SKYSTUDIO

WORKFLOW

This technical white paper explains the design of SkyStudio Agents: autonomous agents that select appropriate workflows, utilize tools effectively, and learn continuously to achieve a user-defined goal.

SKYMOD

Prepared for
SKYMOD TECHNOLOGY

Table of Contents

Introduction

SkyStudio Workflow

In the corporate world, “automation” still often means a step-by-step process based on manual approvals and copying rules. Workflows are fragmented; data lives in different systems, decisions live in people’s memories, and tool integrations are brittle. The result: high manual effort, inconsistent quality, delayed decisions, and workloads that can’t scale.
SkyStudio Workflow is designed to close this gap: it runs your processes with an agent-based (agentic) architecture designed on a natural-language-enabled visual 2D canvas; it delivers not merely a chat that “produces answers,” but a system that pursues a goal and takes action.

The Invisible Barriers to Full Automation

In enterprises, the main barriers to full automation are the continued dependence of decisions on people, brittle integrations, and information disconnects. **Steps such as SLA (Service Level Agreement)** and condition checks, data validation, or who to escalate to mostly rely on personal experience; this reduces consistency and speed. Because many systems are connected with **separately written scripts, integrations easily break with every change.** Moreover, even if the right information exists in documents, it cannot be incorporated into the flow as “context,” so decisions proceed not in a data-driven manner but in a fragmented way. As the number of steps and variations increases, operations become unmanageable, and quality and scalability are severely compromised.

Design your workflows on a visual canvas, and scale and automate AI agents while keeping your data secure.

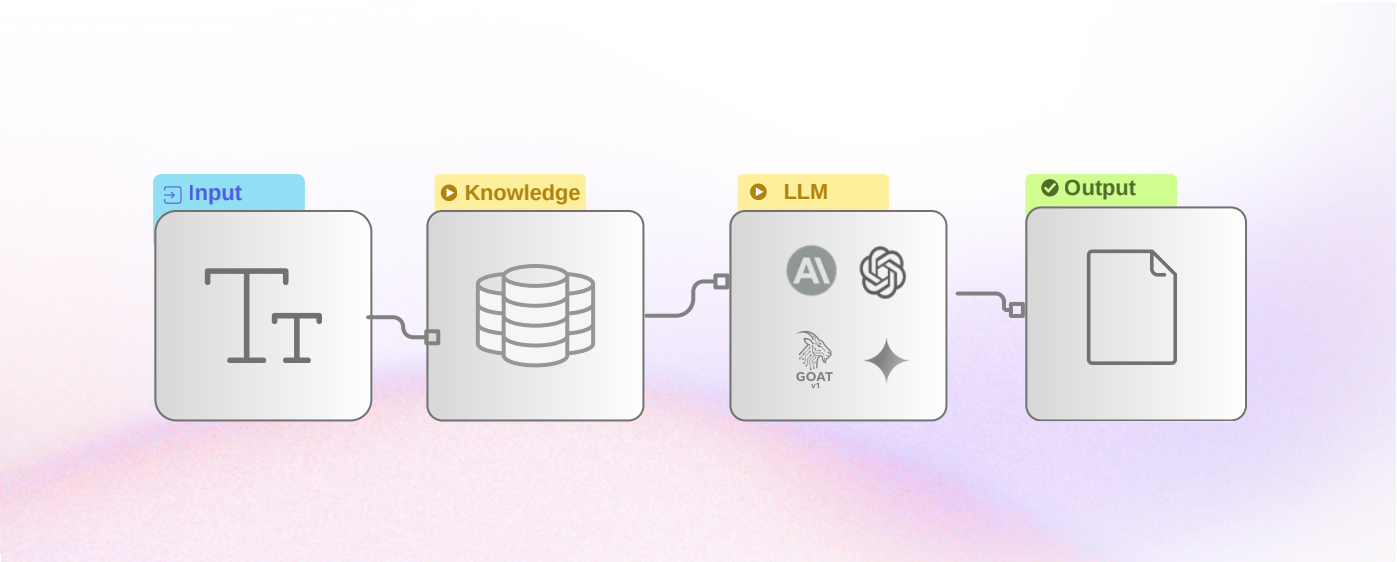

Features and Benefits

- **Visual Canvas & Drag-and-Drop:** SkyStudio's visual canvas lets you build workflows by dragging and dropping nodes—no code required. This approach provides genuine workflow control, deep model flexibility, and one-click integration management. On the canvas, you can combine LLMs, APIs, and rules within the same agent and visually organize every step.

- **Agent Design with Natural Language:** By simply describing the task, you can assemble the right workflow; the system leverages your existing process knowledge to design the agent's logic and safely carry out data-driven automation.

- **Automation & Orchestration:** SkyStudio brings together enterprise data, actions, and third-party components to select the most appropriate tools, triggering tasks on a schedule or event-basis. As a result, large-scale processes become easy to manage.

- **Flexibility & Customization:** Choose among different LLMs or integrate your own model. Customize names, logos, backgrounds, info boxes, and alerts to align with corporate branding and regulatory requirements.

- **Security & Compliance:** The platform treats data security as a top priority and provides an architecture aligned with standards such as SOC 2 and GDPR. It operates on infrastructure with ISO 27001 and SOC 2 Type II certifications; access is governed through single-tenant options, SSO, and granular permission controls. If desired, data can be deployed in your own cloud or data center.

- **Integrations & Connectivity:** With a broad integration ecosystem that indexes hundreds of apps, you can pull documents from sources such as SharePoint, OneDrive, Notion, Confluence, and Google Drive, and interact with agents via Slack or Microsoft Teams. You can unify data on a single platform through APIs or your own vector database. Thanks to the one-click integration architecture, it's easy to connect with 100+ tools. For enterprise systems like CRM/ERP/ITSM, it provides rich connectors for both reading (RAG) and writing (actions), together with strong security/governance and production-grade operational features.

- **Human-in-the-Loop Controls & Version Tracking:** With built-in memory, feedback, and versioning features, you can require human approval for critical operations, prevent errors, and track version history. Agent-level monitoring and logs enhance the auditability of processes.

SkyStudio Workflow: Enterprise AI-Powered Workflows

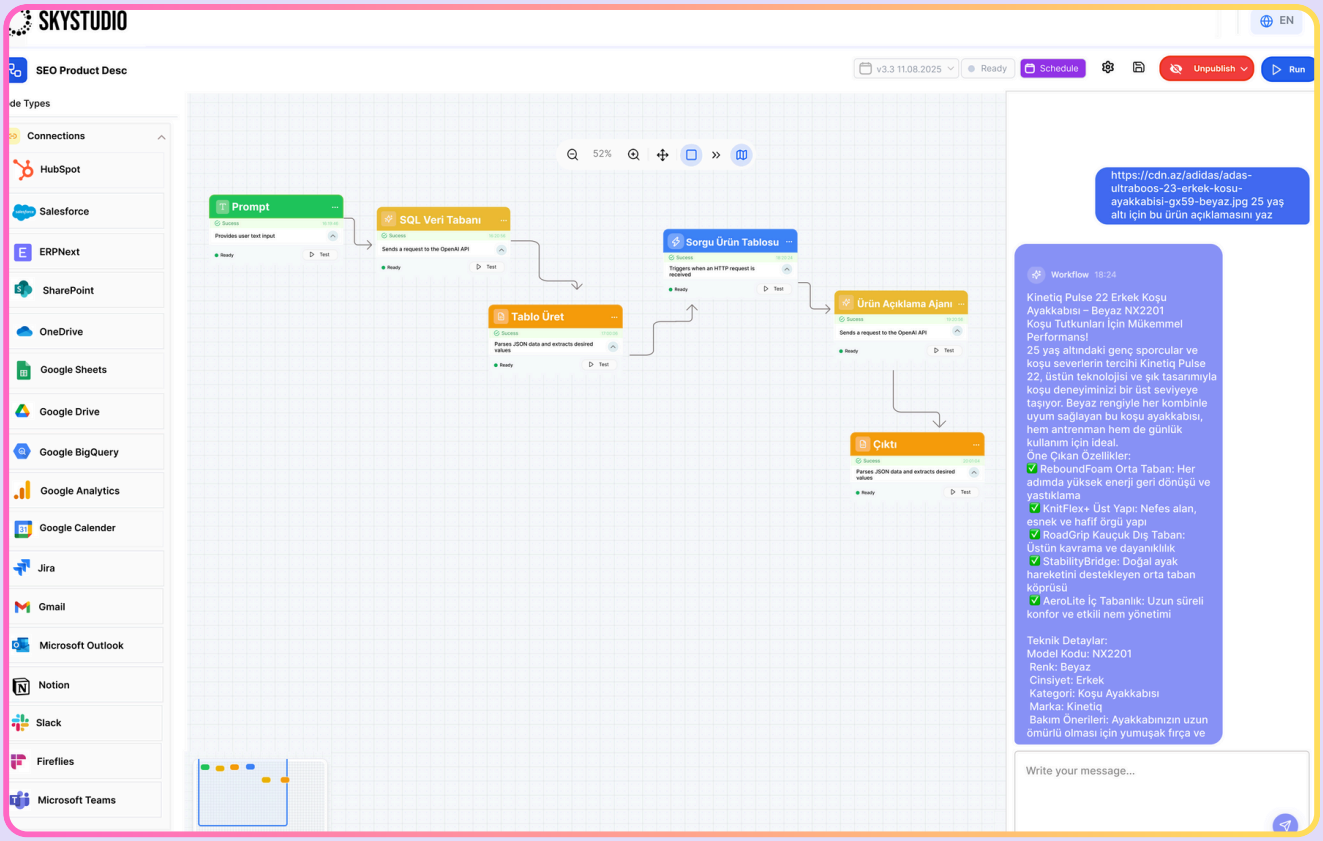

How It Works

1. Define the Process: Identify the business process you want to automate and gather the relevant data.

2. Design on the Canvas: Connect LLM modules, data sources, API actions, and control points by drag-and-drop. When you describe the task in natural language, the system uses your process knowledge to automatically compose the flow; you remain in full control.

3. Set Up Integrations & Security: Connect to apps such as HubSpot, Slack, Google Drive, and OneDrive with one-click integrations. Manage access with SSO and role-based permissions. If you prefer, host the platform in your own cloud to preserve data sovereignty.

4. Test & Deploy: Test the workflow; add human-in-the-loop steps; manage versions with feedback. When ready, roll it out to your employees or customers; operations run automatically and generate recommendations5

5. Monitoring & Optimization: Track performance with agent logs and analytics. Update the design as needed, add new integrations, or improve models. An expert team is with you at every step.

Experience the Workflow of the Future Today

SkyStudio Workflow brings together the power of a visual canvas and natural language to transform enterprise workflows. If you want to scale AI capabilities across your workforce and automate your processes while keeping your data secure, meet SkyStudio.

Build Your Own Agent (Workflow)

Key Features

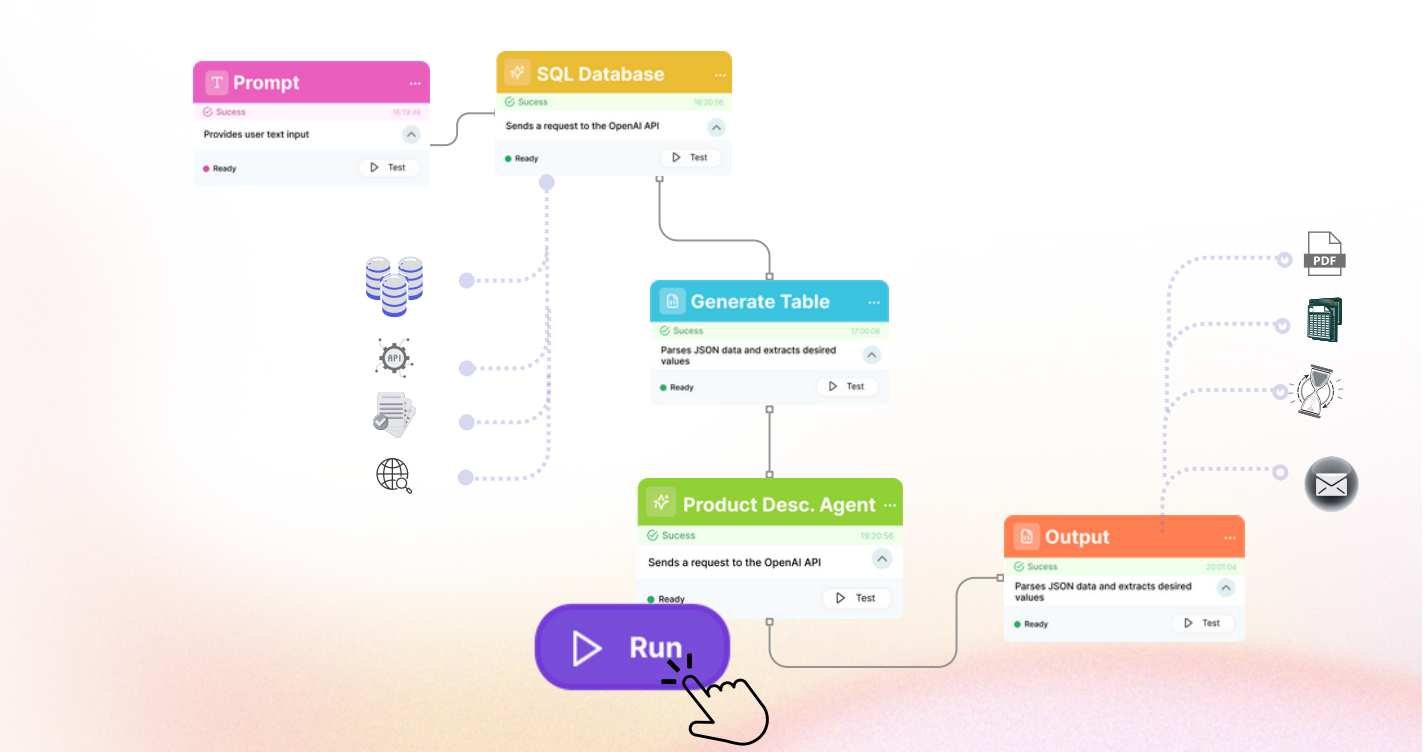

RAG

In SkyStudio Workflow, RAG will be the primary pattern for grounding answers in the most up-to-date and verifiable content rather than leaving them to “training intuition”: we ingest–parse–chunk–embed the data and store it in a vector database, turn the query into an embedding, and retrieve the most relevant passages to provide context to the LLM. This approach is fast, reduces the risk of erroneous actions, and extends across multiple data types.

The vector database can query millions of embeddings with very low latency, enabling **“search by meaning, not just keywords.”** With hybrid search, we find semantically close results in content chunks **(chunk search)** while preserving exact term matching in document titles/tags **(doc name search).**

We also address the recall/precision balance with “retrieve & re-rank”: first we take a broad candidate set, then select the most relevant context with an LLM or a reranker.

For more complex tasks, **Agentic RAG** comes into play: the agent establishes a multi-step plan with a ReAct-style reasoning–action loop, calls multiple tools in sequence, and updates its plan with the results obtained (e.g., web/intranet search), thus elevating a plain “retrieve-generate” flow into real problem solving.

In scenarios where relationship knowledge is critical, we use **GraphRAG** to leverage concept connections over a knowledge graph; this strengthens multi-hop inference beyond simple passage matching.

Data Layer

URL scraping (web scraping): Retrieves page content via HTTP requests or a headless browser; adheres to robots.txt/sitemap rules, manages pagination and rate limiting, and processes dynamic JS content.

API integrations (REST/GraphQL/SOAP): Authorizes with OAuth2/API keys; enables secure, reliable data retrieval with pagination, rate limiting, and error/retry policies.

Veritabanlarına erişim (SQL/NoSQL): Postgres/MySQL/MSSQL/Oracle ile okuma-yazma; Mongo/Redis gibi NoSQL kaynaklara bağlanma; connection pooling ve şema-eşleme ile tutarlı veri akışı kurar.

Enterprise application connectors (SaaS connectors): Retrieves permission-aware data from systems such as Google Workspace, Microsoft 365 (SharePoint/OneDrive), Confluence/Jira, and Salesforce/HubSpot.

Document extraction & OCR: Extracts text, tables, and diagrams from PDF/DOCX/PPTX/XLSX/HTML/MD documents; detects fields and tables in scanned content with OCR.

Email & calendar ingestion: Retrieves email bodies and attachments via IMAP/Gmail API/Exchange; processes ICS data to link meeting and attendee information into the workflow.

Webhook & event ingestion: Securely accepts HTTP callbacks from external systems; applies signature/verification and rate limiting.

Scheduling & Batch Processing: Runs periodic ingestion and backfills via cron/schedulers; respects maintenance windows.

Kimlik, izin & gizli anahtar yönetimi: RBAC/ABAC ile izin-farkındalıklı getirme yapar; Vault/Secret Manager ve IP allowlist/denylist ile güvenliği güçlendirir.

Identity, Authorization & Secrets Management: Enforces permission-aware access with RBAC/ABAC; hardens security using Vault/Secret Manager and IP allowlists/denylists.

Compliance & Ethics: Ensures compliance with KVKK/GDPR, copyright/licensing and robots.txt directives; enforces data residency, audit trails, and data retention policies.

Protocols & Standards: Standardized connectivity to tools and data sources via MCP (Model Context Protocol); agent-to-agent (A2A) messaging.

Agentic Orchestration

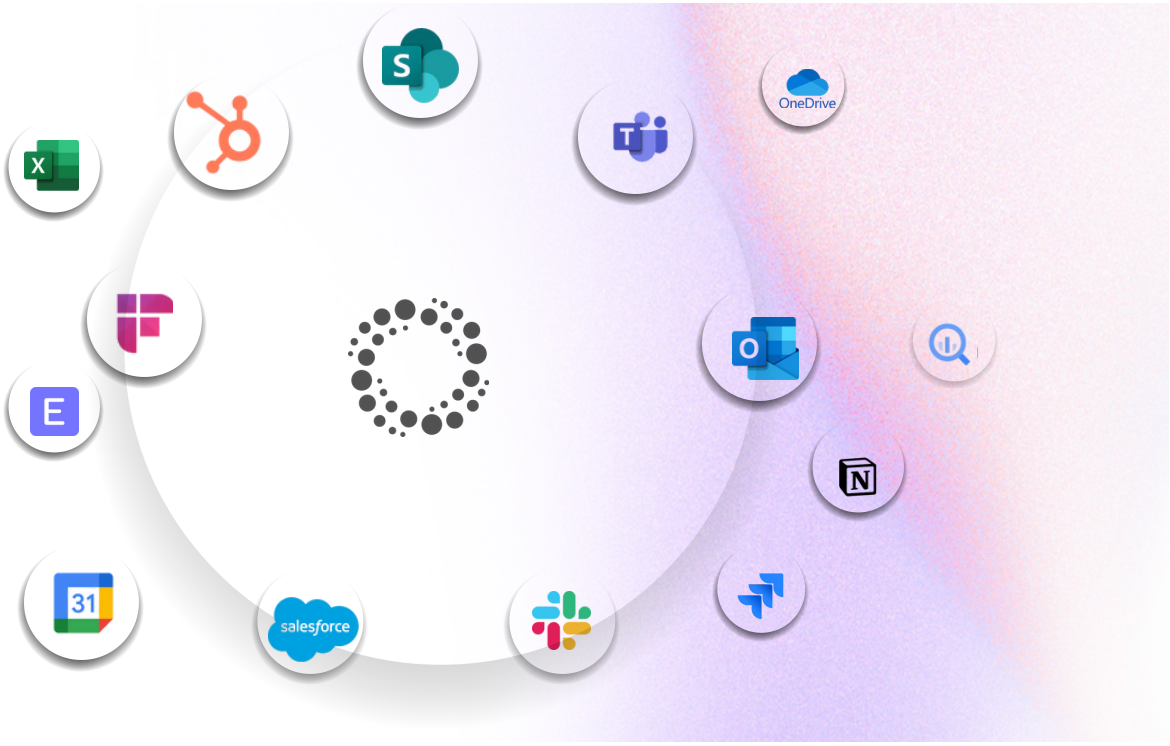

Agent-based Orchestration (ReAct)

Agent-based Orchestration (ReAct) is a model in which an agent operates in a **Reason → Act → Observe** loop to achieve a goal.
Reason : The agent analyzes user intent and the current context, derives a plan, and decides which tool (API/DB/search, etc.) to invoke.
Act : It calls the selected tool in line with defined input–output schemas (e.g., retrieve documents via RAG, read a record from the CRM, send an email).
Observe : It writes the tool output to memory, updates its plan, and repeats the loop if needed. The loop ends when the goal is met, a confidence threshold is reached, the step/latency/cost budget is exhausted, or a human-in-the-loop intervention is required.

Unlike ordinary single-turn chat, this approach combines multi-step reasoning with tool use; as a result, it not only generates answers but also executes actions (e.g., opening tickets, updating records, producing reports).

ReAct in SkyStudio Workflow:

- Safeguarded with guardrails (max steps, time/cost budgets, allowlisted tool set, schema validation).
- Backed by layered memory (RAG for retrieved context, working memory for steps, transaction log/ACID guarantees for critical operations).
- Robust error handling (on timeout/tool failure: retry → alternate strategy → human handoff).

Brief example: The user asks, “Why is my return request pending?”
The agent plans (Reason), fetches the order from the CRM and the policy document via RAG (Act), then evaluates the conditions (Observe). If the conditions are met, it opens the return record and sends a confirmation message to the customer; otherwise, it escalates automatically. Throughout, every step is logged for auditability and budget limits are enforced.

Integrations & MCP Support

SkyStudio Workflow offers a broad integration portfolio that enables agents to read data and perform actions (write) securely and at scale in enterprise scenarios. At the core, tools and data sources are connected under a standard contract via MCP (Model Context Protocol); where MCP is unavailable, bridges via **OpenAPI, GraphQL, JDBC/ODBC, or gRPC** are used. All retrievals and actions are executed in a permission-aware manner (RBAC/ABAC, with row/field-level controls). Defaults include audit logging, data residency controls **(KVKK/GDPR)**, and identity federation **(OAuth2/OIDC/SAML/SCIM).** For on-premises needs, options include **VPC peering, VPN, allowlists, and webhook/CDC (Change Data Capture).**

MCP-Native Operation:

SkyStudio operates both as an **MCP clien**t (consumes tools/resources/prompts hosted on external MCP servers) and an **MCP server** (exposes its own tools/prompts to external agents); schema validation, version/policy management, and permission delegation are built in.

Integration Breadth (Enterprise-Focused):

- **CRM:** Salesforce, Microsoft Dynamics 365, HubSpot (read: customers/sales/activities; write: lead/opportunity/task).
- **ERP:** SAP S/4HANA, Oracle E-Business Suite, NetSuite (read: inventory/orders/invoices; write: orders/pricing/delivery notes).
- **ITSM & Project:** ServiceNow, Jira/JSM, Confluence (read: knowledge base/tickets; write: open/update tickets, comment, reassign).
- **Productivity & DMS:** Google Workspace, Microsoft 365 (Drive/SharePoint/OneDrive/Outlook/Teams/Slack).
- **Data & Analytics:** Snowflake, BigQuery, Redshift, Databricks;
- **DBs:** PostgreSQL, MySQL, SQL Server, Oracle;
- **NoSQL/Cache:** MongoDB, Redis.
- **Storage & File:** S3, GCS, Azure Blob, SFTP/FTP (versioning, large files, MIME/OCR).

Features

- **Read (RAG) integrations:** Documents, email/chat/knowledge bases, and data warehouses are indexed in a permission-aware manner; context is assembled via a hybrid of vector search + keyword search + knowledge graph, with a reranker.
- **Write (action) integrations:** Executes idempotent create/update/close operations in CRM/ERP/ITSM; safeguarded with ACID transactional logging, compensation steps, and dry-run support.
- **Security & governance:** OAuth2/OIDC/SAML/SCIM; secrets management (Vault/Secret Manager); IP allowlists/VPN; encryption in transit/at rest; PII masking/DLP; audit trails and retention policies.
- **Multi-agent (enterprise orchestration):** A2A enables inter-agent delegation and collaboration; supervised automation with roles/quotas/approval flows (human-in-the-loop) and policy guardrails.

In short: SkyStudio uses MCP as a “universal adapter,” while providing rich connectors for both reading (RAG) and writing (actions) across enterprise systems like CRM/ERP/ITSM—backed by strong security/governance and production-grade operations. This enables agents not only to find the right information but also to update the right system correctly, automating the process end to end.

LLM Selection in SkyStudio Workflow

Principle: SkyStudio does not rely on a single model. Through a multi-LLM (multi-model) strategy and task-based routing, it balances quality, latency, and cost; each request is routed to the right model based on its content. Enterprise requirements (KVKK/GDPR, data residency, access policies) and SLOs (accuracy, p95 latency, cost/turn) are first-class inputs to the selection.

Model Classes

- **Reasoning LLM:** Multi-step planning, ReAct + tool invocation, complex business rules.
- **Speed/Throughput (Fast/Flash) LLM:** Short Q&A, summarization, form filling; low latency and cost.
- **Vision/Multimodal LLM:** Reading PDFs/tables/diagrams; extracting fields/tables from screenshots and documents.
- **Function Calling / JSON-Strict LLM:** For integrations requiring strict schema fidelity (writes to CRM/ERP/ITSM).
- **Local / On-prem LLM (TR-focused):** For data residency and isolated networks; a local model such as Skymod/GOAT.
- **Embedding Model:** Vectorization for hybrid retrieval (vector search + keywords).
- **Reranker (Cross-encoder):** High-accuracy ranking of retrieved passages.

Routing Policy (Example)

- **Simple Q&A / short summaries** → Fast LLM
- **Planning + multiple tools / ReAct** → Reasoning LLM
- **PDF/Table/Diagram reading** → Multimodal LLM (+ table extraction tool)
- **CRM/ERP/ITSM writes (schema-critical)** → JSON-strict LLM (+ schema validation/guardrails)
- **Long ID / numeric matching** → Lexical + Reranker for context → then Fast/JSON-strict LLM
- **Sensitive/data-residency needs** → Local on-prem LLM (Skymod/GOAT)

Guardrails & Resilience

Schema/regex validation, tool allowlists, max-steps/max-cost budgets, automatic rollback (compensation), fallback paths (cloud → on-prem/local), and caching (prompt/response/embedding). To reduce false positives, LLM-judge plus human-in-the-loop (HITL) approval gates.

Evaluation & Continuous Improvement

Offline test suites (Q&A, data analysis, action/write), live A/B experiments, and metrics including answer exactness (especially serial/SKU/ID matching), grounding accuracy, function-call success rate, p95 latency, and cost per turn. Judge variance is monitored, and routing/policy tuning is performed using usage data.

Summary: Rather than a "single model," SkyStudio adopts the right-model-for-the-job approach: a portfolio of Reasoning + Fast + Multimodal + JSON-strict + Local LLMs, combined with hybrid retrieval and robust guardrails, achieves both accuracy and SLO targets at enterprise scale.

Memory

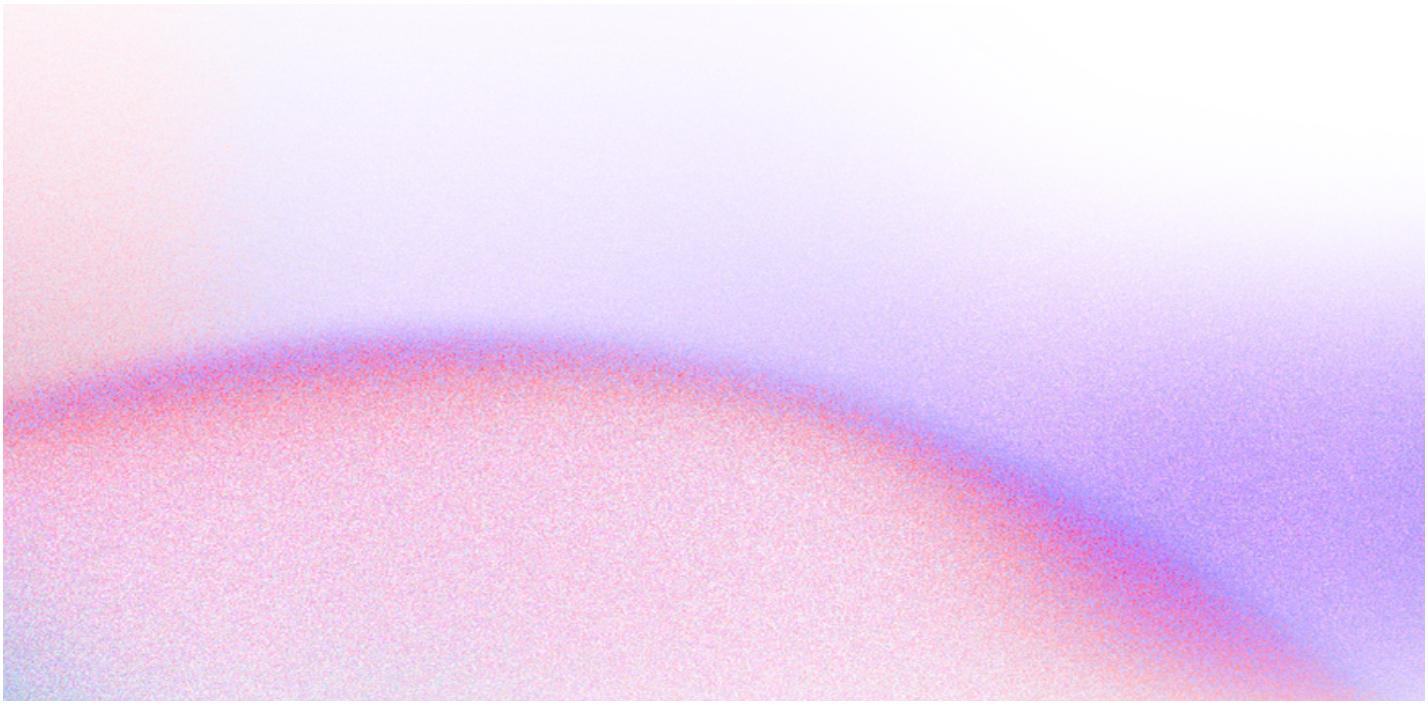

SkyStudio Workflow uses two primary memory approaches in tandem to ensure agents both think with the right context and deliver a consistent, auditable experience at enterprise scale: **Conversation Memory** (persistent conversation memory) and **Sliding Window + Input Memory** (ephemeral working memory). The former securely preserves user-specific long-term habits/preferences, while the latter feeds the most up-to-date context to the model each turn with low latency.

Conversation Memory:

Maintains durable user context across multiple sessions (preferences, style, language, frequently used data sources, approved identities). Messages/interactions are stored with event timestamps and metadata (channel, topic, project IDs).

Sliding Window + Input Memory:

On each request, the model receives the last k turns (rolling window) plus the current user input and—if needed—the most recent tool outputs. The window auto-adjusts to the token budget; when the limit is exceeded, earlier steps are summarized (role-aware summarization) with source attributions preserved. For dense artifacts such as tables or code, smart clipping (snippet extraction) is applied and critical regions are pinned (pinned system context).

Working with RAG (Grounded Context):

Conversation/Sliding Window memories work in concert with the long-term knowledge store: content retrieved from documents, emails, knowledge bases, and data warehouses is added to the context via RAG, allowing the model to combine personal/project context with verifiable sources. Hybrid retrieval and a reranker are used for search.

Conversation Memory provides personalization and continuity, while Sliding Window + Input Memory guarantees up-to-date accuracy and low latency. Together with RAG, SkyStudio agents blend personal/project context with fresh, verifiable information to produce safe, measurable, production-grade outputs.

Example Agent Scenarios for Enterprises

Returns / Complaint Processing (CRM + Policy Check)

The agent validates the contract, creates records in CRM/ERP, provisions accounts, and sends a welcome email escalating to a human only if policy checks fail.

Proactive IT Monitoring & Auto-Remediation (SRE/DevOps)

İzleme alarmıyla tetiklenir; Cloud Logging'den logları toplar, gerekirse GKE'de ilgili pod'u yeniden başlatır ve sonucu Slack'e bildirir.

Stock / Warranty Lookup in Customer Support (Real-time)

Functions like check_inventory and check_warranty_status return real-time inventory and warranty status, which is then shared with the customer instantly.

Automated Customer Onboarding (Multi-step Orchestration)

Opens the account (via API), sends a welcome email, verifies in the database that the first step is complete; if not, triggers a reminder.

Lead Qualification in Sales (CRM Enrichment + Assignee Decision)

The agent enriches a new lead via external APIs and the internal CRM; it checks whether the lead is an existing customer, then decides whether to assign it to a senior sales rep or add it to a nurture sequence.

Enterprise Content Search & Team Automation (Agentspace)

Connects to apps like SharePoint, Google Workspace, and Jira to perform multimodal search across multiple data sources; teams use a no-code interface to build their own agents and automate tasks such as meeting agendas and summaries.

Example Agent Scenarios for Enterprises

Calendar Automation from Email Context

Agents with A2A capabilities use Gmail context to automatically schedule a Zoom meeting, update Google Calendar, and notify participants.

Knowledge-Base Support (RAG / Vector Search)

Product descriptions, warranty details, and FAQs are converted into embeddings and stored in a vector database; semantic search returns more relevant answers.

IT destek/bug triage (Jira entegrasyonu)

“Software Bug Assistant” tipi ajan; kullanıcıyı doğrular, kod tabanında arama yapar ve uygun proje yönetim sisteminde (örn. Jira) ticket açar.

Compliance / Knowledge Base / Audit-Focused Search (Managed RAG)

Opens the account (via API), sends a welcome email, verifies in the database that the first step has been completed; if not, triggers a reminder.

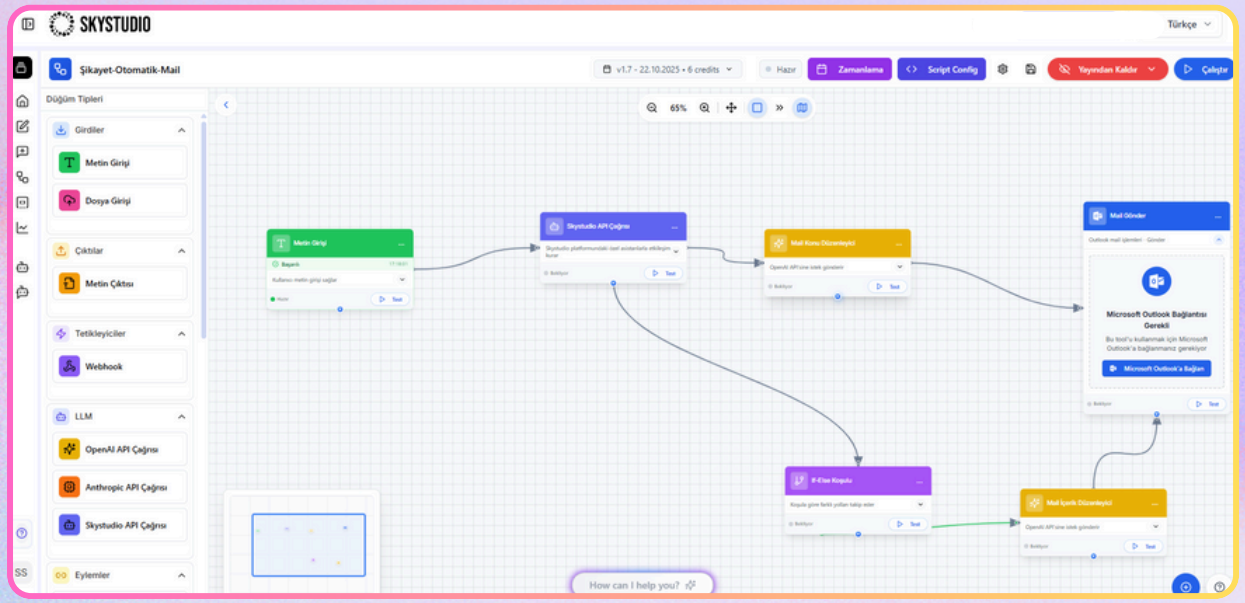

Conclusion

SkyStudio is a pioneering platform that empowers AI-driven business workflows. With a vision to make enterprise structures more efficient and agile, it enables organizations to manage their day-to-day operations. This paper examines the design of the agents developed by SkyStudio. These agents improve their processes by selecting workflows aligned with user goals, utilizing tools, and learning continuously. By engaging with SkyStudio's distinctive approach, you will explore the frontiers of enterprise innovation.

SKYMOD
TECHNOLOGY

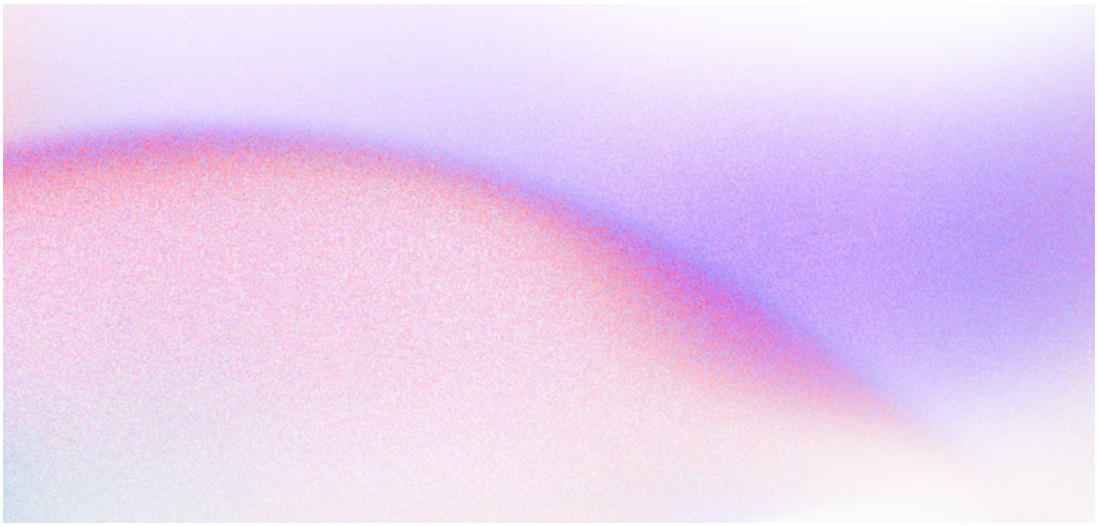

Have a question?

Contact us 👆

Website

www.skymod.tech

Email

hello@skymod.tech

Linkedin

Skymod Tech