



# Tech Product AI

How fast-moving B2B enterprises are driving AI transformation

# 1. Executive Summary

## AI ROI for B2B Tech

B2B technology enterprises face both opportunity and risk in AI adoption. McKinsey projects \$2.6–\$4.4 trillion in annual productivity gains from generative AI, yet value only materializes when AI is embedded into core workflows rather than layered on as disconnected tools.

This paper outlines how Tech Product AI delivers measurable impact across three critical domains:

### Sales AI

Equips every rep with real-time technical expertise, competitive positioning, and customer intelligence. The result: faster cycles, higher win rates, and millions in incremental market cap.

### Support AI

Transforms technical support into a revenue driver, reducing service costs, expanding coverage to long-tail customers, and accelerating time-to-revenue with faster, more accurate answers.

### Team AI

Breaks down silos of tribal knowledge, accelerating onboarding, product content creation, and go-to-market execution.

From our work with global enterprises, five requirements consistently predict success:

1. **Deep Product Knowledge** – AI that understands your products, workflows, and stakeholders.
2. **Cross-Use Case Leverage** – One platform powering sales, support, and teams.
3. **Workflow Integration** – AI embedded in CRM, ticketing, and collaboration tools.
4. **Visibility & Control** – Real-time performance monitoring, transparency, and governance.
5. **Enterprise-Grade Security** – SOC 2, GDPR, RBAC, tenant isolation, and flexible deployment options.

The path to value is iterative: define measurable goals, retain control of proprietary data, and deploy AI where teams already work. Companies that unify AI across functions—rather than fragmenting it into isolated pilots—scale faster, reduce costs, and drive competitive advantage.

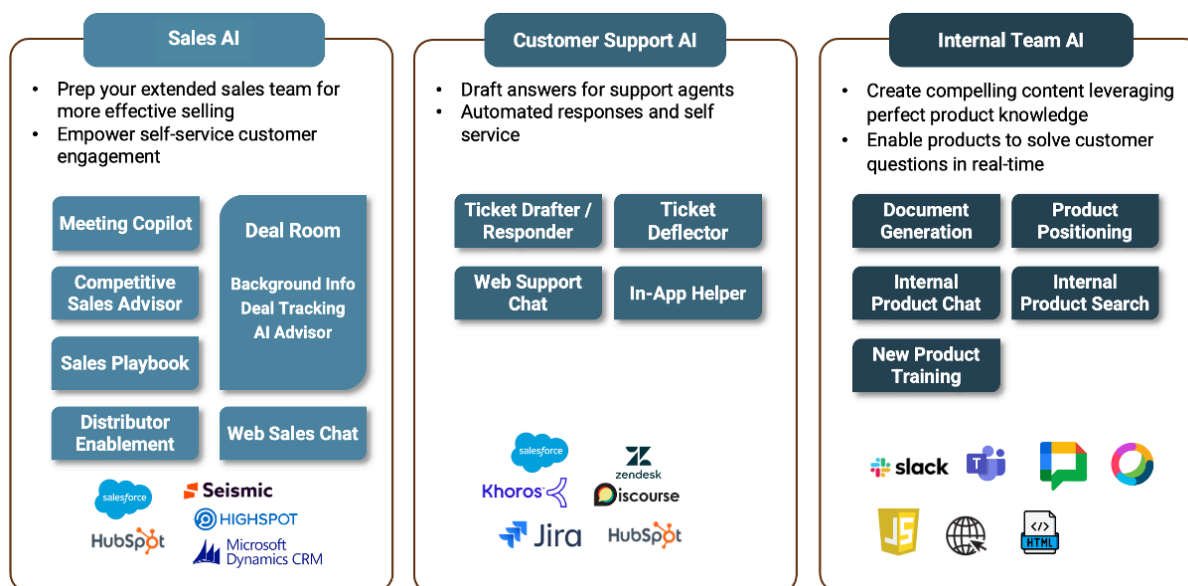
AI is not about replacing tools or people — it's about making teams more effective. An example is provided as a case study in this whitepaper. B2B tech enterprises that move quickly, safely, and with the right platform strategy will capture outsized growth.

## 2. AI Driving ROI for B2B Tech

B2B tech companies have a lot to gain from AI and it is one of the top items in every boardroom. Companies are recognizing that AI can help accelerate their business or it can leave them behind if others adopt faster. The challenge for large enterprises is how to do this safely with the teams they have and with the people they have so they don't get eclipsed by competitors or new market entrants.

McKinsey estimates generative AI could add \$2.6 trillion to \$4.4 trillion annually in global productivity gains. But unlocking that value requires embedding AI into core workflows – not bolting on disconnected tools.

While AI will help many different aspects of business – including product development, product design, finance legal HR and operations – this white paper talks about how AI can help B2B tech companies improve their customer communications and internal product understanding with use cases across sales support, and teams.



*Tech Product AI Use Cases*

### Sales AI: Accelerating Growth, Capturing Market Share

In industries like semiconductors, networking, and industrial systems, a single design win can be worth millions – yet sales cycles stretch 6–12 months. Salespeople are

given the challenge of understanding a broad range of products, understanding the customer application, understanding the customer's situation, AND understanding the competition. For a startup salesperson that may be viable, but for a B2B tech enterprise salesperson, that is difficult or impossible. They lean on people across the organization to help.

Sales AI transforms this by:

- Giving real-time access to product technical expertise. Like an Applications Engineer in every salesperson's pocket
- Providing real-time, deal-specific competitive sales guides to position your products to win
- Providing deep customer background so salespeople understand the situation, the customer, and the individual stakeholders

Together this together boosts win rates, accelerates sales cycles, drives cross selling, and ultimately drives compelling ROI. One global semiconductor customer estimated that optimizing cross-sell motions with Sales AI could generate \$1 million per day in incremental market cap.

## **Support AI: Reducing Cost, Driving Customer Delight**

In some industries technical support is seen as a cost center, but in B2B tech companies it is often a value driver. Pre-sales support is critical to winning a deal, and post-sales technical support is critical to helping the customer succeed (ultimately driving revenue).

AI Support helps to achieve all of the goals of today's support leaders:

- Accelerating deals with faster, more accurate answers for customers
- Supporting long-tail customers that have not been economical to serve before
- Reducing cost of service delivery

In industries where post-sales experience determines time to revenue and repeat business, Support AI isn't optional. It's a competitive edge.

## Team AI: Accelerating Innovation, Amplifying Product Impact

Behind every product launch, spec update, or market push is a maze of tribal knowledge, documentation, and internal collaboration. That's where Team AI comes in:

- Accelerates internal ramp and training—cutting time to proficiency for engineers, PMMs, and channel partners
- Improves product content velocity, helping to generate datasheets, positioning briefs, and battlecards
- Surfaces the best answers across teams, enabling product marketing, support ops, and engineering to work from a unified knowledge layer

This isn't just internal productivity — it's faster product-market fit, cleaner launches, and smarter go-to-market execution.

### 3. Requirements: What Enterprises Need to Get Right

AI isn't just a new tool—it's a new way of working. But realizing its full potential depends on how you approach it. B2B enterprises need to move quickly, learn safely, and scale intelligently across functions—without compromising trust, governance, or performance.

The right approach balances speed and control. It enables teams to ship real use cases in days or weeks, then improve through iteration. It prioritizes workflow fit, shared learning across departments, and data ownership, while avoiding the costs and risks of retraining models or managing fragmented tools.

Based on our work with multi-billion-dollar enterprises, successful AI deployments consistently address these five fundamental requirements:

#### 3.1 Deep Product Knowledge: AI That Speaks Your Language

Generic AI fails in product-centric organizations because it lacks the specificity and context needed to support complex workflows. Enterprises must prioritize systems that:

- Speak fluently about their products, SKUs, lifecycle, and support history
- Understand the difference between how a sales engineer and a support agent asks a question
- Adapt to existing tools—CRM, ticketing, team chat—without requiring teams to switch context

### Why this matters

Without deep product knowledge, AI simply becomes another unused tool. With it, AI becomes a seamless, expert assistant embedded in the business. Studies show that generic AI chatbots fail to resolve roughly **70% of technical inquiries** because they lack deep knowledge of the company's specific products.

### Real-world example

During a recent meeting with a major semiconductor manufacturer, their team highlighted the challenge of supporting a 'long tail' market through distributors. They have tens of thousands of product documents, community forums dating back decades, GitHub repositories, and complex technical specifications that require deep understanding. Generic AI tools cannot navigate this complexity effectively.

### What to look for

AI platforms that use retrieval-augmented generation (RAG) to access your product data securely at query time, avoiding the need to train models on your proprietary information.

## 3.2 Leveraged Across Use Cases: One Platform, Multiple Solutions

AI's greatest value comes when insights from one part of the business inform the others. Yet many organizations run isolated pilots—one in support, one in sales, another for internal training—with no shared infrastructure or learnings.

A unified approach allows teams to:

- Reuse knowledge sources across domains
- Share usage patterns, feedback, and improvements
- Improve performance systemwide with every interaction

### Why this matters

A shared foundation—spanning knowledge sources, security rules, and feedback loops—enables faster time-to-value and more consistent performance. Fragmentation slows everything down.

### Real-world example

The same semiconductor company identified three critical use cases: (1) internal chatbots for navigating complex applications and product portfolios, (2) technical support for long-tail customers who purchase through distribution, and (3) product discovery tools to help customers choose between similar products. A unified platform addresses all three without requiring separate implementations.

### What to look for

A support AI surfacing FAQs can help sales teams better understand customer pain points. Internal AI chatbots that learn from product docs and customer tickets can accelerate onboarding across teams.

## 3.3 Integrates to Existing Tools: Workflow-First, Not Tool-First

The most successful AI deployments fit seamlessly into existing workflows rather than forcing teams to adopt new tools. This means:

- Native connectors for CRM, ticketing, team chat, and knowledge systems
- APIs that plug into existing dashboards and reporting tools
- Single sign-on and role-based access that respects current security models

### Why this matters

Integration friction kills adoption. When AI requires teams to switch contexts or learn new interfaces, usage drops dramatically. The goal is to make AI feel like a natural extension of tools teams already use.

### Real-world example

The semiconductor company already has existing enterprise platforms for sales (Salesforce), support (Jira Service Management), and distributor platforms (homegrown web portal). They need solutions that integrate with their existing tools rather than replacing them.

### What to look for

Pre-built connectors for Salesforce, Zendesk, Slack, Teams, SharePoint, and other enterprise staples, plus robust APIs for custom integrations.

## 3.4 Visibility & Control: Trust Through Transparency

No AI deployment will reach scale if teams can't see what's happening and control the outcomes. This requires:

- Real-time dashboards showing answer quality, usage patterns, and performance metrics
- AI Performance Management tools that grade every response and flag potential issues
- Configurable confidence thresholds and human-in-the-loop workflows
- Audit trails for compliance and governance requirements

### Why this matters

Without visibility, teams can't trust the AI. Without control, they can't manage risk. Both are essential for enterprise adoption.

### Real-world example

The semiconductor company's technical support chatbot needs to handle everything from basic datasheet lookups to advanced debugging and code generation. They require confidence scoring, conflict detection (when different documents contradict each other), and the ability to identify knowledge gaps. They also need guardrails for policy questions (like pricing) and technical constraints.

### What to look for

Platforms that provide continuous monitoring of AI performance, with the ability to adjust confidence thresholds, route low-confidence queries to humans, and maintain detailed logs for compliance.

## 3.5 Enterprise Security: Built for Scale and Compliance

Enterprise AI must meet the same security standards as any other enterprise software, with additional considerations for AI-specific risks:

- SOC 2 Type II and GDPR support from day one
- Tenant isolation across all architecture layers
- Role-based access control (RBAC) with audit logs and data classification
- Zero data retention options and on-prem deployment paths for sensitive industries
- Encryption at rest and in transit, with key management controls

### Why this matters

Security isn't just a checklist -- it's a foundation for adoption. Without it, trust erodes and projects stall. Procurement and IT teams require provable controls before approving any AI deployment.

### Real-world example

The semiconductor company has a rep portal that generates confidential, watermarked documents for customers with NDAs. They need AI solutions that can handle sensitive information while maintaining proper access controls and audit trails.

### What to look for

Platforms that offer multiple deployment models (SaaS, dedicated VPC, full on-prem), with clear compliance documentation and security architecture details.

## 4. Additional Success Factors

### Define Clear, Measurable Goals That Enable Fast, Safe Learning

Enterprises often delay AI adoption by waiting for perfect alignment or a single killer app. In reality, the most successful deployments begin with modest, clearly scoped goals that enable teams to learn and iterate quickly.

Examples include: - Cut ticket response time by 25% for a specific product line - Deflect 30% of common support inquiries from the website - Improve sales call prep efficiency by 40% in one territory - Reduce time-to-first-response on internal product questions by half

These goals should be easy to measure, tied to a real workflow, and achievable in 30–90 days. The priority is not to get it all right on day one—it's to start learning now, and do so safely and visibly.

### Control Your Data -- Don't Train Models on It

One of the most important architectural decisions is how AI handles your proprietary

information. Too many platforms offer fine-tuning or model training as a path to accuracy—but this often introduces data leakage, compliance risk, and long iteration cycles.

Instead, enterprises should choose platforms that: - Use retrieval-based generation (RAG) to access data securely at query time - Avoid training or storing your data in third-party models - Let you retain full control over what's ingested, indexed, and surfaced - Offer clear, auditable isolation between tenants, regions, and user roles

This approach enables learning and improvement without giving up control—and keeps data governance aligned with internal policies and external regulations.

## Address Real-World Enterprise Challenges

Based on direct customer feedback, enterprises face several specific challenges that AI solutions must address:

### Product Discovery & Selection

Customers increasingly need help choosing between similar products (e.g., 'Should I use BMI 330 or BMI 270 for my application?'). Sales reps and distributors often lack the expertise to answer these questions effectively.

### Technical Support Complexity

Support needs range from basic datasheet lookups to advanced debugging, code generation, and troubleshooting. The AI must handle the full spectrum while maintaining accuracy and providing appropriate escalation paths.

### Knowledge Management at Scale

Large enterprises accumulate massive amounts of technical documentation, community discussions, GitHub repositories, and internal knowledge bases. The AI must ingest, organize, and surface this information intelligently.

### Policy & Compliance Guardrails

AI must distinguish between factual questions and policy questions, applying appropriate business rules (e.g., pricing policies, technical constraints, safety guidelines).

## 5. Summary: Requirements That Predict Success

Requirement Area	What to Get Right
Deep Product Knowledge	AI that understands your catalog, specs, and support context
Cross-Use Case Leverage	Unified platform that supports learning across departments
Existing Tool Integration	Fit AI into existing workflows and tools, not the other way around
Visibility & Control	Real-time monitoring and configurable AI Performance Management
Enterprise Security	SOC 2, GDPR, RBAC, audit trails, and tenant isolation from day one
Strategy	Start with narrow, measurable goals that encourage fast, safe learning
Data Ownership	Retain full control of your data—do not train models on it
Real-World Challenges	Address product discovery, technical complexity, and policy compliance

Companies that deploy unified, product-aware AI—and do so with control, speed, and security—learn faster, perform better, and scale more confidently than those who wait. The key is choosing a platform that addresses all five critical requirements while enabling rapid iteration and measurable ROI.

**Remember:** The goal is not to replace your existing tools or processes, but to enhance them with AI capabilities that make your teams more effective. As one enterprise customer put it: "We don't want to rip and replace our CRM, ticketing system, or internal tools. We want AI where our team works."

## 6. Case Study: Liquid Instruments Uses One AI to Power Support, Sales, and Product Experience

Liquid Instruments, a leader in the test and measurement industry, designs precision tools used by engineers, researchers, and educators in fields ranging from aerospace to academia. Their flagship product, Moku:Lab, replaces multiple traditional instruments with a single, software-defined device—allowing users to configure oscilloscopes, spectrum analyzers, waveform generators, and more through a unified interface.

This innovation in instrumentation brought with it a new challenge: how to scale support and sales for a technically demanding global user base without overloading the team or sacrificing quality.

Their answer: a single, unified AI platform—powered by ept AI—capable of serving customers, employees, and prospects across the enterprise.

### One AI, Three Functions

Liquid Instruments began their AI journey in customer support, but quickly expanded to sales enablement and customer-facing product help—all using the same underlying platform and knowledge base. This approach minimized redundancy and ensured consistent, high-quality answers across every touchpoint.

#### Company Profile



**Type:** Innovative, fast-growing startup in Test & Measurement

**Customers:** Scientists and engineers in enterprise, defense, academic, and aerospace

**Requirements:** Accurate answers to complex technical questions

**Goal:** Improve customer experiences while accelerating growth

**Results:** AI across sales, support, teams, and in desktop & iPad apps

## Technical Support

Integrated into HubSpot, the AI initially launched as a co-pilot for support engineers. It provided expert-level answers drawn from manuals, guides, and documentation—reducing response time and increasing ticket quality.

## Sales Enablement

The same AI was deployed to support the sales team, still in HubSpot, but with a new voice and tone. Instead of engineer-speak, it delivered persuasive, benefit-driven answers tailored to prospects. Sales engineers could now respond faster and more confidently—even when navigating highly technical inquiries.

## Product Interface Help

In the most ambitious phase, the AI was embedded directly into the product, allowing end users to get real-time answers without opening a support ticket. Following a structured feedback sprint, the AI was fine-tuned for customer-facing use, leading to an immediate and measurable reduction in inbound questions.

*'ept AI enables users to quickly find answers in our Help tool, and turbocharges both our sales and support teams.'*

— Daniel Shaddock, CEO, Liquid Instruments

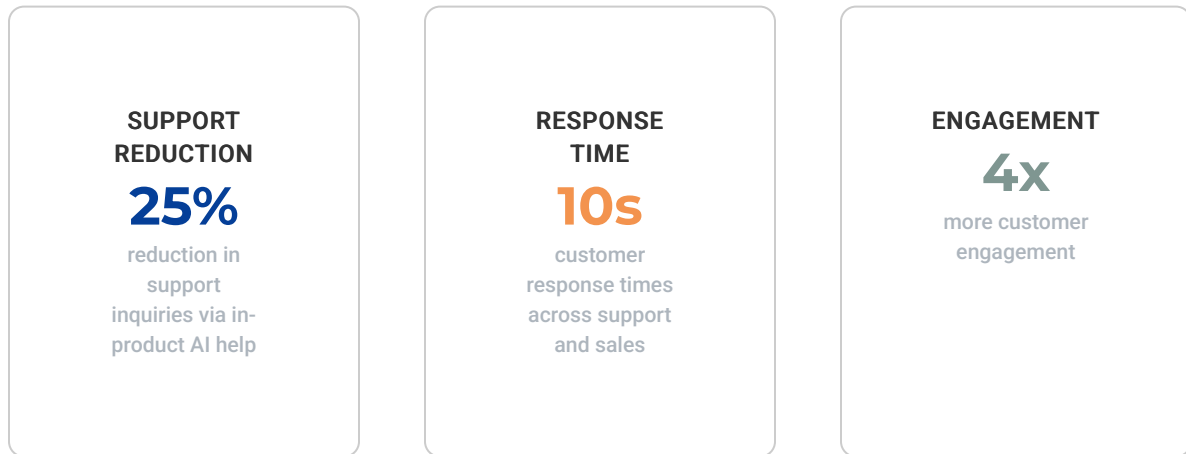
## Why It Worked

What made this deployment successful wasn't just the technology—it was the platform strategy. Rather than investing in separate tools for each function, Liquid Instruments used one AI that could:

- Adapt its voice to match support, sales, or marketing
- Grow smarter across use cases, with improvements in one area benefiting others

- Learn from their existing content, without requiring new documentation or model training
- Integrate natively with the tools their teams already used

This unified approach led to faster deployment, lower overhead, and better outcomes company-wide.



## A Foundation for Expansion

Encouraged by their results, Liquid Instruments is now expanding their AI footprint:

- Adding AI-powered answers to their website homepage to support prospective customers
- Using the platform to generate product marketing content like brochures and datasheets

They're not switching tools or retraining new models—they're simply unlocking more value from the same AI foundation.

# Ready to Transform Your Company?

See how Tech Product AI can accelerate growth.

**Schedule Your Demo**