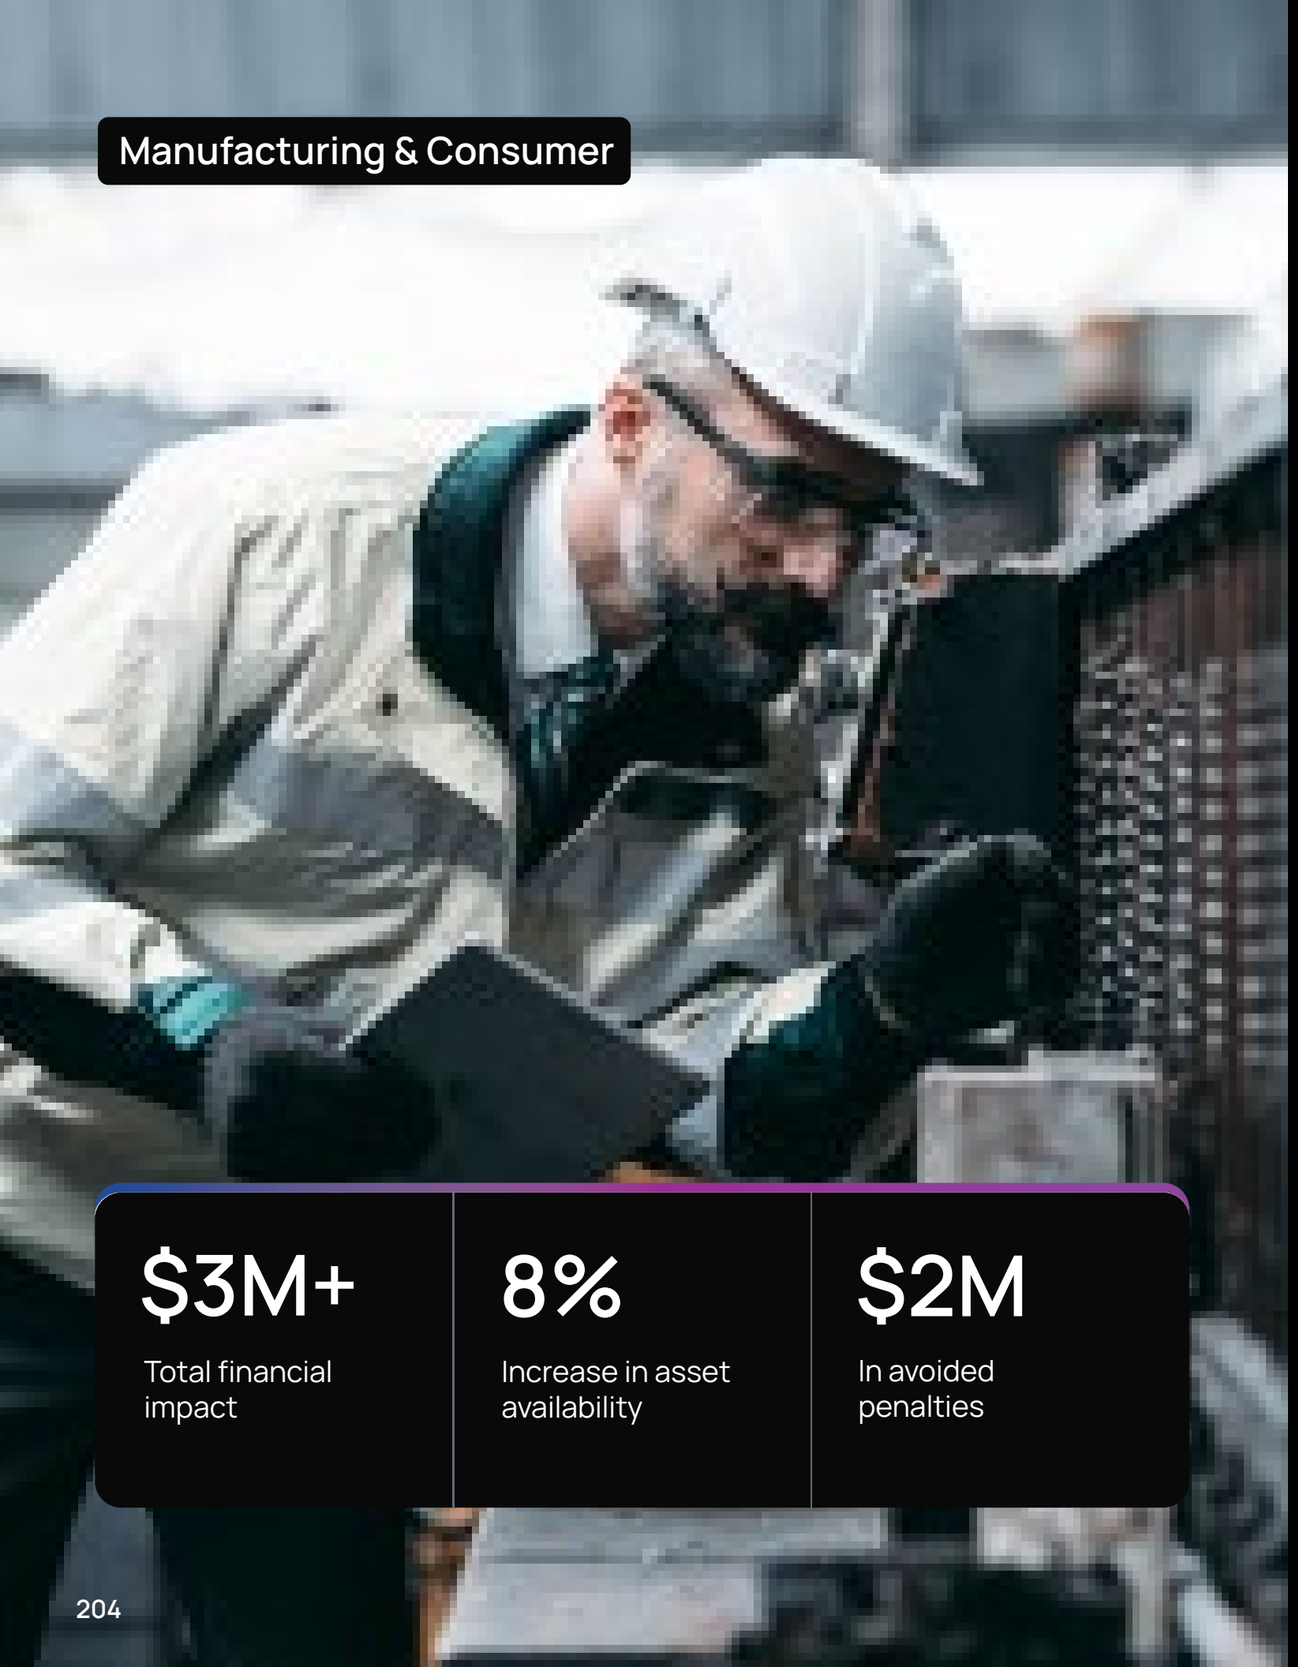

Client

An American Multinational Conglomerate Founded in the 1800s

Catching battery issues earlier to protect uptime, safety, and margin

Battery storage systems were losing availability to degradation, hidden faults, and costly downtime. The client needed earlier visibility into risk so it could prevent penalties, improve maintenance timing, and protect safe operations.

Solution

- Built AI/ML models to predict battery degradation and remaining useful life.
- Delivered early fault detection to surface performance and safety risks sooner.
- Implemented a high-frequency IoT telemetry pipeline for real-time visibility.
- Provided dashboards and predictive alerts tailored to utility-scale BESS operations.

Benefits

- Created $3M+ in total financial impact, including $1M in maintenance savings and $2M in avoided penalties.
- Increased asset availability by 8%.
- Recorded zero fire-safety incidents through earlier fault detection and better cycle optimization.
- Shifted operations from reactive maintenance to more proactive, data-driven decisions.

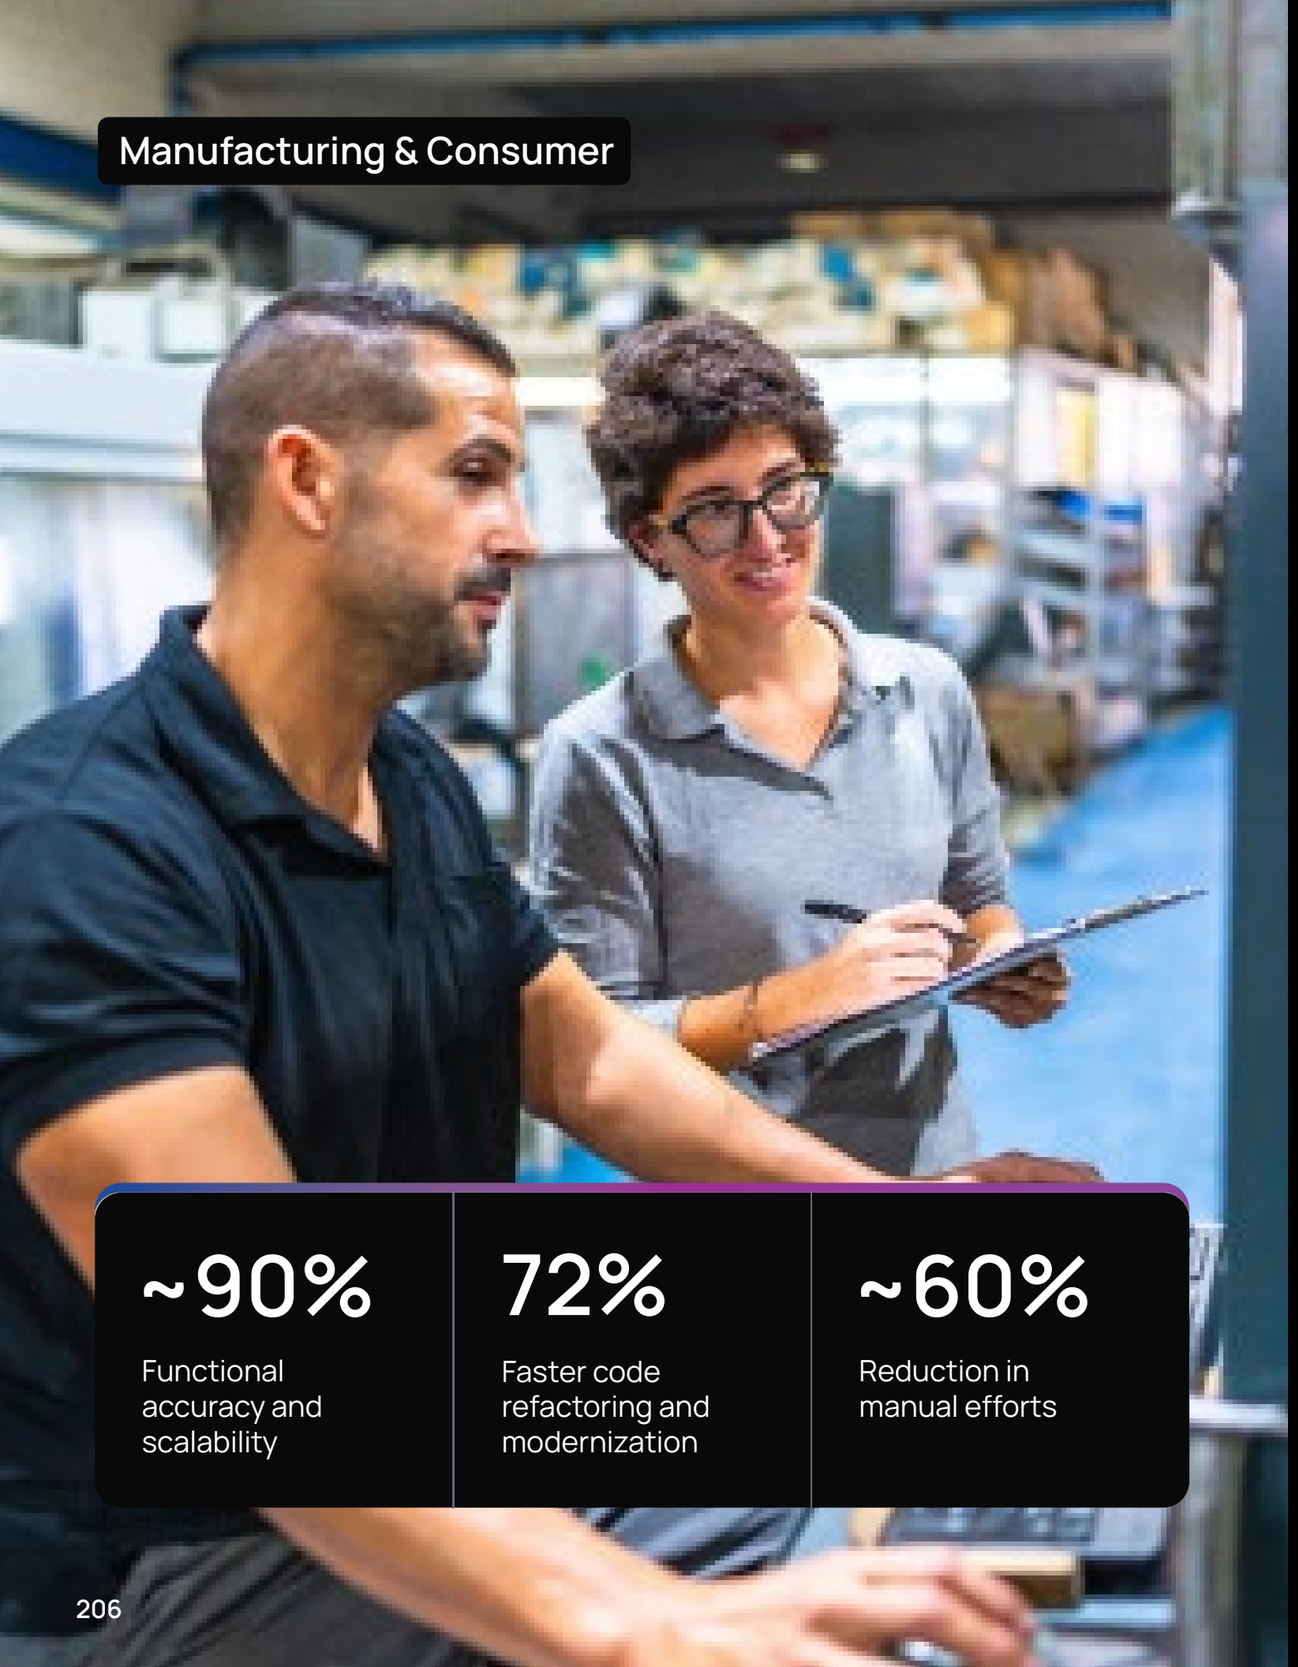

Client

An American Business Service Provider

Extracting hidden business rules before they block modernization

Critical business rules were buried in opaque legacy systems, making modernization slower, riskier, and far too manual. The client needed a faster way to uncover, validate, and reuse the logic driving mission-critical operations.

Solution

- Used Hexaware's RapidX® framework to support end-to-end AI-led modernization.
- Automated COBOL-to-Java conversion to reduce manual refactoring effort.
- Performed deep legacy analysis to identify and interpret decades-old business logic.
- Externalized business and exception rules into clearer, more reusable structures.

Benefits

- Delivered about 90% functional accuracy and scalability in the targeted transformation.
- Accelerated code refactoring and modernization by 72%.
- Reduced manual effort by about 60%.
- Lowered technical debt and modernization risk while preserving business logic.

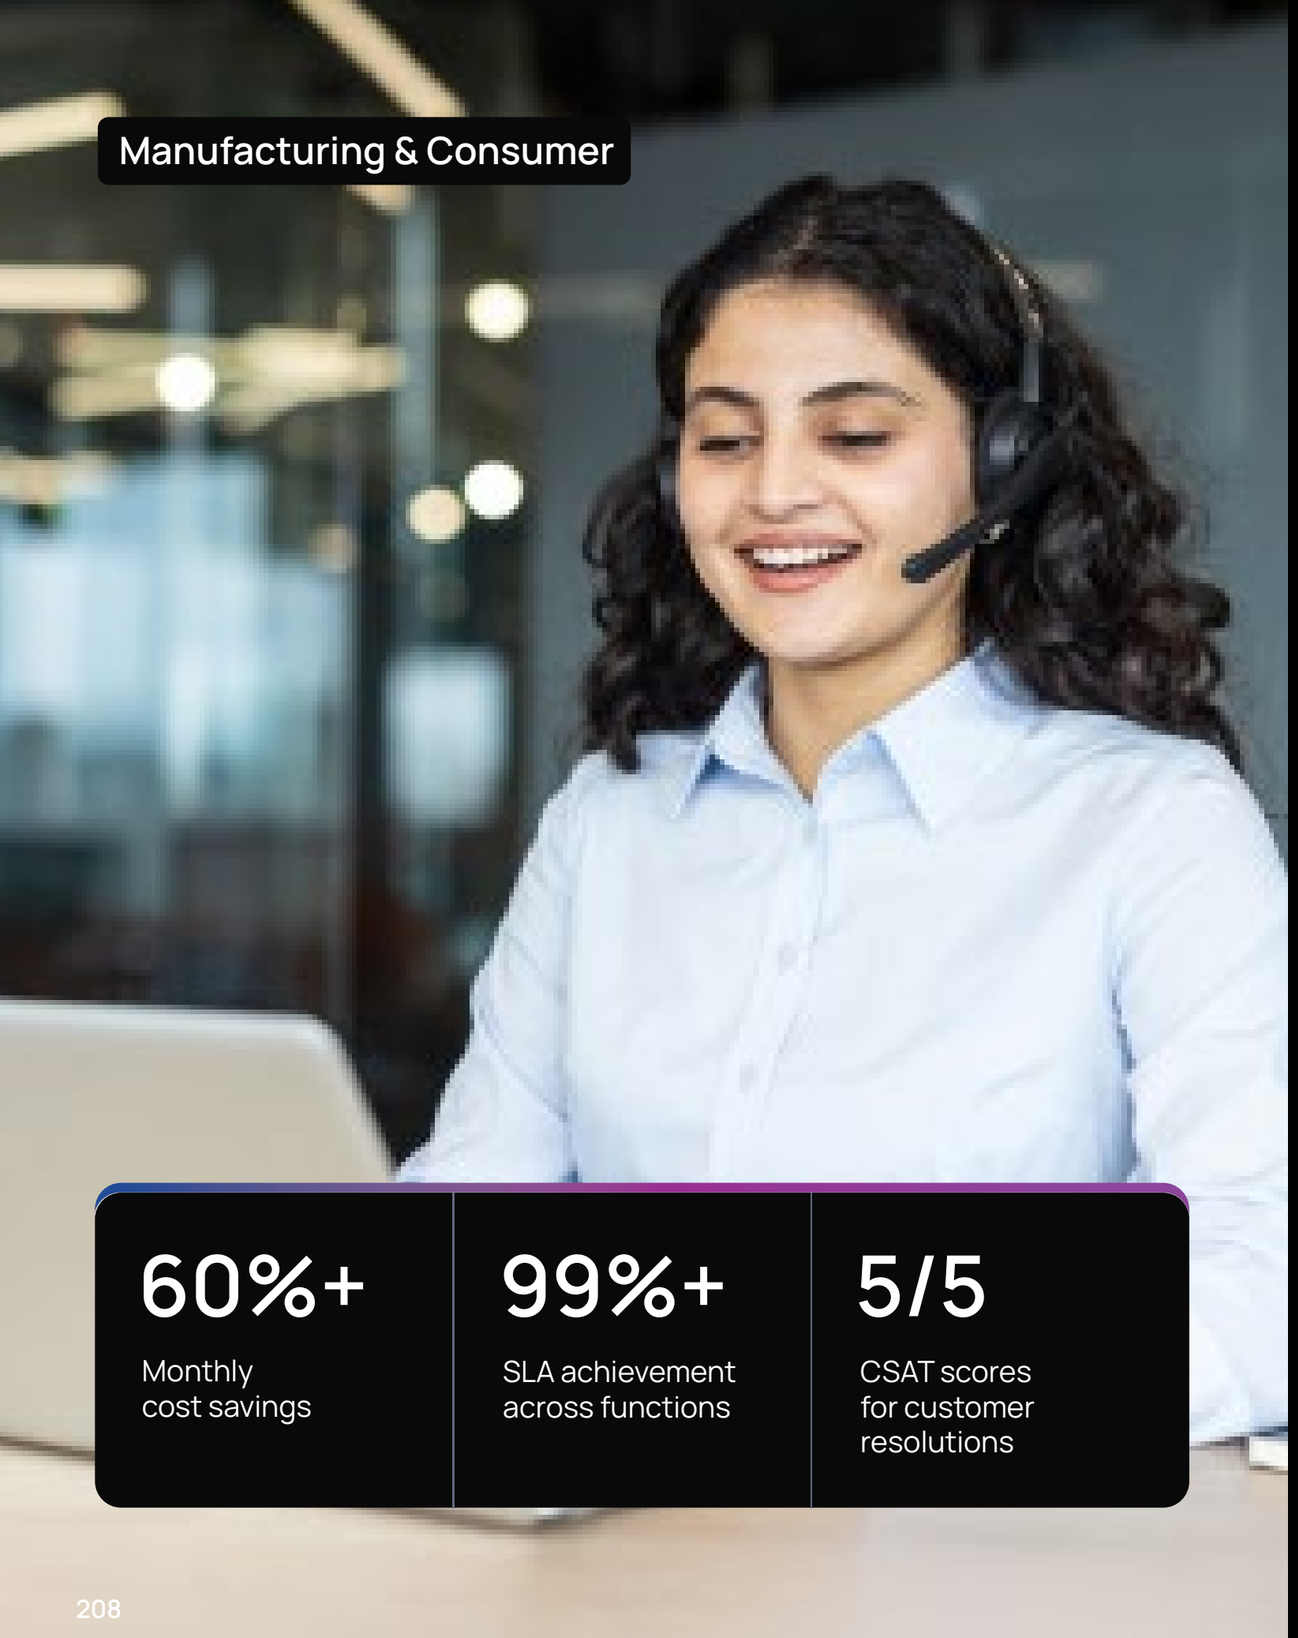

60%+
Monthly cost savings

99%+
SLA achievement across functions

5/5
CSAT scores for customer resolutions

Client

A Leading Nordic Telecom Provider

Enabling multilingual customer support with a compliant AI translation model

Addressed rising customer operations costs, limited automation, and the challenge of extending chat support beyond standard business hours without increasing native-language staffing costs.

Solution

- Built a real-time, two-way AI translation solution embedded into the client's contact center platform, enabling English-speaking agents to support live customer conversations across languages.
- Designed the solution around strict regulatory requirements, with hosting inside the client's environment, PII masking, message-level traceability, audit trails, and compliance reporting aligned to PTS and GDPR obligations.
- Preserved interaction quality through language detection, domain-specific glossaries, Swedish product-term controls, and linguistic audits supported by a Swedish-speaking SME in the delivery team.
- Delivered technology deployment, agent onboarding, training, and live operations in parallel within a compressed three-month timeline, with Hexaware owning SLAs from day one.

Benefits

- Achieved 60%+ monthly cost savings.
- Maintained 99%+ SLA achievement across functions.
- Handled 2x volume surge without disruption.

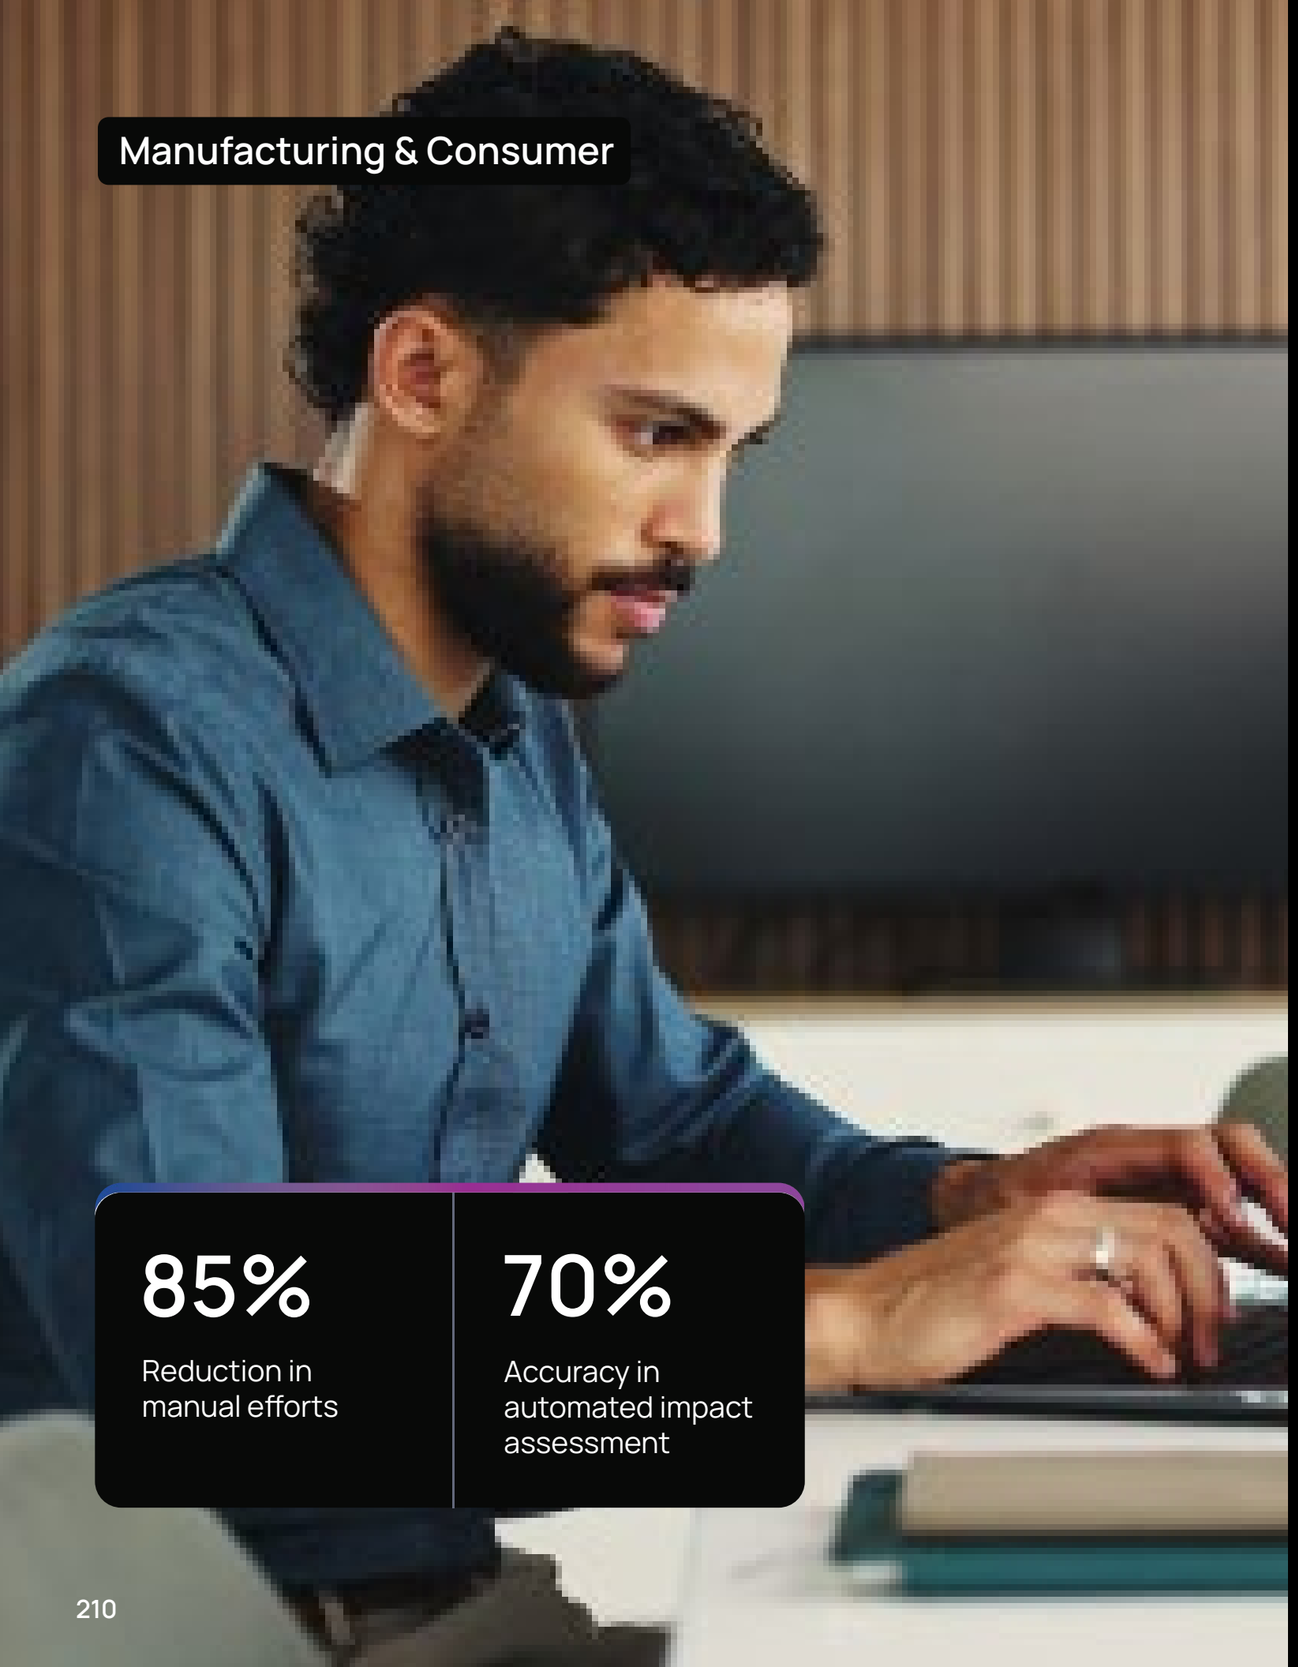

Client

An American Multinational Conglomerate Founded in the 1800s

Reducing compliance-heavy design change work without lowering rigor

Regulated design changes were slowed by document-heavy assessments and repeated review cycles that pulled engineering capacity away from core product work. The client needed to cut the burden without weakening control.

Solution

- Embedded an agentic AI assistant into the software development lifecycle (SDLC) for design-change analysis.
- Automated completion of a 145-question impact assessment.
- Used agentic retrieval augmented generation (RAG) to retrieve technical, historical, and compliance context from engineering artifacts.
- Flagged likely downstream document updates while keeping humans in the validation loop.

Benefits

- Reduced manual effort by 85% across the pilot workflow.
- Reached 70% accuracy in automated impact assessment generation, supporting reliable human review.
- Saved several weeks in end-to-end change cycles.
- Freed engineering teams from repetitive compliance tasks so they could focus more on development work.

85%	~70%	50%+
Relevance and grounding accuracy	Reduction in manual context gathering	Acceleration in first-draft outputs

Client

A US Corporation Specializing in Semiconductor, Display, and Solar Manufacturing Equipment

Creating a reusable AI factory, developers can keep building on

Developers were losing time across fragmented systems, low-quality scanned documents, and too much context switching. The client needed a reusable AI foundation that could support everyday engineering work and larger modernization programs alike.

Solution

- Delivered a modular agentic AI stack with reusable capabilities for software development lifecycle (SDLC) workflows.
- Added developer-focused agents to reduce system switching and accelerate issue resolution.
- Built a modernization accelerator for monolith-to-microservice analysis and conversion.
- Established an Azure-based retrieval augmented generation (RAG) and optical character recognition (OCR) foundation for engineering documents and artifacts.

Benefits

- Achieved 85% relevance and grounding accuracy.
- Reduced manual context-gathering effort by about 70% for key developer workflows.
- Accelerated first-draft outputs for modernization work by 50%+.
- Delivered sub-five-second response times and 90%+ OCR accuracy in targeted use cases.

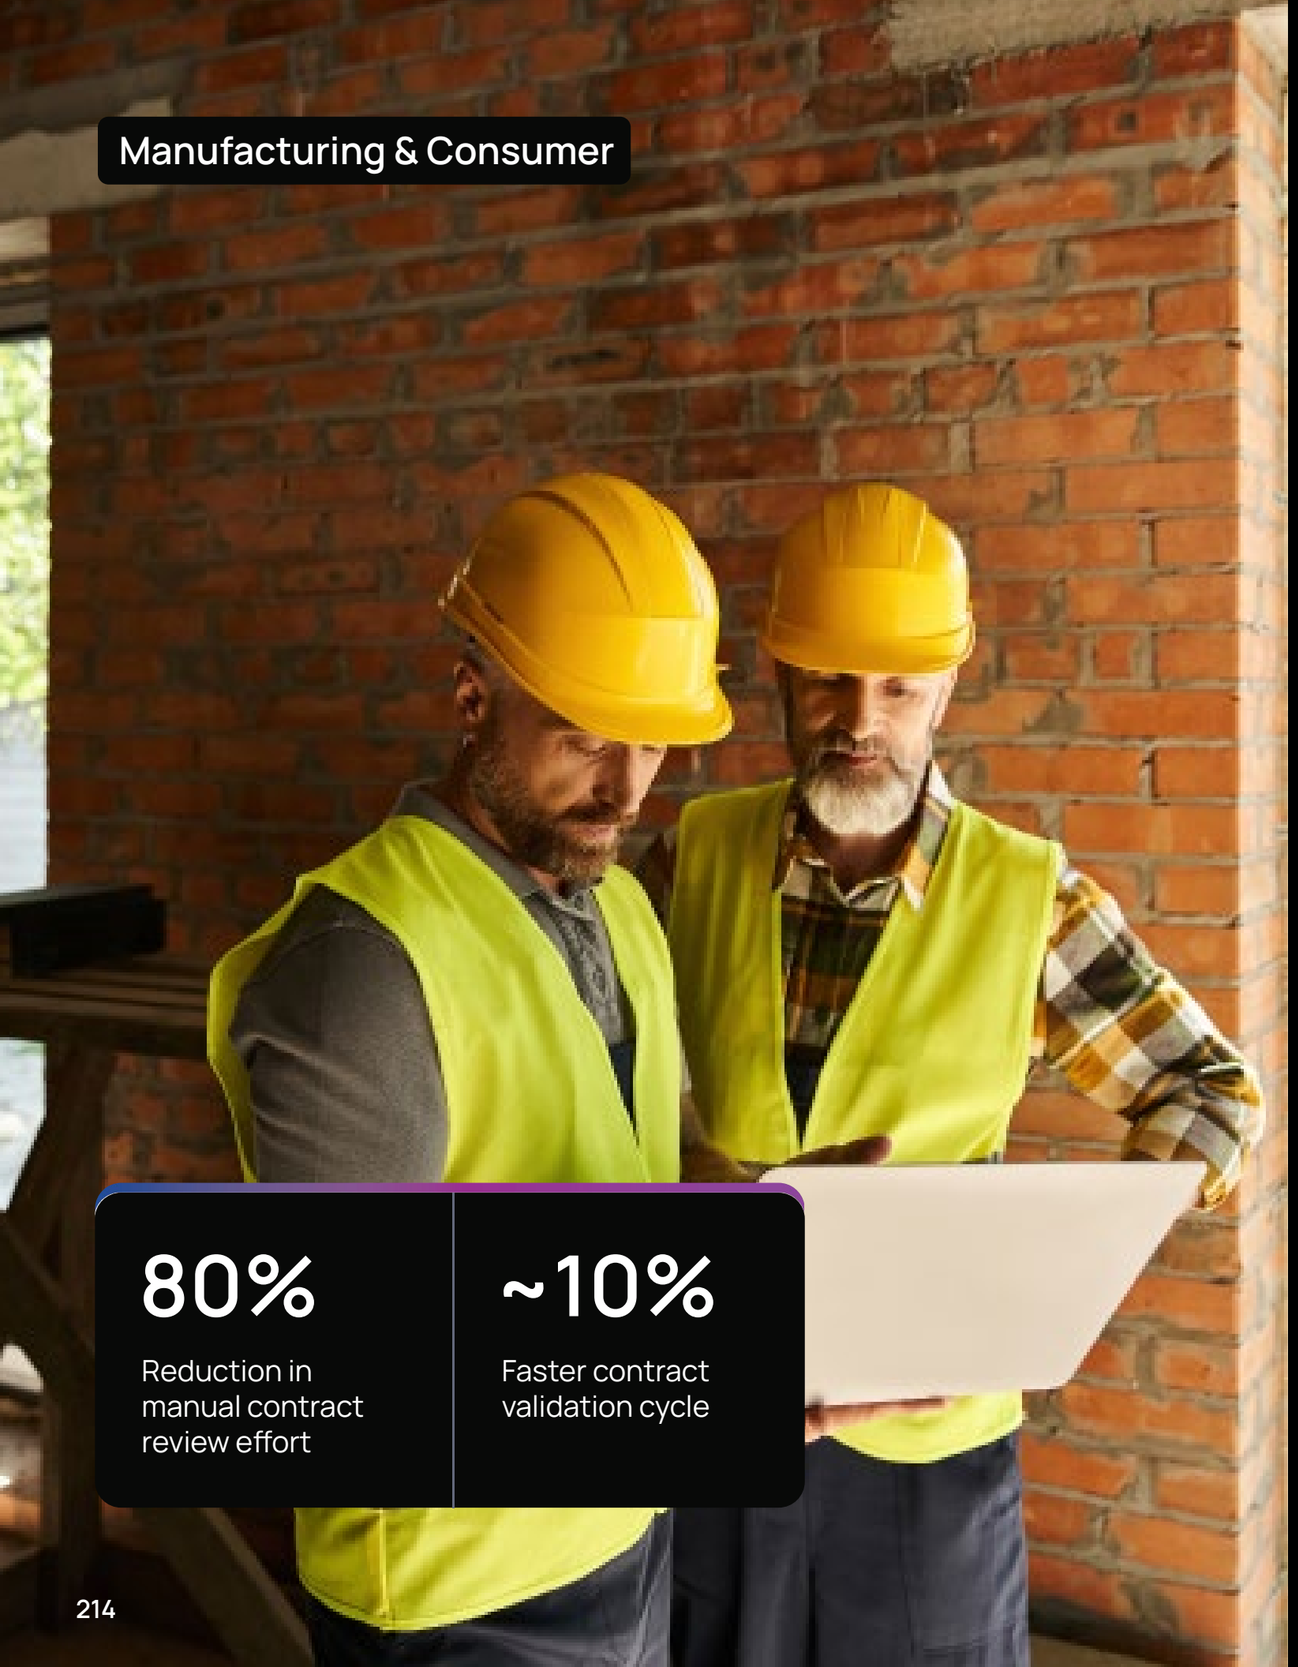

Client

A Global Construction Firm Delivering Large-Scale Commercial and Infrastructure Projects

Scaling contract compliance without scaling manual legal review

Thousands of vendor agreements were creating slow subcontractor onboarding and inconsistent compliance checks. The client needed a more automated process that could keep projects moving without increasing legal exposure.

Solution

- Automated contract analysis with AI-driven clause extraction and compliance validation.
- Used agentic workflows to classify contracts and route exceptions.
- Integrated with procurement and document systems for end-to-end workflow continuity.
- Applied playbook-based validation and transparent recommendations to support auditability.

Benefits

- Reduced manual contract review effort by about 80%.
- Accelerated contract validation cycles by roughly 10x.
- Improved subcontractor onboarding speed and project mobilization timelines.
- Freed legal and procurement teams to focus more on strategic negotiation and risk management.

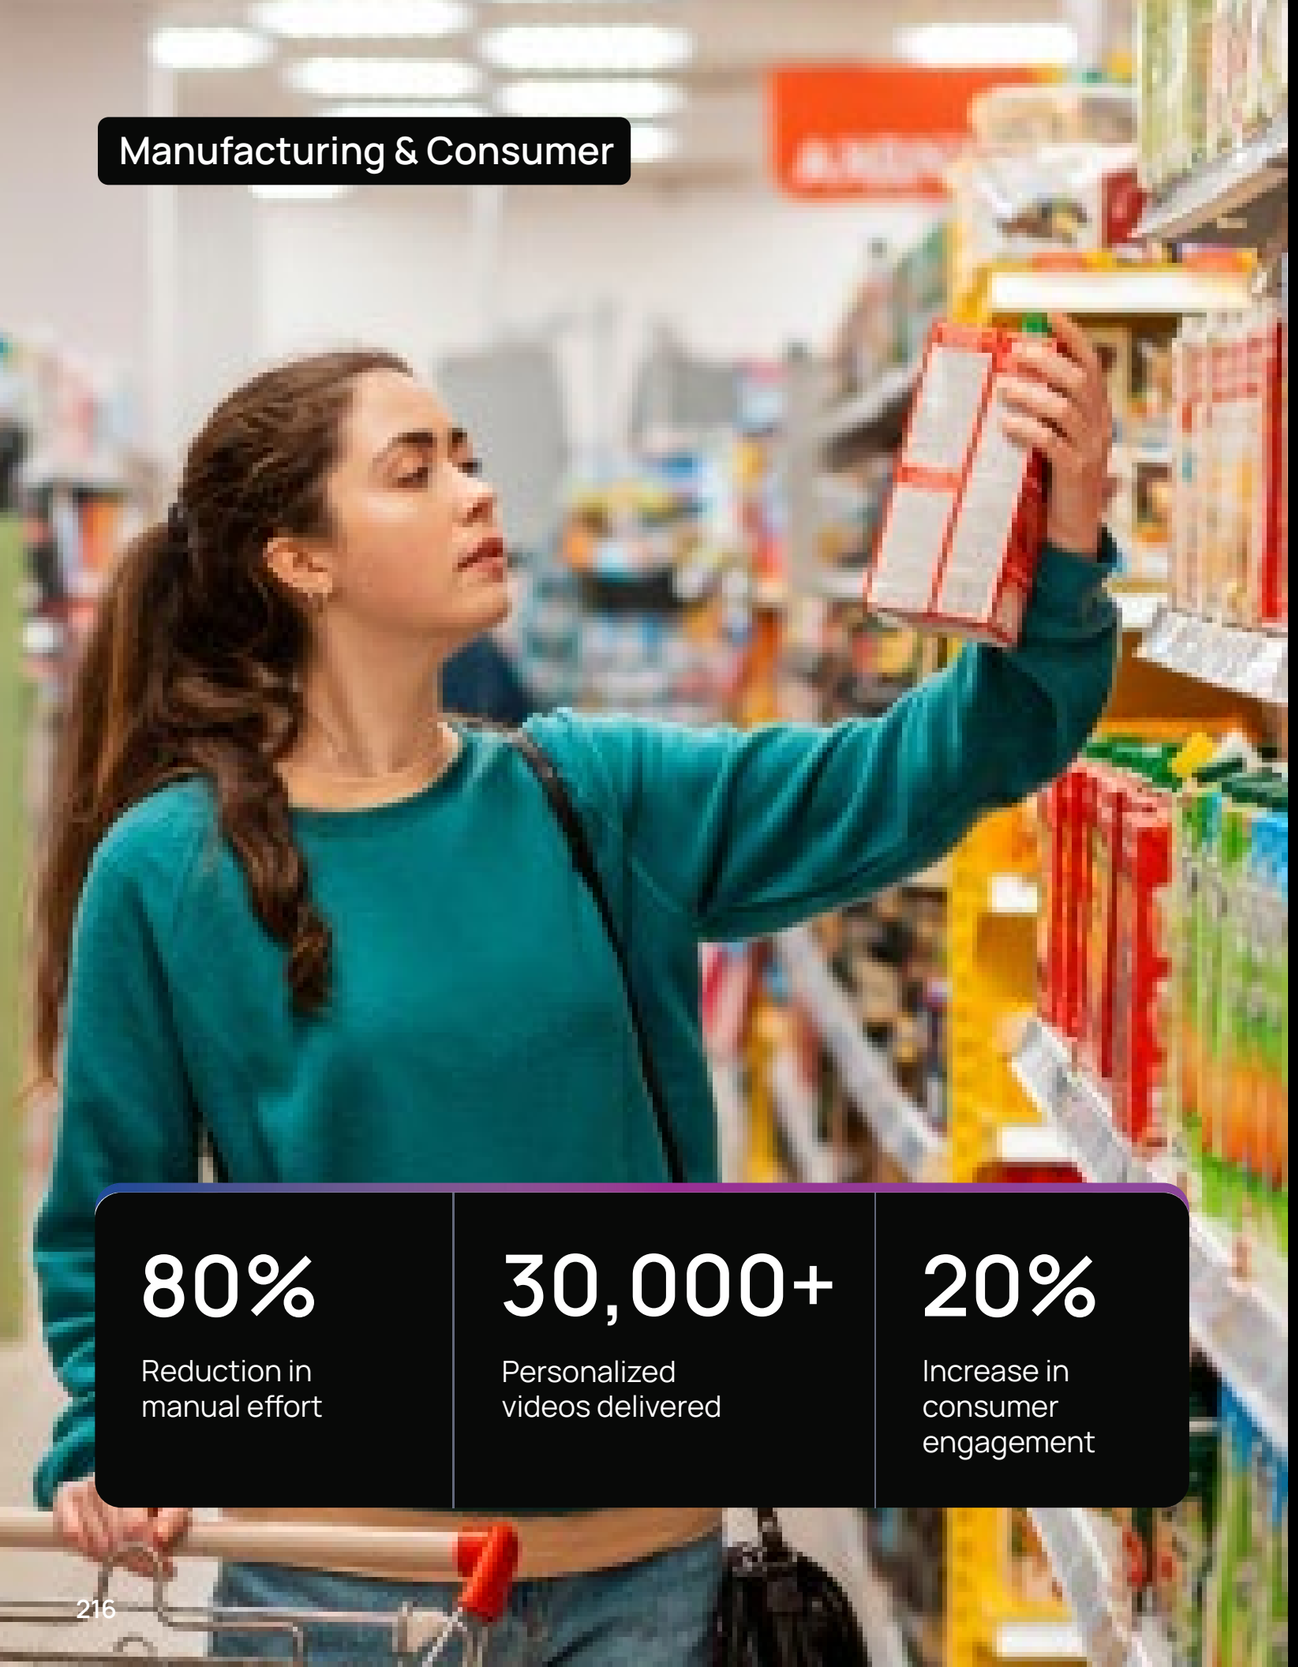

Client

A British Multinational Consumer Packaged Goods Company

Delivering personalized video at market scale, not studio pace

The client wanted to create culturally relevant, personalized video content across diverse markets without the usual production cost and localization delays. The goal was higher engagement with a model that could scale across brands and regions.

Solution

- Built a GenAI engine that generated personalized videos instantly through WhatsApp.
- Trained AI avatars on local celebrities to make content more culturally resonant.
- Enabled name-based and preference-based personalization across languages and markets.
- Deployed a reusable Azure-based architecture integrated with the WhatsApp Business API.

Benefits

- Reduced manual effort by 80%.
- Delivered 30,000+ personalized videos with a 99%+ success rate.
- Increased consumer engagement by 20%.
- Improved brand connection by 40% across key growth markets.

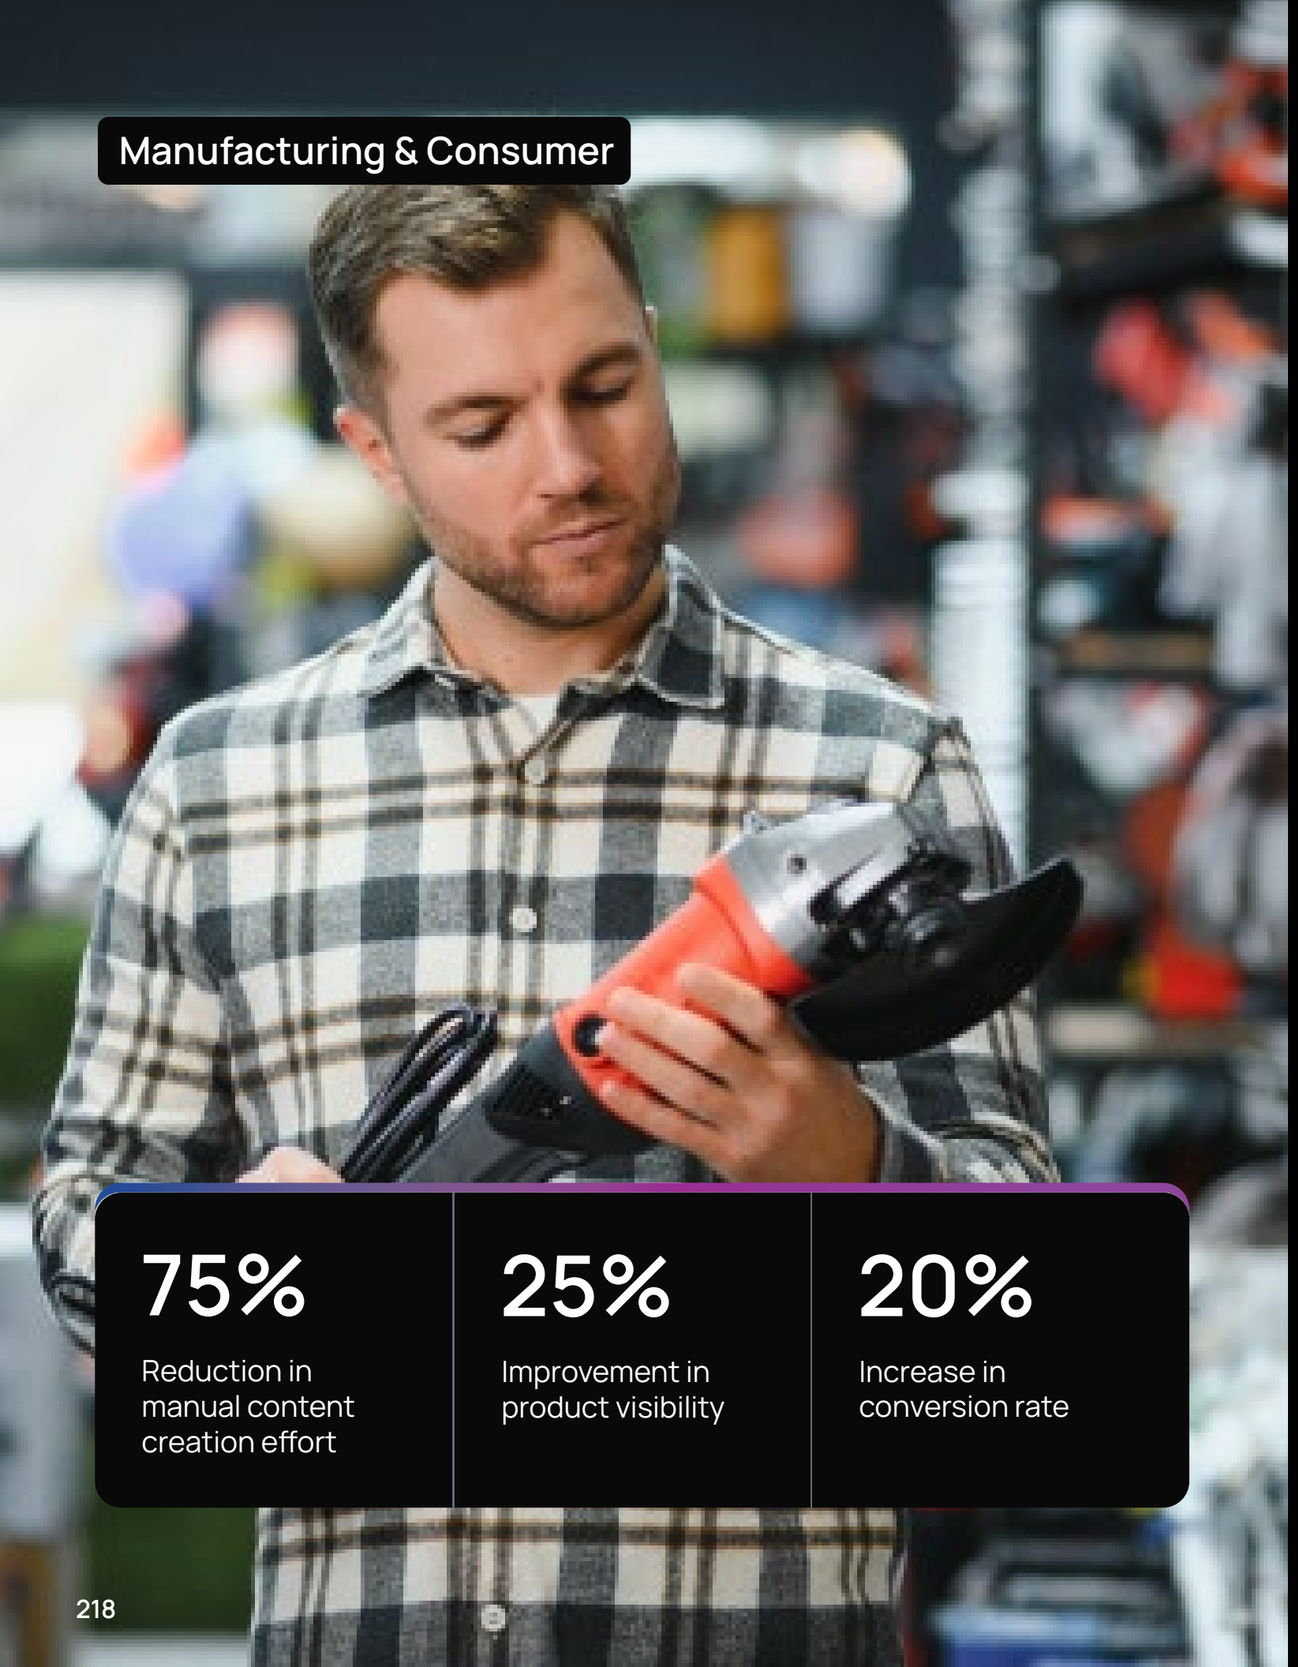

Client

A Leading Global Wholesaler of Home Improvement and Hardware Products

Refreshing product content fast enough to improve discovery and conversion

Incomplete and outdated product descriptions were making it harder for customers to find the right products and harder for the business to scale high-quality content across a huge catalog. The client needed a repeatable content engine, not more manual copy work.

Solution

- Built a GenAI product description engine on GCP using PaLM 2.
- Consolidated and enriched metadata from Oracle product data hub (PDH), AWS, and GCP sources.
- Fine-tuned models across 19 departments and 3,000+ subclasses.
- Generated separate SEO-optimized descriptions for B2B and B2C audiences and fed results back into PDH.

Benefits

- Reduced manual content creation effort by 75%.
- Improved product visibility by 25% through SEO-optimized descriptions.
- Increased conversion rate by 20% with higher-quality AI-generated content.
- Improved digital experience with more complete, relevant product information.

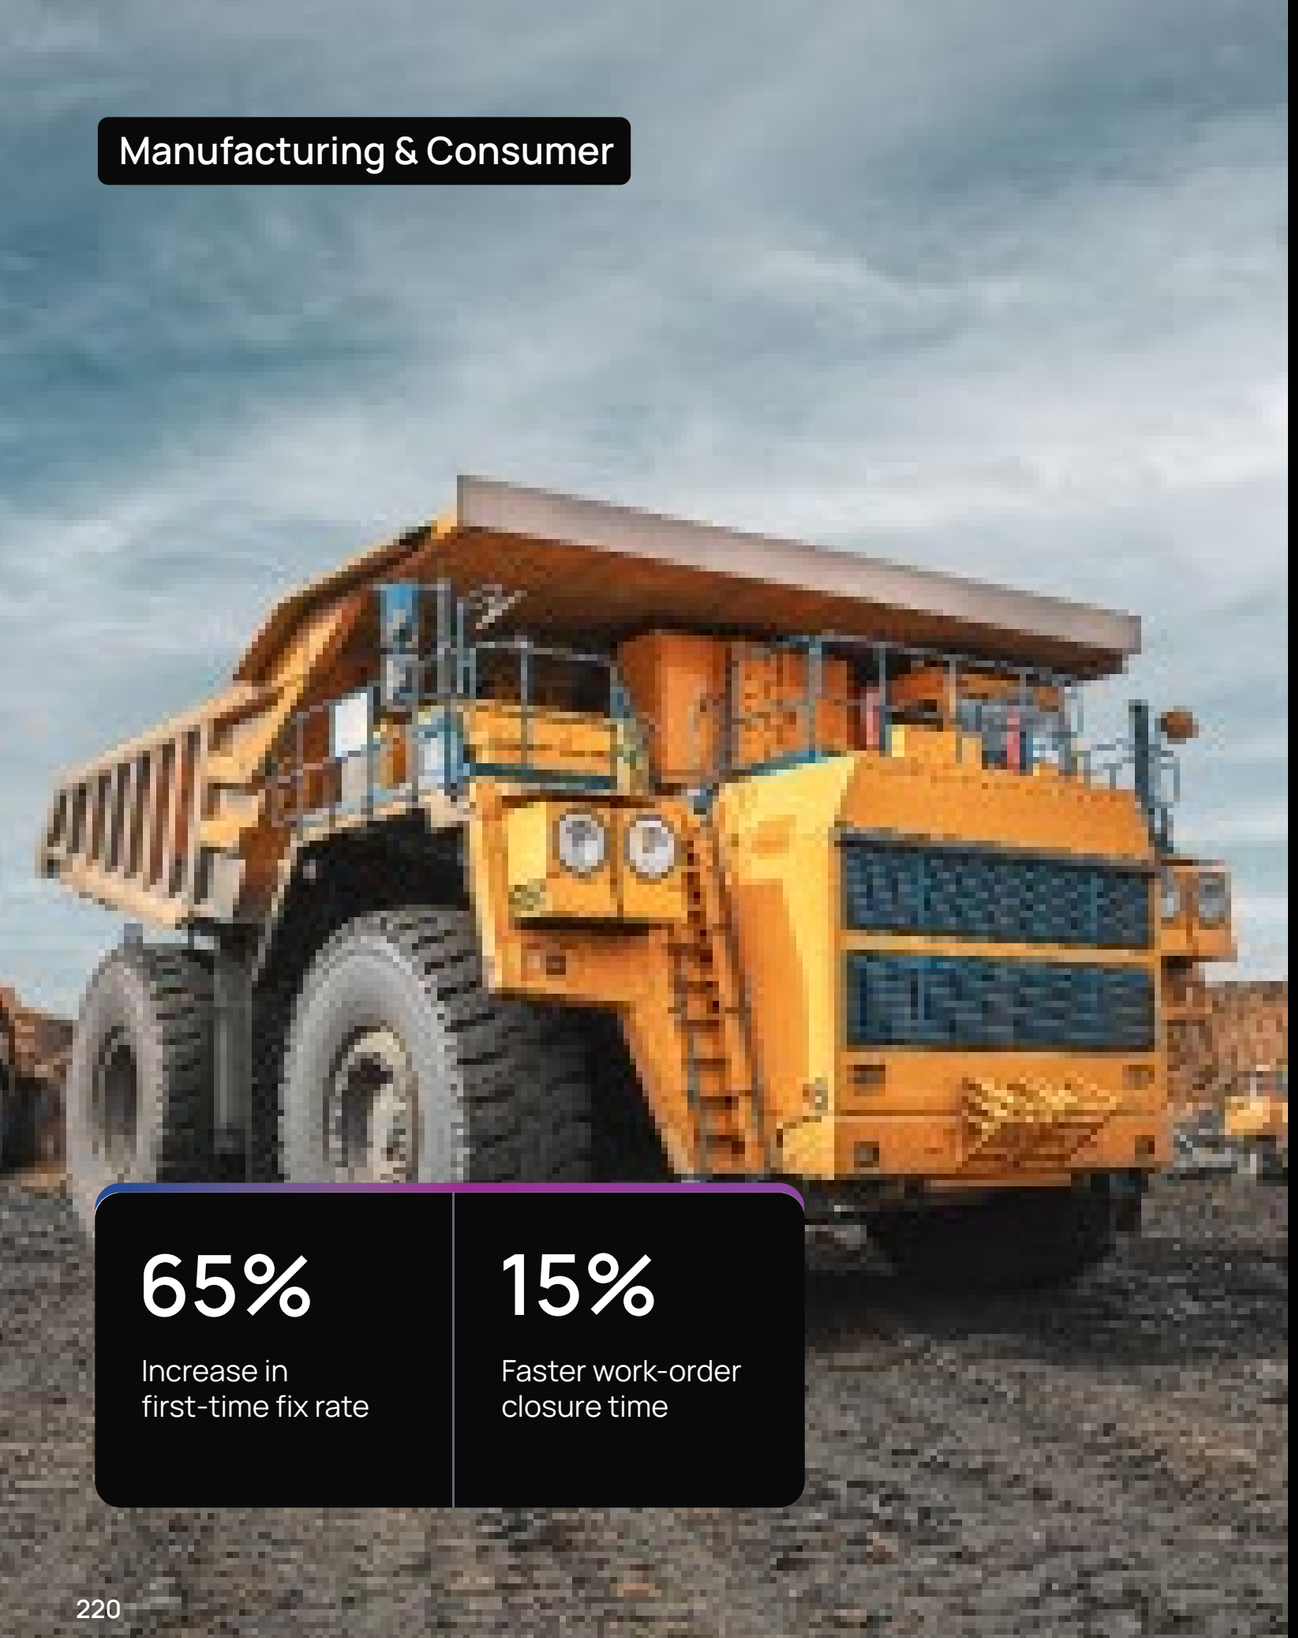

Client

An American Mining Company

Breaking knowledge silos so field teams can fix more on the first visit

Operational knowledge was scattered across telemetry, documents, and enterprise systems, slowing decision-making and increasing reliance on specialists. The client needed a more direct way for operators and maintenance teams to get grounded answers when timing mattered.

Solution

- Unified telemetry, enterprise data, and technical documents into a natural-language decision interface.
- Used agentic retrieval augmented generation (RAG) with Azure OpenAI to deliver grounded, source-linked responses.
- Added persona-based access controls plus human-in-the-loop validation to support trust and governance.
- Built a semantic asset layer and optimized data pipelines for large-scale telemetry environments.

Benefits

- Increased first-time fix rate by 65%.
- Reduced work-order closure time by 15% for maintenance technicians.
- Cut decision delays by giving teams faster access to grounded operational insight.
- Reduced reliance on specialists, dashboards, and manual reporting cycles.

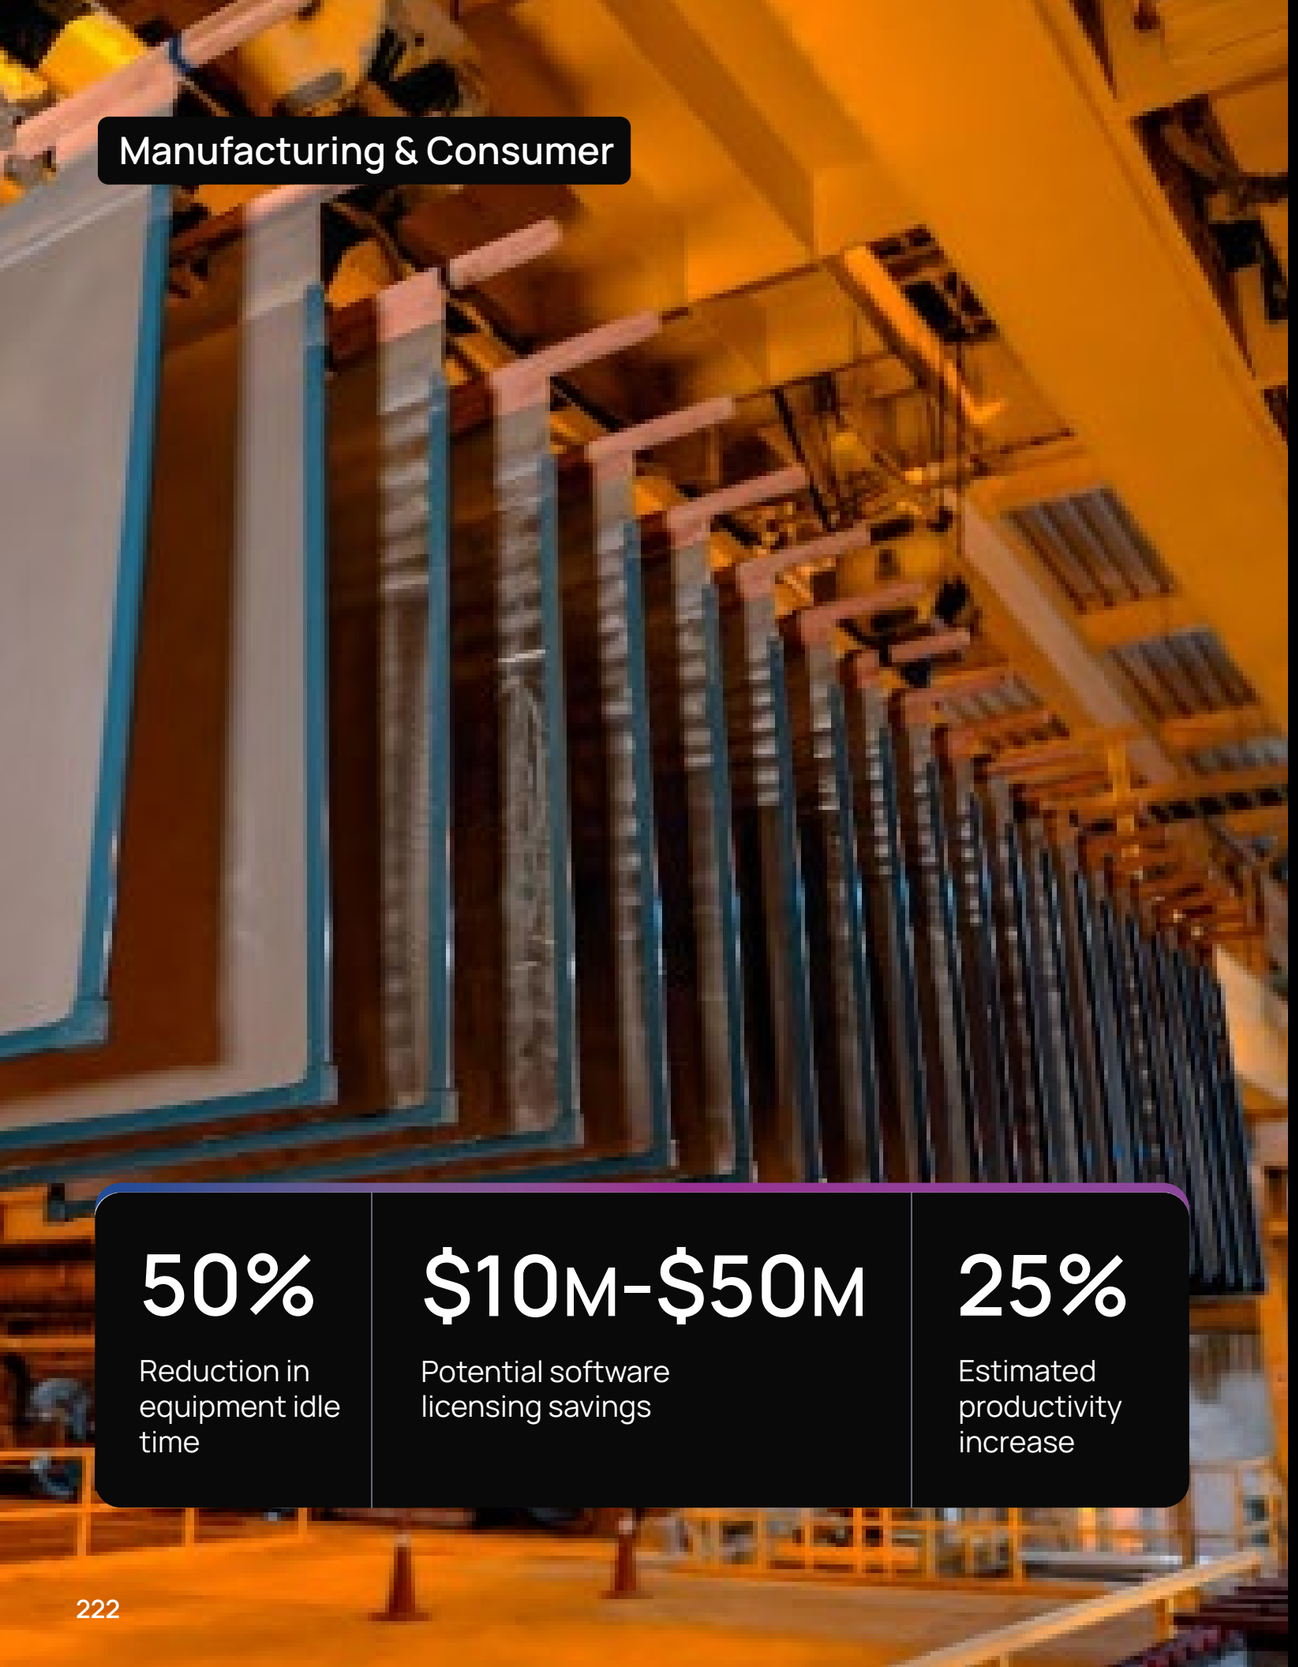

Client

A Global Mining Enterprise

Using AI and IoT to keep heavy equipment moving

The client wanted more real-time visibility into mining operations and a better way to predict maintenance needs before downtime spread. The aim was to improve safety, productivity, and decision speed across complex equipment environments.

Solution

- Applied IoT and AI analytics to deliver near-real-time operational intelligence.
- Supported safety and SOP compliance with advanced driver assistance and stratified alarms.
- Used GenAI to create current-state and forecast-driven operational models.
- Extended value with an AI copilot to support process optimization across the organization.

Benefits

- Reduced equipment idle time by 50%.
- Delivered an estimated 25% increase in productivity.
- Created potential software licensing savings estimated at $10M–$50M.
- Improved decision-making speed through richer operational visibility and predictive maintenance insights.

Client

A Global Leader Providing Semiconductor Manufacturing Equipment And Materials Engineering Solutions

Helping legal and procurement teams find contract risk before it slows the business

Contract data was scattered across systems and hard to compare, which meant longer reviews, more legal risk, and slower decision-making. The client needed a more searchable, dependable way to work through contract changes at scale.

Solution

- Centralized contracts and amendments in a unified secure repository.
- Automated extraction, comparison, and summarization of contract changes.
- Enabled natural-language search, contextual Q&A, and chronological amendment tracking.
- Flagged clause deviations and supported role-based access and auditability across enterprise systems.

Benefits

- Achieved 99.9% contract extraction accuracy.
- Eliminated large amounts of manual comparison effort through AI-driven review.
- Reduced compliance risk with faster clause discovery and deviation alerts.
- Freed legal teams to spend more time on strategic work instead of repetitive comparison.

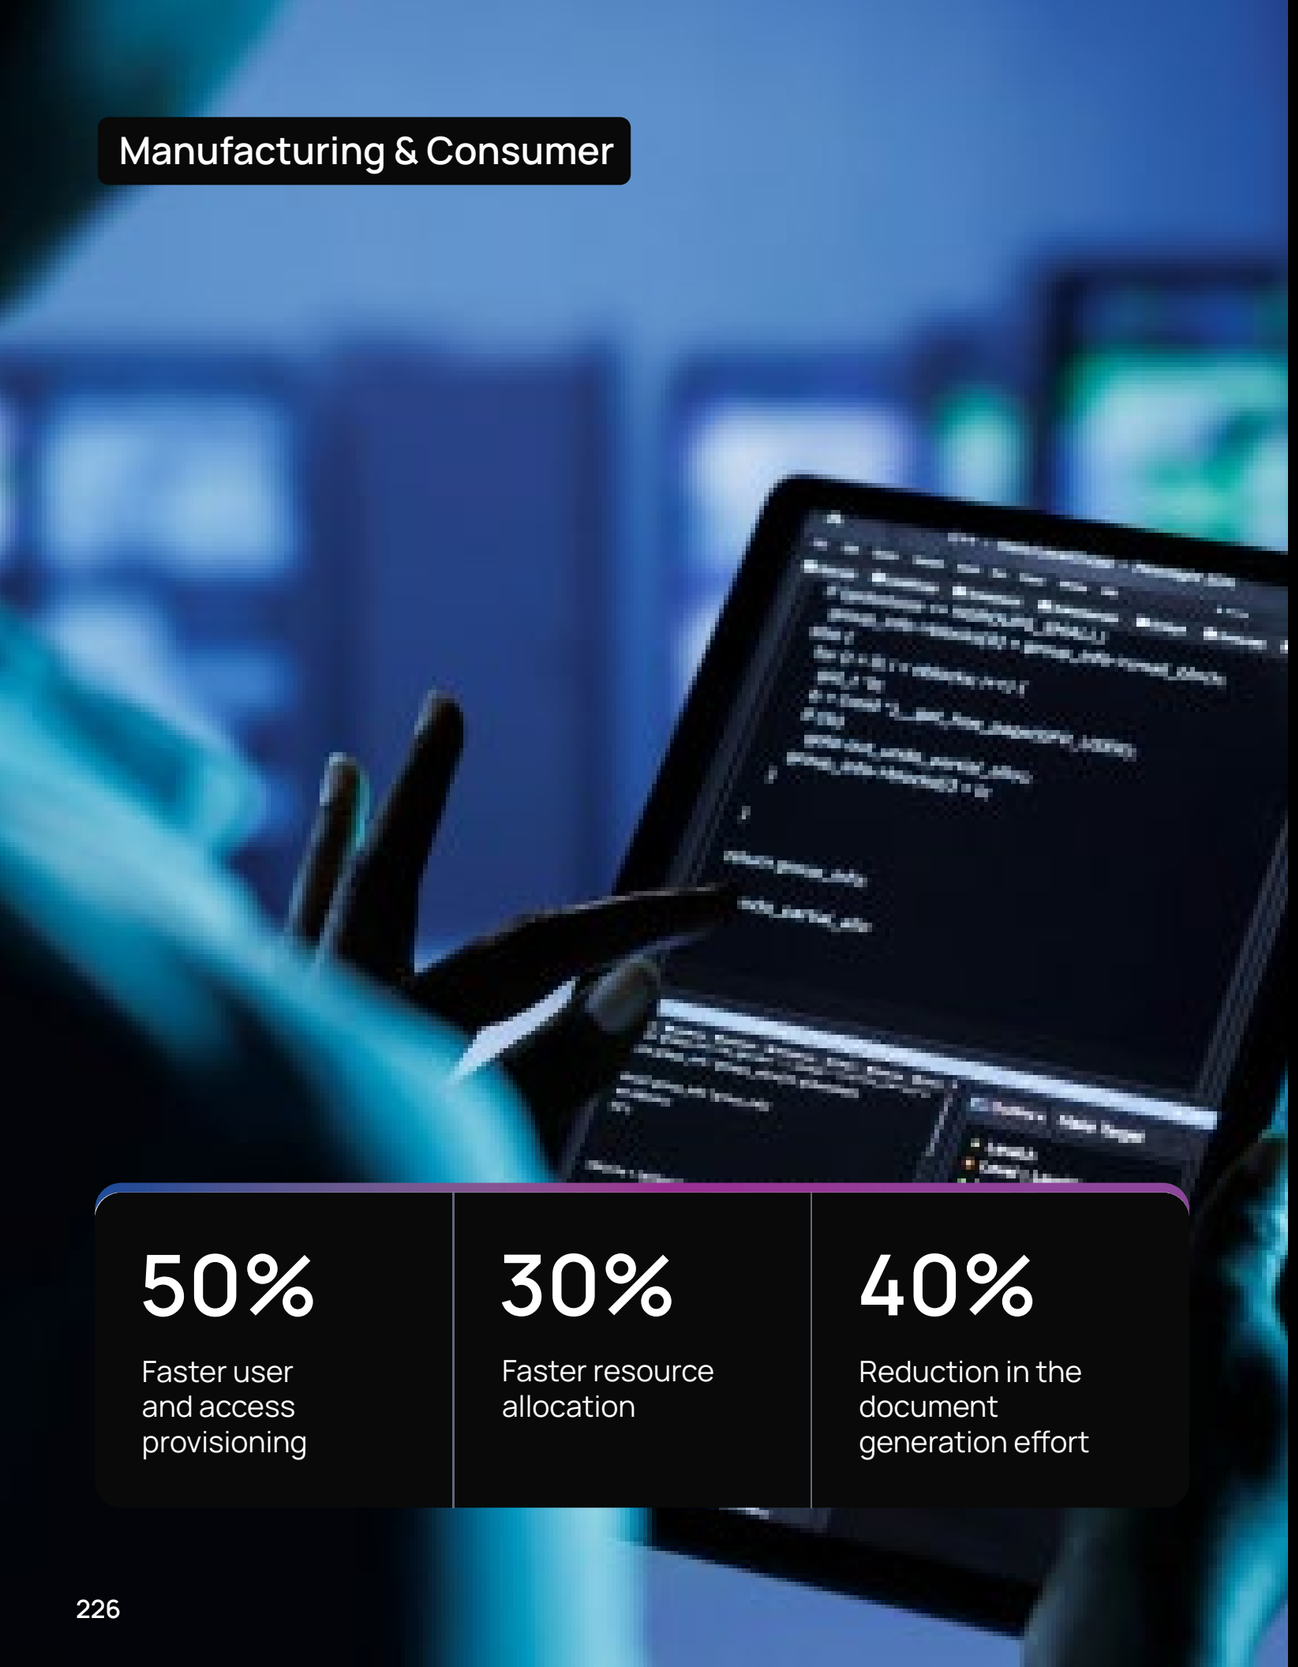

Client

A Fortune 500 Global Conglomerate

Modernizing workforce workflows inside Workday, not around it

Fragmented HR processes, security inconsistencies, and legacy constraints were slowing workforce operations across regions. The client needed faster, more governed execution without creating more disconnected tooling.

Solution

- Built leadership nomination workflows with transparent criteria and audit-ready tracking.
- Digitized rotation assignments with structured routing, notifications, and program profiles.
- Automated Day-1 access provisioning and centralized security role administration.
- Used Workday Extend, AI Gateway, and AWS services to support classification, routing, and validation.

Benefits

- Accelerated user and access provisioning by 50%.
- Reduced document generation effort by 40%.
- Improved resource allocation speed by 30%.
- Strengthened compliance with embedded governance, audit logs, and controlled security models.

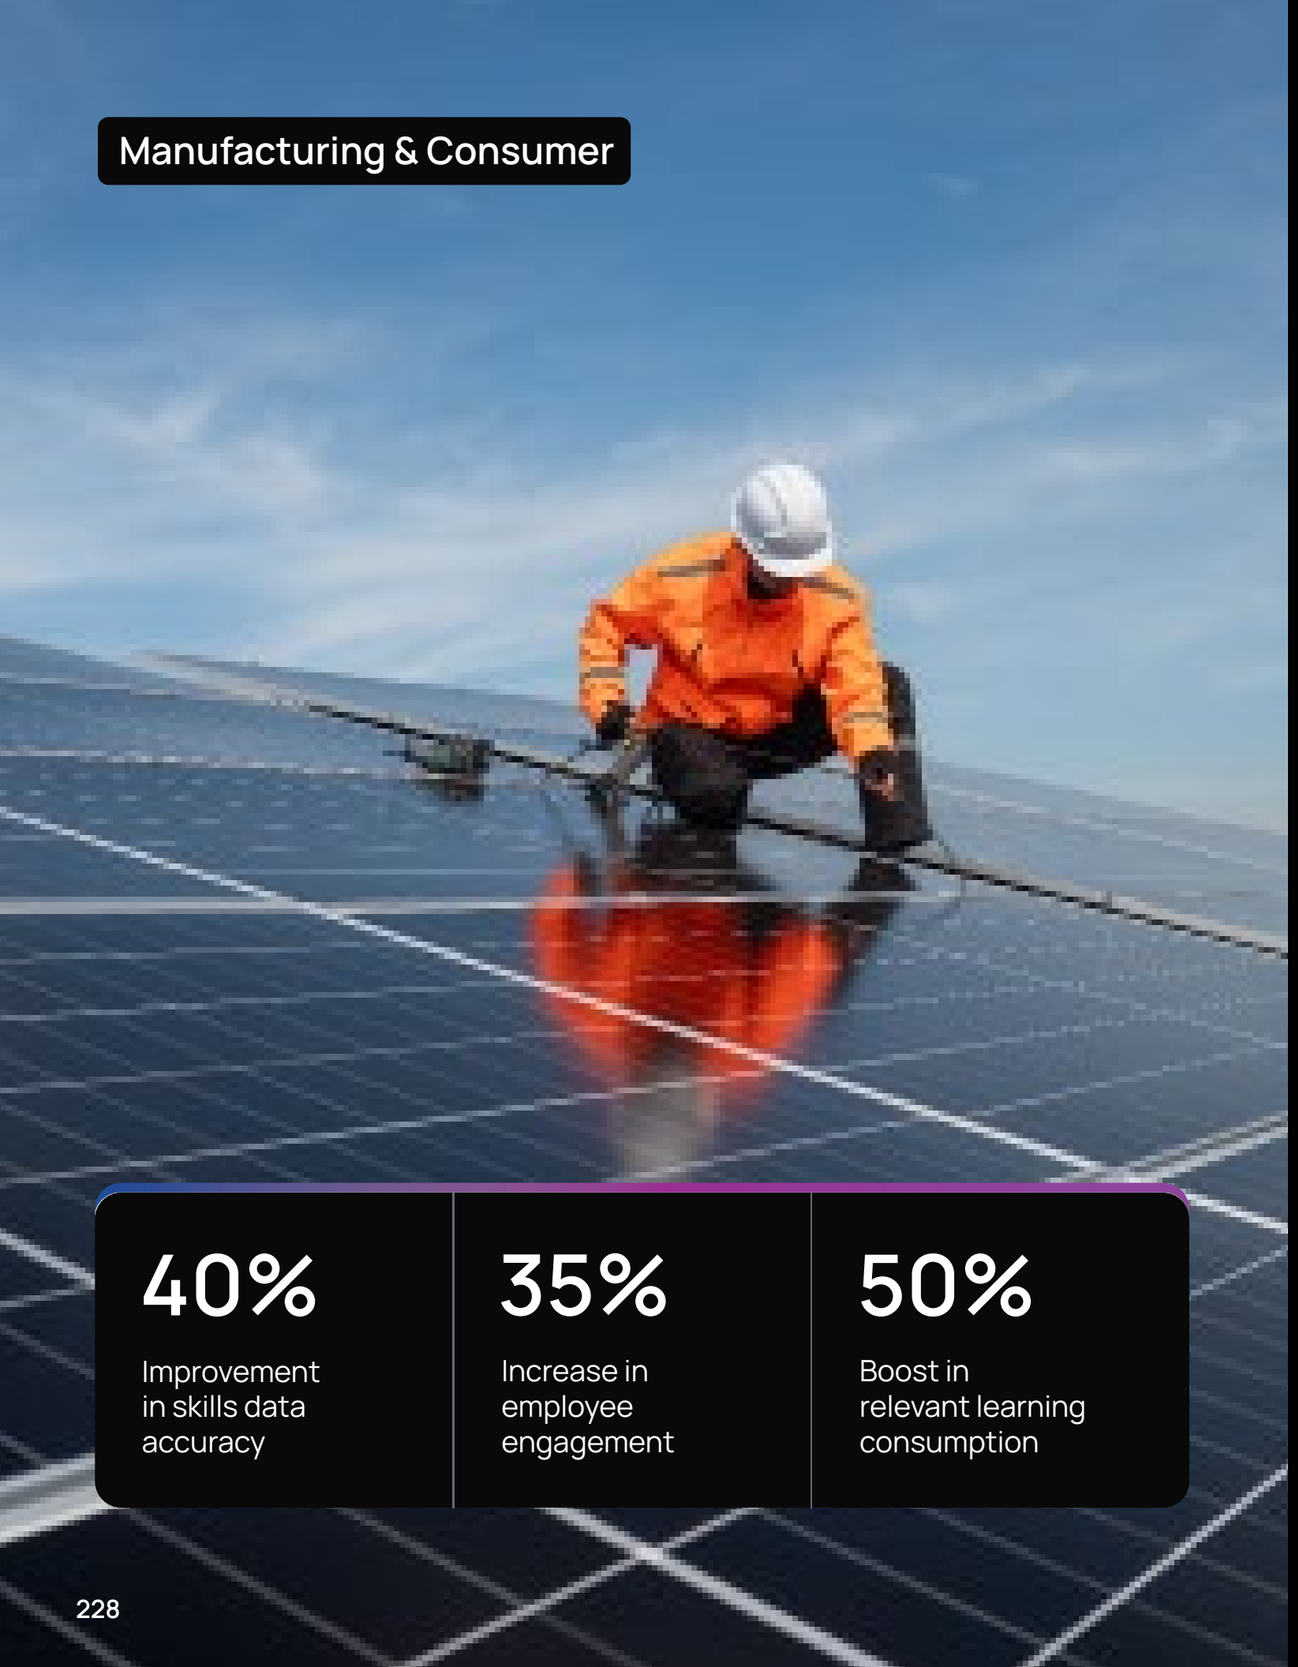

Client

A Global Solar Technology Leader Operating Across 13 countries

Building a skills engine that helps people grow and leaders plan better

Fragmented HR systems and inconsistent skills data were limiting workforce visibility and slowing decisions. The client needed a more connected skills ecosystem that would support employee development and business planning at the same time.

Solution

- Created an AI-curated skills architecture to standardize and enrich skills data globally.
- Built a unified skills center with real-time visibility into strengths, gaps, and development needs.
- Enabled personalized, skills-based career paths and proficiency recommendations.
- Connected learning, performance, and job profiles into a closed-loop system.

Benefits

- Improved skills data accuracy by 40%.
- Boosted relevant learning consumption by 50%.
- Increased employee engagement in development activity by 35%.
- Accelerated internal mobility decisions by 30–40% while reducing manual HR effort by 25%.

Client

A Leading Provider of Turnkey Infrastructure Solutions for Electric, Gas, and Telecom Networks

Sending technicians out better prepared the first time

Hazardous environments, incomplete work packages, and repeat visits were slowing resolution and raising cost. The client needed a safer, more reliable way to prepare technicians before dispatch.

Solution

- Implemented an agentic retrieval augmented generation (RAG) system to orchestrate hazard guidance, compliance steps, and parts planning.
- Connected hybrid retrieval with a semantic knowledge layer across assets, hazards, and historical resolutions.
- Converted technician know-how into reusable institutional knowledge.
- Generated task-ready work packages and linked them to inventory intelligence and governed workflows.

Benefits

- Improved issue resolution time by 12%.
- Raised technician readiness with more complete hazard controls, compliance steps, and parts requirements.
- Increased first-time resolution probability by reducing incomplete dispatches and repeat visits.
- Improved safety and lowered cost-to-serve through better field preparation.

Client

A Global Leader in Semiconductor Manufacturing

Unlocking 45 full-time employees (FTEs) with AI-powered manufacturing

Manual triaging caused 58% downtime in three key steps, leading to $7.2M annual revenue loss and consuming engineering capacity.

Solution

- Designed and implemented an AI-powered command center.
- Deployed edge AI agents for sub-second issue resolution without cloud delays.
- Introduced a business value framework for scalable, multi-site AI adoption.
- Integrated manufacturing execution system (MES), enterprise resource planning (ERP), supervisory control and data acquisition (SCADA), and information Technology Service Management (ITSM) systems to deliver 10x ticket throughput.

Benefits

- **Mean time to resolution (MTTR) reduction:** From 4–6 hours to under 20 minutes.
- **45 FTEs unlocked:** Redeployed engineers saved 35% of floor capacity.
- **10:1–30:1 ROI**: Achieved within 12–18 months, with sustained operational gains.

Zero
Spend on annual maintenance

67%
Improvement in operational efficiency

3X
Improvement in system performance

Client

A Global Mass Data Storage Innovator

Accelerating digital transformation with zero license modernization

The client faced high Oracle license costs (>$4M annually), vendor lock-in, over-provisioning, and underutilized features across 10,000 virtual machines (VMs), hindering modernization and redirecting resources from strategic investments.

Solution

Hexaware implemented the Amaze® for Modernization platform to:

- **Accelerate transformation:** Reduced a 3-year timeline to 1 year with 67% faster delivery.
- **Automate migration:** Achieved 95% Oracle-to-PostgreSQL automation, cutting manual effort by 75%.
- **Optimize performance:** Re-architected 5000+ SQL queries for read/write segregation, ensuring 2x-3x better runtime performance.
- **Eliminate license cost:** Transitioned to a Zero-License model, removing recurring Oracle spend.
- **Minimize risk:** Delivered seamless transformation with controlled workload redirection and rollback planning.

Benefits

- **Cost Savings:** Achieved $0 annual maintenance spend, redirecting funds to strategic initiatives.
- **Operational Efficiency:** Compressed timeline by 67%, delivering modernization in 1 year instead of 3.
- **Performance Gains:** Improved system performance by 2x-3x, ensuring scalability and sustainability.

Digital IT Operations

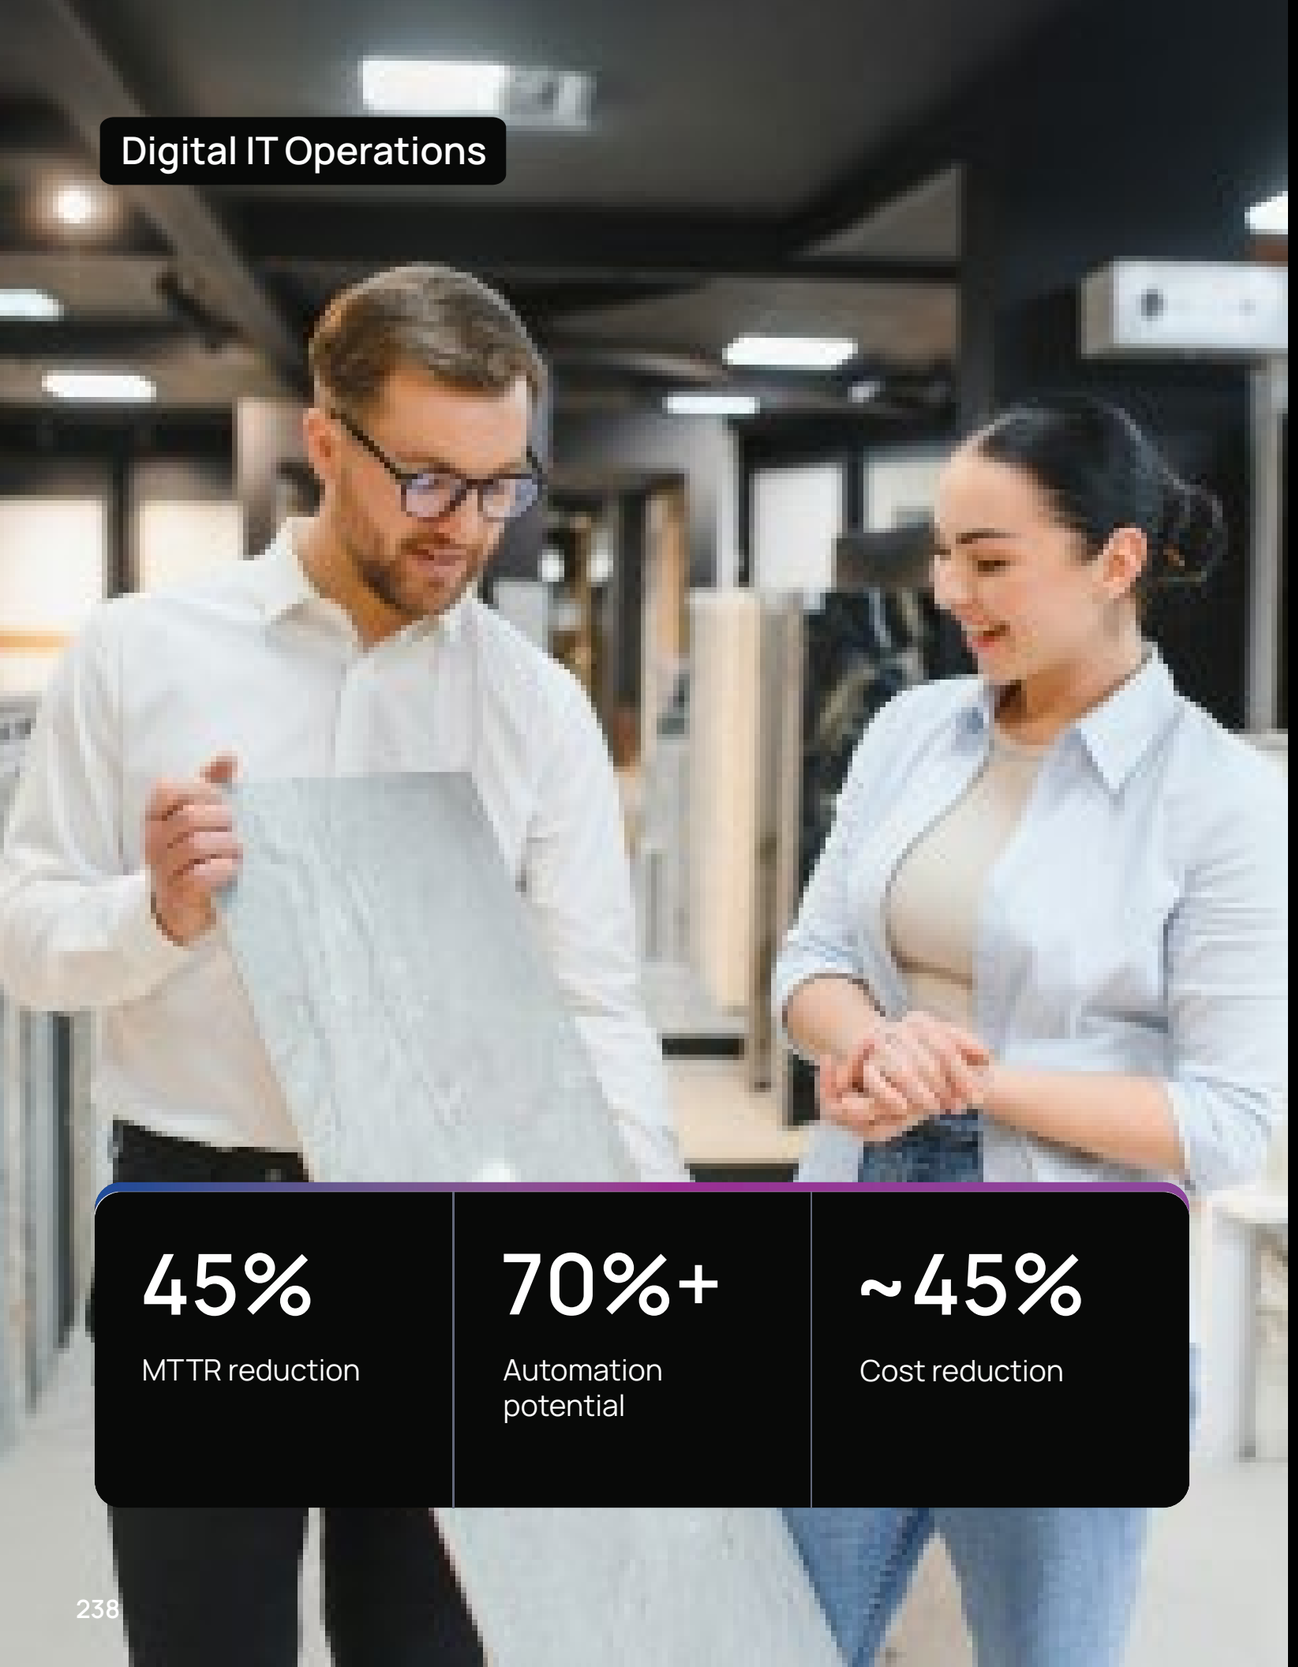

Client

A British Home Improvement Retailer

Closing the automation gap with AI that can reason across operations

Rigid scripts and siloed tools had left too much operational work outside the reach of automation. The client needed a more adaptive AIOps model that could reason across towers, reduce manual load, and keep improving over time.

Solution

- Unified operational knowledge through ontology-driven data consolidation.
- Deployed persona-aligned AI agents tailored to different operational roles.
- Used the DEVAL Engine for continuous learning and self-improvement.
- Added telemetry-aware triage and self-healing capabilities to move from assisted to more autonomous operations.

Benefits

- Reduced mean time to repair (MTTR) by about 30–45%.
- Lowered costs by about 35–45%.
- Increased automation potential to 70%+ from roughly 35%.
- Improved engineer productivity by 30–50% while reducing routine ticket demand by 40–60%.

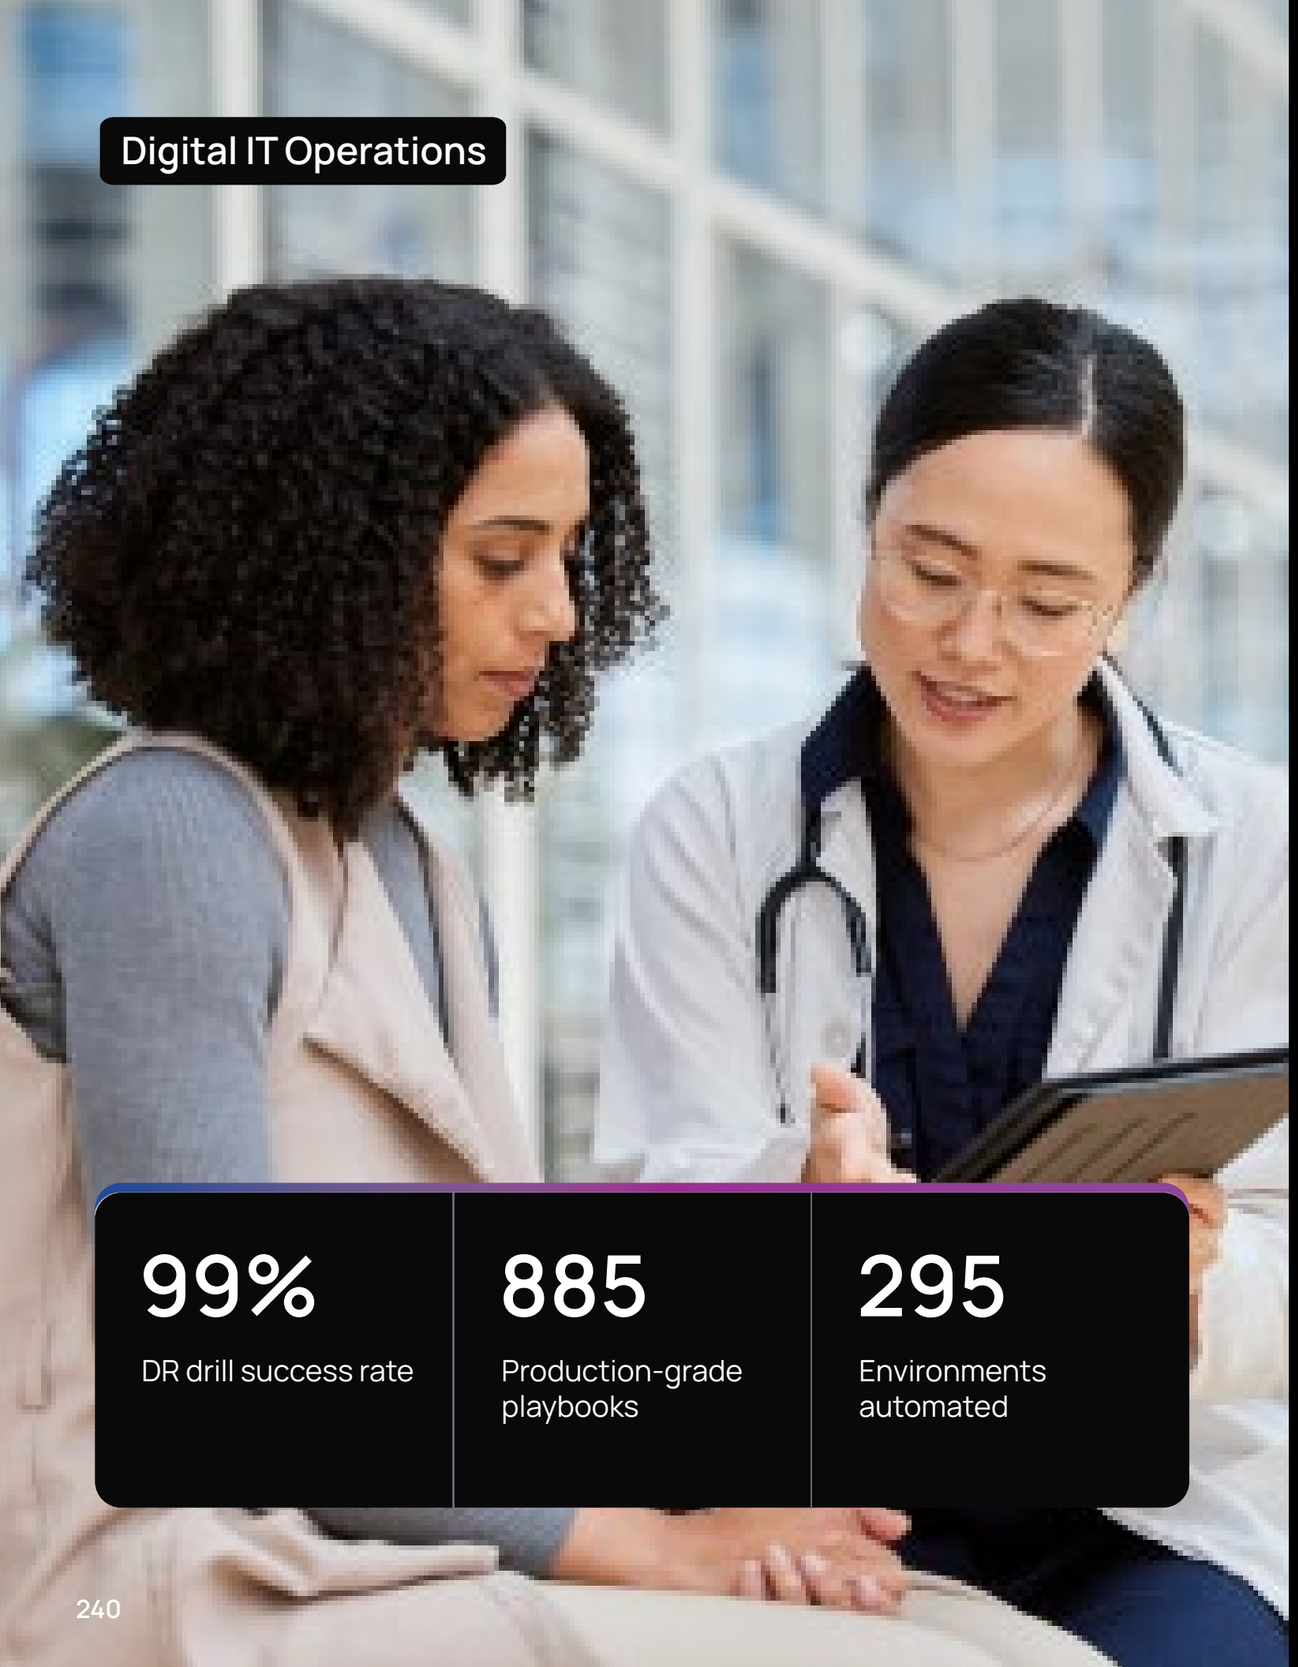

Client

A Regional US-Based Health-Insurance Provider Serving Employers and Individuals

Replacing fragile disaster recovery (DR) tooling with automation, the business can trust

The client needed to move away from a costly, brittle disaster recovery platform without introducing new operational risk. The objective was straightforward: faster, more reliable recovery and fewer licensing burdens.

Solution

- Re-engineered IBM RO XML workflows into a new Ansible-based disaster recovery platform.
- Automated recovery across MSSQL, IBM MQ, IBM ACE/IIB, and Windows/Linux environments.
- Used LLMs to accelerate XML translation, code generation, documentation, and reverse engineering.
- Built resilient, self-healing playbooks for consistent DR execution at scale.

Benefits

- Achieved a 99% DR drill success rate.
- Delivered the fastest DR drill in the client's history at 4 hours 15 minutes.
- Generated 885 production-grade playbooks across 295 enterprise environments.
- Removed recurring IBM RO license dependency with a one-time build that improved long-term economics.

Client

Digital & IT Service Provider of a Leading European Railway Organization

Turning a fragmented SharePoint estate into an AI-powered knowledge platform with natural language search

With over 100,000 documents spread across disconnected SharePoint portals, employees spent excessive time manually navigating repositories with no natural language search capability. Manual classification created compliance risks, inconsistent metadata, and audit gaps, while IT teams were overwhelmed with access requests and governance issues.

Solution

- Deployed an AI-powered SharePoint Orchestrator built on RAG, deeply integrated with Microsoft 365, Graph API, and Purview.
- Natural language chat interface embedded in SharePoint, replacing manual folder navigation entirely.
- Context-aware responses with source citations, ranking, and permission-filtered results via Microsoft Graph.
- Automated document classification, metadata tagging, and policy enforcement using Microsoft Purview.

Benefits

- Instant search over 100K+ documents with no performance degradation.
- Fewer IT support tickets through self-service, permission-aware document access.
- Audit-ready traceability on every AI-generated response via source citations.
- Strengthened compliance posture through automated policy enforcement and tagging.

HEXAWARE

Scan to access the digital edition of this series.

About Hexaware

Hexaware is a global technology and business process services company. Every day, Hexawarians wake up with a singular purpose: to create smiles through great people and technology. With this purpose gaining momentum, we are well on our way to realizing our vision of being the most loved digital transformation partner in the world. We also seek to protect the planet and build a better tomorrow for our customers, employees, partners, investors, and the communities in which we operate.

With offices across the world, we empower enterprises to realize digital transformation at scale and speed by partnering with them to build, transform, run, and optimize their technology and business processes.

Learn more about Hexaware at www.hexaware.com.