

# Local vs Cloud-Based MCP Servers

MCP servers can run on your own machine or in the cloud. Each approach has different strengths — and we build both.



## Local MCP Server

Runs on the user's machine

A local MCP server runs as a process on your computer. It starts when the AI client (like Claude Desktop or Cursor) launches it, and communicates over standard input/output (stdio). No cloud infrastructure, no deployment pipeline, no server costs.

- ✓ Zero infrastructure — runs on any developer's machine
- ✓ No organization-scoped deployment required
- ✓ Works with Claude Desktop, Cursor, Claude Code, and other MCP clients
- ✓ Data stays local — API calls go directly from your machine
- ✓ Fast iteration — update the server, restart the client, done
- ✓ User-level authentication (each user uses their own credentials)
- ⚠ Must be installed on each individual machine
- ⚠ Only available when the machine is running



## Cloud-Based MCP Server

Hosted on your infrastructure

A cloud-based MCP server runs on Azure, AWS, or your own infrastructure. It uses HTTP/SSE transport instead of stdio, making it accessible to any authorized client over the network. Ideal for teams that need centralized control.

- ✓ Centralized — one deployment serves the entire team
- ✓ Always available — not tied to any individual's machine
- ✓ Centralized logging, monitoring, and access control
- ✓ Can integrate with service accounts and managed identities
- ✓ Supports SSE transport for real-time streaming
- ✓ Scalable across departments and teams
- ⚠ Requires cloud infrastructure and hosting
- ⚠ Needs organization-level deployment and security review

## Local Servers with Claude Code — No IT Department Required

One of the most powerful aspects of local MCP servers is that they work with **Claude Code** and similar AI-powered development environments without any organization-scoped deployment or IT approval.

A developer or power user can install a local MCP server on their own machine, configure it in their Claude Code or Cursor settings, and immediately start using AI with live business data. There is no Azure deployment, no admin consent flow, no infrastructure to maintain.

The server authenticates using the user's own credentials (OAuth, API keys, or tokens stored locally), so security permissions are enforced by the target systems themselves — the MCP server simply acts as a passthrough that structures the data for AI consumption.



**Getting started takes minutes, not weeks.** Install the server, add it to your Claude Code config, and you're connected. No tickets, no waiting for provisioning, no cloud costs.