

SignalWire AI Gateway (SWAIG) and How it Compares to MCP

Introduction

The emergence of the Model Context Protocol (MCP) has sparked excitement across the AI and developer ecosystems. Positioned as a forward-looking standard for enabling large language models (LLMs) to securely access external tools, data, and context, MCP proposes a structured, extensible framework that emphasizes modularity, consent, and safety. It introduces concepts like resources, tools, and prompts, bound together by role-based interactions that could allow AI assistants to perform meaningful actions.

But while MCP defines an ambitious blueprint, SignalWire has already built and operationalized this vision. SWAIG, the SignalWire AI Gateway, is a mature, production-grade protocol and runtime system powering real-world AI agents across telecom channels at scale. From programmable voice workflows to real-time messaging orchestration, SWAIG is not just a concept. It is the infrastructure behind live applications.

This blog explores why SWAIG stands as the realized model of an AI Gateway, showcasing practical features, declarative simplicity, and battle-tested execution environments that MCP has yet to reach. MCP paints a compelling vision of how AI integrations might work; SWAIG demonstrates how they already do: cleanly, scalably, and securely.

The Core Philosophy

Both SWAIG and MCP share a foundational goal: to move AI beyond passive response generation and toward safe, structured interaction with the world. An AI Gateway Protocol must allow LLMs to execute actions with contextual awareness and operational control.

MCP introduces this through a layered, declarative model composed of resources (data), tools (functions), and prompts (task templates). It separates system roles-Host, Client, Server-and suggests negotiation mechanisms for capability discovery and permission granting. The framework promotes modularity and abstraction, and in theory, it supports general-purpose AI integrations.

However, in practice, MCP remains early-stage. Public specifications exist, and reference SDKs are emerging, but its deployment is largely confined to IDE assistants and localized environments. Despite claims of vendor neutrality, most implementations remain experimental or non-secure (e.g., Cursor, Sourcegraph, Windsurf). There is no clear path to secure, real-time operational deployment in production-grade applications.

SWAIG, by contrast, is fully implemented and running live across SignalWire's programmable voice and messaging infrastructure. It is not a toolkit for building an integration stack. It is the integration stack. SWAIG defines executable functions, validates arguments, manages consent and scope, and routes real-time media-all through a declarative interface embedded in SignalWire's agent markup language, SWML.

Rather than conceptual roles, SWAIG delivers operational behavior. Its functions can be serverless (templated calls to APIs) or routed to HTTP services, and they are invoked directly in conversation flows without glue code or orchestration layers. The agent understands how and when to call them, what to do with the results, and how to proceed in the interaction.

Design and Simplicity

MCP emphasizes a flexible, modular architecture. However, its design requires considerable scaffolding: multiple abstract layers (resources, tools, prompts, metadata, context), permission negotiation, and external orchestration for runtime behavior. This creates overhead and friction for developers trying to build production agents.

SWAIG simplifies this by doing less. Its primitives are compact and expressive: **function**, **parameters**, **data_map**, **output**, and **action**. These map directly to behavior and require no separate discovery mechanism or orchestration layer. One YAML or JSON block defines the interface, validation, conversational response, and downstream actions.

Example: Serverless Function with Data Map

Unset

```
- function: get_weather
  purpose: Get the current weather based on the user's location
  argument:
    type: object
    properties:
      location:
        type: string
        description: The name of the city or zip code
  data_map:
    expressions:
      - pattern: ".*"
        string: "${args.location}"
  output:
    response: "Weather data retrieved."
```

```
action:
  - SWML:
      version: "1.0.0"
      sections:
        main:
          - send_sms:
              to_number: "+15554441234"
              region: "us"
              body: "Here's the weather for ${args.location}"
              from_number: "+15554441234"
```

Example: Hosted Function with Webhook

Unset

```
- function: get_weather
  description: To determine the current weather for a given
location.
  parameters:
    type: object
    properties:
      location:
        type: string
        description: The location to check
  fillers:
    en-US:
      - "Checking the weather for you."
  web_hook_url:
https://api_key:secret_token@your-swaig-server.com/get_weather
```

Expected server response:

Unset

```
{
  "response": "It's currently 72°F with clear skies in Austin,
Texas.",
```

```
"action": {
  "set_meta_data": {
    "last_checked_location": "Austin"
  }
}
```

In this model:

- **response** shapes the next AI turn
- **action** defines operational behavior (e.g., say, transfer, metadata)

No external discovery. No runtime glue. All validated and run in context.

Integration in the Real World

MCP's primary integrations today are with IDE copilots and developer tools. These environments are constrained, localized, and generally asynchronous. They are valuable but limited in operational scope.

SWAIG is deployed in live telecom flows: PSTN, SIP, WebRTC, and SMS. A SWAIG agent can receive a phone call, parse speech via STT, call a serverless function, respond via TTS, and take action—all within a single call session. It runs real-time, across channels, with no intermediary orchestration.

Real-World Example: Appointment Scheduler

- Define **get_availability** as a hosted or serverless SWAIG function
- Ask the user for a time
- Use **data_map** to query the booking API
- Let user pick a time
- Use **send_sms** to confirm

No backend logic is required beyond the API call. No IVR scripting. No third-party orchestration.

SWAIG is already in use for:

- Voice bots answering questions and routing leads
- SMS agents scheduling appointments or collecting preferences

- Cross-channel assistants combining messaging, voice, and data
 - Real-time agents integrated with CRMs and operational backends
-

Security and Operational Trust

MCP focuses on user-centric permission and abstraction but does not define how runtime environments should enforce execution constraints. This leaves security as an implementation detail-flexible, but inconsistent and risky in operational settings.

SWAIG functions operate within a tightly controlled execution model:

- JSON Schema defines argument validation
- Only a fixed set of actions are permitted
- All functions execute within a call/session context
- Timeout, barge-in, and state transitions are managed by the platform
- Transitions and actions are logged and auditable

This yields a deterministic, traceable behavior model where AI actions are predictable and safe.

Examples of allowed actions:

- `say` or `response`: conversational output
- `SWML`: telephony instructions (transfer, hold, record)
- `set_meta_data`, `toggle_functions`, `stop`, etc.

The result is a structured, declarative, and secure environment for AI operations-critical in regulated industries or sensitive applications.

Scalability and Deployment

MCP aspires to be language-agnostic and modular. However, it lacks a shared runtime, unified execution model, or deployment guidance. Today, it runs mostly in browser tools or developer IDEs without scale benchmarks or secure endpoint integration.

SWAIG is already deployed across SignalWire's telecom-grade infrastructure. It inherits FreeSWITCH's performance and reliability, and has been used in thousands of concurrent voice sessions.

Teams can:

- Start serverless (API-only functions)
- Deploy hosted SWAIG servers (Heroku, Fly.io, AWS)

- Mix real-time calls, SMS, and backend logic
- Reuse functions across multiple agents
- Scale incrementally and predictably

This composability allows fast iteration without compromising structure or safety.

Built-in Tool Hosting and Discovery

A unique feature of SWAIG is its ability to dynamically fetch tools from external sources via HTTP POST. Developers can:

- Host SWAIG-compatible tools on their own server
- Expose functions via authenticated POST endpoints
- Register tool URLs in SWML agents declaratively
- Enable auto-discovery with explicit permission control

This is conceptually similar to MCP's vision of tool discovery, but significantly simpler. No negotiation protocol, no capability graph. Just set the endpoint, define the schema, and let the agent invoke it.

It is declarative, scalable, and secure-without overengineering.

How MCP Could Learn from SWAIG

MCP's structure is thoughtful, but early implementations lack the constraints and integration patterns required for production readiness. It could evolve faster by adopting elements from SWAIG:

1. **Data-first execution models**
 - Instead of resource abstractions, SWAIG uses `data_map` blocks to bind inputs to live data.
2. **Declarative action flows**
 - SWAIG returns typed, structured `action` blocks that eliminate ambiguity.
3. **Lifecycle integration**
 - SWAIG functions execute within voice/messaging sessions and manage state transitions.
4. **Unified format**
 - SWAIG functions behave the same whether serverless or hosted-minimizing mental overhead.

MCP's principles are solid, but operational grounding is what transforms protocols into platforms. SWAIG illustrates that real-world design requires execution, not just abstraction.

Operational Milestones

SWAIG has been running in production since early 2023, integrated directly into SignalWire's global telecom infrastructure. It has handled Thousands of concurrent real-time voice sessions.

This operational maturity is not speculative. It reflects learnings and refinements from live deployments in regulated, customer-facing environments.

Conclusion

MCP is a promising attempt to define the future of LLM-driven tools. It is modular, extensible, and aligned with privacy-first principles. But while it proposes how AI Gateways might work, SWAIG proves how they already do work, today.

SWAIG is declarative, enforceable, composable, and live in production. It eliminates unnecessary complexity, runs directly in the SignalWire platform, and powers real-time voice and messaging AI agents.

It is not a concept. It is a framework. It closes the loop between intent and action, between specification and behavior, between architecture and execution.

For organizations looking to operationalize AI agents today, not in theory, SWAIG provides the shortest path from design to deployment. It is the blueprint for operational AI agents. And it is already live.

Examples and More Information

To explore SWAIG and SWML in practice, see the following live resources:

- [SignalWire AI Platform](#)
- [AI Workshop Examples by @briankwest](#)
- [SignalWire Digital Employees Reference](#)
- [Official SWML AI Guides and Documentation](#)