Understanding MCP

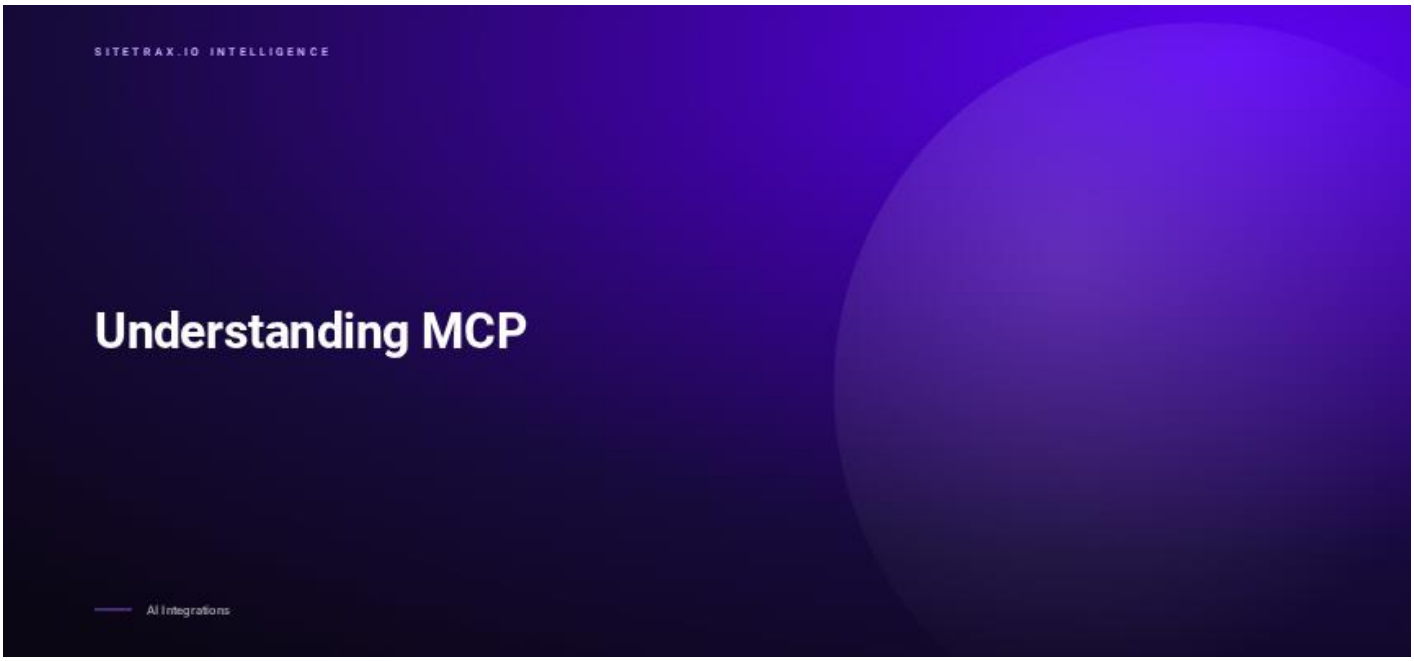

Model Context Protocol, usually shortened to MCP, is an open standard that lets an AI assistant connect to an outside system and use it. SiteTrax.io implements it so that assistants can work with your operational data without anyone exporting it first.

Why it exists

Before MCP, getting an assistant to answer a question about your yard meant copying data into a prompt or building a bespoke integration for each assistant. Both approaches age badly. A copied extract is stale the moment it is pasted, and a bespoke integration has to be rebuilt for the next tool.

MCP replaces that with one connection the assistant understands. SiteTrax.io publishes what it can do once, and any compatible assistant can use it.

What it means in practice

Without MCP	With MCP
Export a report, paste it into a chat	The assistant can retrieve current records when it calls a SiteTrax.io tool

Without MCP	With MCP
The extract reflects one moment and one filter	A new tool call retrieves data for its requested scope and time period; a reply without one may use earlier context
The assistant cannot check its own answer	The assistant can retrieve the supporting record and permitted image links
Sharing the extract shares everything in it	Access follows the signed-in SiteTrax.io account

How SiteTrax.io implements it

- **It is a remote server.** The endpoint is `https://mcp.sitetrax.io/mcp`. An assistant that only supports servers running on your own machine cannot reach it.
- **Authorization is OAuth-based.** You approve the assistant's access on the SiteTrax.io authorization screen using an active account with verified email. Application login alone does not grant an assistant access, and publishing an Intelligence profile is optional for ordinary MCP queries. No SiteTrax.io password is entered into the assistant.
- **Permissions are scoped.** Read access covers operational queries, selected published profile context, notification lists and delivery history. Creating an alert or digest, or cancelling one, needs read access plus an additional write scope. A client may request it during initial connection or later.
- **Analysis happens on the server.** SiteTrax.io calculates supported totals, groupings and durations itself, so the assistant does not need a large raw-data download for those calculations. It can separately retrieve permitted supporting records.
- **SiteTrax.io stays the system of record.** Operational records remain in SiteTrax.io. The assistant retrieves authorized results and can manage supported personal notifications when separately permitted and confirmed.

What it does not do

- It does not widen your SiteTrax.io permissions. A project not shared with you stays out of reach.
- It does not require a bulk export or a synchronized index. However, the records, summaries and any permitted evidence returned for a query are sent to the assistant provider you choose. Follow your organization's policies for using that provider.
- It does not make an unverified answer verified. Interpretation rules still apply.

See Account Access and OAuth for the account and authorization flow. For the full technical contract, including transport, scopes and response semantics, see the SiteTrax.io MCP Server reference.

Revision #3
Created 2026-09-04 15:18:23 UTC by SiteTrax.io Admin
Updated 2026-09-22 07:13:48 UTC by SiteTrax.io Admin