

*This preview is an initial hypothesis based on public website information and the inputs provided. It is not a full AI strategy. The next step is to validate assumptions and define a pilot with clear success metrics.*

## EXECUTIVE SNAPSHOT

**Confidence: Medium**

The company appears to be in the early-pilot stage of AI adoption: copilot tools are in use across individual functions, but value is fragmented and not yet tied to executive KPIs. The strongest near-term lever is to consolidate effort around quotation responsiveness and engineering knowledge access — both close to revenue and measurable inside one quarter.

Primary KPI (quote cycle time) is concrete and measurable — strong basis for a scoped pilot.

Company size and configure-to-order model are well-matched to the identified opportunity areas.

Systems landscape and data architecture require validation before a pilot can be fully specified.

AI maturity level suggests enthusiasm in individual functions, but no enterprise-wide governance yet confirmed.

## CURRENT-STATE HYPOTHESIS

### Operating context

Mid-sized Swiss precision manufacturer with a configure-to-order model. Technical complexity is high; each quote requires engineering input. Customer base spans industrial OEMs and tier-1 suppliers across the DACH region. Microsoft 365 is the primary productivity environment.

### Likely bottleneck

Quotation cycle time is the primary constraint: sales engineers hand-search technical documentation, re-draft from prior quotes, and escalate to senior engineers for sign-off — a process that likely takes days where hours would win more business.

### AI adoption pattern

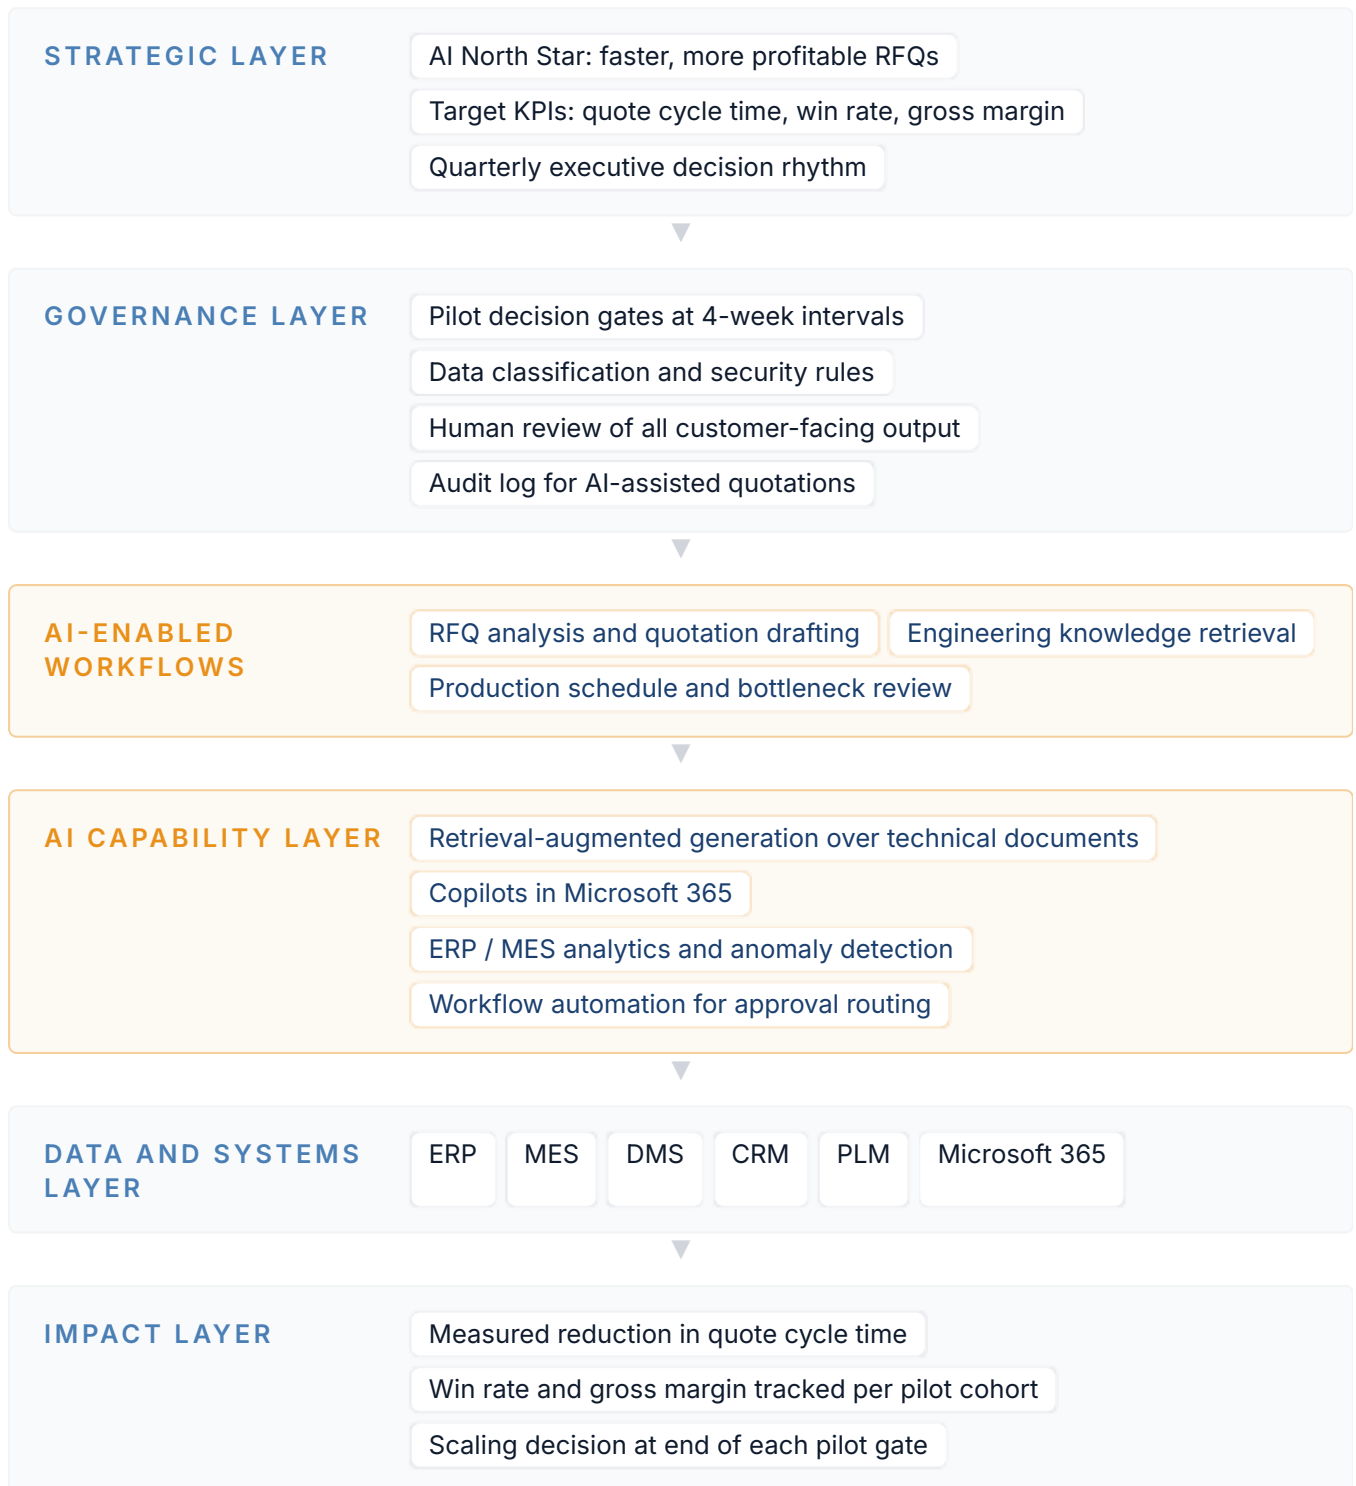
Individual experimentation stage: Copilot is in use for drafting and summarisation in some functions, but there is no shared prompt library, no retrieval over company documents, and no measurement of AI's impact on output quality or speed.

### Strategic risk

Competitors with faster configure-to-order responses will erode win rate before internal inefficiency becomes visible in the P&L. The risk is not technological — it is the cost of fragmented adoption without a North Star tied to revenue.

## AI-ENABLED FUTURE STATE

### AI-enabled future state — precision manufacturing



## AI OPPORTUNITY AREAS

### OPPORTUNITY 1 · STRATEGIC FIT 5/5

#### AI-assisted quotation and technical sales support

Use AI to accelerate RFQ analysis, draft quotation inputs, retrieve technical documentation and prepare margin-aware proposal drafts that the sales engineer reviews and signs off. Every customer-facing output remains human-approved.

**Business value:** Shorter quote cycle times, more consistent proposals and improved sales responsiveness — directly tied to win rate and gross margin.

Complexity: **Medium** Inherent risk: **Low** Controllability: **High**

KPIs: Quote cycle time · Win rate · Gross margin · Sales response time

*Why chosen: Directly addresses the stated primary KPI (quote cycle time) and is close to revenue. The configure-to-order model means every efficiency gain here compounds across the full order pipeline.*

### OPPORTUNITY 2 · STRATEGIC FIT 4/5

#### Engineering knowledge assistant

Make tacit engineering knowledge searchable: prior project files, technical specs, supplier datasheets and internal best practices, exposed through a controlled assistant grounded in the company's own documents.

**Business value:** Reduces dependency on a few senior experts, shortens onboarding time and protects know-how as the workforce evolves.

Complexity: **Medium** Inherent risk: **Low** Controllability: **High**

KPIs: Search and retrieval time · Onboarding time for new engineers · Issue resolution time · Senior engineer escalation rate

*Why chosen: Addresses the knowledge-concentration risk that surfaced across multiple pain points. Complements the quotation pilot by reusing the same document retrieval infrastructure.*

### OPPORTUNITY 3 · STRATEGIC FIT 3/5

#### Production planning and bottleneck intelligence

Use AI to surface schedule risks, suggest sequencing alternatives and explain bottlenecks in plain language for plant managers and supervisors — shifting the weekly review from reactive diagnosis to forward-looking decision-making.

**Business value:** Better lead-time predictability and earlier escalation of capacity issues before they reach the customer.

Complexity: **High** Inherent risk: **Medium** Controllability: **Medium**

KPIs: Schedule adherence · Lead time · OEE · Throughput · Customer on-time delivery

*Why chosen: High strategic value but dependent on ERP and MES data quality — recommended as Phase 2 after the quotation pilot establishes AI governance foundations.*

## WORKFLOW COMPARISONS

### RFQ ANALYSIS AND QUOTATION DRAFTING

#### Today

A sales engineer receives an RFQ, manually searches the DMS for similar past projects, extracts specifications by hand, and drafts a quotation document from scratch or from an outdated template. Escalation to a senior engineer for technical validation adds one to two days.

→ **Faster quote turnaround and more consistent proposal quality**

#### AI-enabled

The engineer submits the RFQ to an AI assistant that retrieves the three most relevant prior projects, summarises the key specification differences, and produces a structured draft. The engineer reviews, adjusts pricing and approves. Senior escalation is reserved for genuinely novel configurations.

### ENGINEERING KNOWLEDGE RETRIEVAL

#### Today

Junior engineers and new hires rely on colleagues' availability to locate design decisions, tolerance standards and supplier constraints. Know-how is embedded in individual expertise and scattered across project folders with inconsistent naming conventions.

→ **Reduced knowledge dependency on key individuals**

#### AI-enabled

A retrieval-augmented assistant indexes technical documentation, past project files and datasheets. Engineers query in plain language and receive sourced answers with document references. Senior engineer availability is freed for high-judgement tasks.

### PRODUCTION SCHEDULE AND BOTTLENECK REVIEW

#### Today

Plant managers review schedule adherence in a weekly meeting using manually compiled spreadsheet reports. Bottlenecks are identified reactively, often after a customer delivery has already been affected.

→ **Earlier escalation of capacity risks before customer impact**

#### AI-enabled

An AI layer over the ERP and MES surfaces emerging schedule risks daily in plain language, flags the most likely root cause and suggests sequencing alternatives. The weekly review shifts from diagnosis to decision.

## RECOMMENDED FIRST PILOT

### AI-assisted RFQ analysis and quotation drafting pilot

It is directly tied to the primary KPI (quote cycle time), close to revenue, visible to management and suitable for a controlled pilot on one product line. It does not require a full enterprise AI platform and can demonstrate measurable value within 8–12 weeks.

#### Scope

One product line, one customer segment (e.g. DACH industrial OEMs). AI supports RFQ summarisation, prior-project retrieval, technical-document search and structured draft preparation. Sales engineers remain the sole decision-makers; AI prepares, they review and approve.

#### Duration

8–12 weeks

#### Target users

Sales engineers assigned to the selected product line, Technical sales support, Pricing and margin analyst (review role), Engineering lead (escalation and validation)

#### Required data

Historical RFQs and accepted quotations (last 24 months), Technical product documentation and datasheets, Pricing and discount guidelines, Standard terms and conditions, Prior project files and engineering notes (DMS)

#### Success metrics

- Measurable reduction in quotation preparation time vs. agreed baseline
- Proposal consistency score (internal peer review) above agreed threshold
- Sales-team adoption rate above 70% within 6 weeks
- Zero unreviewed AI output reaching a customer
- Senior engineer escalation rate reduced by a measurable margin

#### Main risks to manage

- Inconsistent or outdated source documents degrade retrieval quality
- Sales engineers use AI draft without review, creating commercial or quality risk
- Pilot is treated as a tool rollout rather than a workflow change — adoption stalls
- Data access or security approval delays the start of the pilot

*Decision gate: Scale to the next product line only if the pilot shows measurable time savings, clear user acceptance, and no unacceptable quality or commercial risk findings at the 8-week review.*

---

## KPI IMPACT MAP

### Primary KPI: Quote cycle time

#### Efficiency

- Sales engineer time per quotation
- Engineering escalation rate
- Manual reporting effort (production)

#### Quality

- Proposal consistency and completeness
- Customer satisfaction with responsiveness
- Internal peer-review score on quotations

#### Strategic agility

- Time to onboard a new sales engineer to full productivity
- Speed of commercial repricing when input costs change
- Ability to scale sales capacity without proportional headcount

#### Growth

- Win rate on competitive RFQs
- Pipeline conversion rate
- Revenue from faster quote-to-order cycles

#### Risk

- Margin discipline on complex configurations
- Knowledge-concentration risk (key-person dependency)
- On-time delivery rate

---

## 5-PHASE ADOPTION ROADMAP

1

### Strategic AI North Star

*Establish a clear, measurable AI ambition tied to quote cycle time and win rate that the executive team owns and sponsors.*

For a configure-to-order precision manufacturer, the North Star should anchor AI to commercial velocity: faster, more profitable quotations. Frame it as a business outcome — not a technology programme — and secure executive alignment on the primary KPI before any tool is selected.

- Executive workshop to agree the primary commercial KPI and AI ambition
- Current-state baseline measurement of quotation cycle time
- Stakeholder mapping: who must sponsor, who must adopt
- Document the North Star in one page — shared with all pilot participants

**Output: One-page AI North Star aligned with executive team and commercial leadership**

2

### Opportunity discovery

*Identify and rank AI use cases by strategic fit, data availability and implementation risk — confirming that quotation support is the right first pilot.*

The most promising opportunity areas for this company are quotation support, engineering knowledge access and production planning intelligence — ranked by proximity to revenue and data availability. Discovery should validate which data is accessible before committing to a pilot.

- Map current RFQ and quotation workflow in detail — identify handoffs and bottlenecks
- Assess availability and quality of historical RFQ, documentation and project data
- Interview sales engineers and technical leads on pain points and workarounds
- Score and rank the three opportunity areas against agreed criteria

**Output: Prioritised opportunity portfolio with strategic-fit scores and data-readiness assessment**

3

### Use-case prioritisation

*Produce a signed-off pilot brief that defines scope, metrics, governance and the go / no-go decision*

*gate.*

Select the RFQ and quotation support pilot as the first use case. Define the exact scope — one product line, one customer segment — and set the baseline metrics before any tool is deployed. Do not proceed to tooling until the business case is written and approved.

- Draft and agree the pilot brief with sales, engineering and management
- Define and measure the cycle-time baseline for the selected product line
- Agree governance rules: human review requirements, data access, audit log
- Select and configure the AI tooling — scoped to pilot users only

**Output: One pilot brief with KPIs, success thresholds, governance rules and a clear decision gate**

4

#### **Pilot design and validation**

*Validate that AI-assisted quotation drafting reduces cycle time and is adopted by the sales team — and surface any workflow, data or governance issues before scaling.*

Run the RFQ support pilot with a small group of sales engineers on the selected product line. Measure against the agreed baseline at week 4 and week 8. Human review of every customer-facing output is non-negotiable throughout. Use the pilot to learn, not just to validate a hypothesis.

- Onboard pilot users with structured training on the workflow change (not just the tool)
- Run weekly check-ins to catch adoption blockers early
- Measure cycle time, adoption rate and quality score at weeks 4 and 8
- Document lessons learned and prepare the scaling recommendation

**Output: Pilot results report vs. agreed KPI baseline, with a scaling recommendation**

5

#### **Scaling and adoption**

*Embed AI-assisted quotation and knowledge retrieval across the commercial and engineering functions, with clear ownership, governance and a roadmap to the production planning intelligence use case.*

If the pilot validates the hypothesis, extend to the remaining product lines and integrate the engineering knowledge assistant as the second use case. Establish a centre of excellence with a named AI owner, a governance committee and a quarterly KPI review — this is when AI shifts from a project to an operating capability.

- Extend quotation support to all product lines with a structured change management programme
- Launch the engineering knowledge assistant, reusing the document retrieval infrastructure
- Establish quarterly AI KPI review with the executive team
- Define the data and integration requirements for the production planning use case

**Output: Scaling plan with governance structure, training programme and quarterly KPI rhythm**

---

## ASSUMPTIONS & VALIDATION QUESTIONS

### ASSUMPTIONS

1. The company sells configure-to-order or technically complex products where every quotation requires engineering input.
2. Quote cycle time is a meaningful executive KPI and a baseline measurement is feasible within the discovery phase.
3. Historical RFQ and quotation data (last 24 months) is accessible and sufficiently structured for a retrieval pilot.
4. AI tools are already approved or can be approved within the pilot timeframe (e.g. Microsoft 365 Copilot or an equivalent).
5. Data security and access controls for technical documentation can be confirmed before the pilot starts.
6. The system landscape (ERP, DMS, CRM, MES) is as described — integration scope requires validation in discovery.

### VALIDATION QUESTIONS

1. What is the current average cycle time from RFQ receipt to quotation submission, and how is it measured today?
2. Which data sources (historical quotes, technical documents, pricing guidelines) are accessible and in what format?
3. Where do sales engineers and technical staff lose the most time in the current quotation process?
4. What AI tools are already approved or in active use, and what governance constraints apply to new tools?
5. Who owns the decision to pilot and scale AI in commercial and engineering functions — and is that sponsor in the room?

---

### NEXT STEPS

This preview is a starting hypothesis. In a short expert session, AI4Leaders can validate the opportunity areas, sharpen the business case and define the most promising next step — so you leave with a decision, not a presentation.

#### Expert-guided AI strategy session

A focused 90-minute session with an AI4Leaders consultant to validate this hypothesis against your actual workflow, data landscape and commercial priorities — and leave with a pilot brief you can act on immediately.

#### Self-guided with Polaris

Use the Polaris platform to guide your team through the five-phase AI journey at your own pace, with structured milestones, decision frameworks and optional expert check-ins at key gates.