

Autonomous AI Sales Representative For an Australian Consulting Company

How building an autonomous AI virtual worker enabled an Australian consulting firm to process **20% more leads** and boost upselling/cross-selling by **19%** while reducing cost per lead by **15%**.



Business challenge

With AI agents being marketed as a magic wand to wave away employees' daily grind, businesses have quite high expectations of their potential value. Low-code platforms make it look like building AI agents that automate entire workflows is a matter of a few prompts and several clicks. In reality, platform constraints often block the integration required for those much-coveted outcomes. That's where our client got stuck.

An **Australian consulting company** with an in-house IT team but no solid internal AI expertise tried experimenting with **Microsoft Copilot Studio** to fit their mostly MS-first software environment. Yet, the devil turned out to be in those non-Microsoft integration details. Copilot Studio's native connectors covered MS apps only, so anything outside the stack called for custom builds the in-house team didn't have the expertise to deliver.

Microsoft-bound agents made little sense, as a typical day of the client's sales representatives meant bouncing between Outlook and Teams for communication, HubSpot for client and deal records, OneDrive for shared documents, Jira for project tracking, Miro for brainstorming, Power BI for reporting, and LinkedIn for prospecting and messaging.

Three dead ends of Microsoft-centric agents in a heterogeneous stack were:

Static context

As agents couldn't track changes across the full toolset and refresh memory on their own, staff members had to **manually** collect updates from multiple apps and feed the agents.

Prompt sensitivity

Sales reps had to remember which prompts triggered the right responses, adding another layer of **mental burden** on top of managing pipelines.

Passive mode

Initiating the communication was always the humans' duty, which also took its toll on the mental workload and became **exhausting** for sales reps.

While the promise of AI agents was powerful, the reality looked more like **AI crutches that only added to the employees' cognitive load**.

Time went by, but the much-anticipated productivity boost never showed. Still, that setback helped the client **solidify their vision of a helpful AI agent** as one that can:

- "Walk" within their multi-vendor software ecosystem freely to pull together fragmented information into a dynamic, up-to-date project context
- Join the conversations of their core human team to make proactive suggestions and then act on them

With this idea in mind, they started looking for a tech partner to build a sales-focused AI agent. The bar was high: the client expected more from an **AI vendor** than acquaintance with popular AI frameworks, they wanted deep, hands-on experience. Instinctools, with their own agent-building platform, proved to be the right match.

Solution

Instead of forcing a client to adapt to a rigid AI framework, we used our very own technology-agnostic **GENE** platform to build a truly adaptive **autonomous sales representative** that each team member can configure to match their personal workflow.

- By default, the virtual worker is connected to the company's Confluence, Jira, HubSpot, Power BI, Miro, and OneDrive.
- Additionally, employees could drop the AI worker into any Teams chats and channels and give it access to their LinkedIn profile.

What is GENE?

GENE is **instinctools infrastructure middleware** for building secure-by-design AI agents and multi-agent systems faster than with bare-bones AI frameworks.

What did we leverage in GENE to build a better agent?

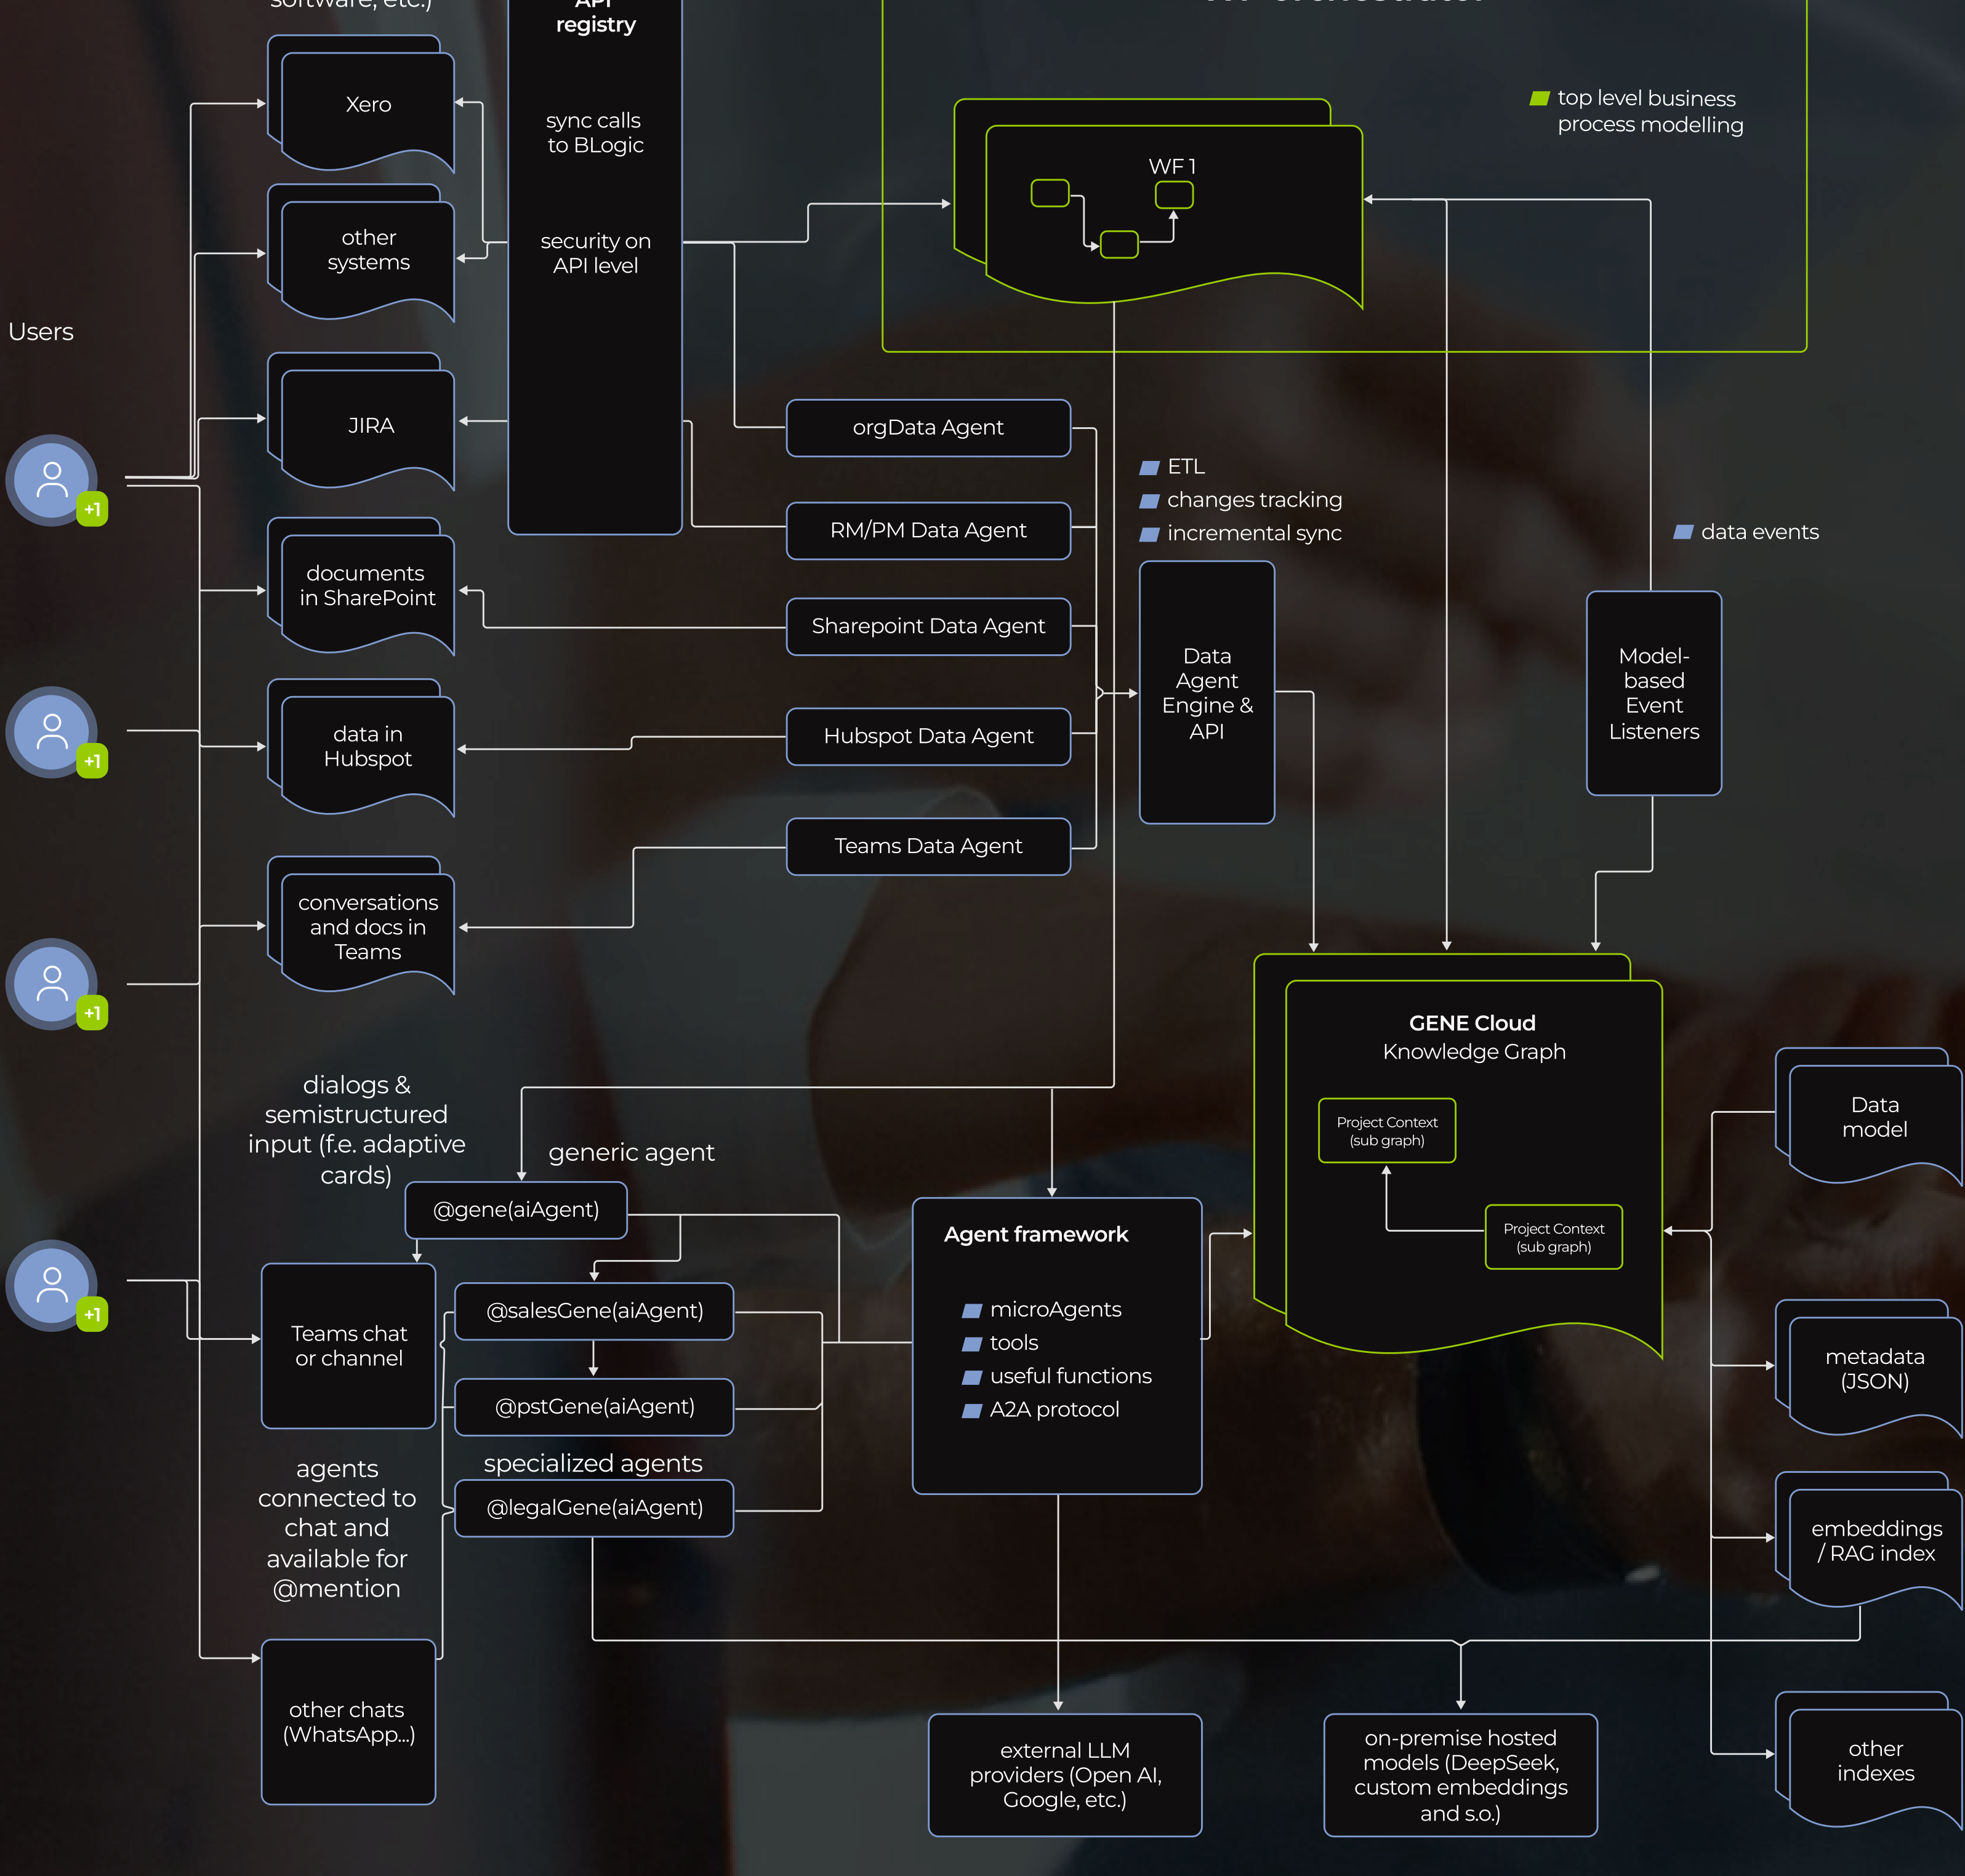
1. A single workspace for open-source AI frameworks like crewAI, LangGraph, LangChain, and others.
2. Built-in responsible AI mechanisms, including bias detection and mitigation, AI governance, and AI regulation compliance.
3. A library of pre-vetted APIs for a hitch-free agent integration into the client's software ecosystem.

Compared to building AI agents from scratch, GENE enabled our AI engineers to develop and integrate reliable, scalable agents **70% faster**. What typically requires 8-12 weeks with just frameworks alone, was condensed into **2-3 weeks with GENE**.

Here's how GENE addressed the client's past frustrations with a constrained low-code platform:

1. Providing built-in agents that automatically maintain tasks, workflows, and project context up to date, so users don't have to.
2. Eliminating repetitive prompting thanks to pre-configured workflows and agents' ability to "listen" to connected tools, allowing for context-aware actions before being asked.
3. Presenting a wide range of whitelisted APIs for AI agents to solve problems independently.

Under the hood of the virtual worker there is a **multi-agent system**.



Knowing the stakes, the client **put data safety and overall software security front and center**. Our **dedicated team** did too.

Standard measures, such as mandatory user authentication, access control, activity tracking, and data encryption at rest and in motion, were table stakes. We built on that with GENE's **AI-specific safeguards**, aligned with NIST AI RMF, Google's SAIF, OWASP AI, and other AI regulation frameworks.

Robust secret management

OAuth tokens, database credentials, API keys, and other secrets are stored in an isolated vault inaccessible to agents themselves. We implemented automated credential rotation, with a standard 30-day cycle that can be tightened to daily for highly regulated industries.

Agent-to-agent (A2A) communication protocol

It allowed agents to work together without exposing their memory, proprietary logic, or tools to each other.

Whitelisted AI tools

We strictly supervise which tools agents can use and what they can do. For example, the virtual worker can create a Jira ticket based on a conversation in Teams, but only human reps can delete it.

Before

- Microsoft Copilot Studio couldn't deliver the tangible value the client was after
- Agents rely on a static, manually updated context
- Agents add to the sales reps' cognitive load
- Passive AI is overreliant on humans to push it along

After


- Instinctools' technology-agnostic GENE platform made it possible to develop an autonomous sales representative that covered all the client's needs
- The AI sales representative keeps a living, self-updating context
- The autonomous virtual worker offloads routine tasks from humans
- Proactive AI agent acts on context and gets things done without interrupting human sales reps

Business value

- **Faster** agent development **without security trade-offs** thanks to reliance on a solid infrastructure middleware instead of bare-bones AI frameworks
- **+20%** leads processed per sales representative
- **+19% more** upsell/cross-sell initiatives
- **-15%** in cost per lead
- **-10 hours** of manual work per employee weekly

Multiplier effect

This success story points to a not-so-distant future where AI workers are graduating from assistants to full-scale teammates operating within your workflows. You can start with AI virtual workers as **personal secretaries** for staff members to ease their mental load. And then scale with **autonomous AI employees** contributing to the projects' progress on par with your human team. Best of all, you **get reliable employees on board right when you need them, omitting a time- and budget-draining hiring process**.



Do you have a **similar project idea**?