1,000 assistants in your pocket

The boring and practical revolution that enables commercial teams in complex BtB companies to **harness the power of GenAI**

August 2025

Whitepaper

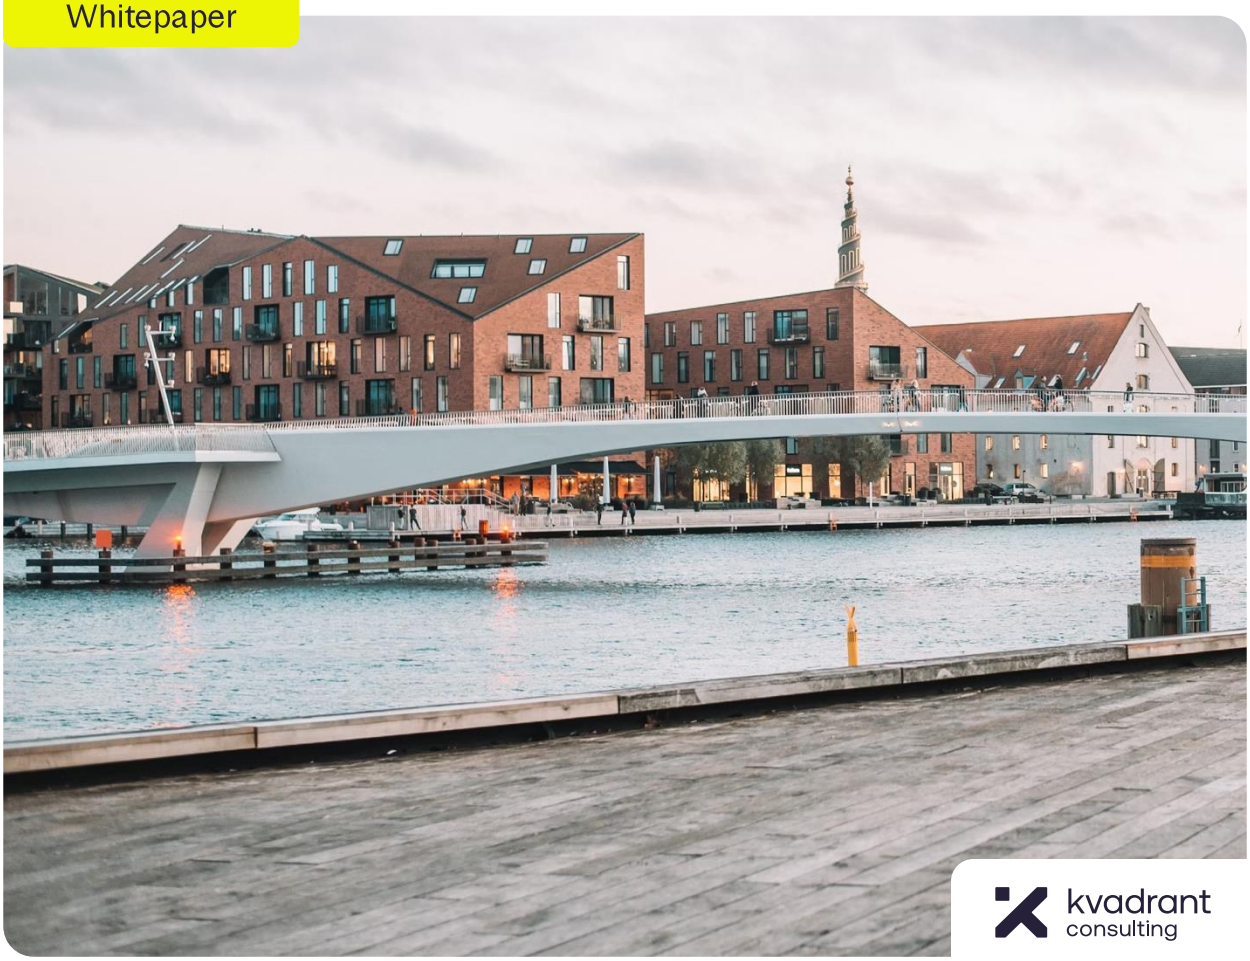

kvadrant consulting

About Kvadrant Consulting

In B2B, buying behaviour is shifting from mainly in-person interactions to a blend of digital and physical touchpoints. Kvadrant Consulting helps companies capitalise on this change and turn their commercial engines into lasting competitive advantages.

We support clients as they assess the commercial potential of acquisitions, choose the right growth paths, and build the capabilities needed to deliver on them.

Today, more than 35 professionals operate from our Copenhagen office across three practices: Transactions & Strategy, Commercial Excellence, and Go-to-Market Excellence.

About the authors

Mikkel Bach-Andersen

Senior Partner, Co-leader of Go-to-Market Practice

Mikkel has more than 15 years of experience of practicing and consulting within B2B marketing & sales. He has a functional focus on value proposition design, commercial operating model transformations, organizational design, commercial & marketing strategy, product launches and deal acceleration. Co-author of *The Campfire Launch*

Nils Lindner Koch

Partner, Co-leader of Go-to-Market Practice

Nils is a seasoned marketing, sales, and positioning practitioner with a strategic mind and more than 15 years of experience leading end-to-end brand and marketing transformations and implementation programs for leading B2B and B2C companies.

Hannibal Herforth

Manager, Go-to-Market Practice

Hannibal combines a technical background in machine learning and AI with deep commercial expertise. He has a functional focus on lead engine development, GTM operations, data analysis, and implementing AI/LLMs in the commercial organization.


kvadrant
consulting

At a Glance

Three years after the first wave of enterprise experiments with generative AI, commercial leaders are asking a new question: **how do we turn scattered pilots into repeatable, large-scale transformation?**

LIMITED ROI FROM FIRST-GEN EXPERIMENTS

Most companies have adopted ChatGPT-style tools and run dozens of proofs of concept, yet fewer than one-third report meaningful, bottom-line impact. Adoption is high, but value is trapped in isolated pilots and one-off use cases.

THE NEAR-TERM PRIZE LIES IN A SWARM OF NARROW AI-ASSISTANTS

Winning firms are shifting from chasing one “super-agent” to cultivating hundreds of task-specific AI Assistants embedded in everyday apps. This approach scales faster, is proven to create value, and frees teams to focus on strategy, creativity, and customer relationships.

TO MOVE BEYOND THE GRASSROOT LEVEL REQUIRES STRUCTURE AND PORTFOLIO THINKING

Moving beyond the grassroots level and low adoption requires treating each AI assistant as a mini-product, with clear ownership, defined guardrails, and usage metrics. It’s important to take a portfolio approach to the full suite of assistants by continuously pruning and optimising existing ones while releasing new ones.

Reality punches back as the GenAI hype is slowing down

Generative AI has exploded onto the scene in 2024-25, with **unprecedented interest, and commercial functions like marketing & sales, have been quick to adopt this new technology.**

Beneath the hype, the picture is sobering. Many firms find only **modest returns so far**, with ROI and scaling lagging.

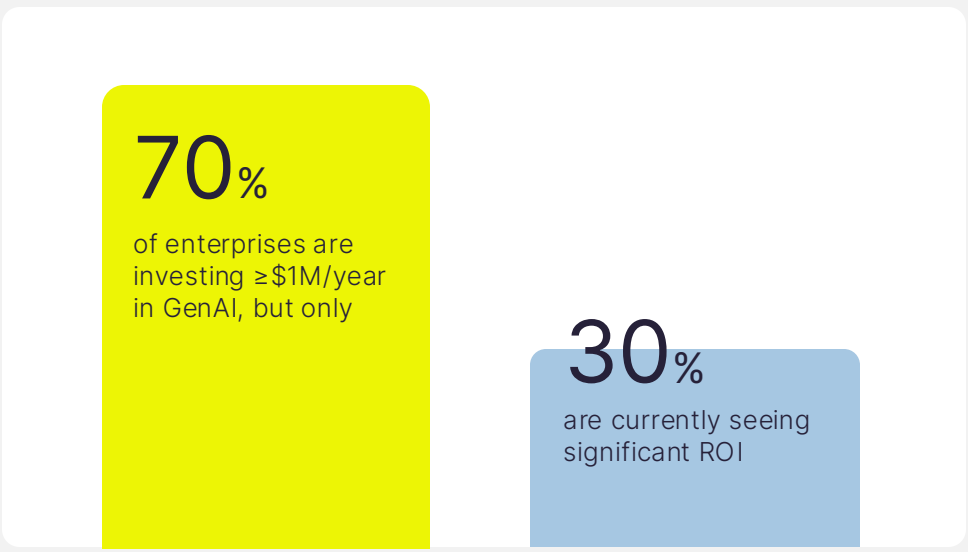

In practice most early projects never leave pilot stage: a Deloitte study notes most companies run **20 or fewer experiments** and **expect only 30% of them to scale within 6 months.** It is clear that GenAI is not a silver bullet and that the great productivity claims has yet to materialize for most companies.

Our observations from the field mirror these findings: Great ambitions and strong commitment to testing things out. Followed by big issues when it comes to getting things in production and scaling GenAI across the commercial organization. What's more large parts of commercial organizations use out-of-the-box GenAI platforms like ChatGPT, Midjourney or MS Co-pilot very little if at all.

Despite these lukewarm (at best!) results we are witnessing a boring and somewhat silent GenAI revolution in some commercial organizations. A revolution that is happening right under the eyes of executive teams and far from polished strategy documents and key notes.

Let's unpack it.

Authors´ note: This is not a technical essay on GenAI's frontier possibilities, we have absolutely no understanding of those, nor a highlight reel of what nimble Silicon Valley start-ups are tinkering with. Instead, it offers a grounded view of what is actually being used, and creating measurable value, inside the commercial functions of large, complex, global B2B enterprises.

The boring GenAI revolution: focus on practical assistants, not magic agents

Many industry pundits have fixated on the “agentic” AI, autonomous bots that can execute complex, multi-step workflows with very little to no human involvement, as the defining breakthrough. Social media is full of slogans like “Go agentic or die!”. You can’t open LinkedIn or a report from one of the big consultancies without getting serious FOMO. But our experience is that for commercial organizations in large global B2B companies, the terrain is very different. They are often characterized by poor data quality, stringent security demands and an organization is made up of thousands of employees, many of whom are not digital natives.

What's more: Too many “moon-shot” GenAI initiatives never goes into production because IT is busy. Or if they do they end up morphing into old-fashioned IT and master-data-management programmes. Once pilots move beyond the sandbox, teams discover that customer data records are filled with errors and duplications, missing product hierarchies, poor integrations, and security shortcoming. Budget shifts to data cleansing and API plumbing, timelines stretch, and business sponsors lose patience.

Key point: The cutting-edge autonomous agents that work in a 10-person startup may misfire or stall in a regulated enterprise with high data and tech complexity.

The right level of agency for your organization

The commercial hype cycle around AI has been dominated by talk of “fully autonomous agents”. AI agents capable of orchestrating complex workflows across systems with no human oversight.

On paper, this sounds like the future. But for large B2B companies, the potential risk of giving AI too much agency is simply too high. You can’t just fully automate your outbound motion, as you’ll risk your reputation to both existing and new customers.

Most commercial functions should not start by aiming for full autonomy. Instead, they should focus on finding the low-risk and high-impact use-cases: AI Assistants that sit inside the tools people already use and support tasks that are high-frequency, rule-based, and time-consuming. These assistants can be deployed rapidly, governed with light-touch controls, and improved continuously based on user feedback.

The Risk–Agency Trade-Off

The relationship between risk and agency for enterprise companies

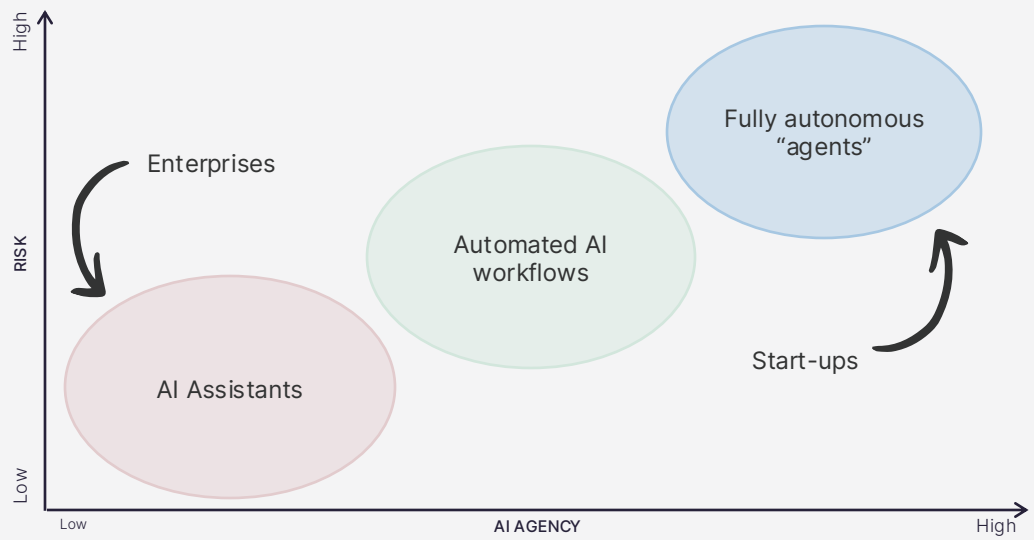

Besides raising baseline AI fluency, we believe that the most value creating approach for most commercial organizations should not entail chasing fairy tale autonomous "AI agents" but rather **to develop "swarms" of narrow AI assistants solving for a specific task.** And then later move on to more autonomous and technically challenging agents.

What is more interesting is that this is already happening in a small scale in in a lot of companies. Yet this boring revolution is happening at the "factory floor", far from top-down enterprise level strategies and glossy vision statements.

How ask-based assistants are different from autonomous agents

While "AI agent" has no universally accepted definition, it always implies a system with some degree of autonomy in completing a task. In practice, the label spans everything from a single-prompt chatbot to a fully orchestrated platform that plans, decides, and acts across shifting conditions.

To keep this discussion concrete, we'll use two anchor points: **task-based assistants** and **fully autonomous agents** that in reality operate at a spectrum.

Attribute	Task-Based Assistants	Fully Autonomous Agent
Typical scope	One narrow, repetitive micro-task (e.g., draft a follow-up email)	End-to-end workflow (e.g., qualify, nurture, and route a lead)
How it works	One-shot prompt → LLM returns draft → human tweaks	Plans a multi-step workflow → calls tools & data → iterates until goal met
Data needed	Text templates, small data snippets; often a single system of record	Multiple live systems (CRM, ERP, calendars, APIs) plus unstructured data
Build effort	Hours to a few days with low-/no-code tools	Weeks to months; requires orchestration, testing, and DevOps support
Governance	Light review/approve loop; human remains "in the chair"	Continuous monitoring or fallback rules; humans only intervene on exceptions
Best used for	Speeding up everyday grunt work and freeing time	Automating complex, cross-system outcomes with minimal human touch

The case for a “swarm” of assistants

Modern commercial teams are drowning in admin work and repetitive, mindless work. Consider sales: frontline reps now **spend just 28 percent of a typical week actually selling**; the remaining 72 percent vanishes into CRM updates, quote building, and internal meetings according to research by Salesforce.

Marketing fares no better. In a survey done by the Marketing AI Institute 1,800 marketers were asked what they most want from AI, their overwhelming Nº 1 answer (80 percent) was “**reduce time spent on repetitive tasks**.”

No, AI assistants won’t empower you to rethink you GTM engine, add anything close to “intelligence” or allow you to totally hand off entire workflows to the machines without human involvement. But they will liberate a lot of time and raise bassline quality in a lot of key output.

Time that can be reclaimed and redirected towards thinking, tinkering and reimagining.

Sources:
Salesforce, *State of Sales, 5th Edition*
Marketing AI Institute, *The 2024 State of Marketing AI Report*

Best Practices

What Makes a Great AI-Assistant Use Case?

Look for tasks that are…	Why they're ideal for an assistant
Repetitive & time-consuming (daily, weekly, or always "one more")	Saves humans from mind-numbing drudgery → frees capacity for higher-value work
Low in fulfilment (little creativity, low professional satisfaction)	Off-loading "grunt work" boosts morale and lets people invest energy where it matters
Easy to judge as "right vs. wrong" (clear quality/accuracy criteria, yes/no checks)	Enables quick automated or human Quality Control loops → confidence builds, errors caught fast
High-volume or high-frequency (hundreds/thousands per month across the org)	Small per-task savings compound into large, trackable business impact
Rule-based & data-ready (inputs exist in CRM, ERP, docs, templates)	LLM can follow well-defined steps and pull structured data without heavy integration
Low- to medium-risk (no regulatory show-stopper if the first draft is imperfect)	Let's teams pilot, iterate, and prove value quickly before tackling riskier domains

Yet, despite the potential of this "boring revolution", it often fails to move beyond the grassroot level in many commercial organizations, because:

- Assistants are built in the different corners of the organizations
- Assistants aren't properly maintained & governed
- Different assistants solving the same use cases are mushrooming across the organizations

There is no adoption strategy, and the assistant are mostly picked up by the few front runners, not the broader organization.

The limits of AI Assistants

AI assistants are nowhere “intelligent”. They execute within a set framework and clear boundaries:

- **They lock in today’s processes.** Assistants accelerate whatever workflow exists. if the workflow is flawed, the flaws scale faster.
- **No strategic imagination.** Assistants answer questions that users ask. Business-model reinvention and category-breaking ideas still require people.
- **Productivity ≠ capability.** Hours saved on drafting or data pulls don’t automatically build deeper selling, negotiation, or innovation skills; leadership must reinvest the freed time.
- **Data debt compounds**. Assistants inherit your taxonomy, templates, and CRM hygiene. The issues stemming from poor data quality or outdated playbooks will remain.
- **Risk of cultural complacency.** Easy automation can lull teams into “autocomplete mode.” Curiosity, critical thinking, and cross-functional dialogue still need active cultivation.

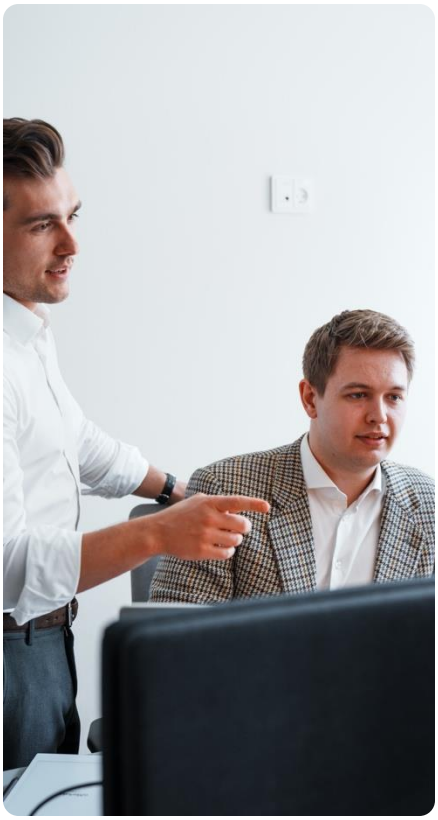

Bottom line: AI assistants are force multipliers for efficiency, not engines of innovation. Humans must still rethink the work, steward data, and lead change.

How to go from grassroot to scale

To move beyond the tinkering phase, treat each assistant like a mini product: owned by the function that uses it, anchored in a proven methodology, and incrementally improved.

For example, a sales rep might have a personal “deal assistant” that drafts proposal outlines from a master template; a marketer might have a “campaign composer” that generates a landing page, email, and social posts from a brief. While they don’t score high in sexiness, these tools can often be built with **light coding or even no-code GPT builders**, don’t come with major security concerns, and they typically leverage just a few data sources (CRM fields, product docs, brand guidelines).

The resulting productivity gains are immediate: employees reclaim hours from drafting and data-gathering tasks, and instead focus on the creative, relational parts of their jobs.

Over time a successful GenAI program could have **many of these assistants**, hence the phrase “age of 1000 assistants.”

To reach that future potential, five infrastructure pieces are crucial:

1

Robust methodology & guardrails.

Every AI assistant must be built on clear rules and structured inputs/outputs. You must define “what good looks like” for that task (e.g., components of a great pitch email, value proposition or press release) and encode those guidelines into the tool (via prompts, templates, or example libraries). Without this, outputs from universal tools Like ChatGPT and Co-pilot vary wildly.

Equally important, build in **AI guardrails** at every step; checks that ensure brand tone is consistent, sensitive content is blocked, compliance/legal standards are met, and factual accuracy is verified (LLMs will produce mistakes if unchecked). Tools like OpenAI’s Assistants API or Microsoft’s Teams AI library can help embed these safeguards (they support moderation filters etc.), but **human oversight** is still required.

2

Clear ownership and governance

Accountability must live in the business function, not in a distant IT queue. The people who deeply understand the task should own the assistant: **they define requirements, validate outputs, and drive user adoption**.

Behind the scenes, a central AI team or center of excellence (even a handful of ML engineers or "AI forward-deployed" experts) can support them, providing tools, enforcing guardrails, and helping with integration. Think of it like a product team: product owners (in marketing, sales, etc.) plus a small tech team.

Meanwhile, corporate should set up a light governance framework: guidelines for how assistants are named, how data privacy is handled, and metrics to track (usage, accuracy, etc.). **Policies should cover the entire lifecycle**, from scoping a use case to sunsetting outdated assistants.

3

Accessible where people work

An AI assistant is worthless if employees can't find or trust it. Centralize them in familiar interfaces. Many companies create an "**AI hub**" or "**GPT store**", a searchable intranet page or Teams channel listing all approved assistants with usage guides.

Embedding assistants into tools like MS Teams or Slack can also boost use: for instance, a team chat bot where a seller simply types "@pressBot generate press release" and gets a draft. ChatGPT Enterprise and Microsoft's Copilot also offer admin consoles and single-sign-on so that IT can roll out "approved GPTs" organization-wide with security.

Embrace those enterprise-grade channels.

4

Ongoing portfolio management

Once a suite of assistants exists, regularly prune and grow it. Not every pilot will work; periodically retire assistants with low usage or high error rates.

Conversely, **empower frontline teams to experiment** and spin up new ideas:

- Host "AI innovation days" for sales/marketing staff to propose assistant concepts.
- Create an idea-review board (with reps from IT, legal, and the business functions) to vet and fund the most promising prototypes.

Also, **encourage a culture of sharing**:

- Publicize which assistants succeeded and why and invite other teams to copy or adapt them.

5

People and change management

Finally, make AI empowering, not threatening.

> According to research by Prosci, a change management consultancy, **63% of organizations find that human factors (resistance, uncertainty) are the top barriers to AI success.**
>
> To avoid that, emphasize that assistants are like virtual colleagues who take orders, not replacements.
>
> For example, explain to a copyeditor that the AI will draft pieces, but they'll review and refine them, freeing them to focus on strategy and creativity.

Train leaders and managers to communicate a clear vision ("This technology will help you do your job better") and provide lots of examples of AI freeing people from mundane work. Encourage an internal "AI champions" network: employees who are enthusiastic and willing to mentor others.

Above all, invest in communications and training. Regularly solicit feedback: are people finding the assistants helpful, confusing, or missing something? Use that to iterate on both the tech and the rollout plan.

A day in the life of: Maria, Marketing Manager *Empowered by AI Assistants*

08:00 **Morning Metrics**
Over coffee Maria opens Teams and meets **Morning Metrics**, a one-page digest of overnight performance. A red tag on LinkedIn CTR links straight to a suggested A/B image test. Five minutes later she already knows which lever to pull.

09:15 **Brief → Assets in one go**
Back at her desk she drops a two-paragraph launch brief into Campaign Composer. In under a minute, it hands back:

- A landing-page draft with brand tone and legal footer baked in
- Two nurture-email variants, each with pre-filled UTM tags
- Seven LinkedIn snippets sized to her region's character limit

Maria changes a headline, tweaks a CTA, and ships the pack to brand review. What once cost two hours is finished in fifteen minutes.

11:30 **Instant localization**
The German sales lead asks for DACH-ready collateral. Maria pastes the EN landing-page URL into Localization Buddy. It flips currency symbols, switches to formal "Sie", and inserts a local warranty line. A quick skim and she approves.

14:00 **Webinar deck, asset-hunt free**
With tomorrow's webinar looming she types "/prep webinar" to Event Prep Bot. It pulls the freshest case study, latest analyst stat, and correct product hero image into a 15-slide skeleton. Maria's afternoon is spent crafting the narrative, not scavenging files.

16:30 **Exec snapshot without spreadsheets**
Before logging off, she clicks Performance Snapshot in PowerPoint. Yesterday's funnel numbers, attribution chart, and a call-out ("Product B leads convert 22 % faster") appear on a single slide—ready for Monday's budget debate.

Overall payoff
~6 hours saved, collateral stays on-brand, and Maria finally books time for competitive-message testing instead of late-night copy edits..

A day in the life of: Lars, Enterprise Sales Rep *Empowered by AI Assistants*

07:30 **Coffee with a call plan**
On the commute Lars opens **Daily Briefing** in his mobile CRM. It:
- Ranks today's deals by health score
- Surfaces an unread procurement email flagged "respond ASAP"
- Suggests three talking points drawn from the prospect's latest 10-K

By the time the train stops he knows exactly who to call first and why.

09:00 **Account research in a click**
Five minutes before a discovery call he taps **Account Researcher**. It scans LinkedIn moves, SEC filings, and industry news, then drops three fresh insights into the meeting record. Instant credibility without rabbit-holing the web.

10:30 **Notes & follow-up on autopilot**
During the call **Note Scribe** transcribes, tags objections, and timestamps next steps. As soon as the meeting ends **Follow-Up Email Creator** drafts a recap, attaches the transcript PDF, and tees up a calendar invite for next week. Lars glances, adds a personal line, and hits Send.

12:00 **Quote in minutes**
He clicks "Generate proposal." **Deal Desk Assistant** pulls SKUs and discount policy, then produces:
- A work sheet with price table
- A ready-to-sign order form

The prospect receives a package before lunch.

16:00 **Pipeline hygiene in one command**
Just before wrap-up he types "/clean pipeline" to Forecast Cleaner in Slack. It auto-moves stale deals, updates close dates, and prompts for missing next-step notes. No more Friday night data chores.

Overall payoff
~30 % less admin, quotes go out same-day, and Lars reinvests the freed hours into live calls and strategic account mapping

Organizational Setup for Scaling AI Assistants

Realizing the vision of embedding hundreds of task-specific AI assistants across a complex B2B commercial organization requires a purpose-built organizational setup. This setup should balance bottom-up experimentation with lightweight, business-led governance.

The core of this structure lives inside the **commercial function**, supported by enabling corporate functions (IT, Legal, HR).

Commercial function

AI assistants are only effective if built close to the user and owned by those who understand the workflow. Therefore, the commercial function must host the central engine for identifying, developing, scaling, and managing assistants.

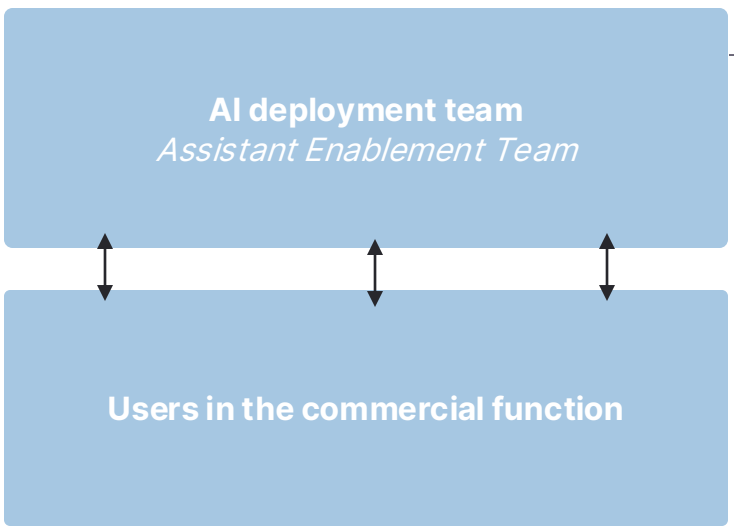

Corporate functions

The commercial function leads, but corporate functions provide essential guardrails and infrastructure.

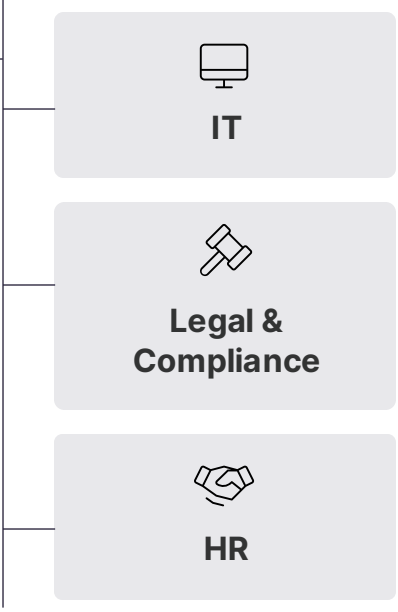

AI Deployment Team

A cross-functional team, ideally embedded within the commercial organization (e.g. marketing, sales ops, enablement), acts as the "product team" for the assistant portfolio.

Key roles include:

Product owners responsible for identifying high-impact use cases, driving adoption, collecting feedback, and owning the lifecycle of assistants in their area.

Embedded AI Specialists support quick builds and iteration, integrating lightweight tech while following security and governance principles.

AI Program Manager responsible for identifying high-impact use cases, driving adoption, collecting feedback, and owning the lifecycle of assistants in their area.

Training & Adoption responsible for identifying high-impact use cases, driving adoption, collecting feedback, and owning the lifecycle of assistants in their area.

Key responsibilities:

- Use case pipeline management and prioritization
- Prototyping and piloting new assistants
- Driving rollout and user onboarding
- Governance adherence and version control
- Sunsetting underperforming tools
- Reporting usage, satisfaction, and ROInt" model.

Users in the commercial function

Frontline commercial staff (e.g. sales reps, marketers, bid managers) play a **crucial dual role** as:

Beneficiaries of assistants who gain back time and productivity

Co-creators who identify opportunities, stress test prototypes, and signal value/fit

Key contributions:

- Suggesting new assistant ideas during “innovation jams”
- Participating in user testing and feedback loops
- Advocating internally through a “champion network”
- Reporting blockers or misfits in real-world workflows
- Creating **AI champions** in each sub-function (e.g., field sales, product marketing) is critical to scale adoption.

IT

Supports scalable deployment while safeguarding data and systems.

- Provides access to platforms (e.g., Microsoft Copilot, OpenAI, internal GPT builders)
- Manages data integrations and API access (CRM, ERP, etc.)
- Enforces security standards and SSO
- Helps maintain central “GPT Store” or AI hub

Legal & Compliance

Ensures assistants adhere to regulatory and internal standards.

- Reviews use cases involving personal or sensitive data
- Defines guardrails for privacy, copyright, or discriminatory content
- Approves content moderation logic (e.g., banned words or claims)

HR / L&D

Ensures the workforce has the capabilities and confidence to use AI.

- Partners with Commercial to raise AI fluency through onboarding, micro-learning, and leadership engagement
- Embeds GenAI readiness into capability frameworks and learning journeys
- Supports change management and communication planning

How to get going and keep going: build structure and focus on adoption

The path to realizing its full potential inside large commercial organizations is less straightforward.

Unlocking the value of “a thousand assistants” requires a structured, business-led journey - one that starts with the right use cases, builds momentum through early wins, and scales through governance, adoption, and continuous improvement.

Below, we outline a pragmatic four-phase approach that has can help you move from isolated pilots to production and adoption at scale:

Phase 1: Discover and Prioritize

Phase 2: Build and Test Useful Tools

Phase 3: Establish the Foundations for Scale

Phase 4: Drive Adoption and Continuous Improvement

Phase 1

Discover and prioritize

The first step is to identify **where AI can create immediate, practical value** without requiring complex integrations or large datasets. This typically begins with structured discovery: short interviews and workshops across key teams to surface repetitive, manual tasks that are **high-volume, rule-based, low risk and currently underserved by automation**. These tasks often live in the "long tail" of operational work, drafting internal emails, assembling campaign assets, or preparing FAQ responses for townhalls.

The goal is **not** to chase the most complex or cutting-edge use cases, but rather to find those that are low-friction and high-impact. Once a range of candidate tasks has been mapped, they should be prioritized using simple but robust criteria: **impact potential, feasibility given current tools and data, and urgency or pain level for users**. Each shortlisted use case should have a clear business owner and success metrics. This sets the foundation for **focused experimentation with real accountability**.

Phase 2

Build and test useful tools

With a set of validated use cases, organizations can move quickly to prototype and test narrow AI assistants built around real tasks. These early tools often take the form of prompt templates or lightweight apps built on existing platforms such as **Microsoft Copilot, internal LLM deployments, off the shelf tools or no-code GPT builders.** The emphasis should be on integrating AI into live workflows, **not** simulated environments or one-off demos.

For example, teams may build assistants that draft press releases from briefing documents, generate first-pass email and landing page copy, support Q&A preparation for leadership events, translate and adapt content for local markets, or pull CRM and product data into customer-facing materials. These tools are not designed to be perfect out of the gate, but **they must be usable, testable, and embedded in day-to-day work**.

Rapid iteration is key. Collect feedback systematically, refine prompts and outputs based on real user input, and begin to shape a repeatable development playbook. The aim is to demonstrate tangible utility early and build internal confidence in the approach.

Phase 3

Establish the foundations for scale

As the number of assistants grows, so does the need for **consistency, governance, and shared infrastructure.** Without some degree of structure, organizations risk fragmentation: duplicative tools, variable quality, unclear accountability, and escalating compliance risk. Yet the solution is **not** to centralize everything under IT or impose rigid controls. Instead, the goal should be to establish lightweight governance that enables scale while preserving speed and flexibility.

This includes **defining clear ownership for each assistant** within the business function that uses it, ensuring that those who understand the task best remain accountable for its performance and ongoing relevance. Editorial guardrails and prompt libraries can help encode organizational standards (such as brand voice or compliance language) into AI outputs. A simple internal hub or "assistant store" should catalogue available tools, explain their use cases, and provide onboarding guidance to new users.

Equally important is **building a network of internal champions, early adopters across functions who can support their peers, help with onboarding, and surface new opportunities**. These individuals often become the bridge between technical enablers and business users, helping scale adoption from within.

Phase 4

Drive adoption and continuous improvement

Ultimately, the success of AI in commercial work depends less on the sophistication of individual tools and more on their **adoption and sustained impact**. This means investing in change management and portfolio evolution over time.

Driving adoption starts with embedding assistants in the tools and channels people already use, whether that's Microsoft Teams, Slack, or internal dashboards. Assistants should be one click or one command away from the user's workflow. Teams should receive **targeted onboarding, including walkthroughs, use cases, and success stories that demonstrate real impact**. Usage data should be monitored not only to measure ROI but to identify friction points or underutilized tools that may require redesign or retirement.

In parallel, organizations should encourage bottom-up innovation. Internal ideation sessions or "AI assistant jams" can source new use case ideas directly from frontline employees. A small cross-functional review board can help evaluate, fund, and prioritize these for rapid development. Over time, **the assistant portfolio should be actively managed-retiring outdated tools, updating logic and prompts, and standardizing what works across regions or teams.**

Phase 4

Crucially, **people must remain at the center of the transformation**. That includes reinforcing the message that AI is an enabler, *not a threat*, freeing employees to focus on judgment, creativity, and relationship-building. Managers should be trained to communicate the “why” behind the rollout, while power users should be supported and recognized for coaching peers and sharing knowledge.

Feedback loops, both technical and cultural, are essential to ensure that assistants remain relevant and that AI becomes a trusted part of how work gets done.

In short, the journey from pilot to pervasive AI doesn’t require reinventing the company, but it does require intention. Find useful tasks, build narrowly and fast, support with just enough structure, and **let the people closest to the value lead the way**.

Closing remarks

Mastering the age of a thousand assistants is not a moon-shot IT project - it is much more akin to a disciplined march.

Yes, it demands patience, because **each assistant must be treated like a mini product with clear owners, guard-rails, and an adoption plan**. And it requires diligence, because success rides on constant pruning, retraining, and re-communicating as the portfolio evolves. But the reward is transformative due to the stacking nature of the assistants: when routine drafting, data-gathering, and formatting chores are handed to purpose-built AI helpers, commercial teams suddenly gain back the scarcest resource in business: **Attention and time**.

With that reclaimed attention, sales and marketing professionals can double down on the uniquely human levers of growth: **crafting sharper value propositions, nurturing customer trust, and spotting market shifts before the competition**. Managers become orchestrators of creativity rather than reviewers and context providers, while frontline employees move from "doing the work" to directing the work.

Getting there, however, is not a technology race; it is an execution race. Companies that pair bottom-up experimentation with lightweight governance, embed assistants where people already work, and invest in broad AI fluency will compound their wins far faster than those chasing a single "*magic agent*." The silent revolution is already under way on the factory floor far from the eyes of the executives.

Now is the moment to give it **structure, scale and staying power**.

Any further questions?

Feel free to reach out to:

Mikkel Bach Andersen

mba@kvadrant.dk

+45 40411331

Nils Lindner Koch

nlk@kvadrant.dk

+45 29691495

Hannibal Herforth

hhe@kvadrant.dk

+45 29808685

kvadrant
consulting

Grow smarter