

MCP DEV SUMMIT · APRIL 3

Path to V2 for MCP SDKs

Max Isbey · Anthropic · MCP Python SDK maintainer





Max Isbey

Python SDK maintainer

FROM New Zealand

PREVIOUSLY Rocket Lab — Operations Flight Software Engineer / AI stuff

NOW MCP Python SDK · Anthropic

FOCUS Open Source MCP Python SDK Maintainer

Today

01

Why horizontal scaling is painful

Transport history and the horizontal scaling problem

02

MRTR

Multi-round-trip requests and compatibility layers

03

What we changed

FastMCP rename, low-level cleanup, pluggable transports

04

Migration Eval Framework


Automated migration guide testing

The original transports

November 2024

Standard I/O

Local, subprocess



✓ Bidirectional · simple · still most common

HTTP + SSE deprecated

Remote, persistent connection



✗ Connection drops → session lost · 1 conn per client

Streamable HTTP

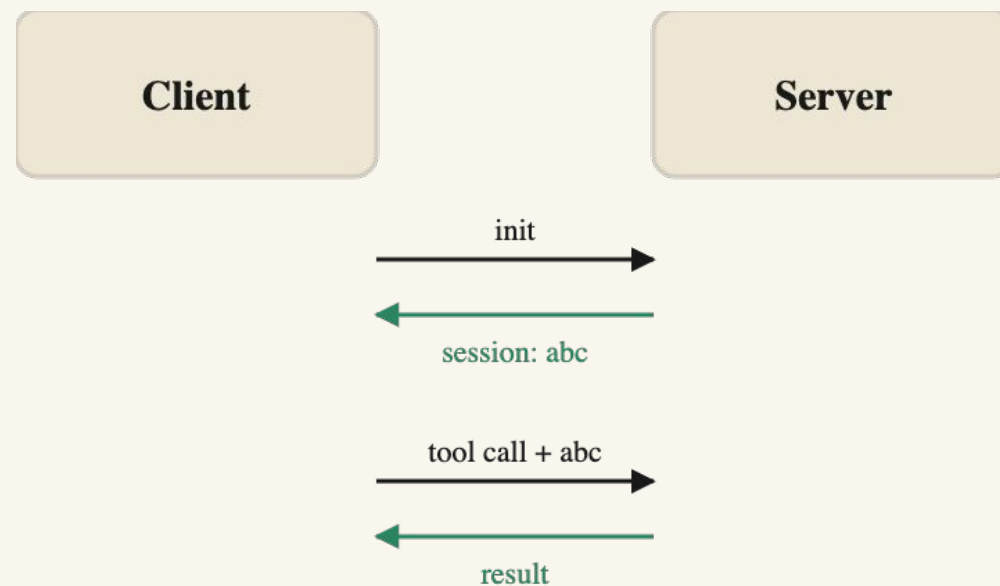
March 2025 — replaces HTTP+SSE

What changed

Requests and responses are individual HTTP requests. No more persistent SSE stream required. Results come back as the response to the same request.

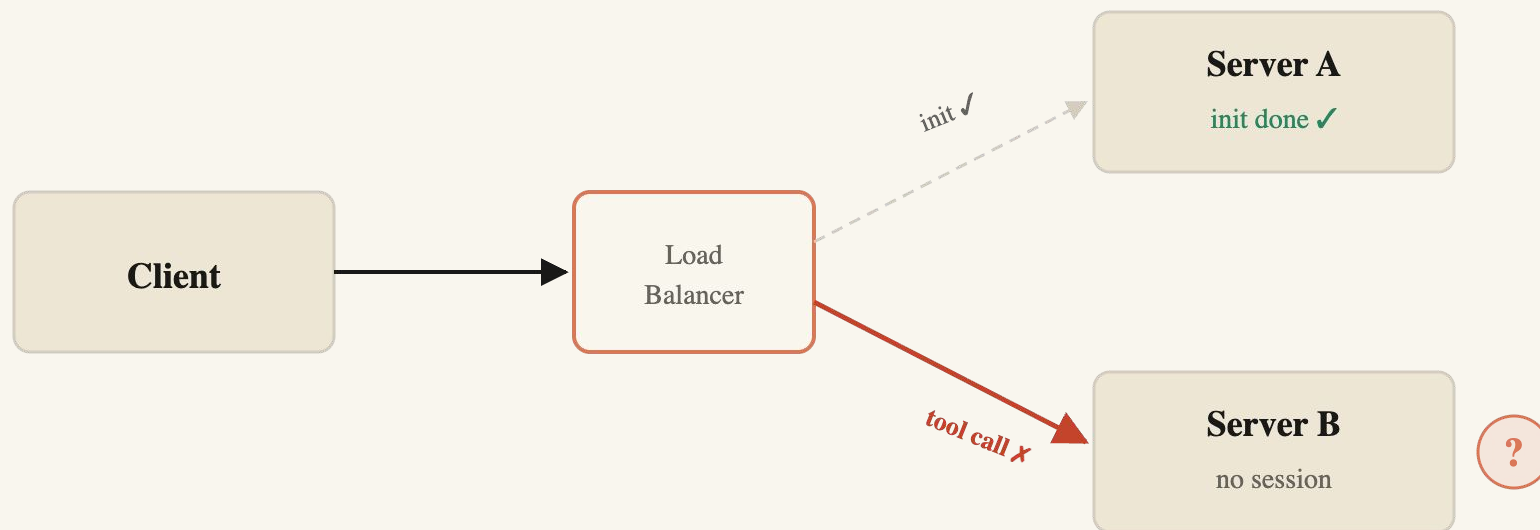
Session IDs

Now an official spec concept (previously just SDK convention).
Server assigns on init, client includes in all subsequent requests.



Each arrow = separate HTTP request

The horizontal scaling problem



Elicitation

```
1  @mcp.tool()
2  async def delete_folder(path: str, ctx: Context) -> str:
3      files = list_files(path)
4
5      if files:
6          result = await ctx.elicit("Folder has files. Delete anyway?")
7          if not result.confirmed:
8              return "Cancelled"
9
10     delete(path)
11     return f"Deleted {path}"
```

How elicitation works



Elicitation breaks at scale



The unofficial fix

Unofficial stateless mode (SDK convention)

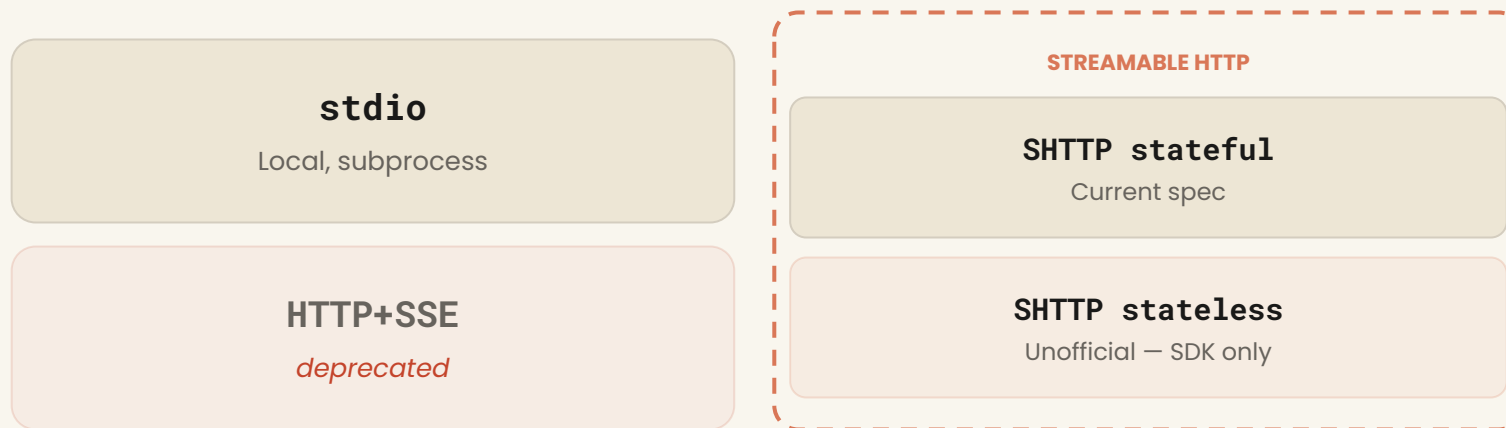
- Initialization is optional — skip the handshake
- No session IDs assigned
- Tool calls work without prior init
- Unlocks simple actions on scaled deployments

But...

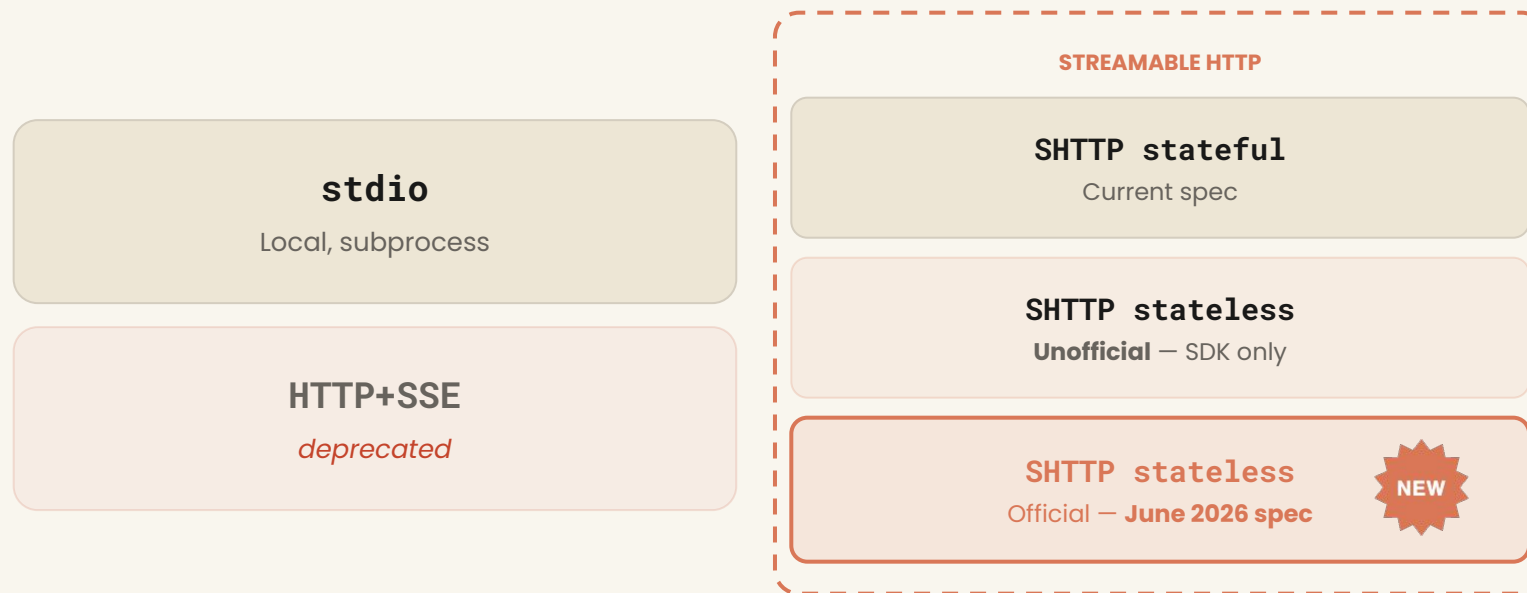
- No session ID → can't correlate elicitation responses
- **Elicitation and sampling won't work**
- Only solves half the problem
- Not in the official spec (*technically not spec compliant*)

That's what the June 2026 spec is trying to solve →

Recap: 3–4 transports



Recap: 3–5 transports



All three Streamable HTTP variants technically share the same name

Multi Round-Trip Requests

BEFORE — SSE stream

- Client → tool call

Server opens SSE stream •

stream held open ∞

Server → "confirm?" •

- **Client → new HTTP request**

✗ Wrong server. Stream on another box.

AFTER — MRTR

ROUND 1

- Client → tool call

IncompleteResult + "confirm?" •

✓ Request ends cleanly

ROUND 2

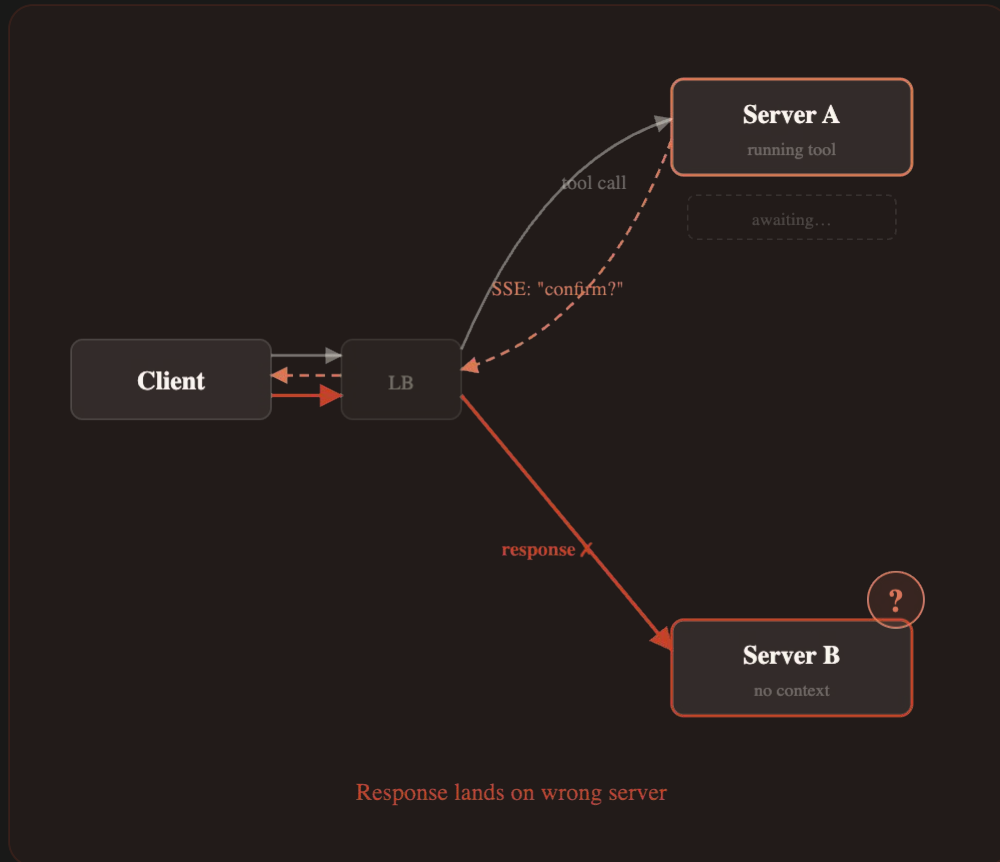
- Client → same call + answer

Server → final result •

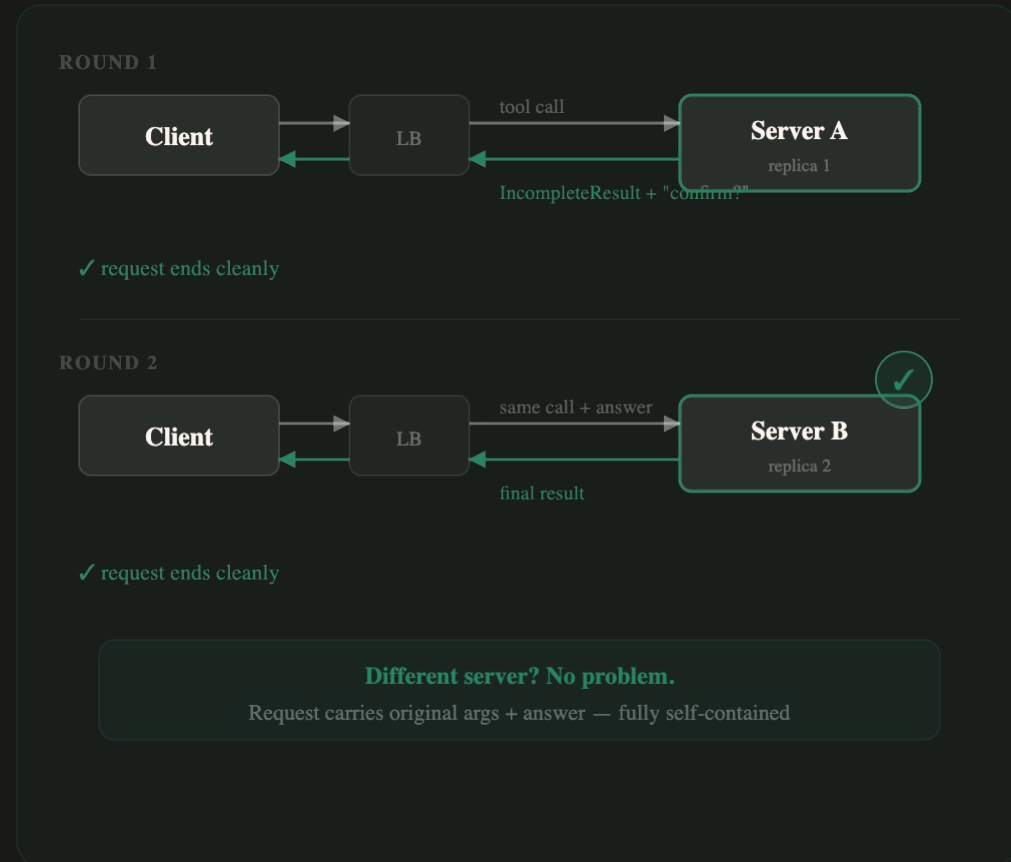
✓ Request ends cleanly

MRTR solves horizontal scaling

BEFORE — SSE elicitation



AFTER — MRTR



Elicitation today

```
1  @mcp.tool()
2  async def delete_folder(path: str, ctx: Context) -> str:
3      files = list_files(path)
4
5      if files:
6          result = await ctx.elicit("Folder has files. Delete anyway?")
7          if not result.confirmed:
8              return "Cancelled"
9
10     delete(path)
11     return f"Deleted {path}"
```

Sequential. The await blocks. If you know Python, you can read this.

With MRTR

```
1  @mcp.tool()
2  def delete_folder(path: str, params) -> Result | IncompleteResult:
3      files = list_files(path)          # ← runs every round
4
5      if files:
6          if "confirm" not in params.responses:
7              return IncompleteResult(
8                  ask="Folder has files. Delete anyway?"
9              )
10
11         if not params.responses["confirm"]:
12             return "Cancelled"
13
14         delete(path)
15         return f"Deleted {path}"
```

Function re-runs from the top each round. Side effects fire every time.

Your code doesn't change - Compatibility layer

unchanged

```
1  @mcp.tool()
2  async def delete_folder(path: str, ctx: Context) -> str:
3      files = list_files(path)
4
5      if files:
6          result = await ctx.elicit("Folder has files. Delete anyway?")
7          if not result.confirmed:
8              return "Cancelled"
9
10     delete(path)
11     return f"Deleted {path}"
```

SDK parks the coroutine. Complexity stays in the dispatch layer.

SDK compatibility layer

WHAT YOU WRITE

```


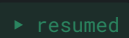
1  @mcp.tool()
2  async def delete_folder(path: str, ctx: Context) -> str:
3      files = list_files(path)
4
5      if files:
6          result = await ctx.elicit("Folder has files. Delete anyway?")
7          if not result.confirmed:
8              return "Cancelled"
9
10     delete(path)
11     return f"Deleted {path}"

```

Familiar sequential code — no re-entrancy

under
the hood

WHAT THE SDK DOES

- Function runs**
list_files() executes normally
- Hits await ctx.elicit()**
SDK in memory **parks coroutine**  frozen
- Sends IncompleteResult**
HTTP request ends cleanly → client gets elicitation
- Client re-calls with answer**
same tool call + elicitation response
- SDK coroutine resumes** 
Function continues from the line with the result `await`

Other things we have changed

Renamed to avoid confusion with external FastMCP

FastMCP (name collision)

✓ **MCP**Server

Low-level: decorators → explicit constructor args

@server.call_tool()

✓ **Server**(call_tool=fn)

Types cleanup across low-level interface

Inconsistent/untyped


✓ **Fully typed handlers**

Other things we have changed

Renamed to avoid confusion with external FastMCP

FastMCP (name collision)

✓ MCPServer ✨



Low-level: decorators → explicit constructor args

@server.call_tool()

✓ Server(call_tool=fn)

Types cleanup across low-level interface

Inconsistent/untyped

✓ Fully typed handlers

Pluggable transports

Top layer

MCP Types

CallToolRequest, CallToolResult...

Middle layer

Message format

JSON-RPC (default) · Protobuf · ...

Bottom layer

Transport

stdio · SHTTP · gRPC · WebSocket...

"Pluggable transports" usually means swapping both bottom layers.

gRPC with protobuf over JSON-RPC makes no sense — you'd smuggle JSONRPC as a string and lose protobuf's whole value proposition.

Solution: Dispatcher pattern — swap both layers

The Dispatcher pattern


BEFORE

BaseSession

MCP semantics
 initialize(), call_tool(), list_tools()... progress tokens, cancellation, validation
 19 methods

Wire protocol tangled
 JSON-RPC wrap, ID correlation, receive loop, stream management

Want gRPC? Reimplement the whole session.



AFTER

python-sdk PR #2320

BaseSession

MCP semantics only
 initialize(), call_tool(), list_tools()...
19 methods – unchanged

↓

Dispatcher 5 methods

JSON-RPC (default) gRPC Protobuf ...

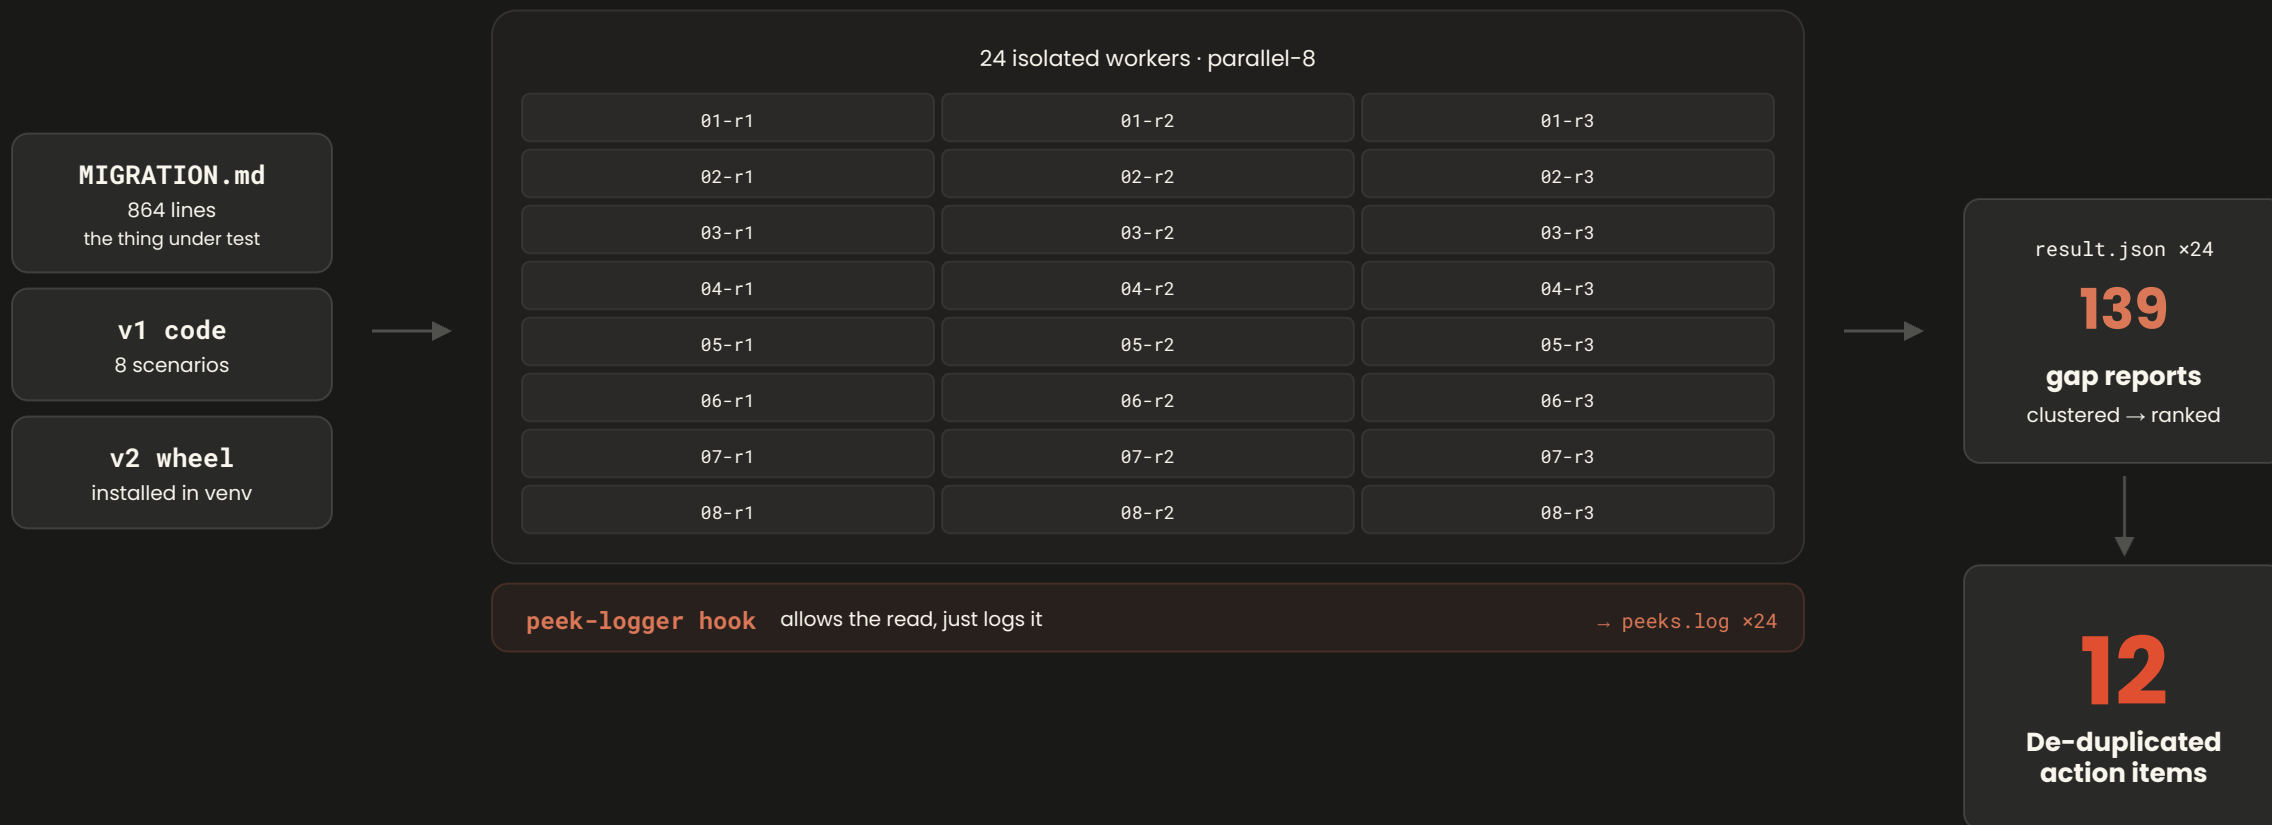
send_request · send_notification · send_response set_handlers · run

Implement 5 methods → all 19 MCP methods work for free

VALIDATION

Migration documentation eval framework

The eval fleet



The 24 runs

20/24 passed 5/24 peeked 139 individual reports

	r1	r2	r3
01 - examples-fastmcp	⚠️	⚠️	⚠️
02 - examples-snippets	✓	✓	✓
03 - examples-servers	✓	✓	✓
04 - examples-clients	✓	✓	✓
05 - lowlevel-heavy	✓	✓	✓
06 - fastmcp-power	✗	✗	✗
07 - client-heavy	✓	✓	✓
08 - type-edges	✓	⚠️	✓

■ passed, no peek
 ■ passed, peeked
 ■ verify failed

Row 01 — all 3 peeked
 Consistent doc hole: two undocumented removals

Row 06 — all 3 failed verify
 Migrations were correct — verify.sh used a v1-only attribute. All 3 workers diagnosed it independently.

“DOC GAP: migration guide does NOT call out the camelCase → snake_case field naming change. I only figured it out from the runtime error's suggestion.”

– Claude worker 07-r2, notes.md

21/24 runs · 7 of 8 scenarios

CLOSE

What's next – Python SDK

BETA — Q2 2026

ROADMAP

Dispatcher extraction

Draft PR

Auth rework

In design

Stateless by default

Pending June spec

MRTR

Pending June spec

Docs

Ongoing

[github.com/modelcontextprotocol/
python-sdk](https://github.com/modelcontextprotocol/python-sdk)

Maintainers

@maxisbey

@Kludex

@felixweinberger

@pcarleton

Full stable release planned for the next major
MCP spec release

Feedback? Please file an issue on GitHub

CLOSE

TypeScript SDK

"Any runtime. Any validator. Less boilerplate."

Runs everywhere

Node, Bun, Deno, Cloudflare Workers — runtime auto-detected, zero config

Bring your own stack

Zod, Valibot, Arktype, or raw JSON Schema. Hono, Next.js, SvelteKit natively.

Less boilerplate

Flat imports, fully typed handlers, ctx helpers — no deep paths or schema imports

Enterprise-ready

SSO built in, redesigned errors, protocol extensions, aligned with next spec

Migration on autopilot

Point your coding agent at the bundled migration skill — upgrading is a prompt

v2 ALPHA — OUT NOW

[github.com/modelcontextprotocol/
typescript-sdk](https://github.com/modelcontextprotocol/typescript-sdk)

Maintainers

[@felixweinberger](#) [@mattzcarey](#)

[@kkonstantinov](#) [@pcarleton](#)

Stable release planned for the next major MCP spec release

Feedback? Please file an issue on GitHub or upvote existing ones

Thanks

github.com/modelcontextprotocol/python-sdk

