OPTI
Smarter, Safer, Scalable

Google Cloud Partner

/preview/

AI Recommendations, Upsell, and Rules

From Algorithms and Guardrails to the Launch of UCP

March 2026

Guide #1 in the **AI for B2B 2026** Series

The **2026 AI in B2B** Report: Engineering, Not Magic

The Six Guides

- ○ **AI Recommendations, upsell and rules**
- // **Integrating precise data with AI**
- △ **Multimodal input**
- ◇ **Hybrid search and assistants (chatbots)**
- ⬠ **Continuous optimization and reporting**
- ⬡ **Security, auditing, and standards**

Why this series?

In 2026, most companies still rely on legacy software (e.g., legacy ERPs) as their core systems. However, customers and internal teams now expect e-commerce-level speed and intelligent interactions. Hybrid architecture (AI + classic tech) is the only realistic option for deploying AI at scale across production, distribution, and B2B sales. *See opti-software.com* ↗

About the Author

Marian Călborean is the founder of OPTI Software, with 20 years of experience in software architecture, from ERP integration and business workflows to B2B automation. He holds a PhD in Logic (2020, University of Bucharest) and was a Fulbright Visiting Scholar (2023, The Graduate Center, City University of New York), focusing on the intersection of formal logic and language. He teaches courses on *Logic and Software* and *Philosophy of Computing*.

As a practitioner, he leads software projects for companies in EU, Israel, and US. **OPTI Software** is a Google Cloud Partner and a HubSpot Solution Partner, holding ISO 9001 and ISO 27001 certifications, as well as ISO 42001. *Learn more on opti-software.com* ↗

Executive Summary

ERPs are powerful storage engines full of valuable data, but they don't know how to sell. Sales reps often overlook promotions and management's sales strategies. The leadership loses control over discounting policies and upsell opportunities. *How can AI help?*

B2B AI is engineering, not magic. This guide outlines a pragmatic approach: hybrid architecture. You don't need to replace your systems. Instead, an intelligence layer is built on top of existing data to drive new sales, while remaining under the strict control of business rules: minimum margins, stock levels, technical compatibility, and permissions.

To achieve rapid revenue impact, we recommend selecting 2–3 sales modules from the 10 proposed in *Ch. 3*, such as: *Smart Substitutes* (replacements when out of stock), *Cross-sell* (increasing AOV with complementary products), or *Intelligent Replenishment*.

***Company data* is a gold mine,** provided it is cleaned and correlated across ERP, CRM, e-commerce, WMS, and AI systems. This AI integration can be implemented ***on-prem*.**

There are *two technically legitimate* paths compared in this guide:

⚙ **Build (Custom)** for companies with proprietary data that need control over the source code.

☁ **Managed (Cloud Native)** for companies seeking rapid speed-to-market (2–3 months) by leveraging pre-trained technologies from Google.

How to Read This Guide:

- **CEO / Sales Director:** Start with *Ch. 1 (The "Why")* and *Ch. 3 (What We Deliver)*.
- **CTO / IT Architect**: Go straight to *Ch. 2 (How It Works)* and *4 (The Cookbook)*.
- **CFO / Financial Director**: Analyze the cost model in *Ch. 5 (The "When")*.
- Those interested in **emerging technologies** can explore *AI in B2B: What's New*.

Contents

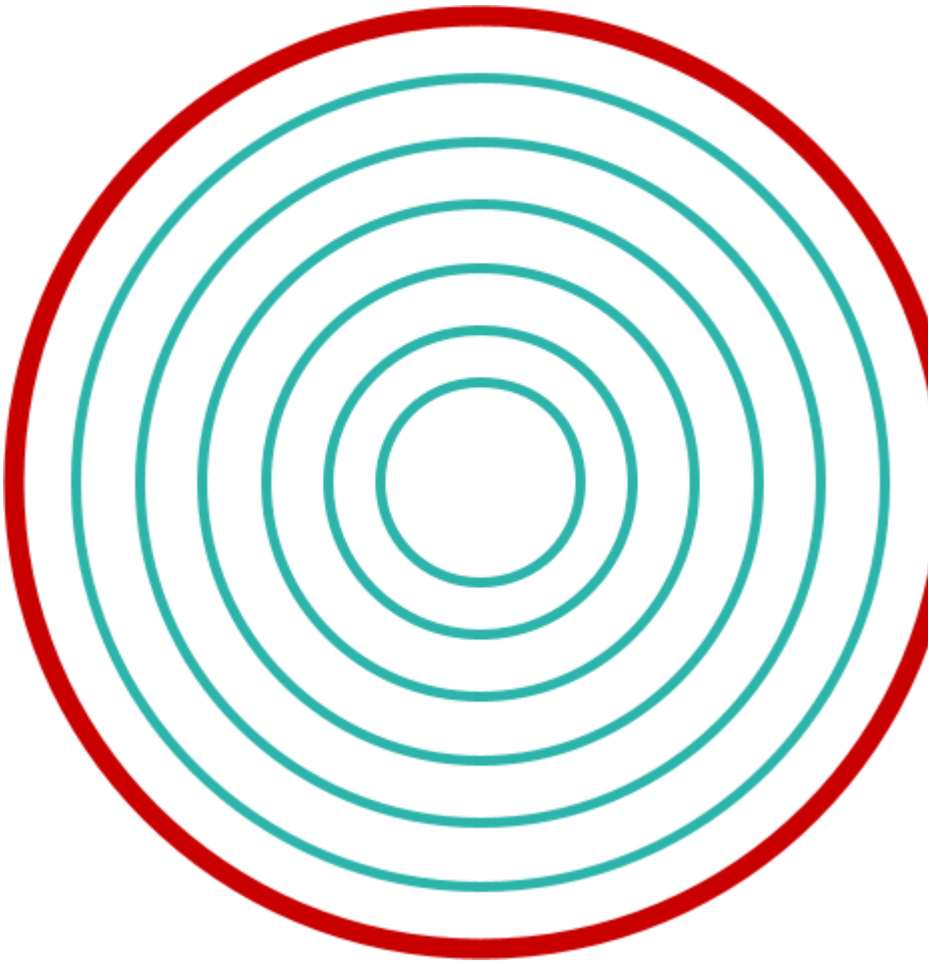

Chapter 01

BUSINESS CASE

Why AI for sales?

/fragments/

INDICATORS 1.1

Why personalization drives profit

Since AI adoption is accelerating, in 2026 **personalization and sales-oriented AI is becoming a competitive necessity.**

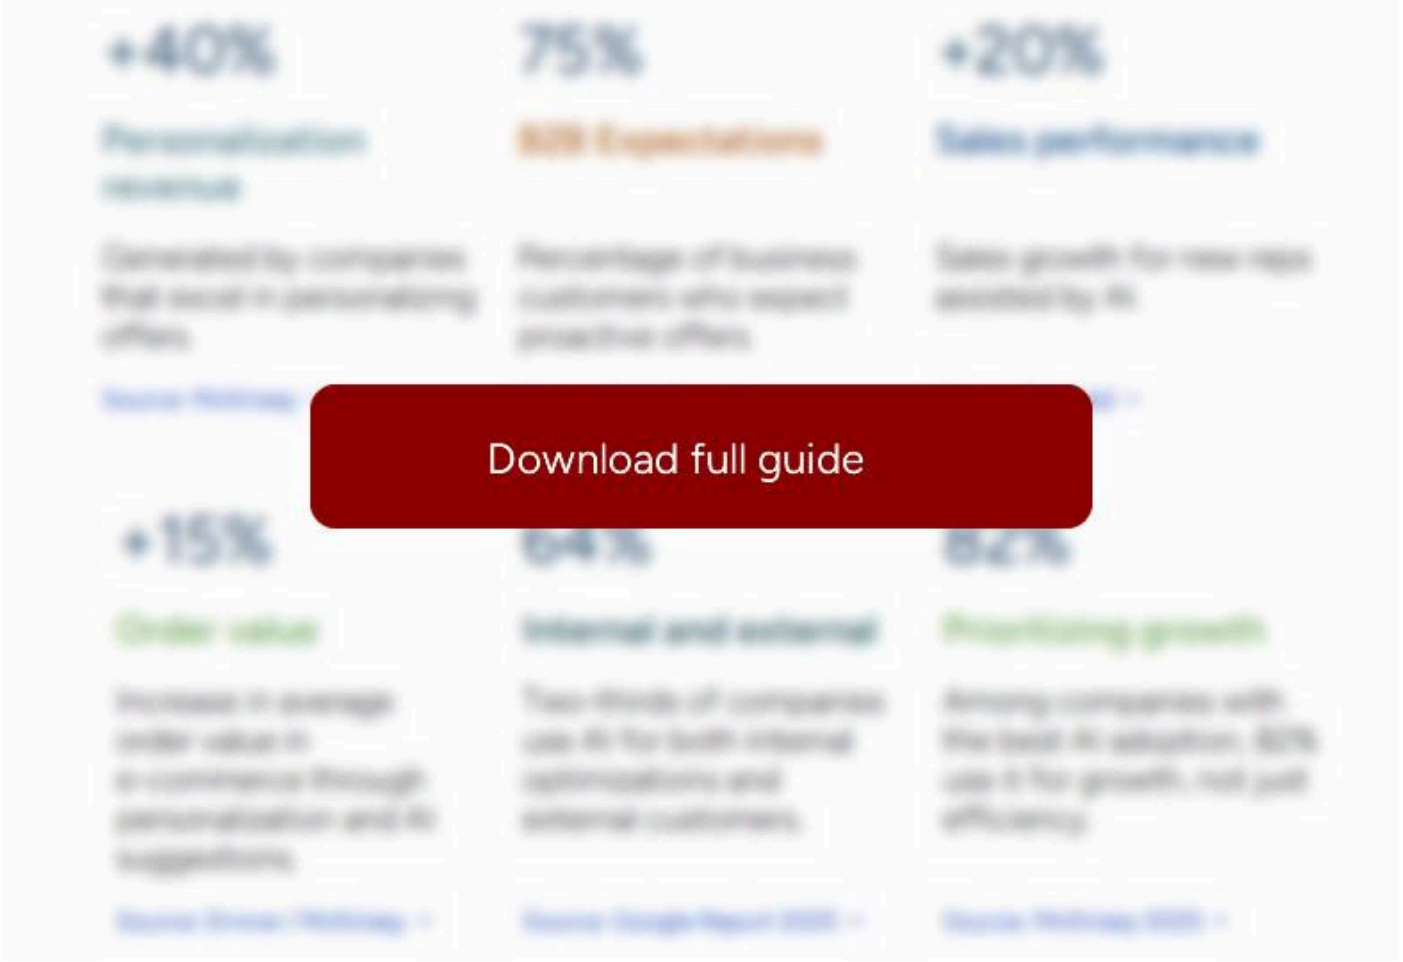

What do these figures say? To **justify the investment** (ROI) in AI, one can start with **sales**, especially for junior reps, with a focus on **growth**, not just efficiency (e.g., cost cutting).

CASE STUDY

1.4

How does human-AI collaboration look in sales?

Increasing order value without human labor

1.4.1

Situation: Junior reps sell basic equipment (low margin) and constantly forget to sell installation accessories (high margin).

Solution: The system detects the lack of a critical accessory in the cart based on a predefined rule and alerts the rep before submission.

Additional Opportunity: The system also displays other products usually bought together with these two, and the rep can add them.

Result: Immediate increase in order value, measured by AOV per segment.

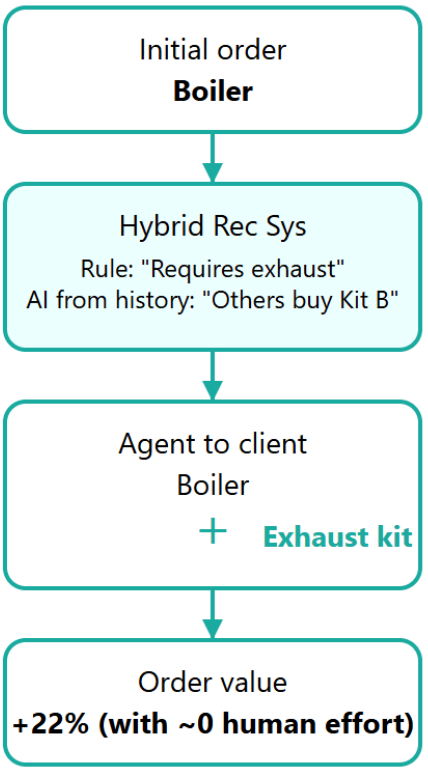

BROWNFIELD CONTEXT 1.5

Rust-Belt AI for contract pricing compliance and liability protection

Fig. 1.3: *B2B distribution market context in US 2025-2026*

If AI hallucinates a price or promises the wrong delivery date, the company may get sued for breach of contract. Just as importantly, guardrails are needed to ensure the client gets their specific negotiated price, not just a discount. Controlled AI can increase sales while protecting the company from financial risks.

1. The retirement cliff and turnover

Recruiting technical reps who understand products, compatibilities, and alternatives is hard. In retail and distribution, experienced specialists are retiring, according to a *Mercer study*. A junior rep needs time to learn a large catalog. Until then, they will sell only what the customer asks for, missing out on opportunities.

What does AI change? It can turn know-how into **institutional knowledge**: it suggests to reps (especially new ones) what to sell according to company policies.

As we've shown, B2B sales grow with artificial intelligence. The rest of the guide explains the AI technologies, available deliverables, and implementation methodologies, including several decision frameworks for management.

The bottom line:

B2B sales can grow if the company invests in a *hybrid AI architecture*: by connecting AI to company data from its core systems (e.g., ERP, WMS).

Later in the guide:

- ↘ Commercial rule examples in Ch. 2.5.1
- ↘ Four industries and ten concrete deliverables in Ch. 3
- ↘ What a minimal complete execution looks like in Ch. 4
- ↘ Implementation plan and budgeting in Ch. 5
- ↘ News: AI Agents and BOM generation
- ↘ Resources, studies, and technical glossary

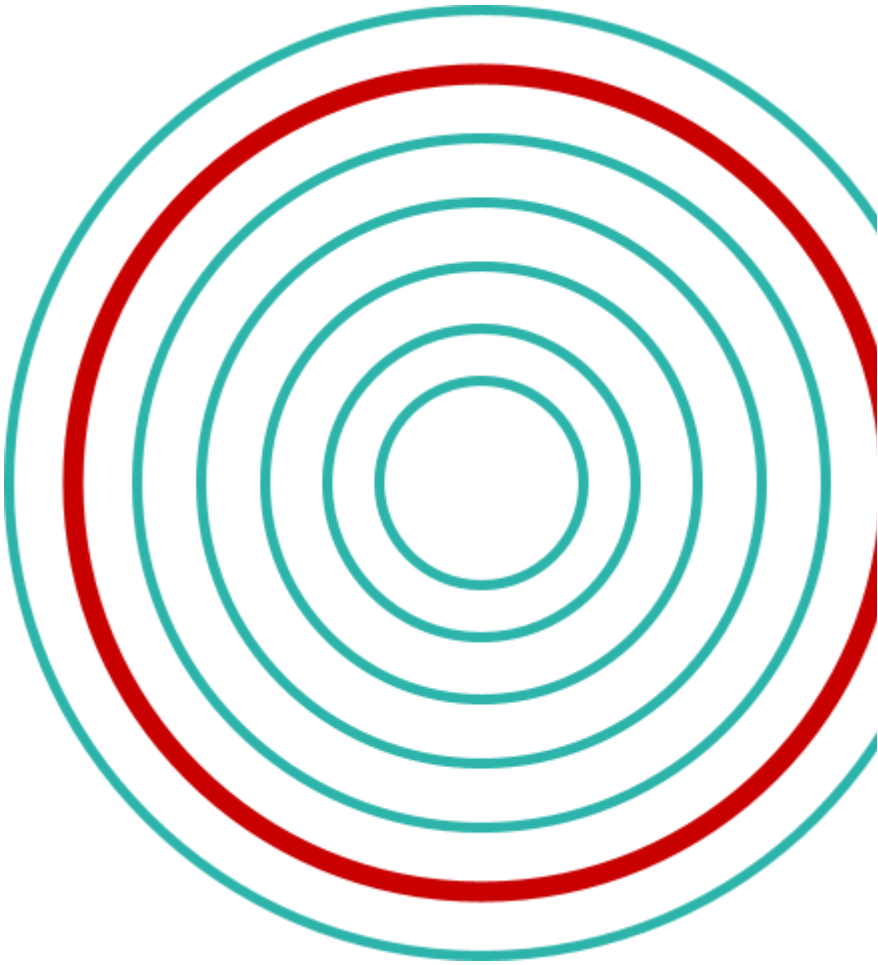

Chapter 02

THE TECHNOLOGY

How does it work?

/fragments/

Did you know that B2B has some advantages when it comes to implementing AI in sales? Basket sizes are large, products usually have unique codes, and there are many recurring transactions. Before debunking myths in *Ch. 3*, here we will examine the **technology.** The non-technical reader can skip to the following chapters.

Contents:

CHAPTER PLAN

Introduction

2.0

This chapter covers the ***evolution*** of recommendation systems, **objectives,** measurable AI and the **architecture** of top-tier solutions. Then it **compares the Build** (custom) approach with **Managed** (native cloud solution) and presents the diagram of a ***hybrid "privacy-first"*** implementation from our projects. Finally, we discuss the challenges of implementing AI for recommendations: the enforcement of ***business rules*** and ***the data contract*** that guarantees the quality of AI recommendations.

As a **technological introduction**, this is the **funnel architecture** diagram for a recommendation system:

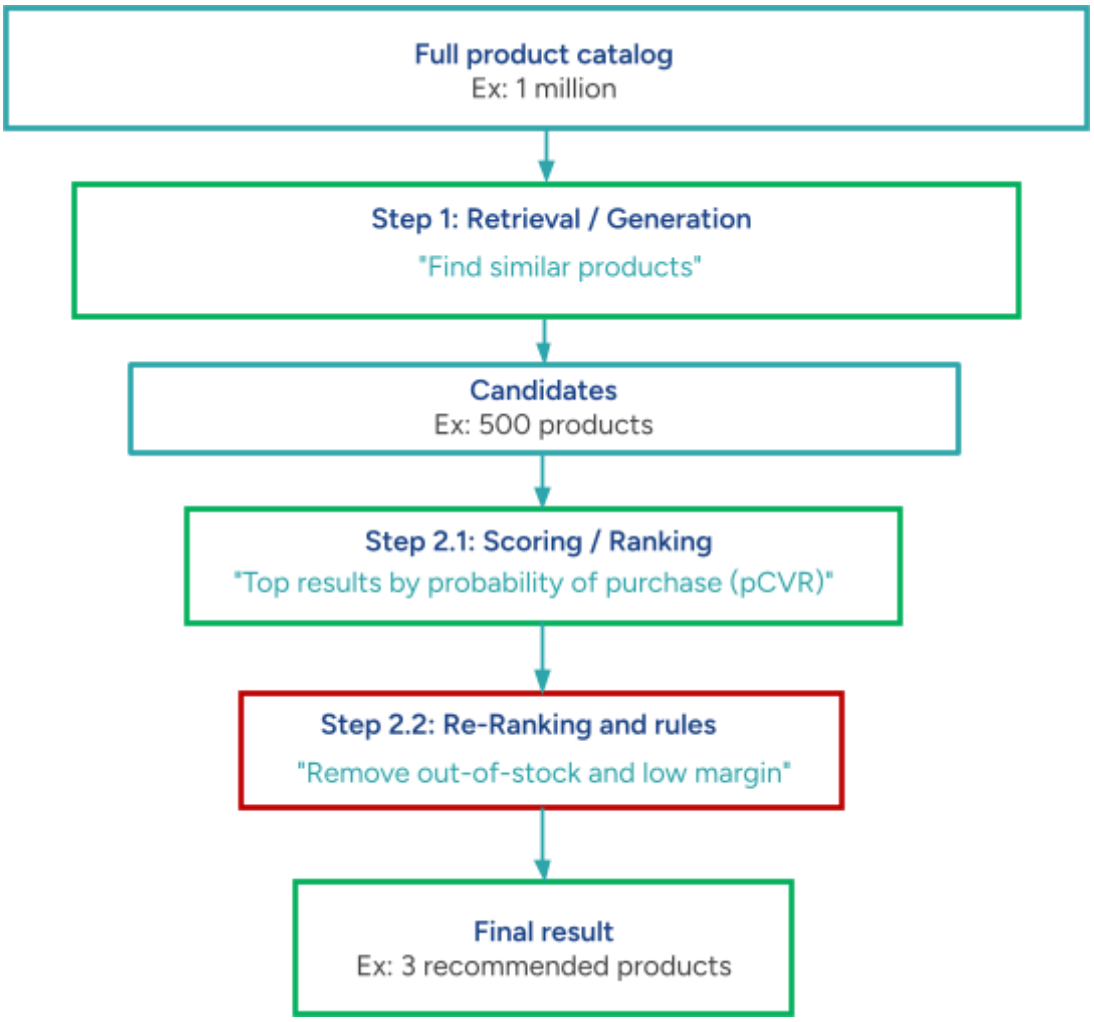

Fig. 2.1*: AI Recommendation Pipeline: finding candidates, ranking them, and applying business rules*

Generations	*Decision Basis*	*Advantages*	*Risks*
Gen 0	**Memory**	Excellent context	Loss of know-how with staff turnover. Difficult training.
Gen 1	**Static Rules**	Total control	Impossible maintenance at high volume.
Gen 2	**Statistics**	Automation	Cold Start: Does not recommend new products or to new customers.
Gen 3 (predictive)	**Deep Learning**	Intent prediction. Semantic search. Chatbots (LLM-enhanced).	Requires pre-trained models (e.g., Google) or tuning on sufficient data.
Gen 4 (agentic)	**Generative Deep Learning**	Reasoning, performs actions.	Not the first choice for a scalable and stable recommendation system.

***Legend**: Generations of artificial intelligence.*

Context vs. prompt and the predictive-generative combination

In **predictive** AI (Gen 3), models operate on **structured context**: products, events, customers, rules. ***Context engineering*** (selecting and cleaning attributes, defining correct events, guardrails) has a major impact on result quality.

The prompt as a tone control instrument only appears in applications built on top of the predictive foundation with the addition of a **generative** function (Gen 4). For example, in *Vertex AI Search for commerce*, you can create chatbots that separate search intent from purchase intent by connecting in the new *Conversational* mode (↘ *see Guide #4*). The Conversational agent will run over the predictive base that associated those products.

For business:

Criterion	⚙ *Build: agnostic architecture (custom)*	☁ *Managed: Vertex AI Search for commerce (Cloud Native)*
Infrastructure	Usually in the **cloud** (including Google Cloud) via technologies like Databricks, on a pay-as-you-go basis	**Google Cloud,** primarily paid per business event (queries, predictions, training).
New products (Cold Start)	**Requires effort.** Possible with the usage of a Feature Store.	**Functional**. Uses metadata, even images and model pre-training.
Learning types	**Individual implementation.** Any model can be developed.	**Pre-defined models and objectives.** Requires custom implementation for extra features.
Business rules	**Requires custom code** ("IF brand = X THEN...").	**Partially configurable** via the Visual console / JSON for Boost, Bury, Filter.
Critical dependency	The implementation team. Usually, you have **access to the source code.**	**Mature technology, without access to source code.** Specialized in product relevance.
Maintenance	**Constant monitoring** of models (for drift). You can do debugging.	**Automatic tuning** of models.
Do they allow chatbots or assistants?	Of course, ↘ *see Guide #4*	

CASE STUDY

↔ Hybrid "privacy-first" architecture

2.4.4

For companies with strict security requirements or on-prem ERP systems, we show, based on OPTI Software projects, what a centralized **middleware layer** for commercial truth can look like.

It can function in the cloud or on-prem and connects selectively to AI functions, whether managed by Google or custom via *Model Garden*.

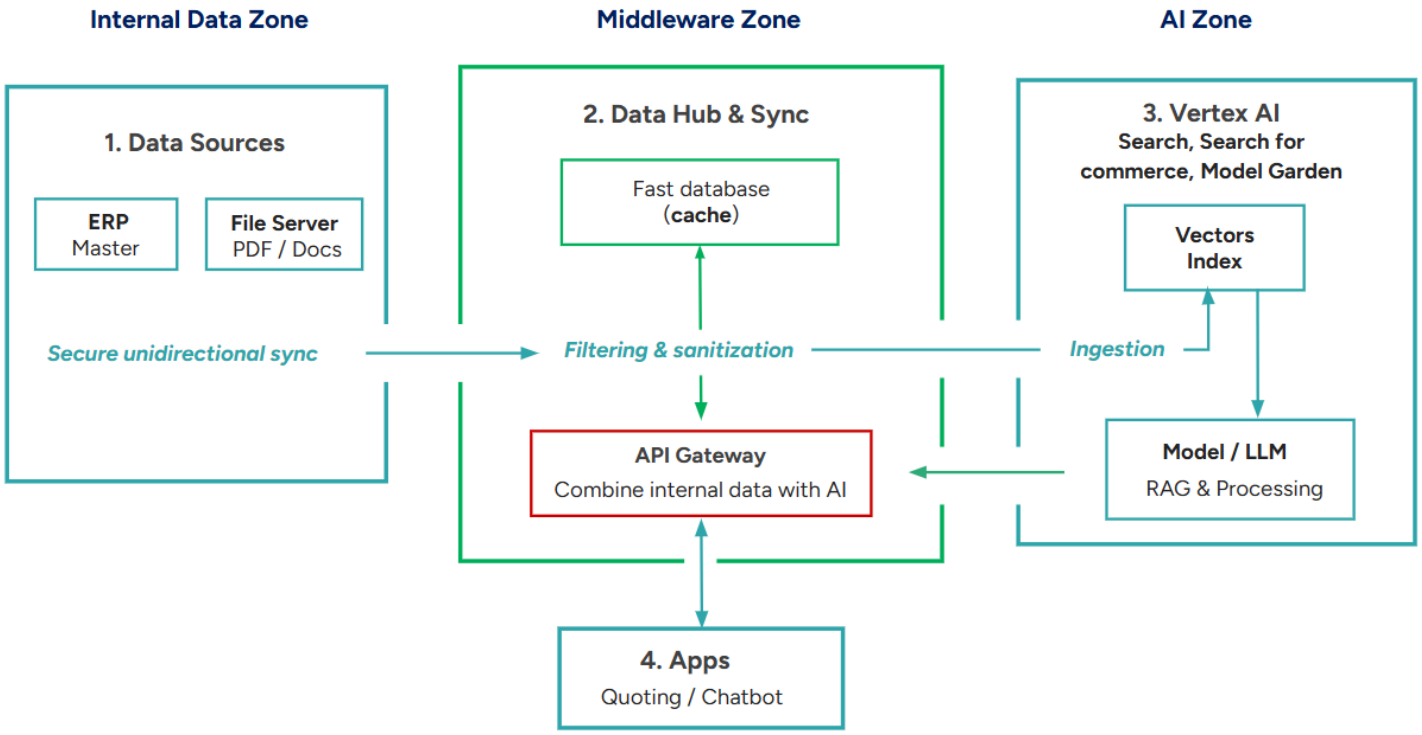

Fig. 2.5: Middleware as a single source of commercial truth, isolated for security

The first variant involves modifying the shopping cart and the ordering module. The second variant involves deep integration with automated commercial flows.

In addition to these challenges, there are others related to scalability, latency, security, and monitoring, depending on the complexity of the company's workflows. But AI technology for B2B sales **exists**, as shown in this chapter.

The bottom line:

2.7

In 2026, the best sales rep is not AI, but humans assisted by AI. We have analyzed the technological engines for AI sales: vector search, re-ranking, guardrails, and the data contract. Companies that clean their data and define their commercial rules will be the ones that benefit from AI in sales, both in the Build (custom) and Managed (Cloud Native) versions, depending on the company's choice.

Which concrete applications can be used by sales reps?

↘ *See ten individual deliverables in Ch. 3*

↘ *See Glossary and Resources for technical terms*

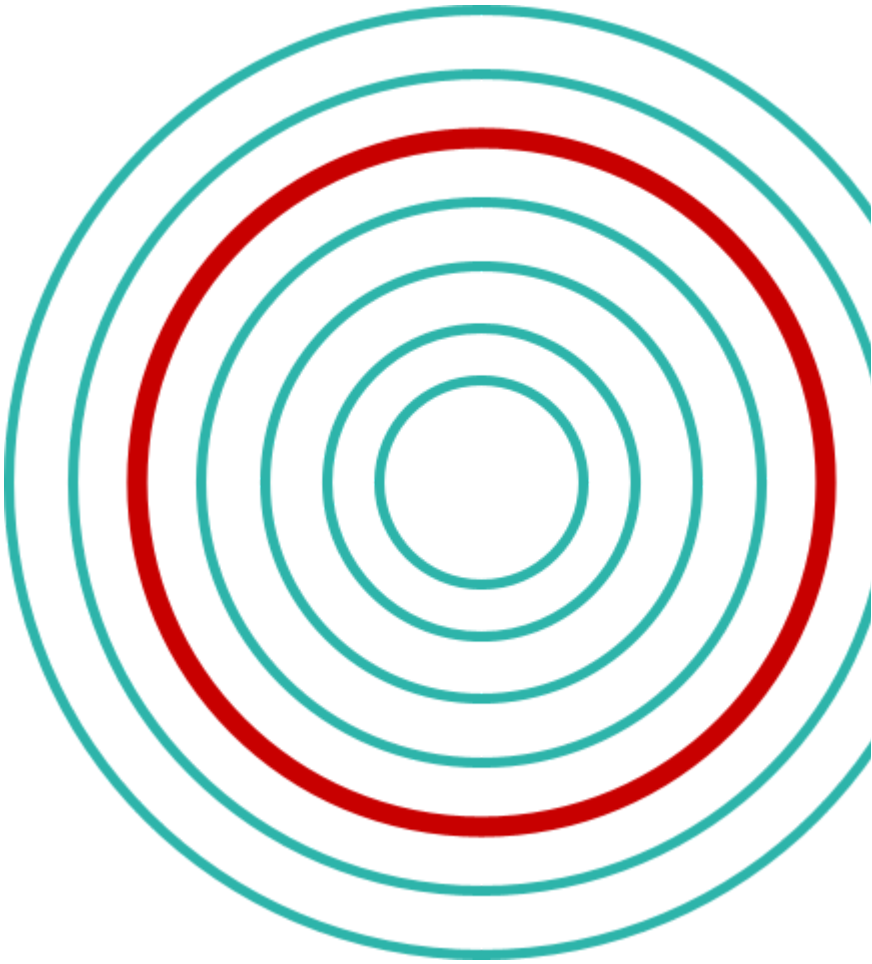

Chapter 03

FROM MYTHS TO CONCRETE APPLICATIONS

What can be delivered?

/fragments/

Did you know where AI projects fail? For the Enterprise sector in general, the governance structure and permissions represent the steepest obstacle. For mid-market companies, the obstacle is usually securing access to fundamental data sources (e.g., ERP). Still, let us not be pessimistic. We will present 4 industries and 10 concrete deliverables for increasing B2B sales with AI.

Contents:

Here are **four ideal industries** for AI implementation via hybrid architecture.

❒ Wholesale: Technical and IT distribution

Why?	**Varied margin in a competitive market.** A laptop has a 2% margin, but the bag and mouse have 40%.
AI role:	AI, through a recommendation system, "understands" the ideal basket and suggests accessories. It can offer dynamic discounts in a controlled manner.
KPI:	Increasing AOV (average order value).
Risks:	Irrelevant recommendations or discounts presented without brand image protection.

✚ Medical: Pharmaceutical equipment and consumables

Why?	**Recurring orders in a sensitive and regulated context.** Baskets are a mix of products that must respect protocols and compatibility. There are batches and expiration dates.
AI role:	Suggests consumables based on installed equipment and expirations. Builds compatible bundles per procedure type. Can offer discounts based on expiration date, FEFO logic (first expired, first out).
KPI:	Reduced quoting errors, optimized stock rotation.
Risks:	Justified reluctance in a regulated context. Fear of technical incompatibilities.

Saving the sale 3.3.2

⇄ Smart Substitutes Engine

Intelligent product substitution when out of stock, for faster rotation, or for delivery from a single location.

Why?	Saving the sale when stock is 0 (e.g., auto, IT, construction, pharma) due to supply issues or fast distribution. Transport optimization with multiple warehouses (WMS).
AI role:	Suggests many possible equivalents (different brand, similar specifications) available for immediate delivery, which are then checked by deterministic rules.
KPI:	Conversion rate for out-of-stock products (direct or mediated e-commerce). Inventory turnover for alternatives. Lost sales (rejected quotes).
Risks:	Recommending a technically incompatible product (maximum return risk). WMS integration complexity for transport optimization.
Flow:	AI is used to suggest the closest products, filtered strictly by compatibility rules (requires clean, normalized data). Then the margin and delivery distance are optimized.

Architecture diagram

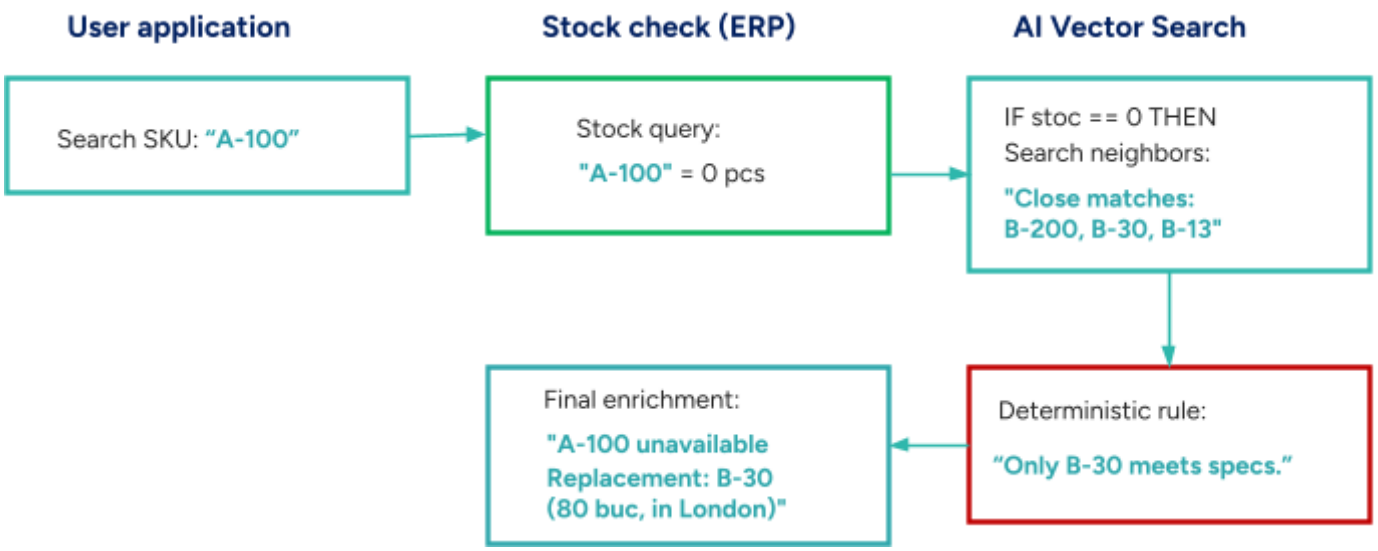

Fig. 3.3: Architecture of a Smart Substitutes engine for intelligent replacement

Checklist: What to examine for AI implementation

- ✔ Minimum data and required data volume
- ✔ Technical compatibility rules
- ✔ Data cleanliness and normalization
- ✔ Stock rotation and lost offers

The bottom line:

Setting myths aside, B2B also has advantages over B2C for AI implementation. These ten deliverables can be built, each with a **mandatory risk and profitability analysis.**

For example, **AI temporal analysis** is very powerful for seasonality (e.g., agriculture), and **dynamic price optimization** can be ideal if there is strict control and compliance verification.

The role of AI in B2B sales is to bring plausible options forward and reduce search time. Implementation follows the commercial strategy and receives **continuous feedback (positive and negative)** to learn the company's niche.

↘ *See in Guide #5. Optimization and reporting*

Want to see what a minimal complete implementation plan looks like? *See ↘ Ch. 4*

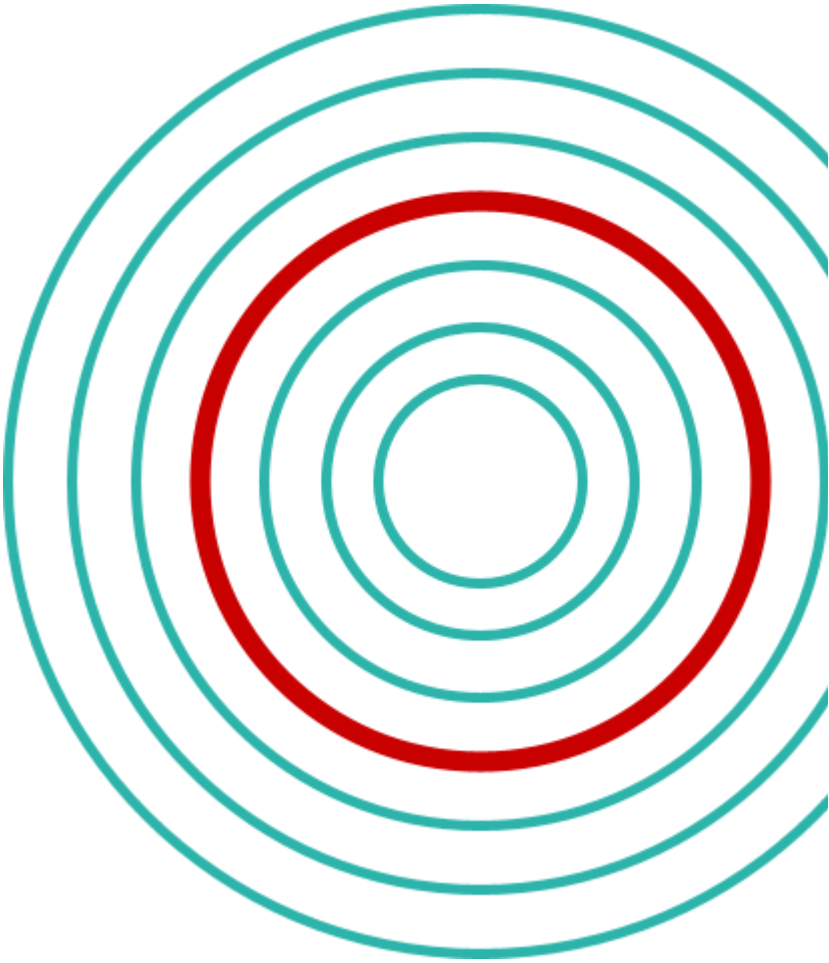

Chapter 04

A MORE DETAILED "HOW"

Technical **Cookbook**

/fragments/

Did you know that AI systems are also built by programmers? At least for the moment, programmers are still needed to integrate AI with existing processes and software. This chapter is dedicated to those who want to see small snippets of such an integration. If you are a decision-maker, the minimum (but complete) implementation flow in 4.1 is for you. Or you can skip to the *next chapter*.

Contents:

Guardrails: post-processing rules over the AI score 4.2.1

The AI has produced a **raw score** for each candidate, and guardrails must ensure we do not sell products below the minimum margin or without stock, and that we don't leave the UI empty if **APIs** fail.

Example B1: Filtering in the recommendation pipeline (Python)

Snippet: pipeline that receives candidates with scores and applies guardrails before sending to UI:

```
from dataclasses import dataclass

@dataclass
class Candidate:
    sku: str
    score_ai: float
    stock: int
    margin_pct: float
    forbidden: bool = False

def apply_guardrails(
    candidates: list[Candidate],
    min_margin_pct: float,
    min_score_ai: float,
    max_items: int = 10,
) -> list[Candidate]:

    # 1. filter out prohibited products, out-of-stock items
    # or those below the minimum margin
    filtered = [
        c for c in candidates
        if not c.forbidden
        and c.stock > 0
        and c.margin_pct >= min_margin_pct
        and c.score_ai >= min_score_ai
    ]

    # 2. If the list is too short after filtering,
    # supplement it with internal fallback (get_top_selling_products)
    if len(filtered) < max_items:
        missing = max_items - len(filtered)
        exclude_skus = {c.sku for c in filtered}
        filtered.extend(get_top_selling_products(exclude_skus=exclude_skus,
limit=missing))

    # 3. sort by AI score and limit results
    filtered.sort(key=lambda c: c.score_ai, reverse=True)
    return filtered[:max_items]
```

The bottom line:

4.4

A recommendation system can be implemented either in the ⚙ Build version (custom), or in the ☁ Managed (Cloud Native) version, with a trade-off between low-level control and implementation/maintenance speed:

- ↘ In the ⚙ Build version, rules are code (e.g., Python, Java, Go).
- ↘ In the ☁ Managed version, rules generally move from code into **declarative configurations** and **data engineering**.
- ↘ In the ☁ Managed version, the model is managed by Google, and developers focus on data synchronization, defining guardrails, and potential extensions (e.g., temporal forecasts discussed in *Ch. 2.1.4*).

If you are evaluating the possibility of implementing one of the two versions in-house, you can follow the project steps recommended above. The ⚙ Build version is recommended if you have an ML development team and your own data. The ☁ Managed version is recommended for faster delivery.

↘ *See the cost matrix and our methodology in Ch. 5*

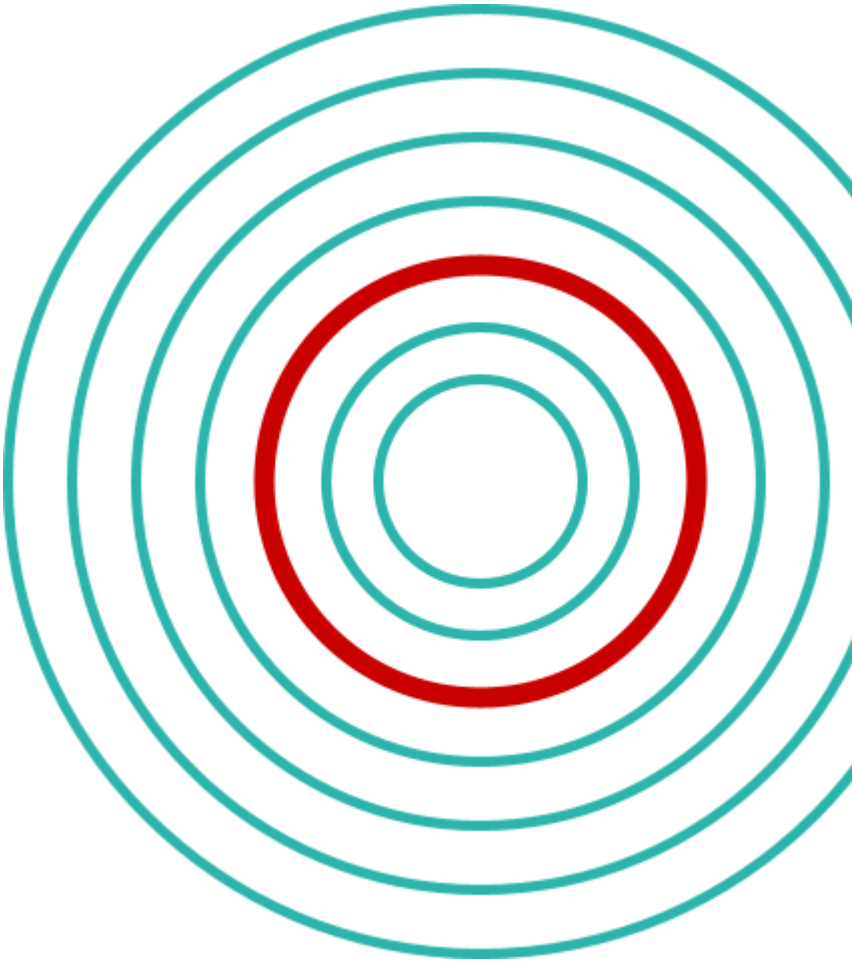

Chapter 05

NEXT STEPS

When can I start growing the business?

/fragments/

This chapter is written for management: CEO, CFO, CTO, Sales Director. It synthesizes the messages of Chapters 1-4 and puts forward the decisions which need to be made, the resources which are required and the cost structure of a B2B AI implementation. We also list a few guarantees of Google Cloud and the methodology we at OPTI Software use to deliver quality.

Contents:

Note: OPTI Software is a Google Cloud and HubSpot Solution partner. The guides express our position and have not been endorsed or supported by Google.

GUIDE SUMMARY

Principles for decision

5.1

As a summary of the previous chapters, the most important ideas are:

AI for sales is a pipeline

It is not a single model, but a sequence of steps: Data → Events → Candidates → Ranking → Guardrails → UX → Feedback loop.

B2B has advantages in adoption

Large baskets and recurring transactions make AI recommendations and automated quoting possible.

Hybrid architecture is the safest choice

The combination of Deep Learning, strict business rules, and clean data ensures growth without redesigning the company's IT infrastructure.

Build if you have plenty of data and a solid team

The Build (custom) approach is recommended for organizations with machine-learning teams and solid data.

Managed if you want quick results

The Managed (Cloud Native) approach is recommended for fast results with less infrastructure.

Data quality and KPIs before anything else

AI cannot work magic on dirty data. Establishing measurable KPIs helps the organization in adoption.

Evolution of roles

AI transforms the sales rep into a negotiator and consultant.

Here are the resources needed for success and how you can calculate your costs.

Setup: Initial Construction

Where are the architecture and development days consumed?

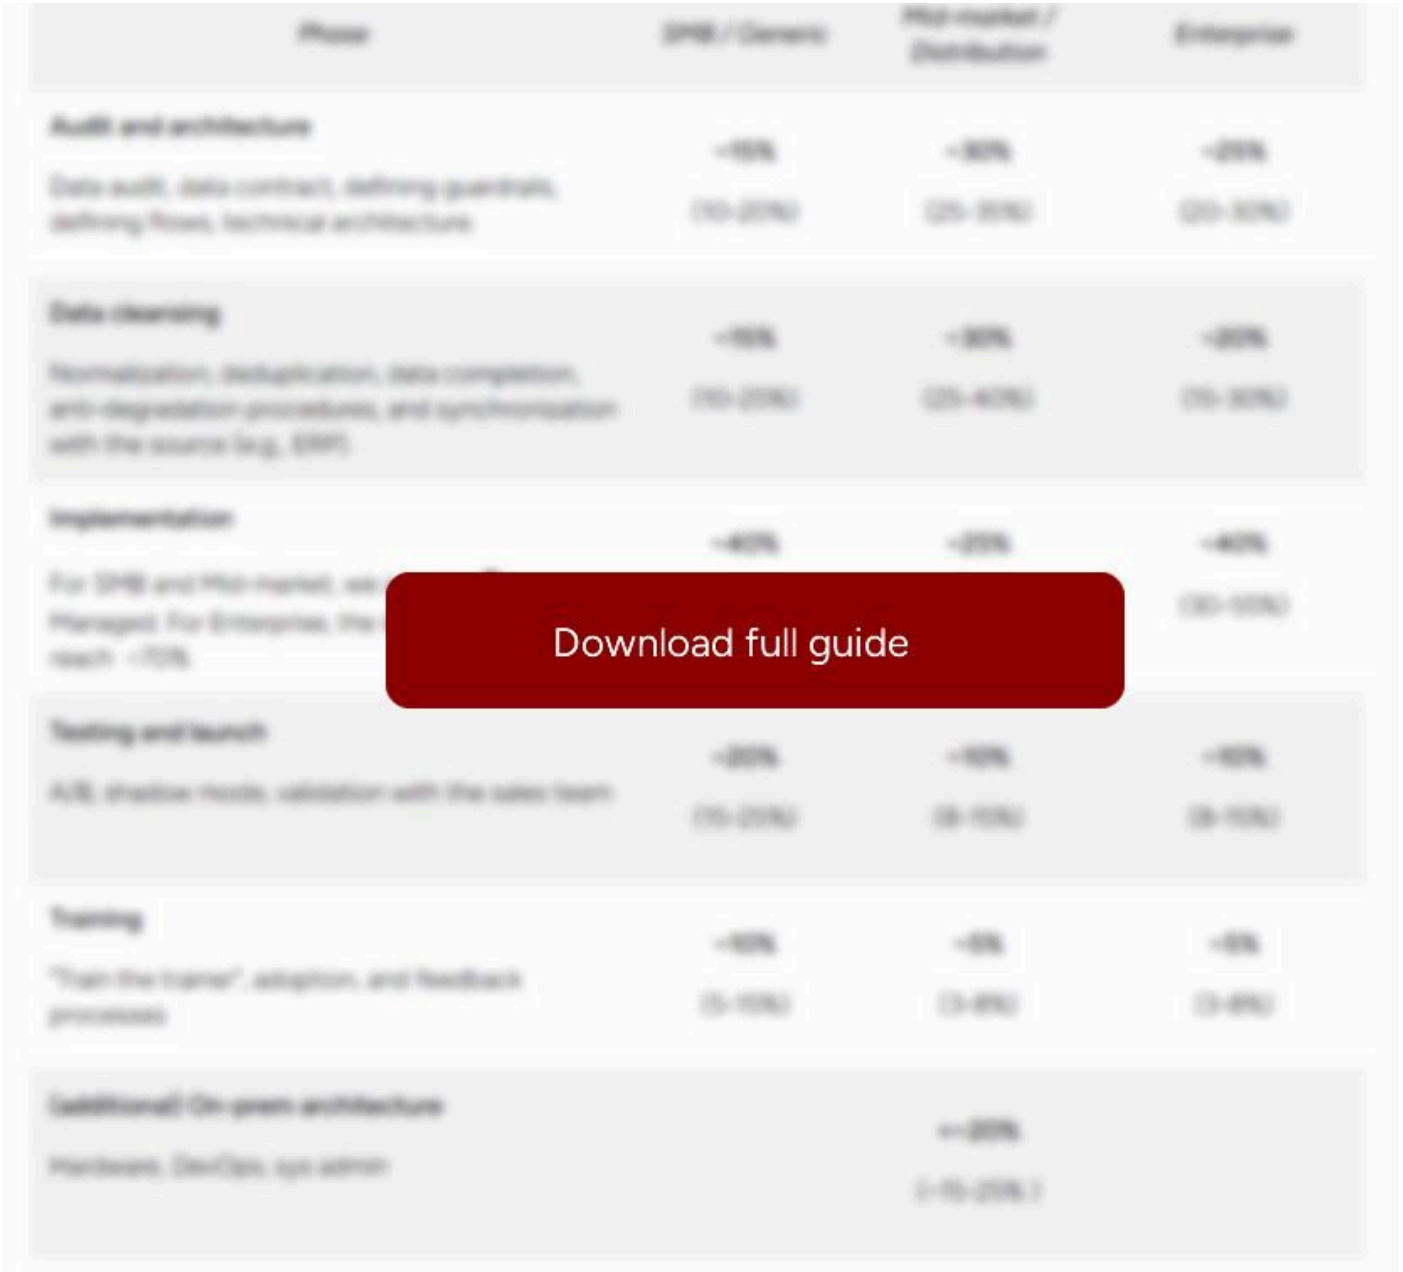

Final recommendations:

5.4

This guide has described a structured way of adopting AI for company growth in 2026.

Returning to the metaphor from *Ch. 2.5.2*, **imagine explaining to a person for the first time what you have in your company. Without abbreviations, without the procedural shortcuts developed over time, and without rushing. When you have finished the explanation, you are ready for AI, and you have additionally gained operational clarity.**

The ↔ hybrid architecture will remain valid even when models, vendors, or channels change. It can deliver the benefits of AI with the stability characteristic of B2B.

↘ *See technologies for 2027 in Ch. 6.*

For resources used and terms, See ↘ Resources and Glossary.

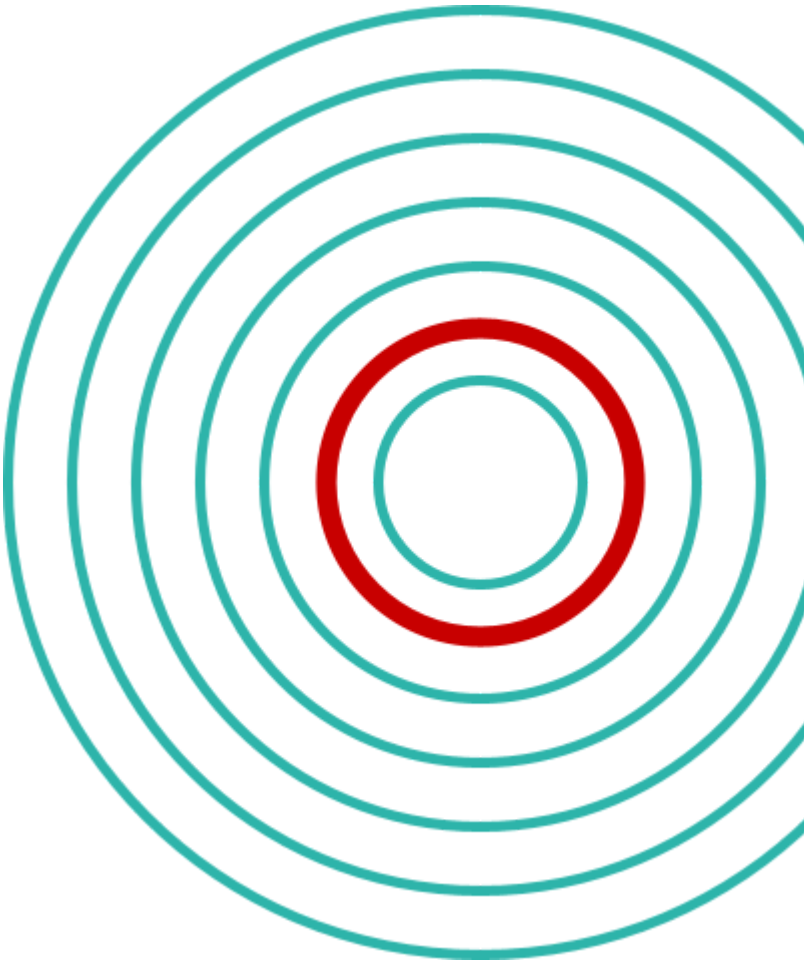

Horizon 2027

New AI tech for B2B sales

6.0

OPTI
Smarter, Safer, Scalable

Guide #1 - AI Recommendations, Upsell, and Rules
Download the complete guide (172 pages) at
www.opti-software.com/ai-2026/0904

In the full version:

- ↘ +9 concrete AI deliverables and architecture diagrams
- ↘ +19 code examples
- ↘ 8 steps for data cleansing
- ↘ AI Budgeting: Capex vs. Opex
- ↘ AI Agents for 2027

Download the full manual