What if your software acted as a proactive partner? How-to guide to Agentic AI

By Allegra Patrizi
allegrapatrizi.com

What if your software acted as a proactive partner?
How-to guide to Agentic AI

By Allegra Patrizi, *allegrapatrizi.com*, May 2025

Introduction: Care, Dare, Share

Care. Dare. Share. Why I started over to build what's Next

I've spent over two decades in the corporate world - at the highest levels. Under my leadership, the first fully digital banks in the Netherlands was developed and thrived to #4 in the country, and Virgin Money was celebrated for award-winning AI innovation. I know how to manage complexity, drive large-scale change, and deliver outcomes in highly regulated environments.

But despite all that, I felt something was missing.

I saw how quickly the world was shifting - especially with AI - and how far many big companies still are from truly keeping up. Customers today are using AI-powered tools in their daily lives. They expect speed, personalization, and magic. And they don't understand why their banks or insurers or many other providers, for that matter, can't offer the same. Frankly, sometimes I don't either.

So I decided to step away from the boardrooms and go back to the roots. I wanted to learn, build, and understand the new wave - not as an executive giving big directions from the top, but as a practitioner, builder, and entrepreneur starting from zero.

But I didn't just do it for the technology and the understanding. I did it because it aligned with how I've always believed leadership should work.

Care. Dare. Share.

Care means putting people first - customers, teams, communities. Understanding what they need, what they dream of, what frustrates them. It's what guided me through every transformation I've led - and it's why I wanted to get closer to them.

Dare is about having the courage to challenge the status quo. To test new ideas. To choose sympathetic bold over hesitating bland. In this new AI era, there's so much noise - but also so much possibility. I wanted to see what's actually possible, what's truly risky, and what just needs a leader willing to start.

And **Share** is the commitment to transparency. To take accountability when things go well or wrong. To open up about the process, not just the outcome. To treat learning not as failure, but as fuel. As I share this experiment, I'm doing so publicly, openly - because I believe that's the only way to move forward together.

If we want our companies to be more agile, more human, more innovative - then leaders need to step into the unknown, not just talk about it. We need to get curious again.

I decided to walk my talk: **Care, Dare, and Share**.

Purpose and structure of this guide

Artificial Intelligence (AI) has evolved from simple data-processing algorithms to sophisticated systems capable of autonomous decision-making, goal-directed behavior, and proactive collaboration. This advanced form of AI, known as Agentic AI, is transforming industries. Agentic AI systems think, act, and learn like digital colleagues, executing tasks independently and adapting to dynamic environments.

This hands-on guide explores what defines Agentic AI, why it matters, and how to create robust, ethical, and effective agentic systems. First, we start with a few foundational concepts then we go into the specifics of how we applied things in "real life" in our micro-company, running three business lines.

Our approach is not a polished enterprise-grade playbook. But it shows what's possible today when you combine creativity, curiosity, and a can-do mindset. This has big implications:

1. For smaller organisations: it shows how many processes could be automated, from the value creating commercial and marketing ones (e.g. lead generation, social marketing, price setting, etc.) to those tedious but necessary administrative processes (VAT declaration, reporting, etc.), bringing benefits in terms of increased revenues and/or reduced profits and /or greater accuracy
2. For large organizations, it offers a glimpse of how lightweight, fast-moving experiments (of course with appropriate guardrails) can seed scalable, future-ready solutions. It opens the possibility for innovative, very targeted deployments done by and with business people, close to where it matters the most for a given team or process. And even more importantly, it shows how, beyond the obvious use-cases of AI for commercial and cost/operational purpose, possibly one of the biggest use-cases is

around *amplification* of key managers – Ie finding ways to provide more leverage and support to those 50-200 people that in every organisation are the absolute "heroes" or pillars of the organisation, the go-to people for everything, the unavoidable critical nodes in your value chain, but who are also overwhelmed, overworked and often become a bottle-neck.

Our approach was simple: one person, no coding background, four AI agents - and together, we're developing or running three business lines.

1. AI for Senior Leaders

Strategic insights and decision support for executives. The first product, Claridora, is already live and growing. More tools are on the way.

2. AI for Small Businesses

Automation and performance enhancing tools that cut costs, boost sales, and simplify admin work - fast, affordable, and practical.

3. AI for Everyday Life

Targeted AI helpers for daily routines for niche segments - social, entertainment, etc.

This guide explains what Agentic AI is and how to build it, using my experience running a micro-company with AI agents as CFO, CTO, CMO, and CLO. It's designed for:

- **Enterprise Leaders**: To boost efficiency and innovation.
- **Product Teams**: To create intelligent, user-focused features.
- **Developers**: To understand agentic design fundamentals, if you are not an expert.

Foundational concepts

Imagine software that doesn't just follow instructions but takes initiative, makes decisions, and collaborates like a digital teammate. This is Agentic AI, where artificial intelligence evolves from a passive tool into an autonomous partner.

What is Agentic AI?

Agentic AI systems perceive their environment or are receptive to triggers, process inputs, make informed decisions, and take actions to achieve specific objectives without constant human intervention. Unlike traditional AI, which operates within fixed rules or predictive models, or chatbots that provide scripted responses, Agentic AI is adaptive, proactive, and goal-driven, emulating human-like actions.

Key characteristics

- **Proactivity**: Initiates tasks, e.g., resolving issues before escalation.
- **Contextual awareness**: Retains history for informed decisions.

- **Goal-oriented**: Plans actions to meet objectives, e.g., optimizing workflows.
- **Adaptability**: Learns from feedback to improve.

Difference between Agentic AI and traditional AI and Chatbots

- **Autonomy:** Acts independently vs relies on user prompts
- **Memory:** Retains context and history vs blank sheet
- **Goal-Setting:** Plans toward objectives vs no outcome awareness
- **Flexibility:** adapts dynamically v static learning

While chatbots excel at reactive, scripted interactions, Agentic AI agents act as digital teammates, orchestrating workflows and collaborating with humans or other systems.

Why develop Agentic AI?

The benefits of Agentic AI are pretty obvious:

- **Proactive Problem-Solving**: Addresses issues early, e.g., IT faults.
- **Enhanced Productivity**: Automates tasks, freeing humans for strategy.
- **Scalability**: Handles complex processes in healthcare, logistics, etc.
- **Personalized Experiences**: Delivers context-aware interactions.
- **24/7 Availability**: Ensures continuous operations.
- **Data Insights**: Drives smarter decisions.

What types of Agentic AI are there?

Agentic AI can be categorized by their level of sophistication and the type of broad use one makes of them.

In terms of sophistication levels, we go from the most basic level of automation with very limited room for manoeuvring to large span of control agents:

Level 1—Reactive Agents

Reactive agents respond to immediate inputs without memory or past knowledge, following preset rules. Seen in basic chatbots that handle predictable tasks, reducing repetitive business work.

Level 2—Task-Specialized Agents

These agents excel in specific domains, like fraud detection or medical diagnostics, using rule-based systems. They optimize tasks such as e-commerce recommendations or delivery routes, requiring clear problem definitions and expert collaboration.

Level 3—Context-Aware Agents

Context-aware agents process complex, dynamic data, adapting to new scenarios using machine learning or LLM advancements.

Level 4—Socially Savvy Agents

These agents understand (basic) human emotions and mental states, drawing from cognitive psychology's theory of mind. They enhance customer interactions and negotiations, requiring investment in affective computing and ethical guidelines.

Level 5—Self-Reflective Agents

Self-reflective agents analyze their decisions and improve autonomously. They need robust feedback systems.

Level 6—Generalized Intelligence Agents

Artificial general intelligence (AGI) systems handle diverse tasks across domains. Recent language model progress suggests AGI potential, requiring businesses to integrate varied data for strategic alignment.

Level 7—Superintelligent Agents

Superintelligent agents surpass human capabilities in many domains. They could solve global issues but pose ethical dilemmas, demanding a reimagining of business and societal frameworks.

With today's tools and technologies, the bulk of the more advanced Agentic AIs sit at level 3-' with some rare forays into 5. Levels 6 and 7 are still in the dream camp, but things are changing rapidly.

The higher the level, the higher the risk of data leaks, manipulation of the Agentic AI, reputational manage in case of fail and the need for stringent ethical and legal guardrails, cyber security systems, use of multiple different "brains" to avoid crowd-thinking and hallucinations.

In terms of broad categories of use:

🔍 Agentic RAG - Knowledge agents that don't just retrieve info, but vet it, reason over it, and ground it in real context.

🛠️ Workflow Agents - Orchestrators that automate multi-step business processes, triggered by events, APIs, or UI actions.

💻 Coding Agents - From debugging to full repo reasoning, they're accelerating dev teams and solo builders alike.

🧰 Tool-Based Agents - Purpose-built, reliable task performers embedded directly into your stack.

🖱️ Computer Use Agents - Agents that act like users: navigating UIs, filling forms, clicking buttons - no API needed.

🗣️ Voice Agents – AI on the phone line, transforming support, sales, and operations with natural dialogue.

Core elements of Agentic AI

Agentic AIs are built on 6 foundational components to enable autonomy and effectiveness:

1. **Perception**: Captures inputs via sensors, APIs triggers, or NLP.
   Value: Enables proactiveness and real-time communication with the world.
2. **Brain**: Uses LLMs or reinforcement learning for optimal actions.
   Value: Enables autonomy and supports thoughtful responses.
3. **Memory**: Stores interactions via short term or long term memory systems (and within these there can be separate memory functions such as the simple immediate context memory in-built in most LLMs today to much more refined and complex semantic memory, interaction memory, etc.).
   Value: Ensures continuity and enables learning.
4. **Action**: Executes tasks through tools by using APIs or actuators.
   Value: Delivers tangible results.
5. **Ethics**: Aligns with societal and legal standards.
   Value: Builds trust.
6. **Feedback loops**: Analyzes past performance and identifies possible enhancements – to be then often submitted to human supervisor for approval.
   Value: Improves continuously and fine tunes behavior.

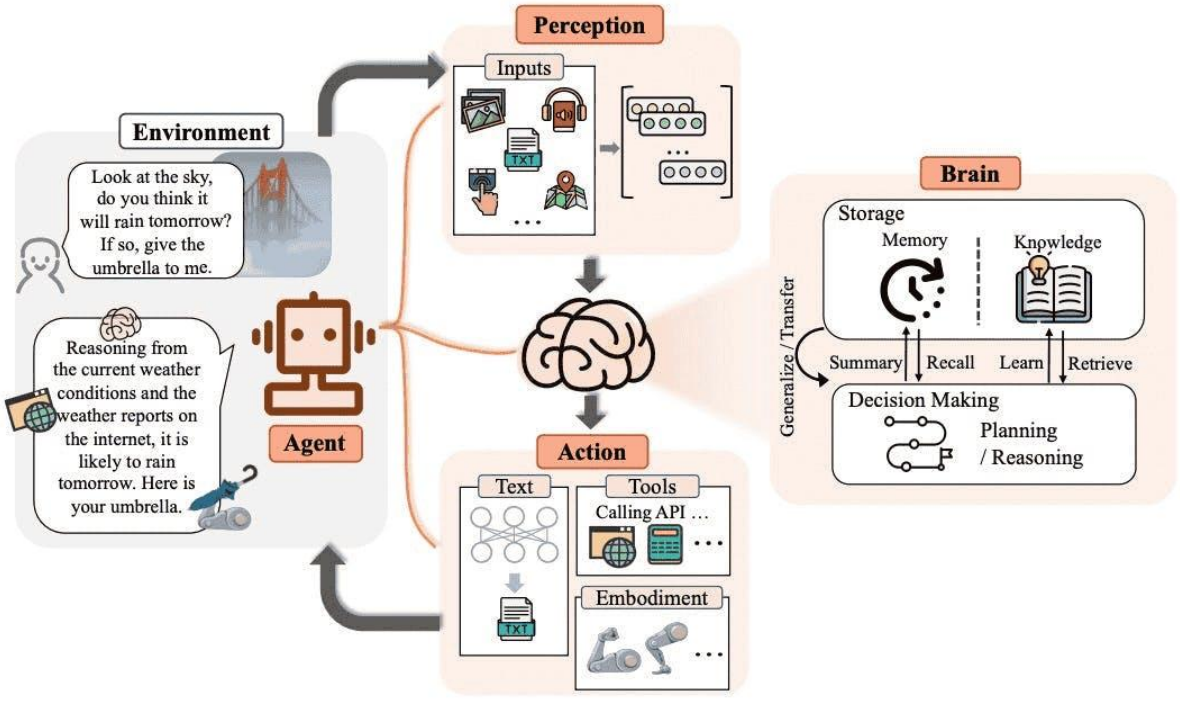

AI agent core components – Credits: from *Intro of AI agent & AI agent project summary* by Henry Heng Luo Medium.com

Where to deploy Agentic AI?

The choice of where to deploy Agentic Ais is critical for both small and large companies to maximize benefits and minimize risks.

In small companies, where resources are limited, deploying agentic AI in high-impact areas like customer service or inventory management can enhance efficiency, reduce costs, and improve customer satisfaction without overextending budgets. However, poor choices - such as automating complex, human-centric tasks like creative marketing - can lead to errors, customer alienation, and wasted investment.

For large companies, the stakes are higher due to scale. Deploying agentic AI in areas like supply chain optimization or data-driven decision-making can streamline operations and drive competitive advantage. Yet, missteps, such as automating sensitive HR processes without oversight, can result in ethical issues, bias amplification, or regulatory violations.

Across both scales, choosing deployment areas requires aligning AI capabilities with business goals, ensuring robust oversight, and prioritizing tasks where automation enhances human efforts rather than replaces critical judgment. Thoughtful selection ensures AI drives growth, trust, and innovation while avoiding costly pitfalls.

My advice? Don’ t start with an overly ambitious approach.

Why? Because:

1. Despite the hype, most companies are just getting started. You have time. This is a marathon, not a sprint.

2. Managers and advisors don’t yet have the hands-on intuition needed for realistic, impactful large scale business cases. Right now, most "grand business cases" are just assumptions dressed up with fancy words, tools and numbers.

3. 70% of long-term success will come down to one thing: internal adoption by humans. If employees or customers don't embrace it, even the best Agentic AI will remain an expensive piece of useless software.

So our advice is to start small and pragmatic. Pick a few simple, clear use-cases (5-20). Give yourself 6 months. Get wins. Build trust inside your organization.

Here’s our simple first-pass framework. For each potential use-case idea, answer YES or NO:

1. **Impact:** Will it grow revenue, cut costs, improve quality, or solve a real daily pain point (e.g. bottle necks, constraints, relief of pressure on a critical role, tedious process or other)? (Yes +2, No -2)

2.**Complexity:** Can you buy/subscribe a tool that's >90% trained? (Yes +1, No -2)

3. **Team**: Do you have the multi-disciplinary team (from business AND tech/IT AND compliance/legal) to develop it and can they be removed from what they are currently doing? (Yes +1, No -2)

4. **Data**: Is the specific data you need for it >80% clean and structured and can it be ringfenced enough to avoid risk of large data leaks or data corruption? (Yes +1, No -6)

5. **Integration**: Does it require deep integration into the legacy system that can not be easily robotised with scraping tools? (Yes -4, No +1)

6. **Reputation**: If the AI tool hallucinates, or performs its task not fully consistently, is this going to be a significant embarrassment? (Yes -5, No +1)

7. **Bias/Explainability:** Is it necessary to be able to explain every outcome (e.g. due to regulatory)? (Yes -5, No 1)

8. **Influence**: If it succeeds, will it be relevant and helpful for more than 30% (employees, or departments, or countries) of our company? ? (Yes +2, No -2)

Once you have your first ranking, dive deeper into the top candidates, using these same dimensions but with more nuance on the pros and cons/metrics.

The whole idea is to adopt a progressive approach. **Build. Succeed. Showcase. Repeat.**

Principles in building Agentic AI

Once you have selected the use cases that make sense for you, it is essential to apply these principles in building the right Agentic AIs.

We have developed these principes based on our extensive experience, both in the large companies I have led, as well as the consulting situations I have been involved in and the recent development of the 4 AI agents described in the second part of this document.

1. **The more basic the better:** Adopt the lowest possible level of Agent that can do the task (if Level 1 can do it go for it)
2. **Don't go Swiss-knife:** give each agent a small number of tools, ideally 1-2 to avoid giving too much latitude of choice and therefore potential error. This entails that the tasks done by one human may need to be split among several Ai Agents
3. **Trust but verify**: ensure multiple intermediate end-products so that you can monitor at each step of the process the decisions taken, the actions performed, the data transformations etc.
4. **The more autonomous and capable the Agentic AI, the harder the pre-work will be**: Agentic AIs are like uber-motivated toddlers, with infinite stamina and memory. Ie their intellectual capabilities are far beyond their experience and common sense, as if a child was born reading, writing and with perfect knowledge of all of human knowledge, but obviously being just born he knows nothing of the world around him, cultural approaches, does and don't s etc. That means that if you target for a Level 3 or above, you will need to very thoroughly describe the culture, the values, the beliefs of your company, the red lines never to trespass, all the possible trade-offs that the Agentic AI may face and how to resolve it, etc. You will also need to foresee increasingly sophisticated systems of long term memory (including type of memory, what to store and what not, how to refresh the memory over time etc.) and feedback loops mechanisms. And you will need extensive testing in all circumstances, including attempting manipulation of the Agentic AI by badly intentioned users.

5. **More brains is better than one**: for use-cases where the tasks at hand involve judgement/creativity/customer facing end-products, AND where the cost of an AI fail (hallucination, inconsistency, wrong answer, etc.), it is essential to break down an Agentic AI into at least three sub-agents, all with powered by *different* brains, (e.g. three different LLM models) one proposer, one challenger, one assessor so that the end-products gets put forward only if the three of them agree. Various academic studies show that the combination of multiple perspectives leads to order of magnitudes better reasoning and fewer hallucinations. And the combination of the three outperforms significantly the abilities of any one of the individual brains.
6. **Keep human supervision**: human supervision is critical at important steps of the process, which are the steps where risky things can occur (e.g. going out to clients, performing payments, transferring data externally, etc.), or where the Agentic AI wants to innovate in terms of process, data usage or other things. In addition human escalation must also be foreseen to avoid the Agentic AI stalls when processes are different than expected or some, sometimes minuscule detail, changes e.g. the famous example of the Carnegie Mellon virtual company that failed on an unforeseen pop-up window.

Essential tools and technologies

Developing Agentic AI requires a good ecosystem of tools and technologies. The larger the enterprise the more tailoring to autonomy and scalability is essential. The choice is wide and growing. These are some of the tools for the more sophisticated actors:

- **Data processing**: Python (Pandas, NumPy), Alteryx for quality data.
- **Machine learning**: TensorFlow, PyTorch, OpenAI Gym for training.
- **NLP**: SpaCy, Hugging Face Transformers for language tasks.
- **Simulation**: Unity ML-Agents, Gazebo for safe testing.
- **Cloud**: AWS, Azure for scalable computing.
- **Ethics**: IBM AI Fairness 360, Google What-If Tool for compliance.
- **Agent frameworks**: LangChain, Autogen, AutoGPT, CrewAI, Codex for development.

By combining these tools with strategic planning and ethical oversight, developers can create powerful, responsible Agentic AI systems.

But if, like us, you are a small player the list is very different and many of these functions are combined into certain tools (see next chapter for details).

Steps to create Agentic AI

Building an Agentic AI system involves a structured process to ensure alignment with goals and operational success:

1. **Define objectives**: Set clear goals, e.g., automate support.
2. **Collect data**: Gather and preprocess relevant datasets.
3. **Select algorithms**: LLM, deep learning (NLP), etc.

4. **Design architecture**: Build modular systems with perception/triggers, reasoning, memory, and action.
5. **Define goals, behaviours, usable tools and train system**: Define operating responsibilities, behaviours, tools available for each agent, use simulations and real data for robustness.
6. **Configure cross-functional communication and decision-making frameworks**: Define interactions, decision-making remits and enable learning from outcomes.
7. **Integrate systems**: Connect to CRMs, APIs for actions.
8. **Test and deploy:** Launch, track, and update performance.
9. **Monitor maintain**: Regularly reassess performance, review tools.

Best practices/tips for developing Agentic AI

To maximize success, follow these best practices:

- **Start small**: Begin with a specific, high-value task before scaling to complex applications.
- **Ensure transparency**: Design agents to explain their decision-making process clearly to users.
- **Implement guardrails from the start**: Include fail-safes and ethical protocols to prevent unintended outcomes.
- **Retain human supervision**: Use a human-in-the-loop approach for critical or sensitive decisions, or when the agent is stuck.
- **Choose wisely**: Select tools and frameworks that match your project's needs and scale. And remember, all LLMs are not created equal for specific tasks.
- **Ensure data quality**: Feed agents clean, structured, and relevant data.
- **Monitor performance**: Track metrics to identify areas for improvement.
- **Test realistically**: Deploy in live environments to validate performance.

Challenges in developing Agentic AI

Despite its potential, Agentic AI development faces several hurdles:

- **Complexity**: Building autonomous systems requires advanced expertise and integration of multiple (legacy) components.
- **Data dependency**: Performance hinges on the quality, variety, and volume of data.
- **Ethical issues**: Ensuring fairness, explainability, and compliance with legal standards is critical.

- **Resource intensity**: Development and maintenance demand significant computational and financial resources.
- **Accuracy risks**: Agents may produce errors or "hallucinations" without proper validation.
- **Security**: Protecting sensitive data requires robust encryption and compliance measures.
- **Autonomy control**: Clear boundaries must be set to prevent overreach or unintended actions.

Addressing these challenges through careful planning, validation layers, and ethical frameworks is essential for success.

And beyond these technical elements, the success of failure of any AI experiment in today's organisations is still largely defined by two non-technical factors: the acceptance by employees and the avoidance of bad reputational surprises.

The future of Agentic AI

Agentic AI is poised to redefine industries as advancements in quantum computing, neuromorphic hardware, and AI frameworks enhance its capabilities. These systems will tackle tasks traditionally reserved for humans, such as creative problem-solving, empathetic communication, and strategic decision-making. From intelligent workflows to digital co-pilots, Agentic AI is evolving into a cornerstone of business innovation.

Practical implementation: running a micro-company with AI Agents

This section details how I, a non-coder, used four AI agents to manage three business lines: AI for Senior Leaders (e.g., Claridora), Small Businesses, and Everyday Life. While not enterprise-grade, it shows what's possible with creativity and practical tools.

In order to make this guide helpful to as many people as possible, while we focus on the choices and methods we have taken, we also try to show the more general set of options one could go for, in limited scale experiments as this one.

Please note all software suggestions are just for illustration or reflect my choices at the specific time and circumstances I chose, and are subject to change. I have no affiliation to any tool. All prices are only indicative.

Phase 1: setup

Step 1: Define your business objectives

The initial challenge was defining the roles of these AI agents. Should they be organized by business line or modelled after an executive committee (e.g., CFO, CTO)? I opted for the latter to ensure adaptability to new or evolving business lines and to establish clear top-level accountabilities. Additionally, I enabled these top level agents to utilize specialized AI sub-

agents. So they act as coordination agents of a number of sub-agents. Unlike a traditional company, where top agents (CFO, CTO, etc.) oversee specific functions, sub-agents are organized by skill (e.g., researcher, data analyst) and can be leveraged by any top agent. This structure prioritizes flexibility over a rigid chain of command, which is less critical for AI agents than for human teams.

To establish AI agents as Chief Financial Officer (CFO), Chief Technology Officer (CTO), Chief Marketing Officer (CMO), and Chief Legal Officer (CLO), begin by defining your business and operational requirements:

- **Document**: Outlined business models, and "how we do things here" a brief of values, do 's and don'ts, beliefs etc that are the backbone of the culture of the company.
- **Roles**:
    - **CFO**: Budgeting, bookkeeping, tax compliance.
    - **CTO**: Tech infrastructure, cybersecurity, coding.
    - **CMO**: Marketing, content, analytics.
    - **CLO**: Contracts, compliance, IP protection.
- **Decisions**: Identify decisions each role can make autonomously (e.g., routine expense approvals) versus those requiring human approval (e.g., investments). For example in our case any expense above a low limit must go through me. Same for any email sent to clients or public posting or change to websites. As the system improves I may release some of these, but we are at the start.
- **Boundaries**: Set boundaries for AI authority, such as financial limits or content approval requirements, to maintain oversight.

An interesting trick we used in developing these important guidelines is reverse prompt engineering. For example we fed a number of recent papers, recent marketing posts etc that I had developed to an LLM, asking it to define the briefs, the templates, etc. that one could deduct from those artefacts.

We also used different LLMs to identify situations where our briefs would be challenged or would not be sufficient to give the right guidance of behavior of the Agentic AI. This allowed us to really strengthen them, which is essential as we were going for Level 3 Agentic AIs in several applications.

Step 2: Collect data

In our case this was a very simple steps as we were stating from essentially no data. And often for small companies this largely a solvable problem due to the limited sheer volume of data.

For large companies, however, this is often a very difficult problem, as layers and layers of data have accumulated with inaccuracies. The possible approaches to slve this are beyond this paper. But in all cases, it is essential to deploy AI agent on largely complete and robust data.

Step 3: Select your AI Agent implementation approach

Choose an implementation approach based on your technical expertise, budget, and customization needs:

Option A: Custom-built LLM Systems

Leverage APIs from OpenAI, Anthropic, or xAI's Grok 3 for tailored solutions.

Requires programming skills or a developer.

Offers high customization but higher setup costs (~$5,000-$15,000 plus monthly API fees).

Option B: Low-Code AI Platforms - this is what we went for

Use platforms like N8N, Make.com, Relevance or AI workflow tools for faster setup. Complement with simple fine tuning of LLM (e.g. though OpenAI developer platform)

Requires moderate technical skills.

Balances flexibility and cost (~$100-$500/month).

Option C: Specialized Business AI Solutions

Adopt tools like Gleans.ai (for finance), Jasper AI (for marketing), or Rocket Lawyer (for legal).

Minimal technical setup needed but less customizable.

Cost-effective (~$50-$300/month per tool).

In my view, for small businesses, a hybrid approach combining low-code platforms and specialized tools is often most practical (so option B with here and there C for certain tools).

Step 4: Set up your architecture and core tech stack

Establish a centralized infrastructure to support your AI agents:

Central Communication Hub

Use Slack, Whattsapp, Telegram, Microsoft Teams, or Notion as a command center for AI interactions and notifications.

Create dedicated channels or sections for each role (e.g., each of CFO, CMO et; have a dedicated channel – like you would actually do a with a real CFO, CTO, etc.).

Document/Knowledge Management

Set up cloud storage (Google Drive, OneDrive) with organized folders for financial reports, contracts, marketing assets, and tech documentation.

Implement version control for critical documents to track changes.

Use Notion or ClickUp for a centralized knowledge base (we chose Notion).

Authentication & Security

Secure API keys and credentials using a password manager like 1Password or LastPass.

Enable two-factor authentication (2FA) across all platforms, and I own the phone that receives the codes for the most critical applications.

Use cybersecurity tools like CrowdStrike for endpoint protection.

Phase 2: Configuring individual executives, their goals, behaviours, tools

This is obviously the critical phase where the AI agents get actually built. It is essential to create extensive prompts/briefs to explain very clearly to the "Brain" (ie LLM model) what are the acceptable behaviors, the do s and don't s, the paradigms/beliefs the business I s built on etc.

It is also essential that in this world tools evolve very very quickly and you should constantly be willing to reassess whether the tool being used is optimal.

Last but not least, contrary to popular belief, LLMs are not created equal. They each have different strengths and weaknesses, both in abilities (language, reasoning, etc) but also in styles and prices. So here as well, accepting to evolve, as LLMs evolve.

Curve of performance of different LLMs. Credits: Llm-stats.com. For detailed ranking refer to Llm-stats.com

Step 5.a: Setting up your AI CFO

Core responsibilities:

Financial planning, expense tracking, bookkeeping, cash flow monitoring, financial reporting, tax compliance, and investment analysis.

Implementation steps and essential tools:

Connect accounting software (e.g. QuickBooks Online ~$50/month, Xero ~$13-$70/month) to your AI system for transaction imports and reporting.

Configure LLM (in our case Gemini) with custom prompts for financial analysis, budgeting, and tax compliance.

Set up automated monitoring for key metrics (e.g., cash flow, spending trends) using Make.com or N8N.

Create templates for financial statements, budgets, and tax reports in QuickBooks or Xero.

Define approval thresholds (e.g., expenses >EUR100 require human (me) review).

LLM for forecasting and analysis (prices vary significantly depending on tokens used).

Sample of one CFO prompt:

You are the CFO for [Company Name]. Using QuickBooks data:

1. Analyze current cash position and forecast for the next 30 days.

2. Identify unusual spending trends and suggest cost-saving measures.

3. Recommend budget adjustments to meet [financial goals].

4. Flag upcoming tax deadlines for [jurisdiction].

Constraints: [e.g., max $5-10,000 marketing spend].

Step 5.b: Setting up your AI CTO

Core responsibilities:

Managing tech infrastructure, software development, cybersecurity, technical problem-solving, and evaluating new technologies.

Implementation steps and essential tools:

Document your tech stack (e.g., AWS, GitHub) with access credentials in a secure knowledge base.

Deploy infrastructure on AWS, Google Cloud, or Azure (~$5-$50/month for small setups).

Use GitHub Copilot (~$10/month) for coding assistance and LLM (Grok or Claude Sonnet in our case) for step by step coding help.

Set up cybersecurity with CrowdStrike or Norton (~$60-$150/year) and configure security scans.

Automate monitoring of performance and alerts using Make.com, N8N et. and tools like Datadog (~$15/month).

Sample of one CTO Prompt:

As CTO of [Company Name], review our AWS infrastructure:

1. Identify security vulnerabilities or performance bottlenecks.

2. Suggest cost optimizations for EC2 instances.

3. Evaluate [new tool] for integration with our stack: [tech stack].

4. Propose updates to our technology roadmap.

Issues: [e.g., recent downtime].

Step 5.c: Setting up your AI CMO

Core responsibilities:

Managing marketing campaigns, social media content creation, websites review and benchmarking, lead generation, customer analytics, brand consistency, and market research.

Implementation steps and essential tools:

Create brand guidelines (voice, visuals) in a shared document.

Set up marketing tools: Canva Pro ($12.99/month) for visuals, Taplio or Hypefury or....(very wide choice available) for social media depending if focus on Linked in, Tik Tok or others, and Mailchimp (~$20/month) for email campaigns, Typeshare for long form text, if relevant.

Configure analytics dashboards in Google Analytics for traffic and campaign performance.

Use LLM (in our case Grok or Open AI Chat GPT depending on how edgy we want to be) for content generation (blogs, ads, social posts).

Automate content scheduling and reporting with Make.com or N8N.

Sample of one CMO Prompt:

As CMO of [Company Name], analyze our [last campaign] data:

1. Evaluate performance (clicks, conversions).

2. Recommend adjustments to improve ROI.

3. Draft three social media posts for next week aligned with [brand voice].

4. Identify new audience segments for targeting.

Goals: [e.g., increase conversions by 10%].

Step 5.d: Setting up your AI CLO

Core Responsibilities:

Contract management, regulatory compliance, intellectual property protection, risk assessment, and policy development.

Implementation steps and essential tools:

Create a database of legal obligations and deadlines in Notion or Google Drive.

Use DocuSign ($25/month) for e-signatures and Rocket Lawyer ($39.99/month) or a tool designed for your jurisdiction for legal templates (NDAs, agreements).

Configure LLM (for us ChatGPT) for drafting and reviewing contracts with compliance checks.

Set up compliance monitoring with LLM to track regulatory changes.

Automate contract storage and renewal alerts using Mae.com or N8N.

Sample of one CLO Prompt:

As CLO for [Company Name], review this [contract/situation]:

1. Identify legal risks or compliance issues under [regulations, e.g., GDPR].

2. Suggest revisions to protect our interests.

3. Draft an implementation plan for changes.

4. Assess regulatory implications in [jurisdiction].

Risk tolerance: [low/medium/high].

Phase 3: integration and workflow set up

Step 6: Configure cross-functional communication and decision-making frameworks

Set up a daily "executive meeting" in Slack or Notion where AI agents share updates (e.g., CFO reports cash flow, CMO shares campaign metrics). For a bit of fun, I have even given them names, and an AI generated face: FofAI (CFO), TotAI (CTO), MomAI (CMO), LalAI (CLO).

Create decision-making flowcharts for cross-departmental issues (e.g., launching a new product).

Build escalation protocols for human intervention (e.g., critical financial decisions or simply if any of the agents is stuck on an unexpected blockage – such as a pop up etc we have a human on call to react and unstuck them).

Implement context-sharing via a shared knowledge base so agents access relevant data (e.g., CMO uses CFO's budget data).

Define decision authority for each agent (e.g., CFO approves expenses <EUR100, CMO adjusts ad spend within budget).

Create templates for decision types (e.g., financial approvals, campaign changes) to enable recording of history and then ex-post analyses in several weeks/months.

Log all AI decisions automatically in Notion or ClickUp for review.

Set "red flag" triggers (e.g., unusual spending, security alerts) requiring human input.

Phase 4: implementation

Step 7.a: Integrate systems - Knowledge Management System

Build a central knowledge base in Notion or Google Drive accessible to all agents.

Automate documentation of decisions and outcomes (e.g., via Zapier to Notion).

Implement version control for key documents (e.g., contracts, budgets).

Schedule routine updates to keep the knowledge base current.

Step 7.b: Integrate systems - implementation method

Select an approach based on your skills:

Non-Technical (Code-Free):

Use Make.com or N8N to connect tools that must be connected together.

Assess whether to subscribe to specialized AI tools (e.g., Jasper for CMO, Rocket Lawyer for CLO).

Manage via Slack or Notion with pre-built templates.

Technical (Custom Development):

Host agents on AWS/Azure with Grok 3 or OpenAI APIs.

Build a database for memory/context (e.g., PostgreSQL).

Create custom dashboards to monitor performance.

Hybrid (Moderate Skill):

Combine no-code tools (Make.com, N8N, Notion) with simple Python/JavaScript scripts.

Step 7.c: Integrate systems – automations, notifications, approvals

Data flow automation

Use N8N and Make.com to connect software to each of the AI Agents (e.g. QuickBooks to CFO agent, Google Analytics to CMO, GitHub to CTO, and DocuSign to CLO).

Notification systems

Configure alerts for critical events (e.g., low cash flow, security breaches) via Slack.

Schedule daily/weekly reports from each agent to Notion or Slack.

Approval workflows

Build approval processes for decisions exceeding AI authority (e.g., via Notion forms).

Set up emergency override protocols for urgent issues.

Step 8: Test and deploy your AI executive team

Run simulated scenarios (e.g., CFO budgets for a product launch, CMO designs a campaign).

Test cross-functional collaboration (e.g., CTO and CMO align on website updates).

Verify escalation protocols (e.g., human alerted for high-risk decisions).

Conduct a full "business day" simulation to ensure system stability.

For example

Morning Briefing*:*

8:00 AM: AI agents send overnight summaries to Slack (CFO: sales report, CMO: ad performance, CTO: website uptime, CLO: contract status).

8:30 AM: You review summaries, set priorities, and approve recommendations in Notion.

Daily Operations:

CFO: Monitors sales and expenses in QuickBooks, flags overspending, and forecasts cash flow.

CTO: Oversees AWS-hosted website, uses GitHub Copilot to fix bugs, and scans for security issues.

CMO: Publishes social posts via Hootsuite, optimizes ads with Google Analytics, and drafts blogs with Grok 3.

CLO: Reviews customer contracts in DocuSign, monitors GDPR compliance, and files trademarks.

Decision Handling:

Routine: AI handles autonomously (e.g., CFO categorizes expenses, CMO schedules posts).

Medium-impact: AI proposes, you approve (e.g., CTO suggests new tool, CLO revises contract).

Critical: AI alerts you immediately (e.g., CFO flags low cash, CLO detects regulatory issue).

End-of-Day:

Agents compile summaries and metrics in Notion.

You review and provide feedback via Slack for continuous improvement.

Phase 5: Ongoing management

Step 9.a: Monitoring and regular maintenance

Review each agent's performance weekly via Notion or Slack summaries.

Update prompts and decision parameters based on business changes.

Audit tool performance monthly with LLM (either one or on a rotation basis) (e.g., "Suggest optimizations for QuickBooks usage").

Conduct an "all-hands" system review monthly to assess integration.

Step 9.b: Continuous improvement

Document successful and failed decisions to refine agent performance.

Use feedback loops to improve outcomes.

Test new AI features (e.g., Grok 3's DeepSearch mode for complex tasks).

Update the knowledge base with industry trends and regulatory changes.

Legal and ethical considerations

Data Privacy: Ensure tools comply with GDPR, CCPA, or jurisdictional laws (CLO agent verifies).

AI Limitations: Review critical financial and legal outputs manually, as AI cannot fully replace human judgment.

Liability: Use Rocket Lawyer to draft terms-of-service agreements limiting liability for AI-driven decisions.

Maintenance and growth

Monthly Check-ins: small audit with help of an LLM (depending on the area audited) to challenge and optimize tools usage.

Tool upgrades: upgrade plans (e.g., QuickBooks Advanced, AWS enterprise) as revenue grows.

Human backup: where necessary, hire freelancers for complex tasks (e.g., tax audits, lawsuits).

By following this guide, you can deploy a cost-effective, scalable AI executive team to manage your business with minimal human oversight, keeping you in control of strategic decisions.

No doubt, this is only the start of the journey.

As our learning progresses, we will keep on updating this how-to guide.

We welcome your feedback, ideas and comments.

If you want to contact us, please go to

- ➔ *allegrapatrizi.com*
- ➔ *Linked in Allegra Patrizi Barberini*