



WHITE PAPER

# The Importance of Being Open for AI

## Why AI Agility, Innovation, and Enterprise Value Depend on Escaping the Walled Gardens of Legacy CCaaS Vendors

### Table of Contents

**Introduction:**  
The Architecture of Control .....1

**Section 1:**  
A Strategic Inflection Point in Enterprise AI .....2

**Section 2:**  
What Does It Mean to Be Open for AI? .....3

**Section 3:**  
The Walled Gardens of Legacy CCaaS Vendors .....4

**Section 4**  
Why Openness Wins .....5

**Section 5:**  
Avaya Infinity in Action .....6

**Section 6:**  
The Risks of Closed AI Platforms .....6

**Section 7:**  
The Future Is Open .....7

**Section 8:**  
Why Avaya Is the Strategic Choice .....7

**Appendix:**  
RFP Questions to Evaluate AI Openness .....8

### Introduction: The Architecture of Control

In 1995, AOL was booming. Millions of Americans were discovering the internet through its curated channels and chirping dial-up tones. For a brief moment, it felt like AOL was the internet. Safe. Predictable. Packaged.

But just beyond its tidy borders was something messy and magnificent: the open web. No single company could control it. And that, as history proved, was exactly the point.

Today, enterprise AI stands at a similar crossroads.


The familiar path is lined with comforting promises: one vendor, one platform, one solution. Everything pre-integrated. Everything is branded. Everything is "agentic." But dig deeper, and you find a different truth: these environments are designed not for freedom, but for containment. AI becomes a walled garden, lush but limited.

Another path is emerging. Less polished but infinitely more powerful. This path prioritizes open protocols over proprietary APIs, collaboration over consolidation, and modularity over monoculture. It doesn't ask who owns the stack. It asks: who should?

The question facing every enterprise is no longer "Do we adopt AI?" That decision has already been made. The question now is: Who controls the intelligence of your business?

This paper explores why that question matters more than any product demo or roadmap slide. We unpack how closed architectures quietly trade your future agility for short-term convenience, and how open, protocol-based orchestration restores control, fosters innovation, and protects your most valuable asset: your ability to evolve.

This is not a paper about features. It is a paper about freedom.



**Vendors are now competing to control the orchestration layer that governs AI decisions, the context layer that stores memory and intent, the model layer that generates responses, and the integration layer that connects systems.**

## Section 1: A Strategic Inflection Point in Enterprise AI

Enterprise communications has entered a new competitive era. For decades, contact center innovation followed a predictable path: add a channel, improve routing, optimize workforce management, and layer in analytics. Each advance improved efficiency but did not fundamentally change the architecture of customer engagement.

Generative AI changes everything. What was once deterministic and rule-based is now probabilistic and adaptive. What was once a workflow is now an intelligent system. What was once a queue is now a network of collaborating agents.

Vendors are now competing to control the orchestration layer that governs AI decisions, the context layer that stores memory and intent, the model layer that generates responses, and the integration layer that connects systems. This is no longer a feature battle. It is an architectural battle.

### Two Diverging Paths

As generative AI moves from experimentation to deployment, two architectural philosophies are emerging.

**Path One:** Closed, Vertically Integrated AI Stacks. Some vendors are consolidating control by building proprietary orchestration layers tightly coupled with their own AI models, marketplaces, and cloud infrastructure. These environments promise simplicity but deliver limited model choice, restricted integration flexibility, dependence on a single innovation cycle, and escalating switching costs.

**Path Two:** Open, Protocol-Based AI Orchestration. Other vendors are embracing interoperability. Rather than controlling every layer, they standardize how layers communicate. Open architectures separate AI models from orchestration, orchestration from memory, memory from data sources, and data sources from the deployment location. This separation enables enterprises to evolve without rebuilding, shifting power from the platform vendor to the platform user.

### The Control vs. Innovation Paradox

At the center of the platform race lies a strategic paradox: the more control a vendor claims, the less innovation it enables for its customers. Closed ecosystems appear attractive at first. One interface. Pre-trained models. Tight integration. But the cost is clear: you cannot evolve faster than the vendor does. If a new open-weight model outperforms your vendor's LLM, you cannot easily switch. If your data requires stricter sovereignty, you may be forced to choose between compliance and innovation.

Avaya has taken a different path. Rather than building a proprietary AI layer, Avaya natively incorporates the Model Context Protocol (MCP) industry standard protocol, enabling any AI-governed tool, in any environment, to be governed by the customer. Customers decide which models to use. AI agents collaborate across tools and departments. Context is preserved across cloud, on-prem, and hybrid systems. Innovation decisions remain in the hands of the enterprise.

### The Five Pillars of Open AI Architecture

- 1. Open Protocols.** Standardized communication rules that allow AI agents, systems, and memory layers to coordinate regardless of origin. MCP is a prime example of how AI agents can share intent, memory, and tools in a vendor-neutral format.
- 2. Model Neutrality.** The ability to switch LLMs or agent models without re-architecting the stack, supporting flexibility across closed LLMs (such as GPT-4 or Claude) and open-weight models (such as LLaMA or Mistral).
- 3. Agent-Agnostic Orchestration.** Support for agents from any vendor without favoring the vendor's own tools, ensuring a best-of-breed approach across domains and customer journeys.
- 4. Interoperable APIs.** APIs that follow universal standards like JSON-RPC 2.0, making tool and system connections plug-and-play rather than bespoke.
- 5. Context Continuity**—the ability for different agents to share context, memory, and history across systems. Without shared memory, AI becomes fragmented. With it, agents become collaborative.

## Section 2: What Does It Mean to Be Open for AI?

"Open" has become one of the most misused terms in the enterprise AI lexicon. Nearly every vendor claims some version of openness, but few deliver it in a way that truly empowers the customer. This is not about whether a model is open-source. It is about whether your architecture gives you agency.

An AI system is open if the customer controls the key levers of intelligence: which models are used, where memory is stored, and how or which agents are permitted to act. If a vendor limits any of these, the system is not truly open.

This is why Avaya's definition of 'open' is not just a marketing claim—it is an objective standard backed by the industry. As a member of the Agentic AI Foundation (AAIF), a Linux Foundation initiative, Avaya is committed to open-source governance. The Model Context Protocol (MCP) we champion is not controlled by a single vendor; it is an open standard governed by a global coalition. This ensures that the architecture prioritizing your agency remains neutral, transparent, and driven by community consensus rather than proprietary lock-in.


Maturity Level	Characteristics	Risk Profile
Closed	One-vendor LLM, black-box orchestration, fixed tools	High vendor lock-in
API-Connected	Some third-party tools via proprietary APIs	Medium flexibility
Open Protocol-Based	Agentic orchestration, model-neutral, context sharing	High adaptability, future-proofed

Most vendors remain in the first or second tier, marketing "open" while architecting closed. Only those with true protocol support, like MCP, operate in the third.

### Avaya Infinity: Designed for Openness from Day One

MCP is a standardized protocol that connects AI agents to memory, tools, and external systems using a modular plug-in architecture. It enables model neutrality, context continuity, deterministic agent orchestration, and enterprise governance with defined trust boundaries and policy enforcement.

MCP is not middleware. It is a paradigm shift that turns static workflows into adaptive intelligence frameworks. It enables Tandem Care, a model in which AI initiates, guides, and supports an interaction, then seamlessly passes it to a human or another AI without loss of memory, intent, or emotional context. And unlike competitors who treat MCP as a marketing label, Avaya has made it an architectural backbone native to the Avaya Infinity platform.



**Avaya's open MCP architecture prioritizes transparency, composability, and co-innovation, allowing enterprises to retain sovereignty over their data and flexibly compose AI experiences with any LLM or engine they choose.**

### **Section 3: The Walled Gardens of Legacy CCaaS Vendors**

As enterprises explore AI-powered transformation, many vendors promise openness while quietly erecting walls. These walls are not made of code, but of architecture, licensing, and integration limits. Three dominant approaches illustrate the problem.

#### **Vendor Type X: Organic Walled Garden Development**

This archetype builds rather than partners, creating a highly curated, vertically integrated system. Orchestration, memory, and models are developed in-house. External AI tools are limited to pre-approved APIs. The vendor's marketplace reinforces a "native-first" mindset. The term "agentic orchestration" suggests a breakthrough. Still, in practice, agent autonomy is limited to the vendor's cloud, and AI agents can feel constrained when embedded only as an add-on rather than as a native, first-class enterprise. They may be deploying a dynamic agent framework while actually embedding their CX strategy within a single vendor's roadmap.

#### **Vendor Type Y: Walled Garden via Acquisition**


This archetype centralizes intelligence within a proprietary AI framework, recently amplified by strategic acquisitions of conversational AI platforms. Bots are pre-trained, pre-packaged, and pre-positioned. Customers are discouraged from bringing their own models or agents. There is no agent-agnostic orchestration, no persistent memory that flows across tools or vendors, and no open protocols at the architectural layer. This stack is frequently flagged for high upfront costs, disjointed user interfaces, and rigid license models.

#### **Vendor Type Z: Holistic Hardware-Software Control**

This archetype requires vendor-managed UC and CC layers, vendor-branded AI orchestration, and vendor-certified infrastructure. APIs are available, but true orchestration flexibility is gated by full-stack alignment. This vendor often positions itself as "open," citing API access and partner tools, but this is interface-level openness, not architectural openness. You can build on top of them as long as you use their base. You can integrate AI as long as it conforms to their framework.

#### **The Common Thread: Breaking the Black Box**

All three models present elegant façades but distance customers from the raw materials of transformation: data, models, context, and code. Opaque model behavior means businesses cannot see how AI decisions are made. Stunted innovation means AI capabilities are only as flexible as the vendor roadmap allows. Data captivity means valuable interaction data may be locked inside proprietary systems, inaccessible for reuse. By contrast, Avaya's open MCP architecture prioritizes transparency, composability, and co-innovation, allowing enterprises to retain sovereignty over their data and flexibly compose AI experiences with any LLM or engine they choose.



**With Avaya's MCP-native Infinity platform, new capabilities can be layered in rather than locked in. What takes months in closed systems can take weeks or days in open ones.**

## Section 4: Why Openness Wins

### Openness Drives Innovation

Open AI architectures create a foundation where innovation can emerge from any direction. With standardized protocols like MCP, enterprises can plug in new tools without ripping apart the foundation. New AI models or tools can be deployed without disrupting existing workflows. An open architecture turns every compatible tool into a potential advantage. Instead of relying solely on one vendor's innovation pipeline, organizations benefit from a global AI ecosystem.

Consider a financial services enterprise deploying a fraud-detection bot, a real-time location service, and a custom policy rules engine. In a closed system, this requires vendor approval, APIs to be created or purchased, and a lengthy rollout. In an open system like Avaya's, each function can be exposed via an MCP-compliant server, requiring no vendor coordination, and can be deployed in weeks rather than months.

### Openness Protects Enterprise Control

When vendors control your orchestration layer, your memory, and your models, they don't just host your AI. They own your ability to evolve. With an open architecture based on MCP, enterprises choose which LLMs power which agents, where and how memory is stored, which tools are orchestrated in what sequence, and which agents act under what rules. Your AI stack becomes modular, portable, and observable rather than monolithic, locked in, and opaque.

### Openness Accelerates Time to Value

Closed architectures slow time to value by locking enterprises into sequential workflows: waiting for vendor roadmap support, navigating custom integration requirements, and dealing with multi-month onboarding cycles. With Avaya's MCP-native Infinity platform, new capabilities can be layered in rather than locked in. Rapid integration through standard interfaces, immediate interoperability with existing systems, and reuse of context and logic structures across workflows mean that what takes months in closed systems can take weeks or days in open ecosystems.

### Openness enables and models the behavior of AI

Proprietary, closed-loop AI systems often act as black boxes. Their decision trees are hidden, their training data locked away, and their model's behavior unpredictable. Open AI systems with MCP allow for transparent, trainable inputs so auditors can trace how prior inputs influenced decisions, explainable reasoning paths aligned with enterprise policies, and retrainable, swappable models so biases can be corrected without vendor lock-in. These capabilities turn AI from a black box into a glass box, essential for ethical deployment.

### Openness Future-Proofs Your Investment

Foundation models are evolving monthly. New capabilities such as agentic workflows, synthetic memory, and multimodal reasoning are emerging rapidly. Platforms like Avaya Infinity, with MCP at the core, are model-agnostic (swap or add models without rewriting workflows), composable and modular (add new agents or tools over time), and scalable across use cases from contact centers to field service. Your AI investment remains valuable even as the landscape shifts.

Furthermore, because Avaya aligns with the Agentic AI Foundation, enterprises are permanently shielded from 'protocol drift' or forced migrations. The underlying standards evolve through broad industry consensus, not unilateral vendor decisions, guaranteeing that your AI foundation will remain universally compatible as new breakthroughs emerge.

## Section 5: Avaya Infinity in Action

### Use Case Concepts

**Hybrid Healthcare Agent.** A healthcare provider leverages Avaya Infinity's MCP to coordinate a medical history summary that adheres to human nurses for high-priority escalations. Each agent accesses the patient's case file in real time via a central memory layer that adheres to open standards. No data is duplicated or lost in handoffs. Expected impact: faster resolution, fewer callbacks, higher staff satisfaction.

**Financial Services Copilot.** A global bank uses Avaya Infinity to orchestrate a private, sovereign-hosted LLM trained on regulatory policy, a policy database agent that provides explainable decision trees, a Specialized Action Model for transaction anomaly detection, and a human compliance officer with a secure AI copilot. New models can be layered in over time without rearchitecting the system. Expected impact: faster decisions, zero governance gaps in regulatory audits.

**Retail Customer Journey.** A major retailer harmonizes voice IVR agents on AWS Lex, chatbots on Meta's LLaMA 3, store associate copilots, and loyalty agents using historical data. Unlike platforms that silo data by channel, Avaya's open architecture ensures context persists across all systems. A customer who starts a return online can walk into the store and be greeted by an associate already briefed on the interaction.

### Architectural Advantage

Avaya Infinity separates and modularizes each layer of the AI stack. The memory layer maintains long-term state shared across agents via open MCP protocols. The orchestration layer governs routing with transparent, model-agnostic policies. The agent layer executes tasks through plug-and-play MCP-compliant tools. The interface layer provides modular, multichannel, context-aware integrations.

At the heart of this architecture lies the persistent context layer, enabling context to follow a customer across channels, sessions, and agents. Unlike traditional platforms that silo memory per bot, Avaya's open memory model stores context once and shares it across bots. This eliminates data reentry during handoffs, preserves sentiment and history, and coordinates action between agents to reduce errors and duplication.

### Partner Enablement

Avaya Infinity is built to empower the entire ecosystem. System integrators can design multi-agent workflows using MCP without waiting for vendor updates. Innovation labs can prototype agents that plug into live workflows without disrupting operations. Enterprise customers can swap models, deploy internal tools, or integrate their own CRM without creating fragile code. Avaya's ecosystem already includes AI partners such as Google and Microsoft, specialty providers such as Verint, and regional innovation labs experimenting with multi-agent workflows in healthcare, telecom, and the public sector.

## Section 6: The Risks of Closed AI Platforms

**Governance Gaps.** Legacy vendors promote compliance promises, but proprietary orchestration logic is buried in vendor-controlled microservices, decision trees are undocumented, context flows are invisible, and audit trails are limited. In healthcare, finance, and government, this triggers regulatory violations, legal liability, brand damage, and operational bottlenecks. Avaya Infinity provides editable orchestration policies, fully auditable decision trails, explicit context boundaries, granular trust permissions, and built-in rollback options.

**Lack of Portability.** In closed systems, memory is vendor-bound, agents are proprietary, workflows are hardcoded, and APIs are custom. Attempting to migrate can take 12 to 18 months. With Avaya and MCP, memory is portable and standards-based, agents are compatible across vendors, orchestration is rules-based and externalized, and deployments span cloud, on-prem, and sovereign environments.

**TCO Spiral.** Closed AI vendors promise everything "out of the box," but dependencies create escalating costs. Bundled agent pricing, token-based model lock-in, custom integration taxes, and upgrade complexity compound year over year. **Research**



# While competitors announce MCP support in limited, cloud-only, contact-center-only forms, Avaya's MCP is embedded, extensible, and enterprise-wide, spanning cloud, hybrid, on-prem, and sovereign cloud deployments.

from MIT Sloan found that proprietary models cost 6x more per million tokens than open-weight models while delivering only marginal accuracy improvements over time, with estimated global savings of \$24.8B annually from shifting workloads where appropriate\*. Avaya decouples model usage from vendor fees, lets enterprises bring their own LLM, minimizes integration costs with MCP servers, and supports hybrid deployments to control infrastructure costs.

Regulatory Risk. GDPR mandates the right to explanation for automated decisions. HIPAA requires traceable handling of personal health information. The EU AI Act classifies certain AI applications as "high risk" and demands full transparency. If you can't explain it, you can't defend it. Avaya Infinity addresses regulatory risk at the architecture level: all AI actions are routed through a policy-governed orchestration engine, context boundaries are explicitly defined, and every action is timestamped, logged, and traceable. New regulatory requirements can be adopted without vendor dependence.

## Section 7: The Future Is Open

### The Rise of AI Interoperability Protocols

Just as HTTP transformed the internet and SMTP unlocked global email, a new class of AI interoperability protocols is defining how intelligent agents, models, and tools coordinate across vendors and environments. The Agentic AI Foundation (AAIF) is institutionalizing these standards. Backed by nearly 150 leading organizations—including foundational players like AWS, Google, and Microsoft—the AAIF ensures that protocols like MCP provide model-agnostic orchestration, context persistence, and vendor-neutral interoperability. Enterprise platforms are already adopting it, and SI partners, government sectors, and customers are demanding future-proof architectures.

### What Enterprises Must Demand

**Every enterprise should require:** full control over data, memory, and orchestration logic; vendor-agnostic integration of any LLM, SAM, or tool without vendor approval; protocol-defined

interoperability rather than API-only access; transparent and predictable cost structures; human override capability with logged, explainable, and reversible AI actions; cloud, hybrid, and on-prem deployment flexibility; and multi-agent orchestration with continuous integration of emerging open standards.

**The critical test question:** If a breakthrough model is released next month, how fast can we use it without reengineering our workflows?

## Section 8: Why Avaya Is the Strategic Choice

**First-to-Market with MCP.** Avaya did not just participate in the MCP conversation. It initiated it. By being first-to-market with MCP as a native capability of Avaya Infinity, Avaya created not a feature but a foundation. While competitors announce MCP support in limited, cloud-only, contact-center-only forms, Avaya's MCP is embedded, extensible, and enterprise-wide, spanning cloud, hybrid, on-prem, and sovereign cloud deployments.

**Foundational Leadership.** Avaya did not just adopt open protocols; we joined the Agentic AI Foundation to actively shape the future of AI interoperability alongside the world's leading technology creators. By choosing Avaya Infinity, enterprises are not just buying a platform. They are tapping into a massive, foundation-led ecosystem that outpaces the R&D capabilities of any isolated walled garden.

**Trusted, Explainable AI.** MCP provides observable logic paths that enterprise teams can inspect, adjust, and audit. Every action taken by an AI agent can be traced back to a defined policy. Memory boundaries are explicitly defined. Trust levels are protocol-enforced. This transforms AI from a risk factor into a governance asset, especially critical in regulated industries where explainability is mandated.

**Proven Enterprise Fit.** Avaya's architecture is tested in complex, regulated, multi-modal environments where failure is not an option. From Tandem Care models in healthcare to financial services orchestration to sovereign government

\* Frank Nagle and Daniel Yue, "The Latent Role of Open Models in the AI Economy," Working Paper, MIT Initiative on the Digital Economy, November 2025. Available on SSRN at [papers.ssrn.com/sol3/papers.cfm?abstract\\_id=5767103](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5767103).

deployments, Avaya supports multilingual orchestration, hybrid deployment flexibility, and integration with industry-specific tools. Performance at scale, cost predictability, and interoperability with existing systems reduce risk and accelerate time-to-value.

**Empowering, Not Replacing, Humans.** While competitors race toward full autonomy, Avaya champions Tandem Care, where AI agents handle intake, summarization, and real-time insights while human agents manage empathy, escalation, and judgment. MCP enables memory and context to persist across both roles, ensuring continuity. Rather than being sidelined by automation, agents become strategic participants in an AI-augmented workflow. The goal is not fewer humans. It is more empowered humans, backed by AI that understands them, supports them, and never replaces them.

## Appendix: RFP Questions to Evaluate AI Openness

### Orchestration and Agent Flexibility

1. Can your orchestration layer support AI agents from multiple vendors across different domains?
2. Can we use third-party AI agents alongside your native tools without custom development or additional licensing?
3. Is orchestration logic transparent, policy-driven, and editable by our teams?

### Governance, Auditability, and Control

4. Can we set policy-defined trust boundaries between AI agents, memory, and execution layers?
5. Can every AI interaction be logged, traced, and audited across human and machine workflows?
6. Do we retain full control over orchestration logic, memory state, and agent permissions?

### Model Neutrality and Integration

7. Can we swap or add models without rewriting the orchestration layer?
8. Does the platform work across hybrid, on-prem, and sovereign cloud with consistent orchestration?
9. Can your AI layer integrate with both modern SaaS tools and legacy systems via open protocols?

### Context Continuity and Outcomes

10. Can the platform maintain persistent context across channels, agents, and systems without re-entry?
11. Is the memory system decoupled from orchestration logic and model execution?
12. Do AI copilots enhance rather than replace human agents through tandem collaboration?

Avaya Infinity with MCP was built from the ground up to answer yes to every one of these questions. The future of enterprise AI belongs to those who demand more.

## About Avaya

Avaya is a global enterprise software leader that helps the world's largest organizations and government agencies forge unbreakable customer connections. The Avaya Infinity® platform is built to unify fragmented experiences, equipping enterprises to evolve their contact centers into connection centers and strengthen relationships that create business value. Learn more at [www.avaya.com](http://www.avaya.com).

