WHITE PAPER

The 2026 AI Inflection Series - Chapter 15

Why MCP Becomes the Control Plane for Enterprise AI

> The next phase of enterprise AI will not be decided by model choice alone. It will be decided by who controls how agents connect, authenticate, act, log, and scale.

Model Layer
Where most enterprise AI investment is focused today

Protocol Layer
Where control, governance, and portability are actually decided

Systems Layer
ERP, CRM, databases, identity, cloud services, internal tools

By: Logan Sivanasen | 21 April 2026

EXECUTIVE SUMMARY

The Control Layer Is the Real Strategic Decision

Most leaders still discuss enterprise AI at the model layer. They compare vendors, copilots, latency, reasoning quality, and cost per token. That view is now too narrow. Model Context Protocol, or MCP, is the emerging standard that lets AI systems connect to external tools, data, and actions through a more consistent interface. Once agents start touching real systems, the real issue shifts from intelligence to control.

That is why MCP matters. It brings structure to the connection layer between AI and enterprise execution, moving the conversation beyond prompt quality toward identity, governance, auditability, and scalability.

This chapter argues that MCP is becoming the control plane for enterprise AI. Not because it replaces every API or proprietary connector, but because it gives enterprises a cleaner way to standardize access, enforce policy, and reduce switching friction as the model landscape changes. The strategic question is no longer only which model sits on top. It is who governs the protocol layer underneath it.

Three Core Propositions

The model gets the headlines.

The protocol decides who stays in control.

The control plane decides whether AI scales safely.

Open Standard

Donated to Linux Foundation AAIF, December 2025

Enterprise Roadmap

Governance, identity, gateways, and audit trails named explicitly in March 2026

Vendor Adoption

Microsoft, Cloudflare, Block, and Google building actively around the protocol

Strategic Inflection

From developer tooling to enterprise infrastructure in under 18 months

CHAPTER 15 · SECTION 1

Leaders Think They Are Buying Models. They Are Inheriting an Interface Layer.

The Model-Centric Buying Assumption

The first mistake most enterprises make is assuming AI scale is a model problem. It is not. Model quality still matters, but the moment an agent needs access to files, databases, tickets, identities, workflows, or business logic, the problem shifts. The enterprise no longer manages a chatbot. It manages a live interface between intelligence and execution.

MCP formalizes that interface through a host-client-server architecture. Hosts connect to servers through dedicated clients, and the protocol creates a standard way for tools, resources, and prompts to move into the agent environment. In business terms, it means the path between AI and enterprise action is starting to standardize.

The Control-Centric Architecture Reality

Standardization changes power. In the old model, each vendor built its own connectors, auth patterns, and runtime assumptions. That model produced duplication, fragmentation, and weak portability. In the new model, the enterprise starts to gain leverage at the protocol layer.

The winners will not be the firms with the flashiest demo. The winners will be the ones that make the connection layer governable, observable, and replaceable. That distinction is not philosophical. It has direct implications for procurement, architecture decisions, and long-term vendor negotiating position.

ⓘ The enterprise does not just choose a model. It chooses an integration architecture that will outlive any single model vendor. The protocol layer is where that durability is built or lost.

CHAPTER 15 · SECTION 2

What MCP Is, in Business Language

MCP is the standard connection layer between AI applications and external systems. Anthropic defined it as an open standard for secure, two-way connections between AI-powered tools and data sources. In simple terms, MCP gives enterprises a common way to expose capabilities to agents without rebuilding one-off integrations for every model, assistant, or workflow.

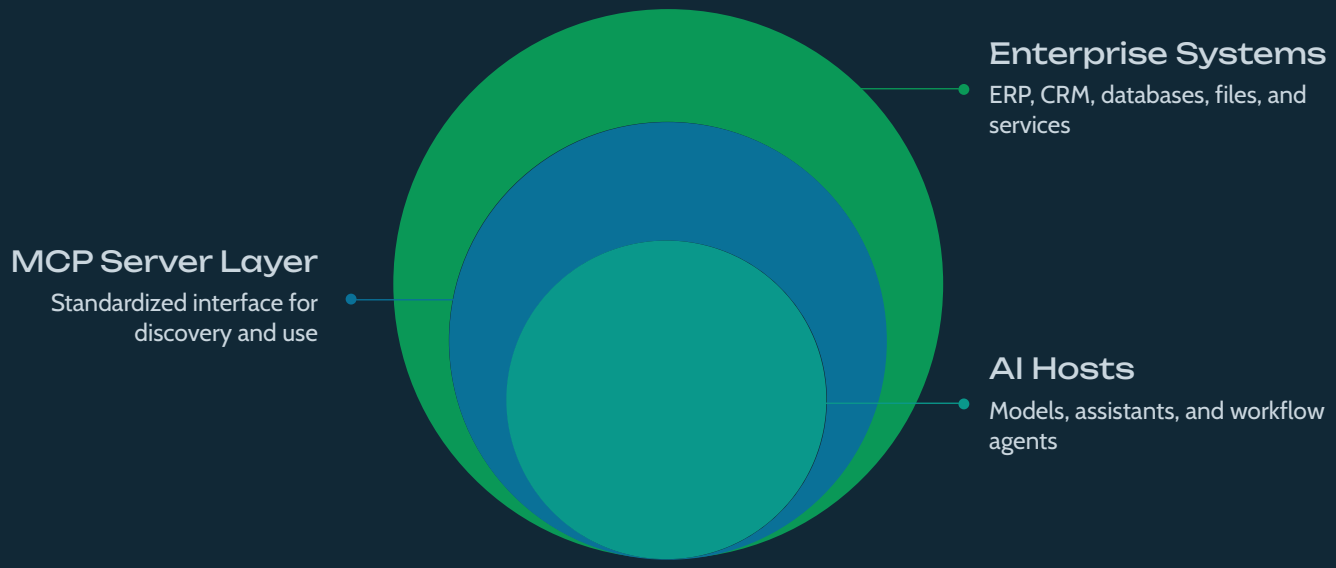

That does not mean MCP replaces APIs. The underlying systems still run on APIs, databases, services, and internal logic. MCP sits above them as a standardized interface for discovery and use. That distinction matters. APIs define the back-end capability. MCP defines how AI systems discover, understand, and call that capability in a more uniform way.

What MCP Provides

- A common discovery mechanism for tools, resources, and prompts
- Standardized connection patterns across different model providers
- A stable interface that survives model swaps without full rewiring
- A foundation for identity, policy, and audit layers to attach

What MCP Does Not Replace

- The underlying APIs and database interfaces that power enterprise systems
- Identity and authorization infrastructure already in place
- Governance policy and operating procedures for sensitive data
- The security review process for new service integrations

Why MCP Matters Now: A Strategic Timeline

Timing is what turns a technical standard into a strategic one. The transition from developer experiment to enterprise infrastructure happened in under 18 months. Each milestone below represents a structural shift, not an incremental product update.

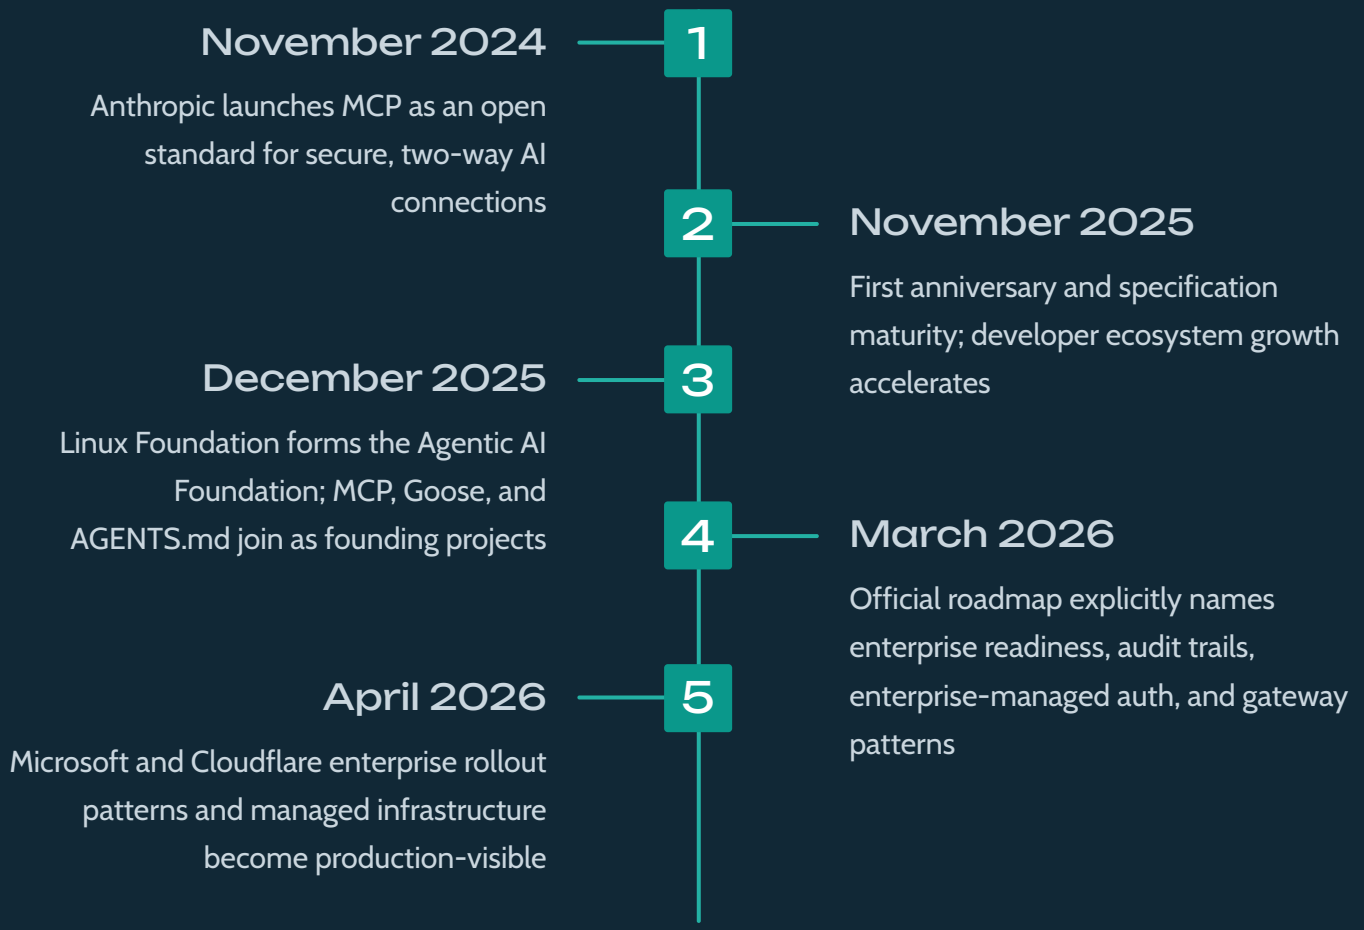

The Linux Foundation move matters strategically. It shifts MCP from one company's innovation into shared ecosystem infrastructure under neutral governance. Enterprises evaluating long-term standards adoption can now treat MCP the way they treat HTTP, OAuth, or OpenID Connect: as infrastructure that will not be owned or withdrawn by a single commercial actor.

The December 2025 Linux Foundation formation is the single most important governance event in the MCP timeline. It removes single-vendor dependency risk and positions MCP for enterprise standards committees and procurement frameworks.

The Old World: AI Without a Shared Protocol Layer

Before: Fragmented by Default

Before a shared protocol layer, enterprise AI usually followed a messy path. Every assistant needed custom connectors. Every tool exposed actions in a different shape. Auth rules varied by product. Audit trails lived in different places or did not exist. Vendor switching meant new integration work. Security teams lost line of sight once agents started calling mixed systems across business units.

That is not a temporary inconvenience. It becomes structural drag at scale. Each new agent deployment adds integration debt. Each model switch triggers a rebuild. Each audit request exposes gaps. The cost compounds quietly until a security incident or a compliance review makes it impossible to ignore.

After: Governance Gets a Foundation

Microsoft's governance guidance now makes this clear in plain terms. It calls for organizational accountability, a single registry, a unique identity for every agent, consistent policy enforcement across agent platforms, and continuous visibility into agent activity and tools.

That language matters because it shows where enterprise practice is heading. The market is moving from "let teams experiment" to "register, govern, and observe every agent pathway." MCP fits this shift because a standard connection layer gives governance somewhere stable to attach. Registry discipline, policy enforcement, and audit logging all become far more tractable when the connection layer follows a shared specification.

Scattered Connectors

Custom integration code per vendor, per team, per deployment

Duplicated Auth

Each product manages its own credential and permission model

Unclear Logs

Audit trails fragmented across products with no unified view

No Portability

Every model swap triggers full integration rebuild and revalidation

CHAPTER 15 · SECTION 5

From Integration Layer to Control Plane

A control plane governs how systems discover resources, authenticate requests, apply policy, observe behavior, and enforce change. That is where MCP is heading. The official roadmap now points toward end-to-end visibility, SSO-integrated enterprise-managed auth, defined gateway and proxy behavior, and portable configuration across clients. Those are classic control-plane concerns, not developer tooling concerns.

Agent Host

Models and user-facing interfaces

MCP Client

Standardized connection initiator within the host

Gateway / Registry

Central enforcement, logging, and approval layer

Authorization / Identity

Enterprise IdP, scoped tokens, consent gates

MCP Servers

Approved capability endpoints with metadata contracts

Enterprise Systems

ERP, CRM, files, databases, services, internal tools

The Phase Shift in Enterprise AI Leadership

In the first phase of enterprise AI, leaders focused on intelligence. Which model reasons best. Which copilot saves the most time. Which use case demonstrates the clearest ROI. Those are still valid questions. But they are increasingly insufficient as agents begin acting inside live business systems.

In the next phase, the competitive advantage will go to enterprises that focus on governable action. The model will keep changing. The connection layer will need to stay stable enough for identity, approvals, observability, revocation, and policy. That is why the protocol layer matters more than many leaders realize today.

> The model speaks. The protocol decides what the model is allowed to touch.

CHAPTER 15 · SECTION 6

Why Protocol Matters More Than Model Choice Over Time

Models will keep changing faster than enterprise systems do. Enterprises do not rebuild ERP, CRM, compliance pathways, identity systems, and core data stores every quarter. They do switch model providers, orchestration logic, and user-facing interfaces. That makes portability a strategic issue, not just a technical convenience.

Anthropic's framing of MCP as a universal way to connect AI systems to external tools exists for this reason. Standardization lowers duplicated integration effort and reduces dependence on one model or one interface pattern. That is not marketing language. It is a direct commercial benefit that compounds over time as the number of deployed agents grows.

The business implication is straightforward. If the enterprise connection layer stays proprietary and fragmented, every model shift becomes a rebuild tax. The integration work done for Model Provider A does not transfer to Model Provider B. The compliance review done for one agent runtime has to be repeated for the next. That is not a hypothetical cost. It is the current reality for enterprises that moved fast without protocol discipline.

If the connection layer becomes more standardized, the enterprise keeps more control while model vendors compete above it. That is where staying power sits. Not in betting on one model forever, but in owning a governed and portable action layer beneath changing models.

High Volatility

Model layer: providers, versions, reasoning patterns, pricing

Medium Volatility

Orchestration and interface layer: copilots, agents, UX patterns

Lower Volatility

Protocol and policy layer: where strategic control compounds

Stable Foundation

Enterprise systems: ERP, CRM, identity, data, compliance

ⓘ Strategic control rises as you move downward in the stack. Enterprises that invest in the protocol layer retain leverage regardless of which model wins.

CHAPTER 15 · SECTION 7

The Security Battle Starts Here

The protocol that enables scale also expands the attack surface. The failure modes are already documented. Microsoft's February 2026 guidance names them directly: tool poisoning or shadowing, silent swaps in tool metadata, and environments with no strong sandboxing around edits or code execution. These are not theoretical vulnerabilities. They are live design issues that enterprise security teams need to address before agents reach production scale.

Known Failure Modes

Tool Poisoning and Shadowing

Malicious or compromised tools masquerading as approved capabilities within the MCP server environment

Silent Metadata Drift

Tool descriptions, endpoints, or behavioral contracts changing without triggering a security review

Proxy Consent Failure

Confused deputy problems in MCP proxy servers when static client IDs, dynamic registration, and consent cookies mix without per-client controls

Over-Broad Tokens

Long-lived or excessively scoped credentials granting agents access beyond operational need

Hidden Third-Party Egress

Agents routing data or requests to undisclosed external endpoints through poorly reviewed MCP servers

Operating Response Requirements

Gateway Placement

Place remote MCP servers behind an API gateway; never allow direct unmediated agent-to-server connections in production

Least-Privilege Tokens

Use short-lived, minimally scoped tokens for every agent action; make revocation immediate and automatic

Metadata Snapshotting

Snapshot tool metadata at connect time and compare continuously against approved contracts to detect silent changes

Egress Pinning

Restrict outbound destinations to an approved allowlist; flag unexpected destinations automatically

Pre-Publication Review

Complete security, privacy, and responsible AI reviews before any MCP server is published or promoted to production

MCP's official security guidance warns specifically about confused deputy problems in proxy configurations. Every proxy server, every registry entry, and every metadata change path needs explicit policy review. Default trust assumptions will create real vulnerabilities at enterprise scale.

CHAPTER 15 · SECTION 8

Identity and Authorization Become Non-Negotiable

MCP's authorization guidance follows OAuth 2.1 patterns and strongly recommends authorization whenever servers expose user data, admin actions, or enterprise-sensitive operations. The official enterprise-managed authorization extension goes further. It introduces a delegated flow where the enterprise identity provider acts as the intermediary between client and server, maintains a registry of approved MCP servers and policies, and supports single sign-on through existing identity systems.

That is a major signal. Identity is moving from an integration afterthought into first-class MCP governance. The enterprise-managed authorization extension is not optional infrastructure for mature deployments. It is the baseline posture for any enterprise deploying agents at scale. Without it, every MCP server becomes a potential weak point where identity discipline breaks down.

Microsoft is moving in the same direction with purpose-built velocity. Entra now describes itself as the identity control plane for securing AI systems and positions agent identities as purpose-built constructs for authentication, authorization, governance, and audit at enterprise scale.

Identity Governance Failure Modes to Avoid

Agent Sprawl

Agents deployed without registered identities, owners, or audit records

Over-Privileged Agents

Credentials scoped wider than any specific task requires

Weak Ownership

No named human accountable for agent behavior, access, or revocation

Shadow Integration

MCP without enterprise identity discipline reproduces the shadow IT problem at agent scale

CHAPTER 15 · SECTION 9

Gateways, Registries, and Portals Are Where Control Will Sit

The 2026 roadmap points directly at gateway and proxy patterns, authorization propagation, session semantics, and configuration portability. Microsoft's operational guidance recommends API gateways as stable choke points for authentication, authorization, validation, rate limiting, and logging. This is not an emerging best practice. It is the direction the entire market is signaling simultaneously.

Central Registry

Azure API Center now supports maintaining an inventory of local and remote MCP servers, giving stakeholders a single discovery portal with governance metadata attached to every entry.

Managed Gateways

Microsoft recommends API gateways as the primary enforcement layer. Cloudflare has launched managed remote MCP servers with OAuth, DLP scanning, policy inspection, and access logging behind a single endpoint.

Curated Portals

Cloudflare's MCP server portals centralize multiple servers behind one endpoint with access policies, logging, and optional DLP. Enterprises can curate tool exposure rather than allowing broad direct access.

Approval Workflows

Serious enterprises will need the ability to pause, re-review, and revoke server access when tool capabilities change. The registry-plus-gateway pattern makes that operationally tractable at scale.

ⓘ Cloudflare's portal design, Microsoft's inventory-first governance model, and the MCP roadmap all converge on the same conclusion: the control plane will sit in managed intermediary layers, not in unmanaged direct agent-to-server access.

Commercial Proof Points: The Market Is Already Building Around This Layer

The enterprise infrastructure case for MCP is no longer theoretical. Three major commercial operators are building production systems around the protocol layer right now. Their choices reveal where the market is heading before the analyst consensus catches up.

Microsoft

Copilot Studio supports extending agents with MCP tools and resources. Microsoft Foundry lists MCP and A2A among enterprise controls and treats remote and local MCP servers as first-class tool types. Windows now includes an On-device Agent Registry for secure discovery and management of MCP connectors with admin controls, containment, and auditability. Microsoft also ships product-specific MCP surfaces including SQL MCP Server and Dynamics 365 ERP MCP Server, both aimed at controlled access to live business systems.

Cloudflare

Cloudflare runs a catalog of managed remote MCP servers with OAuth support and now offers portals that centralize multiple servers, expose curated tools, and route traffic through logging and optional DLP inspection. That is not hobbyist tooling. That is enterprise control infrastructure surfacing around agent connectivity at internet scale.

Block

Block offers one of the strongest live operating examples. At its 2025 Investor Day, Block stated Goose integrates with many services through MCP, had already connected approximately 150 services, and was being used across functions to analyze internal sales data, build dashboards, manage ticketing work, create documents, edit video assets, and prepare diagnoses and code changes for bug reports. That is a large commercial operator using a general agent with a growing MCP-connected service ecosystem to automate internal work at scale.

MCP and A2A Solve Different Problems

Leaders need a clean mental model here. MCP and A2A are complementary, not interchangeable. Google's A2A announcement explicitly states that A2A complements MCP. Many strategy discussions still collapse every open agent standard into one vague category. That is a mistake. Enterprises need both boundaries clear to make sound architectural and governance decisions.

MCP: Agent-to-System

Primary Purpose: How an agent connects to tools, APIs, data sources, and prompts in a standardized, governable way.

Core Governance Concern: What can this agent access, under what identity, with what scope, and with what audit trail?

Enterprise Example: Agent securely queries the ERP system, retrieves customer data from CRM, reads from a file store, and triggers a ticketing workflow, all through registered, policy-gated MCP servers.

A2A: Agent-to-Agent

Primary Purpose: How multiple agents coordinate work across systems and organizational boundaries in structured, interoperable ways.

Core Governance Concern: How do agents exchange tasks, delegate authority, and maintain accountability across multi-agent workflows?

Enterprise Example: A research agent delegates a subtask to a writing agent, which routes output to a review agent before a final delivery agent publishes the result, all with verifiable task handoffs.

When the question is "How do multiple agents coordinate work?" think A2A. When the question is "How does this agent securely access ERP, SQL, files, or workflows?" think MCP. These are different architectural decisions, different governance controls, and different risk surfaces. Conflating them produces incomplete designs in both dimensions.

Distinct Standards

MCP and A2A address different layers of the agent architecture stack

Complementary Use

Production enterprise systems will use both, with clear boundaries between them

Different Risk Surfaces

Security and governance models for each standard address distinct threat vectors

Separate Procurement

Evaluate each standard independently against your enterprise use cases and risk posture

CHAPTER 15 · SECTION 12

The Enterprise Operating Model for MCP

Enterprises should implement MCP governance through six structured moves. Each move builds on the last. None is optional for organizations deploying agents into live business systems at any meaningful scale.

Establish a Trusted Registry

Every server should have an owner, purpose, risk tier, version history, and approval record before it receives production traffic. Microsoft's guidance is direct on this point: you cannot govern what you cannot see. Azure API Center and Microsoft's broader agent governance model reinforce this registry-first pattern. Build the registry before scaling deployments.

Put Identity at the Center

Require a distinct identity for every agent and use enterprise-managed authorization for server access wherever possible. Scope tokens tightly and make revocation immediate. Both MCP's auth guidance and Microsoft's Entra material support this direction. Agent identity is not a feature to add later. It is a prerequisite for any serious governance posture.

Route Traffic Through Gateways or Portals

Use them for authentication, policy enforcement, rate limiting, logging, and inspection. Direct unmediated agent-to-server connections are appropriate only for the lowest-risk read-only scenarios in controlled environments. Microsoft and Cloudflare both demonstrate this pattern in production deployments.

Separate Read From Write

Low-risk retrieval flows should never share the same approval posture as destructive actions, data mutation, or code execution. Microsoft's own MCP governance pattern relies on consent gating for side-effecting actions and enforces confirmation before writes. That separation needs to be explicit in your classification scheme, not assumed.

Standardize Observability

Log what the agent requested, what the server did, what data was touched, what token scope applied, and whether the action matched the approved contract. The MCP roadmap explicitly calls for end-to-end audit trails in forms enterprises feed into existing logging and compliance pipelines. Observability is the foundation of post-incident review, compliance evidence, and continuous improvement.

Build a Re-Review Loop for Capability Drift

Tool descriptions, prompts, server endpoints, and metadata will change after initial approval. The enterprise needs automated checks for drift and clear pause-and-review pathways when the change touches riskier actions. Microsoft's connect-time contract snapshotting is a strong operational pattern. Build the review trigger before you discover why you needed it.

CHAPTER 15 · SECTION 13

The MCP Scorecard Leaders Should Track

If MCP is becoming infrastructure, leaders need a way to manage it like infrastructure. The following scorecard organizes measurement into five operational dimensions. These measures align with the direction already visible in the MCP roadmap and Microsoft's governance material, and they give senior leaders a way to manage protocol maturity without falling into shallow adoption metrics.

Coverage

- Share of agent tool access routed through approved MCP paths
- Share of MCP servers registered in the central inventory
- Share of production MCP traffic behind a managed gateway or portal

Control

- Share of servers using enterprise-managed authorization
- Share of high-risk tools with explicit consent gating
- Share of servers with named owner, runbook, and retirement path

Security

- Metadata drift incidents detected per review cycle
- Blocked unsafe actions by server and action class
- Session anomalies and policy violations by owner

Portability

- Time to onboard a new compliant MCP server from intake to production
- Time to move a workload from one model host to another without rebuilding tool access
- Share of configuration reused across multiple host environments

Operations

- Logging completeness rate across registered servers
- Mean time to approve a new server through the intake process
- Mean time to revoke or rotate access following a trigger event
- Time from drift detection to completed re-review

These five dimensions give the board, CISO, and CTO a governance view of MCP maturity that sits alongside standard IT infrastructure reporting. Treat protocol governance as infrastructure management, not as a developer activity.

5

Scorecard Dimensions

Coverage, Control, Security, Portability, Operations

6

Operating Moves

Registry, Identity, Gateway, Action Controls, Observability, Drift Management

150

MCP Integrations

Block's live MCP-connected service count as of 2025 Investor Day

18mo

Time to Infrastructure

From Anthropic launch to Linux Foundation neutral governance

CHAPTER 15 · SECTION 14

Sector Implications

MCP's commercial opportunity and governance burden distribute differently across sectors. The protocol does not remove management complexity. It raises the standard for it. Each sector faces a distinct combination of upside and risk surface as agents move from pilot to production.

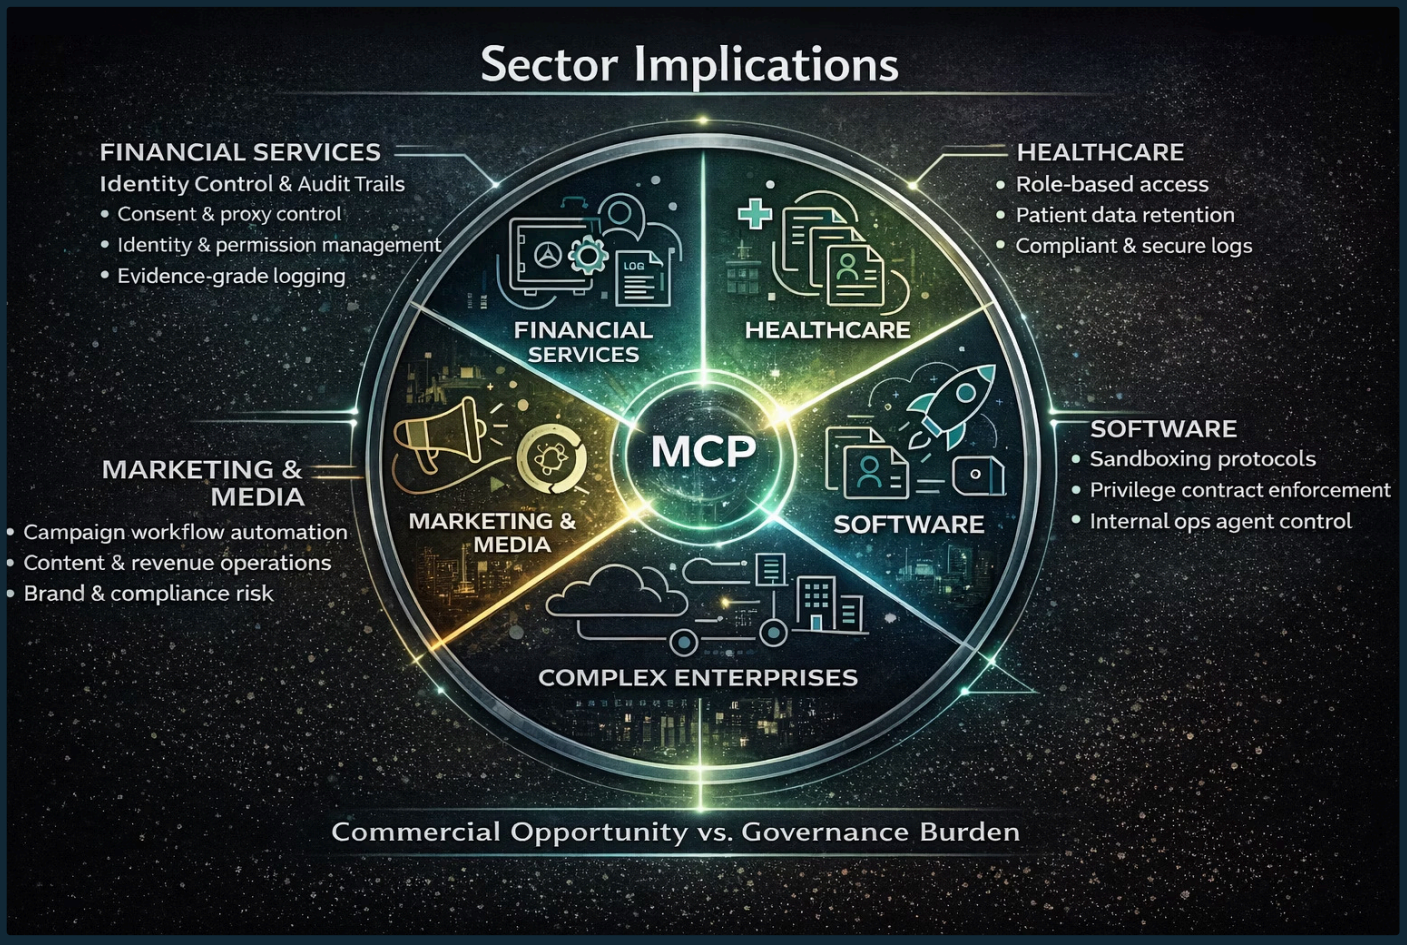

The common thread across all sectors is that governance maturity must scale in direct proportion to agent reach. Every gain in agent connectivity increases the need for registry discipline, identity control, gateway enforcement, and observable action trails. Sectors that move fast on deployment without moving equally fast on governance will face the same structural debt that shadow IT created, but with significantly higher velocity and broader blast radius when controls fail.

Where the Commercial Opportunity Is Clearest

- Connector simplification across large multi-system enterprise environments
- Workflow automation in marketing, operations, and internal developer tooling
- Model portability that reduces rebuild costs as the provider landscape evolves
- Audit readiness for regulated industries with existing logging infrastructure

Where the Governance Burden Is Highest

- Financial services: data movement, consent chains, and audit evidence requirements
- Healthcare: patient data boundaries and role-based access enforcement
- Any sector with destructive write actions or code execution in agent scope
- Multi-cloud enterprises with heterogeneous identity and policy infrastructure

LEADERSHIP TAKEAWAY

The Question Has Changed

Most AI strategies still center on models, copilots, and use cases. That is yesterday's framing. The next strategic decision sits one layer lower. Who controls the connection layer between intelligence and execution. Who approves access. Who owns the registry. Which pathways stay visible. Which actions trigger consent. Which architectures preserve portability when models change.

> In the first phase of enterprise AI, leaders asked which model to use. In the next phase, the better question is who controls the protocol that lets the model act.

MCP sits at the center of those questions because the market is already building the supporting layers around it. Neutral governance through the Linux Foundation. Enterprise-focused roadmap priorities named explicitly in March 2026. Identity extensions, gateways, registries, and managed portals shipping from Microsoft and Cloudflare. Product-native MCP surfaces in ERP, CRM, and developer tooling. Block demonstrating a live operating model at commercial scale. The infrastructure is not speculative. It is here.

The enterprise leaders who will have the clearest position in 2027 are the ones who start treating the protocol layer as infrastructure today. That means building registries before they are needed. Enforcing identity before agents proliferate. Establishing gateway patterns before audit requests arrive. Running the scorecard before the board asks for it. The window for getting ahead of this is still open. It will not stay open indefinitely.

Own the Registry

Build the central server inventory before agent sprawl makes it reactive

Enforce the Protocol

Make MCP governance a first-class infrastructure program, not a developer afterthought

Measure What Matters

Run the five-dimension scorecard as standard infrastructure reporting from day one

CHAPTER 15 · DOWNLOADABLE ASSETS

Supplementary Tools and Templates

The following assets extend the framework in this chapter into operational form. Each is designed for direct use by enterprise architecture, security, and governance teams. Formats and purposes are described below for procurement and distribution planning.

MCP Readiness Scorecard

A practical executive scorecard to assess whether your organization is ready to govern MCP at scale across registry maturity, identity, gateway control, logging, and risk oversight.

MCP Control Plane Canvas

A one-page strategic canvas to map the full MCP operating environment, including hosts, clients, servers, gateways, identity layers, approval points, and audit paths.

Approved MCP Server Intake Template

A standardized intake document to review, assess, and approve any new MCP server before it enters pilot or production use.

MCP Risk Classification Matrix

A decision tool to classify MCP-connected tools and servers by risk level, action type, data sensitivity, and required control measures.

MCP Governance Policy Starter

A policy starter document to help organizations define ownership, approval rules, access standards, logging expectations, review cadence, and revocation processes for MCP environments.

MCP vs A2A Executive Explainer

A concise leadership explainer that clarifies the difference between MCP and A2A, helping decision-makers understand where each standard fits within the enterprise AI stack.

90-Day MCP Adoption Plan

A focused implementation roadmap that helps leaders move from experimentation to controlled adoption across inventory, identity, gateways, certification, and observability.

CHAPTER 15 · REFERENCES

Citations and Sources

All references below correspond to claims, frameworks, and direct language cited in this chapter. Where available, sources are linked to their original publication locations.

Anthropic and MCP Specification

Anthropic. "Introducing the Model Context Protocol." November 25, 2024.
Model Context Protocol. "Architecture Overview."
Model Context Protocol. "The 2026 MCP Roadmap." March 9, 2026.
Model Context Protocol. "Understanding Authorization in MCP."
Model Context Protocol. "Enterprise-Managed Authorization."
Model Context Protocol. "Security Best Practices."

Microsoft Governance and Security

Microsoft. "Protecting AI conversations at Microsoft with Model Context Protocol security and governance." February 12, 2026.
Microsoft Learn. "Governance and security for AI agents across the organization."
Microsoft Learn. "What is Microsoft Entra Agent ID?"
Microsoft Learn. "Microsoft Entra security for AI overview."
Microsoft Learn. "Inventory and Discover MCP Servers in Your API Center."
Microsoft Learn. "Model Context Protocol on Windows overview."
Microsoft Learn. "SQL MCP Server overview."
Microsoft Learn. "Use Model Context Protocol for finance and operations apps."
Microsoft Learn. "Agent tools overview for Microsoft Foundry Agent Service."
Microsoft Learn. "What is Microsoft Foundry?"

Linux Foundation and Ecosystem Governance

Linux Foundation. "Formation of the Agentic AI Foundation." December 9, 2025.

Cloudflare Infrastructure

Cloudflare Docs. "Cloudflare's own MCP servers."
Cloudflare Docs. "MCP server portals."
Cloudflare Docs. "Securing MCP servers."

Google and Agent Standards

Google Developers Blog. "Announcing the Agent2Agent Protocol."

Block Operating Model

Block Investor Day 2025 Transcript. Goose and MCP operating model references.
Block Engineering. "How We Red-Teamed Our Own AI Agent."

The 2026 AI Inflection Series is published for enterprise technology and security leaders, CTOs, AI platform leads, and strategic investors. Chapter 15 covers the Model Context Protocol as emerging control-plane infrastructure for enterprise AI. All citations reflect sources available as of publication date.