



MCP
Dev Summit
Mumbai

The 5 Wrong Reasons to Build an MCP Server (And What to Do Instead)

Daniel Oh - Developer Advocate at IBM (@danieloh30)

Java Champion | CNCF Ambassador | CNCF TAG DevEx Co-chair | Microsoft MVP

AI Landscape

THE AGENTIC BOOM



The explosion of the "Agentic IDE" ecosystem.

5,000+ community MCP servers listed on central registries.

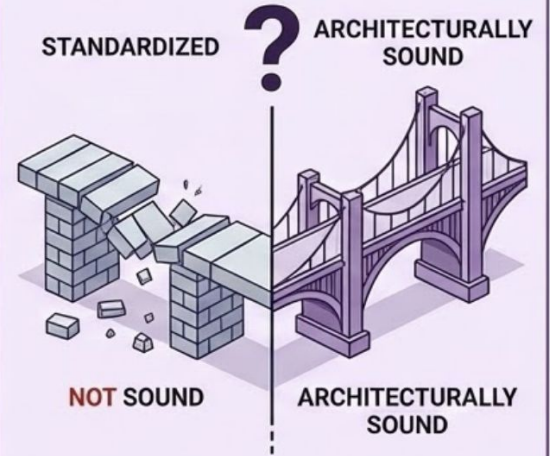
CORPORATE REALITY & RISKS



The corporate pressure to "AI-enable" every internal API.

The risk of the "Everything looks like a nail" syndrome.

STANDARDIZATION VS. SOUNDNESS

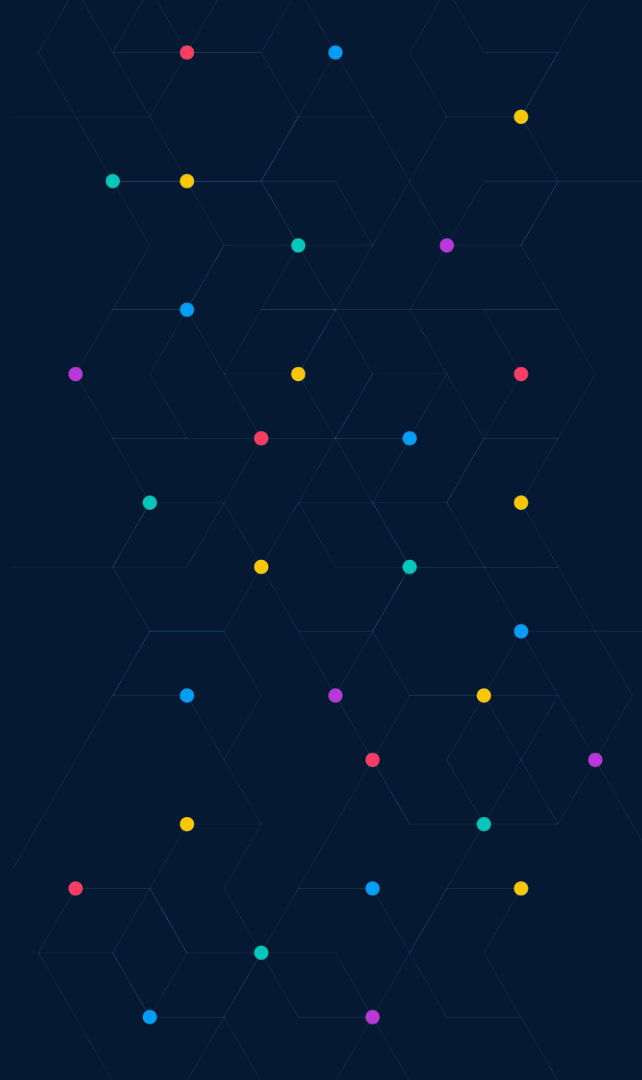


Why "Standardized" does not always mean "Architecturally Sound."

No is the Right Answer



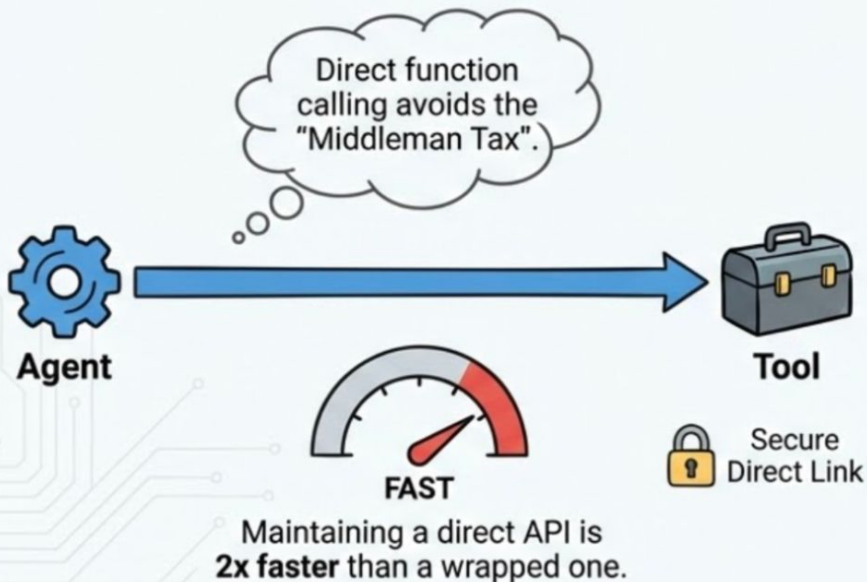
MCP
Dev Summit
Mumbai



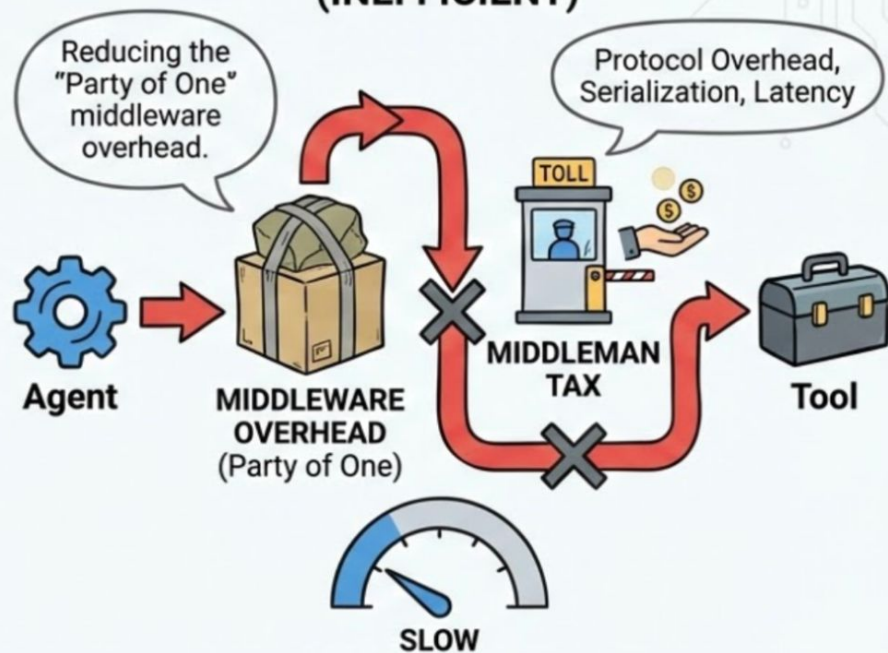
Wrong Reason #1: Single Consumer Trap

One Agent + One Tool = Zero reason for a protocol!

DIRECT FUNCTION CALLING (EFFICIENT)

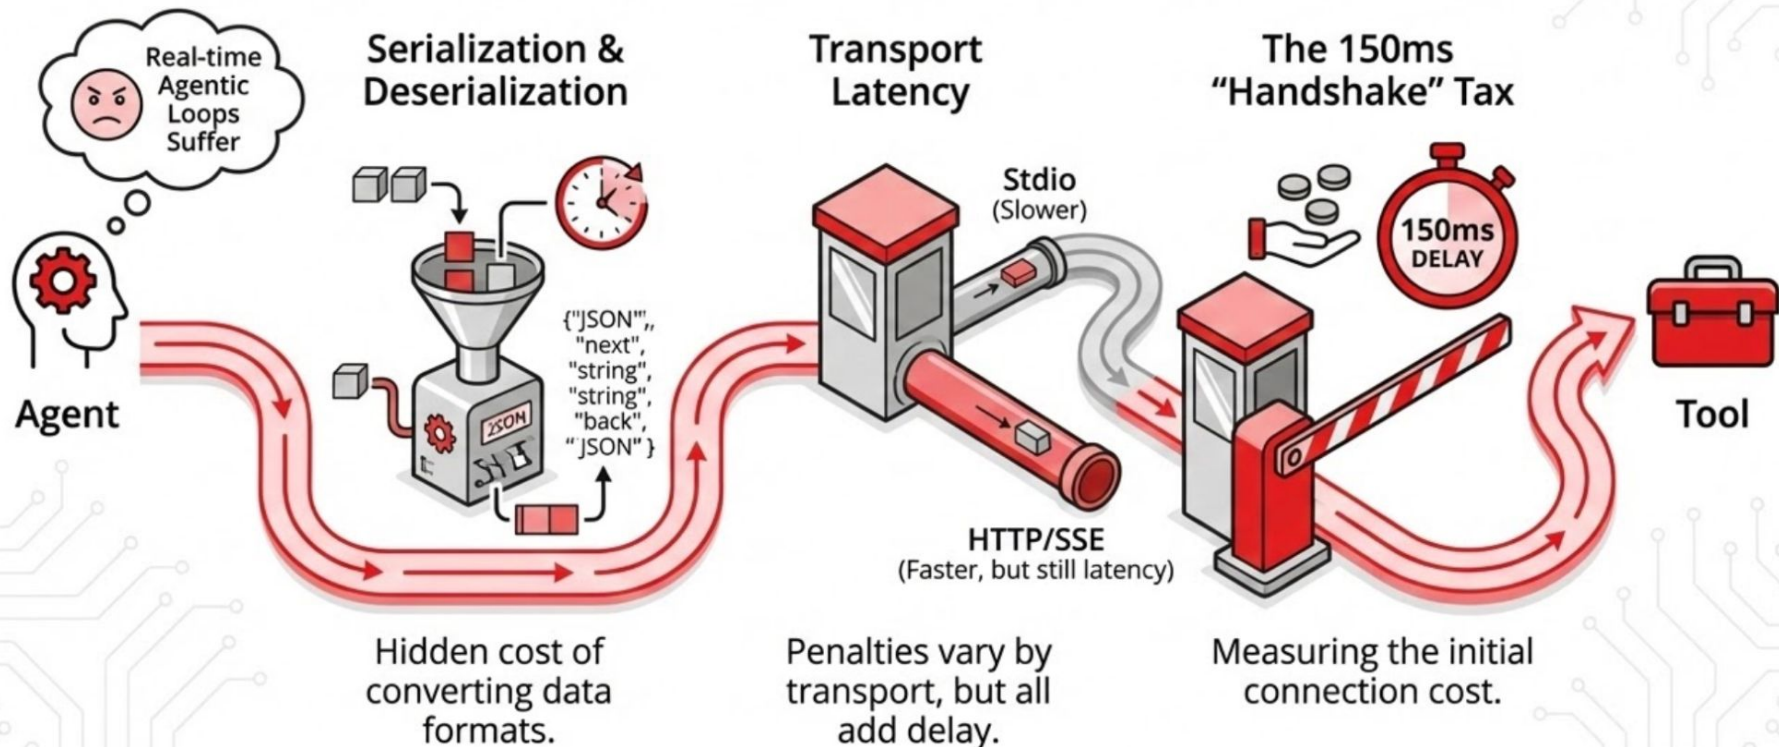


PROTOCOL WRAPPER (INEFFICIENT)



Wrong Reason #2: Performance Myth

JSON-RPC is a convenience, not a performance optimization!!



Wrong Reason #3: Security Fallacy

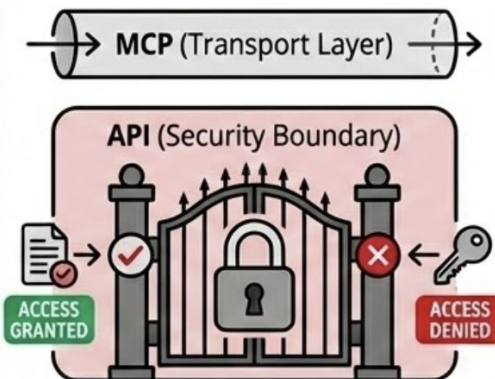
MCP is a transport layer, not a security boundary

The Danger of Blind Trust



Trusting the LLM client blindly.
Increasing the 'Prompt Injection' surface area via tools.

Authorization Belongs at the API Level



MCP is just a **transport mechanism**.
API level is the correct place for authorization.

The "Exposed Internal Tool" Nightmare



Exposed Internal Tool nightmare scenario.

Wrong Reason #4: Documentation Replacement

MCP won't fix a poorly designed, undocumented API

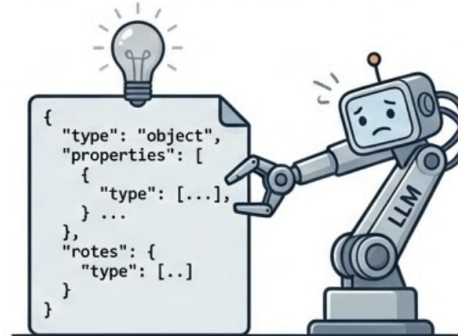
- LLMs still require high-quality semantic context to work.



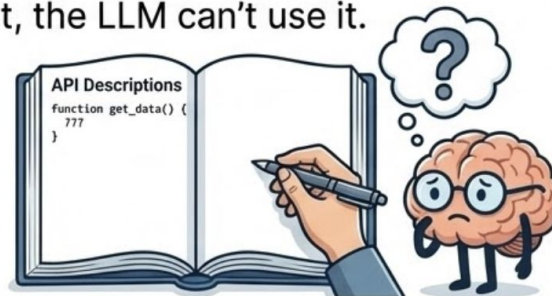
- “Garbage In, Garbage Out”: Standardizing a mess just makes it a standard mess.



- The JSON Schema is not “Auto-Intelligence.”



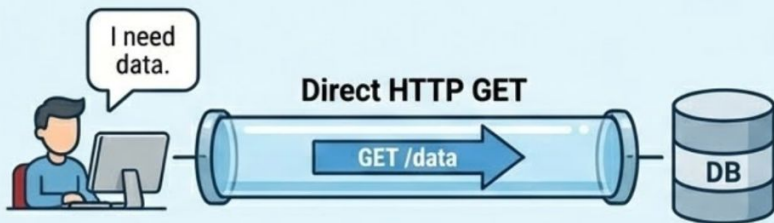
- Descriptions are code: If you can't describe it, the LLM can't use it.



Wrong Reason #5: Simple CRUD Operations

HTTP verbs are already a "Standard Protocol"!

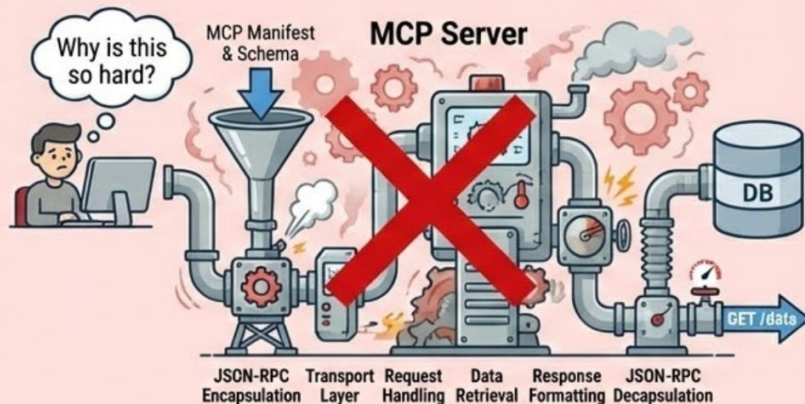
The Simple Way: Standard Protocol



- ✓ HTTP verbs (GET, POST, etc.) are the standard. cURL or shell scripts are sufficient "tools". Direct, efficient, and easy to maintain.

```
$ curl https://api.example.com/data
```

The Wrong Way: MCP Overkill



- ✗ Wrapping a single endpoint in a server manifest is unnecessary. High maintenance burden for a single endpoint. Creates "Interface Bloat" for basic tasks.

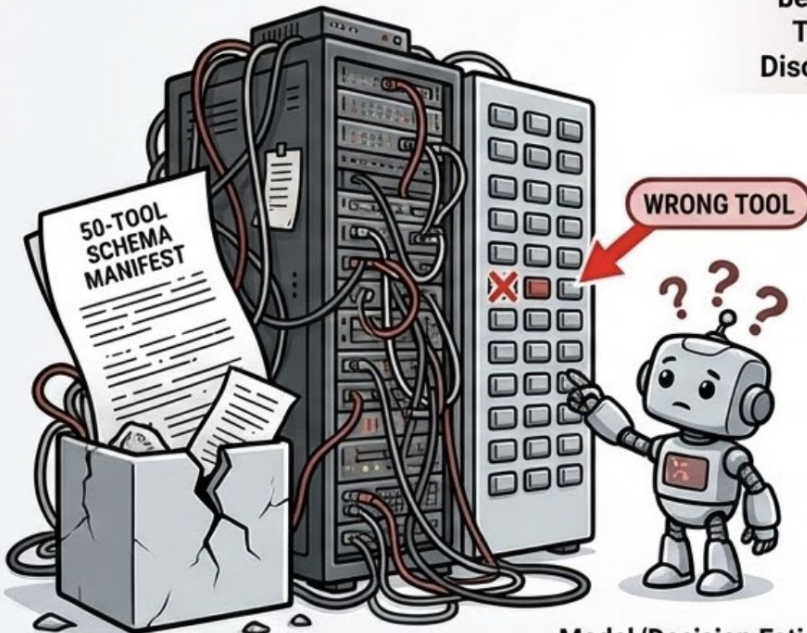
Architecture & Battle Scars



MCP
Dev Summit
Mumbai

Case Study: Mega Server Mistake

GOD-SERVER (MEGA SERVER MISTAKE)

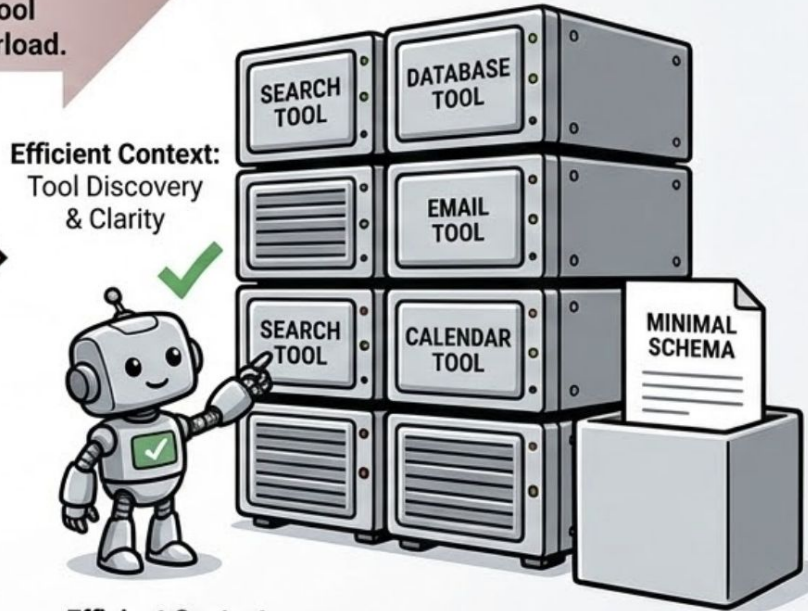


CONTEXT WINDOW BROKEN
(TOKEN BLOAT)

Model 'Decision Fatigue':
Picking the wrong tool
30% of the time

Why 'Micro-Servers' are usually
better than 'God-Servers'.
The Importance of Tool
Discovery vs. Tool Overload.

MICRO-SERVERS (BETTER APPROACH)



Efficient Context:
Tool Discovery
& Clarity

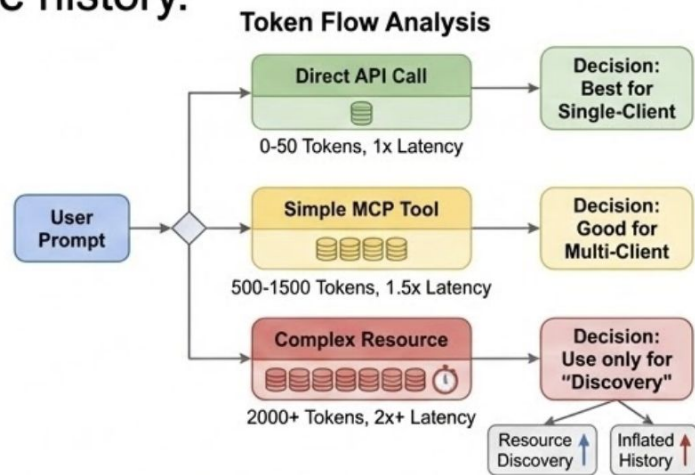
Efficient Context:
Tool Discovery & Clarity

EFFICIENT
CONTEXT WINDOW

Token Efficiency Table

- Quantitative comparison of Tool schemas vs. Function signatures.
- Impact of “Instructional” prompts within the MCP server.
- The price of multi-step “Agentic Loops.”
- How “Resource Discovery” inflates the message history.
- Strategies for “Token-Lean” MCP design.

Token Overhead & Latency Comparison			
Interaction Type	Est. Token Overhead	Latency Factor	Decision
Direct API Call	0 - 50	1x	Best for Single-Client
Simple MCP Tool	500 - 1,500	1.5x	Good for Multi-Client
Complex Resource	2,000+	2x+	Use only for “Discovery”



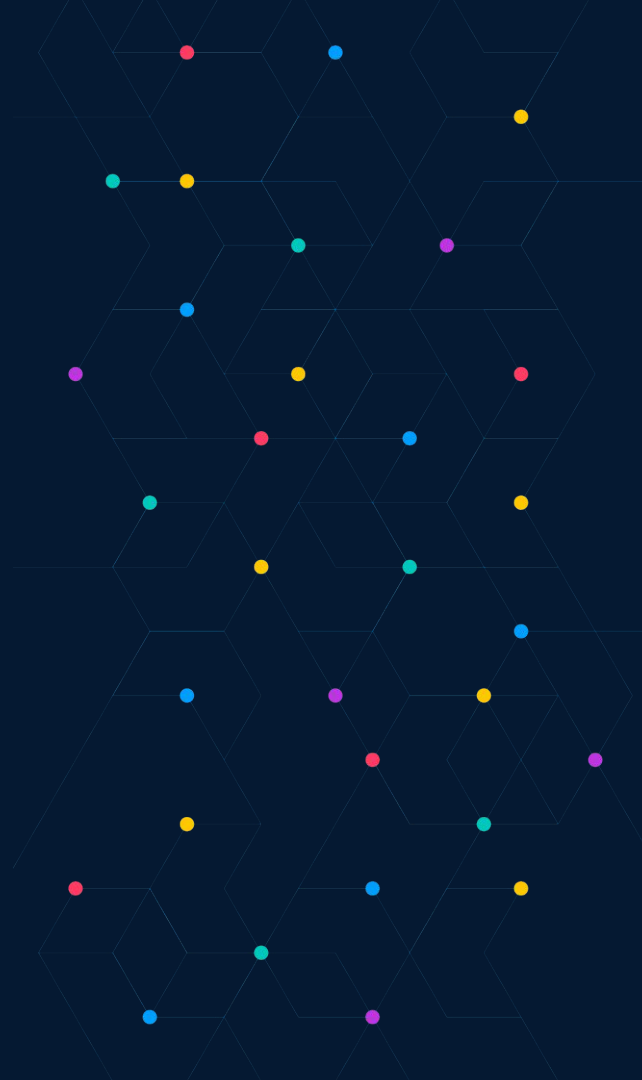
Token-Lean Strategies

- 🔧 **Optimize Tool Schemas** (Minimize Tokens)
- 📄 **Avoid Instructional Prompts** (Reduce Context)
- 🔄 **Limit Agentic Loops** (Control Steps)
- 🗑️ **Use Resource Discovery Sparingly** (Manage History)

Worth It Framework



MCP
Dev Summit
Mumbai



5-Point "Painless" Checklist



Distribution:
2+ clients?



Discovery:
Hierarchical
data?



Interop:
Cross-language/
Cross-team?



Context:
High-density
info?



Stability:
Is the API
schema settled?



Painless
Integration

Live Demo: Over-Engineered Helpdesk

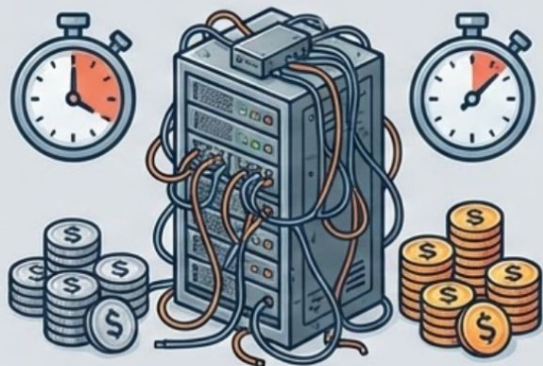


MCP
Dev Summit
Mumbai



Scenario: The "Over-Engineered" Helpdesk Server.

Over-Engineered Server



Latency & Token Use:
High / Inefficient.



The "Aha!" Moment:
When MCP finally
pays for itself.

Clean SDK



Latency & Token Use:
Low / Optimized.

Tech Stack



MCP SDK



Quarkus



LangChain4j



OpenAI

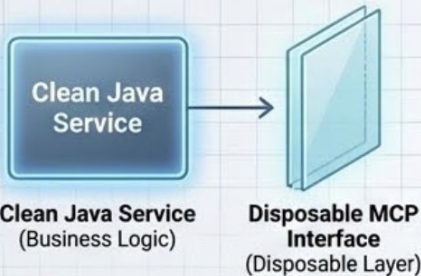


Expectations

Setting expectations for
"Live Code" failures.

Conclusion & The Pragmatic Path Forward

ARCHITECT FOR AGILITY

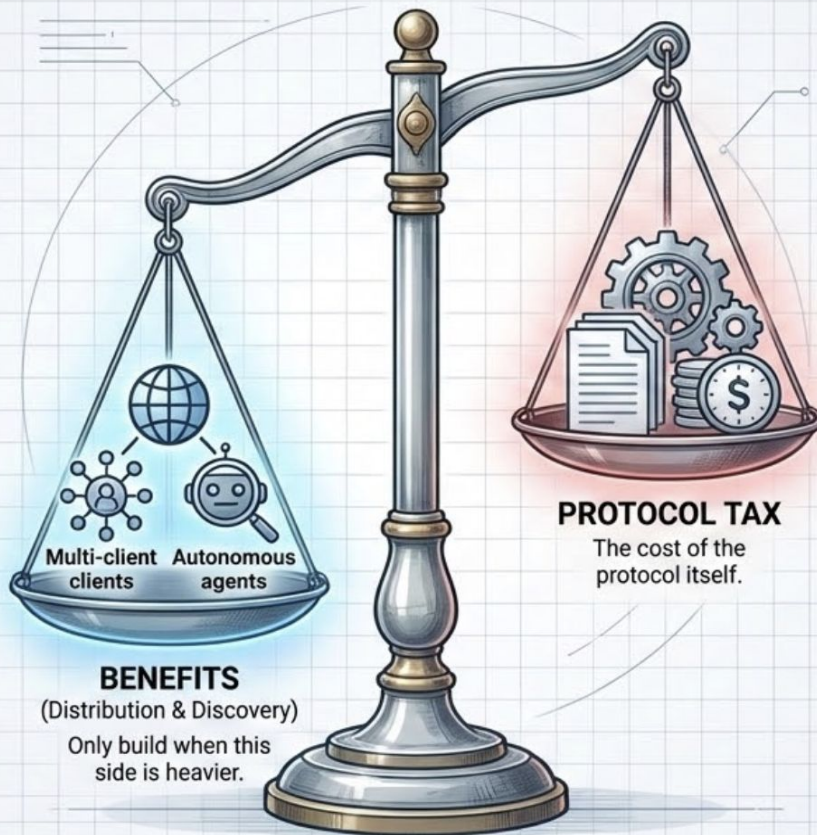


Keep business logic in clean, protocol-agnostic Java services; treat MCP as a "disposable" interface layer.

FOCUS ON THE "WHY"



Does the LLM actually need a dynamic tool registry, or just a well-documented API endpoint?

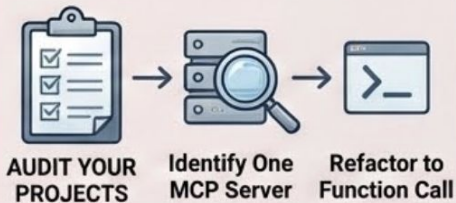


STANDARDIZE WITH INTENT



Don't let the "Standard" label fool you. A 10-line SDK is often superior to a 500-line JSON-RPC wrapper.

FINAL CALL TO ACTION

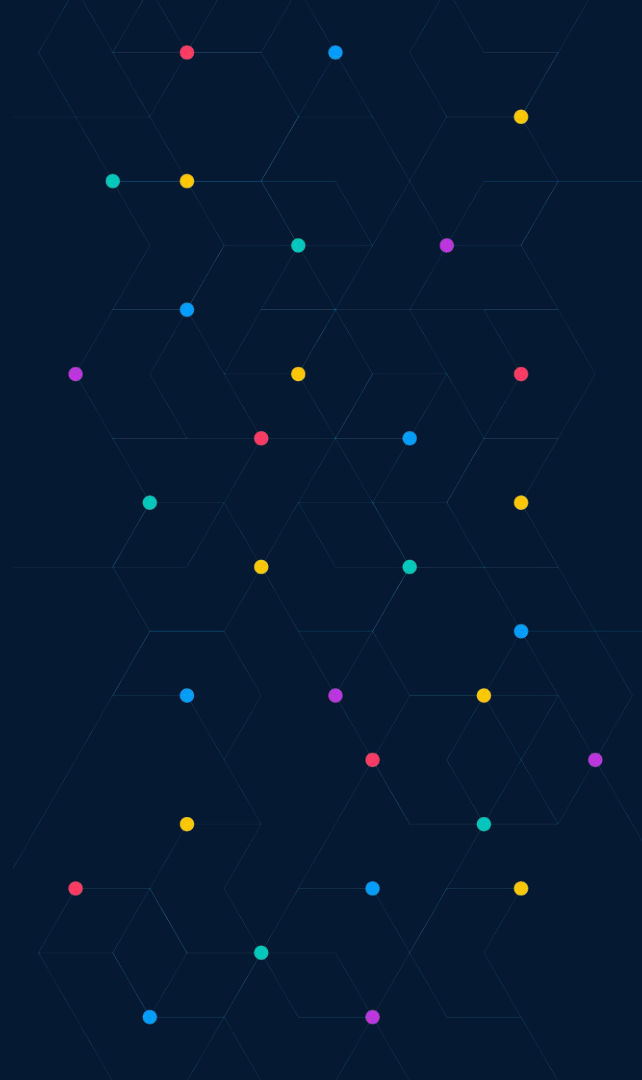


Audit your current "Agentic" projects today—identify one MCP server that should have been a simple function call and refactor it.



MCP
Dev Summit
Mumbai

THANK YOU!
QUESTION?





MCP
Dev Summit
Mumbai

