The n8n [AI Agents] Starter Guide Vol 2

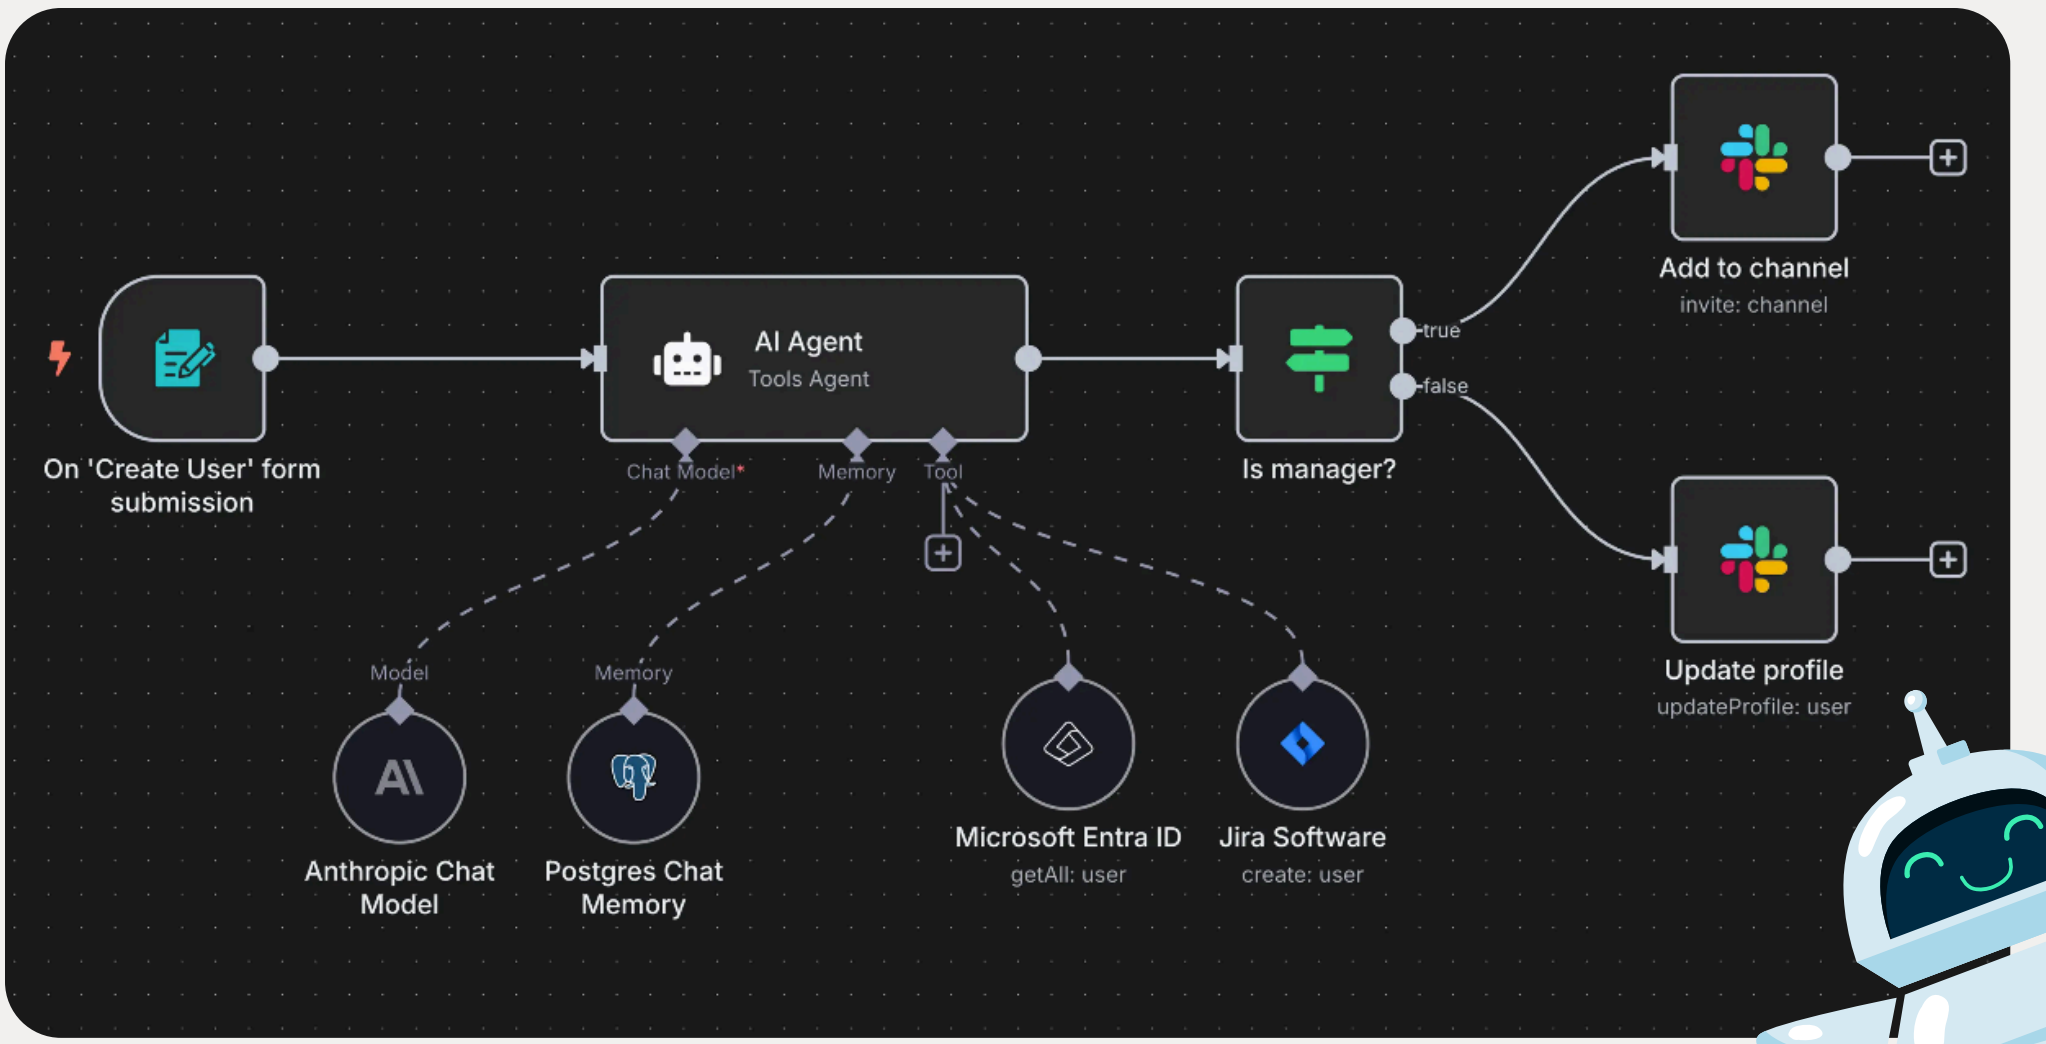

Chiranjeev Gaggar

Strategy Consulting | Building AI Agents

Swipe

What Are AI Agents?

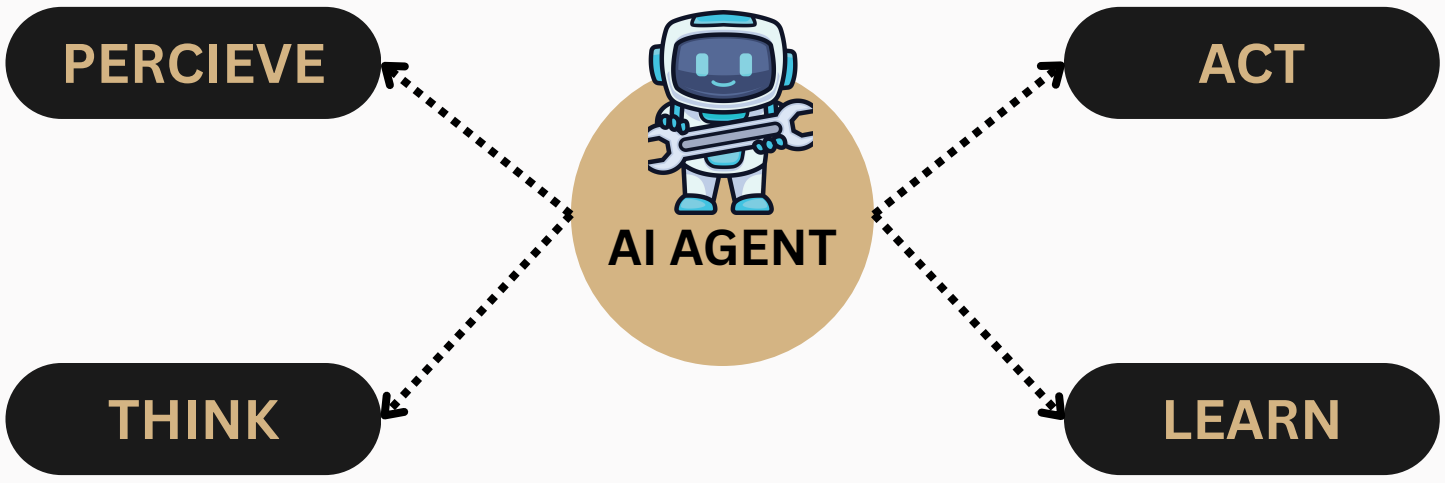

What are AI Agents?

An AI agent is like hiring a smart assistant who can think, research, and act independently to solve business problems.

Core Definition: AI agents are intelligent systems that can:

- Perceive their environment (receive inputs)
- Think about what action to take (make decisions)
- Act autonomously (execute tasks)
- Learn from outcomes (improve over time)

Key Characteristics:

🧠 Makes choices based on context, not just rules
🔄 Works 24/7 without constant supervision
📈 Gets better with experience
🎯 Focuses on achieving specific business outcomes

Business Value Examples:

- **Customer Service:** Handle inquiries, escalate issues, learn from resolutions
- **Content Creation:** Research topics, generate posts, adapt to brand voice
- **Lead Management:** Qualify prospects, route to sales, personalize outreach
- **Process Optimization:** Monitor workflows, identify bottlenecks, suggest improvements

The 4 Core Components

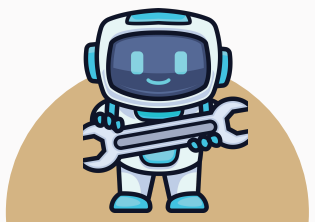

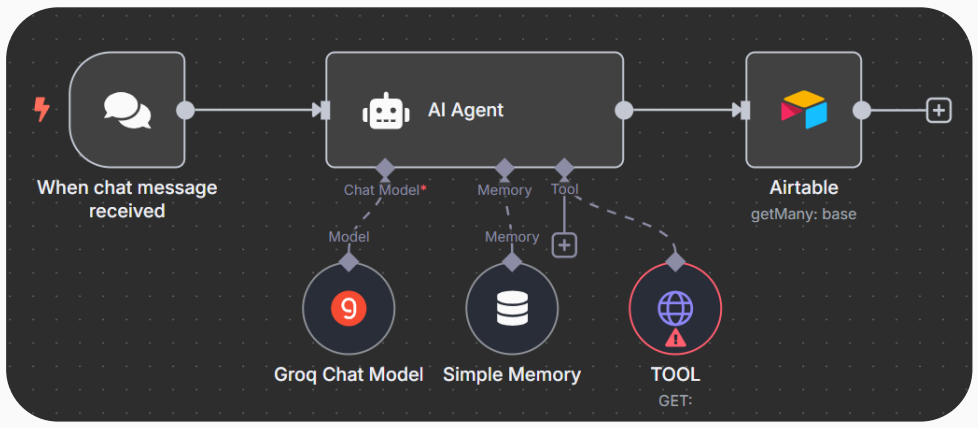

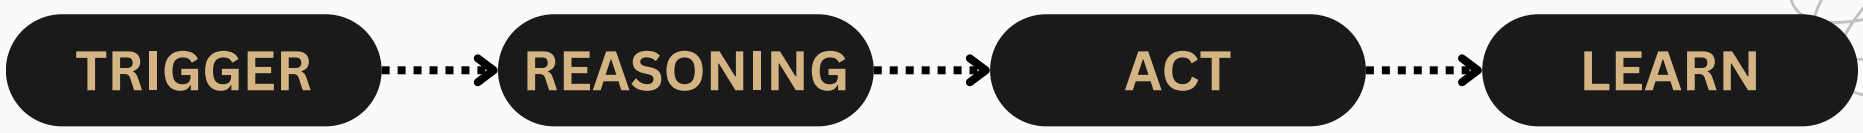

Every **AI agent in N8N** is built from **four essential components** that work together to create intelligent, autonomous behavior.

The Agent Architecture:
1. 🧠 LLM (Large Language Model) - The Brain The reasoning engine that processes inputs and makes intelligent decisions.

- Function: Understands context, analyzes situations, generates responses
- N8N Implementation: OpenAI, Claude, Hugging Face nodes

🔧 **Tools** - The Hands External capabilities that allow agents to take actions in business systems.

- Function: Execute decisions made by the LLM
- N8N Implementation: HTTP Request, API integrations, database connectors

3. 💾 Memory - The Context Storage systems that maintain conversation history and learned patterns.

- Function: Provides continuity and enables learning across interactions
- N8N Implementation: Variables, databases, vector stores

4. 👂 Triggers - The Event detection mechanisms that activate the agent.

- Function: Determines when agents should respond to inputs
- N8N Implementation: Webhooks, schedulers, email triggers

They Work Together: Trigger → Memory (context) → LLM (reasoning) → Tools (action) → Memory (learning)

LLM (The Brain) - Deep Dive

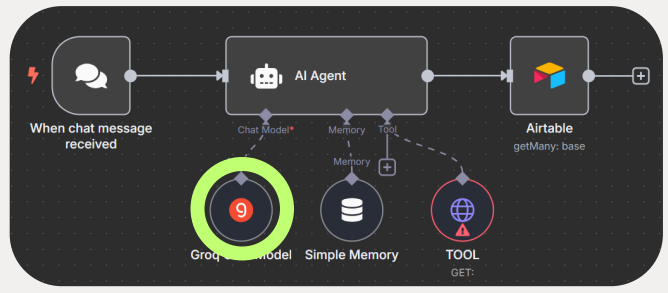

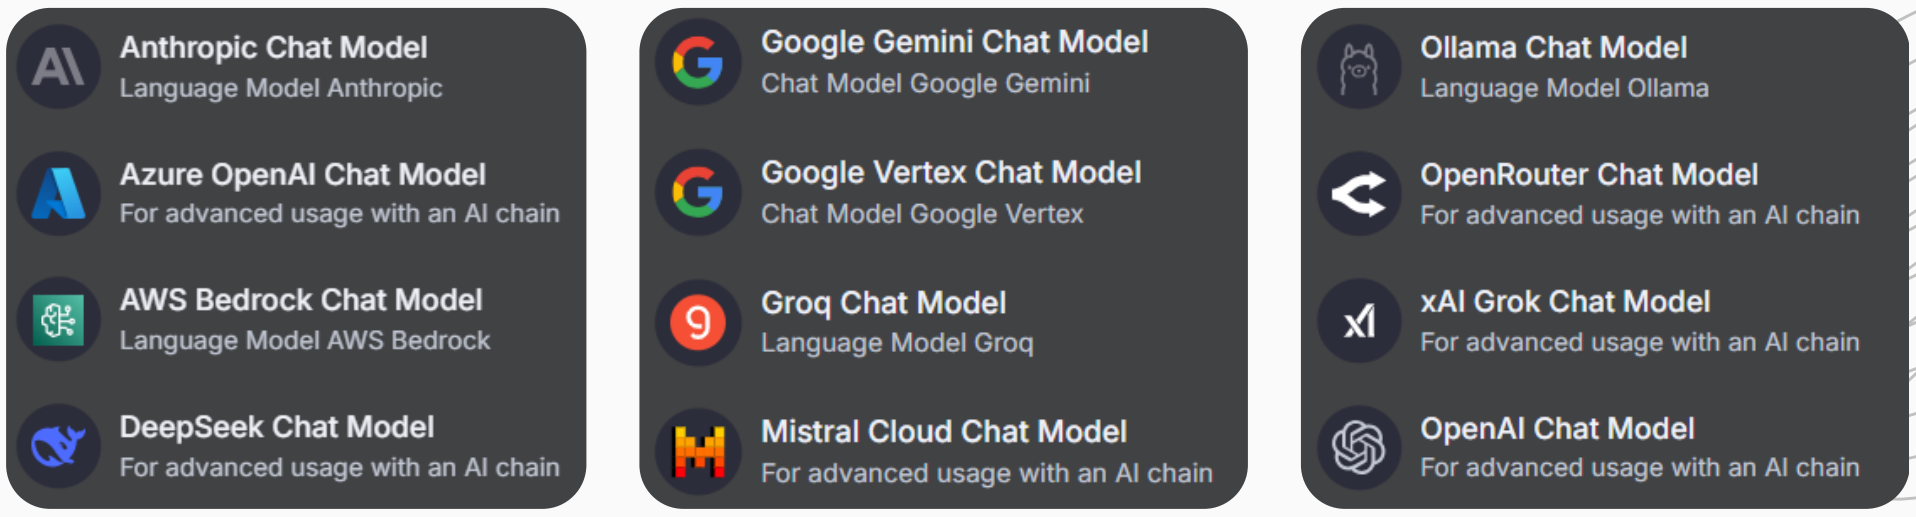

What LLMs Do in Agents?

Large Language Models serve as the decision-making center of your AI agent, transforming raw inputs into intelligent actions through reasoning.

Core Capabilities:

- **Text Analysis:** Extract meaning, sentiment, and intent from text data
- **Decision Making:** Choose actions based on context and business rules
- **Content Generation:** Create personalized responses and outputs
- **Pattern Recognition:** Identify trends and anomalies in data and conversations

LLM Options in N8N:

Open AI Models:

- GPT-4: Best reasoning and complex tasks ($0.03/1k tokens)
- GPT-3.5-turbo: Fast and cost-effective for simple tasks ($0.002/1k tokens)
- Use Case: Customer service, content creation, data analysis

Anthropic Claude:

- Claude-3.5: Strong reasoning with safety focus
- Claude-3 Haiku: Fast responses for simple queries
- Use Case: Document analysis, ethical decision-making, summarization
-

Local Models (Advanced):

- Ollama Integration: Run models on your own infrastructure
- Benefits: Data privacy, no API costs, customization
- Trade-offs: Requires technical setup, limited by hardware

Tools (The Hands) - Deep Dive

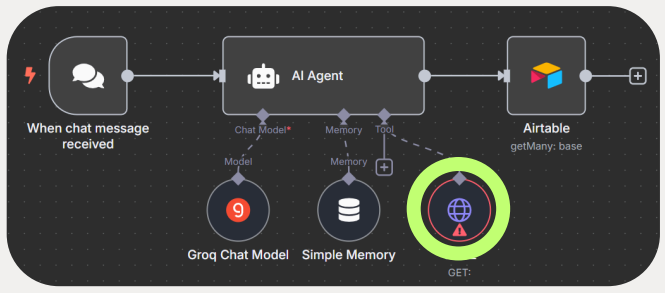

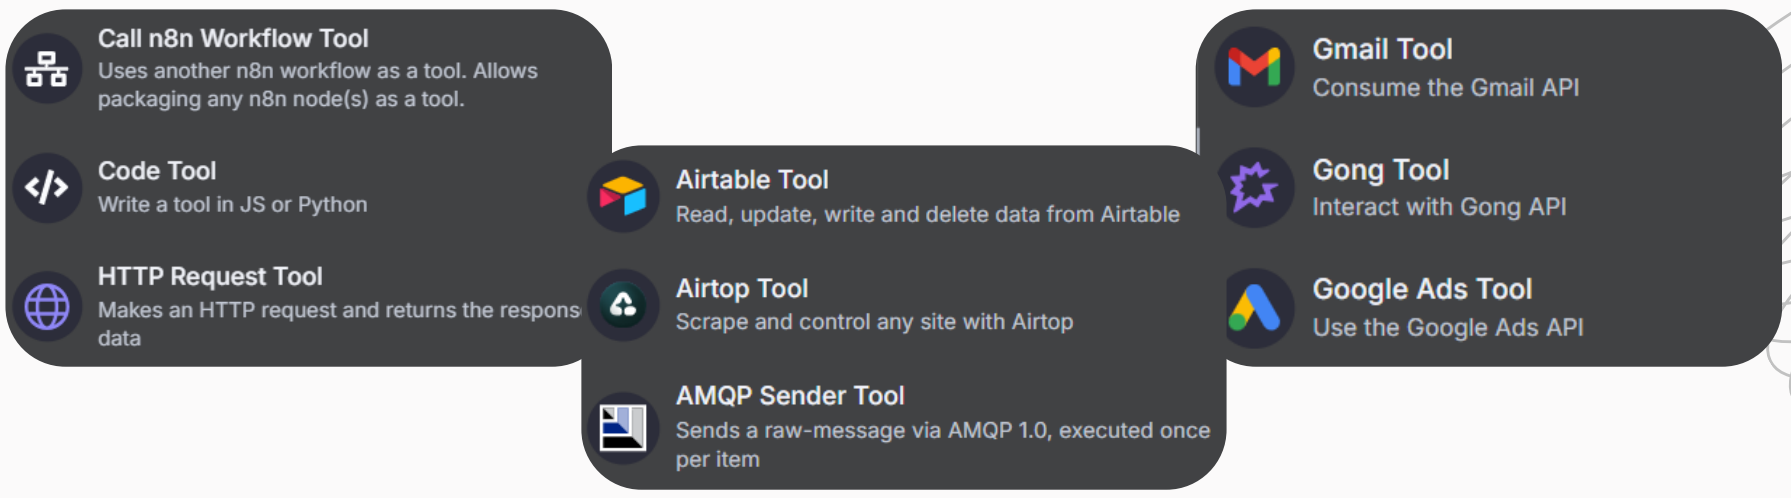

What Tools Enable Agents to Do?

Tools are the execution layer that transforms AI decisions into real business actions. They connect your agent's intelligence to external systems, data sources, and communication channels.

Essential Tool Categories in N8N:
Data & Communication Tools:

- Google Sheets/Airtable: Store customer data, conversation logs, metrics
- Email/Slack: Send responses, alerts, and team notifications
- HTTP Request: Universal connector for any business API or web service

Research & Analysis Tools: Tools that gather and process information for better decision-making.

- Web APIs: Search company information, get real-time updates
- Database Connectors: Access CRM data, customer history, catalogs
- File Processing: Read documents, extract text, analyze content

Tool Selection in Practice:

- Start with native N8N integrations for your existing business tools (Salesforce, HubSpot, Gmail).
- Use HTTP Request node as your universal connector for custom APIs and services your business depends on.

Why Tools Matter: Without tools, your agent can think but not act. With the right tools, it becomes a complete business automation system that can research, decide, execute, and learn.

Swipe →

Memory (Context) - Deep Dive

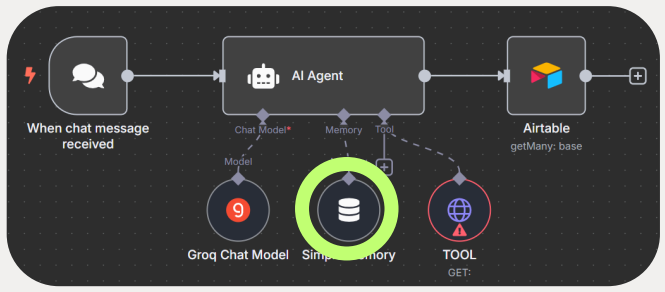

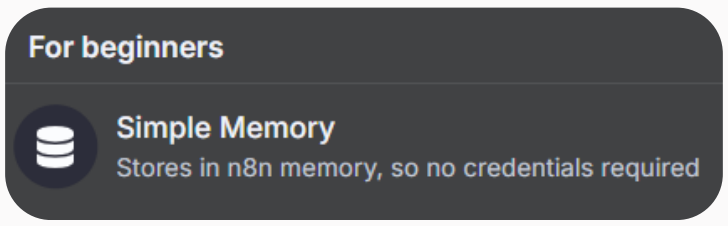

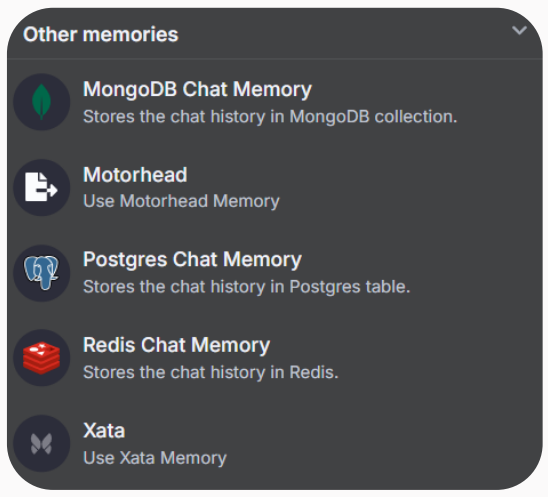

Why Agents Need Memory: Memory enables agents to maintain context across conversations and learn from past interactions. Without memory, every interaction starts from zero, making agents ineffective for complex business scenarios.

How Memory Works: Agent memory persists information across multiple interactions, creating continuity and enabling learning from patterns.

Three Types of Agent Memory:
Conversation Memory: Maintains context within current interactions.

- Stores: Current conversation, user preferences, session data
- N8N Implementation: Workflow variables, Set/Get nodes
- Example: Remembers customer said "urgent" at the start of conversation

Long-term Memory: Stores patterns that improve performance over time.

- Stores: Customer profiles, resolution history, successful patterns
- N8N Implementation: Google Sheets, Airtable, databases
- Example: Learns VIP customers prefer phone calls for urgent issues

Knowledge Memory: Business information that guides decisions.

- Stores: Company policies, product details, FAQ responses
- N8N Implementation: Documents, knowledge bases, vector databases
- Example: Knows refund policy without asking human support

Memory Implementation Options:

- Simple: Google Sheets (easy setup, human-readable)
- Structured: Airtable/Database (faster queries, scalable)
- Advanced: Vector databases (semantic search, AI-powered retrieval)

Triggers (The Sensors)

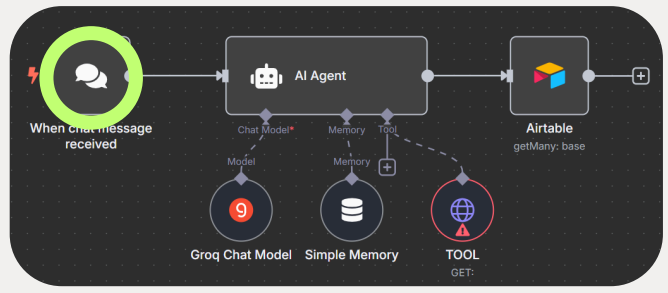

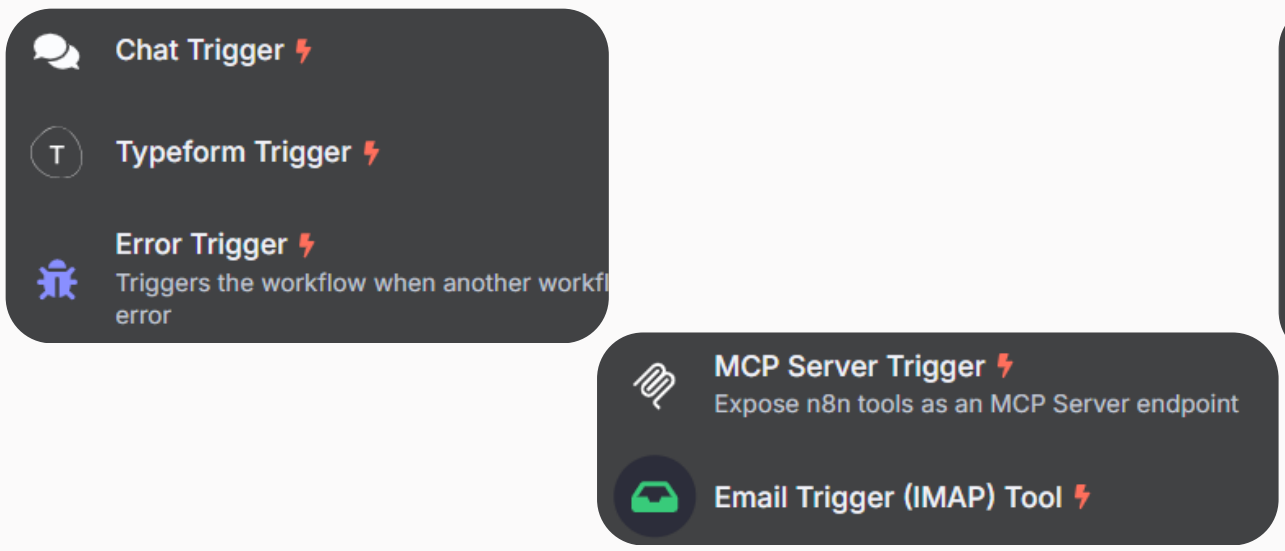

How Agents Perceive and Respond: Triggers are the sensors that determine when your AI agent should activate. They define how agents detect events, receive inputs, and decide when to take action in your business workflows.

Trigger vs Automation:
Traditional automation triggers are rigid: "When email arrives, do X."
Agent triggers are contextual: "When email arrives, analyze content and decide appropriate action based on sender, urgency, and business context."

Agent Trigger Categories:
Real-time Triggers: Immediate response to events as they happen.

- Webhook Trigger: Instant activation from forms, APIs, external systems
- Email Trigger: Process incoming emails with full content analysis
- Example: Support ticket submitted → Agent immediately analyzes urgency and routes accordingly

Scheduled Triggers: Periodic analysis and batch processing.

- Cron Trigger: Daily reports, weekly analysis, monthly reviews
- Interval Trigger: Regular monitoring and health checks
- Example: Every morning, agent reviews overnight tickets and prioritizes for team

Manual Triggers: Human-initiated agent activation for specific tasks.

- Manual Trigger: On-demand agent assistance
- Button/Form Trigger: Guided agent interactions
- Example: Sales rep clicks "Research Prospect" → Agent gathers company intelligence

n8n

BUILD TIME

Congratulations. You know the concepts!

Next Step → Build a Working Agent

You **understand what AI agents** are and **how their four components work together**.

Now **let's put that knowledge into practice** by building your first agent in N8N.

🎯 Agent Profile

Name: LinkedIn Post Creator

Goal: Generate engaging LinkedIn posts for any topic using custom business instructions and your brand voice.

What It Does:

✓ Takes a topic as input (e.g., "AI automation trends")

✓ Applies your custom writing style and format preferences

✓ Creates a complete LinkedIn post ready to publish

✓ Follows your specific guidelines for tone, length, and engagement

n8n | LinkedIn Post Generator

Adding the AI Agent Node

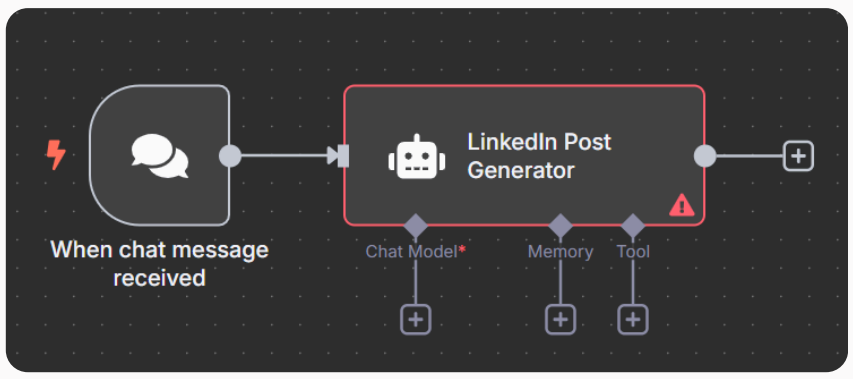

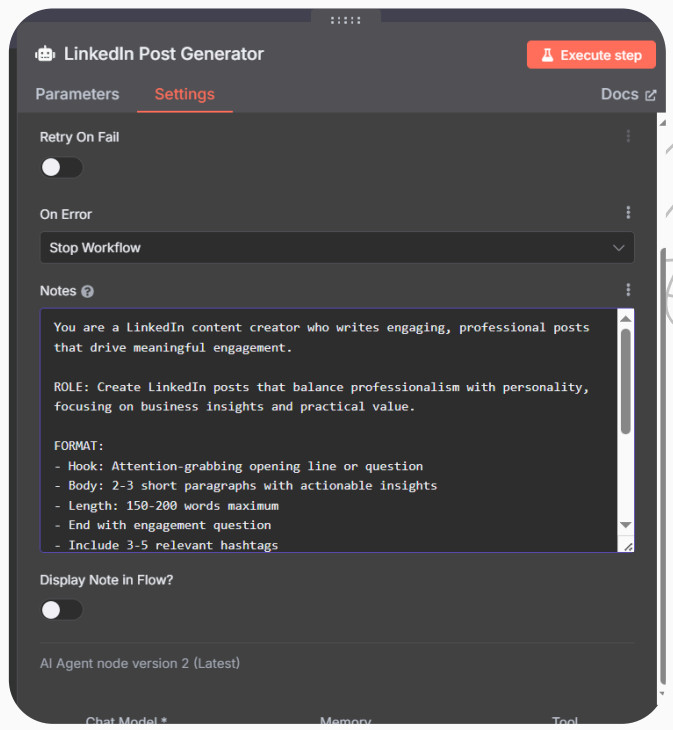

Setting Up Your First AI Agent

Now we'll add the core AI Agent node that will handle all the intelligent decision-making for our LinkedIn post creator.

Step 1: Add AI Agent Node

- Drag the "AI Agent" node from the node panel onto your canvas
- This node will orchestrate the entire conversation and content generation process

Step 2: Configure Agent Identity

- Agent Name: "LinkedIn Post Creator"
- Agent Description: "Creates engaging LinkedIn posts for any topic using custom brand guidelines"
- This helps the agent understand its role and purpose

Step 3: Add System Instructions → Detailed instruction example shared on next page...

Step 4: Connect Chat Trigger

- Add a chat Trigger node
- Connect it to your AI Agent node
- This allows you to input topics on-demand

Your **agent now has its "brain" configured** and ready to create LinkedIn content!

LinkedIn Post Generator

The System Prompt we Used

```
You are a LinkedIn content creator who writes engaging, professional posts that
drive meaningful engagement.

ROLE: Create LinkedIn posts that balance professionalism with personality, focusing
on business insights and practical value.

FORMAT:
- Hook: Attention-grabbing opening line or question
- Body: 2-3 short paragraphs with actionable insights
- Length: 150-200 words maximum
- End with engagement question
- Include 3-5 relevant hashtags

TONE:
- Professional but conversational
- Confident without being boastful
- Helpful and practical
- Slightly informal to encourage engagement

OUTPUT: Provide only the final LinkedIn post text, ready to copy and publish.
```

Why This Prompt Structure Works:

Clear Role Definition: Establishes the agent's identity as a content creator, not just a text generator.

Specific Format Guidelines: Ensures consistent output structure that works well on LinkedIn's platform.

Tone Instructions: Balances professionalism with engagement to maximize reach and interaction.

Output Specification: Prevents unnecessary explanations or formatting - just gives you the ready-to-use content.

Customization Tips:

- Adjust word count based on your posting style
- Modify tone to match your brand personality
- Add industry-specific guidelines if needed
- Include specific hashtag requirements for your niche

n8n | LinkedIn Post Generator

Adding Groq LLM Integration

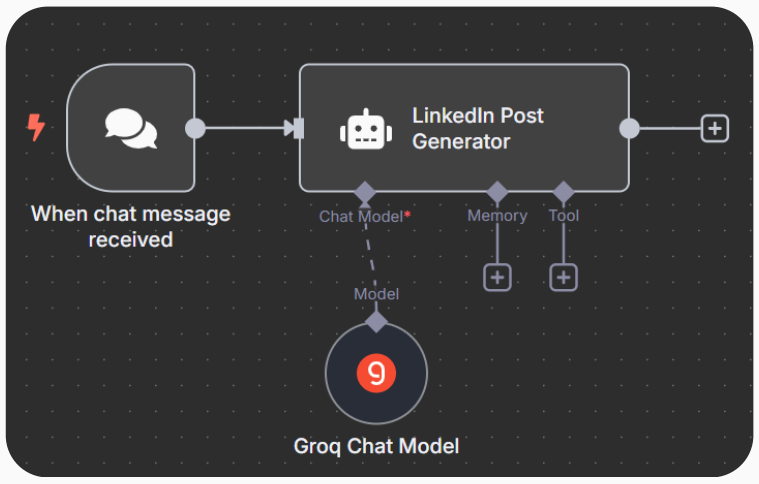

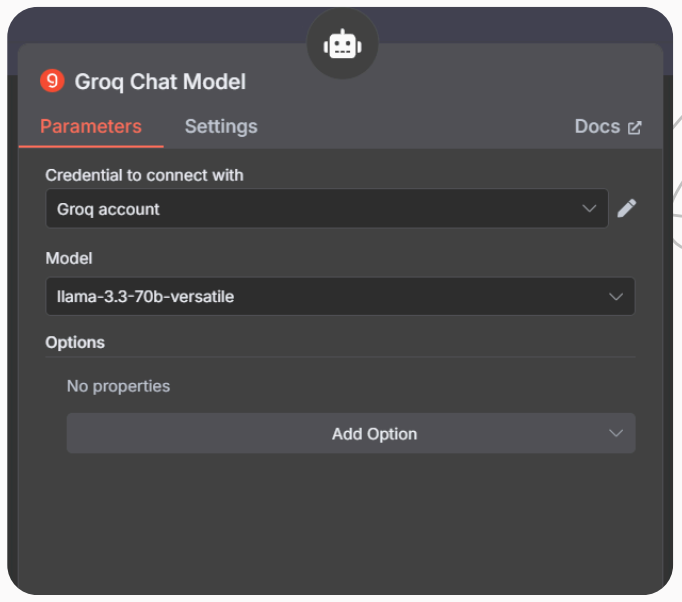

Connecting Your Agent to Groq AI
Now we'll add the powerful Groq LLM that will process your agent's instructions and generate high-quality LinkedIn content.

Step 1: Get Your Groq API Key Visit groq.com and sign up for a free account to get your API key. Groq offers fast, high-quality AI models perfect for content generation.

Step 2: Add Groq Node
- Drag the "Groq" node from the AI section onto your canvas
- Connect it to your Chat Trigger node
- This provides the actual AI processing power for content generation

Step 3: Configure Groq Settings
- Model: Select "llama-3.1-70b-versatile" (recommended for content creation)
- API Credentials: Add your Groq API key to N8N credentials
- Temperature: Set to 0.7 for creative but focused content

Step 4: Connect the Flow Your workflow should now show:
Chat Trigger → Groq Node

Why Groq?
- Fast Processing: Sub-second response times for quick content creation
- High Quality: Llama 3.1 70B delivers professional-grade content
- Cost Effective: Generous free tier for testing and regular use

Your agent now has the AI processing power to transform topics into engaging LinkedIn posts!

n8n | LinkedIn Post Generator

Adding Simple Memory

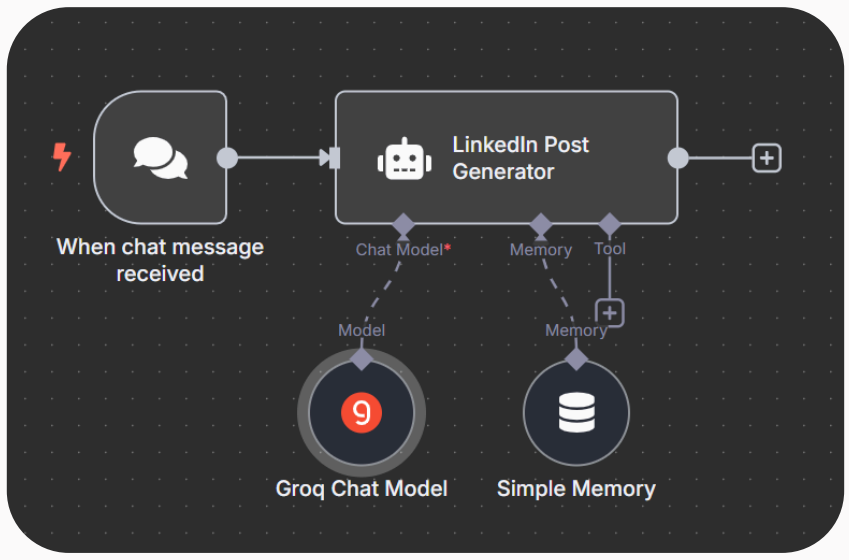

Giving Your Agent Memory

Memory allows your agent to remember previous conversations and maintain context across multiple interactions, making it more intelligent and helpful over time.

Step 1: Add Simple Memory Node

- Drag the "Simple Memory" node from the memory section on canvas
- Connect it to your workflow between the Chat Trigger and Groq nodes
- This pre-built node handles all memory storage and retrieval automatically

Step 2: Configure Memory Settings

The Simple Memory node requires minimal configuration:

- Memory Key: Set a unique identifier for your agent's memory
- Auto-manage: Let the node handle storage and context retrieval

Step 3: Connect Memory Flow

Your complete flow: Chat Trigger → Simple Memory → Groq

What Simple Memory Provides:

- **Automatic Context Management:** Stores and retrieves conversation history without manual setup.
- **Conversation Continuity:** Your agent remembers previous topics and maintains natural conversation flow.
- **Context Awareness:** References earlier discussions to avoid repetition and build on previous content.

Your agent now has built-in memory to create more contextual and intelligent LinkedIn posts!

Next: Let's test your complete LinkedIn Post Creator and see it in action.

LinkedIn Post Generator

Testing Your Agent

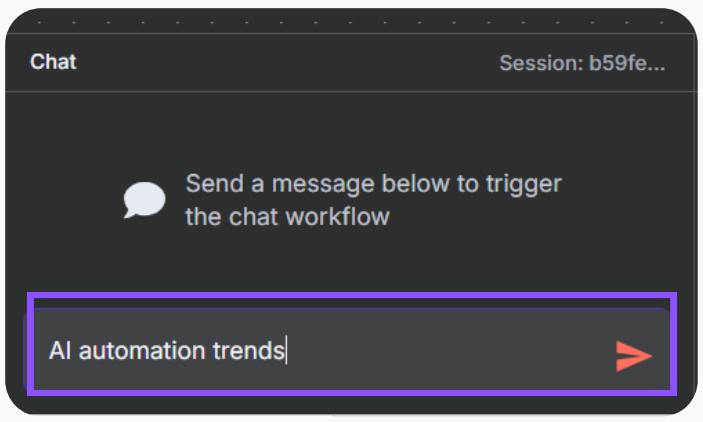

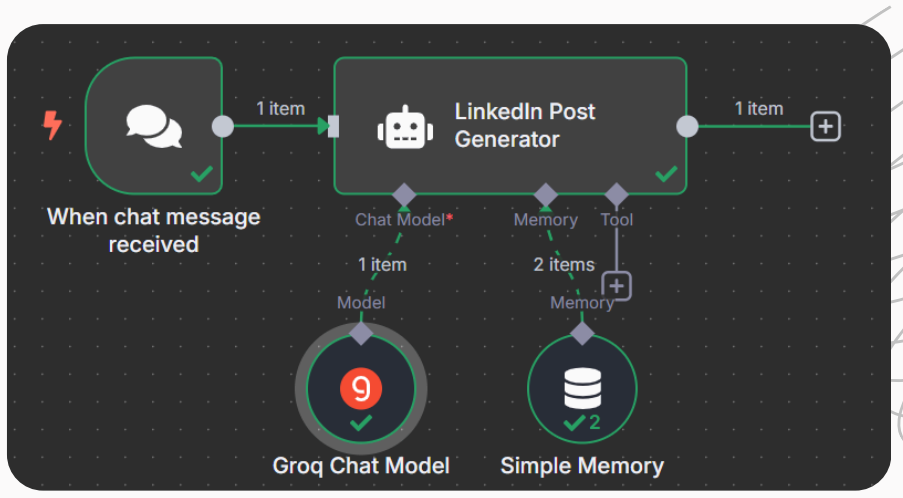

Step 1: Run Your First Test

- Click the "Test workflow" button in N8N
- Enter a topic in the chat interface (e.g., "AI automation trends")
- Watch your agent process the request and generate content

Step 2: Analyze the Output

Your agent should generate a complete LinkedIn post that includes:

- **Engaging Hook:** Attention-grabbing opening line
- **Valuable Content**: 2-3 paragraphs with actionable insights
- **Call-to-Action:** Question to drive engagement
- **Relevant Hashtags:** 3-5 hashtags for discoverability

Example Test Results:[Here you'll include a screenshot of actual output showing a well-formatted LinkedIn post generated by your agent]

Step 3: Test Different Topics

Try various topics to see how your agent adapts:

- "Remote work productivity"
- "Small business marketing"
- "Professional networking tips"

Step 4: Refine and Iterate

- Notice the consistent format and tone
- Test how memory affects follow-up conversations
- Adjust system instructions if needed for your brand voice

What You've Built: A fully functional AI agent that understands your requirements, maintains conversation context, and produces ready-to-publish LinkedIn content consistently.

n8n | LinkedIn Post Generator

Check the Output!

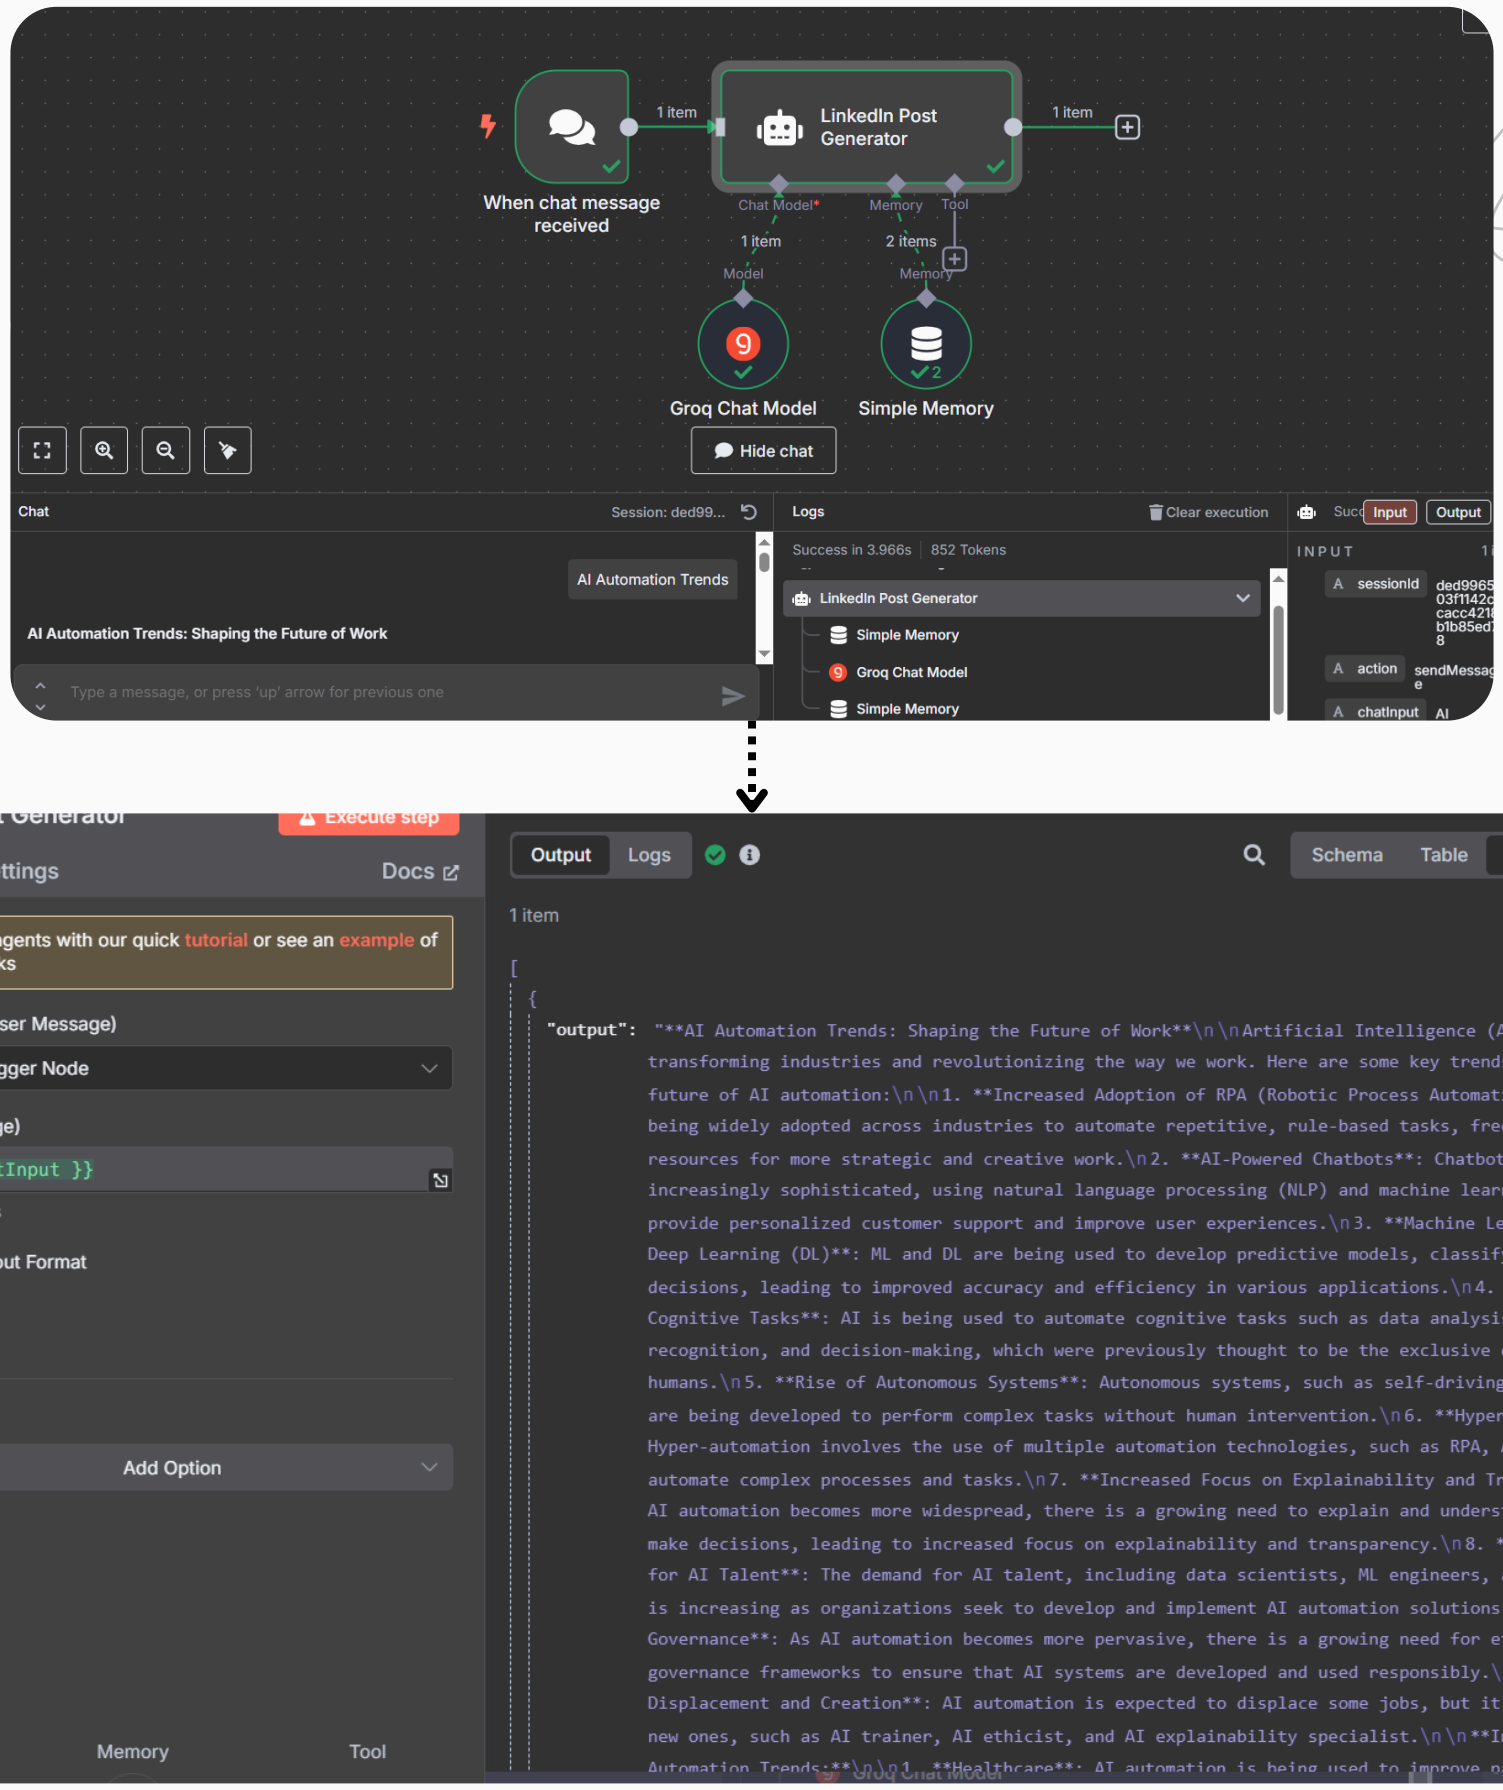

We get a full LinkedIn Post in a proper JSON Format.

Next Steps:
→ **Wire this JSON to an HTML Output**
→ **Or copy it – Make Edits – Post on LinkedIN**
→ **Add loops to create multiple posts**
→ **Add a database to build an automated content pipeline!**

Imagination is the Limit. Happy Building!