

THE MODEL CONTEXT PROTOCOL

Standardizing the Agentic Nervous System

Chris Siller | Engineering Kiosk Meetup | Innsbruck, 15.01.2026

Tonight's Agenda



- 1 Evolution from Chatbots to Agents and its Challenges
- 2 Introducing MCP
- 3 A little MCP Role Play
- 4 Why You Should Care about MCP
- 5 Resources

Beyond the Chatbot: The Rise of Agency

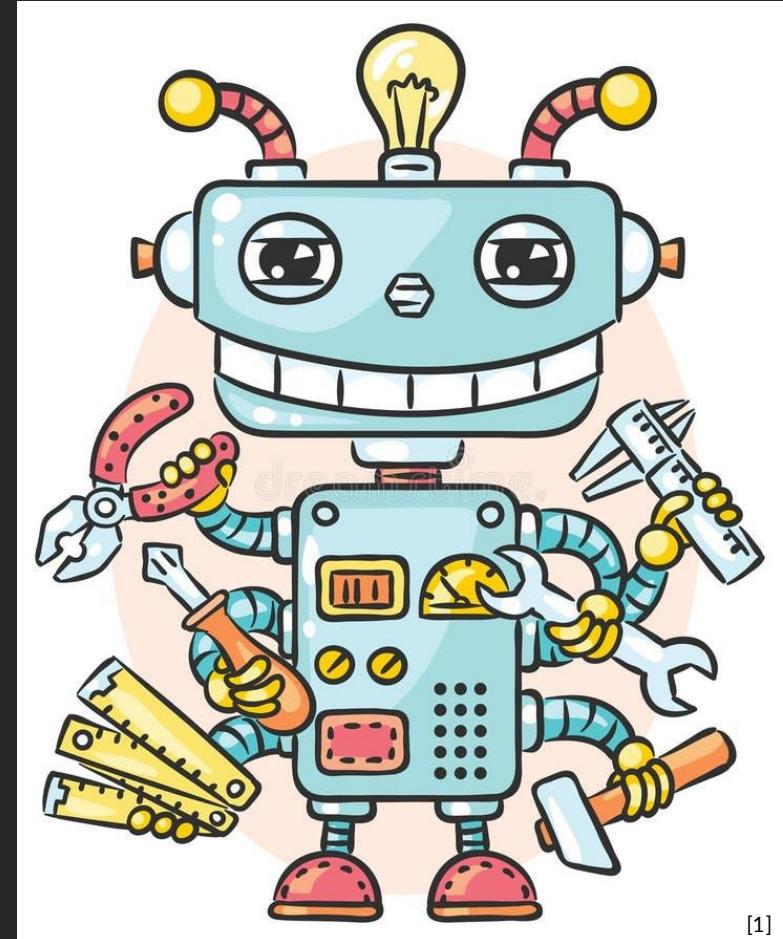
From Chatbots to Agents:

We have transitioned from passive Chatbots to Autonomous Agents. The fundamental difference is simple: Agents don't just talk; they ACT.

The Need for Tools:

Agents solve complex problems by accessing specific data and environments, but they cannot do this alone. To bridge the gap between reasoning and reality, they require Tools—the "hands" that allow them to query databases, execute code, etc..

But how do we connect the tools to the robots?



The "Tool Integration Hellscape"

N × M

The Complexity Trap

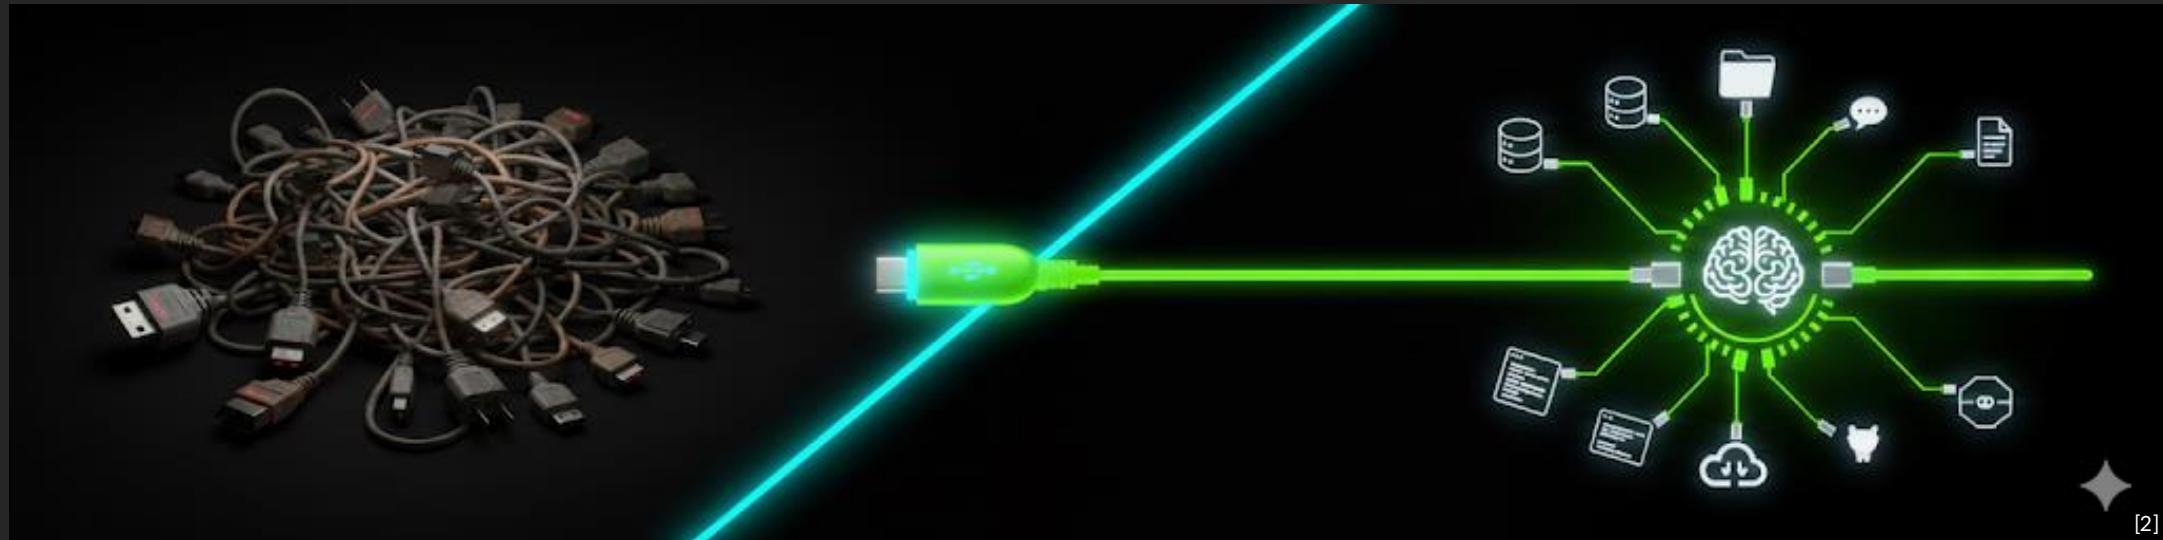
Non-Standardized Tool Integration faces a variety of Challenges:

- **Exponential Integration Cost**
Every new tool must be integrated separately with every model → **N tools × M models**.
- **Brittle Architectures**
Custom wrappers multiply quickly, creating hard-to-maintain, tightly coupled “spaghetti” systems.
- Each integration reimplements authentication, request/response parsing, retries, and error handling.
- **Zero Portability**
Tools are bound to specific models and runtimes, blocking reuse across the ecosystem.
- **Scaling Becomes Unsustainable**
Adding a single tool or model increases system-wide complexity, not capability.

MCP: A USB-C Moment for AI?

We are currently in the "proprietary charger" era of AI. Before USB-C, every device had a unique plug; today, every agent has a unique connector.

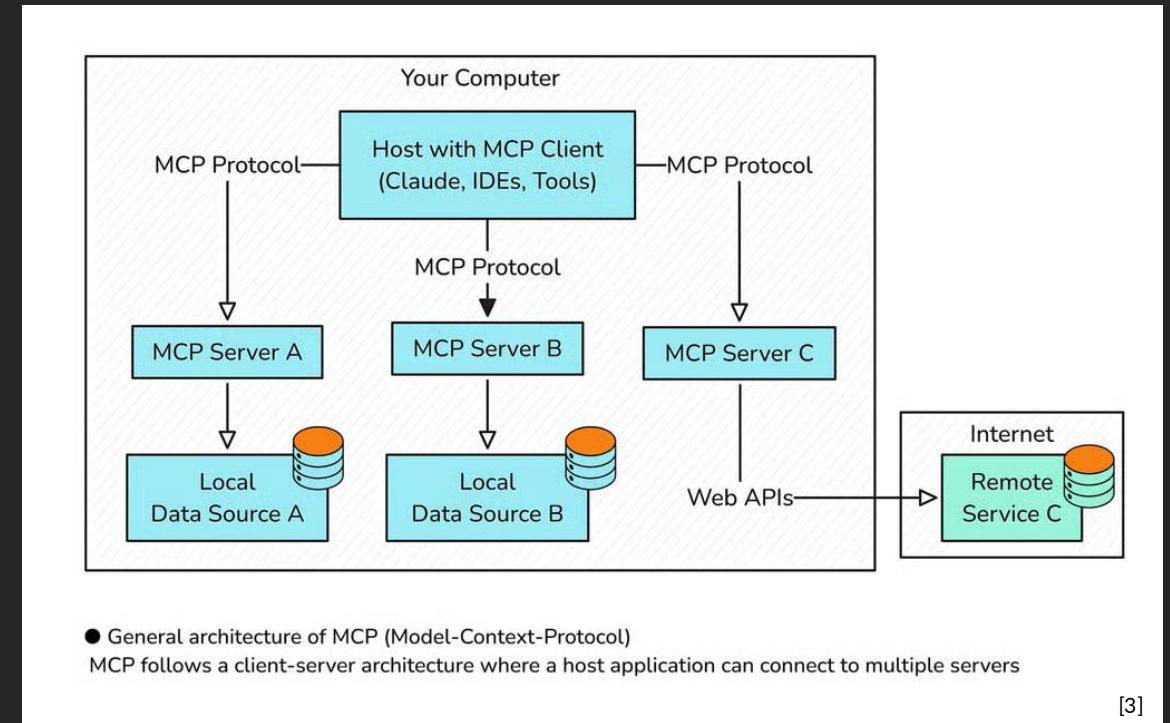
MCP serves as the universal port. It decouples the reasoning engine from the data source, allowing any compliant server to provide context to any compliant host instantly.



MCP: Standardizing the AI Handshake

The Model Context Protocol (MCP) establishes a universal interface between AI applications and data sources.

- + Decoupled Architecture: Separates reasoning (Host) from implementation (Server).
- + Stateful Connections: Maintains 1:1 state between clients and servers
- + Universal Scaling: Enables N models to talk to M tools via one standard.



The Three Pillars of MCP



MCP Host

The orchestrator. It is the top-level AI application—like Claude Desktop or an IDE—that manages the user interface and the LLM reasoning loop. The Host does not talk to tools directly; it coordinates one or more Clients to bridge the gap between the model's intent and the system's data.



MCP Client

The dedicated connector. It lives within the Host and maintains a 1:1 relationship with a specific Server. The Client's job is to translate the high-level needs of the Host into protocol-compliant JSON-RPC requests, handling the negotiation of capabilities and securing the transport tunnel.



MCP Server

The source of truth. It is a lightweight program that exposes specific capabilities—Tools, Resources, or Prompts—to the Client. Servers abstract away the complexity of the underlying data (like a database or API), providing a standardized interface that any MCP-compatible Host can understand.

Server-Side Primitives

Server-side primitives are the three standardized "building blocks" that an MCP server offers to an AI model to give it agency and context



Tools (The Hands): Executable functions. They allow the AI to *act* on the world.

Example: create_jira_ticket, execute_python_code, query_database.

Resources (The Knowledge): Read-only data. They provide the AI with specific *context*.

Example: app_logs, database_schema, documentation_file.

Prompts (The Instructions): Reusable templates. They tell the AI *how* to handle a specific scenario.

Example: senior_dev_code_review_template, customer_support_tone_guidelines.

In short: **Tools** let the model act, **Resources** let it see data, and **Prompts** tell it how to think.

Client-Side Capabilities

Closing the Loop:

MCP isn't just one-way.
Clients expose primitives
that allow Servers to "reach
back" into the Host
environment for
intelligence and interaction.



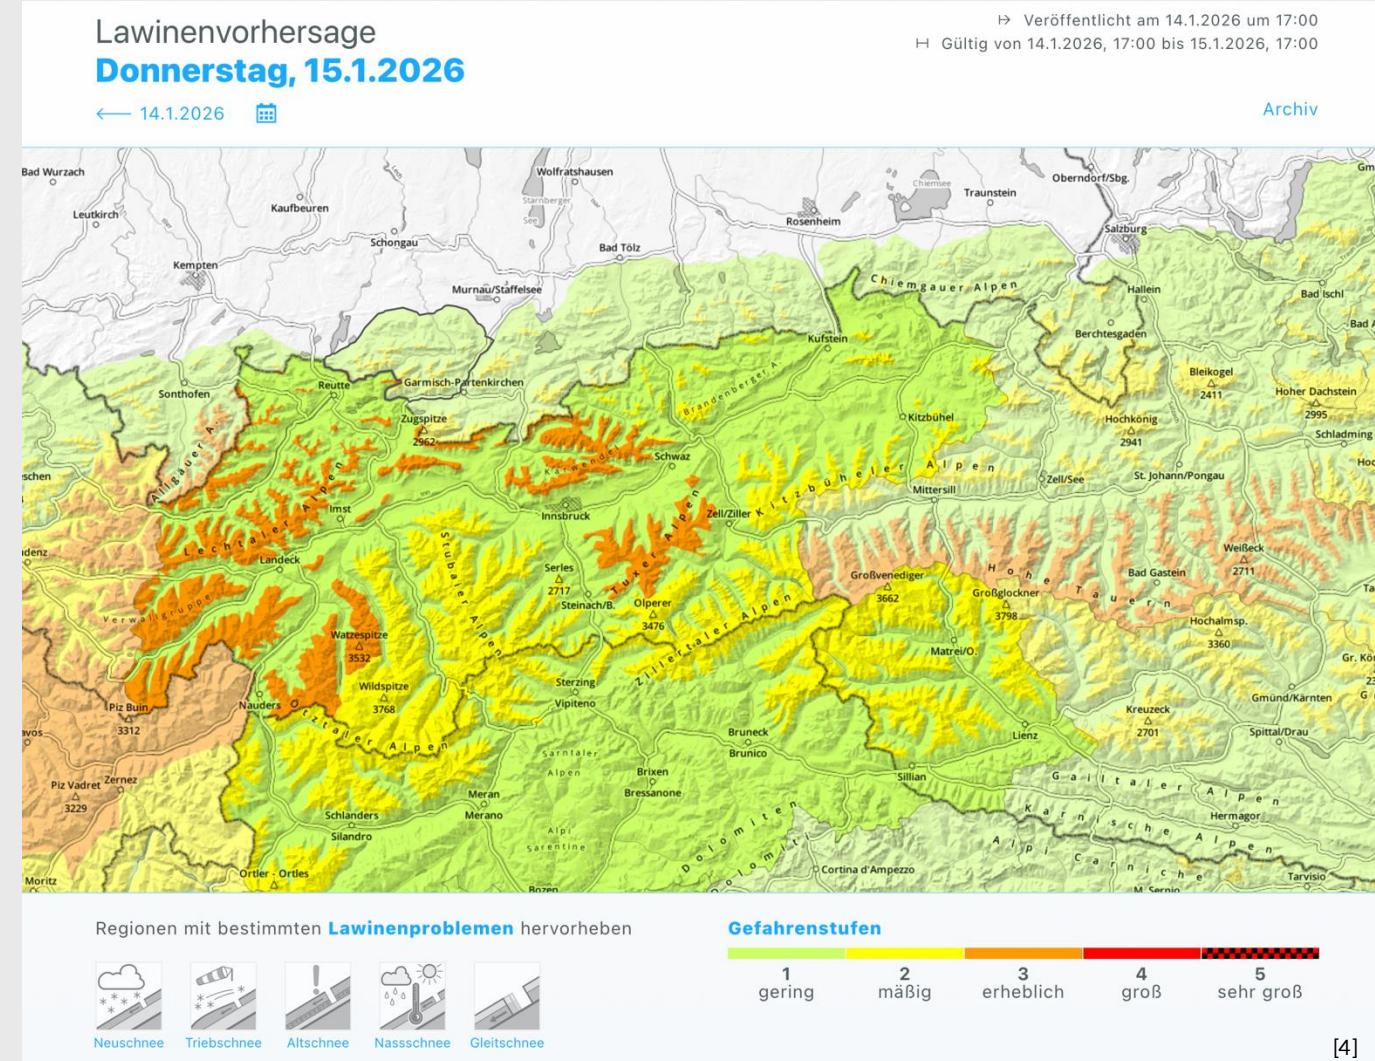
Sampling: Servers can request LLM completions **via the host**, using the host's model, policies, and context.

Elicitation: Servers can ask the host to obtain user input or confirmation when required.

Logging: Servers can emit standardized telemetry, traces, and debug information back to the host.

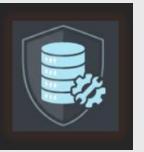
MCP Roleplay

Asking AI about avalanche conditions isn't casual—it's high-stakes.
A hallucinated safety report isn't a typo; it's a fatal error.
We need authoritative "eyes".
The model needs to stop guessing and start knowing.



What if we had a
Lawinenwarndienst MCP Server?

MCP Roleplay



(user)

(host)

(client)

(server)

Safe to Skitour On Friday ?
(instead of working)

What tools you got?

Where do you wanna go?

Is_route_advisable
w/ these parameters: ...

IBK Nordkette, N Aspect

call is_route_advisable
Params: Loc, Elev, Aspect

advisable: false, dangerlevel: 3,
explanation: ...

not safe to skitour in
Nordkette on Friday

```
{  
  "jsonrpc": "2.0",  
  "id": "1",  
  "result": {  
    "tools": [  
      {  
        "name": "is_route_advisable",  
        "description": "Evaluates mountain route safety based on terrain and time.",  
        "inputSchema": {  
          "type": "object",  
          "properties": {  
            "point": {  
              "type": "object",  
              "properties": {  
                "lat": {  
                  "type": "number"  
                },  
                "lng": {  
                  "type": "number"  
                }  
              }  
            },  
            "required": [  
              "lat",  
              "lng"  
            ]  
          },  
          "elevation": {  
            "type": "integer",  
            "description": "Meters above sea level"  
          },  
          "aspect": {  
            "type": "string",  
            "enum": ["N", "NE", "E", "SE", "S", "SW", "W", "NW"]  
          },  
          "epoch_time": {  
            "type": "integer",  
            "description": "Unix timestamp for temporal conditions"  
          }  
        },  
        "required": [  
          "point",  
          "elevation"  
        ]  
      }  
    ]  
  }  
}
```

```
{  
  "jsonrpc": "2.0",  
  "id": "2",  
  "result": {  
    "content": [  
      {  
        "type": "text",  
        "text": {  
          "advisable": false,  
          "danger_level": 3,  
          "aspect": "N",  
          "trend": "increasing",  
          "primary_hazard": "wind_slab",  
          "critical_threshold": "2200m",  
          "valid_until": 1737019858,  
          "advisory": "NOT RECOMMENDED. Recent snowfall combined with strong southerly winds has created significant wind slabs on north-facing slopes. High trigger probability by single winter sports participants."  
        }  
      },  
      {  
        "isError": false  
      }  
    ]  
  }  
}
```

```
{  
  "jsonrpc": "2.0",  
  "id": "1",  
  "method": "tools/list",  
  "params": {}  
}
```

```
{  
  "jsonrpc": "2.0",  
  "id": "2",  
  "method": "tools/call",  
  "params": {  
    "name": "is_route_advisable",  
    "arguments": {  
      "point": {  
        "lat": 47.3039,  
        "lng": 11.3833  
      },  
      "elevation": 2300,  
      "aspect": "N",  
      "epoch_time": 1736947200  
    }  
  }  
}
```

Why should You care about MCP?*



[5]

*if you are not the one who has to build
the wrappers?

Why you should care about MCP (The Sovereignty Gap)

The Problem:

The "Walled Garden" Trap

Big Tech wants to own the "Brain" and the "Hands." If you use their proprietary connectors, you are locked into their ecosystem forever. If they raise prices or the model gets "dumber," you are stuck.



A Solution:

Decoupling with MCP

- + **Exchange the Brain:** Don't like Gemini anymore? Plug your tools into Claude or a local Llama instance. The "Hands" (your data/tools) stay the same.
- + **Build Independently:** Your company data stays in your control, or even on your premise. You provide the MCP Server; the AI model just visits it.
- + **Vendor Insurance:** An open standard means you aren't at the mercy of a single API provider's roadmap.

So What Now ?

- + Future-Proof Your Tech
- + Take Ownership of your Data Ecosystem
- + Maintain the freedom to switch “brains”
without rebuilding your operational layer
- + Take your cool ideas and build cool Agent
Systems

Resources and Links

Resources:

<https://www.anthropic.com/news/model-context-protocol>
<https://modelcontextprotocol.io/>

Images:

[1]Our Robots now have tools : <https://www.dreamstime.com/stock-illustration-cute-robot-six-hands-holding-different-working-tools-cartoon-image59923784>
[2]Nano Banana via <https://gemini.google.com/>
[3]<https://sspark.genspark.ai/cfimages?u1=xA%2FxnxDwlqlNql6mlBhRyMYxX0d8FV3wUzFwlSmzL%2Bgdq0NwRg%2BbbU7kQ16%2FE%2BPYA7hfVr%2F5nByCatu0XKclyo%2BeQCW7wyZrl%2F307rmRBVSi3P2JVPG%2F9omzLvDYCt69ygn9YTjp%2F1798ECOpW1Dvcr2gKjh1WKMpLgMHYYb3Sz5v63XMyu5iSeESI649y53pfkDe%2By6lWAfsBxWTU27joIKhgYrcVKDxD%2BgsppA%3D&u2=hjOC7m8wgHosCRal&width=1024>
[4]<https://lawinen.report/bulletin/latest>
[5]<https://www.google.com/url?sa=t&source=web&rct=j&url=https%3A%2F%2Ftenor.com%2Fsearch%2Fbut-why-gifs&ved=0CBUQjRxqFwoTCNCuhMD5jZIDFQAAAAAdAAAAABAH&opi=89978449>