What is the Model Context Protocol (MCP)?

1 November 2026 · 5 min read · blogs.iictp.org/what-is-model-context-protocol

The Model Context Protocol (MCP) is an open protocol for connecting AI applications to tools and data sources in a standard way. Instead of writing a custom integration for each pairing, a tool exposes an MCP server once and any MCP-capable application can use it.

TL;DR The short version

- The Model Context Protocol (MCP) is an open protocol for connecting AI applications to tools and data in a standard way.
- Each application and tool implements it once, turning an N-by-M integration problem into N plus M.
- Do not connect any MCP server blindly; start with a well-known server against test data and read the tool descriptions the model sees.

Read online, listen to it, or change the text size: https://blogs.iictp.org/what-is-model-context-protocol

How does MCP work?

An AI application, called the host, runs an MCP client that talks to one or more MCP servers. Each server describes what it offers, the client discovers that and presents it to the model, and the model can then ask for a tool to be used.

What can an MCP server offer?

It can offer tools, which are actions the model may request, such as searching a ticket system. It can offer resources, which are data the application can read, such as a file or a record. It can also offer prompts, which are reusable templates a user can choose.

Why does it matter?

It turns an N-by-M integration problem into N plus M, because each application and each tool implements the protocol once. It also puts tool access in one place where it can be reviewed, instead of scattered through application code.

Is it safe to connect any MCP server?

No. A server is software that runs with real access, so treat a third-party server as untrusted code until it has been reviewed. Give it the least privilege it needs, keep credentials out of prompts, and require human confirmation for actions that change things. Content a server returns can also carry hidden instructions; see what is prompt injection.

How do you get started?

Run a well-known server locally against a test dataset, and read the tool descriptions the model will actually see. Limit its permissions, widen them only when needed, and log every tool call so you can see what the model did.

Related credential: AI Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Model Context Protocol. Specification and documentation. modelcontextprotocol.io.
2. OWASP. Top 10 for LLM Applications. genai.owasp.org.

Published by IICTP, an independent technical credential body, on 1 November 2026. The current version of this post is at https://blogs.iictp.org/what-is-model-context-protocol. More guides: blogs.iictp.org.