

Think.AI — Platform Capability Comparison

Red indicates a meaningful gap vs. Think.AI for enterprise buyers. Accurate as of June 2026.

	Think.AI	Copilot Studio	n8n	LangGraph	Custom dev
WHO BUILDS IT					
Code or no-code?	No-code	No-code Power Platform authoring environment	Low-code Visual builder; code required for real workflows	Code Python / JS	Code Ad-hoc, from scratch
Engineers required?	No Platform training only	No Power Platform training required	Yes JS/Python for anything real	Yes ML/AI engineers	Yes Dedicated AI/ML team
How is intent specified?	Natural language No platform ceiling	Visual topic + flow designer NL authoring available in Agent Builder	Visual node graph + code	Code (Python/JS)	Code (ad-hoc)
Expressive power	Rich NL — no platform ceiling	Moderate Bounded by Power Platform connector library	Moderate Node library + custom code	Rich Full code control	Rich Full code control
LLM Agnostic	Yes Any LLM in your tenant	Primarily GPT-4 via Azure OpenAI; limited third-party model support via Azure AI Foundry	Yes — any LLM via API	Yes — framework is model-independent	Yes — you choose
WORKFLOW FIT					
Type of work	Repeatable human workflows Process or SOP-driven, variable inputs, judgment at key steps or throughout	Conversational agents + task automation M365-native; expanding to external systems via MCP	Rule-based automation between systems AI available at specific nodes	Dynamic agentic tasks Developer designs the judgment/deterministic mix	Whatever you build No default orientation
Good for workflows where human judgment is...	Some or Heavy	None → Some	None → Some	Some → Heavy	None → Heavy

	Think.AI	Copilot Studio	n8n	LangGraph	Custom dev
RELIABILITY					
Who detects / fixes errors?	The platform	Mostly the platform	Your developers	Mostly your developers	Your developers
Published benchmark	99% AI reliability ITSM benchmark, Feb 2026	None No published AI output benchmark	None You implement reliability	None You implement reliability	None You implement reliability
Time to reliable production	~3 months	Unknown	Months Depends on dev team	1 year+ <10% make it	1 year – never <10% make it
DEPLOYMENT & DATA SOVEREIGNTY					
Deployment model	Managed private instance In your hyperscaler tenant	Microsoft-hosted Data in your M365 tenant; geographic residency controls (US/EU); processed on Microsoft infrastructure	Self-hosted or n8n cloud Self-hosted: full data sovereignty, you run it; Cloud: data on n8n servers	Multiple options Cloud SaaS, BYOC (your VPC, AWS only), or fully self-hosted enterprise	Your infra You own and run everything
Who owns security & compliance?	The platform Within your tenant boundary — your security controls apply	Microsoft + you Microsoft Purview governs; your M365 tenant boundary; requires compliance configuration	Your team Full control if self-hosted; n8n team if cloud	Your team You configure and maintain all controls	Your team You own everything
Who incorporates AI updates?	The platform	Microsoft	Your developers	Your developers	Your developers
SYSTEM REACH					
Works with any enterprise system?	Yes Any Application via MCP	Yes MCP enables external APIs and third-party systems beyond M365	Yes 400+ integrations incl. SAP (partnership May 2026)	Yes Any system via code	Yes Any system via code
AGENTIC RUNTIME / HARNESS					
What is built in?	Fully built in and platform-managed Reliability, error handling, human oversight, AI updates, deployment — all platform responsibility	Harness built in Memory, multi-agent, error recovery, human-in-the-loop — within Microsoft infrastructure	Components available; you assemble Error handling, retry, memory nodes, approval gates — you configure and wire everything	Low-level primitives only State management, memory, graph execution — reliability and deployment built by your team	Nothing built in Every layer built from scratch by your team
SCALE ECONOMICS					

	Think.AI	Copilot Studio	n8n	LangGraph	Custom dev
Scaling to new workflows / apps	Sub-linear & compounding Reusable primitives — each new workflow faster than the last	Consumption grows Per message / per credit — costs scale with usage	Linear Each new workflow needs dev time	Linear Each agent needs engineering cycles	Linear Each workflow built from scratch
Time to initial demo	A day to an hour	Hours to days	Days to a week	A week to a day	A few weeks to a week

Copilot Studio MCP GA March 2026. LangGraph BYOC available on AWS only. Think.AI supports Claude and other LLMs deployed in your tenant.