



2 MIN READ

Ayfie Explores MCP: Making Enterprise AI Accessible to Every Business



 Herman Sjöberg (<https://www.ayfie.com/impact/blog/author/herman-sjoberg>) : Nov 5, 2025 9:00:00 AM



Yesterday, Ayfie announced that we're bringing MCP-powered solutions to market in the first half of 2026. And I believe this will fundamentally expand who gets to benefit from enterprise AI.

The results we've achieved so far have blown us away. We're talking about making advanced enterprise AI accessible to small and medium sized businesses — and doing it in a way that's faster, cheaper, and more practical than we've seen using RAG alone.

What is MCP, and why should you care?

Let me paint you a picture. Imagine you're running a 30-person consulting firm. You've got client data in your CRM, invoices in your accounting software, project updates scattered across email and SharePoint, and proposals living in various folders. When a client asks for a status update, you're toggling between five different systems trying to piece together the story.

Now imagine asking your AI assistant: "What's the status of the Johnson project, including outstanding invoices and next deliverables?" And getting a complete, accurate answer in seconds, pulled directly from all those systems.

That's MCP.

Model Context Protocol is essentially a standardized way for AI language models to connect directly to your existing software systems. Think of it as a universal translator that lets your AI speak fluently with Tripletex, Visma, SharePoint, and dozens of other business planning software you already use.

Traditional RAG systems require you to essentially copy all your data into a specialized database, set up pipelines, and maintain ongoing synchronization, it can support on prem and can screen through super old databases. RAG is great, and so far the best and only option for large enterprises with more legacy and compliance.

MCP takes a different path. Instead of moving your data, it lets the AI connect directly to where your data already lives. The language model reaches out to your existing systems in real-time, retrieves what it needs, and gives you an answer.

For small businesses, this is transformative. You don't need a data engineering team. You don't need to rebuild your tech stack. You just connect the systems you already use.

The token problem (and how we solved it)

Now, I need to get a bit technical for a moment, this is where things get really interesting.

When we started testing MCP, we ran into a significant bottleneck. Standard API implementations were returning massive amounts of data, sometimes up to 500,000 tokens per request. Tokens are the text segments that AI models use to process information, and they're how most AI services charge you.

Processing that much data was expensive, slow, and created too many queries, it was completely unnecessary.

Think about it this way: if you ask "What did we invoice Johnson Corp last month?" you don't need the AI to read through every invoice you've ever created. You need it to find the specific, relevant information.

That's where Ayfie Indexing Layer comes in, inspired by our 15 years of experience in RAG (Retrieval Augmented Generation).

We've developed technology that sits between the MCP connection and your systems. It intelligently filters and organizes the data before it reaches the language model. Instead of 500,000 tokens, we're delivering around 5,000 tokens per request.

The results speak for themselves:

- 99% reduction in token usage
- 10x faster response times
- Up to 90% lower costs

This isn't just an incremental improvement. It's the difference between a solution that's theoretically possible and one that's actually practical for everyday business use.

What we're building toward

Our planned launch in the first half of 2026 will open a door that's been closed to too many businesses.

With MCP and Ayfie Indexing Layer, Ayfie can now offer a cost-efficient and quick-to-implement solution that makes advanced AI accessible even to businesses with fewer than 50 employees – a segment representing over 95% of Norwegian limited companies.

Between now and our official launch, we'll continue to refine the technology, expand our integrations, and work with our partners to make sure we're building something that truly solves real problems.

And if you have questions or ideas about how this technology could help your business, reach out (</company/contact-us>).



Share

Share 0