ivalua

Powering Procurement Transformation with Autonomous AI Agents

From One-and-Done Features to Autonomous Agents

In today's swiftly changing technological landscape, autonomous agents powered by Generative AI are not just a trend; they represent a transformative shift in how organizations approach efficiency and automation. At Ivalua, we recognize the immense potential of these technologies to revolutionize procurement processes, enabling unparalleled levels of productivity and strategic insight.

Traditionally, Generative AI features have operated on a «one-and-done» basis—delivering outputs that serve primarily as drafts or prototypes. However, as Large Language Models (LLMs) and LLM orchestration frameworks advance, we see Generative AI systems becoming increasingly adept at contextual reasoning and dynamic information retrieval. This evolution allows for the orchestration of autonomous agents that can collaborate seamlessly, validate each other's outputs, and engage users in meaningful ways. The result? A new standard of cognitive applications that elevate automation and efficiency to unprecedented heights.

Our product strategy at Ivalua is not merely to sprinkle in isolated AI features across our platform. Instead, we are embedding Generative AI at the very core of our solutions, transforming them into intelligent agents that function as the backbone of procurement operations. These agents will automate mundane tasks in real-time, allowing procurement professionals to focus on strategic decision-making and value creation.

As we navigate this exciting frontier, we are committed to building a robust framework that supports the next generation of autonomous agents. By leveraging multi-modal LLMs, we aim to create a system that adapts to the complexities of modern business environments, ensuring that our clients can respond swiftly to changing data, policies, and market conditions.

Market Trends Enabling the Evolution of Autonomous Agents

The emergence of autonomous agents powered by Generative AI in procurement is shaped by several key trends, each supported by real-world data and insights:

- **Data Processing and Scenario Analysis:** Generative AI is being used to crunch large data sets and process scenario-based results, reducing the need for complex manual interventions. This helps in increasing efficiencies and generating actionable insights based on historical trends, demand profiles, and supplier performance.
- **Rapid Improvement and Adoption:** The reasoning capabilities of generative AI models are improving rapidly, and the speed of these improvements is unprecedented. This rapid advancement is driving quicker adoption and integration into procurement processes.
- **Automation and Efficiency:** A significant portion of procurement tasks can be automated using generative AI. This includes automating RFI/RFP/RFQ generation, spend analytics, and contract lifecycle management. It's estimated that 50-80% of current procurement work can be automated or shifted to self-service models.
- **Enhanced Decision-Making:** Generative AI aids in handling broad, layered questions and drawing implications that lead to well-informed decisions. It integrates reasoning capabilities, which are crucial for making complex procurement decisions.
- **Integration with Existing Systems:** Generative AI is being integrated with digital procurement backbones, such as source-to-pay platforms, to connect internal data (e.g., spend categories, suppliers) with external data (e.g., category risk, supplier risk) and enhance functionalities like supplier discovery, contract management, and spend analysis. This integration helps in forming a comprehensive view of procurement operations.
- **Multi-Agent Frameworks:** Companies are developing multi-agent frameworks that allow various AI agents to work collaboratively, improving procurement and supply chain operations. These frameworks are expected to bring significant advancements in 2025.
- **Conversational and User-Friendly Interfaces:** The simplicity of conversational user interfaces for generative AI tools has driven their adoption, making it easier for procurement professionals to leverage these technologies without extensive training.

These trends collectively signify a profound transformation in procurement, driven by the capabilities of generative AI and autonomous agents, promising increased efficiency, strategic value, and higher returns on investment.

Definition of Autonomous Agents and Ivalua's Delivery

Autonomous agents powered by Generative AI are sophisticated systems designed to perform complex tasks with minimal human intervention. They leverage advanced reasoning capabilities through next-generation Large Language Models (LLMs) to analyze data, evaluate options, and make informed decisions. These agents operate within a multi-layered framework that integrates various data sources and orchestrates workflows, enabling them to adapt to real-world conditions and respond dynamically to changing business environments, without a static or prescribed configuration.

The Ivalua Platform has singularly accelerated the development and release of autonomous AI agents because of its unified, no-code, and privacy-first multi-instance architecture:

- **Data Security:** Multi-instance architecture isolates data, protecting sensitive information better than shared, multi-tenant setups.
- **No-code Configurability:** Multi-instance systems with no-code tools allow tailored AI workflows, unlike shared environments that limit flexibility.
- **Unified Data Model:** A single data model enables seamless, real-time AI processes and ensures accuracy across procurement tasks.

As such, with these features, the Ivalua Platform enables the secure, adaptable, and reliable deployment of autonomous AI agents for future-ready procurement organizations.

Ivalua's Commitment to Autonomous Agents

Ivalua is at the forefront of this evolution, embedding autonomous agents into its no-code platform to automate 40-70% of procurement tasks and streamlining processes such as RFI/RFP generation, intake management, category intelligence, and contract lifecycle management. The power of these agents lies in their ability to adapt to rapidly changing data, policies, and business contexts.

Here's how Ivalua is extending its no-code, unified multi-instance platform to deliver on the promise of autonomous agents:

- **Core Integration:** Generative AI is now central to the Ivalua Platform, providing all Ivalua solutions, with a common architecture and web services enabling a comprehensive AI-driven business operating system rather than a collection of isolated features.
- **Improved Decision-Making:** With the integration of advanced LLMs, Ivalua enhances decision-making capabilities, allowing for deeper reasoning and analysis that lead to better-informed procurement strategies, deeply embedded into data transformations, workflows, applicative clickstream and conversational experiences.
- **Dynamic AI Systems:** The AI orchestration layer effectively manages complex reasoning and data processing flows, enabling the agents to break down tasks, solicit or retrieve contextual information on-the-fly, and provide context-driven responses in real-time.
- **User-Friendly Interfaces:** Ivalua's conversational agents simplify the interaction with AI technologies, making it accessible for procurement professionals without the need for extensive training and opening the way for comprehensive, mobile-friendly, voice-driven conversational procurement.
- **Flexible Framework:** The strength of the orchestration service is that it is LLM-independent and allows for the integration of each organization's preferred generative AI models, ensuring compatibility with various corporate strategies and IT environments.

Concretely, to support the development of configurable and autonomous agents, the Ivalua Platform has evolved across three main layers:

- The **LLM (and Tools)** layer: Next generation Large Language Models (LLMs), such as the OpenAI GPT-o1 series, are trained to reason–as opposed to the first generation of LLMs that could mimic and spit out learned patterns from the past based on their training data. Instead, newer LLM systems are able to stop and think at inference-time, which means they are able to generate a range of possibilities, evaluate those possibilities and consider outcomes, and make decisions to better answer the question with logic, self-reflection, and structure. They represent the next frontier of Generative AI technology. As we see the evolution of those LLMs, we can expect more deliberate and deeper reasoning and inference, producing more structured and logical answers and analyses.

- The **Customer Application** layer powered by conversational business agents (or assistants): This layer deals with the messy, real world of workflows, business rules, and enterprise data. This is not something that generic foundation LLMs are trained to harness well. In this layer, intelligent business agents, powered by Generative AI, are able to execute specific task flows and achieve specific business objectives, like creating contracts, validating supplier documents, adding qualified suppliers to sourcing events, or even launching and managing quick bid events to support simple purchase requests, all done in a way that is as automated as possible. To achieve this objective, we allow configurators to implement their own agents via conversational skills for IVA (Ivalua's Intelligent Virtual Assistant). IVA skills are configurable sequences of instructions that allow IVA to solicit input from the user, update the database, navigate the user to a relevant page, search for information in documents, database tables, as well as leveraging its own generative AI functions to create, extract, translate, summarize, or analyze content. While the sequence of instructions was initially scripted like the logical and programmatic flow of an application wizard, Generative AI has now enabled the dynamic generation of contextual instructions that create more relevant and intuitive conversational experiences.
- The **AI orchestration** layer that turns reactive agents and enterprise knowledge into autonomous ones via proper planning and execution.

  While reasoning LLMs can take the time to generate structured thoughts via chains of thought leveraging their intrinsic knowledge and inference-time self-reflection, they have no knowledge of confidential enterprise knowledge, business logic, and workflows. At the same time, IVA skills can follow a logical sequence of instructions based on business logic and proprietary enterprise information, but they rely on relatively deterministic processing flows and data structures defined in advance. They lack the organic ability to define their own execution flow by dynamically breaking the larger problem into a series of intermediate steps and finding contextually relevant knowledge across internal and external sources to enable more precise action. This is where the AI orchestration layer comes into play to provide a reliable cognitive bridge between the messy world of skills, workflows, business rules, documents, and data with the general-purpose reasoning of foundation LLMs.

  The AI orchestration layer turns static and reactive conversational business agents into dynamic and autonomous ones. By leveraging the reasoning capabilities of more recent LLMs, the orchestration service breaks down complex requests into smaller information retrieval, document processing, content generation, and quality control tasks, which, in combination, define a step-by-step execution plan that produces a better response.

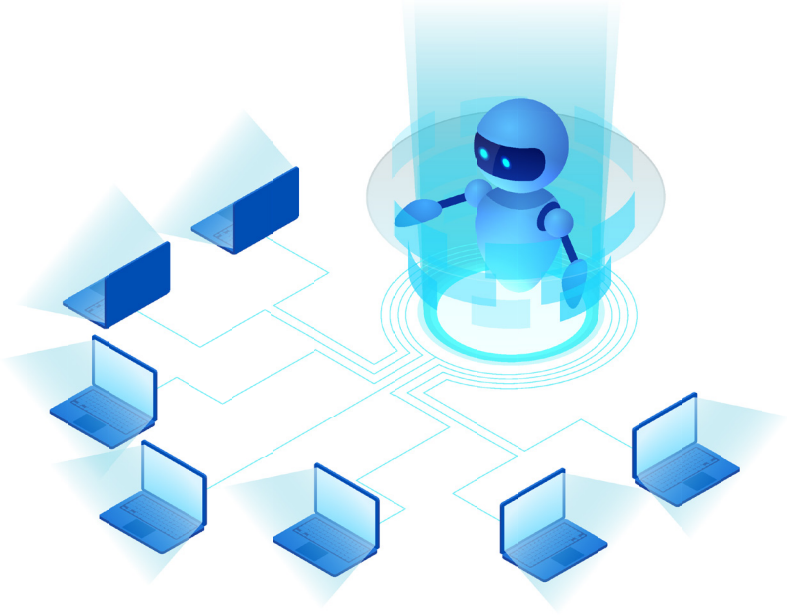

The orchestration layer implements three types of functions:

- An **Information Search and Retrieval** function able to retrieve real-time content from internal data sources (such as a vector store and database tables) and external data sources (such as the internet or third-party application APIs) to ground LLM responses with precise and contextually relevant knowledge.
- A **Document Processing** function able to extract structured information, classify, translate, summarize, combine, and compare content across multiple documents.
- A **Quality Control** function able to verify and validate the quality of the processing steps via self-check, error detection, and self-consistency techniques.

Figure 1: Ivalua Generative AI Framework

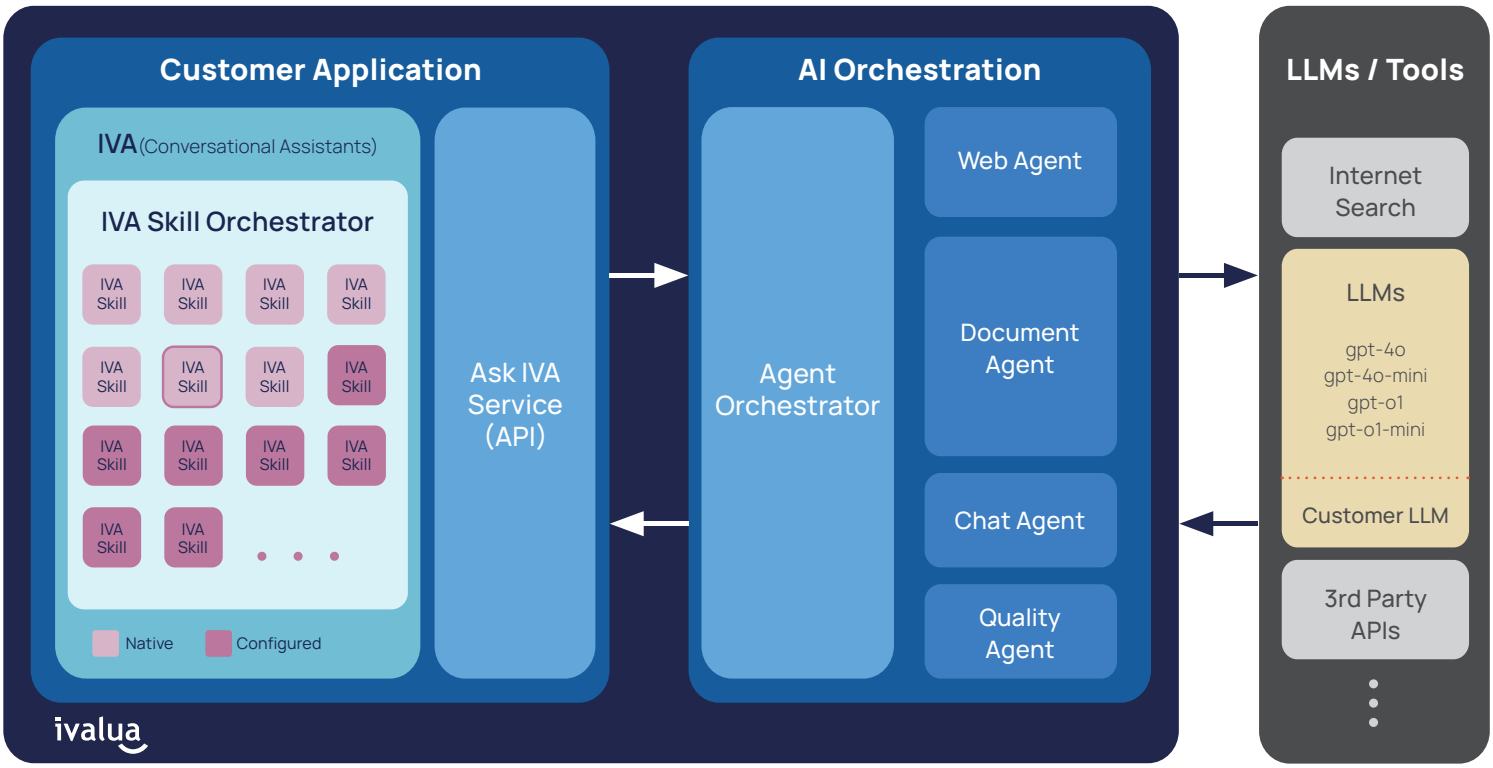

As a result, the system can support smarter and more advanced tasks, particularly those involving multiple document-related tasks, web searches, or a combination of both. For instance, it can extract relevant sections of a contract and then check whether they comply with internal procurement policies or country-specific regulations.

With these capabilities, the orchestration layer which separates the application agents from the backend LLMs, not only improves the overall quality and relevance of the response, but also makes the overall application system LLM-independent, allowing customers to plug their own system of LLMs. That way Ivalua as a platform provides solutions that are compatible with the enterprise generative AI strategy and leverage the corporate investment in specific LLM technology by their IT organization.

Ivalua's Commitment to Autonomous Agents

Our vision is not just theoretical; it is already being implemented with concrete solutions:

- **Intake Management:** IVA can understand user needs, guide them through relevant forms with context-specific questions, and automatically fill out the forms. This streamlines the intake process, reducing time and effort for users. This approach also sets the stage for a 100% voice-enabled, mobile-friendly conversational intake flow.
- **Conversational Analytics:** Users can ask about key performance indicators and receive detailed breakdowns by organization, commodity, supplier, or over time. The system provides conversational, real-time insights based on user queries.
- **Supplier Onboarding:** Our agents welcome suppliers, request specific information and documents, and guide them through the onboarding process. This ensures a smooth and efficient onboarding experience via the Ivalua Supplier Portal.
- **Contract Lifecycle Management (CLM):** IVA can search through tens of thousands of documents stored in the contract repository, answer user questions, and run detailed analyses and comparisons across contracts. This includes leveraging external references like regulations, laws, and public cases found on the internet.
- **External Workforce Management:** IVA helps write job descriptions, analyze candidate resumes, prepare interview questionnaires, and summarize interview notes. This ensures that organizations can efficiently manage their external workforce.

Figure 2: The New Source-to-Pay Experience: Conversational, Agentic, Autonomous

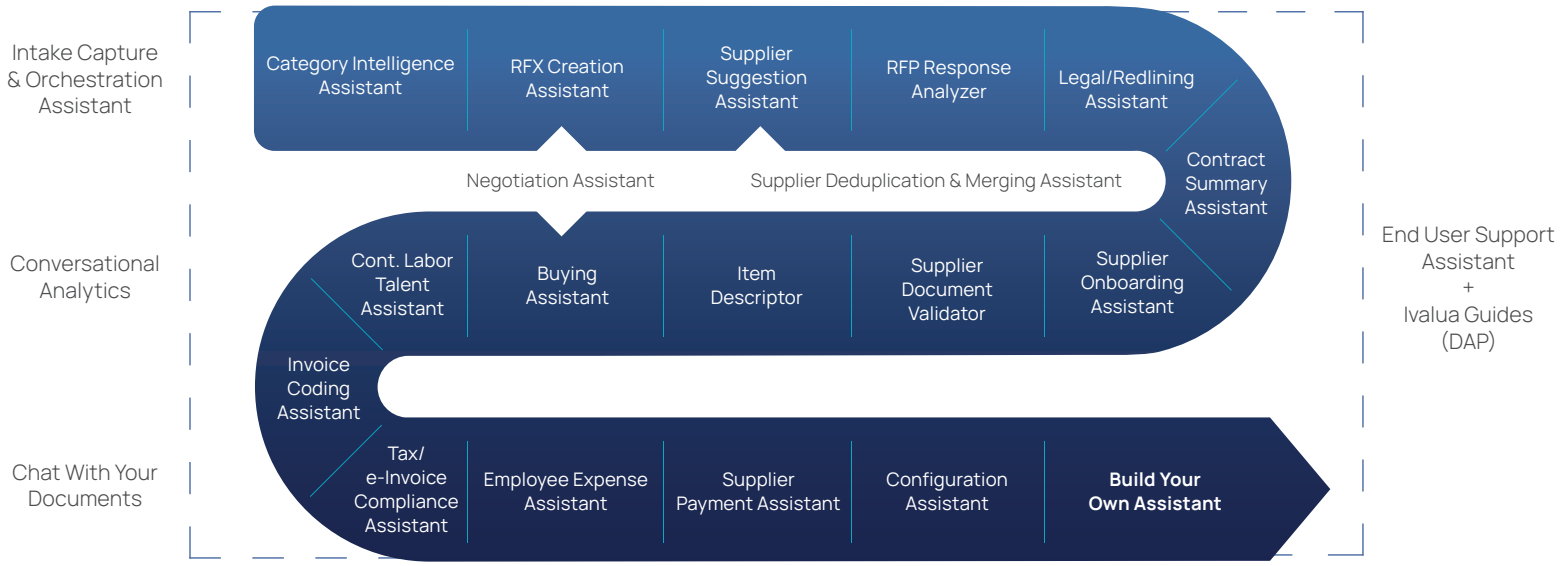

Those are just a few examples of what we have been able to achieve with an agentic framework for autonomous agents. It's only the beginning as opportunities to configure new autonomous agents come up in every discussion with customers and partners.

Conclusion

The evolution of the Ivalua platform to support configurable and autonomous agents powered by Generative AI represents a significant leap in sourcing and procurement technology. By integrating Generative AI at the core of its solutions, Ivalua is not merely adding incremental features but is fundamentally transforming how business processes are automated and optimized. The platform's layered approach—comprising intelligent business agents, advanced LLMs, and a robust orchestration layer—ensures that these AI-driven solutions are both powerful and flexible.

This strategic shift towards a more integrated and intelligent system enables Ivalua to deliver greater value, efficiency, and automation to its customers. Practical applications, such as enhanced intake management, conversational analytics, supplier onboarding, contract lifecycle management, and external workforce management, demonstrate the tangible benefits and readiness of these innovations. As the platform continues to evolve, it promises to remain at the forefront of leveraging Generative AI to meet the complex needs of modern enterprises, ensuring compatibility with diverse LLM technologies and aligning with corporate AI strategies.

AUTHOR

Pascal Bensoussan

Chief Product Officer

Pascal brings 23 years of experience leading product strategy and development at Enterprise SaaS companies and driving rapid growth by delivering innovative, highly-differentiated products and solutions. His solutions address high-scale and high-stakes enterprise challenges involving complex technology, big data, machine learning, and a large ecosystem of partners. A powerful advocate for the potential of technology to transform organizations, Bensoussan is a frequent keynote speaker and board member at several start-ups. He holds an M.S. in Engineering-Economic Systems from Stanford University, an M.S. in Economics from ENPC (France), and a B.S. in Mathematics from Ecole Polytechnique (France).

About Ivalua

Ivalua is a leading provider of cloud-based, AI-powered Spend Management software. Our unified Source-to-Pay platform empowers businesses to effectively manage all categories of spend and all suppliers, increasing profitability, improving sustainability, lowering risk and boosting employee productivity. We are trusted by hundreds of the world's most admired brands and recognized as a leader by Gartner and other analysts. Learn more at www.ivalua.com.

Follow us on LinkedIn and @Ivalua.