Technology Update

MCP 2.0

Transforming DQL-Ready Data into AI-Ready Systems

Edited by **MOHAN DAS VISWAM**

Artificial Intelligence (AI) systems—particularly Large Language Models (LLMs)—have made remarkable advances in understanding, interpreting, and generating human language. These models are now widely deployed across domains for tasks such as information retrieval, document summarization, decision support, content generation, and conversational user interfaces. Their ability to reason over vast amounts of text and respond in natural language has significantly enhanced productivity and accessibility in both consumer and professional settings.

However, despite these advances, LLMs fundamentally operate as isolated computational systems. By design, they lack a native, standardized, and secure mechanism to interact directly with external tools, live databases, enterprise applications, APIs, or operational workflows. As a result, while LLMs can recommend or describe actions, they cannot reliably execute them within real-world systems without extensive external scaffolding.

This limitation severely constrains the practical deployment of AI in real-world and enterprise environments, where access to real-time data, controlled system actions, and compliance with organizational policies are critical. Existing integration approaches typically rely on custom-built connectors, bespoke middleware, or tightly coupled interfaces. These solutions are often brittle, difficult to scale across multiple models or tools, expensive to maintain, and introduce significant risks related to security, access control, auditing, and long-term governance.

MCP is an open standard that allows AI models to securely connect with external tools, data, and systems. By replacing complex, custom integrations with a unified client–server architecture, MCP enables safe, scalable, and governed AI interactions. It transforms AI from a standalone language model into a reliable, action-capable system suitable for enterprise and mission-critical use.

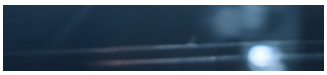

The Model Context Protocol (MCP) addresses these challenges by introducing an open, standardized communication framework that enables AI models to interact with external systems in a secure, governed, and interoperable manner. By serving as a common interface between AI models and real-world tools, data sources, and workflows, MCP eliminates the need for ad-hoc integrations and enforces a clear separation between AI reasoning and system execution. This architectural approach allows AI systems to move beyond passive language understanding and evolve into reliable, action-oriented applications—capable of operating within enterprise-grade constraints while maintaining trust, security, and scalability.

The Core Problem

Understanding the N × M Problem

Before MCP, integrating AI systems with external tools suffered from a major scalability issue known as the N* × M* problem:

**N: Number of AI models (e.g., GPT, Claude, Gemini)*
**M: Number of tools or data sources (e.g., APIs, databases, CRMs, file systems)*

Each AI model required a custom integration for every external tool.

Example:

- 3 AI models × 5 tools = 15 separate integrations
- High maintenance overhead
- Increased cost
- Greater risk of errors
- Significant security exposure

MCP's Architectural Solution

MCP replaces this complexity with a 1 × M or N × 1 architecture:

- Tools implement MCP once
- Any MCP-compatible AI model can immediately use them

This approach dramatically simplifies integration, improves maintainability, and strengthens security.

What is the MCP?

MCP is an open, standardized communication protocol that enables AI models to securely interact with external systems, tools, and data sources. Built on JSON-RPC 2.0, MCP defines a consistent client–server framework through which AI applications can discover and invoke external capabilities in a controlled manner.

MCP provides a unified way to expose tools, data resources, and structured prompts to AI models, while maintaining a clear separation between AI reasoning and system execution. AI models do not directly access infrastructure or APIs; instead, all interactions are mediated through MCP servers that validate and govern each request.

By connecting AI models to authoritative, real-time data and controlled actions, MCP helps reduce unreliable outputs and enables dependable, task-oriented AI behavior. This makes MCP particularly suitable for enterprise and government environments where security, auditability, and scalability are essential.

Niladri Bihari Mohanty
Scientist – D
niladri.mohanty@nic.in

Nikhil Kumar
Scientific Officer/ Engineer - SB
nikhil.kumar27@nic.in

MCP Architecture Overview

MCP Client

The MCP Client represents the AI-facing component of the Model Context Protocol. It is typically embedded within an AI application, agent framework, or LLM runtime and is responsible for managing interactions between the AI model and external systems.

Using a standardized JSON-RPC 2.0 interface, the MCP Client sends structured requests to MCP Servers to access tools, retrieve data, or execute predefined actions. It supports dynamic discovery of available capabilities, allowing the AI application to adapt at runtime without relying on hard-coded integrations.

Importantly, the MCP Client does not directly interact with infrastructure or external APIs. All real-world actions are delegated to MCP Servers, ensuring a strict separation between AI reasoning and system execution.

MCP Server

The MCP Server acts as the secure and authoritative gateway between AI models and real-world systems. It exposes well-defined and discoverable capabilities while preventing direct access to underlying infrastructure.

The server is responsible for validating requests, enforcing permissions, and executing approved actions initiated by AI models. It can encapsulate a wide range of backend systems, including APIs, databases, file systems, business logic, and internal services, and present them in a standardized, AI-ready form.

By centralizing control and execution, the MCP Server enables organizations to integrate AI into existing architectures without redesigning systems, while maintaining strong security, governance, and auditability.

Communication Layer

The Communication Layer forms the foundation of the Model Context Protocol and is built on JSON-RPC 2.0, a lightweight and widely adopted standard for structured communication.

This layer enables reliable, bidirectional communication between MCP Clients and Servers in a language-agnostic manner, allowing implementations across diverse platforms and programming environments. It introduces only minimal abstractions for tools, resources, and prompts, ensuring simplicity and ease of adoption.

Through its standardized and extensible design, the communication layer ensures interoperability, scalability, and long-term stability for AI-system interactions.

Security and Governance in MCP

Security is a first-class principle in the Model Context Protocol. Its key features are:

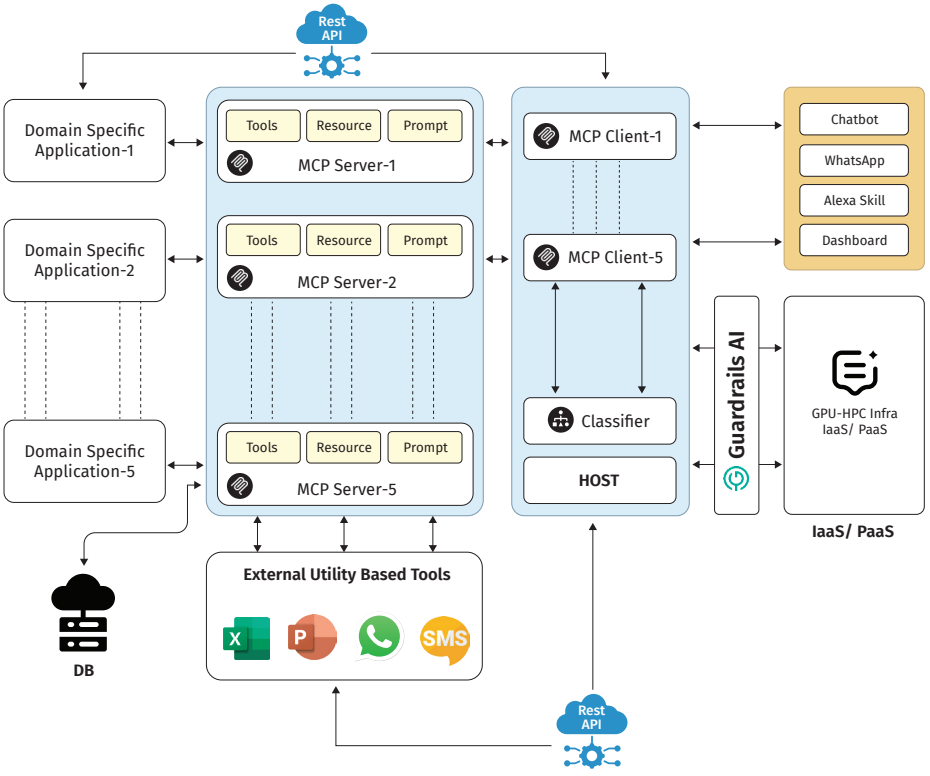

▲ Fig 11.1 **MCP Architecture Overview**

- Role-Based Access Control (RBAC)
- Capability scoping
- No direct system access by LLMs
- Server-side validation of all actions
- Auditable execution trails

At no point does the AI model directly interact with infrastructure. All access is mediated through controlled MCP interfaces, ensuring compliance and trust.

The Future of MCP

The Model Context Protocol (MCP) is emerging as a foundational infrastructure layer for next-generation AI systems, enabling the development of autonomous agents that can plan, reason, and execute tasks across multiple steps. By providing a standardized and governed interface to external tools and systems, MCP allows AI models to operate reliably within real-world environments rather than responding to isolated prompts.

MCP also enables robust multi-step reasoning, allowing AI systems to iteratively interact with data, tools, and workflows while maintaining a strict separation between reasoning and execution. This structured approach ensures controlled access to authoritative data and predictable system behavior.

At the enterprise level, MCP serves as a core building block for scalable, secure, and interoperable AI platforms. By replacing fragmented integrations with a unified protocol, it allows organizations to deploy AI across systems and vendors without compromising governance or compliance. As AI evolves from conversational chatbots to action-oriented "do-bots," MCP provides the essential infrastructure to safely connect intelligence with execution—transforming AI into a trusted, action-capable participant in real-world workflows.

Way Forward

MCP represents a decisive shift in how artificial intelligence systems engage with the real world. By resolving long-standing challenges related to integration complexity, security governance, and scalability, it establishes the foundational infrastructure needed to move beyond isolated language models. MCP elevates AI from a passive source of insight to a secure, reliable, and action-capable agent, capable of interacting with live systems, data, and workflows under well-defined controls. As AI adoption expands into mission-critical domains, MCP stands out as a key enabling standard—allowing intelligence to operate with trust, precision, and real-world impact.

Contact for more details

State Informatics Officer
NIC Odisha State Centre
Unit-IV, Sachivalaya Marg
Bhubaneswar, Odisha-751001
Email: sio-ori@nic.in