



AI-Powered Intelligent Automation for Anaptyss

Transforming Operations Through n8n + AI Integration

1. Executive Summary

This proposal presents a strategic plan to integrate **AI-powered automation** across our operations using **n8n** — an enterprise-grade, open-source workflow automation platform trusted by companies like Delivery Hero (saving 200+ hours/month) and StepStone (running 200+ mission-critical workflows).

By combining n8n with AI models (from free local models to enterprise-grade LLMs), we can build **personalized AI agents**, automate the entire **Microsoft 365 ecosystem**, enable **intelligent data analysis**, and streamline **project management** — delivering **measurable ROI within weeks**.

What this unlocks:

- Personalized AI agents tailored for every team and task — working 24/7
- Complete Microsoft 365 integration — Outlook, Teams, Excel, SharePoint, OneDrive, Calendar all AI-powered
- Intelligent data analysis that turns raw data into insights automatically
- Project management on autopilot — task assignment, progress tracking, status reports, and follow-ups
- Full data sovereignty — self-hosted, with sensitive data never leaving our infrastructure

2. Why n8n is the Right Platform

Platform Capabilities

Capability	What It Means for Us
500+ pre-built integrations	Connect every tool we use — MS 365, Jira, Asana, Trello, Slack, CRM, databases
AI-native (LangChain built-in)	Build AI agents, RAG pipelines, and multi-agent systems visually — no heavy coding
Self-hosted / On-premise	Full data sovereignty — our data stays on our servers
SOC 2 audited	Enterprise security with SSO, RBAC, audit logs, encrypted credentials
Open source (174k+ GitHub stars)	Full transparency, no vendor lock-in, massive community
MCP Protocol support	Future-proof architecture — standardized AI tool connectivity
Visual drag-and-drop builder	Business teams can build and modify workflows without developers
Handles 220 executions/second	Production-grade performance for high-volume operations

Cost Advantage

Platform	Monthly Cost (comparable usage)
n8n (self-hosted)	Free (Community Edition)
n8n Cloud (Pro)	~\$50/month
Zapier	\$500+/month
Microsoft Power Automate	\$15/user/month × headcount
Custom development	\$10,000+/month (developer salaries)

3. AI Model Strategy: From Local to Enterprise

n8n supports a **flexible, tiered AI model architecture** — we choose the right model for each task:



Tier 1 — Local Models (via Ollama) — \$0 Cost

- **Models:** Llama 3.2, DeepSeek R1, Mistral, Phi-4
- **How:** Install Ollama on our servers → connect to n8n via API
- **Data stays 100% on our servers** — zero external API calls
- **Best for:** Internal document summarization, email drafting, data classification, sensitive data processing
- **Hardware:** Runs on existing servers — 8GB RAM handles 7B parameter models comfortably

♣ Tier 2 — High-Speed Cloud APIs — Low Cost

- **Models:** Groq (ultra-fast), Google Gemini Flash, Claude Haiku, GPT-4o Mini
- **Cost:** \$0.10–\$1 per million tokens (extremely economical)
- **Best for:** Customer-facing content, higher-quality analysis, real-time chat agents



Tier 3 — Premium Enterprise Models — For Complex Tasks

- **Models:** OpenAI GPT-4o, Anthropic Claude Sonnet/Opus, Google Gemini Pro
- **Cost:** \$3–15 per million tokens
- **Best for:** Complex multi-step reasoning, strategic analysis, advanced data interpretation

Recommended Strategy

Deploy **Tier 1 locally for all internal tasks** (zero cost, maximum privacy). Use **Tier 2/3 selectively** for complex or customer-facing tasks. This hybrid approach optimizes both cost and data security.

4. Projected ROI & Benefits

Quantified Time Savings (Conservative Estimates)

Process	Current Time	After AI Automation	Reduction	Annual Hours Saved (10-person team)
Email triage, drafting & follow-ups	2 hrs/day/person	15 min/day	87%	4,375 hrs
Data analysis & insight generation	3 hrs/task	15 min/task	92%	2,800 hrs
Excel report/dashboard preparation	2 hrs/report	10 min/report	92%	1,500 hrs
Meeting scheduling, prep & follow-up	1 hr/day/person	5 min/day	92%	2,375 hrs
Document summarization & search	45 min/document	2 min/document	96%	2,100 hrs
Project status tracking & reporting	1.5 hrs/day	10 min/day	89%	1,825 hrs

Process	Current Time	After AI Automation	Reduction	Annual Hours Saved (10-person team)
Task assignment & follow-up	45 min/day	5 min/day	89%	1,040 hrs
Total Annual Hours Saved				~16,000 hrs

Financial Impact

Metric	Projected Value
Hours saved annually	16,000+ hours
Cost equivalent (@ ₹800/hr avg.)	₹1.28 Cr saved/year
Employee productivity gain	3–5 hrs/day freed for strategic work
Data analysis turnaround	Hours → minutes (real-time insights)
Project visibility	100% real-time — zero manual status updates
Platform cost (Year 1)	₹2–5 Lakhs (infrastructure + cloud APIs)
ROI	25x–60x return on investment

Strategic Benefits

- **Competitive edge:** AI-native operations position us ahead of industry peers
 - **Talent retention:** Employees focus on high-value strategic work, not manual grunt work
 - **Scalability:** Handle growing workload without proportionally increasing headcount
 - **Speed:** Decisions backed by real-time data analysis instead of week-old reports
 - **Consistency:** AI agents deliver the same quality 24/7 — no fatigue, no missed steps
-

5. **Microsoft 365 Deep Integration (HIGH PRIORITY)**

This is the backbone — turning our entire MS ecosystem into an AI-powered productivity engine.

Outlook — Intelligent Email Operations

- **AI Email Triage:** Every incoming email auto-classified by priority (urgent / important / FYI / spam), department, and action type (needs reply / needs action / informational)
- **Smart Auto-Drafts:** AI reads the full email thread, understands context, and drafts a response in your writing style — you review and send with one click

- **Bulk Email Processing:** Monday morning inbox overload? AI summarizes all unread emails into a priority digest with recommended actions
- **Follow-up Tracker:** AI detects emails that haven't received a response after X days → sends polite auto follow-ups or reminds you
- **Email-to-Task:** AI reads an email containing a task request → automatically creates a task in Jira/Asana/Trello with correct assignee, deadline, and context
- **Smart Signatures & Templates:** AI selects the right email template based on recipient and context, personalizes it with relevant data
- **Attachment Intelligence:** Incoming attachment → AI auto-summarizes the document and adds the summary to the email thread



Excel — Automated Data Analysis & Live Dashboards

- **Auto-Updating Dashboards:** n8n pulls data from databases, APIs, CRM → AI processes, analyzes & formats → automatically populates Excel sheets on schedule (hourly/daily/weekly)
- **Natural Language Analysis:** Ask in plain English — "What was our best performing segment last quarter?" → AI queries data, generates answer + chart in Excel
- **Anomaly Detection:** AI continuously scans data in Excel, highlights outliers and unusual patterns with color-coded conditional formatting and notes
- **Data Consolidation:** Multiple data sources (different CSVs, databases, API responses) → AI merges, cleans, de-duplicates → updates a master Excel workbook

- **Auto Pivot Tables & Charts:** AI generates pivot tables, charts, and visualizations based on the data patterns it detects or your natural language request
- **Trend Analysis:** AI identifies trends, correlations, and patterns across time-series data → writes narrative insights alongside the numbers
- **Workflow:** Database/API → n8n transformation → AI analysis → Excel update via Microsoft Graph API → Teams/Email notification with key findings



Calendar & Scheduling — AI Meeting Management

- **Smart Scheduling:** Tell AI "Schedule a 30-min call with Ravi, Priya, and the design team this week" → AI checks all calendars → finds optimal slot → sends invites with agenda
- **Meeting Prep Packets:** 30 minutes before each meeting, AI auto-compiles relevant documents, recent emails with attendees, last meeting notes, and pending action items → sends to your inbox
- **Post-Meeting Intelligence:** Meeting ends → AI transcribes recording (Whisper AI) → extracts action items → creates tasks in project management tool → sends summary to all participants
- **Recurring Meeting Reports:** For weekly standups/reviews — AI auto-generates progress report from project management data and shares it before the meeting
- **Conflict Resolution:** Double-booked? AI suggests which meeting to reschedule based on priority, offers alternative times, and handles the rescheduling communication

Teams — AI-Powered Collaboration Hub

- **Company Knowledge Bot:** Embed an AI chatbot in Teams that answers questions about company policies, processes, product info, project status — pulling answers from your document knowledge base (RAG)
- **Channel Summarizer:** Missed 200 messages in a Teams channel? AI gives you a 5-line summary with key decisions and action items
- **Smart Notifications:** AI filters Teams notifications — only alerts you about messages that actually need your attention, silences noise
- **Meeting Bot:** AI joins your Teams meetings, takes notes in real-time, and creates a structured summary with action items immediately after
- **Command Bot:** Type commands like "@AI create a task for Amit to review the proposal by Friday" → bot creates the task directly in your project management tool

SharePoint & OneDrive — Document Intelligence

- **Auto-Processing:** New document uploaded → AI reads, summarizes, extracts key data, classifies by type, and routes notification to relevant team
- **Smart Search:** Instead of browsing folders — ask "Find the vendor agreement we signed in November" → AI searches across all documents semantically
- **Version Comparison:** AI compares document versions, highlights all changes, and summarizes what's different in plain language
- **Template Generation:** AI auto-generates documents from templates using data pulled from CRM, databases, or project management tools

- **Knowledge Auto-Indexing:** Every new document automatically indexed into the RAG vector database — making it instantly searchable by the AI chatbot
-

6. Intelligent Data Analysis

What n8n + AI Enables

Turn raw data into decisions — automatically, continuously, and in plain language.

Data Analysis Workflows

Workflow	What Happens
Scheduled Data Pulls	n8n auto-fetches data from databases, APIs, SaaS tools on schedule → AI analyzes → outputs to Excel/Dashboard
Anomaly Alerts	AI monitors data streams → detects unusual patterns → sends instant alert with context to Teams/Email
Trend Reports	AI runs time-series analysis → identifies trends, seasonality, correlations → generates narrative report + charts
Cross-Source Analysis	AI pulls data from 5+ sources → merges → finds insights you'd miss looking at each source separately
Ad-hoc Analysis	Ask in natural language via Teams bot: "Compare this month's performance vs last quarter" → AI queries, analyzes, responds

Workflow	What Happens
Predictive Insights	AI identifies patterns and provides forward-looking insights — "Based on current trends, X metric is projected to Y"
Data Quality Monitoring	AI scans incoming data for inconsistencies, missing fields, format errors → auto-corrects or flags for review

Supported Data Sources (via n8n's 500+ integrations)

Databases (MySQL, PostgreSQL, MongoDB) • APIs • Google Sheets • Excel files • CSV/JSON files • CRM systems • SaaS platforms • Web scraping • Email data • SharePoint lists • Custom data sources via HTTP requests

7. Project Management Automation

The Problem

Project managers spend hours on status updates, task tracking, follow-ups, and reporting — time that should go into actual project delivery.

AI-Powered Project Management Workflows

Automation	How It Works
Auto Task Creation	Email/Teams message/meeting notes → AI extracts tasks → creates them in Jira/Asana/Trello with assignee, deadline, priority, description
Smart Task Assignment	AI suggests best assignee based on workload, skills, and past performance

Automation	How It Works
Progress Tracking	AI monitors task statuses across all projects → generates real-time progress dashboard in Excel/Teams
Automated Stand-up Reports	AI pulls data from project tool every morning → generates "What was done / What's planned / Blockers" report → posts to Teams channel
Deadline Alerts	Task approaching deadline with no progress? → AI sends escalating reminders: gentle nudge → urgent alert → manager notification
Sprint/Milestone Reports	End of sprint → AI auto-generates completion report with metrics, burndown analysis, and highlights
Cross-Project Visibility	AI aggregates status from multiple projects → creates executive summary dashboard showing health of all projects
Meeting → Action Items Pipeline	Meeting recorded → AI transcribes → extracts action items → creates tasks → assigns → tracks → reports
Risk Detection	AI identifies projects at risk (too many overdue tasks, blocked items, resource conflicts) → alerts PM and leadership
Resource Allocation Insights	AI analyzes who's overloaded vs. underutilized → suggests rebalancing

Supported Project Management Tools

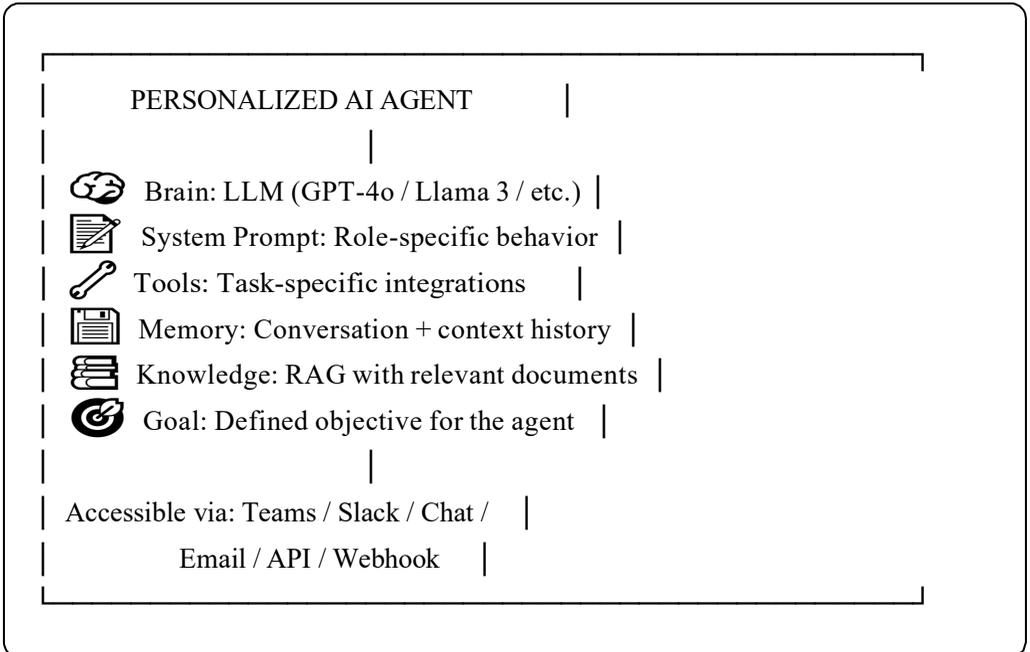
Jira • Asana • Trello • Monday.com • ClickUp • Notion • Linear • GitHub Issues

8. Personalized AI Agents (HIGH PRIORITY)

What Are Personalized AI Agents?

Each agent is a **purpose-built digital assistant** designed for a **specific role, team, or task**. It has its own personality, instructions, tools, memory, and knowledge base. It's not a generic chatbot — it's a specialist.

How Personalized Agents Work in n8n



Agent Catalog — Ready to Build

 **Personal Productivity Agent (per employee)**

- **Role:** Your personal AI assistant that knows your schedule, tasks, preferences, and communication style
- **Can do:** Draft emails in your tone • manage your calendar • summarize your unread messages • remind you of pending tasks • prep you before meetings • track your action items
- **Knowledge:** Your email patterns, calendar, task list, writing style
- **Access:** Teams DM, Slack DM, or dedicated chat interface



Project Manager Agent

- **Role:** Automated PM that keeps projects on track
- **Can do:** Create tasks from conversations • send deadline reminders • generate sprint reports • update status dashboards • identify blocked tasks • post daily stand-up summaries
- **Tools:** Jira/Asana/Trello API, Excel, Teams, Calendar, Email
- **Knowledge:** Project plans, timelines, team assignments, historical velocity



Data Analyst Agent

- **Role:** On-demand data analyst available 24/7
- **Can do:** Query databases • generate charts • identify trends • create Excel reports • answer data questions in plain English • detect anomalies
- **Tools:** Database connectors, Excel API, charting tools, Python/JS code execution
- **Knowledge:** Database schemas, business metrics definitions, historical benchmarks



Communication Agent

- **Role:** Handles all routine communication tasks
- **Can do:** Draft emails • compose client updates • write internal announcements • create newsletters • manage follow-ups • summarize long email threads
- **Tools:** Outlook, Teams, CRM, contact database
- **Knowledge:** Company tone guides, templates, recipient preferences, communication history



Research Agent

- **Role:** Gathers and synthesizes information from multiple sources
- **Can do:** Search the web • analyze competitor activity • summarize articles/reports • compile market intelligence • generate briefing documents
- **Tools:** Web search, news APIs, document readers, summarization engine
- **Knowledge:** Industry context, company focus areas, past research



Onboarding & Training Agent

- **Role:** Guides new employees through onboarding and answers questions
- **Can do:** Walk through onboarding steps • answer policy questions • explain processes • recommend training materials • track onboarding progress
- **Tools:** HRMS, Knowledge base (RAG), Calendar, Email, Teams

- **Knowledge:** All company policies, procedures, org structure, tools documentation

✂ IT Support Agent

- **Role:** First-line IT helpdesk that resolves common issues instantly
- **Can do:** Troubleshoot common problems • guide users through solutions • create support tickets • escalate complex issues • track resolution status
- **Tools:** Ticketing system, Knowledge base, Teams, Email
- **Knowledge:** IT documentation, common issue database, troubleshooting guides



Document Agent

- **Role:** Handles all document-related intelligence
- **Can do:** Summarize documents • extract key information • compare versions • search across document library • generate documents from templates • answer questions about any company document
- **Tools:** SharePoint, OneDrive, RAG vector database, PDF/DOCX processors
- **Knowledge:** Entire company document library via RAG

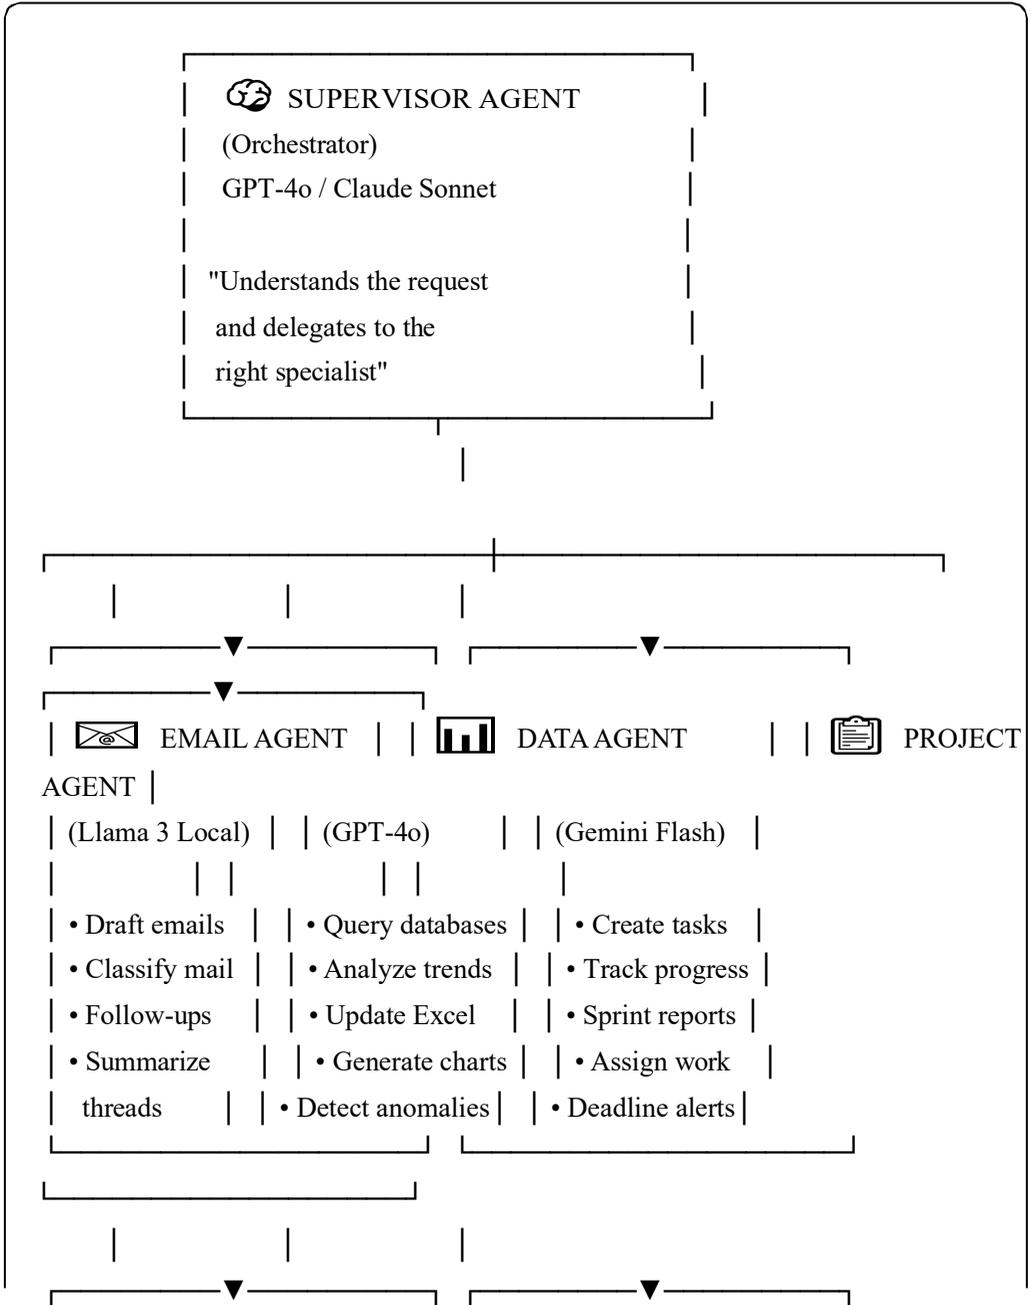
9. Multi-Level Agent Architecture

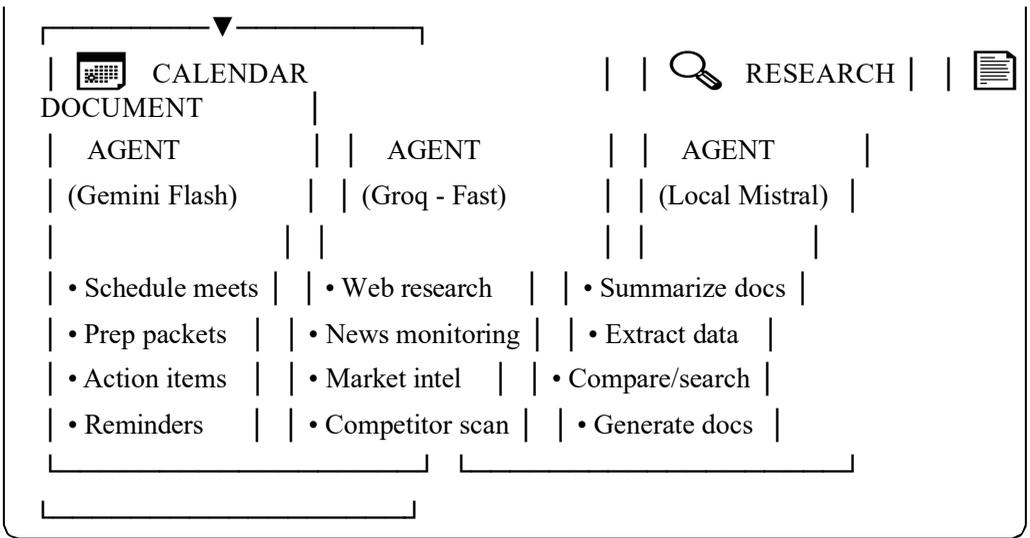
Why Multi-Agent?

A single agent trying to do everything becomes unreliable. The multi-agent

approach uses **specialized agents that collaborate** — each expert in their domain, coordinated by a supervisor.

Architecture: Hierarchical Multi-Agent System





How It Works — Real Example

User says in Teams: *"Prepare for my meeting with Acme Corp tomorrow"*

1. **Supervisor Agent** receives the message → understands it needs multiple things
2. **Delegates to Calendar Agent:** Find tomorrow's meeting details, attendees, time
3. **Delegates to Email Agent:** Pull all recent email threads with Acme Corp contacts
4. **Delegates to Document Agent:** Find any contracts, proposals, or presentations related to Acme Corp
5. **Delegates to Research Agent:** Get latest news about Acme Corp
6. **Delegates to Data Agent:** Pull Acme Corp's account data from CRM
7. **Supervisor synthesizes everything** → delivers a comprehensive meeting prep packet to your inbox/Teams

All of this happens in **under 2 minutes**, automatically.

Multi-Agent Architecture Patterns

Pattern	How It Works	Best For
Hierarchical	Supervisor delegates to worker agents	Complex multi-step requests (meeting prep, onboarding)
Sequential Pipeline	Each agent processes then hands off to next	Step-by-step processes (data → analyze → report → distribute)
Parallel Execution	Multiple agents work simultaneously	Speed-critical tasks (simultaneous research + data pull + email scan)
Dynamic Handoff	Context-based routing to the right agent	Chat interfaces (route to PM agent vs. data agent vs. email agent based on question)

Cost Optimization Through Multi-Agent

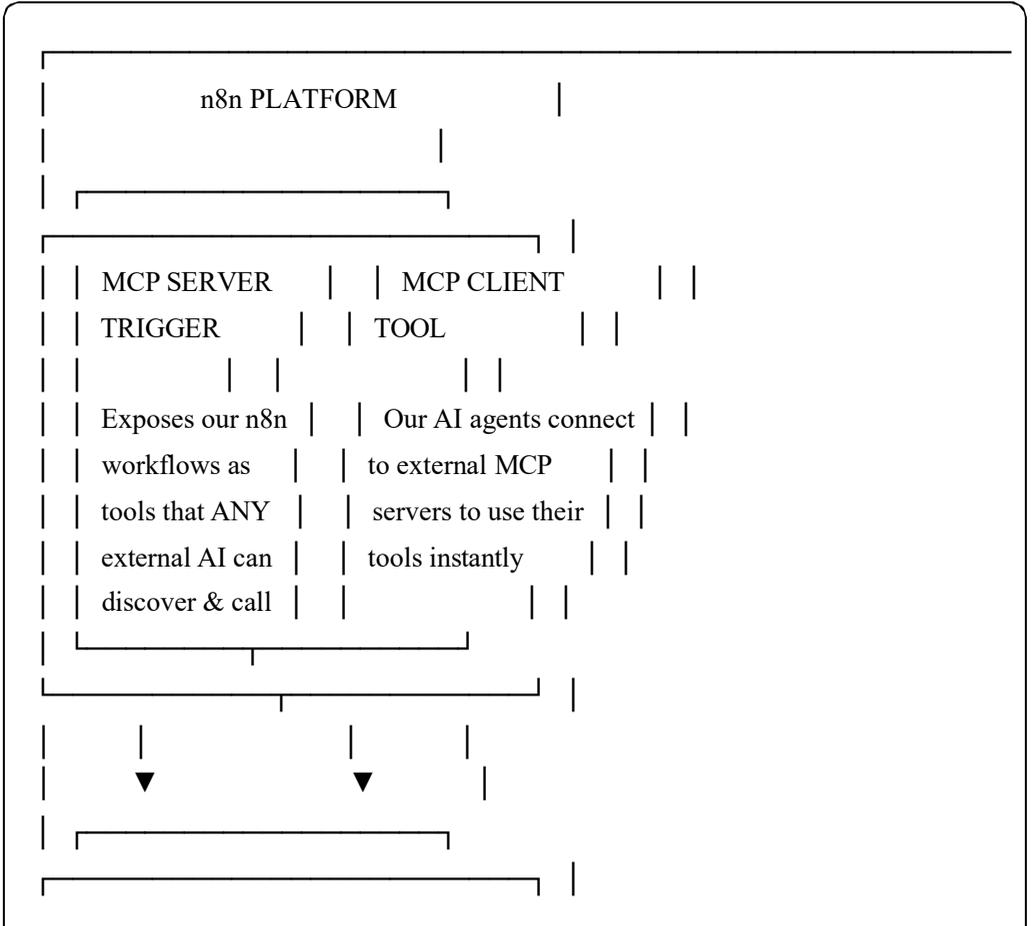
- **Simple tasks** (email classification, reminders) → **free local models** (Llama 3)
- **Medium tasks** (report writing, scheduling) → **low-cost cloud models** (Gemini Flash, Groq)
- **Complex tasks** (data analysis, strategic synthesis) → **premium models** (GPT-4o, Claude)
- Result: **60–80% cost reduction** compared to using premium models for everything

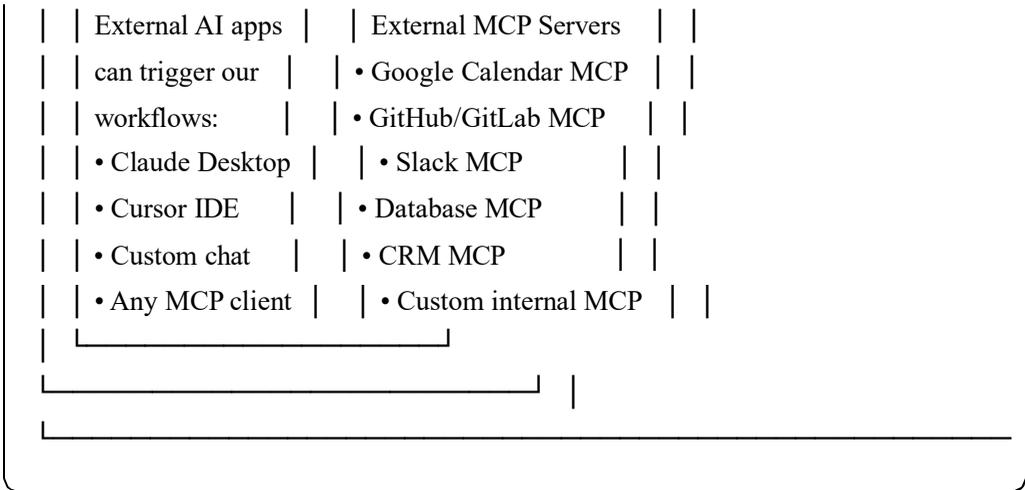
10. 🖱️ MCP Server Architecture (Model Context Protocol)

What is MCP?

MCP is an **open standard** (by Anthropic) that acts as **"USB-C for AI"** — a universal protocol standardizing how AI models connect to tools, data sources, and services. Instead of building custom integrations for every AI-tool combination, MCP provides **one plug that works everywhere**.

n8n + MCP Architecture





Why MCP is a Game-Changer

Without MCP

With MCP

Build custom integration for every AI + tool combo

One standard protocol — plug and play

Adding a new tool = rebuild workflows

New tool added once = all agents can instantly use it

Each AI agent needs separate tool configs

Tools are auto-discoverable by all connected agents

Fragile, hard to maintain

Future-proof, scalable, maintainable

Our MCP Use Cases

- n8n as MCP Server:** Expose our internal workflows (project status check, data query, report generator, task creator) as tools — anyone using

Claude Desktop, Cursor, or any AI assistant can trigger them via natural language

- **n8n as MCP Client:** Our AI agents connect to external MCP servers for calendar management, code repositories, communication tools, and third-party data
 - **Bidirectional ecosystem:** Internal and external AI tools talk to each other seamlessly through a standardized protocol
-

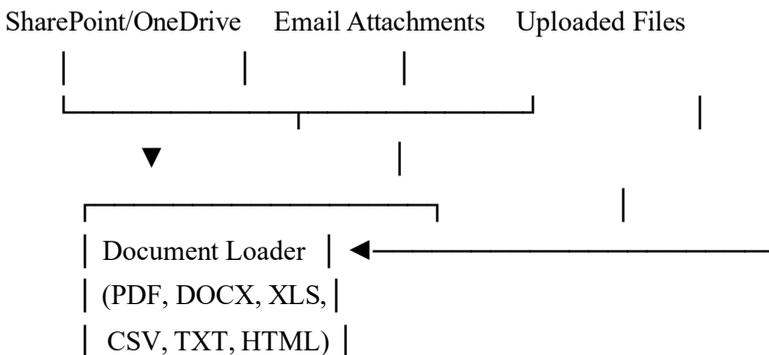
11. 📖 RAG Pipeline for Document Intelligence

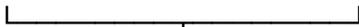
What is RAG?

Retrieval-Augmented Generation (RAG) enables AI to answer questions using **our own company documents** instead of general training data. The AI retrieves relevant information from our knowledge base and generates accurate, source-grounded responses — no hallucination.

RAG Architecture in n8n

PHASE 1: DOCUMENT INGESTION (runs automatically)





Text Splitter | Splits into optimized chunks
(500-1000 | for accurate retrieval
tokens/chunk) |



Embedding Model | Converts text → vector
(OpenAI Ada or | (numerical representation
local model) | of meaning)



Vector Database | Stores & indexes all
(Qdrant / | document embeddings for
Supabase / | lightning-fast semantic
PGVector) | search

PHASE 2: INTELLIGENT RETRIEVAL (when someone asks a question)

User Question (via Teams Bot / Chat / Email)



Query Embedding | Same model converts question
into a vector



| Similarity Search | Finds the most relevant
| (Vector DB) | document chunks



| Reranker | Re-scores results for
| (Cohere/local) | maximum accuracy



| LLM Response | AI generates answer using
| Generation | retrieved documents as context



Accurate Answer + Source Document References

What We Feed Into RAG

- **Company policies** — HR handbook, IT policies, operational procedures, leave policy
- **Project documentation** — Requirements, specs, architecture docs, meeting notes
- **Product knowledge** — Feature documentation, pricing, API docs, user guides
- **Process manuals** — SOPs, workflow guides, best practices
- **Training materials** — Onboarding guides, tutorials, presentations
- **Historical data** — Past reports, decision records, archived communications

RAG Use Cases

Use Case	Example
Policy Q&A	"What's our remote work policy?" → Instant answer from HR docs
Project Context	"What were the key decisions in the Alpha project kickoff?" → AI finds meeting notes
Process Lookup	"How do I request access to the production server?" → Step-by-step from IT docs
Document Search	"Find the vendor agreement we discussed in October" → Semantic search across all docs
Cross-Document Insights	"What are all the deadlines mentioned across active project docs?" → AI scans everything
Onboarding Support	New hire asks "How do I set up my dev environment?" → Instant guided answer

Advanced: Agentic RAG

Beyond basic RAG, n8n supports **Agentic RAG** — the AI agent **intelligently decides** how to answer:

- **Policy question?** → Search vector database
- **Data question?** → Run SQL query on database directly
- **Need full document?** → Retrieve complete file instead of chunks

- **Cross-source question?** → Search vector DB + database + email + project tool and synthesize
-

12. Complete Technology Stack

INTERACTION LAYER

MS Teams Bot | Outlook | Slack | Web Chat | API

ORCHESTRATION LAYER

n8n Platform

(Workflows + AI Agents + Multi-Agent + MCP)

AI & INTELLIGENCE LAYER

Local: Ollama (Llama 3, Mistral, DeepSeek)

Cloud: OpenAI, Claude, Gemini, Groq

RAG: Vector DB (Qdrant / Supabase / PGVector)

Framework: LangChain (built into n8n natively)

INTEGRATION LAYER

Microsoft 365 (Outlook, Teams, Excel, SharePoint)

Project Tools (Jira, Asana, Trello, ClickUp)

Databases | CRM | APIs | 500+ connectors

INFRASTRUCTURE LAYER

Docker/Kubernetes | On-Premise Servers | SSL/TLS

Encrypted Credentials | Audit Logging | RBAC

13. Implementation Roadmap

Phase 1: Quick Wins (Weeks 1–3) — Zero Cost

Objective: Demonstrate immediate, visible value

Deliverable	Impact
n8n self-hosted setup (Docker)	Platform live and ready
Ollama + Llama 3 local AI installed	Free AI capabilities running
Email auto-classification & smart draft replies	2+ hrs/day saved per person
Document summarization bot (RAG prototype)	45 min → 2 min per document
Auto Excel dashboard from database	Eliminate manual report prep
Basic project task creation from email/chat	Zero manual task entry

Phase 2: Full MS 365 + Agent Deployment (Weeks 4–

8) Objective: AI-powered Microsoft ecosystem + specialized agents

Deliverable	Impact
Complete Outlook automation (triage, drafts, follow-ups)	Email on autopilot
Teams AI chatbot (company knowledge base via RAG)	Instant answers for everyone
Auto-updating Excel dashboards via Microsoft Graph API	Real-time data, zero manual work
Calendar AI (smart scheduling, meeting prep, action items)	Meetings fully managed
SharePoint document intelligence pipeline	Auto-process every upload
Project Manager Agent (Jira/Asana integration)	PM on autopilot
Data Analyst Agent (database + Excel)	On-demand data insights
Personal Productivity Agent (per-user assistant)	Every employee gets an AI assistant
Vector database setup (Qdrant) for full RAG pipeline	Company-wide document intelligence

Phase 3: Multi-Agent & Enterprise Scale (Months 3–6)

Objective: Advanced architecture, full organization deployment

Deliverable	Impact
Multi-level agent architecture (Supervisor + 6 sub-agents)	Autonomous complex task handling
MCP server integration (bidirectional)	Future-proof AI tool ecosystem
Advanced Agentic RAG with reranking	Best-in-class document intelligence
Full project management automation suite	Complete PM autopilot
Cross-department agent deployment	Every team has specialized AI
Enterprise n8n license (SSO, RBAC, audit)	Production-grade governance
Internal AI champions trained (3–5 people)	Self-sustaining, growing capability

14. What We Need to Get Started

Requirement	Details
Management approval	Green light for Phase 1 — zero cost proof of concept
Server access	One VM/server for n8n + Ollama (existing infrastructure)
MS 365 API permissions	IT to enable Microsoft Graph API access for Outlook, Excel, Teams, SharePoint

Requirement	Details
Project tool API access	API keys for Jira/Asana/Trello (whichever we use)
Pilot team	5–10 people from one department to test and provide feedback
2-week demo slot	We'll present a working prototype with live, running automations

The bottom line: Every hour our team spends manually sorting emails, updating Excel sheets, chasing status updates, searching for documents, and scheduling meetings is an hour stolen from innovation, strategy, and growth. n8n + AI gives every employee a personalized AI assistant, turns our Microsoft ecosystem into an intelligent productivity engine, and puts data insights at everyone's fingertips — starting at zero cost, scaling to enterprise, with our data fully under our control.

Prepared by: Ritesh Singh Bisht

Date: February 2026