Agentic AI Foundation

The Agentic AI Foundation: Open Standards and the Sovereign Agentic Stack

Stewarding the Linux Foundation Ecosystem and Seeding Singapore's Community



The Macro Shift: From Inference to Autonomous Loops

Agentic AI Foundation

Phase 1: Static Inference

User Prompt

↓

LLM Processing

↓

Text Response

Stateless, deterministic, single-turn execution.

Phase 2: Autonomous Multi-Agent Loops

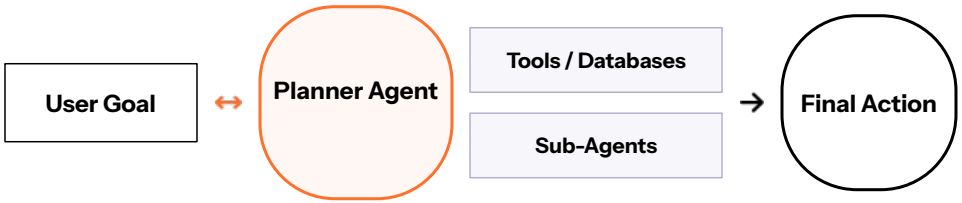

Stateful reasoning, multi-turn tool invocation, non-deterministic environmental impact.

As enterprise AI transitions to autonomous loops, brittle Python glue code and proprietary runtimes create severe operational liabilities.

The 3-Layer Open Agentic Stack

Agentic AI Foundation

Layer 3: Agentic Interaction & Governance (AAIF Standards)

MCP | A2A Protocol | AGENTS.md | Goose

- Standardizes agent-to-tool (MCP) and agent-to-agent (A2A) interactions under neutral Linux Foundation governance [AAIF].
- Decouples agent cognitive reasoning from proprietary host environments and vendor APIs.

Layer 2: Cloud-Native Substrate & Governance (CNCF Orchestration)

Kubernetes | KServe | llm-d Fabric | Envoy AI Gateway | vCluster / DRA

- Offloads operational tasks—token rate-limiting, KV-cache routing, and confidential sandboxing—from Python app code into platform primitives.
- Enforces zero-trust workload identity (SPIFFE) and multi-tenant GPU isolation.

Layer 1: Distributed Compute & Hardware Abstraction (PyTorch Foundation)

PyTorch Core | Ray | vLLM Engine | Triton | OpenReg

- Unifies model training, distributed orchestration, and high-throughput inference serving under open governance.
- Delivers cross-architecture enablement (NVIDIA CUDA, AMD ROCm, Intel, NPUs) to eliminate silicon lock-in.

The Strategic Architecture Shift

Moving away from fragmented, ad-hoc Python wrappers and proprietary API rentals toward a connected, vendor-neutral operational fabric that guarantees sovereignty across any cloud and accelerator.

Enterprise Agentic Bottlenecks

Agentic AI Foundation

01 Legacy Data Friction

Mainframes and ERPs lack modern APIs, forcing Change Data Capture (CDC) mirrors that add latency and stale context.

02 Machine-Speed Execution Risk

Compromised reasoning triggers unauthorized production mutations; post-hoc audit logs cannot intercept damage in flight.

03 Shadow Amnesia & Credential Sprawl

Unregistered MCP tools retain standing credentials and live database connections outside enterprise governance perimeters.

04 Multi-Agent Accountability

Valid single-agent decisions combine across handoffs into invalid outcomes, breaking single-owner RBAC and audit models.

Reference:
Enterprise Transitions to the Agentic Era (LF Research, September 2026)

The Anchor: The Agentic AI Foundation (AAIF)

Agentic AI Foundation

Hosted by the Linux Foundation

Neutral open-source governance framework enabling cross-industry collaboration and vendor-neutral IP management.

Neutral IP Stewardship

Ensures open specifications, patent non-assertion, and trademark protection across foundational protocols.

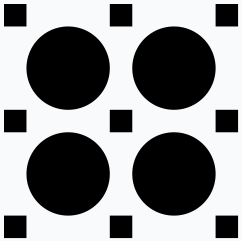

300+

Member Organizations

Global coalition of cloud providers, AI labs, enterprise adopters, and chipmakers.

10+

Foundational Projects

Core protocols & runtimes including MCP, A2A, AGENTS.md, Goose, and Agent Router.

AAIF provides a neutral, multi-stakeholder foundation to ensure the core protocols of agentic AI remain open, interoperable, and free from proprietary runtime lock-in.

AAIF Foundational Projects

Agentic AI Foundation

Context & Capabilities

AGENTS.md
Simple format for guiding coding agents

Agent Skills
Open format for reusable agent expertise

Runtimes & Environments

Goose
Extensible AI agent going beyond code suggestions

OpenSandbox
General purpose sandbox platform for AI

OpenEnv
Open environments for training AI agents

Routing & Gateways

Agent Router
Open AI gateway on Envoy for models & MCP tools

agentgateway
Secure, scalable connections across agents & APIs

Core Protocols

Model Context Protocol (MCP)
Integration protocol for LLMs and external data

Agent2Agent (A2A)
Open protocol for multi-agent communication

Open Responses
Shared LLM API format across providers

How Open Protocols / Standards Help

Agentic AI Foundation

Eliminates N×M Custom Glue

MCP defines a typed JSON-RPC contract connecting any agent runtime to tools and context stores.

Enforces Policy Beneath Models

MCP Interceptors and gateways validate tool calls pre-execution instead of trusting model self-policing.

Prevents Vendor Lock-In

Open governance under AAIF ensures interoperable discovery, identity, and hybrid local/frontier routing.

GOOSE	Agent runtime & loop; routes 80–90% routine tasks to local models
A2A	Cross-vendor agent-to-agent task delegation and state handoffs
MCP	Agent-to-tool JSON-RPC interface with runtime message Interceptors
AGNTCY	Cryptographic agent identity, service discovery, and trace telemetry
GATEWAY	**Baseline control: pre-execution policy, RBAC, token metering & audit**

Model Context Protocol

Agentic AI Foundation

MCP Client

The AI Assistant / IDE — Requests context or actions

JSON-RPC 2.0 over stdio or SSE

MCP Host

The coordinating environment — Manages lifecycle, enforces safety policies, aggregates context

JSON-RPC 2.0 over stdio or SSE

MCP Server

External Enterprise Tools & Databases — Exposes specific tools without sharing full application memory

TRADITIONAL INTEGRATION · M × N

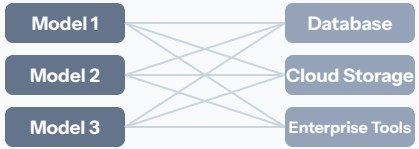

Complex & Fragile

Requires custom code for every pair (9 total connections).

UNIFIED MCP STANDARD · M + N

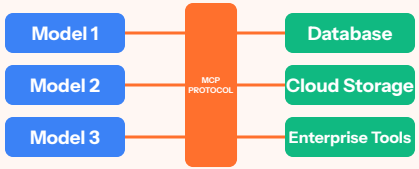

Scalable & Standardized

Integrate once per entity (6 total connections).

Replaces the M×N custom integration nightmare with a unified M+N architecture. Vertical tool execution is standardized and secure.

MCP Specification 2026-07-28: The Cloud-Native Evolution

Agentic AI Foundation

RELEASE CONTEXT: Published July 28, 2026. Marks MCP's transition from a session-bound desktop protocol to an enterprise-grade stateless HTTP substrate.

Stateless Core & L7 Routing

- Removed **initialize/initialized** handshake and **Mcp-Session-Id** header (SEP-2575, SEP-2567). Runs behind round-robin load balancers.
- Protocol version, client ID, and capabilities travel in per-request **_meta**; **server/discover** enables optional pre-flight checks.
- Native **Mcp-Method** and **Mcp-Name** HTTP headers enable Layer 7 routing and CEL authorization on gateways (**Envoy**, **Agent Router**, **agentgateway**).
- Tool caching hints (**ttlMs**, **cacheScope**) and W3C Trace Context (**traceparent**, **baggage**) in **_meta** for OpenTelemetry.

Asynchronous Flows & Ext.

- Multi Round-Trip Requests (**MRTR**) replace long-lived SSE streams for sampling and input elicitations via **resultType: 'input_required'**.
- Formal **Extensions Framework** with reverse-DNS identifiers (**io.modelcontextprotocol/***) and decoupled release cycles.
- **Tasks Extension** (SEP-2663) standardizes long-running jobs via task handles and **tasks/get** (**tasks/list** removed for security).
- **MCP Apps** (SEP-1865) brings pre-fetched, server-rendered UI components in sandboxed iframes over standard JSON-RPC.

Enterprise Auth & Governance

- Hardened OAuth/OIDC with **RFC 9207 iss validation** to block mix-up attacks across multi-server topologies.
- Shift from Dynamic Client Registration to **Client ID Metadata Documents (CIMD)**; **application_type** declarations fix desktop localhost redirects.
- Full **JSON Schema 2020-12** support for tool inputs/outputs (**oneOf**, **anyOf**, **$ref**) and JSON-RPC error **-32602**.
- Strict **Tier 1 SDK** conformance suite requirements across TypeScript, Python, Go, and C#.

12-MONTH DEPRECATION WINDOW (Active → Deprecated → Removed):

Roots (use tool params/URIs) · **Sampling** (use provider APIs) · **Logging** (use OpenTelemetry) · **Legacy HTTP+SSE** (use Streamable HTTP)

Protocol Deep Dive: Agent2Agent (A2A)

Agentic AI Foundation

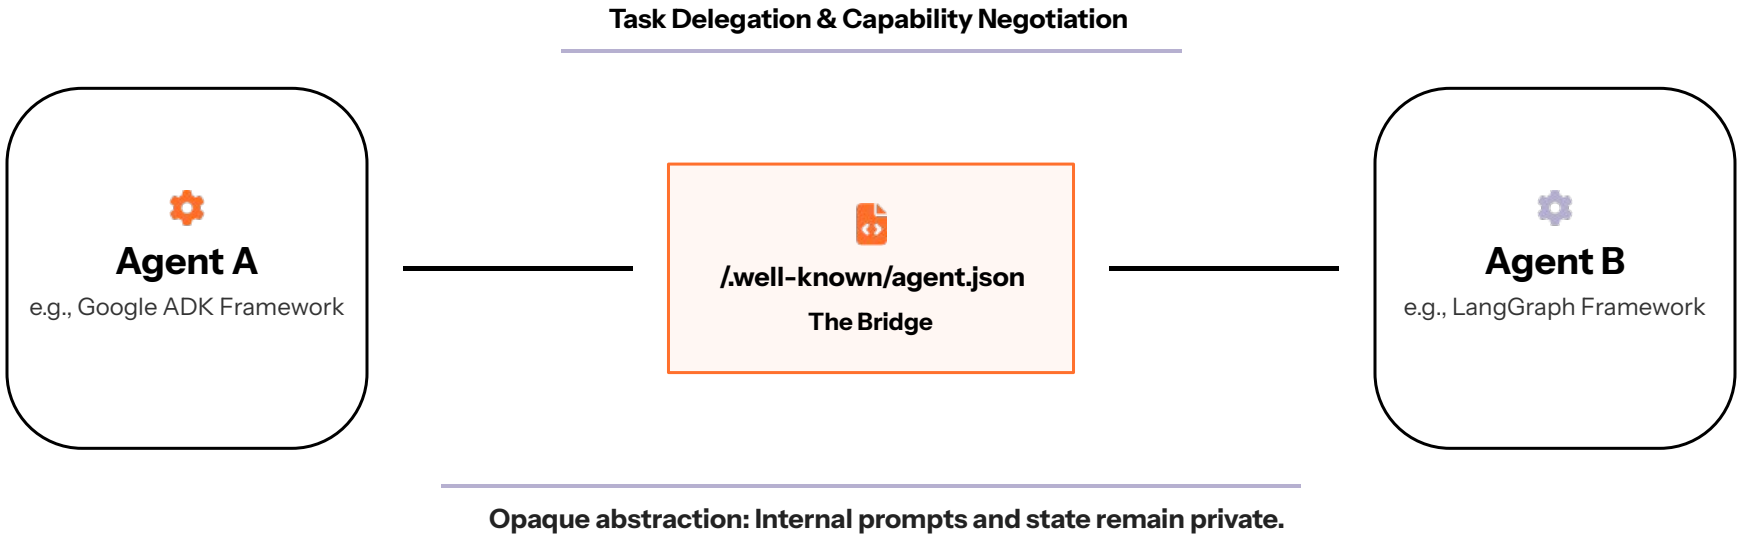

Enables autonomous cross-organizational collaboration without exposing internal proprietary code or memory state.

Protocol Deep Dive: Open Responses

Agentic AI Foundation

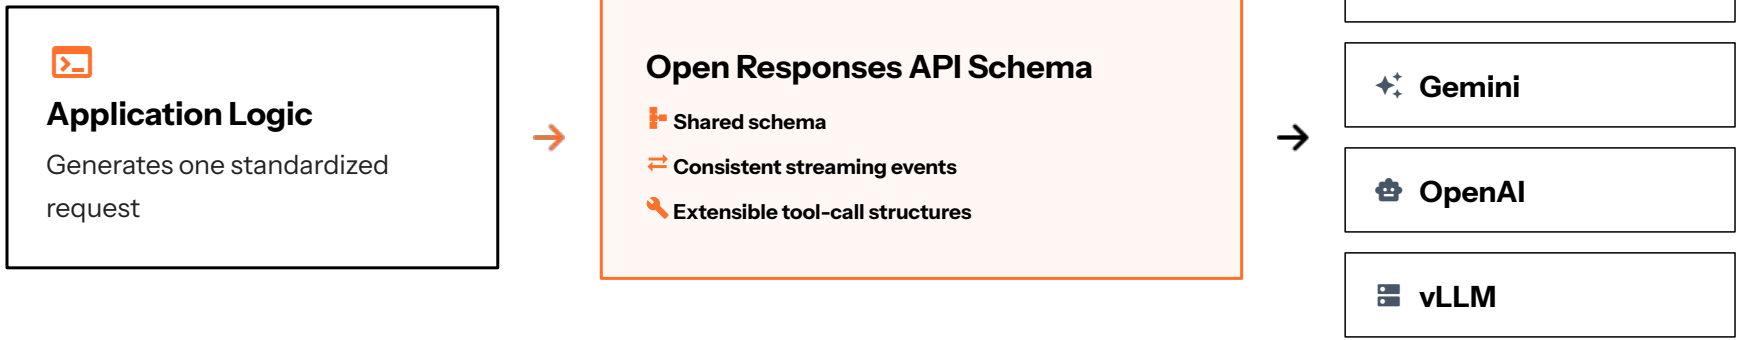

Developers write model invocation logic once and execute requests across disparate backend providers with zero custom translation layers.

The Agentic Protocol Matrix

Agentic AI Foundation

Dimension	MCP	A2A	Open Responses
Directionality	North-South (Agent-to-Tool)	East-West (Agent-to-Agent)	Provider-Agnostic (App-to-LLM API)
Architectural Role	Client-Host-Server topology	Peer-to-Peer system exchange	Universal Inference Interface
Capability Discovery	Dynamic exchange via server/discover	Static metadata via Agent Card	N/A (Standardized Request/Response)
Primary Target	Deterministic function invocation	Goal-oriented task delegation	Inference and token generation

Each standard governs a distinct operational boundary, eliminating overlapping protocols and ensuring modular architecture.

Unresolved Architectural & Policy Gaps

Agentic AI Foundation

Domain	What Standards Solve Today	Remaining Production Gap
Stack & Legacy Coverage	MCP standardizes tool calls across cloud-native APIs (1 of 5 layers).	Non-API legacy systems still require costly CDC pipelines; A2A and AGNTCY maturity lags MCP.
Protocol Ergonomics	MCP enforces typed JSON-RPC schemas and deterministic tool contracts.	High coding overhead vs. declarative Markdown Skills; pattern convergence remains unsettled.
Multi-Agent Liability	Single-agent execution maps cleanly to one authorizing human owner.	Emergent failures across A2A chains and multi-principal delegation lack attribution models.
Policy Scale & Ops	Gateways and MCP Interceptors block policy violations pre-execution.	Authoring, testing, and automating deterministic rules across thousands of agents is still manual.

Reference:
Enterprise Transitions to the Agentic Era (LF Research, September 2026)

Bridging Global Standards to Singapore's Model Governance Framework

Agentic AI Foundation

Singapore IMDA MGF v1.5 (Policy)	The Protocol Bridge (Enforcement Plumbing)	AAIF & Cloud-Native (Code Execution)
1. Assess & Bound Risks Upfront · Least-privilege tool access · Cryptographic agent identity	`SPIFFE/SPIRE SVIDs + CEL Policies` `Filters Mcp-Method / Mcp-Name at L7 layer`	**Agent Router, agentgateway & MCP Core** CIMD & RFC 9207 authentication tokens
2. Meaningful Human Accountability · Fail-closed gates for high-impact actions · Guard against automation bias	`MCP Multi Round-Trip (MRTR) + Apps` `Pauses loop with resultType: 'input_required'`	**A2A Protocol v1.0, MCP Apps & Goose** Human-in-the-loop inter-agent delegation
3. Technical Controls & Processes · Hardware isolation of code execution · Continuous auditing & tracing logs	`MicroVM Sandboxing + W3C Trace Context` `gVisor/Kata execution + distributed tracing (telemetry)`	**OpenSandbox, K8s agent-sandbox** OpenEnv & Project Moonshot integration
4. Enable End-User Responsibility · Upfront AI disclosure & boundaries · Versioned human runbooks & skills	`Git-Versioned Declarative Manifests` `JSON Schema 2020-12 tool schemas & CI/CD`	**AGENTS.md & Agent Skills (SKILL.md)** Open Responses & SEA-Guard standards

From Principles to Plumbing: Enforcing Singapore's Model AI Governance Framework via Open, Vendor-Neutral AAIF Protocols.

Ref: IMDA MGF for Agentic AI — https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf

The Singapore Sovereign Ecosystem

Agentic AI Foundation

TOP LAYER
Assurance & Evaluation

AI Verify's Project Moonshot integrated directly into OpenEnv loops for red-teaming.

AI Verify Project Moonshot OpenEnv Red-Teaming

MIDDLE LAYER
Safety & Routing

SEA-Guard and GovTech's AI Guardian (Litmus/Sentinel) acting as policy enforcers inside agentgateway.

SEA-Guard GovTech AI Guardian agentgateway

BOTTOM LAYER
Regional Compute

AI Singapore's SEA-LION v4 (multimodal) and MERLION-3 models exposed via Open Responses.

AI Singapore SEA-LION v4 MERLION-3 Open Responses

Open standards allow Singapore to swap out generic global models for culturally-aware, sovereign components without breaking the enterprise architecture.

Actionable Pathways for Singapore Community

Agentic AI Foundation

1 Inaugurate a Local Chapter

Bridge platform engineers and AI developers by co-locating AAIF meetups with CNCF Singapore and KSUG.AI.

2 Build Sovereign MCP Connectors

Develop open-source MCP tool servers for regional data APIs (e.g., Data.gov.sg, LTA DataMall) and payment infrastructure (PayNow).

3 Integrate Sovereign Models

Connect Southeast Asian open foundation models into MCP and A2A workflows to address localized dialects and regulatory nuances.

4 Operationalize MGF Blueprints

Package declarative MCP allowlists and OpenSandbox policies so SMEs can implement IMDA's governance out of the box.

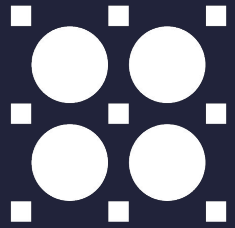

Agentic AI Foundation

Where the **agentic stack** is **being built**.