MCP SERVER TEST REPORT

Microsoft Learn MCP

https://learn.microsoft.com/api/mcp

TESTED	Wednesday, November 11, 2026, 1:00 AM PT
TRANSPORT	Streamable HTTP · None. Public. No API key.
RUNS	5 (latency p50 over all runs; schema and tool-call bodies from run 1)
SERVER	Microsoft Learn MCP Server · version 1.0.0
PROTOCOL	2025-06-18
LISTING	https://agentscouthq.com/tools/microsoft-learn-mcp.html



OVERALL SCORE

KEY FINDINGS

Every connection and tool call succeeded on all 5 runs, and each one finished in under 1 second.

1. Mark query as required on microsoft_docs_search, and remove "default": null. — about 5 min, one-line schema change
2. Stop advertising empty surfaces — or fill them. — rough estimate: a small capability change
3. Add outputSchema to microsoft_docs_fetch, and enrich serverInfo. — rough estimate: a small metadata change

PROTOCOL AND CONNECTIVITY	CORRECTNESS AND RELIABILITY	LATENCY	SCHEMA AND TOOL QUALITY	CAPABILITY HONESTY AND METADATA
25 / 25	25 / 25	14 / 15	16 / 20	10 / 15

A strong, public, first-party documentation server. Every tool call succeeded, in well under a second. A handful of schema and capability-honesty issues keep the score short of 100 — and they are fixable.

What this means for you

You can rely on this server for everyday use. It is a public Microsoft service, and every documentation search, page fetch, and code-sample lookup we ran came back successfully.

It is fast, too: connecting and every request finished in well under a second.

What holds the score at 90, not 100, is incomplete self-description. The server says it offers extra menus (prompts and resources), but both lists were empty, and two of the tool definitions are missing small but important details.

Make the three fixes listed on the cover, then book a free retest so we can show the score improving.

Verdict

Microsoft Learn MCP is a strong, public, first-party server. It negotiates protocol 2025-06-18, issues a session, lists three well-described read-only tools, and returned useful documentation and code samples on every call, in well under a second. The daily AgentScout HTTP reachability check already had this listing as live (homepage and GitHub both HTTP 200, checked November 10, 2026, 7:03 AM PT).

The score is not 100. The server advertises prompts and resources, then returns empty lists for both. microsoft_docs_search does not mark its query argument as required. microsoft_docs_fetch has no outputSchema, unlike the other two tools. serverInfo is only a name and version. Those are fixable.

Scorecard

AREA	POINTS	WHY
Protocol and connectivity	25 / 25	Initialize HTTP 200, protocol 2025-06-18 accepted, session header set, notification accepted, 5/5 connects.
Correctness and reliability	25 / 25	3/3 tool calls succeeded on every run. Search, fetch, and code-sample results were substantial and on-topic. Zero errors.
Latency	14 / 15	Connect p50 145.2 ms. Tool calls p50 272.7–515.6 ms. One notifications/initialized outlier at 551.5 ms.
Schema and tool quality	16 / 20	Excellent titles, descriptions, annotations (readOnlyHint, idempotentHint, destructiveHint: false). Search query is not required and has "default": null. Fetch has no outputSchema.
Capability honesty and metadata	10 / 15	Advertises prompts and resources with listChanged: true, but both lists are empty. serverInfo has only name and version (no title, website URL, or description). Initialize instructions themselves are excellent.
Total	**90 / 100**	

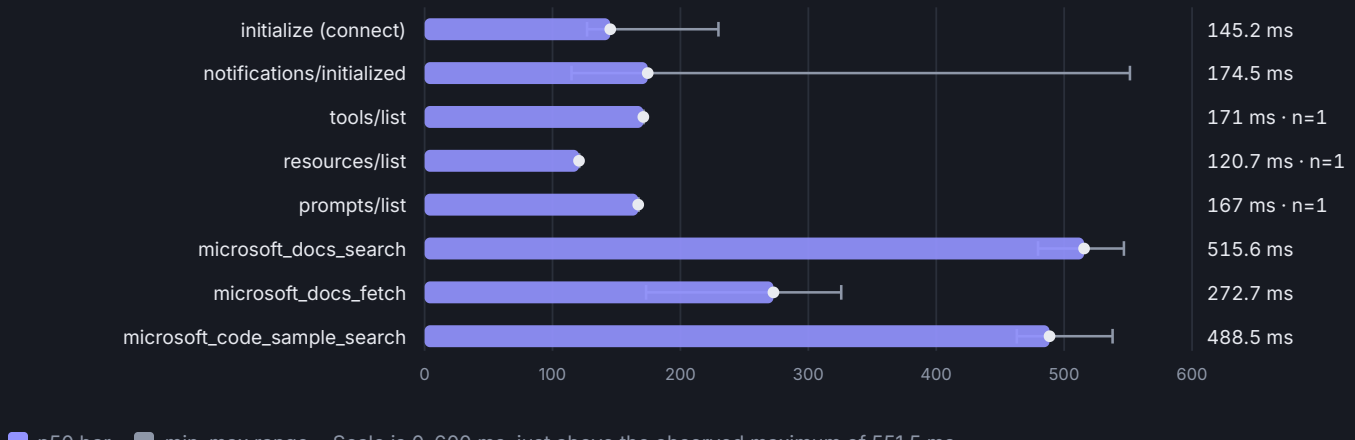
3	0	0
TOOLS LISTED	RESOURCES LISTED	PROMPTS LISTED

What we connected to

Listed name	Microsoft Learn MCP	AgentScout listing	https://agentscouthq.com/tools/microsoft-learn-mcp.html
Website in the catalog	https://learn.microsoft.com/en-us/	GitHub	https://github.com/MicrosoftDocs/mcp
MCP endpoint tested	https://learn.microsoft.com/api/mcp	Transport	Streamable HTTP (Accept: application/json, text/event-stream)
Auth	None. Public. No API key.	serverInfo	name: Microsoft Learn MCP Server · version: 1.0.0
Negotiated protocol	2025-06-18	Daily HTTP check	November 10, 2026, 7:03 AM PT · url 200, github 200, snapshot label OK
Client	agentscout-mcp-test-sample 0.1.0	User-Agent	AgentScoutMCPTest/0.1 (+https://agentscouthq.com/about.html)

Latency (p50, with min–max range)

Median over 5 runs unless marked n=1 (tools/list, resources/list, and prompts/list were recorded on run 1 only). Connect and every tool call stayed under one second. The bar and dot are p50; the whisker is the observed min–max, in milliseconds from the transcript.



Success rate by protocol step

Success means HTTP 200 or 202 and no transport, JSON-RPC, or tool-level error. resources/list and prompts/list returned empty arrays (HTTP 200); that honesty issue is scored in the checks table, not as a failed call.



Exact p50, min, and max figures are in the appendix at the end of this report.

Checks

16 checks · 10 pass · 4 warn · 2 fail. Status reflects the written scorecard; every evidence figure is from the transcript.

AREA	CHECK	RESULT	EVIDENCE
Protocol and connectivity	initialize succeeds and negotiates protocol 2025-06-18	PASS	HTTP 200 on 5/5 runs. Negotiated protocol 2025-06-18. Session header set (value not shown).
	notifications/initialized accepted	PASS	HTTP 202 on 5/5 runs.
	tools/list returns tools	PASS	HTTP 200, 3 well-described tools (run 1).
Correctness and reliability	microsoft_docs_search succeeds	PASS	5/5 runs, zero errors. Results on-topic.
	microsoft_docs_fetch succeeds	PASS	5/5 runs, zero errors. Markdown starts at "# Azure Functions HTTP trigger".
	microsoft_code_sample_search succeeds	PASS	5/5 runs, zero errors. Official C# samples include an HttpExample HttpTrigger function.
	No write or destructive tools offered or called	PASS	All three tools are annotated read-only, idempotent, and not destructive.
Latency	Connect and tool-call p50 under one second	PASS	initialize p50 145.2 ms (min 127 ms, max 229.7 ms, n=5). Tool-call p50 272.7 ms–515.6 ms.
	notifications/initialized consistency	WARN	p50 174.5 ms (min 114.9 ms, max 551.5 ms, n=5). One outlier at 551.5 ms on run 3 (still under one second, HTTP 202).
Schema and tool quality	Titles, descriptions, and annotations	PASS	All three tools have a title, a detailed description, and readOnlyHint / idempotentHint / destructiveHint: false.
	microsoft_docs_search marks query as required	WARN	query exists and is well described, but it is omitted from required and has "default": null.
	microsoft_docs_fetch declares outputSchema	WARN	Search and code-sample search both declare an outputSchema. Fetch does not.
Capability honesty and metadata	Advertised prompts match prompts/list	FAIL	initialize advertises prompts.listChanged: true, but prompts/list returned {"prompts": []} with HTTP 200.
	Advertised resources match resources/list	FAIL	initialize advertises resources.listChanged: true, but resources/list returned {"resources": []} with HTTP 200.
	serverInfo includes title, website URL, and description	WARN	serverInfo has only name ("Microsoft Learn MCP Server") and version (1.0.0).
	initialize instructions are useful	PASS	The instructions correctly describe a three-step workflow: search for breadth, code-sample search for examples, fetch for depth.

Tools, resources, prompts

Tools: 3. Resources: 0. Prompts: 0. No schema issues rose to “missing description” or “missing inputSchema”. The issues are the un-required search query, the missing fetch outputSchema, and the advertised-but-empty prompt and resource surfaces.

TOOL	TITLE	ANNOTATIONS	INPUT	OUTPUTSCHEMA	NOTES
microsoft_docs_search	Microsoft Docs Search	read-only, idempotent, not destructive	query (string) — not required; "default": null	Yes	Has outputSchema. Description is detailed and tells the client to follow up with fetch.
microsoft_code_sample_search	Microsoft Code Sample Search	read-only, idempotent, not destructive	query (string) — required; language (string) — not required; "default": null	Yes	Has outputSchema. Language enum documented in the description (csharp, python, etc.).
microsoft_docs_fetch	Microsoft Docs Fetch	read-only, idempotent, not destructive	url (string) — required	No	No outputSchema. Description states the URL must be a microsoft.com HTML docs page.

Tool-call results (run 1, all read-only)

The same three calls succeeded on runs 2–5 with no errors. Search and code-sample results are returned both as stringified JSON inside content[].text and as structuredContent. Fetch returns Markdown in content only. Times below are the run 1 HTTP times from the transcript.

TOOL	ARGUMENTS	HTTP TIME	RESULT
microsoft_docs_search	query: Azure Functions HTTP trigger	538.4 ms	Success. 10 structured results plus a 47,546-character payload (~47.5 KB) of official Learn chunks. First hit: “Azure Functions HTTP trigger (programming-language-powershell)”.
microsoft_docs_fetch	url: https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger	272.7 ms	Success. Full page converted to Markdown (62,521 characters, ~62.5 KB), starting at “# Azure Functions HTTP trigger”.
microsoft_code_sample_search	query: Azure Functions HTTP trigger C#	537.9 ms	Success. 10 official C# samples (34,029 characters, ~34.0 KB), including an HttpExample HttpTrigger function.

Top 3 fixes

These are the changes that would close the largest remaining gaps in the score. Each one is a schema or capability change — no behavior change is required for the tools that already work.

1 Mark query as required on microsoft_docs_search, and remove "default": null.

The property exists and is well described, but it is omitted from required and is given "default": null. A client can legally call the tool with no query. The other two tools correctly require their main argument. This is a one-line schema change that prevents empty searches.

1. Open the inputSchema for microsoft_docs_search.
2. Add query to the required array, matching microsoft_code_sample_search and microsoft_docs_fetch.
3. Remove "default": null from the query property.
4. Redeploy and confirm that a call with no query is rejected.

2 Stop advertising empty surfaces — or fill them.

initialize capabilities include prompts.listChanged and resources.listChanged, but prompts/list and resources/list both returned empty arrays with HTTP 200. Clients may offer empty menus, or assume list-changed notifications that never matter.

1. Decide whether prompts and resources are in scope for this server.
2. Option A — fill them: add one or two prompts that wrap the documented workflow (for example “Search Microsoft docs” and “Fetch a Learn page”), and/or a resource template for Learn URLs.
3. Option B — be honest: omit prompts and resources from initialize capabilities until there is something to list.
4. Redeploy and confirm that prompts/list and resources/list match the advertised capabilities.

3 Add outputSchema to microsoft_docs_fetch, and enrich serverInfo.

Search and code-sample search both declare an outputSchema. Fetch is the tool that returns the largest body, and it only has unstructured content. serverInfo is only a name and version. A small metadata change makes this clearer for registries, catalogs, and typed clients.

1. Declare an outputSchema on microsoft_docs_fetch (for example Markdown text plus the canonical URL).
2. Add title, websiteUrl, and a one-line description to serverInfo. A website URL here should be https://learn.microsoft.com/ (or the GitHub repo).
3. Optional: add icons so MCP client UIs can display the server consistently.
4. Redeploy and confirm typed clients can validate the fetch result.

Also worth fixing

These two items did not cost as many points as the top three fixes. They are still worth fixing.

1. serverInfo is thin. Only name and version. Context7, by comparison, also returns title, websiteUrl, and description on initialize. Optional icons would help MCP client UIs.
2. Search results are double-encoded as text. The content[0].text value is a JSON string of the same results already present in structuredContent. Harmless, but messier for clients that only read text content.

Share a “Tested by AgentScout” badge

Add this badge to your site or GitHub README. It links to this server’s AgentScout listing. It is a dated test result from this report, not an endorsement.



This badge is a dated test result, not an AgentScout endorsement. Scores can change as the server is updated.

HTML

The image URL will be filled in at launch. The link already points at the public listing.

```
<a href="https://agentscouthq.com/tools/microsoft-learn-mcp.html">
  <img src="(badge image URL added at launch)"
      alt="Tested by AgentScout · 90/100 · Nov 11, 2026. Dated test result,
not an endorsement."
      width="440">
</a>
```

Markdown

Paste this into a README. Same listing link, same dated-result meaning.

```
[![Tested by AgentScout · 90/100 · Nov 11, 2026](badge image URL added
at launch)](https://agentscouthq.com/tools/microsoft-learn-mcp.html)
```

NEXT STEP

Book a free retest after you fix these

After you make the three fixes, book a free retest. We will run the same checks again and show whether the score improved.

https://agentscouthq.com/ (booking link added at launch)

How this was tested

1. Reused AgentScout's existing daily MCP reachability result (HTTP HEAD-then-GET of the listed website and GitHub URL — this does not speak MCP).
2. Connected for real over Streamable HTTP: JSON-RPC initialize, then notifications/initialized, then tools/list, resources/list, and prompts/list.
3. Called three safe, read-only tools with harmless documentation queries. No write, send, pay, or delete tools were present or called.
4. Repeated initialize and the same three tool calls for five runs and took the median (p50) latency.
5. Linted tool schemas (descriptions, required fields, types, annotations).

Client: agentscout-mcp-test-sample 0.1.0. User-Agent: AgentScoutMCPTest/0.1 (+https://agentscouthq.com/about.html). Raw JSON-RPC transcript (all five runs, including full tool-call bodies) is the source of every latency number and schema claim in this report.

initialize capabilities returned

```
{
  "logging": {},
  "prompts": {
    "listChanged": true
  },
  "resources": {
    "listChanged": true
  },
  "tools": {
    "listChanged": true
  }
}
```

The server also returns a long instructions block that correctly describes a three-step workflow: search for breadth, code-sample search for examples, fetch for depth.

Reproducible commands

These speak MCP. They do not submit anything, and they only read documentation.

```
# initialize (expect HTTP 200, text/event-stream, mcp-session-id)
curl -sS -D - -o /tmp/mslearn-init.sse \
  -X POST 'https://learn.microsoft.com/api/mcp' \
  -H 'Accept: application/json, text/event-stream' \
  -H 'Content-Type: application/json' \
  -H 'MCP-Protocol-Version: 2025-06-18' \
  -H 'User-Agent: AgentScoutMCPTest/0.1 (+https://agentscouthq.com/about.html)' \
  --data '{"jsonrpc":"2.0","id":1,"method":"initialize","params":{"protocolVersion":"2025-06-18","capabilities":{},"clientInfo":{"name":"agentscout-mcp-test-
sample","version":"0.1.0"}}}'

# Then: notifications/initialized (send the session header from step 1),
# tools/list, resources/list, prompts/list, and the three read-only tool calls:
#   microsoft_docs_search        {"query":"Azure Functions HTTP trigger"}
#   microsoft_docs_fetch         {"url":"https://learn.microsoft.com/en-us/azure/azure-functions/functions-bindings-http-webhook-trigger"}
#   microsoft_code_sample_search  {"query":"Azure Functions HTTP trigger C#"}
```

What this sample does not include

- Authenticated or paid MCP servers (OAuth, API keys, x402).
- stdio / npx servers (not needed here; this endpoint is public HTTP).
- Write-path or destructive tool calls (none were offered).
- Changing the AgentScout daily check, the listing page, or Microsoft's server.

Disclaimer

This is a point-in-time test of a public MCP endpoint. All protocol calls were read-only. Results can change as the server is updated. The existing daily AgentScout check only requests the listed website and GitHub URL over HTTP; this report is the protocol-level step on top of that. Nothing in this sample was published, emailed, or changed on agentscouthq.com.

Test date: Wednesday, November 11, 2026, 1:00 AM PT. Time zone: America/Los_Angeles (Pacific Time).

Appendix: Exact latency figures

Median over 5 runs unless noted as run 1 only. All values are in milliseconds, taken from the transcript. Connect and every tool call stayed under one second.

STEP	P50	MIN	MAX	N
initialize (connect)	**145.2 ms**	127 ms	229.7 ms	5
notifications/initialized	**174.5 ms**	114.9 ms	551.5 ms	5
tools/list	**171 ms**	171 ms	171 ms	1 (run 1 only)
resources/list	**120.7 ms**	120.7 ms	120.7 ms	1 (run 1 only)
prompts/list	**167 ms**	167 ms	167 ms	1 (run 1 only)
microsoft_docs_search	**515.6 ms**	479.6 ms	546.7 ms	5
microsoft_docs_fetch	**272.7 ms**	173.2 ms	325.7 ms	5
microsoft_code_sample_search	**488.5 ms**	463 ms	537.9 ms	5