

- # AI Expertise & Services

Enabling innovation with GenAI, ML, and AI-powered operations.

AI Expertise at Techbar

10+

Years of experience in AI, ML & data-driven solutions

15+

AI & GenAI projects across industries

1M+

Of end users reached through AI-powered platforms

200+

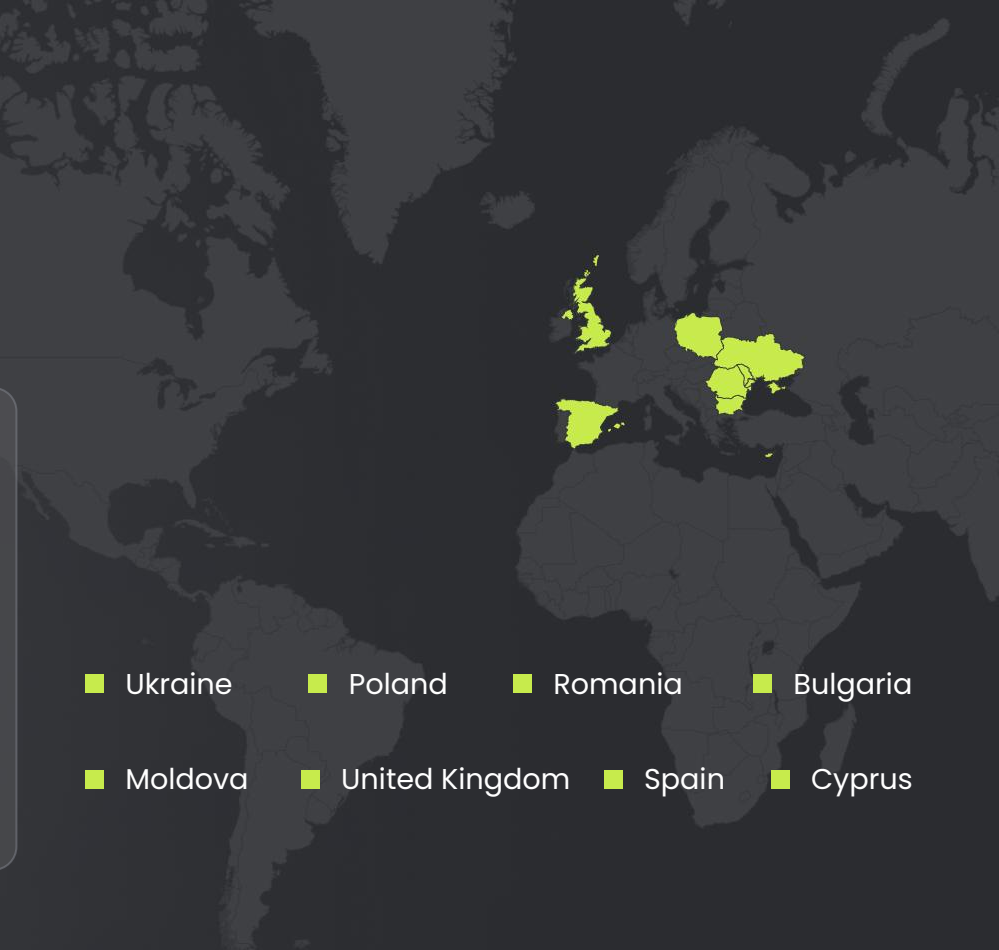
Available AI/ML experts, data scientists, and engineers

5+

Active AI clients leveraging GenAI and ML solutions

80%

Senior engineers with domain-specific and AI/ML expertise

- 
- Ukraine
 - Poland
 - Romania
 - Bulgaria
 - Moldova
 - United Kingdom
 - Spain
 - Cyprus

Unlock the value of AI with Techbar

Address key challenges to enable scalable, secure, and high-impact AI adoption.



Data Quality & Silos
Unstructured, fragmented, and biased data reduce model accuracy.



High Infra Costs
Training and running AI/GenAI models requires expensive compute and storage.



Limited Business Alignment
AI pilots often fail to scale due to unclear ROI and lack of integration.



Trust & Compliance Risks
Bias, hallucinations, and regulatory concerns block enterprise adoption.



Operational Complexity
Models degrade without MLOps/LLMOPs, monitoring, and governance.

Key Challenges

How we can help

We enable our clients to build and scale AI solutions with robust data foundations, GenAI and ML expertise, and enterprise-grade platforms and operations.

Benefits

Reliable AI Models

Accurate, unbiased, explainable predictions and generations

Cost-Optimized Infrastructure

Efficient cloud and hybrid deployments tailored to workloads

Business-Driven Use Cases

Solutions aligned with measurable ROI and industry needs

Responsible AI

Governance, security, and compliance by design

Faster Innovation

Proven expertise and streamlined delivery to accelerate AI adoption

Our AI Expertise Areas in Detail

We support our clients with scalable, secure, and business-driven AI solutions.



Generative AI

- Copilots and chatbots for employees and customers
- Automated generation of text, images, and code
- Agentic AI for multi-step task automation and external system orchestration
- RAG-based integration with enterprise data
- Multimodal experiences with text, voice, and video



ML & Data Science

- Demand, risk, and failure prediction
- Recommendation systems to drive engagement
- Computer vision for image and video analytics
- NLP for search, classification, and insights
- Fraud detection and anomaly monitoring



AI Platforms & Ops

- MLOps/LLMOps for full model lifecycle management
- Monitoring of AI quality and performance
- Security, compliance, and governance by design
- Cost-optimized cloud and on-prem deployments
- Scalable production-grade AI infrastructure


Techbar Across AI Value Chain



AI Foundations
(Data, infrastructure, governance)



AI Models
(ML, GenAI, multimodal)



AI in Action
(Enterprise apps, copilots, domain solutions)

Data & Capture

- Data collection (text, images, audio, video, logs)
- Data labeling & curation
- Synthetic data generation

Ingest & Prepare

- Data pipelines & integration
- Cleaning, normalization, enrichment
- Metadata & embeddings

Model Development

- Classical ML (prediction, classification, CV, NLP)
- GenAI (LLMs, multimodal, RAG)
- Training, fine-tuning, evaluation

Deploy & Operate

- MLOps / LLMOps
- Monitoring & observability
- Security, governance, compliance

Scale & Monetize

- Embedding AI into business processes
- AI-driven personalization & automation
- ROI tracking & optimization

Deliver & Engage

- AI copilots & chatbots
- Agentic AI for autonomous workflows
- Domain-specific AI apps
- AI-powered decision support

Responsible & Trusted AI

What makes our team stronger

Our certifications validate hands-on expertise in **cloud infrastructure and deployment** — driving **real-world, enterprise-grade solutions**.

Cloud



AWS Certified
Solutions Architect
Professional



Google Cloud
Certified —
Professional Cloud
Architect



Azure Solutions
Architect Expert

DevOps & Infra



AWS Certified
DevOps Engineer —
Professional



Certified Kubernetes
Administrator



HashiCorp
Terraform Associate
Certification

ML/AI



Google Cloud
Professional
Machine Learning
Engineer



AWS Certified
Machine Learning —
Specialty



Microsoft Certified:
Azure AI Engineer
Associate

Data



Databricks Data
Engineer
Professional

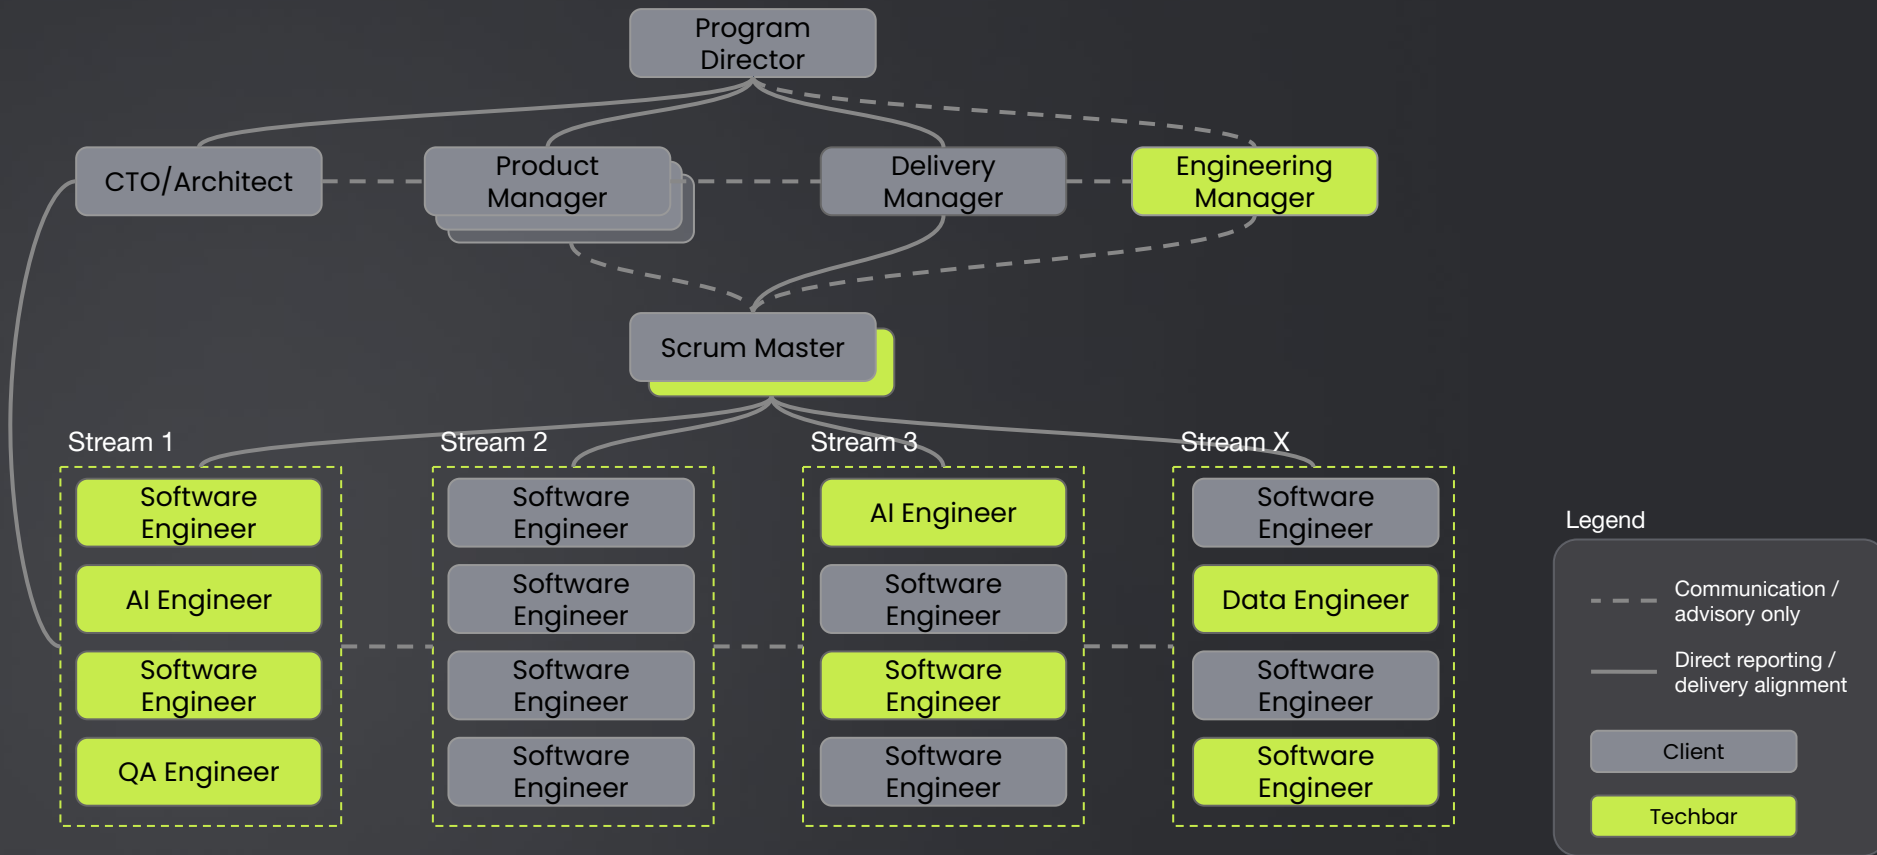


Google Data
Engineer
Professional



Azure Data Engineer
Associate

How we integrate with your team



Why Techbar Engineering Teams?



Techbar provides senior-led engineering squads that embed directly into the product or platform team, working as an extension of your internal capability.

Unlike traditional outsourcing, we focus on product thinking, and team continuity.



Engagement models

One team. Multiple ways to work together. Whether you need extra hands or full ownership, we adapt to your workflow, not the other way around.

Model	Use case	Commitment
Staff Augmentation	You need extra hands quickly	Flexible
Dedicated Teams	You need a team to build	Sustained
End-to-End Delivery	We fully own the process	Structured

Outcome-driven delivery model

How we work

- Agile from day one
- Transparent planning & reporting
- Shared ownership across teams

1

Discovery

We align on goals, scope, and team structure

2

Team Setup

We onboard the right talent – fast and with purpose

3

Delivery

We build iteratively, focusing on speed and real value

4

Support

We stay with you to help you scale and improve

GenAI Copilot for Customer Support

Opportunity

A mid-sized SaaS company providing HR & payroll software for SMEs faced increasing pressure on its support team. With 20k+ end users across multiple countries, agents were overloaded with repetitive questions (password resets, tax form issues, onboarding help). Knowledge was fragmented across Zendesk Help Center, Confluence, and Google Drive, which slowed responses and hurt customer satisfaction.

Value created

- Delivered an internal GenAI copilot integrated directly into Zendesk.
- Connected to Confluence and Google Drive via RAG to surface the latest policies and guides.
- Automated ticket summarization and categorization for faster triage.
- Introduced tone control and guardrails for safe, compliant responses.

Impact

- Support agents handled tickets ~30% faster thanks to suggested replies.
- Small pilot group quickly adopted the tool and rated it highly useful in daily work.
- Customer satisfaction (CSAT) improved as responses became more consistent.
- Lower workload on support team freed time for handling complex issues.



GenAI Copilot for Customer Support



Technical Implementation

- GenAI copilot powered by Azure OpenAI (GPT-4).
- LangGraph orchestration with RAG pipeline.
- Pinecone vector database for embeddings.
- Integrated with Zendesk API (ticket workflow), Confluence API, and Google Drive API.
- Prompt templates with built-in tone and safety filters.
- Observability with Prometheus + Grafana dashboards.



Engineering Management

- Delivered pilot in 3 months, production rollout in 5 months.
- Agile delivery with 2-week sprints and close client collaboration.
- Automated regression testing of prompts and RAG pipeline.
- Regular demos and feedback sessions with support team.



Technologies Used

Azure OpenAI (GPT-4) | LangGraph | Pinecone | FastAPI | Zendesk API |
Confluence API | Google Drive API | Azure Kubernetes Service (AKS) |
Azure Key Vault



ML-Driven Demand Forecasting for a Retailer

Opportunity

A US-based retail chain with 50+ physical stores and an e-commerce channel relied on manual forecasting. Managers exported sales and stock data from SAP ERP into Excel, applied simple formulas or last year's figures, and adjusted based on intuition. Promotions were considered inconsistently. As a result, the retailer faced overstock on some SKUs, stockouts on others, and lost revenue.

Value created

- Built an ML-driven demand forecasting system at SKU and store level.
- Integrated sales, inventory, promotions, and e-commerce data.
- Delivered forecasts into Power BI dashboards for planners.
- Provided an API integration with ERP for automated replenishment.

Impact

- Forecast accuracy improved by 20% vs. Excel-based methods.
- Excess inventory reduced, cutting storage costs.
- Stockouts decreased, driving uplift in sales of top products.
- Clear, actionable insights for planners without technical skills.



ML-Driven Demand Forecasting for a Retailer



Technical Implementation

- SAP Business One (ERP), Shopify (e-commerce), internal promotions system as data sources.
- Existing PostgreSQL DWH leveraged for storage.
- Pipelines built with Azure Data Factory + Airflow for ETL.
- LightGBM, Prophet, combined in Azure ML, retrained weekly.
- Forecasts exposed via REST API and Power BI dashboards.



Engineering Management

- Full rollout in 6 months across all stores and e-commerce.
- Collaboration with supply chain and category teams for validation.
- Continuous improvement with monthly accuracy reviews.



Technologies Used

SAP Business One | Shopify API | Azure Data Factory | Apache Airflow |
Azure ML | LightGBM, Prophet | PostgreSQL | Power BI



LLMOps Platform for Financial Services

Opportunity

A fintech payments company was experimenting with GenAI pilots (internal assistants, customer-facing chatbots).

While initial results looked promising, scaling faced roadblocks:

- API usage costs were rising unpredictably.
- No visibility into model performance or quality of responses.
- Security and compliance risks from sensitive financial data.
- Lack of governance over who used which models and how.

Value created

- Designed and deployed an LLMOps platform to standardize model usage.
- Integrated monitoring for latency, cost, and response quality.
- Implemented access control and role-based access control tied to corporate identity (Azure AD).
- Set up centralized logging, auditing, and compliance dashboards.

Impact

- 25% reduction in GenAI infrastructure/API costs through caching and routing rules.
- Improved governance so only authorized users and apps could access models.
- Transparency for leadership via dashboards (usage, cost, quality).
- Foundation for safe scaling of GenAI use cases across multiple business units.



LLMOps Platform for Financial Services



Technical Implementation

- Connected to Azure OpenAI Service (GPT-4).
- Orchestrated with LangGraph + FastAPI service layer.
- Prometheus + Grafana for latency, error rate, token usage.
- Cost tracking with Azure Monitor and internal dashboards.
- Access via Azure AD SSO, role-based access control, and API gateway.
- Logging & compliance with ELK stack (Elasticsearch, Logstash, Kibana).



Engineering Management

- The platform deployed in 6 months with 5+ use cases.
- Agile delivery with 2-week sprints and regular compliance reviews.
- Close collaboration with IT security and risk teams.



Technologies Used

Azure OpenAI | LangGraph | FastAPI | Prometheus | Grafana | Azure Monitor | Azure AD | API Gateway | ELK Stack



Personalized Music Discovery for a Niche Streaming Service

Opportunity

A niche music streaming service offered a flat catalog with no personalization. All users saw the same playlists and recommendations, which led to low engagement and poor user retention. The client needed a lightweight recommendation engine to increase listening time, repeat sessions, and create more relevant user experiences without building a heavy data science infrastructure.

Value created

- Designed and implemented a recommendation service using collaborative filtering and content-based features (tags, genres, artist similarity).
- Added A/B testing framework to experiment with personalized playlists and new music discovery flows.
- Integrated listening analytics (session length, skips, favorites) into the recommendation pipeline.
- Delivered a web dashboard for the editorial team to monitor playlist performance and adjust recommendation rules.

Impact

- +18% increase in average listening time per user.
- +12% growth in repeat sessions within 3 months.
- Higher satisfaction reported in user surveys, especially among new listeners.



Personalized Music Discovery for a Niche Streaming Service



Technical Implementation

- Built a microservice in Python with FastAPI for recommendations.
- Collaborative filtering model trained on anonymized user-listening data.
- Content-based enrichment from track metadata (tags, genres, artist similarity).
- Event tracking pipeline (listening events collected via Kafka - stored in PostgreSQL + Redis for fast lookups).



Engineering Management

- Discovery phase to align business KPIs (listening time, retention) with recommendation goals.
- Delivered in phases (baseline recommendations - A/B testing - editorial dashboard).
- Agile sprints with close coordination between developers and client's product team.



Technologies Used

Python (FastAPI) | Collaborative Filtering (ALS) | Embeddings | PostgreSQL (pgvector) | Redis | Kafka | React | AWS (S3, Lambda, ECS)



AI-Powered ID Document Recognition (KYC)

Opportunity


A European fintech offering digital wallets and SME payments faced bottlenecks in customer onboarding. Manual KYC verification (passport, driver's license, ID cards) was slow, error-prone, and required significant back-office resources. This increased onboarding costs and frustrated new customers.

Value created

- Implemented an AI-based KYC document recognition system.
- Used OCR to extract key fields (name, DOB, ID number, expiration date).
- Applied CNN-based classification to detect document type and flag suspicious formats.
- Integrated with existing onboarding workflow via API.

Impact

- Onboarding time reduced from 24h to 30 minutes for standard cases.
- Manual review load cut significantly.
- Improved fraud detection with automatic checks for document validity.
- Smoother customer experience, higher conversion at signup.



AI-Powered ID Document Recognition (KYC)



Technical Implementation

- OCR with Azure AI document intelligence for text extraction.
- CNN model (ResNet-based) trained on labeled ID/passport datasets for classification.
- Data pipelines built with Python + FastAPI, connected to onboarding service.
- Document metadata stored in PostgreSQL with audit logs.
- Alerts and suspicious cases routed to compliance officers for manual check.



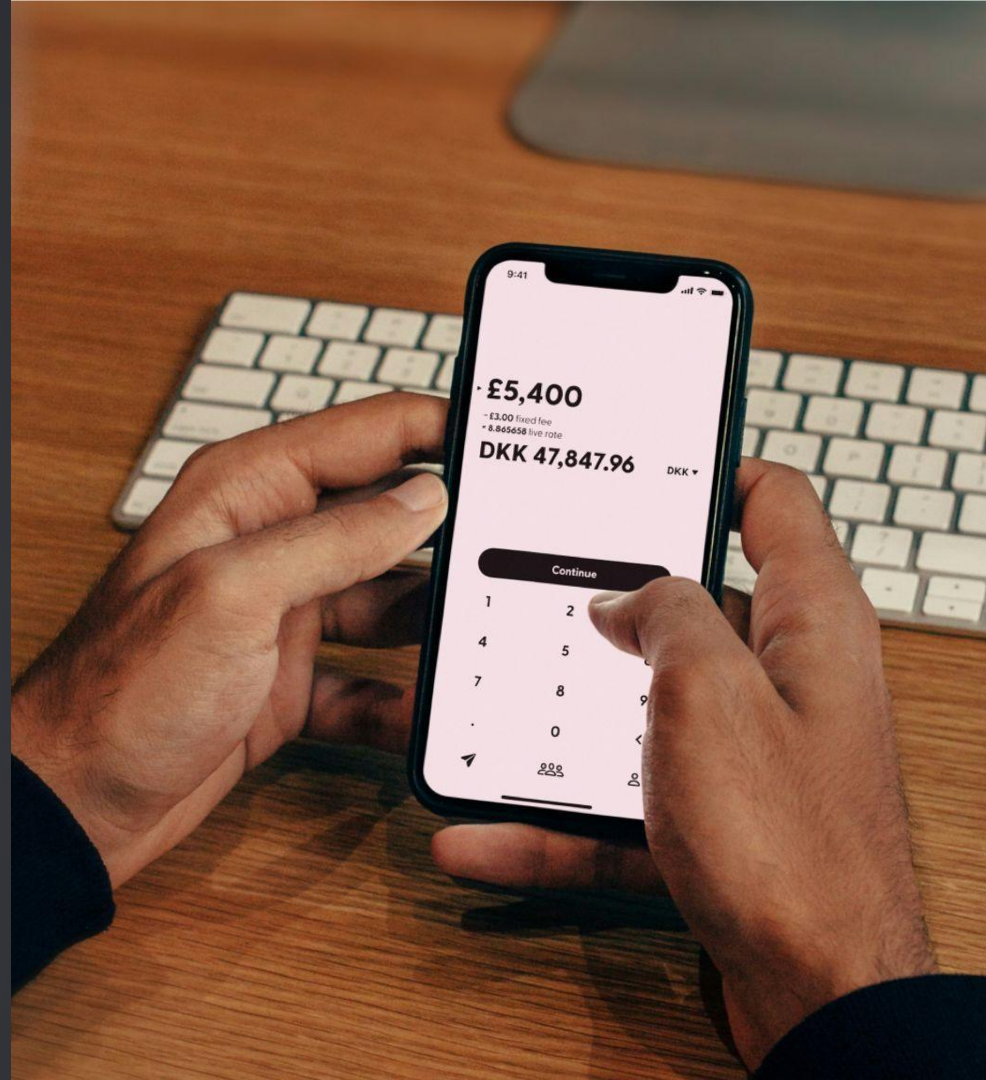
Engineering Management

- Pilot delivered in 3 months (passports and driver's licenses).
- Extended to multiple EU ID formats in under 6 months.
- Agile delivery with close collaboration with compliance and risk teams.
- Regular accuracy benchmarks and retraining with new samples.



Technologies Used

Azure Computer Vision | PyTorch (ResNet) | FastAPI | PostgreSQL | Docker | Power BI



- Case Study: Property Management Platform Modernization

Property Management Platform Modernization

Opportunity

A U.S.-based PropTech company provided a SaaS platform used by residential property management firms to manage communities, maintenance requests, and resident communication.

The platform had grown outdated (built on an early .NET Framework monolith with limited scalability).

Administrators used the system to manage service tickets, maintenance schedules, waste collection, and supplier tasks, while residents accessed a web portal for submitting requests.

Value created

Working jointly with the client's in-house product and architecture teams, Techbar helped rebuild the platform's core modules (Service Requests, Maintenance Workflows, Supplier Management, and Resident Portal) on a modern, scalable foundation.

We migrated the backend from legacy .NET Framework to .NET 7 microservices, redesigned the data model, and implemented a React-based responsive web UI for both property managers and residents.

Impact

- Unified multi-tenant SaaS architecture supporting 50+ property management firms.
- Improved system performance and response times.
- Reduced onboarding time for new residential communities from weeks to days.
- Simplified configuration for property-specific workflows (cleaning, maintenance, waste collection, repairs).



Property Management Platform Modernization



Technical Implementation

- Migrated backend from .NET Framework 4.x monolith to a .NET 7 microservice-based architecture hosted in Azure.
- Implemented API Gateway with authentication and tenant routing for multi-community access.
- Developed React + TypeScript front-end to replace legacy ASP.NET views; unified UI/UX for resident and admin roles.
- Rebuilt data model in PostgreSQL, introduced caching (Redis) and async job processing (RabbitMQ).



Engineering Management

- Joint delivery model with client's architect and product owners managing roadmap priorities.
- Techbar led software development, DevOps, and QA streams.
- Two scrum teams working in 2-week sprints with shared backlog and planning.
- Frequent demos and stakeholder reviews to validate UX and migration progress.



Technologies Used

.NET 7 | React | TypeScript | PostgreSQL | RabbitMQ | Redis | Azure
DevOps | Docker | Kubernetes | API Gateway



techbar
Software Engineering

“ We’re not just building a delivery pipeline, we build engineering partnerships.

Our goal is to be the team you trust with both your code and your roadmap.

- Long-term mindset
- Product thinking
- Real ownership



Ready to build?

Let's connect and explore how we can help.

No back-and-forth delays — just direct, hands-on collaboration.

Email

maria.tolkacheva@techbarsw.com

VP of Partnerships

Maria Tolkacheva

