



FRESCHÉ

MCP SERVERS EXPLAINED:

YOUR GUIDE TO GOVERNED AI AND SELF-SERVICE ANALYTICS

www.freschesolutions.com

What's Included

- 03** **What This Guide Covers**

- 04** **What is an MCP Server?**

- 06** **Why MCP Servers Matter**

- 07** **How MCP Servers Support Analytics**

- 09** **What an MCP Server Looks Like**

- 10** **How MCP Servers Compare to Other AI Tools**


- 11** **What MCP Servers Mean for Your Team**


- 12** **Is a Governed MCP Server Secure?**

- 13** **What You Need Before You Start**

- 13** **Who Should Use a Governed MCP Server**

- 14** **How Long Deployment Takes**





What This Guide Covers

Ask any data leader what they want from AI, and you'll hear a version of the same thing:

"Give my business users trusted answers from our data, without exposing what shouldn't be exposed, losing control, or ripping anything out."

That's a tall order. Traditional BI can't keep up. Public tools like ChatGPT will happily answer, but you can't put regulated data anywhere near them. And custom-built connections? Most are still stuck in pilot a year later.

So, the question isn't whether AI belongs on top of your enterprise data. It's how to do it safely, quickly, and at scale.

That's the gap MCP servers are here to close. In 2026, they're quickly becoming the new standard for secure, self-service AI analytics, whether your data lives on IBM i, Snowflake, Databricks, SQL Server, or a mix of all of the above (including systems of record like SAP or Oracle). MCP servers don't replace your stack. They serve as the connective tissue that delivers business data to end users in a governed and secure manner.


This guide will walk you through what an MCP server really is, how it works, and what it changes about the way your organization turns data into confident decisions.

In this guide, you'll learn:

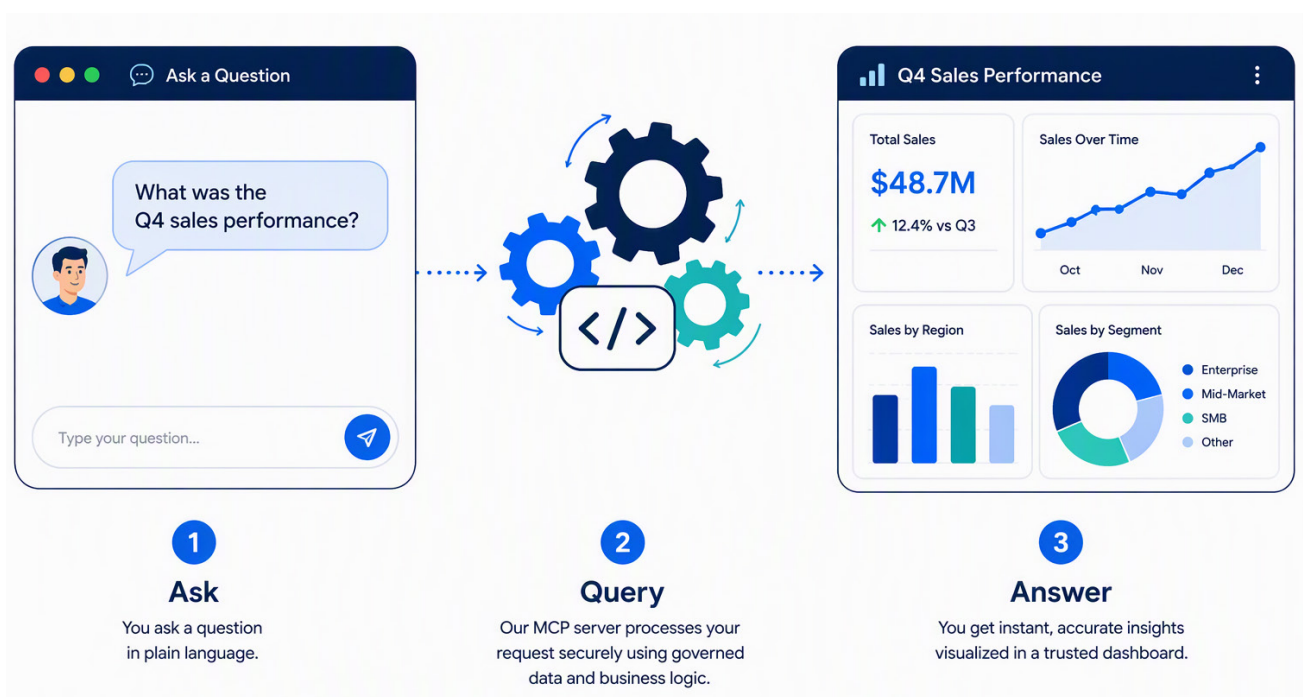
- An overview of MCP servers, in plain language
- Why MCP servers matter right now
- The 4 components of an MCP server
- The architecture required to support secure AI governance
- How the Fresche MCP server differs from Fabric Data Agents, ChatGPT, or warehouse-native AI
- Timelines and next steps for deploying MCP servers in your organization

What is an MCP Server, in Plain Language?

MCP stands for [Model Context Protocol](#), an open standard that lets AI assistants like Claude, Copilot, and ChatGPT connect to your enterprise systems through a well-defined, governed interface.



An MCP server exposes a curated library of “tools.” Each tool answers one specific business question the same way, every time. The AI picks the right tool and calls it, your system runs it, and the result comes back.



Think of it as the difference between asking a question with and without governance:

Without MCP


“Hey ChatGPT, here’s a spreadsheet I exported this morning, do your best.”

The AI works from whatever is in the file and guesses what your figures mean. Fast, ungoverned, and **different every time you ask**.

With MCP

“Hey Claude, give me Q3 bookings vs. target.”

The AI calls a pre-approved tool. Your system runs it against the definitions your finance team has already set and returns the result. **The same answer every time, for every user.**



That’s what a governed MCP tool does. Every tool in the library is pre-built and pre-validated, so the same question returns the same answer for every user, and every call is logged. The AI still does the analysis. The definition underneath it is yours.

An MCP server is where governance gets applied to AI. It reaches only the data and the definitions you have already approved.

Why MCP Servers Matter

You have spent years building data warehouses, reports, and analytics platforms. Your business users still cannot get an answer without going through an analyst, a report developer, or IT.

The problems are familiar:

- ✓ Your business users depend on IT for answers to routine questions.
- ✓ Your BI investment sits unused because it is too complex for the people it was bought for
- ✓ You make decisions before the data arrives
- ✓ Your AI projects stall at proof of concept
- ✓ You are asked to raise productivity without adding headcount
- ✓ No one can tell you how AI answers are governed, secured, or checked

Plenty of AI chat tools can connect to data. Few can deliver accurate, governed reporting at enterprise scale. Point a public LLM at your data without that governance and it can hallucinate, returning answers that look right and are not.

That is what an MCP server is built to solve

How MCP Servers Support Analytics

The Fresche MCP Analytics Platform is built on four services, starting with your data wherever it lives today.

1

Getting Trusted Data Into a Modern Warehouse

Everything starts with getting your enterprise data into a modern, AI-accessible environment without disrupting production systems.

If you are already running [Microsoft Fabric](#), Snowflake, Databricks, or a modern SQL warehouse with clean, modeled data, you are ready for the next layer. If your data still lives in systems of record like IBM i (Db2 for i), SAP, or Oracle, the first step is to replicate it into a modern data platform, where business logic can sit on top of it.

For IBM i organizations, this is often the hardest part of the journey. Moving IBM i data into a modern platform takes more than a standard data pipeline. It usually means custom engineering, significant cost, and months of work.

We have already built and tested that capability: a [pre-built IBM i to Fabric connector](#) delivering near real-time ingestion.

If your warehouse is already in place and trusted, you can go straight to the governed tools layer.

The principle holds throughout: your source systems keep running the business, and the warehouse and MCP handle the analytics.

2

A Governed Tools Library

A curated, growing library of governed tools sits on top of your warehouse. Each tool answers one business question the same way, every time, so every AI assistant and every dashboard works from the same definitions.

Categories today include:

- **Sales & Pipeline:** bookings, win/loss, whitespace, targets
- **Delivery & Capacity:** utilization, bench, forward availability, rates
- **Finance:** recognized revenue, AR/AP, invoices, margin

3


A Conversational Analyst, Through Claude and MCP

Your team asks questions in their own words. Claude connects to the governed tools through MCP, shows its work, and turns the answer into a live dashboard or a written analysis. It is built to extend to Microsoft Copilot and OpenAI models as they add MCP support.

4

A Reporting Portal

One home for trusted reporting: embedded [Power BI plus self-service dashboards](#) your team publishes, each scoped to the viewer.




What an MCP Server Looks Like

A governed MCP server lets you trace every answer back to its source. When a figure shows up in Claude, you can see which tool produced it, who ran it, and what data it touched.

Here is how we build it:

- Your source systems (IBM i, SAP, Oracle, or another system of record) stay exactly where they are. If your data needs replicating into a warehouse, it happens in near real-time, without touching production.
- Your warehouse (Microsoft Fabric, Snowflake, Databricks, or another modern SQL platform) holds the unified, governed data, with row-level security.
- Microsoft Entra ID provides single sign-on and On-Behalf-Of authentication, so every request runs as the person who asked it.
- The Fresche MCP Server, running on Azure Container Apps, exposes the governed tools to Claude.
- The Reporting Portal and Power BI deliver the web app and embedded reports, on Azure Static Web Apps.
- The Governed Metric Layer provides one definition for every figure, shared by Claude and the portal.

A secure, governed architecture



How MCP Servers Compare to Other AI Tools

MCP servers, Fabric Data Agents, direct warehouse connectors, and general-purpose ChatGPT deployments all sound similar. The difference shows up at enterprise scale.

Approach	Strength	Weakness
Microsoft Fabric Data Agents	Native to Fabric	Reaches Fabric only, limited ad-hoc reporting
Snowflake / Databricks native AI	Fast on warehouse-native queries	Reaches its own platform only, no cross-platform governance
ChatGPT (custom builds)	Flexible, widely adopted, strong models	Heavy custom development, governance gaps, hard to move into production, hallucination risk
Custom-built AI analytics	Tailored to you	Expensive, long timelines, hard to maintain and scale
Fresche MCP Analytics Platform	Governed tools, warehouse-agnostic, pre-built IBM i connectivity, live in weeks	Needs a clean, modeled data warehouse, and we get you there

What we bring is a framework that is already built. You reuse the data investment you have already made, and you go live in weeks rather than the months a custom AI project takes.

What MCP Servers Mean for Your Team

Self-service analytics changes more than the speed of reporting. It changes how your whole organization works with data.

Before	With a Governed MCP Server
Submit a request, wait in the BI queue	Ask a question, get the answer now
IT hand-codes reports from raw data	Business users get their own answers
A handful of analysts hold the keys	Far more people can answer their own questions
Static reports, stale on arrival	Living dashboards, always current
Every team's figures disagree	One trusted set of figures across the business

What this delivers:

- Decisions at conversation speed. Minutes from question to answer, while the moment still matters
- One source of truth. Every team, tool, and AI assistant works from the same governed definitions
- Shorter BI backlogs. Your experts stop filling tickets and start driving strategy
- A platform that grows. The tools library expands as we build more

Is a Governed MCP Server Secure?



Yes. If your instinct is to be skeptical about connecting an AI model to your enterprise data, that is the right instinct.

Governance is the foundation here, built in from the first day. Your source systems stay as your systems of record. MCP reads from a governed warehouse layer, so your production workloads stay isolated from AI activity.

Security by design means:

- Your production systems are never touched by AI. MCP queries run against the governed warehouse layer. Production workloads on IBM i, SAP, Oracle, or anything else stay isolated
- Your data stays yours. Everything runs on your Microsoft cloud, in your own tenant. Nothing leaves your environment
- Everyone sees only their own data. Every query runs as the signed-in user through On-Behalf-Of authentication, with row-level security. No shared service account, and no over-exposure
- Your figures agree everywhere. Every AI assistant and the portal call the same governed tools, so a figure is identical wherever you read it
- Nothing goes live unreviewed. Publishing a dashboard across the business is a reviewed, approved step

Compare that with what happens when people have no governed option. Sensitive data gets copied into a public AI tool, outside your controls, with no record of who did it or what left the building. A governed MCP server gives you productivity without that exposure.

What You Need Before You Start

There is exactly one: a data warehouse with clean, modeled, trustworthy data. That can be Microsoft Fabric, Snowflake, Databricks, SQL Server, or another modern SQL warehouse. What matters is that the data is queryable, reliable, and governed.

If your data still lives in a system of record like IBM i, SAP, or Oracle, the first step is getting it into that warehouse, through near real-time replication that leaves production undisturbed. For IBM i organizations, our pre-built IBM i to Fabric connector removes months of custom pipeline engineering. For other systems, the same principle applies: near real-time, low-overhead, governed replication into the warehouse you choose.

Get this right once, and every team, every future tool, and every AI assistant runs on it. Not there yet? A data-readiness assessment starts every engagement.

Who Should Use a Governed MCP Server

The strongest fit is an organization above \$50M in revenue where IT, finance, and operations leaders are asking:

- “How do we give the business self-service analytics without exposing PII or losing control?”
- “We are investing in Microsoft Fabric, or Snowflake, or Databricks. How do we get more business value from it?”
- “How do we get the data in our systems of record ready for AI, without a full migration?”
- “We do not want to rip and replace. How do we bridge what we already have to what comes next?”

This fits especially well for:

- IBM i organizations ready to bring decades of trusted Db2 for i data to modern AI, without disrupting production. Our pre-built IBM i to Fabric connector shortens that path
- Microsoft Fabric organizations ready to get more return from Fabric by adding a governed AI analytics layer their business users will use
- Any organization with a modern data warehouse (Snowflake, Databricks, SQL Server, or a mix) that needs governed, self-service AI without months of custom pipeline work

How Long Deployment Takes

Faster than most people expect. The Fresche MCP Analytics Platform comes in three tiers.

Level	Length	What's Included
Bronze	2 Weeks	Core MCP implementation and training, customer-driven adoption, and 2 weeks of usage-based tool tuning
Silver	4 Weeks	Everything in Bronze, plus tuning and testing of curated tools
Gold	6 Weeks	Everything in Silver, plus the Reporting Portal, extended training, and a structured adoption program across the business

Delivery runs in four steps: Implement, Train, Tune, Embed. Installing the tools is the quick part. The rest of the time goes into making self-service the way your organization works.

Key Takeaways

- A governed MCP server is where your governance rules get applied to AI, so AI only ever reaches what you have approved
- The data is already there. Getting to it is the problem, and that is what MCP solves
- Bridge, don't replace. Your source systems stay in place, on IBM i, SAP, Oracle, or a modern warehouse, and MCP adds the analytics layer
- A governed tools library, rather than raw database access, is what makes AI answers consistent and safe at enterprise scale
- Security is architectural: On-Behalf-Of authentication, row-level security, and your own tenant from day one
- The BI backlog shrinks when business users can answer their own routine questions
- The one prerequisite is a clean, modeled data warehouse. Not there yet? A data-readiness assessment gets you there

Bring Us Your Question

Book a demo with our experts and bring one real question from your business. We'll answer it using governed data and build the dashboard during the session, so you can see exactly how it works.

[Schedule Your Demo](#)

FRESCHÉ

www.freschesolutions.com