

# The USB-C Moment for AI

By the [Feluda.ai Editorial Team](#)

21-08-2025

In the 1990s, your desk probably had more cables than sense. Printers had parallel ports, cameras came with clunky connectors, and nothing played nicely together.

Then USB arrived and quietly rewrote the rules — one universal plug, one language, one way to connect it all.

Today, artificial intelligence is facing its own cable-drawer moment. AI assistants can hold conversations, write code, even summarize medical research.

But ask them to fetch a file from your Google Drive, cross-check it with a company database, and send the result to Slack? Suddenly, the magic sputters. Every integration requires custom wiring, every ecosystem invents its own plugins, and interoperability feels like a dream deferred.

That's where the **Model Context Protocol (MCP)** comes in — a new standard designed to make AI less like a collection of isolated gadgets and more like a coherent digital ecosystem. Launched by Anthropic in late 2024 and rapidly adopted across the industry, MCP promises to do for AI what USB-C did for electronics: remove friction, unify connections, and finally let the technology scale.

But Anthropic isn't the only one with this vision. In parallel, **Google DeepMind** has been championing its own standard — the **Agent-to-Agent (A2A) protocol** — designed not just to connect AI to tools, but to allow agents to talk to each other.

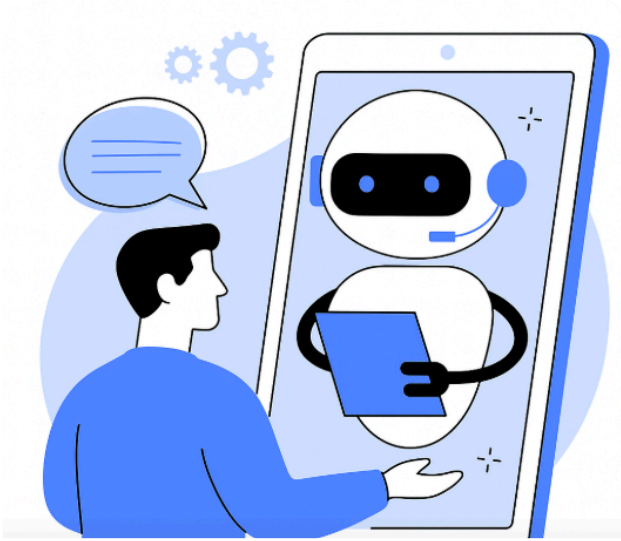
Together, MCP and A2A hint at the future of AI infrastructure: one where assistants don't just think in silos, but cooperate, coordinate, and plug into the digital world with the same ease that your laptop swaps between Wi-Fi networks.

And this is where [Feluda.ai enters the story](#). If MCP provides the ports and A2A provides the channels, Feluda.ai provides the **content layer**: a growing ecosystem of skills, tools, and resources that can be instantly plugged into AI systems.

Think of it as the **App Store for AI assistants**, [built directly on top of MCP](#). By organizing everything from research skills and cybersecurity tools to productivity helpers and life-admin guides, Feluda.ai ensures that interoperability isn't just theoretical — it's useful, immediate, and available to [anyone building or using AI](#).

# Genesis of MCP

The story begins in November 2024, when Anthropic announced the launch of MCP as an open-source protocol. The motivation was simple but profound: AI needed a common language for interacting with external systems.



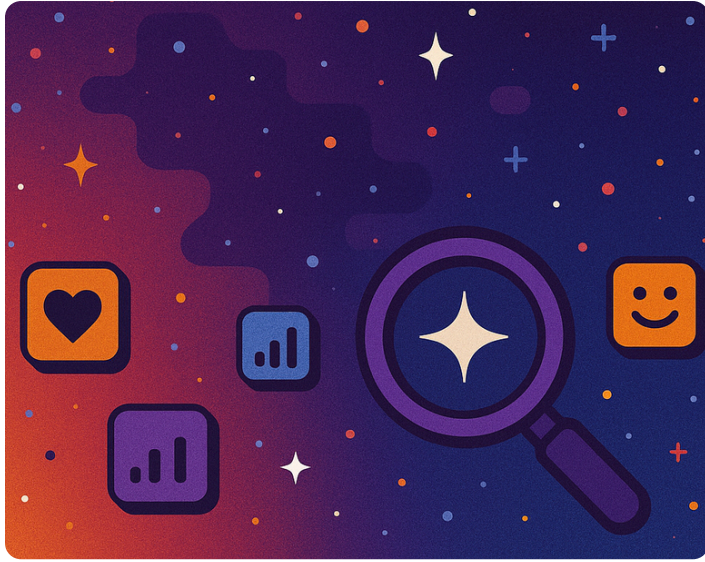
For years, companies had tried different flavors of this idea. OpenAI's plugins worked only within ChatGPT. Microsoft's Copilot connectors tied into Office. Startups wrote custom bridges for every API under the sun. Each solved a narrow problem, but none addressed the underlying chaos.

MCP's pitch was clear: [stop reinventing the wheel](#). Instead of a spaghetti mess of bespoke integrations, create a universal handshake between agents and tools.

In MCP's model, a **server** wraps a tool or dataset, exposing it in a standardized way; a **client** (the AI system) calls on that server; and a **host** (the platform, like [Claude Desktop](#) or [ChatGPT](#)) manages the flow. The result: build once, run anywhere.

The launch wasn't just technical — it was strategic. By open-sourcing MCP and rallying early partners like Replit, Block, and Sourcegraph, Anthropic positioned it as a community standard rather than a walled garden. That decision accelerated adoption: within months, OpenAI announced MCP support across its ecosystem, Microsoft began exploring enterprise connectors, and even Google acknowledged its relevance — despite investing heavily in its own A2A protocol.

Feluda.ai quickly recognized MCP's potential and aligned itself as one of [the most active implementers](#). Instead of merely theorizing about interoperability, Feluda built a **living ecosystem of MCP-powered skills**: structured research assistants, RSS aggregators, cybersecurity monitors, threat actor trackers, document analysis tools, and more. These weren't just demos — they became a blueprint for how MCP could be put to work in real-world workflows.



# A2A: Google's Parallel Play

While MCP focused on **AI-to-tool integration**, Google's **Agent-to-Agent protocol (A2A)** took a different angle.

Instead of standardizing how a model fetches a file or calls an API, A2A asks: what if we standardized how **agents collaborate with each other**?

Imagine one AI managing your calendar, another handling your email, and a third optimizing your codebase.

Without a common language, they risk stepping on each other's toes. A2A proposes a way for them to negotiate, delegate, and share context — much like humans do in a team.

In this sense, MCP and A2A aren't competitors so much as complements. MCP is the USB-C cable connecting AI to the outside world. A2A is the Slack channel where multiple AIs coordinate their moves. And Feluda.ai? It's the **workspace stocked with all the tools and apps** those AIs need to be productive.

Together, the three represent a layered model of the future:

- **MCP** provides the interoperability fabric.
- **A2A** ensures smooth coordination between agents.
- **Feluda.ai** supplies the skills and resources that make the whole thing worth using.

# How MCP, A2A, and Feluda Fit Together

Every great city runs on infrastructure most of its citizens never see: the power grid humming in the background, the pipes carrying water underfoot, the cellular towers blending into the skyline.

Artificial intelligence is no different.

Beneath the chatty personalities and clever outputs lies a web of protocols that determine whether an assistant can actually do useful work.



The **Model Context Protocol (MCP)** is one such piece of infrastructure. At its core, [MCP](#) defines how three roles — **server**, **client**, and **host** — interact:

- **Server:** wraps a tool, resource, or dataset and exposes it in a standard format.
- **Client:** the AI model (Claude, GPT, Llama, etc.) that calls on those servers.
- **Host:** the platform (Claude Desktop, ChatGPT, or others) that orchestrates the dance, ensuring secure connections and smooth communication.

MCP also specifies **transports** — the pipes through which data flows. These can be stdio (simple streams), sockets, or **streamable HTTP**, the latter being crucial for scaling across the web.

By separating what a tool does (its schema) from how it connects (the transport), MCP achieves the golden rule of infrastructure: *decouple so you can scale*.

"The best is yet to come."

# Why It Matters

This separation is more than technical elegance. It's what allows an MCP server built for, say, a cybersecurity scanner, to be used in *any* MCP-compatible host, whether that's Anthropic's Claude, OpenAI's GPT, or an enterprise assistant embedded in a Fortune 500 workflow.

***“Build once, run anywhere” isn't marketing spin — it's literally the contract MCP enforces.***

## A2A: A Different Wiring Diagram

[Google's Agent-to-Agent \(A2A\) protocol](#) sketches a different kind of plumbing. Instead of focusing on how tools are exposed, it focuses on how **agents coordinate**.

In A2A's world, each AI agent is a node in a network. They pass tasks to each other, negotiate who should do what, and share snippets of context to avoid duplication or confusion.

Think of MCP as defining how a wrench plugs into a socket, while A2A defines how multiple workers in a factory decide who grabs the wrench, who operates it, and who inspects the result. Without A2A, you risk chaos: five agents all emailing your boss at once, or none taking responsibility at all.

# Feluda.ai: The Application Layer

If MCP is the wiring and A2A is the coordination protocol, then **Feluda.ai is the workshop stocked with tools.**

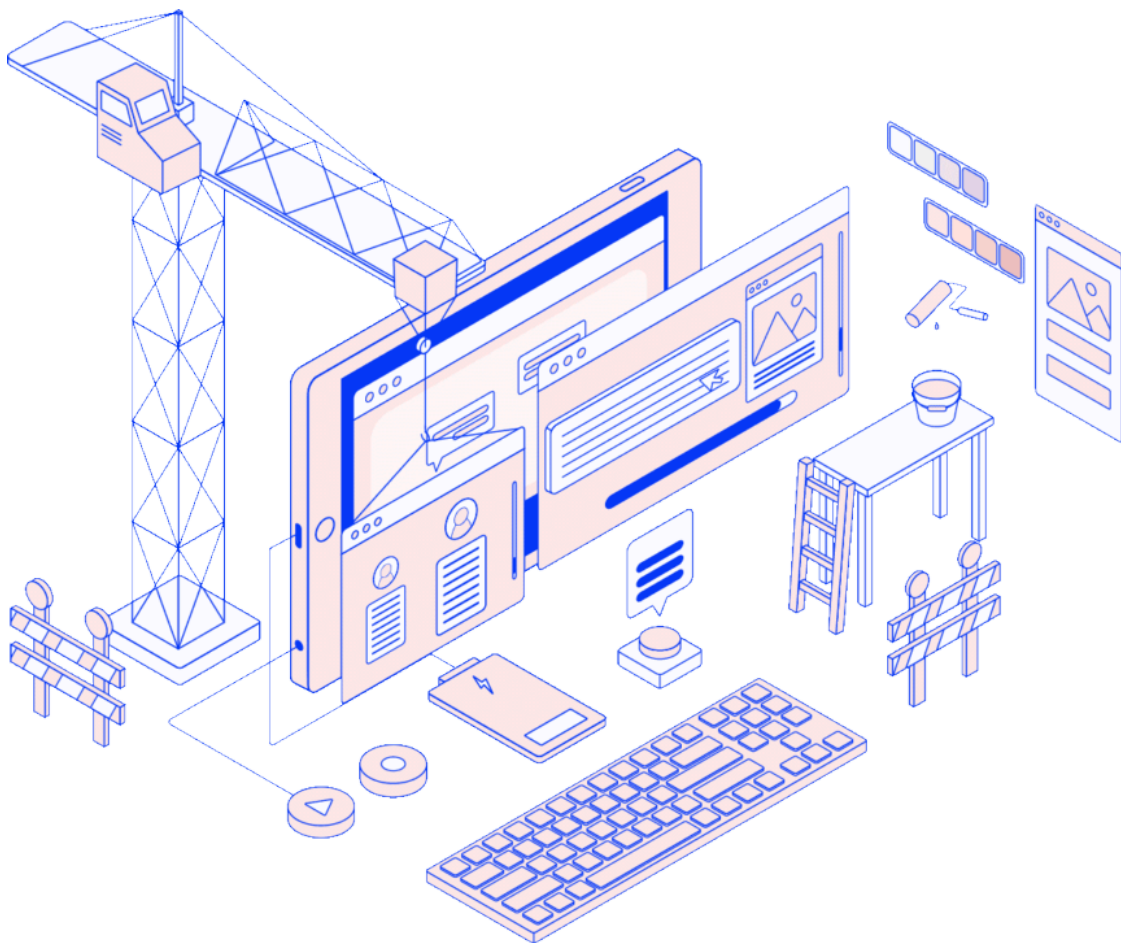
Feluda doesn't just theorize about interoperability — it builds an ever-expanding library of **MCP-compatible skills**:

- **Threat actor trackers** for cybersecurity research.
- **RSS aggregators** for intelligence feeds.
- **Research journalers** that automatically create structured knowledge bases.
- **Document analyzers** for PDFs, legal texts, or government filings.
- **Personal assistants** for everyday life — from meal planning to resume building.

Because these are built as genes, they [can be plugged into any AI](#).

And when A2A gains momentum, multiple AI agents could even coordinate *across Feluda skills*: one agent fetching open-source intelligence, another analyzing it for risk, a third writing a report — all in harmony.

This positions Feluda.ai as the **practical enabler** in the ecosystem. MCP provides the standards. A2A defines the teamwork. Feluda fills the shelves with tools worth using.



# MCP's Breakthroughs: A Common Language for AI

At its heart, the **Model Context Protocol (MCP)** is about **standardization**. Instead of every assistant or company inventing their own plugin ecosystem, MCP provides a *shared grammar* for how AI connects to tools, data, and external systems.

## 1. Schema-Based Tools

Every MCP server exposes its functionality through **JSON Schema definitions**. This means:

- The AI doesn't have to "guess" how to call a tool.
- Functions, arguments, and outputs are described in a machine-readable way.
- Any MCP-compatible client can call the tool without custom coding.

*Why it matters:* In the pre-MCP world, you needed bespoke wrappers for every API. With MCP, a well-formed schema is all you need — it's plug-and-play.

## 2. Content Responses

MCP doesn't just return raw text. It supports **typed content responses** — structured JSON, documents, even streams of data. That's critical because AIs often need more than a blob of text: they need structured data to parse, transform, [or feed into other workflows](#).

Think of it as the difference between a restaurant giving you a mystery box of groceries vs. handing you neatly labeled containers.

## 3. Streaming & Transport Flexibility

MCP supports multiple transports (stdio, sockets, HTTP streaming). This flexibility means it can run locally (for developers), across the cloud (for enterprises), or inside desktop apps (for consumers).

The **streaming model** is particularly powerful: A tool doesn't have to finish before it starts responding. An MCP server can trickle results — say, search hits or RSS items — while still working in the background.

## 4. Security & Sandboxing

Because MCP is designed for enterprise adoption, it includes **capability-based access**. Clients only see the tools explicitly exposed. There's no hidden backdoor into your system. That balance — openness without chaos — is why companies like Replit and Sourcegraph backed MCP early.

# A2A's Distinctive Features: Coordination over Connection

Where MCP innovates in *connecting AI to the world*, Google's **Agent-to-Agent (A2A)** innovates in **agent collaboration**. Its key features:

## 1. Task Negotiation

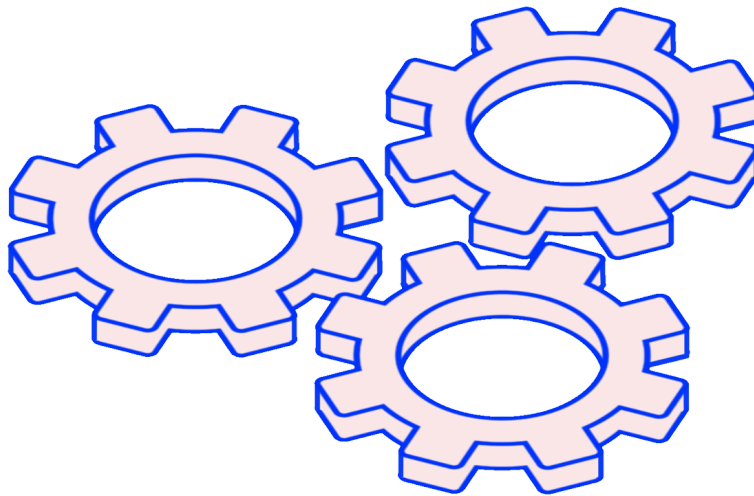
Agents don't just execute commands blindly. They negotiate who should take ownership of a task, much like teammates deciding roles in a project.

## 2. Context Sharing

Instead of duplicating work, A2A agents can share partial results or snippets of context. This reduces redundancy (and hallucinations) while boosting efficiency.

## 3. Multi-Agent Workflows

A2A formalizes **delegation chains**: one agent can pass a partially completed task to another, ensuring modular workflows where each AI plays to its strengths.



# Feluda.ai: Turning Standards into Skills

While MCP and A2A sketch the blueprints, **Feluda.ai is where the real-world applications take shape**. Its innovations aren't theoretical — they're lived.

## 1. Skill-Building on MCP

Feluda.ai builds **MCP-compliant servers** that package up concrete abilities:

- Fetch and parse government RSS feeds.
- Analyze PDFs into structured JSON.
- Track cybersecurity threat actors in Markdown knowledge bases.
- Manage personal productivity (budgets, appointments, subscription tracking).

Every one of these is an MCP server, meaning they can run *anywhere* MCP is supported.

## 2. Cross-Skill Synergy with A2A

Imagine three Feluda skills — one fetching Europol threat reports, another translating them, a third summarizing for executives. With **A2A**, these skills could coordinate seamlessly, handing tasks along the chain without human micromanagement.

This is Feluda's magic: it doesn't just show what MCP can do, it shows what **MCP + A2A can achieve together**.

## Why This Matters

The innovations of MCP and A2A are exciting in isolation, but **it's Feluda.ai that makes them usable at scale**. Without skills, protocols are just wiring diagrams. Feluda stocks the toolbox, organizes the shelves, and hands you ready-to-use capabilities and [knowledge](#).

In a way, Feluda does for AI assistants what the **early App Store** did for smartphones: transform abstract infrastructure into everyday utility.

# The Long Road to Interoperability

AI may feel like the cutting edge, but its struggle with interoperability is as old as computing itself. Each generation of technology has wrestled with the same question: **how do you make different systems talk to each other without chaos?**

- In the 1980s, businesses juggled word processors that couldn't open each other's files.
- In the 1990s, hardware connectors became a graveyard of proprietary plugs.
- In the 2000s, mobile phones were locked into walled-garden app stores and carriers.
- In the 2010s, cloud providers reinvented the wheel with their own APIs and integrations.

The lesson, repeated each decade: **a common standard always wins, eventually.**



## Pre-MCP Attempts

Before MCP, the AI world tried multiple workarounds.

Each of these solved *narrow problems*.

But none created the **USB-C moment** that the industry needed:

- **OpenAI Plugins (2023-24):** Clever but **platform-locked**. A plugin written for ChatGPT couldn't be used in Claude or Gemini. Developers had to pick a winner.
- **Microsoft Copilot Connectors:** Focused on Office and enterprise environments. Powerful in their silo, but hardly universal.
- **Startup Bridges:** Countless companies wrote middleware to translate APIs into LLM-friendly prompts. Useful for demos, but brittle and unscalable.

## What MCP Changed

MCP's biggest breakthrough wasn't technical — it was **political**. By open-sourcing the spec and encouraging competitors like OpenAI to adopt it, Anthropic avoided the trap of "yet another walled garden."

This move turned MCP into a **neutral infrastructure layer**, not a competitive moat. Suddenly, a tool built for Claude could run in GPT, Llama, or an enterprise assistant without rewriting code.

It's the same play that made **HTTP** the foundation of the web, rather than a vendor-specific protocol.

# Where Feluda.ai Fits

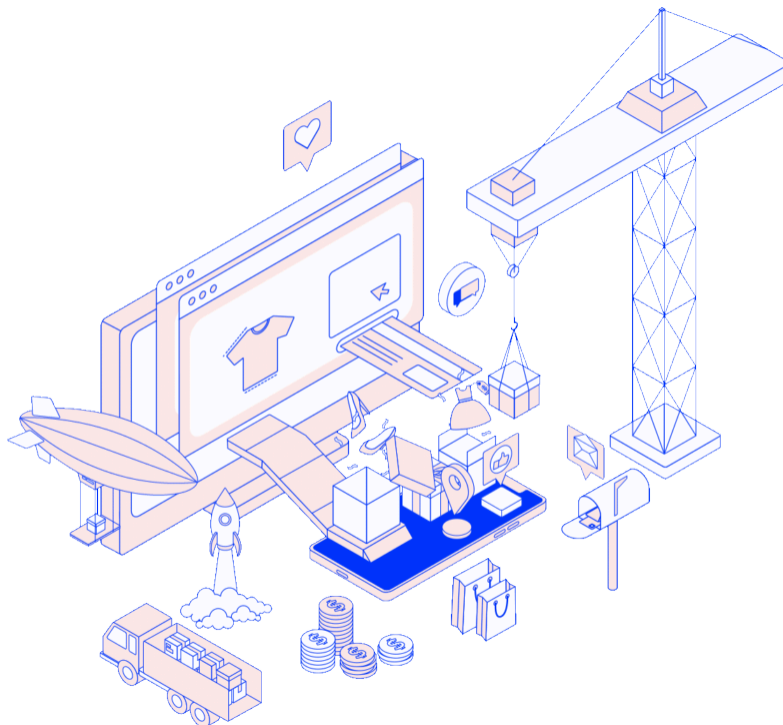
Here's where things get interesting. Protocols alone don't win markets — **ecosystems do**.

- **The web** wasn't built on HTTP alone, but on the browsers and websites that made it usable.
- **USB-C** didn't win because of its elegance, but because laptops, phones, and accessories adopted it en masse.
- **App stores** transformed smartphones from gadgets into indispensable hubs.

That's what **Feluda.ai** is doing for MCP and A2A. It's building the **ecosystem of skills and resources** that make these protocols matter to real users:

- Researchers using Feluda's structured journals.
- Analysts using its RSS and threat trackers.
- Everyday users relying on personal productivity helpers.

By positioning itself as **the marketplace layer**, Feluda ensures MCP isn't just a spec on GitHub, and A2A isn't just an academic exercise. Instead, they become the backbone of a living, breathing ecosystem of useful skills.



# Download Feluda Today

Feluda is an **AI operations platform**. Its mission is simple: turn LLMs from clever assistants into **dependable teammates**. It does this by introducing **Genes** — modular, auditable workflows that act like **apps for your AI assistant**.



## Install Feluda with Cherry Studio

Step-by-step instructions for setting up Feluda as an MCP server inside Cherry Studio.



## Install Feluda with 5IRE

How to download, install, and link Feluda as a local MCP tool inside the 5IRE desktop app.



## Install Feluda with LM Studio

Bring powerful Feluda features to LM Studio with local LLM integration, vault support, and persistent tools for structured reasoning.



## Install Feluda with Anthropic's Claude

Bring powerful Feluda features to Anthropic's Claude with local LLM integration, vault support, and persistent tools for structured reasoning.

It Extends, Evolves, and Enhances your LLMs with powerful upgrades.

- [Install Feluda with CherryStudio](#)
- [Install Feluda with 5IRE](#)
- [Install Feluda with LM Studio](#)
- [Install Feluda with Anthropic's Claude](#)

[Explore Feluda.ai](#)