

autoagent.

AI-Native Business Decision & Execution System

Business Plan 2026



Strategic, Operating & Investment Version

INITIAL MARKET	CUSTOMER PATHS	MODEL
Germany	Existing businesses + entrepreneurs	Modular subscriptions + usage

Prepared for strategic planning, advisory review and investor discussions

Document Status & Interpretation

This document is designed to serve two purposes at the same time: (1) a practical operating plan for building and commercialising Autoagent, and (2) a defensible basis for investor discussions. It intentionally separates verified market facts from planning assumptions, targets and hypotheses.

Label	Meaning
Verified market fact	Supported by an external source listed in the Sources section.
Product principle	A design choice that guides product development.
Hypothesis	A belief that must be tested through pilots and market evidence.
Planning assumption	A temporary numerical input for modelling; not a forecast.
Target	An ambition used to guide execution; not a guarantee.

CORE POSITIONING
Autoagent turns business goals into measurable change.

Contents

- 1. Executive Summary
- 2. The Market Problem — The Decision-to-Outcome Gap
- 3. Why Now
- 4. Why Existing Tools Do Not Fully Solve It
- 5. Autoagent — The Solution
- 6. How Autoagent Works
- 7. Competitive Positioning — Why Autoagent
- 8. Product Architecture & Agent Marketplace
- 9. Target Customers & Market Entry
- 10. Business Model & Pricing
- 11. Go-to-Market Strategy
- 12. Technology, Security & Governance
- 13. Validation & Pilot Strategy
- 14. Market Opportunity
- 15. Roadmap
- 16. Financial Architecture & Funding Strategy
- 17. Investment Thesis
- 18. Risks & Mitigation
- 19. KPIs & Evidence Framework
- 20. Sources & Data References

1. Executive Summary

Businesses have more software, data, SaaS, automation and artificial intelligence than ever before. Yet having more tools does not automatically produce better business outcomes.

Autoagent is being designed to address the gap between business goals and measurable execution. Rather than beginning with a technology choice, it begins with the outcome a business wants to achieve, develops an understanding of the relevant business context, identifies and diagnoses problems or opportunities, compares possible interventions, coordinates controlled execution and measures what actually happened.

PROBLEM
Businesses increasingly have access to AI, software and automation, but often lack a reliable decision layer that can identify what should change, why, which intervention creates the highest expected value, and whether the change actually worked.

SOLUTION
Autoagent is an AI-native Business Decision & Execution System: Goal → Context → Observe → Detect → Diagnose → Prioritise → Select → Execute → Measure → Learn → Repeat.

The commercial entry model is modular. Customers can activate specialised agents according to their needs and pay subscriptions for the agents they use. Over time, the strategic value of the platform is intended to move beyond individual agents toward a shared Intelligence Layer that helps determine whether an agent, an existing SaaS product, a process change, human work, outsourcing or no change at all is the economically best intervention.

Autoagent is not positioned as another generic agent builder, automation platform, CRM or ERP. Existing systems can become execution infrastructure. The long-term ambition is to become the decision layer between business goals and execution.

Strategic question	Autoagent answer
What is the market problem?	The Decision-to-Outcome Gap: businesses struggle to identify the highest-value change and prove its impact.
What is the solution?	A business-first, AI-native system that diagnoses, decides, orchestrates execution and learns from measured outcomes.
Why should customers buy?	To reduce the cost of making the wrong business change — not simply to automate more work.
Why invest?	As AI, agents, SaaS and automation multiply, the coordination and decision problem becomes larger, creating room for a vendor-neutral intelligence layer.

2. The Market Problem — The Decision-to-Outcome Gap

2.1 The problem is larger than labour shortage

Labour scarcity is one pressure on German businesses, but it is not a sufficient foundation for the company thesis. In July 2026, 23.2% of German companies reported a shortage of qualified workers; the burden varied substantially by sector. Autoagent therefore treats labour shortage as an accelerator of the problem, not the problem itself. [S4]

The deeper issue is that every organisation operates with limited time, people, capital and management attention, yet often lacks the visibility and decision capability to allocate that capacity to the highest-value changes.

2.2 The gap in eight layers

Layer	Business problem
Capacity pressure	Time, labour, capital and management attention are limited.
Fragmentation	Data, processes, people and software are distributed across systems and teams.
Visibility gap	Management lacks a complete view of how work actually flows.
Diagnosis gap	Symptoms are visible, but root causes remain uncertain.
Decision gap	It is unclear which intervention has the best economics.
Execution gap	Even the correct intervention can be difficult to implement safely.
Measurement gap	Changes are deployed without rigorous evidence of business impact.
Learning gap	Results are not systematically converted into reusable knowledge for the next decision.

2.3 The problem chain

PROBLEM CHAIN People + Process + Technology → Fragmentation → Visibility Gap → Diagnosis Gap → Decision Gap → Execution Gap → Measurement Gap → Lost Productive Capacity → Higher Cost / Lower Growth / Slower Operations / Worse Customer Experience

2.4 Four words that explain the problem

SEE	UNDERSTAND	ACT	PROVE
What is happening in the business?	Why is it happening?	What should change?	Did the change create economic value?

CENTRAL MARKET STATEMENT

Businesses are not necessarily short of data, software or even AI. They are short of coordinated visibility, decision quality and execution.

3. Why Now

3.1 AI adoption is accelerating

In 2025, 26% of German enterprises with at least 10 employees used artificial intelligence. Adoption reached 23% among firms with 10–49 employees, 36% among firms with 50–249 employees and 57% among firms with at least 250 employees. [S1]

3.2 Adoption friction is not mainly about model availability

Among German enterprises that had considered AI but were not yet using it, the leading reported barriers were lack of knowledge (72%), unclear legal implications (62%), data protection and privacy concerns (60%), incompatibility with existing equipment/software/systems (45%), and difficulties with data availability or quality (44%). Cost was cited by 32%. [S2]

Barrier to AI adoption	Share
Lack of knowledge	72%
Unclear legal implications	62%
Data protection / privacy concerns	60%
Incompatibility with existing systems	45%
Data availability / quality	44%
Cost	32%

3.3 Digitalisation is active, but uneven

KfW reported that 30% of German Mittelstand companies had recently implemented digitalisation projects, with activity falling back to its pre-pandemic level. Digitalisation expenditure in the Mittelstand was reported at approximately €20.0 billion, while the gap between small and larger firms remained substantial. [S3]

WHY NOW

The first AI wave gave businesses intelligence. The next waves give them agents and automation. The growing problem is deciding what those systems should actually do — and proving that it was worth doing.

4. Why Existing Tools Do Not Fully Solve It

The market is crowded, and that is evidence of demand. Enterprise suites, automation platforms, agent builders, vertical AI products and specialised business tools increasingly offer agents, orchestration, integrations, governance, approvals and outcome tracking. Autoagent therefore does not treat those capabilities as unique selling propositions.

Category	Typical starting point	Strength	Unresolved decision question
Enterprise suites	Existing application / department	Deep installed data and workflows	Is this ecosystem the right intervention for the business problem?
Automation platforms	Workflow / process	Powerful execution and integration	What should be automated in the first place?
Agent platforms	Agent / task	Fast creation of AI workers	Which agent, if any, creates the highest business value?
Vertical AI	Industry workflow	Deep domain specialisation	How should cross-functional problems be prioritised across the whole business?
Consulting / services	Transformation project	Human diagnosis and change management	How can diagnosis, execution and learning become continuous and scalable?

STRATEGIC SHIFT
Execution is becoming easier. Decision quality before execution — and evidence after execution — becomes more valuable.

Autoagent should not attempt to rebuild every layer of the technology stack. Existing CRM, ERP, automation, productivity and data systems can be used as execution infrastructure. The platform should remain business-first and technology-agnostic wherever economically and technically feasible.

5. Autoagent — The Solution

DEFINITION
Autoagent is an AI-native Business Decision & Execution System.

The system is intended to help a business understand what should change, choose the best way to change it, execute the change within controlled boundaries, and prove the resulting business impact.

5.1 Business goal as the starting point

The user should not need to begin with a technology request such as “build a sales agent.” A more valuable starting point is an outcome: “increase revenue,” “reduce operating cost,” “increase capacity without adding headcount,” or “improve conversion.” Autoagent then investigates where the real constraint sits.

5.2 Business context

Within explicit permissions, Autoagent may use context from CRM, ERP, email, documents, project management, support, finance, HR, website analytics and internal systems. The purpose is not to become another data warehouse; the purpose is to construct a useful, as-current-as-practical model of the business for decision support.

5.3 Diagnosis before automation

A visible symptom such as delayed order processing does not imply that the company needs an order-processing agent. Root causes might include duplicate data entry, missing integration, approval bottlenecks, incomplete information, capacity imbalance or poor process design. The highest-value intervention may be process redesign, integration, training, a responsibility change, use of existing software, automation, an AI agent, custom software, outsourcing — or no new technology at all.

5.4 Solution engineering

Decision factor	Question
Business impact	How material is the expected outcome?
Cost	What is the full cost of implementation and operation?
Risk	What legal, operational, security or reputational risks exist?
Effort	How much organisational change is required?
Time-to-value	How quickly can value be tested and realised?
Expected ROI	What is the expected economic return and uncertainty?
Security	What permissions and controls are required?
Existing infrastructure	Can the current stack solve the problem without another tool?

PRODUCT PRINCIPLE
Choose the best solution for the business, not the solution the platform wants to sell.

6. How Autoagent Works

Stage	Purpose	Typical output
Business Goal	Define the outcome the company wants	Revenue, cost, capacity, service or growth target
Business Context	Understand constraints and current state	Systems, processes, KPIs, permissions, relevant data
Observe	Collect permitted evidence	Current behaviour and process signals
Detect	Find problems and opportunities	Bottlenecks, anomalies, missed opportunities
Diagnose	Distinguish symptom from root cause	Root-cause hypotheses with evidence and confidence
Prioritise	Rank what matters	Impact / cost / risk / effort ordering
Solution Engineering	Compare intervention options	Human, process, SaaS, agent, automation, outsourcing or none
Execute	Coordinate approved implementation	Tasks performed by people, software

and agents

Measure	Compare before and after	Outcome metrics and actual cost
Learn	Store what worked in what context	Reusable decision evidence
Repeat	Identify the next highest-value change	Continuous improvement loop

6.1 Controlled autonomy

Level	Mode	Description
0	Observe	Read-only observation and analysis.
1	Recommend	Produces recommendations; no execution.
2	Prepare	Prepares actions or artefacts but does not execute.
3	Execute with approval	Executes after explicit human approval.
4	Autonomous within authorised scope	Executes independently only inside pre-approved boundaries, budgets and permissions.

6.2 Measure & learn

Autoagent should not measure success by the number of agents created or workflows automated. Success should be measured by business outcomes such as processing time, error rate, response time, throughput, employee workload, cost, revenue, conversion, customer experience and productive capacity.

NORTH-STAR PRODUCT CONCEPT
Productive Capacity Unlocked

7. Competitive Positioning — Why Autoagent

7.1 Competitive positioning

Agent creation, multi-agent orchestration, workflow automation, RAG, MCP, integrations, governance and cross-system execution are increasingly standard capabilities. Autoagent should therefore compete one level higher: on decision quality and economic evidence.

POSITIONING
Business-first. AI-native. Vendor-neutral.

7.2 The competitive white space

Capability	Autoagent focus
Problem discovery	Which problem or opportunity is economically meaningful?
Root-cause diagnosis	Why is the problem happening?
Intervention selection	Which intervention has the highest expected value?
Controlled execution	How can the change be implemented safely?
Outcome learning	Did the result justify the decision, and what can be learned for next time?

7.3 Why customers should choose Autoagent

CUSTOMER PROMISE
Autoagent is designed to reduce the cost of making the wrong business change.

Businesses lose value not only through inefficient operations, but also through incorrect transformation decisions: unsuitable software, unnecessary integrations, poorly chosen automation, low-value consulting projects, agents without a valid use case, or changes that consume management attention without measurable return.

Before execution, Autoagent asks: “What is the highest-value change?” After execution, it asks: “Did it create the expected value?”

7.4 Potential defensibility

The defensibility thesis should not depend on a particular model, prompt or agent framework. Those layers are likely to become increasingly substitutable. A stronger potential moat is the structured relationship between Business Context → Problem → Root Cause → Predicted Intervention Impact → Chosen Intervention → Actual Outcome.

MOAT HYPOTHESIS
More customers → more diagnosed problems → more interventions → more measured outcomes → better goal-to-solution decisions → higher customer value → stronger retention and expansion.

Important: this is a moat hypothesis, not yet a proven moat. It must be validated through customer pilots and measurable outcomes.

8. Product Architecture & Agent Marketplace

Layer	Role
Autoagent Workspace	Business profile, goals, users, permissions, billing, audit, agent management.
Specialised Agent Marketplace	Modular agents that customers activate according to their needs.

Autoagent Intelligence Layer	Diagnosis, prioritisation, solution selection, orchestration, measurement and learning.
Governance & Security	Permissions, policy controls, audit, cost limits, approval logic and risk controls across all layers.

8.1 Initial agent portfolio

Agent	Primary outcome
Lead Intelligence	Identify, qualify and prioritise genuine sales opportunities rather than simply finding contacts.
Competitor Intelligence	Analyse public market signals, positioning, offers and advertising patterns.
Customer Voice	Extract recurring pains, objections and needs from permitted sources and customer data.
Website Intelligence	Evaluate digital presence against business goals and prioritise improvement opportunities.
Website Builder	Create or improve websites according to the approved business objective.
Social Media Manager	Select relevant channels, plan content, prepare/schedule approved activity and measure results.
Video Repurposing	Convert long-form video into short-form social assets.
Browser Execution	Perform permitted repetitive research and browser-based tasks.
Market Research	Research market structure, customer segments and opportunity hypotheses.
Business Planning	Support founders through structured business planning and launch preparation.
Industry-specific agents	Future vertical agents with domain KPIs, workflows and compliance rules.

8.2 Lead Intelligence as a proof of philosophy

A list of 600 businesses is not 600 qualified opportunities. A company may have a weak website yet no meaningful need for a new one. Lead Intelligence should therefore move from contact discovery toward opportunity diagnosis: Is there evidence of a real problem? Is it material? Is the company a fit? What offer, if any, fits the problem? Is outreach appropriate and legally compliant? A correct “no opportunity” decision is itself valuable because it prevents wasted sales effort.

9. Target Customers & Market Entry

Autoagent has two entry paths that share a common platform but should be validated separately. The company should not assume both paths will receive equal resources; resource allocation should follow evidence from conversion, retention, willingness to pay and unit economics.

9.1 Path A — Business Growth

Target: existing small and mid-sized businesses that have enough operational complexity to benefit from diagnosis and automation, but do not necessarily have large internal transformation teams.

- Primary outcomes: higher productive capacity, lower operating cost, more qualified revenue opportunities, faster execution and better customer experience.
- Initial use cases: lead intelligence, competitor/customer intelligence, digital presence, social media, process opportunity discovery.
- Potential sectors for early testing: professional services, business services, trade/distribution, e-commerce, logistics, information & communication and selected operational SMEs.

9.2 Path B — Business Creation

Target: founders and small teams converting an idea into a commercial business. The proposition is not “AI builds a company automatically.” The proposition is to reduce fragmentation and execution cost across validation, planning, website creation, market research, initial go-to-market and early operating workflows.

- Initial outcomes: faster movement from idea to real market test, lower coordination burden and a more coherent launch process.
- Entry products: market research, business planning, website creation and launch-oriented agents.
- Success metric: time and cost to reach a genuine commercial experiment, not the number of generated documents.

10. Business Model & Pricing

10.1 Revenue architecture

Revenue layer	Model
Autoagent Intelligence	Base subscription for workspace, context, decision support and shared platform capabilities.
Specialised agents	Independent monthly subscriptions selected by the customer.
Usage	Included allowance plus metered overage where compute/data intensity requires it.
Advanced integrations	Higher-tier or enterprise pricing for complex integrations and data connectivity.
Implementation	Optional professional services, clearly separated from software recurring revenue.
Custom / private deployment	Contract pricing for larger customers with specific security or deployment requirements.

10.2 Pricing hypothesis

Plan	Target customer	Planning range
Solo	Founder / very small business	€29–€49 / month base platform
Business	Small company	€99–€199 / month base platform
Growth	Growing SME	€299–€599 / month base platform
Enterprise	Complex organisation	Custom pricing

Planning assumption: these ranges are not final prices. They must be tested against cost-to-serve, willingness to pay and competitive alternatives.

Agent subscriptions are layered on top. A customer might begin with one agent and expand as value is proven. The central expansion KPI is not the number of agents offered, but the number of economically justified agents actively used per customer.

EXPANSION MODEL

Land → Prove Value → Expand

10.3 Key SaaS expansion metrics

Metric	Why it matters
Agents per customer	Shows whether the modular ecosystem creates cross-sell value.
Second-agent adoption	Early evidence that the platform creates value beyond a single tool.
Net Revenue Retention	Measures retention plus expansion within the installed base.
Gross margin per agent	Prevents popular but compute-heavy agents from hiding poor economics.
CAC payback	Tests whether acquisition cost is recovered quickly enough to scale.
Churn	Tests whether the system becomes part of ongoing operations.

11. Go-to-Market Strategy

Autoagent should not initially be marketed as a broad “AI Business Operating System.” Customers buy specific outcomes, not architecture diagrams. The entry offer must therefore be narrow, measurable and connected to a painful business problem.

11.1 Entry offers

- Find high-value sales opportunities — not just contact lists.

- Discover where operational capacity is being lost.
- Understand customer and competitor signals before launching a campaign.
- Launch a business with a coordinated set of research, planning and execution tools.

11.2 Early GTM motion

1. Founder-led sales to a tightly defined set of design partners.
2. Paid pilots instead of free “interest” wherever possible.
3. Use ITmercer service capabilities to open doors, but separate recurring software economics from services.
4. Capture structured evidence from every pilot: problem, context, intervention, cost, outcome and decision quality.
5. Convert successful pilots into repeatable packages and reference cases.

11.3 Why the Lead Finder strategy changes

The acquisition engine should not optimise for the number of businesses found or emails sent. It should optimise for the number of qualified opportunities with evidence of a meaningful problem, a relevant solution and an appropriate outreach path. This reduces wasted sales effort and aligns the GTM engine with Autoagent’s broader philosophy: diagnose before executing.

12. Technology, Security & Governance

12.1 Architectural principles

- Modular agents connected to a shared business context layer.
- Model-agnostic routing where practical to control performance, cost and dependency risk.
- Explicit permission boundaries for every data source and execution capability.
- Evidence-linked recommendations, confidence indicators and human review for material decisions.
- Audit logs for agent actions, data access, approvals and execution history.
- Cost controls and usage budgets at account and agent level.
- Ability to integrate existing tools instead of forcing replacement.

12.2 Privacy and regulation

For the European market, privacy, security and AI governance must be product features rather than afterthoughts. Most AI Act rules became applicable on 2 August 2026, while certain high-risk requirements follow later timelines. Autoagent should classify risk by use case, apply human oversight and appropriate controls, and obtain specialist legal review for regulated or high-impact deployments. [S5]

- Data minimisation and purpose limitation.
- Tenant isolation and role-based access.
- Retention controls and deletion processes.
- No use of customer data for public model training without an appropriate legal and contractual basis.
- Human approval for sensitive actions.
- Auditability, rollback and incident handling.
- Compliance-by-design for future vertical agents.

13. Validation & Pilot Strategy

The business plan deliberately distinguishes differentiation hypotheses from proven advantages. Autoagent should not claim a durable moat until the following questions are answered with real customer evidence.

Hypothesis	Evidence required
The system finds a meaningful cross-functional problem that existing tools or the customer did not easily identify.	Paid pilot + before/after diagnostic comparison.
It recommends a better intervention than a simple technology-first approach.	Alternative comparison and customer decision review.
Expected impact can be estimated with useful accuracy.	Prediction vs. realised outcome.
Execution can occur with appropriate permissions, approvals and rollback.	Audit data, incidents and execution success rate.
Business outcomes can be measured credibly.	Baseline, cost, KPI and measurement period.
Customers return after the first problem is solved.	Renewal, retention and second-agent adoption.

13.1 Pilot design

- 6. Select one business problem with a clear baseline and measurable cost or outcome.
- 7. Document the current process and the customer’s existing alternative.
- 8. Run Autoagent diagnosis and intervention selection.
- 9. Execute at the safest appropriate autonomy level.
- 10. Measure outcome over an agreed period.
- 11. Record total implementation and operating cost.
- 12. Compare expected vs actual impact.
- 13. Decide whether the customer would renew, expand or discontinue.

14. Market Opportunity

The preferred market-sizing approach is bottom-up rather than relying only on broad “AI agent market” forecasts. The initial planning market is German SMEs and selected entrepreneurial users; the product architecture can later support additional European markets.

14.1 Market-sizing method

Layer	Method
TAM	All organisations and entrepreneurs that could economically use the platform or specialised agents over time.
SAM	Initial sectors/geographies where the current product and GTM can realistically serve customers.
SOM	A realistic customer and ARR target over 3–5 years based on acquisition capacity, pricing, churn and expansion.

For investor-grade use, the next iteration of the financial model should connect SAM and SOM directly to verified company counts by size and sector, an explicit ARPA assumption and a customer-acquisition capacity model. This document therefore avoids presenting a speculative market number as a proven forecast.

MARKET-SIZING PRINCIPLE

Use market reports for context; use customer count × realistic ARPA × achievable penetration for the operating model.

15. Roadmap

Delivery approach: an ambitious 12-week target for a focused MVP and first controlled pilot, subject to team capacity, integration access, data quality, governance and customer availability. This is a planning target, not a commitment that the full platform or commercial scale will be completed in 12 weeks.

Stage / indicative window	Objective	Core deliverables	Evidence gate
Weeks 1–2 · Define	Scope a narrow, paid-value use case	Pilot customer, outcome KPI, baseline, permission model	Problem and buyer validated
Weeks 3–6 · Build MVP	Launch a working first use case	Workspace, 1–2 agents, business context, basic governance	End-to-end flow demonstrated
Weeks 7–10 · Pilot	Operate with real design partners	Approvals, diagnosis, selected actions, audit and measurement	Verified baseline vs. outcome
Weeks 11–12 · Decide	Improve and release next iteration	Customer feedback, unit-cost analysis, prioritised backlog	Paid continuation or explicit pivot decision

OPERATING PRINCIPLE

Sell → Measure → Learn → Build deeper. Do not build the full operating system before the market proves the next layer.

16. Financial Architecture & Funding Strategy

16.1 Unit economics first

Metric	Definition / target direction
ARPA	Monthly recurring software revenue per account; should expand as justified agents are added.
AI / API cost	Direct cost of model usage, web execution, data providers and infrastructure by agent.
Software gross margin	Target should be high enough to support SaaS scaling; calculate after all direct variable delivery costs.
CAC	Fully loaded acquisition cost per new paying customer.
CAC payback	Months required for gross profit to recover acquisition cost.

NRR	Retention plus expansion revenue from existing customers.
Burn & runway	Cash consumption and months of operating capacity under base/downside cases.

16.2 Milestones & Evidence — Not a Forecast

Milestone	Target evidence (not revenue forecast)	Decision
Weeks 1–2	Clear ICP, high-value problem, agreed baseline	Proceed only if measurable and economically relevant
Weeks 3–6	Runnable, governed MVP and first usable output	Proceed only after technical validation
Weeks 7–10	Design-partner usage and measurable pilot outcomes	Compare expected vs. observed impact
Weeks 11–12	Renewal or pilot purchase signal and unit cost	Scale / improve / stop based on evidence
After first pilots	Repeatability, retention, agent economics	Expand incrementally; no fixed completion date

The 12-week milestones replace the previous five-year ARR projection. There are no verified revenue forecasts at this stage. A driver-based forecast can be added once paid-pilot data on conversion, pricing, costs and retention becomes available.

16.3 Funding hypothesis

Rather than anchoring prematurely on valuation, financing should be milestone-based. A potential planning hypothesis is €1.5M–€2.0M of pre-seed / seed-bridge capital to fund a milestone-based operating runway, subject to validation, subject to legal structure, current cash position, team plan and pilot evidence.

Use of funds	Planning allocation
Product & engineering	45%
Sales & go-to-market	25%
Security / compliance / data	15%
Infrastructure & AI cost	10%
Legal / administration	5%

FUNDING OBJECTIVE
The purpose of the round is not to “finish Autoagent.” It is to prove a repeatable business model and the core decision-to-outcome thesis.

17. Investment Thesis

The investment thesis is not “AI is growing.” That is true but insufficient. The stronger thesis is structural: as companies adopt more AI models, agents, SaaS products, automation systems and APIs, they face a growing coordination and decision problem.

INVESTMENT THESIS
As execution becomes cheaper and more abundant, the scarce capability becomes deciding what should change, choosing the economically best intervention, executing it within governed boundaries, and learning from the result.

Autoagent is designed to occupy that layer. If the company can prove that its diagnosis and intervention selection lead to better measurable outcomes, and if customers retain and expand after the first use case, the business can develop recurring software revenue, expansion revenue and a growing evidence base about which interventions work in which business contexts.

17.1 Why this could become a large business

- The problem exists across functions and industries, while the initial go-to-market can remain narrow.
- Modular subscriptions lower entry friction and create expansion potential.
- A shared business context can increase the value of additional agents over time.
- Vertical packs can add domain depth without abandoning the common platform.
- Measured intervention-outcome data can create a learning asset that is harder to reproduce than prompts or agent workflows.
- Existing software ecosystems can become execution infrastructure rather than direct replacement targets.

18. Risks & Mitigation

Risk	Why it matters	Mitigation
Crowded AI market	Agent and automation features are rapidly commoditising.	Compete on decision quality, business evidence and vertical learning rather than generic features.
Overlapping visions	Other products may pursue similar “business operating system” narratives.	Validate faster, narrow the wedge, build proprietary evidence and vertical depth.
Low willingness to pay	The problem may sound important but not create budget.	Use paid pilots and explicit ROI/cost-of-problem qualification before broad build-out.
Overbuilding	A broad platform can consume capital before demand is proven.	Wedge-first roadmap with phase gates tied to revenue and evidence.
AI error / hallucination	Incorrect recommendations can damage trust.	Evidence-linked outputs, confidence levels, human approval and limited autonomy.
Data / privacy / legal risk	The platform may process sensitive business or personal data.	Privacy-by-design, legal review, access controls, audit, retention and risk classification.
High variable AI cost	Heavy research/browser/data agents can compress margins.	Model routing, quotas, metered usage, caching and per-agent gross-margin

		tracking.
Long B2B sales cycles	Complex enterprise selling can delay learning.	Begin with modular agents and SME/design-partner pilots; enterprise later.
Conflict with vendor neutrality	If revenue depends only on selling agents, recommendations may appear biased.	Charge for the Intelligence Layer and disclose commercial incentives clearly.
Weak measurement	Before/after changes may be affected by other factors.	Define baseline, measurement period, total cost and appropriate comparison methods in every material pilot.

19. KPIs & Evidence Framework

Area	Primary KPIs
Demand	Paid pilot conversion, qualified opportunity rate, sales-cycle length.
Activation	Time to first useful outcome, successful completion of core task.
Product value	Measured time saved, cost reduced, revenue opportunity created, productive capacity unlocked.
Quality	Recommendation acceptance, human correction rate, agent success rate, incident rate.
Economics	ARPA, gross margin by agent, CAC, CAC payback, contribution margin.
Retention	Logo churn, revenue churn, NRR, repeat usage.
Expansion	Second-agent adoption, active agents per customer, expansion MRR.
Decision thesis	Expected vs actual intervention impact; proportion of recommendations that do not require new technology.
Trust & governance	Approval rate, permission violations, rollback events, audit completeness.

19.1 Evidence ladder

Question	Evidence needed before making the claim
Is the market problem real?	Market research + interviews + paid problem validation.
Does the solution work?	Prototype + architecture + measured pilot outcomes.
Why Autoagent?	Alternative comparison + workflow evidence + ROI / decision-quality evidence.

Can this become a large business?

Repeatable acquisition + retention + expansion + unit economics + credible market model.

20. Sources & Data References

The main narrative intentionally avoids competitor names. Competitive research is maintained as a separate internal benchmark. The following sources support the market data used in this business plan.

S1 — German Federal Statistical Office (Destatis), ICT in enterprises: AI use in German enterprises, reference year 2025. Key figures used: 26% overall; 23% for 10–49 employees; 36% for 50–249; 57% for 250+.

<https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Unternehmen/IKT-in-Unternehmen-IKT-Branche/IKT-U-Erhebung/info.html>

S2 — Destatis, reasons against the use of artificial intelligence technologies, reference year 2025. Key figures used: 72% lack of knowledge; 62% unclear legal implications; 60% privacy/data protection; 45% incompatibility; 44% data availability/quality; 32% cost.

<https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Unternehmen/IKT-in-Unternehmen-IKT-Branche/Tabellen/ikti-gegen-nutzung-kuenstliche-intelligenz.html>

S3 — KfW Research, Digitalisation Report Mittelstand 2025, published 23 April 2026. Key figures used: 30% recently conducted digitalisation projects; approximately €20.0bn digitalisation expenditure; persistent digital gap by firm size.

https://www.kfw.de/%C3%9Cber-die-KfW/Newsroom/Aktuelles/News-Details_891136.html

S4 — ifo Institute, Business Survey, 21 August 2026. Key figure used: 23.2% of German companies reported a shortage of qualified workers in July 2026; 64.7% in legal/tax/auditing.

<https://www.ifo.de/en/facts/2026-08-21/almost-quarter-companies-germany-have-shortage-skilled-workers>

S5 — Council of the European Union, Artificial Intelligence Act timeline. Key point used: majority of AI Act rules applied from 2 August 2026, with later dates for specified high-risk provisions.

<https://www.consilium.europa.eu/en/policies/artificial-intelligence-act/timeline-artificial-intelligence/>

20.1 Final strategic statement

NORTH-STAR QUESTION

What should this business change — and what is the best way to make that change happen?

Autoagent should not be judged by how many agents it can create. It should be judged by whether it can repeatedly identify higher-value change, coordinate execution safely and produce measurable business evidence. That is the product thesis, the operating plan and the investment case.