

Securing The Model Context Protocol (MCP): A Dual-Axis Survey with a Mitigation-Oriented Threat Taxonomy

Anonymous authors

Paper under double-blind review

Abstract

Agentic AI systems are increasingly adopting the Model Context Protocol (MCP) to integrate external tools and data sources. However, MCP’s extensibility and multi-agent design introduce a broader and more complex attack surface. This exposes these systems to a diverse set of security risks. Prior research provides only a partial view of these risks, focusing on either temporal aspects (MCP system lifecycle) or spatial aspects (MCP components) in isolation, without capturing how the associated threats emerge, evolve, and propagate across agentic workflows. This missing linkage obscures causal chains across components and lifecycle phases, hindering timely root cause localization and effective prevention of cross-phase threat cascades. To address this gap, we present a spatio-temporal aligned taxonomy that maps each threat to both its affected MCP component and the lifecycle phase in which it arises. This dual-axis view highlights where threats occur, when they emerge, and how they propagate across stages of deployment and operation. Grounded in this taxonomy, our survey catalogs over 50 MCP-specific threats and classifies each by lifecycle phase and affected component, providing, to our knowledge, one of the broadest phase-by-component mappings of MCP security risks to date. Each threat is also cross-referenced to STRIDE (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege) and MAESTRO (Modeling Adversarial Events for Systematic Threat Representation and Organization) to align with widely adopted threat-modeling frameworks for agentic AI, covering core security properties (via STRIDE) and layers specific to agentic AI (via MAESTRO). We then pair each taxonomy category with actionable controls, explicit verification checks, and runtime signals, supporting structured control prioritization, verification checks, and runtime monitoring across design-time and runtime. We further validate this control-oriented view through a compact benchmark of verifiable checks in a real MCP deployment, showing how representative failures can be detected, localized, and tied to concrete audit artifacts. Together, the taxonomy and its control mappings provide an evidence-informed framework for organizing MCP security priorities, verification checks, and reviewable security evidence across the MCP lifecycle.

The benchmark implementation, scenario definitions, verification scripts, and reproduction instructions are available at <https://anonymous.4open.science/r/mcp-security-benchmark-22C9/>.

1 Introduction

The Model Context Protocol (MCP) is an open standard introduced by Anthropic to enable large language models (LLMs) to interact with external tools, data sources, and computational services in a standardized and extensible manner Anthropic (2024a). It defines a client–host–server architecture and uses JSON-RPC 2.0 (JavaScript Object Notation-Remote Procedure Call) for all messages, exposing a unified interface to invoke remote tools, access files, query databases, and call application programming interfaces (APIs) Li et al. (2025); Kumar et al. (2025b). This JSON-RPC-based interoperability is what makes MCP-based agentic systems flexible and broadly applicable, enabling complex tasks with real-time access to external resources. This flexibility arises because MCP encodes tool invocations and data access as JSON-RPC 2.0 messages over

a unified logical channel between client and server that supports capability discovery and structured argument passing, thereby decoupling the host and LLM runtime from service-specific APIs and integrations. However, because this interoperability relies on cross-boundary calls and shared state, it also creates a broader, more intricate attack surface, elevating exposure to threats including prompt injection, tool poisoning, credential leakage, colluding agent and supply-chain compromise Ferrag et al. (2025); Narajala & Habler (2025); Hasan et al. (2025). Concretely, the shape of that threat surface follows the MCP architecture: the client-side LLM host, the JSON-RPC communication layer, tool-provider servers, and mechanisms for continuous updates and tool registration. Each component opens a distinct pathway for misuse or failure, and thus a potential vector for exploitation. Compounding these component-level vectors, MCP deployments are often dynamic and multi-agent, combining human-in-the-loop and autonomous interactions across decentralized infrastructures. As a result, the likelihood and impact of cross-component threat propagation across deployment stages increase Fang et al. (2025); Hasan et al. (2025). Motivated by these threats, a growing body of MCP-specific security research has emerged, spanning surveys, empirical audits, attack demonstrations, and defense frameworks Yang et al. (2025); Song et al. (2025); Narajala & Habler (2025); Kumar et al. (2025b); Brett (2025); Hou et al. (2025); Radosevich & Halloran (2025); Guo et al. (2025); Hasan et al. (2025); Li & Gao (2025); Zhang et al. (2025). Benchmarking efforts have further strengthened this line of research by evaluating agent safety across real-world MCP servers and diverse attack types, and by introducing a large-scale benchmark for tool-use competency in realistic multi-server MCP workflows Zong et al. (2025); Bandi et al. (2026). Beyond benchmarking, the literature has also expanded at both the protocol and tool levels. Maloyan & Namiot (2026) analyze architectural weaknesses in MCP itself, while Li et al. (2026a) show how implicit tool poisoning in metadata can redirect agents toward high-privilege tools even without invoking the poisoned tool. Taken together, this growing body of work has expanded the evaluation and understanding of MCP-specific security failures and realistic MCP agent behavior, yet most studies still examine only a single axis, either lifecycle stages or architectural components, leaving a gap for a unified spatio-temporal taxonomy that links when threats arise to where they propagate across MCP workflows.

Hou et al. (2025) proposed one of the earliest structured approaches to MCP security by categorizing threats based on the distinct lifecycle phases of an MCP deployment, namely the creation, operation, and update stages. Their lifecycle threat mapping and mitigation analysis focuses on mapping different types of vulnerabilities to these temporal categories. For example, they treat name collisions during server registration, installer spoofing, and backdoor attacks from code-integrity failures as creation-phase threats. In operation, they highlight tool-name conflicts, slash-command overlap, and sandbox escapes. In the update phase, they discuss privilege persistence after authorization changes, redeploying vulnerable versions, and configuration drift. While this lifecycle taxonomy provides a useful temporal abstraction, it lacks the granularity needed to associate threats with specific system components, which limits its utility in tracing the structural impact of a vulnerability. In contrast, Narajala & Habler (2025) propose an architectural perspective on MCP security, organizing threats around the core components of an MCP system: the Host, the Client, and the Server. Within the server, they further analyze sub-modules, including tools, prompts, and resources. Their model maps security threats to these architectural elements, covering issues such as tool poisoning, prompt injection, insecure communication, and hosting environment vulnerabilities, and paired them with mitigation strategies, including scoped access controls, communication hardening, and operational safeguards. Complementing this architecture-centric view, Li & Gao (2025) present the first ecosystem-wide security analysis of real MCP deployments, decomposing the ecosystem into hosts, registries, and servers and empirically demonstrating how weak vetting and missing output verification pose a threat to MCP security. These component-driven approaches enable more localized threat identification by highlighting threats associated with specific elements, such as man-in-the-middle attacks on the JSON-RPC interface, insecure tool execution environments, or rogue tool provider servers. However, their analysis treats these components as static units, without linking them to lifecycle phases. As a result, it becomes challenging to contextualize when a given component-related threat might emerge or propagate during the lifecycle of a system. While all of the aforementioned works contribute significantly to structuring early discussions of MCP security, they also exhibit a shared limitation. They treat lifecycle stages and architectural components as separate silos, which produces a fragmented view of the threat landscape. Thus, it becomes challenging to capture how a vulnerability introduced during one phase of the lifecycle, such as creating a tool with malicious payloads, might later interact with various components like the LLM host, communication channel, or tool registry to




Figure 1: **Creation-to-runtime threat cascade in MCP**, showing how an over-permissive file tool configured during creation is exploited during operation to leak credentials and trigger follow-on tool and server actions, motivating a dual-axis mapping by component and lifecycle. Red arrows and crosses mark the attack path and violations, green arrows show expected interactions, and green dashed circles mark control points.

cause harm during execution. This disjointed framing also limits the ability to anticipate compound threats that evolve across phases, such as those involving delayed activation, staged escalation, or multi-component compromise.

More recently, Jing et al. (2025) introduced the Model Contextual Integrity Protocol (MCIP), a multidimensional scheme that maps runtime threats to axes such as phase (Config, Interaction, Termination), source, scope, violation type, and MAESTRO layer. MCIP offers valuable insights into runtime information-flow violations, particularly in capturing how contextual integrity can be compromised during tool invocation and result integration. While the taxonomy includes setup and termination lifecycle phases, the empirical focus leans toward interaction-phase behaviors and flow-based violations: the experiments target interaction threats, and the paper does not detail sandbox-escape mechanics or comprehensive software-supply-chain hardening beyond the taxonomy. In parallel with lifecycle and component analyses, Guo et al. (2025) emphasize an attack-type taxonomy and empirical evaluation with minimal architectural grounding. They do not map threats by lifecycle stage, and component coverage is partial. Instead, they develop an attack-centric taxonomy and release an MCP attack library (MCPLib), which catalogs MCP-specific attacks in four classes, namely direct tool injection, indirect tool injection, malicious-user attacks, and LLM-inherent attacks. Taken together, these strands advance the field but leave open how to connect attack behaviors to both specific components and specific points in the lifecycle. In practice, however, the most damaging issues often appear as concrete, cross-cutting failures in deployed MCP systems rather than neatly categorized attack types. Now, against this backdrop, recurring failure modes continue to surface in practice: agents over-trust tool descriptions, file-based attacks succeed disproportionately, credentials leak across tool boundaries, and local misconfigurations cascade across components at runtime Xing et al. (2025); Guo et al. (2025); Radosevich & Halloran (2025). For example, large-scale audits of open-source MCP servers report MCP-specific tool poisoning and widespread misconfigurations that persist in real deployments Hasan et al. (2025). These security failure patterns span both where attacks land (host, transport, server, and tool/resource/prompt sub-modules) and when they materialize (creation, operation, update).

Components × Lifecycle threat density							
	Pre-MCP	Creation / Registration	Connection / Initialization	Discovery	Invocation / Execution	Update / Maintenance	Shutdown
Host	1	5	0	7	24	8	3
Client	0	2	4	8	16	3	3
Server	0	7	4	5	15	10	3
Tools	0	2	0	6	6	2	2
Prompts	0	0	0	3	2	0	0
Data Layer	0	0	1	6	4	1	1
Protocol Layer	0	0	5	0	0	0	0
Auth	0	2	4	0	0	3	1

Figure 2: **Threat distribution across components and lifecycle stages.** Each cell reports the count of distinct risks from our threat corpus mapped to the corresponding component–stage pair. Risks tagged with multiple components or stages contribute to each applicable cell.

Without a framework that jointly captures components and lifecycle phases, developers, researchers, and security practitioners struggle to localize root causes, anticipate cross-phase propagation, and select enforceable controls at design time and runtime. For instance, in Fig. 1, an over-permissive file tool configured during creation is later exploited at operation via prompt injection, leading to credential leakage and enabling follow-on actions through other tools and servers; precisely the kind of creation-to-runtime cascade that is difficult to diagnose and contain without a unified view. For clarity, the Host and Client are combined, and the Pre-MCP, Initialization/Connection, and Shutdown phases are omitted, focusing on Creation, Discovery, Invocation/Execution, and Update. Now, a lifecycle-only perspective shows when issues occur (Creation → Invocation/Execution → Update) and supports coarse, phase-scoped checks (Creation: scope review; Execution: output redaction; Update: token rotation), but it cannot localize fixes to specific components or produce component-scoped evidence. A component-only perspective shows where the chain runs (Server: Tools → Transport → Host) and supports control placement (tool-scope lint at Server: Tools; mTLS/log redaction at Transport; revocation at Host), but it does not specify when each control must apply; there is no Creation gate (CI before deploy), no Invocation/Execution-time channel policy, and no Update-phase revocation checklist, so enforcement is non-temporal and audits are non-deterministic. In contrast, the dual-axis view provides both when and where, localizing controls to Creation × Server: Tools (directory allow-lists and scope lint), Invocation/Execution × Transport (JSON-RPC) (mTLS and log redaction on the channel), and Update × Host (token-revocation checks), making the component chain explicit, assigning ownership (tool-server, platform/transport, host teams), and enabling verifiable evidence per component (configuration invariants, transport policy assertions, revocation audits).

Guided by this need for a dual-axis map, we pose four research questions: (i) *What MCP-specific threats arise across core components and lifecycle phases?* (ii) *How do these threats emerge, evolve, and propagate, including the causal patterns that drive cross-phase cascades?* (iii) *Which design-time and runtime controls most effectively disrupt these paths, and how can their enforcement be verified?* and (iv) *How can MCP-specific threats be aligned with established threat-modeling frameworks to support consistent modeling and auditing?* In examining these questions, we find that existing approaches usually address only a restricted

segment of one axis (components or lifecycle phases), resulting in unaddressed threats not only on the alternate axis but also within the unexamined segments of the chosen axis. To address these limitations, this paper introduces a finer-grained taxonomy that assigns when and where each threat emerges. Our approach provides a spatio-temporal view of MCP deployments by systematically integrating architectural (spatial) and temporal dimensions, enabling precise attribution of when and where each threat emerges. Unlike prior work that isolates phases or focuses on specific classes of threats, our taxonomy covers major stages of the MCP lifecycle, from pre-deployment setup and tool installation to runtime execution and system shutdown, yielding a structured mapping of threats to both components and lifecycle stages. We visualize this coverage in Fig. 2, which lists components by rows and lifecycle stages by columns and shows the count of distinct risks in each cell. We omit Resources from the component axis in this figure because it rarely appears as the primary affected component or blast radius in our corpus. When resources are implicated, the effect is captured under the Data layer or the executing tool or server.

Building upon this foundation, our work makes three principal contributions:

1. **Dual-axis threat mapping.** We build a deployment-wide taxonomy of MCP-specific threats by (i) partitioning the attack surface into nine failure modes and (ii) cross-classifying each threat along orthogonal coordinates: failure mode, MCP components, lifecycle phases, and established threat-modeling frameworks. This dual-axis, multi-coordinate scheme localizes where and when each threat manifests and grounds it in widely adopted modeling frameworks; the full mapping is presented in Tables 1–3.
2. **Threat mitigation with verifiable enforcement.** Existing guidance is fragmented and seldom specifies which controls to apply where (component) and when (lifecycle phase), or how to verify that they are enforced. We address this by attaching enforceable controls to each threat category in our dual-axis taxonomy and by pinpointing the component and lifecycle phase at which each control should apply. For every category, we provide implementation patterns such as scoped capability bindings, privilege fences, provenance and integrity checks, and session or isolation guards, along with verification hooks in the form of machine-checkable evidence such as policy assertions, configuration invariants, and runtime signals. We also define pass or fail checks that can be integrated into CI or CD gates and runtime monitors. We further validate this control-oriented view through a compact benchmark of verifiable checks in a real MCP deployment, showing how representative failures can be detected, localized, and tied to concrete audit artifacts. Together, these elements turn high-level guidance into verifiable controls, evidence-producing checks, and concrete audit artifacts rather than description alone.
3. **Expanded threat coverage with operational safeguards.** Prior work leaves multiple individual MCP threats unaddressed or only loosely implied, including lack of verified server identity/PKI, session-isolation failures, agent handoff chains, orchestration drift, agent collusion, and others. We add these threats to the MCP threat surface and locate each by component and lifecycle phase within our dual-axis scheme. Their prevention and detection can then be handled using the proposed threat-mitigation strategy mapping for the corresponding threat category, which supplies testable detection targets and deployable controls.

In addition to these core contributions, we examine recent advances in security frameworks for MCP, identify key limitations of existing approaches, and highlight opportunities for future work. This includes addressing the need for automated threat detection, robust context validation, and secure multi-agent orchestration mechanisms. By offering a broad and operationally grounded taxonomy together with mitigation strategies, our work aims to strengthen the security foundation of tool-augmented AI systems. The remainder of this paper is organized as follows: Section 2 provides background on the MCP architecture and lifecycle, laying the foundation for the subsequent threat modeling. Section 3 describes our dual-axis taxonomy and the methodology used to develop it. It also systematically catalogs the identified threats and classifies them within the taxonomy. Section 4 presents mitigation strategies and security controls for MCP based on the proposed taxonomy. Section 5 provides a structured evaluation of verifiable checks in a compact MCP benchmark. Section 6 reviews emerging frameworks and protocol extensions in MCP security research. Section 7 outlines future directions, and Section 8 concludes the paper.

2 MCP Architecture and Lifecycle

The structure of an MCP deployment is governed by two complementary dimensions that operate in parallel: architecture and lifecycle. The architecture specifies **where** each component resides and **who** owns it. This includes a host that reasons, clients that communicate, servers that execute, and tools that act, data they access, and transport that connect them. The lifecycle defines **when** each stage of operation occurs, including setup and connection, creation and discovery, invocation and execution, update, and shutdown. Tracking only the “where” gives a static view of the system, allowing attacks to evolve over time and bypass controls that are not active in the right phase. Tracking only the “when” captures timing but not responsibility, allowing faults to move across components with no clear owner. Considering both where and when dimensions enables complete traceability across the system. Building on this view, we formalize the architecture and lifecycle as orthogonal coordinates for the taxonomy presented in Section 3.

2.1 MCP Architecture

The Model Context Protocol (MCP) follows a layered *client–host–server* architecture rather than a flat client–server design Model Context Protocol (2025f). In this structure, the host coordinates clients, manages links to servers, and provides the setting required for model reasoning and planning Singh et al. (2025). Each **MCP host** runs one or more **MCP clients**, which are protocol-level components instantiated by the host to communicate with particular **MCP servers**, and each client maintains a dedicated one-to-one connection to a single MCP server Model Context Protocol (2025g). In effect, this separation lets the host add new capabilities by spinning up additional client instances, each paired with a different server. Thus, the same client implementation can then be reused across many servers, with changes usually limited to configuration rather than client code Model Context Protocol (2025h); Microsoft (2025). In practice, this means the system can scale horizontally: an AI agent can extend its abilities simply by connecting to additional servers, each providing its own set of tools or resources. For example, a host can run one client instance that connects to a file management server and a second client instance that connects to a data analytics or scheduling server. The same client implementation is reused in both cases with configuration changes only.

In MCP the roles are cleanly separated, with the host handling reasoning and planning, the client managing communication sessions and enforcing protocol consistency, and the server executing tasks, as shown in Fig 3. Fig. 4 illustrates how MCP components coordinate during a request lifecycle. Working together, these layers form a coordinated pipeline in which the host decides what should happen, the client determines how to communicate it, and the server carries it out. So AI systems using MCP can scale efficiently while orchestrating multiple external tools and data sources without losing control or transparency Model Context Protocol (2025f). Once the host forms a request, the client interprets the intent, establishes a session with the chosen server, negotiates available capabilities, and issues structured method calls. The server completes the operations and returns results, which the client formats and delivers back to the host.

Building on this pipeline, the client and server communicate using JSON-RPC 2.0, a lightweight remote procedure call standard encoded in JSON that provides a structured, predictable message format Wigren & Conservancy (2020). Within MCP this communication is expressed as two layers that work together. The **data layer** defines the structure and meaning of JSON requests and responses, covers the flow of a session such as connection setup, capability discovery, and orderly shutdown, and names the core primitives that implementations rely on, including tools, resources, prompts, and notifications. The **transport layer** carries these JSON messages across different environments. Local deployments often use `stdio` for in process communication, while remote deployments use the Streamable HTTP transport over HTTPS so TLS provides confidentiality and integrity Hitachi Vantara (2023); Model Context Protocol (2025q). In higher assurance settings, mutual TLS can be required so both client and server authenticate before any data is exchanged Cloudflare Learning Center (2024). This separation keeps wire details independent from application semantics and allows one client to interoperate with many servers without changes to the protocol.

Within this communication model, the MCP server exposes callable tools, accessible resources, and reusable prompts, which together define its execution surface. **Tools** represent executable operations such as database queries, file actions, or REST calls that the client can invoke MCP Working Group (2025a); Fei et al. (2025). **Resources** act as structured data endpoints, including files, databases, or live streams, that can




Figure 3: **MCP architecture** - The *Host* encloses three MCP clients (C1-C3). Each client sends required control/data traffic (solid arrows) through its in-process Communication Layer (CL1-CL3) to a dedicated MCP server (S1-S3), which advertises domain-specific *Tools/Resources/Prompts* blocks (TRP1-TRP3). Optional interactions (dashed arrows) include capability discovery and authentication via the external Authorization Server (AS).

be read or updated depending on authorization scope. **Prompts** function as reusable templates that standardize interactions between language models and external systems, encoding reasoning patterns or output structures to maintain consistency across tasks. For instance, a data management server may provide a `process_dataset` tool for batch normalization, a `/data/reports.json` resource containing structured output, and a `summary_template` prompt for generating concise textual reports. By exposing tools, resources, and prompts through a unified JSON-RPC interface Morley (2010), MCP enables clients to discover, negotiate, and invoke capabilities. It also supports isolation, auditability, and policy enforcement when implementations follow the specification’s security guidance Radosevich & Halloran (2025).

Access to these capabilities is governed by OAuth-based authorization. Once a secure HTTPS channel is established, MCP proceeds with authorization at the application layer to manage access to tools and capabilities. An **authorization server** issues access tokens, and the MCP server enforces access to tools and resources using those tokens Model Context Protocol (2025l). This arrangement defines how clients obtain scoped authorization and how they maintain valid tokens throughout a session. In protected deployments, authorization servers implement OAuth 2.1, MCP servers expose OAuth protected resource metadata and support discovery, and clients use that metadata to locate the authorization server Hardt et al. (2024); Model Context Protocol (2025l). When a client connects to a protected MCP server, it initiates an authorization flow to obtain an access token for the server’s resource Posta (2025a). The token conveys the intended audience for the resource, the authorized scope, and an expiry time, and it must accompany each request in the `Authorization:Bearer <token>` header Model Context Protocol (2025l). The server validates the token’s issuer, audience, scope, and lifetime before executing any operation, ensuring that only authorized entities invoke sensitive tools or retrieve protected data. For example, a language-model-based research assistant might request read-only access to a document retrieval server that hosts academic literature. Through OAuth 2.1, it receives a time-limited token with the narrow scope needed for search operations but not




Figure 4: All communication between each MCP client and its paired server occurs through JSON-RPC 2.0 across the transport boundary. The callouts illustrate generic multi-line request and response templates with placeholders, where the response reuses the request id and returns the result object.

modification. Similarly, an enterprise analytics client querying a payroll server could be issued a token scoped to anonymized statistical aggregates, maintaining compliance with corporate privacy and auditing policies.

Overall, the MCP architecture cleanly separates reasoning, communication, and execution. This modularity lets agents orchestrate multiple tools and data sources across distributed environments while supporting strong security when transports, authentication, authorization, and logging are correctly configured. Once connected, a client can enumerate available capabilities, plan composite workflows, and execute coordinated tasks across heterogeneous servers. Such architectures mirror emerging agentic systems that integrate multiple models, APIs, and data endpoints to perform complex reasoning and multi-step task execution Shen et al. (2024).

2.2 MCP Lifecycle

The operational lifecycle of the Model Context Protocol (MCP) defines how an AI agent and its connected infrastructure progress from preparation to decommissioning. Studying this lifecycle is security-critical because each phase exposes distinct attack surfaces, and controls are typically phase tied across secure setup, careful discovery, least-privilege invocation with auditing, ongoing maintenance, and clean shutdown Ross et al. (2022); Tabassi (2023a); Souppaya et al. (2022a). Accordingly, this paper adopts an end-to-end, deployment-oriented lifecycle for MCP systems (*Pre-MCP* → *Creation & Registration* → *Connection & Initialization* → *Discovery* → *Invocation & Execution* → *Update & Maintenance* → *Shutdown & Deregistration*) that integrates and extends two established references Model Context Protocol (2025i); Hou et al. (2025). It refines the official MCP connection lifecycle by emphasizing an *Initialization* → *Operation* → *Shutdown* arc that

covers capability negotiation, routine JSON-RPC calls (e.g., `tools/list`, `resources/list`, `prompts/list`, `tools/call`), and graceful termination at both client and server connections Model Context Protocol (2025i). It also extends the server-centric lifecycle of Hou et al. (2025), which organizes MCP server operations into creation, deployment, operation, and maintenance with 16 activities, by adding pre-protocol preparation, explicit shutdown and deregistration hygiene, and a clear separation of discovery from invocation to support least-privilege analysis and auditing. Figure 5 presents the aforementioned lifecycle stages alongside a generic request flow in MCP. Taken together, this new framing lets phase-bound controls map directly to the corresponding Remote Procedure Calls (RPCs), configurations, and state transitions, setting up the stage-by-stage narrative that follows as we trace MCP from initialization to termination.




Figure 5: **MCP Lifecycle** - A user action in the Host creates an MCP client. If the server requires authorization, the client performs the OAuth 2.1 flow; in the third-party variant, the MCP server obtains a third-party token and then issues its own access token to the client. With a secure JSON-RPC channel established, the client performs capability discovery (`tools/list`, `resources/list`, `prompts/list`) and invokes the desired capability. The MCP server then executes the request against its Tool/Resource/Prompt back-end and streams results back to the client via JSON-RPC notifications. *Dashed arrows indicate optional steps in the sequence.*

The lifecycle begins with the **Pre-MCP Phase**, which covers the preparatory work performed before any protocol-based connection is established. This includes model alignment, dataset preparation, and tool-use familiarization Ouyang et al. (2022); Gebru et al. (2021); Tabassi (2023b). During this stage, developers train and fine-tune models to understand structured tool invocations, define access policies, and register the endpoints or APIs that the MCP server will later expose. For example, in a financial domain, developers may prepare a financial assistant by aligning it with transaction schemas, regulatory terminology, and dataset formats it will later encounter through MCP tools. This preparation does not involve live tool interaction but ensures the model understands how to interpret data such as account summaries, balance sheets, and invoices. In the same phase, schema definitions for tools are created so the MCP server can later expose them in a structured, interpretable form Model Context Protocol (2025p); OpenAI (2025). This alignment

allows the agent, once connected through MCP, to interact with these tools effectively without additional retraining Hou et al. (2025). Once the groundwork is complete, the **Creation and Registration Phase** initializes the server-side ecosystem. MCP servers are instantiated with metadata, including name, version, and supported schema definitions Model Context Protocol (2025c). This phase also includes registering tool manifests, resource URIs, and prompt templates that define the server’s callable GitHub, Inc. (2025); Model Context Protocol (2025e;d). Secure trust anchors, such as certificate pinning or public key registration, are configured so that future clients can verify the server’s authenticity Ahmadi et al. (2025). For instance, a deployment might register a “CI/CD tools” server exposing commands like `build`, `deploy`, and `rollback`, enabling reproducible software automation.

Following registration, the **Connection and Initialization** phase begins when the MCP client sends an `initialize` request, after which the client and server negotiate the protocol version and capabilities Model Context Protocol (2025b); Stytych (2025). Over HTTP transports, authorization endpoints must be served over HTTPS, which encrypts the channel and authenticates the server Model Context Protocol (2025l). Where required, deployments can also add mutual TLS to authenticate the client. Once the transport is secured, the client loads the available tool and resource lists, establishing an operational session McGinley-Stempel (2025). For example, a research assistant might initialize a connection to a document retrieval server, validating its certificate and synchronizing available datasets before use. After establishing a secure channel, the **Discovery Phase** enables dynamic capability enumeration. The client queries the connected server for its registered `tools/list`, `resources/list`, and `prompts/list` endpoints. The server responds with structured metadata describing each callable entity like its name, input format, output type, and required permissions Model Context Protocol (2025p). This phase allows AI agents to self-adjust their reasoning and available actions depending on what tools or datasets are currently registered. For instance, a travel planning agent could discover new `currency_exchange` or `visa_tracking` tools that extend its existing workflow.

Once discovery is complete, the **Invocation and Execution Phase** governs active interaction between the client and the server Model Context Protocol (2025b). Here, the agent issues structured JSON-RPC calls such as `tools/call` or `resources/read` to invoke specific capabilities Model Context Protocol (2025e;p). The server executes operations like querying databases, processing files, or composing responses and returns structured outputs back to the client. The host integrates these results into its reasoning loop, enabling chained or conditional actions Ahmadi et al. (2025). A data analysis agent, for example, might sequentially invoke `fetch_data`, `normalize_records`, and `summarize_report` tools in one continuous workflow. Over time, the **Update and Maintenance Phase** ensures that deployed servers remain stable and secure. Administrators may update tool manifests, rotate credentials, patch vulnerabilities, or retire deprecated endpoints Hasan et al. (2025); Hou et al. (2025). Observability and audit logs are also reviewed to maintain compliance and performance Ithena (2025). For instance, if analytics tools receive requests in a new format, developers can register updated schemas and redeploy compatible endpoints without interrupting existing clients. Finally, the **Shutdown and Deregistration Phase** completes the lifecycle. This stage involves gracefully terminating active sessions, cleaning up allocated resources, and unregistering outdated tools or datasets. Tokens and session keys are invalidated to prevent further invocations Lodderstedt et al. (2013). In distributed deployments, containers or compute instances may be reclaimed to reduce cost and exposure Raveendran (2024). For example, a tax filing assistant might de-register fiscal tools at the end of a financial year to prevent unauthorized reuse.

In summary, the MCP lifecycle establishes a continuous operational loop that blends initialization, execution, and controlled termination. Each stage reinforces the others, ensuring reliability, version traceability, and least-privilege access control across the protocol’s end-to-end operation.

3 Threat Taxonomy for MCP

MCP-based architectures involve tightly coupled interactions among underlying models (e.g., LLMs), agentic abstractions (e.g., agents and workflows), and communication protocols (e.g., tool invocation and orchestration) Model Context Protocol (2025a). Because this interplay expands the attack surface across components and lifecycle stages, we need a systematic way to reason about and mitigate the resulting risks. Figure 6 presents the threat taxonomy used in this work, summarizing the main categories of security risks in MCP




Figure 6: Threat taxonomy categories in MCP

deployments. However, knowing *where* and *when* a threat appears is not sufficient; we must also account for *what* the weakness fundamentally is, since the underlying failure mode is primarily used to generalize fixes, compare incidents, and prioritize controls. This structure is deliberate because a failure-mode-first taxonomy explains why incidents happen, whereas a dual-axis grid of component and lifecycle only describes where and when they surface. If we start from the dual axis alone, causally identical issues get split across multiple component–lifecycle coordinates and duplicated whenever they recur in a different place or phase, so the shared mechanism disappears. For example, consider an over-permissive file tool created during registration that is later exploited by prompt injection during execution, with secrets lingering after a missed revocation during update. Read purely through the dual axis, that single causal chain appears as three separate entries, one in Server Tools at Creation, one in Host at Execution, and one in Host at Update, with different owners and severities, which obscures that one authorization-scope flaw seeded the entire cascade. Grouping first on the root cause and only then annotating each threat with precise where and when coordinates (i) keeps causally similar events together, (ii) avoids combinatorial duplication, and (iii) attaches the right class of controls to the actual mechanism (e.g., scope linting and allow lists for mis-scoping, not just host logging at execution).

More generally, failure modes are stable, while components and phases are contingent on deployment and operational context such as where a tool runs, how it is configured, and when checks are applied. For example, a tool may move from local studio to a remote server, or a check may shift from creation to connection; the grid coordinates change, yet the underlying vulnerability, such as an identity and trust failure from weak server verification, remains the same. Anchoring on failure modes provides invariants that can make evidence, mitigations, and future metric design more reusable across deployments. It also helps avoid double counting and can support later prioritization by keeping severity and frequency reasoning attached to stable failure modes rather than scattered across many cells. This framing matches how teams actually fix things: engineering ships control patterns to defeat mechanisms, for example PKI and mTLS for

identity and trust, signing and provenance for supply chain, and sandboxing and quotas for execution, while SRE and security use the dual-axis tags to decide where to enforce them. In short, failure modes partition by cause so patterns, metrics, and mitigations stay coherent, and the dual axis localizes by context so fixes land precisely. The result is a taxonomy that travels well across deployments and remains surgical in practice: a structured threat taxonomy that links *where* a threat occurs to *when* it emerges in the system’s lifecycle and, crucially, to *what* the weakness fundamentally is. This unified what–where–when view provides the holistic context needed to understand and mitigate MCP vulnerabilities, closing gaps left by approaches that track only location and phase.




Figure 7: One exemplar threat from each of the nine MCP threat categories (§ 3.1–§ 3.9). (a) *Identity & Trust Management* - server-name collision/impersonation. (b) *Malicious or Poisoned Tool Servers* - pre-invocation code injection exfiltrating secrets. (c) *Credential Leakage & Data Exposure* - centralized auth/token store as a single point of compromise. (d) *Supply Chain & Code Integrity* - trojanized dependency/source. (e) *Prompt, Input & Context Injection* - indirect prompt injection via retrieved web content. (f) *Authorization & Access Control Gaps* - missing runtime/context-aware checks. (g) *Execution Environment & Sandbox Risks* - prompt-driven OS command/DoS from an unsandboxed tool. (h) *Workflow & Configuration Risks* - context drift across a multi-step pipeline. (i) *Multi-Agent & Orchestration Risks* - unverified agent handoff/privilege carry-over. Dashed red traces depict the attack flow; icon overlays indicate the involved MCP components.

Furthermore, to ground this framework in evidence, we anchor the taxonomy in field observations so its categories reflect how failures actually occur. We construct it by examining real-world incidents, empirical audits, and prior frameworks to identify recurring failure patterns in MCP, drawing on synthesis pieces and technical write-ups that survey the ecosystem. Our evidence base includes reports of 492 publicly exposed MCP endpoints lacking authentication or encryption Data Science Dojo Staff (2025); Lammerts (2025). We also draw on ecosystem-scale assessments that quantify prevalent weaknesses, with about 43% of analyzed servers exhibiting command-injection or code-execution exposure, 33% permitting unrestricted URL fetches, 22% exposing files, and 66% showing poor secret-handling code smells Raina (2025). In parallel, independent

coverage reports more than 15,000 MCP servers identified via public GitHub code search, which we treat as a prevalence indicator rather than an internet-reachability measurement Prompt Security Team (2025b). Field measurements further document hundreds of servers bound to 0.0.0.0 enabling neighborjacking and widespread misconfiguration, with multiple sources characterizing nearly 500 unauthenticated exposures Naamnih & Ginzburg (2025); Data Science Dojo Staff (2025); Lammerts (2025); Nelson (2025). We also ground the taxonomy in concrete incidents, including a critical RCE in the official MCP Inspector and a separate arbitrary code execution issue in **autogen** Lumelsky (2025); Snyk Security Research (2024). These are a few representative examples among many we examined, and together they underscore recurring, high-impact failure modes in real deployments. Building on these observations, we defined **nine dominant failure-mode categories** that reflect *how* MCP incidents actually arise in practice. They include problems of identity and trust, software supply-chain compromise, prompt/agent interference, malicious tool use (“tool abuse”), data leakage, authorization scope errors, and others. Every MCP-specific threat we encountered falls under one of these broad failure modes. Crucially, we then annotate each threat with the specific component(s) it affects (such as the host, the JSON-RPC communication layer, or a tool server) and the lifecycle phase(s) in which it arises (for instance, during tool creation, at invocation/execution time, or upon update). This dual-axis tagging situates each threat in both space and time without losing the generality of the failure-mode grouping. The result is a taxonomy grounded in real incident types, yet precise in mapping each threat to the exact points in the MCP architecture and lifecycle where it must be addressed. Under the proposed taxonomy, each threat is also cross-referenced with both STRIDE and MAESTRO threat modeling frameworks. STRIDE was chosen because it encapsulates the classical security threat categories (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, and Elevation of Privilege), thereby aligning each threat with fundamental security properties. In contrast, MAESTRO was included as it is purpose-built for agentic AI systems, adding context-aware threat dimensions that traditional models overlook, such as adversarial attacks, data poisoning, model misuse, or misalignment issues. Mapping to both frameworks is essential in the MCP context to ensure that conventional cybersecurity concerns and novel AI-specific risks are addressed in tandem.

Review Protocol and Coding Procedure. To improve reproducibility and reduce selection bias, we followed a structured review protocol to build the threat corpus used in tables 1–3. The objective was to identify MCP-specific threats, MCP-adjacent failure modes with clear transferability, and mitigation approaches relevant to real MCP deployments. Rather than relying on an informal or purely narrative selection process, we defined explicit search, screening, consolidation, and coding steps so that the construction of the taxonomy is transparent and auditable.

We searched a mixed evidence base that included both archival and operational sources. This was necessary because MCP security is an emerging domain, and relevant evidence is distributed across research papers, protocol documents, benchmarks, and incident-focused practitioner reporting. Our search covered the period from November 2024, when Anthropic introduced the Model Context Protocol, through February 2026. The search included Google Scholar, Semantic Scholar, arXiv, the ACM Digital Library, IEEE Xplore, official MCP specifications and documentation, security advisories, benchmark papers, and high-signal industry and practitioner reports. Query templates combined MCP-specific and transfer-oriented terms. Example queries included “*Model Context Protocol security*,” “*MCP threat*,” “*MCP server attack*,” “*MCP prompt injection*,” “*MCP tool poisoning*,” “*agent-tool protocol security*,” and “*multi-agent MCP risk*.” To improve coverage, we also used backward citation tracing from highly relevant papers, benchmark artifacts, and protocol documents.

To account for differences in source quality, retained materials were grouped into broad source categories, including peer-reviewed papers, preprints and workshop papers, official specifications and documentation, security advisories, and practitioner reports. Peer-reviewed and benchmark-based evidence was prioritized when available, while non-archival sources were used mainly to capture emerging incidents, deployment practices, and ecosystem-specific attack patterns not yet reflected in the archival literature. We then screened the collected material in two stages. Titles, abstracts, and executive summaries were first reviewed to remove clearly irrelevant records. Candidate sources that appeared to contain MCP-relevant threat, defense, architectural, or incident evidence were then assessed through full-text review, after which duplicate records were removed before final coding. This process yielded a final survey corpus of 285 sources. Included

material covered direct analyses of MCP threats, defenses, incidents, and design weaknesses. Sources from adjacent ecosystems were also retained when the failure mode clearly transferred to MCP deployments, including package registries, plugin platforms, sandboxes, and multi-agent systems. Additional sources were included when they described protocol or architectural properties implying MCP-relevant risks, or provided mitigation, benchmark, or evaluation evidence relevant to MCP security. Sources were excluded if they were duplicates, discussed generic LLM risks without an MCP-specific mechanism, made purely speculative claims without plausible deployment grounding, or mentioned MCP only in passing without security-relevant analysis.

Deduplication was performed at the level of the underlying failure mechanism rather than the immediate triggering surface. This allowed reports describing the same core issue across different tools, transports, or deployment settings to be merged into a single canonical threat entry while preserving supporting evidence. Each retained source was then coded by failure-mode category, affected MCP component, lifecycle phase, STRIDE label, MAESTRO label, and mitigation linkage. When a source supported multiple threat manifestations, each was coded separately and consolidated when appropriate. Borderline cases and cross-category ambiguities were reviewed jointly by the authors using a written coding schema. The resulting protocol is structured and reproducible, though not intended to be a fully exhaustive systematic review. Instead, it provides a transparent survey process for a fast-moving security domain.

The remainder of this section presents the taxonomy by category (§ 3.1–§ 3.9) and cross-classifies each threat along the orthogonal axes introduced above. Each 3.x subsection opens with a brief primer, followed by a set of threats that share a fixed structure, first describing the MCP-specific failure mechanism, then grounding it in concrete examples or evidence, and finally summarizing the deployment-level impact. Figure 7 provides a per-category overview with one exemplar threat, and tables 1–3 present the complete mapping of threats by category and, for each entry, enumerate the associated component, lifecycle phase, and STRIDE/MAESTRO coordinates.

3.1 Identity and Trust Management

MCP’s decentralized architecture and lack of a global naming authority make identity assurance difficult across clients, servers, and tools. In practice, clients and agents rely on names, manifests, and other lightweight metadata rather than cryptographic proofs, and this weak binding between identity and privilege increases exposure to look-alike endpoints and on-path interception, reducing a client’s ability to distinguish legitimate from adversarial services during registration, discovery, and connection establishment. Four recurring failure modes illustrate how these identity and trust breakdowns arise in practice, each compounding the others through a shared root cause of absent enforced cryptographic identity at the protocol level.

The most direct manifestation is **Server Name Collisions and Impersonation**, where MCP clients may resolve to malicious servers whose names are deliberately crafted to appear confusingly similar to trusted ones (for example, `dropboxuploader` vs. `dropbox-uploader`), exploiting typographical errors or visual similarity during configuration and discovery. Empirical analyses of MCP-style registries and adjacent plugin ecosystems show that weak vetting allows adversaries to publish such lookalike packages and endpoints, creating realistic opportunities for typosquatting and impersonation that are difficult for both human operators and automated agents to distinguish Hou et al. (2025); Posta (2025b); dts-solution (2025), and once a client or agent is bound to an impersonating server, the adversary can intercept or exfiltrate sensitive context, inject malicious tool outputs, and silently alter execution paths, turning a superficial naming confusion into a persistent supply chain foothold and workflow hijack that propagates across multiple MCP hosts and orchestrated agents. This risk is directly amplified by **Mutual Authentication and Encryption Gaps**. MCP’s early specification defined client-server transports such as `stdio` and HTTP without requiring mutual authentication or secure channel enforcement, leaving choices about TLS, certificate validation, and client verification to implementers Hanson & maintainers (2024). This design gap persisted even after OAuth-based authentication support was later introduced, because it remained an optional transport-level extension rather than a baseline requirement. As a result, multiple security reviews have found MCP endpoints exposed on the public internet with little or no authentication, or protected only by server-side TLS without strong client identity checks Model Context Protocol (2025l); TrolEye Security (2025). Independent analyses and best-practice guides now strongly recommend mutual TLS and hardened authentication for MCP deploy-

ments Ramnarayanan (2025); Paktiti (2025); Kanoon (2025). These controls matter because an attacker who controls a network path or a misconfigured gateway can interpose on traffic, impersonate servers, and read or tamper with JSON-RPC exchanges when such protections are missing. As a result, clients may silently bind to attacker-controlled endpoints, disclose sensitive context, and act on manipulated tool outputs. These failures create a pathway for man-in-the-middle attacks, credential theft, and long-lived workflow hijacking that can propagate across multiple hosts and orchestrated agents. Even where transport is secured, the **Lack of Verified Server Identity or PKI** persists because MCP does not yet define a native mechanism for binding servers to cryptographic identities beyond generic HTTPS/TLS and OAuth infrastructure Anthropic (2025); Posta (2025b). As a result, clients and registries often treat servers as trustworthy based on self-reported metadata such as names, descriptions, or repository provenance. Because these signals can be spoofed or cloned, they do not reliably distinguish legitimate servers from adversarial ones during registration and discovery. This increases the risk of impersonation, workflow hijacking, and long-lived malicious footholds across the MCP ecosystem. **Implicit Trust in Unauthenticated Servers** deepens this exposure by allowing unverified endpoints to continue participating in apparently legitimate workflows. Early MCP deployments often ran servers on localhost or within tightly controlled internal networks, so implementers frequently omitted strong authentication and identity checks between clients and servers Spletzer (2025). When those same tools and agents are later exposed over HTTP, cloud infrastructure, or shared internal networks, this inherited trust model becomes a liability. Any process with network reach to the server port can issue tool invocations without being verified Segal et al. (2025); National Vulnerability Database (2025a). As a result, malicious or misconfigured servers can trigger privileged tool executions, escalate access, and exfiltrate sensitive data while workflows continue to appear legitimate to users. Detailed definitions of each threat are provided in Appendix A.1.

Table 1: MCP Threat Taxonomy (Part 1)

Section	Category	Threat	Components	Lifecycle	STRIDE	MAESTRO
§ 3.1	Identity & Trust Management	Server Name Collisions and Impersonation	Transport, Server, Client, Auth. Infra	Connection / Initialization	Spoofing	Security & Compliance
		Mutual Authentication and Encryption Gaps	Transport, Server, Client, Auth. Infra	Connection / Initialization	Spoofing, Information Disclosure	Security & Compliance
		Lack of Verified Server Identity or PKI	Transport, Server, Client, Auth. Infra	Connection / Initialization	Spoofing	Security & Compliance
		Implicit Trust in Unauthenticated Servers	Transport, Server, Client, Auth. Infra	Connection / Initialization	Spoofing	Security & Compliance
§ 3.2	Supply Chain & Code Integrity	Trojanized or Compromised Installers	Server, Host	Creation / Registration	Tampering	Deployment Infrastructure
		Malicious Dependencies or Backdoored Code	Server, Host	Creation / Registration	Tampering	Deployment Infrastructure
		Lack of Integrity Verification	Server, Tools	Creation / Registration, Update / Maintenance	Tampering	Deployment Infrastructure
		Vulnerable Redeployments and Stale Manifests	Server, Host	Update / Maintenance	Tampering	Deployment Infrastructure
		Unverified Package Sources and Automatic Invocation	Server, Host	Creation / Registration	Tampering	Deployment Infrastructure
§ 3.3	Malicious or Poisoned Tool Servers	Tool Shadowing and Name Collisions	Tools, Client, Server, Data layer	Discovery	Spoofing	Agent Frameworks
		Rug Pull via Dynamic Tool Redefinition	Tools, Client, Host, Data layer	Invocation / Execution, Update/Maintenance	Tampering	Agent Frameworks
		Poisoned Tool Metadata	Tools, Host, Client, Data layer	Discovery	Tampering	Agent Frameworks
		Pre-Invocation Code Injection	Tools, Host, Data layer, Client, Server	Invocation / Execution	Tampering	Agent Frameworks
		Slash Command Collisions	Tools, Host, Data layer, Client	Discovery	Spoofing	Agent Frameworks
		Preference Manipulation	Tools, Host, Data layer, Client	Discovery	Spoofing	Agent Frameworks
		ANSI Escape Code Injection	Tools, Host, Data layer, Client	Discovery	Spoofing	Agent Frameworks
§ 3.4	Prompt Input & Context Injection	Poisoned Prompt Templates	Server, Client, Host	Discovery, Invocation / Execution	Tampering	Data Operations
		RADE (Retrieval-based poisoning)	Host, Client, Tools, Server	Invocation / Execution	Tampering	Data Operations

Table 2: MCP Threat Taxonomy (Part 2)

Section	Category	Threat	Components	Lifecycle	STRIDE	MAESTRO
§ 3.4	Prompt Input & Context Injection (cont.)	Transitive-Trust Poisoning	Server, Client, Tools, Data layer, Host	Discovery, Invocation / Execution	Tampering	Data Operations
		Format-Embedded Poisoning	Data layer, Tools, Client, Server	Discovery, Invocation / Execution	Tampering	Data Operations
§ 3.5	Credential Leakage and Data Exposure	Centralized Credential Storage	Data layer, Host, Server, Auth. Infra	Creation / Registration, Update / Maintenance, Invocation / Execution	Information Disclosure, Elevation of Privilege	Security & Compliance
		Insecure Secret Management	Server, Client, Host, Tools	Creation / Registration, Invocation / Execution	Information Disclosure	Security & Compliance
		Overprivileged & Persistent Tokens	Auth. Infra, Server, Client	Creation / Registration, Update / Maintenance, Invocation / Execution	Elevation of Privilege, Information Disclosure	Security & Compliance
		Cross-Tool Data Flows	Host, Client, Server, Tools, Data layer	Invocation / Execution	Information Disclosure	Data Operations
		Emergent Aggregation Risks	Host, Client, Server, Tools, Data layer	Invocation / Execution	Information Disclosure	Data Operations
		Residual Data Exposure	Data layer, Host, Server, Tools, Client	Shutdown / De-Registration	Information Disclosure	Data Operations
		Trojanized MCP Servers	Host, Client, Server, Tools, Data layer	Invocation / Execution	Information Disclosure, Tampering	Data Operations
§ 3.6	Authorization & Access Control Gaps	Binary Permissions	Auth. Infra, Client, Server, Tools	Creation / Registration	Elevation of Privilege	Security & Compliance
		Lack of Risk-Aware Differentiation	Client, Host, Server	Invocation / Execution	Elevation of Privilege	Security & Compliance
		Consent Fatigue and Automation Drift	Host, Client	Invocation / Execution	Elevation of Privilege	Security & Compliance
		No Context-Aware Revalidation	Auth. Infra, Server	Invocation / Execution	Elevation of Privilege	Security & Compliance
		Uneven RBAC Support	Auth. Infra, Server	Creation / Registration	Elevation of Privilege	Security & Compliance
		Orphaned Sessions and Credentials	Auth. Infra, Server, Client, Host	Shutdown / De-Registration	Elevation of Privilege	Security & Compliance
		Confused Deputy / Role Leakage	Auth. Infra, Server, Client, Host	Invocation / Execution	Elevation of Privilege, Information Disclosure	Security & Compliance
		Decentralized & Inconsistent Policy Enforcement	Auth. Infra, Server, Host	Creation / Registration, Update / Maintenance	Elevation of Privilege	Security & Compliance
§ 3.7	Execution Environment & Sandbox Risks	Lack of Execution Isolation	Host, Server	Creation / Registration, Update / Maintenance, Invocation / Execution	Elevation of Privilege, Information Disclosure	Deployment Infrastructure
		Sandbox Escape Vulnerabilities	Host, Server	Update / Maintenance, Invocation / Execution	Elevation of Privilege	Deployment Infrastructure
		Root or Elevated Execution	Host, Server	Creation / Registration, Invocation / Execution	Elevation of Privilege	Deployment Infrastructure

3.2 Supply Chain and Code Integrity Risks

Even if all of the risks discussed under Section 3.1 were fully mitigated, MCP deployments would still face supply chain and code integrity problems, because servers and tools that are correctly identified at the protocol layer can be subverted when their code, dependencies, or build artifacts are compromised. MCP’s reliance on decentralized, third-party tool ecosystems and ad hoc deployment practices means that many deployments operate without centralized vetting or consistent integrity policies for server binaries, manifests, and libraries, and under these conditions malicious packages, compromised dependencies, and outdated or tampered versions can slip into the toolchain and quietly undermine entire workflows even when endpoints appear to be authenticated. Five recurring failure modes rooted in weak provenance tracking and insufficient integrity checks illustrate how these supply chain breakdowns arise in practice, each one widening the attack surface that the previous leaves open. The first and most immediate failure is **Trojanized or Compromised Installers**.

Tools like **Smithery-CLI**, **mcp-get**, and **mcp-installer** streamline MCP server setup, tool registration, and discovery by offering one-command workflows to install MCP servers and configure them for use with agents

Table 3: MCP Threat Taxonomy (Part 3)

Section	Category	Threat	Components	Lifecycle	STRIDE	MAESTRO
§ 3.7	Execution Environment & Sandbox Risks (contd.)	Prompt-Based Remote Code Execution	Host, Server, Client	Invocation / Execution	Tampering, Elevation of Privilege	Deployment Infrastructure
		DoS and Resource Exhaustion	Host, Server	Creation / Registration; Update / Maintenance, Invocation / Execution	Denial of Service	Deployment Infrastructure
		Vulnerable Third-Party Dependencies	Server, Tools	Update / Maintenance; Invocation / Execution	Elevation of Privilege, Information Disclosure	Deployment Infrastructure
		Lateral Movement via Shared Infrastructure	Auth. Infra, Host, Server	Creation / Registration, Update / Maintenance, Invocation / Execution	Elevation of Privilege, Information Disclosure	Agent Ecosystem
§ 3.8	Workflow & Config. Risks	Context Drift and State Inconsistency	Host, Client, Server	Invocation / Execution	Tampering	Agent Ecosystem
		Configuration Drift Across Servers	Host, Server	Update / Maintenance	Elevation of Privilege	Security & Compliance
		Transactional Integrity Gaps	Host, Client, Server	Invocation / Execution	Repudiation, Tampering	Evaluation & Observability
		Patch Gaps and Vulnerable Instances	Server	Update / Maintenance	Elevation of Privilege	Deployment Infrastructure
		Session Isolation in Multi-Agent Environments	Host, Client, Server	Invocation / Execution	Spoofing, Information Disclosure	Agent Ecosystem
		Unsafe Fallbacks and Error Handling	Host, Client, Server	Invocation / Execution	Information Disclosure, DoS	Evaluation & Observability
		Operational Visibility and Logging Gaps	Host, Client, Server	Update / Maintenance	Repudiation	Evaluation & Observability
		Incomplete Tool De-Registration	Auth. Infra, Server, Client, Host	Shutdown / De-Registration	Spoofing, Elevation of Privilege	Security & Compliance
§ 3.9	Emergent Threats from Model Behavior and Human Factors	Model Inversion Attacks	Host (LLM)	Invocation / Execution	Information Disclosure	Foundation Models
		Membership Inference	Host (LLM)	Invocation / Execution	Information Disclosure	Foundation Models
		Pre-MCP Data Poisoning	Host (LLM)	Pre-MCP Model Training	Tampering	Foundation Models
		Agent Collusion	Host, Client, Server	Invocation / Execution	Tampering	Agent Ecosystem
		Orchestration Drift	Host, Client, Server	Invocation / Execution	Tampering	Agent Ecosystem
		Unverified Agent Handoffs	Host, Client, Server	Invocation / Execution	Elevation of Privilege, Tampering	Agent Ecosystem
		Malicious Agent Operators	Host	Invocation / Execution	Tampering	Agent Ecosystem
		Insider Threats and Shadow MCP Servers	Host, Server	Creation / Registration, Update / Maintenance	Tampering, Information Disclosure	Security & Compliance
		Agent Interface Social Engineering	Host	Invocation / Execution	Spoofing	Agent Ecosystem

and hosts Smithery (2025); mcp (2025a;b). If an installer binary, its scripts, or its update channel is trojanized, that same single command can instead pull a backdoored MCP server, register attacker-controlled endpoints under plausible names, and persist those changes in local configuration files and install directories that clients will later reuse Smithery (2025); mcp (2025a;b). Because a single run with these privileges can rewrite trust anchors, pin vulnerable versions, or add rogue endpoints that blend into later discovery, the compromise can become durable and propagate across multiple agents, hosts, and projects that share the same installer-driven environment. Even without installer-level compromise, MCP deployments remain vulnerable through **Malicious Dependencies and Backdoored Code**, since MCP tools distributed through ecosystems such as PyPI and npm inherit the broader software supply chain risks of those platforms including dependency hijacking, typosquatting, and backdoor injection into both direct and transitive dependencies, and empirical studies of malicious npm packages show that attackers frequently abuse transitive dependencies and install-time scripts so that malware runs automatically during installation and can tamper with or exfiltrate sensitive data Sejfia & Schäfer (2022). Recent incidents such as the compromise of **rand-user-agent** demonstrate how seemingly innocuous packages can be repurposed into delivery vectors for Remote Access Trojans Toulas (2025); Eriksen, Charlie (2025), and typosquatting cases like **colorizr** versus **colorama** show

how small, deceptive naming changes can quietly steer operators toward malicious code Meyer (2025); The Hacker News (2025); Cohen & Yerushalmi (2024), meaning that within an MCP deployment, packaging a server or tool around such backdoored dependencies turns each tool invocation into a potential opportunity for remote access, credential theft, or data exfiltration wherever the compromised image is reused.

Both of these failure modes are worsened by a **Lack of Integrity Verification**. MCP currently provides no built-in mechanism to enforce digital signatures or cryptographic hash checks, so clients cannot reliably detect whether tool calls or responses have been tampered with in transit Strobes Security Labs (2025). Reflecting this gap, Red Hat’s MCP security guidance treats servers as ordinary executable code and recommends signing MCP components, verifying dependency integrity, and using hardened build pipelines with SAST and software composition analysis Red Hat Product Security (2025). Because these protections must be added around MCP rather than enforced by the protocol itself, an adversary who can modify a repository, package feed, or network path can silently swap or patch server binaries and alter their responses. Routine installs, auto-updates, and tool invocations can therefore become supply chain and message-tampering attacks that propagate to every client and workflow depending on the compromised server.

Even after vulnerabilities are patched, **Vulnerable Redeployments and Stale Manifests** can quietly downgrade an MCP deployment back to a known vulnerable version when stale configuration files, lockfiles, or deployment manifests continue to reference the old image Ehtesham et al. (2025). Related context-drift failures in real-world systems show that lost or corrupted manifest updates can leave agents and services running conflicting or outdated code across nodes MacVittie (2025). This problem becomes more serious because the same manifests and version pins are often cloned across environments, including developer IDEs, CI pipelines, and shared agent hosts. As a result, a single poisoned or outdated manifest can keep a compromised image in rotation across projects and turn what should have been a localized patch into a system-wide incident. Automated version validation and build failure on hash mismatch are therefore essential controls.

Cutting across all of the above is **Unverified Package Sources and Automatic Invocation**, where tools in MCP are explicitly designed to be model-controlled and hosts translate LLM output into tool calls, so once a server from a registry or package source is configured its tools can be automatically invoked based on their metadata during normal agent workflows Model Context Protocol (2025p); Li & Gao (2025), and a large-scale empirical study of the MCP ecosystem shows that public registries lack vetted server submission mechanisms and that a substantial number of listed servers can be hijacked Li & Gao (2025), while a separate security scan of 1,000 public MCP servers reports that roughly one third expose at least one critical vulnerability Birur (2025), meaning that a single compromised server installed from an unverified source can exfiltrate chat or backend data or trigger unintended actions during otherwise routine workflows. Detailed definitions of each threat are provided in Appendix A.2.

3.3 Malicious or Poisoned Tool Servers

Identity and supply-chain defenses in Section 3.1 and Section 3.2 address *who* we trust and *what* we deploy; however, this threat class concerns *how* tools present and are bound at runtime. At runtime, MCP clients discover tools through server-advertised metadata, route calls based on names and schemas, and execute whatever implementation the server exposes, which means a malicious or manipulated tool server does not need to break authentication or compromise a package to cause harm. It only needs to control what the client sees during discovery, what code runs when a tool is called, or what happens to an already-trusted tool after approval. Six recurring failure modes illustrate how malicious or poisoned tool servers exploit this binding surface in MCP, each targeting a different point in the chain from discovery through execution.

The first and most foundational failure is **Poisoned Tool Metadata**, where adversaries embed hidden prompts in MCP tool descriptions, YAML manifests, or docstrings that the client surfaces verbatim to the LLM during tool discovery and selection, causing these fields to be interpreted as instructions instead of neutral metadata Li et al. (2026b). Such misleading metadata can bias the model toward the poisoned tool, weaken safety checks, or trigger additional tool invocations, and recent work shows this influence persists even when the poisoned tool is never invoked, instead steering the model toward other high-privilege tools through implicit tool poisoning Invariant Labs (2025). A seemingly harmless `add()` tool whose docstring quietly

instructs the model to exfiltrate the `~/.ssh/id_rsa` private key illustrates how any host that auto-registers such a tool can be coerced into leaking SSH keys, configuration files, or API credentials during routine completions Invariant Labs (2025). Building directly on this metadata surface, **Tool Shadowing and Name Collisions** allow adversaries in open MCP registries to pre-register look-alike tool names that differ only by a hyphen, typo, or word order, such as `deploy_model` versus `deploy-model` Yu et al. (2025). These small naming differences can cause string-based discovery logic to surface attacker-controlled endpoints alongside legitimate ones. Security analyses show that such collisions can divert invocations intended for trusted tools to malicious implementations that mimic the expected interface Beretta et al. (2025). As a result, data and execution may be silently redirected even though the invocations still appear to be routine tool calls. Where name collisions misroute calls at discovery, **Pre-Invocation Code Injection** operates earlier still by hiding pre-execution logic in wrapper functions or initialization code. As a result, a call that appears benign may actually execute attacker-controlled instructions first in the host environment Fenstermacher (2025). Other studies also find tools that interpolate unvalidated parameters directly into shell commands such as `os.system(...)`, allowing crafted payloads to result in remote code execution whenever the agent invokes the tool Cross (2025). Figure 8 illustrates how these three failures combine in a single attack chain.




Figure 8: An Example attack combining *Cross-Server Tool Shadowing*, and *Poisoned Tool Metadata*, and *Pre-Invocation Code Injection*. Two MCP servers expose `read_file` with near-identical visible descriptions; the client binds by bare name (e.g., last-seen), so `tools/call` is routed to the malicious server. A hidden comment in the tool description induces a pre-invocation read of `~/.ssh/id_rsa`, and the secret is piggybacked in the response.

Even tools that pass initial inspection remain vulnerable through **Rug Pull via Dynamic Tool Redefinition**. The danger is that a tool which appears trusted at approval time can later be silently replaced with a malicious variant when its server code or manifest is updated in place Song et al. (2025). As a result, agents may continue resolving the same tool name while actually fetching and executing altered logic. Empirical demonstrations show malicious servers uploaded to public aggregators that later trigger harmful actions such as accessing private files or transferring digital assets once users have trusted them Song et al. (2025), and Bhatt et al. (2025) propose ETDI with immutable, versioned tool definitions so that clients bind to specific versions rather than whatever happens to be served under a name at invocation time, because without such protections previously benign servers become system-wide backdoors that hijack file access, API calls, or

transaction flows across all hosts that continue to trust the original registration. Operating at a subtler level, **Preference Manipulation** steers tool selection by shaping what the agent sees during discovery through persuasive descriptions that claim greater safety or capability, curated examples that present the tool as the natural default, and presentation cues such as badges or recommended labels that raise salience Hasan et al. (2026); Model Context Protocol Security (2025); Elastic Security Labs (2025). Unlike poisoned metadata that embeds explicit directives, preference manipulation uses non-imperative phrasing to bias ranking so the attacker’s tool is chosen under normal selection logic, resulting in a persistent bias toward unsafe tools that increases the likelihood of bypassed safeguards without ever triggering an explicit instruction. Rounding out this category, **Command and Function Collisions** arise when multiple MCP tools expose the same slash command or function name, such as `/query` or `/upload`, within a shared namespace Yang et al. (2025); dts-solution (2025). In that setting, one definition can shadow another, causing the client to consistently bind the command to the wrong implementation. Even without malicious intent, ambiguous autocompletion and precedence rules can still route execution to the wrong tool. The result can be information disclosure, privilege escalation, or an effective denial of service against the legitimate command.

Finally, **ANSI Escape Code Injection** targets the human review layer rather than the model. In this attack, an adversary embeds ANSI control sequences in strings shown to humans, such as a tool’s `name` or `description` from `tools/list`, or text returned by `tools/call`. If a CLI or inspector renders these strings without escaping them, the terminal can rewrite the visible screen while the underlying JSON remains unchanged. Public CVEs confirm the feasibility of this kind of display-layer spoofing, in which escape codes can fake listings or hide content (CVE-2024-33899; CVE-2024-36052) nvd (2024a;b). The result can be staged warnings, masked prior lines, and misleading screens that increase the likelihood of unsafe approvals or missed indicators during tool or log review. Detailed definitions of each threat are provided in Appendix A.3.

3.4 Prompt, Input, and Context Injection Attacks

In contrast to Sections 3.1–3.3, which focus on who is trusted, what code is deployed, and how tool servers are provisioned, this class of attacks targets how MCP agents assemble and reuse context at runtime. An agent’s working context is constructed from many channels including user prompts, server-supplied templates, retrieved website content, local files, API responses, and intermediate tool outputs, and because any of these channels can carry hidden instructions, adversaries can plant payloads that hitchhike on otherwise legitimate data and are later reintroduced into the model’s context. Once incorporated, these injected instructions can steer the model to exfiltrate sensitive information, circumvent safeguards, or perform unintended actions that still appear aligned with the surrounding workflow. The underlying failure is insecure context assembly, where heterogeneous inputs are accepted and re-prompted without validation, provenance checks, or isolation.

The most direct instantiation is **Poisoned Prompt Templates**, where servers publish reusable prompt templates that clients automatically inject into the model context, so a compromised server can embed hidden instructions that the model treats as trusted guidance rather than untrusted user input Microsoft (2025b). Prompt-level payloads alone can induce information leakage, credential theft, and tool misuse across open-source agent frameworks Palo Alto Networks Unit 42 (2025), meaning a single hijacked server propagates misbehavior across every host that imports its templates. Extending this surface to retrieved content, **Retrieval-Based Poisoning (RADE)** embeds malicious MCP instructions inside files or online content that retrieval tools ingest, so the LLM treats them as benign context and follows them as part of the user’s query Radosevich & Halloran (2025). The RADE attack demonstrates this concretely, where Claude uses the Chroma MCP server to index a crafted file, then uses filesystem and Slack tools to add SSH keys and exfiltrate API keys to external channels Radosevich & Halloran (2025), and follow-on work shows that existing guardrails often fail to block such attacks even after the model has been explicitly warned Halloran et al. (2025).

Transitive-Trust Poisoning (Fourth-Party Injection) extends this surface one step further, arising when a trusted MCP server forwards content from external systems into the model context without revalidating it, so hidden instructions in upstream data are processed as authoritative tool guidance rather than untrusted input. In a widely analyzed Supabase MCP deployment integrated into the Cursor IDE, a single crafted support ticket prompted the agent running with `service_role` privileges to execute SQL that read sensitive tables and wrote the results back into the same ticket sshh.io (2025); General Analysis (2025), show-




Figure 9: An example illustrating Retrieval-Agent Deception (RADE) attack in MCP. A user asks for a Q3 KPI summary and uploads a crafted `FBA.doc` whose hidden instructions cause the agent, via an MCP file server, to read and return additional internal files, exfiltrating sensitive data beyond the explicit request.

ing that a single poisoned upstream message can drive unsafe tool calls across any workflows or tenants that depend on the same server. Rounding out this category, **Format-Embedded Poisoning (Weaponized Content Channels)** arises when MCP resources such as HTML pages, PDFs, logs, or JSON API responses are surfaced directly into the model context, allowing attackers to hide model-interpretable instructions inside content that appears to be neutral data. Hidden instructions in tool text or WhatsApp messages can cause an MCP-enabled agent to read private files, forward prior message history, and exfiltrate that data to an attacker-controlled endpoint without the user noticing Beurer-Kellner & Fischer (2025). The risk becomes even greater when clients do not constrain which tools the LLM may invoke after interpreting such resources. In that setting, a single poisoned document or message can redirect execution and compromise multiple workflows that ingest the same content channel Willison (2025b). Each threat is defined in detail in Appendix. A.4.

3.5 Credential Leakage and Data Exposure

Prior sections address who MCP agents talk to, what code they run, and how tools and context are bound, and this section addresses what secrets those interactions expose and how they can be lost. MCP servers often act as high-privilege intermediaries holding credentials, tokens, and high-value business content on behalf of multiple tools and services, so weaknesses in how they store, scope, or transmit this information can turn otherwise well-secured workflows into channels for exfiltration and misuse. The resulting risks generally follow three broad patterns. First, poor credential handling can allow compromise of one secret to cascade across multiple downstream systems. Second, data may leak when tools log, transmit, or redirect information beyond its intended scope. Third, aggregation effects can arise when individually benign fragments combine into sensitive disclosures.

The first pattern appears in credential concentration and retention practices that allow a single compromise to unlock many downstream capabilities. **Centralized credential storage** is a foundational failure in this category because MCP servers often consolidate API keys, OAuth tokens, and other service credentials in a single configuration file or secret store. As a result, compromising that host or backing vault through misconfiguration, host malware, or lateral movement can expose valid bearer tokens for every connected

integration. Recent security analyses of MCP deployments describe servers that retain long-lived OAuth tokens for email and enterprise services, warning that a breached server can reuse those tokens to perform arbitrary tool actions and API calls that appear indistinguishable from legitimate agent activity (Pillar Security, 2025). Once attackers control these credentials, downstream APIs treat their traffic as ordinary MCP usage, allowing them to exfiltrate data, send or delete messages, and invoke administrative operations across multiple tools and tenants until credentials are rotated, with blast radius proportional to the number of services consolidated behind that store. **Insecure Secret Management** compounds this by allowing secrets such as API keys and OAuth tokens to leak through operational choices including plaintext configuration files, process environments, debug artifacts, or verbose server logs, and because tool outputs and traces can be fed back into model context, insufficient filtering may propagate credentials into prompts or transcripts where they persist and replicate across agents and workflows Palo Alto Networks (2025). **Overprivileged and Persistent Tokens** widen the blast radius further when tools and servers request broad permission scopes and rely on OAuth tokens that remain valid beyond a single interactive session. Broader OAuth incident reports show that malicious applications can retain access even after password resets or multifactor authentication, because authorized applications and their secrets remain valid until they are explicitly revoked Pillar Security (2025); Proofpoint Threat Research (2025). As a result, a compromised token can become a long-lived capability that allows an adversary to operate across multiple workflows until rotation and revocation fully complete.

Cross-Tool Data Flows shift the exposure from stored secrets to data in motion, where a single hidden directive can steer data from internal tools, resources, or conversation history to external network endpoints within one multi-step transaction. In one documented attack, the assistant URL-encodes the user’s conversation history and sends it as a query parameter to an attacker-controlled domain after retrieving a compromised webpage Palo Alto Networks Unit 42 (2025). In another, a benign-looking server proxies requests to a hidden MCP server, where the model is instructed to exfiltrate environment variables within the same session CyberArk Threat Research (2025). Where cross-tool flows move data actively, **Trojanized MCP Servers** compromise it at the source. A malicious server can capture and modify user queries, tool parameters, file contents, and credential-bearing responses while continuing to advertise the same capabilities to the client. If that server also stores authentication tokens for multiple services, a breach can quickly become a keys-to-the-kingdom scenario, allowing an attacker to invoke tools across email, storage, and calendar APIs on the victim’s behalf Pillar Security (2025). Such malicious servers can also impersonate legitimate integrations such as a chat interface and quietly reroute or tamper with messages while the agent continues to treat the server as trusted Muayad Ali (2025). A single trojanized server can therefore act as a credential theft pivot and poison multi-step workflows simultaneously, causing cascading data leakage and systematic bypass of policy checks.

At a broader scale, **Emergent Aggregation Risks** arise when agents gradually accumulate fragments such as calendar entries, file names, and chat metadata into a shared model context even when each individual tool call appears benign. Work on the mosaic effect shows that records posing little risk individually can reveal identities, routines, or sensitive attributes when combined across sources OSTC (2025). Figure 10 illustrates this directly, where aggregating low-risk outputs from email, calendar, and file servers can expose sensitive project details that no single call would reveal. Because each invocation surfaces only a small fragment, these emergent leaks are difficult to detect since no individual call appears to violate policy. While aggregation risks surface during active tool use, **Residual Data Exposure** persists after the session ends. Tool inputs, outputs, and credentials left in logs, caches, or crash dumps remain recoverable in the absence of sanitization Kissel et al. (2014). In particular, memory dump files generated during crashes may expose cached passwords, encryption keys, and the contents of sensitive files resident in RAM VeraCrypt Project (2015). Thus, an attacker who compromises the host can reconstruct tool inputs and credentials that were intended to be ephemeral. Detailed definitions of each threat are provided in Appendix A.5.

3.6 Authorization and Access Control Gaps

Despite MCP’s growing adoption, its authorization model remains largely binary and session-scoped, lacking granularity, context awareness, and consistency across implementations. Once a tool is approved, the same token can invoke all actions with little contextual constraint and no routine revalidation. Because annotations




Figure 10: An illustration of Emergent aggregation risk in MCP, where an internal agent combines benign looking outputs from multiple MCP servers into a shared context that reveals sensitive information that no single tool response exposes on its own.

are non-binding, the protocol provides no built-in risk tiering, so destructive or high-sensitivity calls do not automatically trigger stronger checks. In human-in-the-loop flows, consent fatigue further weakens least privilege, and in multi-tenant or federated deployments, decentralized policy leads to inconsistent scopes, allow-lists, and logging that creates confused deputy and role leakage paths when per-user scoping is absent. Together with uneven RBAC/ABAC support and missing context-aware rechecks, enforcement becomes unreliable and hard to audit across concrete MCP workflows and deployments.

The most fundamental gap is Binary Permissions, where a single static token issued at approval time can invoke any action a tool exposes. Clients perform no additional per-request or per-action authorization checks, resulting in a coarse-grained, session-wide grant that calls for role- or attribute-based extensions to MCP Narajala & Habler (2025). Practical hardening guidance describes this as a binary approach in which any holder of a valid token can call all available tools, violating least privilege Nikhil (2025), so an attacker who compromises one tool token gains full control over data access, state changes, and destructive operations reachable through the approved tools. Compounding this, **Lack of Risk-Aware Differentiation** means MCP defines no mandatory mechanism for distinguishing high-sensitivity tool calls from low-risk ones, and the specification states that `ToolAnnotations` such as `readOnlyHint` and `destructiveHint` are only advisory hints, warning that clients should never make security-critical decisions based solely on annotations Anthropic (2025c). Recent MCP security frameworks propose Zero Trust-style controls such as just-in-time access and per-request authorization to gate each MCP interaction Narajala & Habler (2025), because without them potentially destructive actions such as writing to production systems are presented at the same privilege level as benign read-only queries, increasing the likelihood that a single misrouted call causes high-impact changes. In human-in-the-loop flows, **Consent Fatigue and Automation Drift** further erode these checks. Repeated permission prompts can desensitize users, causing them to click through approvals or skim long descriptions instead of carefully reviewing what a tool is about to do. Sampling-based attacks further exploit hosts configured to always allow previously approved tools by hiding exfiltration instructions in long story prompts and bypassing human-in-the-loop review entirely CyberArk Threat Research (2025). As a result, sensitive actions such as exfiltrating environment variables or invoking shell access can be executed without fresh user scrutiny.

In shared or federated deployments, **Confused Deputy and Role Leakage** arise when an agent or server with higher privileges is coerced or misrouted into executing actions on behalf of a lower-privileged principal. The MCP authorization specification explicitly warns that servers acting as intermediaries to third-party APIs can be abused through token passthrough and audience validation failures MCP Working Group (2025b). Empirical audits of MCP-enabled agents further show that large language models can be prompted

to chain tools and servers to perform actions such as remote access and credential theft Radosevich & Halloran (2025). Under these conditions, coarse session-wide grants and missing per-user scoping can let seemingly low-risk invocations reach high-privilege capabilities indirectly. This creates opportunities for silent privilege escalation and cross-tenant data exposure while also making responsibility harder to trace in audit logs. Figure 11 illustrates a representative instance of this failure.




Figure 11: An example illustrating the confused deputy problem

Spanning all of the above, **Decentralized and Inconsistent Policy Enforcement** means that each server independently defines its own OAuth scopes, tool allow-lists, logging rules, and update policies, with no central layer to keep configurations aligned. Ecosystem overviews report that this fragmentation leads to security variability and configuration drift, especially in multi-tenant and cloud deployments Emergent Mind (2025); Field, Michael (2025). As a result, two servers exposing the same nominal tool can provide very different protections, and compromising the weakest instance can give an adversary capabilities that users assume are uniformly governed. Even within a single server, **No Context-Aware Revalidation** allows clients to reuse the same OAuth-based access tokens for subsequent tool and resource calls without any mechanism for reauthorizing those calls when runtime context changes MCP Working Group (2025a). Recent work on secure delegation shows that bearer-style OAuth and JWT tokens satisfy identity checks at issuance time but allow any token holder to continue issuing API calls until expiry regardless of changed conditions Goswami (2025). As a result, a token granted in a low-risk context can silently authorize sensitive tool invocations under very different runtime conditions.

Uneven RBAC Support reflects the same fragmentation in authorization design, because the protocol defines no shared RBAC model and leaves access control entirely to individual servers Model Context Protocol (2025n); Zhou (2025). Consequently, some deployments implement role- or attribute-based policies, whereas others expose capabilities behind a single bearer token. As MCP deployments scale, this server-by-server variability produces fragmented authorization that undermines least privilege and makes agent capabilities harder to audit consistently across tools and hosts Posta (2025c). Finally, **Orphaned Sessions and Credentials** persist when hosts and servers shut down or are decommissioned while leaving active sessions, access tokens, or API keys valid. MCP security hardening guides warn that failing to rotate or explicitly invalidate tokens during teardown creates persistent access paths, and they recommend explicit session timeouts together with invalidation of old session identifiers to prevent this class of failure Sadeddin (2025); HiveTrail (2025). Any attacker retaining those secrets can continue calling MCP servers long after shutdown, turning a bounded maintenance window into a long-lived backdoor. Each threat is defined in detail in Appendix. A.6.

3.7 Execution Environment and Sandbox Risks

Sections 3.1–3.6 show how identity, supply-chain, malicious-tool, prompt, data, and authorization gaps can place agents in risky states, and this section turns to the next containment layer, the execution environment. MCP runs local and remote tool logic from shell commands to HTTP calls, which means weaknesses above the runtime quickly translate into machine-level impact if isolation, least-privilege, and resource controls are thin. In practice, insecure defaults and operational drift can expose the host, leak data across boundaries, and degrade availability. Weak provenance and auditing further make the resulting consequences difficult to trace. The risks that follow examine how these dynamics arise in real deployments and why hardening the runtime boundary is necessary for defense in depth.

The most immediate weakness at this boundary is insufficient isolation between MCP tool execution and the host. **Lack of Execution Isolation** is therefore the most immediate gap. In the default STDIO mode, MCP servers run each tool as a local process with the same operating system privileges as the invoking user, and the protocol does not mandate any sandboxing or resource confinement for these processes. A recent security preview notes that such servers inherit host privileges, so a malicious or misconfigured server can access sensitive files, execute arbitrary code, or exfiltrate data without isolation Pellicce (2025). A broader MCP security overview likewise observes that servers typically run with full user-environment privileges and recommends containerization to restore isolation Fenstermacher (2025). As a result, compromise of one unsandboxed server can translate into host-level access, including arbitrary file reads and network connections that affect every workflow and agent depending on that server. Even when MCP servers run inside containers or virtual machines, **Sandbox Escape Vulnerabilities** remain a concern because servers still depend on the underlying container runtime or hypervisor for isolation. Analyses of container escape vulnerabilities such as CVE-2019-5736 in runc show that crafted workloads can overwrite host binaries and gain code execution on the host AWS Compute Blog (2019). An MCP safety audit further demonstrates that once tools can reach the host environment, LLM-driven agents can inject malicious commands into system files and exfiltrate API keys and other secrets from environment variables and configuration files Radosevich & Halloran (2025). In this way, a single server compromise can become a host-level breach that enables persistent backdoors across multiple MCP workflows. This risk becomes even more severe when MCP tools or helper processes run with elevated system privileges. **Root or Elevated Execution** removes the final layer of operating system protection, because running tools and helper processes with root or Administrator privileges means that any tool compromise can immediately give an attacker full control over the host and its log and configuration surfaces. A practitioner-oriented guide lists running MCPs as Root or Administrator among the most common configuration mistakes observed in real deployments SOC Radar (2025). Public advisories for CVE-2025-49596 further describe a remote code execution vulnerability in the MCP Inspector tool that allows unauthenticated attackers to execute arbitrary MCP commands with the same privileges as the inspector process National Vulnerability Database (2025b). As a result, a single exploit or prompt-injection-driven tool chain can escalate to complete host compromise and persistent control over every workflow reachable from that machine.

The runtime boundary can also be breached indirectly through the model itself. For example, **Prompt-Based Remote Code Execution** arises when an agent with access to shell or scripting tools is induced through prompt injection to issue arbitrary operating system commands. The core problem is that the model may interpret injected tool requests as valid high-privilege instructions rather than as untrusted input. An MCP safety study illustrates this risk with a proof-of-concept attack in which an injected prompt causes the agent to append a netcat command to the user’s shell configuration file, creating a reverse shell that reconnects to an attacker-controlled host whenever a new terminal is opened Radosevich & Halloran (2025). In practice, this can enable persistent backdoor installation and credential theft from local files and environment variables across workflows that rely on the affected MCP tooling. Furthermore, runtime failures are not limited to unauthorized execution. They also arise when agents are allowed to consume computational or external service resources without effective limits. **DoS and Resource Exhaustion** capture this failure mode because the protocol itself does not enforce quotas or rate limits on tool invocations. As a result, a misconfigured or malicious agent can repeatedly write large files, flood external APIs, or spawn CPU-heavy tasks until host resources or downstream services degrade or fail Cloudflare (2025). Hardening guidance therefore recommends rate limiting and quotas per user, agent, or session, together with timeouts and

throttling for costly operations, to prevent abusive or runaway workloads Palo Alto Networks Community (2025). In the absence of these controls, a single faulty workflow can exhaust CPU, memory, or third-party quotas and cause denial of service for other tenants.

Runtime exposure also arises through the software stack that MCP servers rely on to execute ordinary tool behavior. Within this software layer, **Vulnerable Third-Party Dependencies** widen the attack surface because MCP servers and tool plug-ins depend on open-source parsers, HTTP clients, and SDKs whose known flaws may be triggered through routine tool use Model Context Protocol Security Working Group (2025). This risk is especially clear in parsing libraries. Unsafe deserialization or entity expansion bugs in JSON, YAML, or XML libraries can be triggered simply by asking an agent to fetch and parse attacker-controlled content. This pattern is illustrated by remote code execution vulnerabilities in widely deployed YAML parsers when they deserialize attacker-supplied data Trauben (2023). A single vulnerable dependency can therefore turn a routine request into remote code execution that compromises the MCP server host and every workflow relying on the affected tool. Finally, **Lateral Movement via Shared Infrastructure** extends the consequences of a single server compromise beyond the initially affected tool. This risk arises when multiple servers share hosts, virtual networks, centralized token stores, or common configuration patterns, so compromising one server can let an attacker reuse shared credentials and network access to pivot into other MCP integrations and internal services that trust the same infrastructure Pillar Security (2025). Security hardening guides therefore recommend isolating instances, segregating secrets, and segmenting network access, precisely because shared infrastructure and common credential stores create practical lateral movement paths once any single server is breached SlowMist Team (2025). Under these conditions, an adversary may harvest OAuth or API tokens intended for other workflows and exfiltrate data from email, storage, and connected services across workflows and tenants. Detailed definitions of each threat are provided in Appendix A.7.

3.8 Workflow and Configuration Risks

MCP simplifies multi-step work across tools and servers, but workflow correctness remains an operational property rather than a protocol guarantee because it depends on consistent naming, versioning, state propagation, and policy interpretation across environments. Small inconsistencies during registration or dispatch can silently change which action runs and where results are routed, and minor state mismatches can accumulate over successive steps until the effective behavior diverges from what operators intended. As deployments grow, heterogeneous server baselines further erode determinism, so the same prompt can produce different behavior across tenants even when workflows appear identical. Weak session isolation, permissive fallbacks, and uneven telemetry further compound the problem by obscuring provenance and making systematic auditing of tool use and results more difficult. The eight recurring failure modes discussed below illustrate how these workflow and configuration breakdowns emerge in practice, with each one amplifying inconsistencies left unresolved by the others.

Among the recurring workflow failures, **Context Drift and State Inconsistency** is especially fundamental because multi-step MCP workflows rely on consistent propagation and update of shared state across tools and servers. If an early step fails silently, only partially commits its results, or is rolled back without compensating other operations, later steps may continue from an inconsistent view of the world rather than the state that actually holds in the underlying systems Chang & Geng (2025); MacVittie (2025). Work on context drift in LLMs further shows that model behavior degrades as the surrounding context naturally diverges from a goal-consistent or pretraining-aligned reference, even when all necessary information remains present, formalizing this drift as a temporal divergence between the model’s outputs and a stable reference distribution Wu et al. (2025); Dongre et al. (2025). In MCP deployments that compose such models into long-running workflows, these inconsistencies and contextual drifts can accumulate across steps and agents, causing downstream tools to act on outdated or contradictory information. Over time, small local update failures can therefore grow into cross-workflow integrity and policy violations that may only be detected long after the fact. The same loss of consistency also appears at the deployment level. **Configuration Drift Across Servers** arises because enterprise MCP deployments rarely consist of a single uniformly hardened server and instead accumulate many instances with small divergences in authentication settings, logging, allowed scopes, and dependency versions. Older or less managed nodes may gradually drift toward permissive configurations

such as anonymous access, broad `allow_all` scopes, and unpatched libraries. A recent MCP risk assessment reports that over 13,000 MCP servers were publicly available on GitHub and describes them as often enabled by default, frequently unauthenticated, and rarely supervised, highlighting how easily shadow MCPs emerge at the edge of organizational visibility Prompt Security Team (2025a). Complementary large-scale analysis of 1,899 open-source MCP servers further finds that 7.2% contain general vulnerabilities and 5.5% exhibit MCP-specific tool-poisoning behavior Hasan et al. (2025). Once configuration drift leaves even a single vulnerable or poisoned server inside an agent’s trust boundary, an attacker can use that outlier as an entry point to harvest secrets, misroute tool calls, and pivot across workflows that implicitly treat all servers in the fleet as equivalently trustworthy.




Figure 12: This is a transactional-integrity gap / atomicity failure - a race-window (TOCTOU-style) policy gap. Because deploy (Server A) and policy apply (Server B) aren't atomic, a hard/soft failure or label-mismatch leaves the pod with default-allow for a window, which an attacker can exploit to reach it (over-privilege window).

Even when individual servers are correctly configured, **Transactional Integrity Gaps** expose workflows to partial failure because HTTP-based MCP servers treat each JSON-RPC request as a stateless call, where the server processes one payload and returns a self-contained response with no protocol-level notion of an atomic multi-step transaction across calls Knostic Team (2025). This limitation pushes rollback and recovery to higher layers. For example, Network Intelligence (2026) places transaction rollback for suspicious sequences of tool calls at the workflow layer, while Chang & Geng (2025) shows that the absence of robust transaction-like guarantees in current LLM planning frameworks leads to inconsistent states and partial failures that must later be repaired. In MCP-based enterprise automations that span multiple services, a single misaligned or aborted run can therefore silently corrupt records or apply access-control changes in one subsystem but not its peers. The result is an integrity or policy violation that may only be discovered after the workflow has completed, as illustrated in Figure 12. Compounding these runtime gaps, **Patch Gaps and Vulnerable Instances** allow MCP servers that reuse reference implementations, connectors, and community-maintained packages to quietly retain exploitable code across registries and internal deployments when centralized patching is absent. Recent reporting by Park (2025) and Lyons (2025) describes a SQL injection flaw in Anthropic’s SQLite MCP server that was forked more than 5,000 times before the repository was archived, with no official patch planned and each fork left to its own maintainer for auditing and remediation. As a result, a single vulnerable MCP server can leave thousands of agents and workflows exposed to persistent remote code execution and data exfiltration even after the original project has been deprecated. These runtime weaknesses become especially consequential when the same infrastructure is shared across users or agents. In multi-tenant deployments, **Session Isolation in Multi-Agent Environments** breaks down when shared MCP servers implement sessions with weak boundaries, so an attacker who can replay, guess, or intercept a session identifier can cause the server to mix contexts between agents and reuse another

user’s memory snapshot. MCP security best practices explicitly warn about such session hijacking scenarios, noting that session identifiers may surface in URLs or logs and that misconfigured session stores can leak data between users and tenants, so session state must be treated as sensitive authentication material Model Context Protocol (2025j). A compromised session can therefore expose private tools, prompts, and memories from one agent to another, leading to cross-tenant disclosure of chat histories, credentials, and retrieved documents across every workflow that shares the affected server.

Beyond isolation failures, operational breakdowns can also emerge in the way MCP systems respond to errors, expose internal state, and retire outdated capabilities. **Unsafe Fallbacks and Error Handling** arise when MCP tools fail and servers return structured error payloads that the host injects back into the model context. In such cases, agents may interpret these error messages as instructions and follow fallback paths that retry operations, call alternative tools, or surface debug information. Recent guidance on MCP error handling documents responses that embed recovery instructions and detailed diagnostic fields, warning that poorly designed handlers can expose internal service details to both the model and end users Barthelet (2025); Hora (2025). Consequently, verbose error responses and improperly handled debug outputs can leak internal URLs, stack traces, and credentials. Attackers who are able to trigger tool failures can exploit these disclosures to exfiltrate data and map internal topology across agents and workflows that share the affected server Hora (2025); Matt (2025).

Across all of these failures, **Operational Visibility and Logging Gaps** make detection harder because effective forensics in MCP deployments requires uniform, high-granularity logs of tool invocations and context updates, yet real-world servers often implement logging inconsistently or omit it entirely. Ecosystem scans report that the vast majority of MCP implementations lack sufficient telemetry to audit what users and agents are doing, leaving operators with substantial blind spots for abuse Sakib (2025). Similarly, Narajala & Habler (2025) frame this weakness as *Insufficient Auditability*, defining it as inadequate logging that restricts detection and investigation of security events, while case studies from enterprise deployments argue that traditional monitoring lacks the context and granularity required for MCP interactions Ithena (2025). As a result, these logging gaps leave operators unable to reliably reconstruct agent behavior or prove what tools ran and what data was accessed across users, workflows, and tenants, turning MCP incidents into opaque failures rather than auditable security events. The same operational weakness also affects the retirement of obsolete capabilities. **Incomplete Tool De-Registration** allows outdated MCP tools or endpoints that are not properly unregistered during shutdown to continue appearing as available capabilities, so agents and orchestration layers may still discover and invoke them after operators believe they have been retired. Research on MCP-aware tool synchronization describes pipelines that treat server tool lists as the single source of truth and include an explicit Remove outdated tools step that deletes entries no longer present on the server, while production MCP server APIs and management CLIs expose operations to remove MCP tool configurations from persistent stores in deployed environments Lumer et al. (2025); OpenSearch Project (2025); Akin (2025); TrustGraph (2025). If these de-registration steps are skipped or fail, stale tool records can linger in indexes and configuration groups, allowing malicious actors to impersonate or re-register those identifiers and route calls through a malicious endpoint across any host or tenant that still trusts the outdated entry Lumer et al. (2025); TrustGraph (2025). Detailed definitions of each threat are provided in Appendix A.8.

3.9 Emergent Threats from Model Behavior and Human Factors

Sections 3.1–3.8 described threats that look static, as if an MCP deployment could be frozen at a single moment and someone could point to the misconfigured tool, missing check, or fragile boundary. Once MCP enabled assistants move into real workflows the picture becomes more fluid. Models respond to long histories of prompts and context rather than a single request, operators adjust safeguards under pressure to ship features, and new agents appear over time to automate tasks that once required direct oversight. Small choices and isolated lapses begin to interact, so that the deployment drifts away from its original design, with sensitive information moving through unexpected channels and effective privileges shifting across tools and agents. This section follows that drift and focuses on emergent risks that grow from model behavior and human factors over time, setting up the privacy, insider, and coordination threats that the rest of this section examines. Detailed definitions of each threat are provided in Appendix A.9 below.

LLM-Centric Privacy Threats. Language models behind MCP agents create privacy risks, especially in deployments where many agents and tools share the same model backend, so probing a single agent can expose information that originated elsewhere Anthropic (2025a; 2024b). In MCP deployments, this surfaces through three main threats: model inversion, membership inference, and pre-MCP data poisoning or backdoored models.

- **Model Inversion Attacks.** Because multiple MCP agents often call into the same model, an attacker can use one agent interface to craft queries that reconstruct sensitive training data or internal examples the model has memorized. Empirical work shows that carefully designed black-box queries can recover private properties or approximate original training records Fredrikson et al. (2015), so a single malicious workflow can leak proprietary data that other tenants and applications may rely on.
- **Membership Inference.** In deployments where MCP agents expose a shared enterprise LLM trained on internal data, an attacker can probe an agent to determine whether particular records, users, or documents contributed to the training corpus. Prior work demonstrates that black-box membership inference can reliably distinguish training points from non-members Shokri et al. (2017), allowing attackers to infer who or what the organization trained its shared model on without ever compromising MCP servers.
- **Pre-MCP Data Poisoning and Backdoored Models.** MCP implementations typically implicitly treat the underlying model as a trusted component, so any covert behavior introduced during pre-training, fine-tuning, or instruction tuning is inherited by every agent that calls it. Studies of data poisoning and instruction-tuned backdoors show that attackers can implant behaviors that only trigger on specific prompts or contexts Anthropic (2024c); Xu et al. (2023), meaning ordinary MCP conversations or tool outputs can activate leaks, override safety instructions, or bias tool selection across all workflows that reuse the compromised model.

Insider and Human-Driven Threats. Insider and human-driven threats arise when people who operate MCP agents and servers weaken security controls through their configuration choices or interactions.

- **Malicious Agent Operators.** Because MCP hosts translate model output directly into tool invocations, operators who wire unvetted servers into agents or grant overly broad permissions effectively expand the agent’s execution authority beyond intended policy Li & Gao (2025). Analyses of MCP ecosystems show that hosts can rely on tool metadata and model output without strong secondary validation, so a malicious or careless operator may cause routine-looking workflows to perform high-privilege actions or tamper with data across the agent ecosystem Guo et al. (2025).
- **Insider Threats and Shadow MCP Servers.** Community guidance and ecosystem studies describe how ungoverned or unauthorized MCP servers can be introduced inside organizations without central oversight, creating “Shadow MCP” instances that access internal resources while bypassing standard monitoring and governance OWASP Foundation (2025b). When employees deploy such servers on organizational infrastructure or register them with agents outside formal review, they can inadvertently expose proprietary data and sensitive operations to abuse, and these paths are difficult for security and compliance teams to track or remediate Li & Gao (2025).
- **Agent Interface Social Engineering.** Prior work on MCP security treats the agent interface and host as an attack surface where adversaries can issue crafted prompts or interaction sequences that steer tool selection and execution Guo et al. (2025). These findings imply that attackers who interact with agents can present seemingly legitimate requests or workflows that drive high-impact tool invocations and data exfiltration under the guise of normal usage, and similar interaction patterns can nudge human operators into approving unsafe actions.

Multi-Agent Coordination and Collusion. As MCP deployments shift from single assistants to graphs of cooperating agents, coordination itself becomes a security-critical surface. Multi-agent topologies can

develop hidden dependencies and emergent behaviors, so a single compromised or misconfigured agent can influence others, bypass local controls, or silently expand system-wide capabilities. In this setting we focus on three coordination risks in MCP-based multi-agent systems: agent collusion, unverified agent handoffs, and orchestration drift.

- **Agent Collusion.** Two or more compromised agents may collaborate to bypass access controls or exfiltrate data stealthily, including via covert channels or steganography in inter-agent messages Motwani et al. (2024). MCP’s support for nested, server-initiated agentic behaviors (*Sampling*) increases the surface for such interactions if oversight is weak Model Context Protocol (2025n).
- **Unverified Agent Handoffs.** When one agent (or MCP server) delegates to another without verifying identity and scopes, it can enable privilege escalation or “confused-deputy” failures; the MCP *Authorization* spec explicitly requires token audience binding and forbids token passthrough to prevent these issues Model Context Protocol (2025l). Anthropic’s multi-agent engineering note likewise emphasizes careful handoffs and coordination boundaries Anthropic (2025b).
- **Orchestration Drift.** Over time, loosely coupled MCP workflows can diverge, accumulate “shadow logic,” or re-enable deprecated behaviors (e.g., tools coming and going) unless guarded by policy and observability. The MCP Tools spec includes change notifications and security considerations, while Anthropic reports emergent coordination issues and drift risks in production multi-agent systems Model Context Protocol (2025p); Anthropic (2025b). Field reports from MCP deployments also document how misconfigurations can silently expand agent capabilities Backslash Security (2025).

The risks described in section 3 show that MCP deployments face a coupled threat surface spanning identity, transport, tools, data flows, and multi-agent behavior, with many threats reinforcing and amplifying one another. Weak authentication or naming collisions can enable poisoned tools, while poor container or dependency hygiene can turn routine parsing into remote code execution. Misaligned lifecycle management or configuration drift further expands the blast radius of a compromise. As a result, MCP operators cannot rely on static policies or one-time audits. Instead, they require continuous, telemetry-driven defenses that track who calls which tool, on what data, and with what downstream effects. The next section moves from taxonomy to countermeasure design, mapping these risks to a defense-in-depth control plane built on secure gateways, zero-trust principles, hardened execution, and continuous monitoring.

4 Mitigation Strategies and Security Controls for MCP

In this section, we translate the threat landscape for MCP deployments into a defense-in-depth control plane. Figure 13 illustrates the overall architecture, where six control families are layered outward from the protected core of MCP tool invocation to enforce security boundaries at multiple levels. Concretely, we group countermeasures into six complementary control families: **Secure Gateway / Firewall**, **Zero-Trust Architecture**, **Containerization & Sandboxing**, **Cryptographic Signing & Registry**, **Fine-Grained Authorization & Policy**, and **Monitoring & Logging**, with the aim of reducing blast radius, ensuring verified execution, and making violations observable and attributable. This control plane aligns with established security frameworks: Zero Trust’s identity and context-sensitive access enforcement (NIST SP 800-207); SLSA supply chain integrity objectives with verifiable provenance; operational guardrails that reflect the OWASP LLM Top 10 classes; and governance/assurance practices consistent with the NIST AI Risk Management Framework Rose et al. (2020); Tabassi (2023b); SLSA Community (2025); OWASP Foundation (2025d).

The ordering of these families follows the typical flow of an MCP request: (i) network-level gateways constrain external traffic; (ii) identity is continually attested via zero-trust principles; (iii) execution units are isolated and their dependencies pinned; (iv) all artifacts and tool interfaces are cryptographically verifiable; (v) runtime privileges are enforced at the narrowest feasible scope; and (vi) telemetry feeds a feedback loop for anomaly detection and policy refinement. Together, these layers form a coherent scaffold in which no single




Figure 13: Six control families form a defense-in-depth architecture for MCP deployments. The outer layers enforce network controls, followed by Zero Trust, containerization and sandboxing, cryptographic signing and registry protections, fine-grained authorization and policy, and monitoring and logging. Together, these controls help contain, detect, and limit security failures.

control is sufficient, but in aggregate they raise the bar for exploitation, limit blast radius, and provide the observability needed for rapid incident response.

For each control family that follows, we will specify: (i) security objectives and threat coverage; (ii) reference architecture patterns suitable for MCP (e.g., identity-aware proxies and default-deny egress for gateways; per-request authentication/authorization for Zero Trust; per-tool isolation for sandboxing); (iii) verifiable controls and how to check them (e.g., signed artifacts with SLSA-conformant provenance and Sigstore/Cosign verification before deployment); and (iv) operational signals and metrics for continuous monitoring (structured logs, policy decision telemetry, anomaly thresholds) to support detection, response, and audit. This structure ensures each mitigation is both principled and testable, bridging high-level policy with concrete enforcement and measurable assurance.

The remainder of this section is as follows: Subsections 4.1–4.6 each develop a control family, stating objectives and MCP-specific patterns. Table 4 links these families to the threat categories defined in Sections 3.1–3.9 of the taxonomy and indicates whether the coverage is primary or indirect, with pointers to concrete checks. Finally, Tables 5 and 6 list the verifiable controls and operational signals for each control family.

4.1 Secure Gateways and Network Controls

In MCP deployments, agents invoke server hosted tools over structured RPC and those calls often produce side effects on the filesystem, the network, or third party APIs. The MCP specifications place the burden for safety at the boundary, which means that servers are expected to validate inputs, enforce access controls, rate limit requests, and sanitize outputs Model Context Protocol (2025k). Authorization guidance further cautions against token passthrough and emphasizes careful separation of token audiences to prevent confused deputy failures Model Context Protocol (2025o). When these requirements are considered alongside well documented prompt injection and insecure output handling risks in LLM applications, transport security, identity assurance, and strict request validation emerge as central security concerns for MCP traffic.

Within this setting, secure gateways and network controls function as primary policy enforcement points for MCP. **Identity aware proxies and gateways** form the outer security boundary by terminating TLS, enforcing encrypted connections to upstream services, applying strong client authentication, and generating

Table 4: Mapping MCP risk categories to mitigation mechanisms across six defense layers. A checkmark indicates a direct control, “indirect” denotes partial or supporting coverage, and “no direct” indicates minimal or no native mitigation.

Risk Category	Secure Gateway /Network Controls	Zero Trust Principles	Containerization & Isolation, and Runtime Defenses	Cryptographic Signing & Registry	Fine-Grained Auth & Policy	Monitoring & Logging
Identity / Trust	✓namespace allow-lists	✓continuous identity checks	indirect	✓verify authenticity	✓per-server trust zones	indirect: detect impersonation
Supply Chain & Code Integrity	✓inspect downloads	indirect: verify components	✓isolate install env	✓prevent tampered packages	indirect: policy gates	indirect: anomaly detection
Prompt / Input Injection	✓block malicious payloads	✓validate inputs/outputs	indirect	indirect: trusts prompt sources	✓block high-risk actions	✓alert prompt attacks
Malicious Tools	✓monitor tool API calls	✓no implicit trust	✓restrict tool actions	✓verify tool code	✓allow/deny lists	✓detect tool misuse
Credential Leakage	✓prevent outbound key leaks	✓least-privilege, short tokens	no direct	no direct	✓scoped API keys	✓alert on key use
authorization (authz) Gaps	✓policy on traffic	✓JIT access	no direct	no direct	✓RBAC, approvals	✓audit trails
Exec / Sandbox Risks	indirect: segmentation limits blast radius	✓micro-segmentation	✓prevent OS escape	no direct	no direct	✓detect sandbox escapes
Workflow / Config Risks	✓uniform network policies	✓config posture checks	indirect: immutable images	✓block unvetted versions	✓central policy	✓flag configuration drift
LLM Privacy Attacks	✓DLP scan on responses	✓least-privilege data paths	no direct	no direct	indirect: limit data exposure	✓alert on PII/inversion
Insider Threats	indirect: logs constrain misuse	✓restrict insider privileges	indirect: limit insider reach	no direct	✓curb insider abuse	✓correlate insider actions
Multi-Agent	indirect: agent allow-lists	✓agent auth, context	indirect: isolate agents	no direct	✓role isolation	✓flag rogue agents

audit quality logs. These transport practices align with MCP guidance, which recommends using TLS for all network transports, enforcing authentication at the transport layer, and validating HTTP origins for remote connections Model Context Protocol (2025q). In practice, zero trust proxies such as Pomerium terminate TLS and can enforce encrypted connections to backend services, ensuring that agent sessions remain confidential and auditable end to end Pomerium (2023). By ensuring that every inbound request is authenticated, encrypted, and recorded before it reaches an MCP server, these gateways reduce the risk of interception, impersonation, and unauthorized access.

Building on that perimeter, **Zero Trust segmentation** constrains which services and networks are permitted to communicate with MCP servers. Rather than allowing broad internal connectivity, segmentation limits lateral movement by isolating MCP components into tightly scoped network zones. Each server and its supporting services should be reachable only from explicitly approved peers, and egress paths should be restricted to required destinations. NIST recommends dividing internal networks into small isolated segments so that traffic entering and leaving each segment can be monitored and controlled Chandramouli (2022). In cloud and Kubernetes environments, operators can further enforce microsegmentation through service mesh controls that apply workload identity and encryption within clusters istio.io (2025); Chandramouli & Butcher (2022). This layered containment makes it significantly more difficult for an attacker to pivot from one compromised MCP component to another. Even within segmented environments, internal traffic must not be implicitly trusted. **Service mesh mutual TLS** extends identity bound encryption to service to service communication inside clusters istio.io (2025); Chandramouli & Butcher (2022). By authenticating

workloads and encrypting their exchanges, service mesh controls ensure that internal MCP calls remain confidential and tied to specific identities rather than to network location alone.

Security hardening must also extend to the hosts themselves. **Host based firewalls and operating system hardening** reinforce network level protections by adopting a strict deny by default posture and minimizing externally exposed services Red Hat, Inc. (2025). Additional operating system level guidance recommends using minimal and hardened installations, enabling kernel level access control mechanisms such as SELinux or AppArmor, disabling unnecessary background services and open ports, and prohibiting direct root logins in enterprise production environments Narajala & Habler (2025). These measures reduce the likelihood that general system vulnerabilities become indirect entry points into MCP workflows.

Finally, security controls must operate at the protocol boundary where individual JSON RPC calls are interpreted. **Application layer gateways and validation** provide this last line of defense by inspecting every request before it reaches tool handlers. Requests should be validated strictly against the MCP schema and rejected if they are malformed, oversized, or out of specification CSF Tools (2025). Additional rules aligned with OWASP input validation guidance can detect suspicious parameter values or invocation patterns that resemble tool poisoning or injection attempts OWASP Foundation (2025c). Granular rate limiting and anti automation controls should also be enforced on all MCP endpoints to reduce abuse and denial of service risks OWASP Foundation (2019; 2023a). Together, these layered controls ensure that MCP requests are encrypted, identity bound, segmented, hardened at the host level, and rigorously validated before reaching sensitive tool execution logic.

4.2 Zero Trust Principles

Many of the MCP specific failures discussed above share a common root. Systems often assume that a component becomes safe once it has been authenticated or placed on a trusted network. Zero Trust Architecture rejects this assumption by refusing to grant implicit trust on the basis of location or prior approval and instead requiring continuous verification of each identity and each action Gambo & Almulhem (2025). In an MCP deployment, this means treating every tool invocation and every data access as a potentially risky operation and reevaluating on each call who is making the request, under what delegation, and with what minimum permissions. When applied consistently, this model directly counters several risks introduced by autonomous agents. Per request verification reduces reliance on long lived bearer tokens, narrowly scoped credentials limit the damage caused by stolen secrets, and ongoing policy enforcement at gateways or authorization services can block harmful tool calls even when the model has already been influenced by prompt hijacking or insider abuse.

Per request identity verification provides the starting point for operationalizing Zero Trust in MCP because every agent call must be treated as a fresh authorization decision rather than an action that inherits trust from an earlier session Ramnarayanan (2025). Each request to an MCP server should therefore be authenticated and authorized independently so that trust is established at the moment of action rather than assumed from prior interaction. In practice, this means that the client presents a user scoped or service scoped token for every request and the MCP proxy or server verifies its signature and critical claims, including issuer, audience, scope, and expiration, against the identity provider Stytych (2025). Requests that lack credentials, carry expired or altered tokens, or violate audience or scope constraints should be denied immediately. This approach prevents stale sessions and long lived bearer tokens from quietly accumulating privilege over time and ensures that access decisions remain tied to the current security context rather than to an outdated trust state. Once identity has been verified on every call, **just in time least privilege access** becomes essential for ensuring that authenticated agents do not retain more authority than they need Curity (2025); ScaleSec (2025). Zero Trust in MCP requires authority to remain minimal and short lived rather than broad and persistent. Tools and data should be accessible only when they are needed and only with the narrowest rights necessary for the current step. An agent may, for example, receive a temporary credential that permits a single database query or a bounded read operation, and that authority should disappear as soon as the task is complete. This design reduces the dangers associated with persistent credentials by limiting access to individual actions rather than entire sessions Ping Identity (2025); Astrix Research Team (2025). As a result, even when credentials leak, an attacker has very little time and very little

scope in which to misuse them. Short lived and narrowly scoped access therefore turns credential exposure from a system wide failure into a much more contained event.

These protections become more effective when they are reinforced by **device and posture checks**. Zero Trust for MCP must evaluate not only who is making a request and what authority accompanies it, but also whether the machine running the agent remains trustworthy enough to participate in the workflow. A remote host should not be accepted simply because it belongs to an internal environment or because it was previously allowed to connect. Before access is granted, the host should demonstrate compliance with a defined security baseline. An MCP gateway can enforce checks on operating system patch status, the presence of endpoint protection, and whether approved agent software is installed. Clients that fail these checks can be denied access or placed in a restricted state so that outdated or compromised machines do not become hidden entry points into MCP tools. This posture aware approach follows Zero Trust guidance that emphasizes continuous monitoring of asset health and removal of access when a device falls out of compliance OWASP Foundation (2021).

Taken together, these controls show why Zero Trust provides a natural foundation for securing MCP systems. Recent MCP specific research reflects this direction directly Ando (2026). The same shift is also becoming visible in industry practice, where major providers such as Google and Microsoft are extending Zero Trust frameworks to cover AI agents in enterprise settings and treating them as first class users or workloads that must satisfy the same authentication, authorization, and policy enforcement requirements as human users and service components Microsoft (2025a); Laurie et al. (2024). Within MCP, this perspective is especially valuable because it replaces static assumptions of trust with security decisions that are revalidated as workflows evolve, components change, and agents act with greater autonomy.

4.3 Containerization, Isolation, and Runtime Defenses

MCP and multi agent deployments should assume that components can be subverted in practice, including compromised tools, unvetted images with supply chain implants, and crafted inputs that steer execution into unsafe behavior. In this setting, containerization and isolation mechanisms aim to minimize blast radius, enable rapid containment and revocation, and preserve high fidelity telemetry for detection and attribution. These objectives are achieved by confining tool and agent workloads to narrowly scoped execution domains with strictly limited ambient privileges across network, filesystem, system calls, identities, and secrets, while simultaneously instrumenting runtime monitoring so that deviations are detected quickly and attributed to concrete principals. To make these mechanisms operational in MCP deployments, the controls described below are structured as a defense-in-depth progression from **prevent** (hard isolation, least privilege) to **contain** (sandbox boundaries) to **detect and respond** (runtime anomaly signals). Implemented together, controls along this progression reduce lateral movement and limit secret exfiltration even under partial compromise, while preserving the operational context needed for swift remediation.

Effective protection begins with **sandbox isolation**, which confines faults to a single execution domain by ensuring that each tool or model runs inside its own container with a deliberately constrained runtime surface. Containers should execute as non root users, unnecessary Linux capabilities should be removed, and the root filesystem should be mounted read only so that privilege seeking behaviors are curtailed and routine writes are prevented Kubernetes (2025b). Because outbound reachability often enables escalation and exfiltration, egress should be restricted to an explicit allowlist so that containers communicate only with endpoints required for their task rather than reaching public networks by default Kubernetes (2024a). Persistence should also be minimized by keeping containers stateless and short lived so that footholds and failed exfiltration attempts disappear with the container lifecycle. To further address supply chain risk, images should be admitted only when they are signed and hash verified prior to deployment Souppaya et al. (2017). These controls ensure that even when a tool is compromised, its operational domain remains tightly bounded. Building on this foundation, **agent confinement** narrows the scope within which the agent runtime itself can operate so that model or tool driven actions remain contained. The agent should run in its own pod with a dedicated service account and the smallest RBAC footprint required for orchestration, which reduces the likelihood that credentials are reused across components Kubernetes (2024b). Host namespace features such as `HOSTPID=FALSE` and `HOSTNETWORK=FALSE` should remain disabled, and `HOSTPATH` mounts should be avoided to limit host exposure Kubernetes (2022; 2025a). Only the minimal volumes and sockets required for

Table 5: MCP controls: verifiable checks and runtime signals (part 1).

Control Family	Verifiable controls & how to check	Operational signals & metrics
Secure Gateway / Firewall	• <i>TLS enforcement</i>: Only strong proto/ciphers; verify via handshake tests; reject downgrades. • <i>Identity-aware proxy</i>: Missing/expired/wrong-audience tokens → 401/403. • <i>Upstream (m)TLS</i>: Remove client cert → handshake fails. • <i>Protocol-aware validation</i>: Malformed/oversized JSON-RPC → blocked (fuzz to test). • <i>Origin validation (HTTP)</i>: Spoofed Origin → denied. • <i>Default-deny egress</i>: Non-allowlisted host/port → blocked.	• Allow/deny counts by reason (unauth, schema, policy). • WAF rule hits (injection patterns, oversize payloads). • TLS handshake failures/downgrades; cert errors. • Blocked egress attempts; new external domains/day. • % requests with complete edge audit (trace/span IDs).
Zero-Trust Architecture	• <i>Per-request AuthN/Z</i>: Expired/replayed token → denied. • <i>Least-privilege, short-lived creds</i>: Inspect scopes/TTL; out-of-scope action → denied. • <i>Continuous verification</i>: Flip device to non-compliant → blocked/quarantined. • <i>Revocation/re-auth</i>: Revoke token; subsequent calls fail until re-auth.	• AuthZ decision logs (who/what/why; policy path). • Token TTL distribution; expired/invalid reject rate. • Posture check failures; deny-rate by resource. • Elevation prompts/approvals per period (spikes flagged).
Containerization & Sandboxing	• <i>Isolated execution</i>: Non-root, readOnlyRootFilesystem, no privilege escalation. • <i>Syscall filtering</i>: Seccomp/AppArmor/SELinux; disallowed syscall → blocked/logged. • <i>Network isolation</i>: Default-deny; off-allowlist egress → blocked. • <i>Trusted images</i>: Cosign/Sigstore verify on pull/start; unsigned → fail closed. • <i>Ephemeral secrets</i>: Scan for hard-coded secrets; reuse JIT secret → fails.	• Seccomp/AppArmor violation counts; HIDS/Falco hits. • Restart/crash loops from policy kills; throttling events. • Blocked egress attempts; inter-namespace anomalies. • % running images with verified signatures/provenance.

MCP input and output should be mounted, and writable paths should be marked **NOEXEC** to discourage code staging. On Kubernetes, this configuration aligns with Pod Security settings that require **RUNASNONROOT**, enforce **ALLOWPRIVILEGEESCALATION=FALSE**, apply a minimal **SECCOMPPROFILE** such as **RUNTIMEDEFAULT** or a tailored alternative, and mandate AppArmor or SELinux policies Kubernetes (2022). Network policies can further restrict the agent to approved MCP servers and internal control endpoints while blocking access to metadata services and public egress by default. In combination, these measures ensure that even if the model is steered toward unsafe actions, the operational impact remains constrained.

While sandbox isolation and agent confinement restrict process scope, **system call hardening** further narrows the kernel interface available to compromised workloads. A deny by default policy for system calls can be enforced using **SECCOMP** or a sandboxing runtime such as gVisor, admitting only the minimal set of calls required by the workload MCP Security Working Group (2025a); SSOJet (2025). In practice this often begins with the **RUNTIMEDEFAULT** profile and proceeds through iterative refinement, first observing violations in audit mode and then enforcing restrictions once stable. High risk operations such as **ptrace**, **mount**, **keyctl**, **bpf**, and **perf_event_open** should be blocked for typical MCP tools Docker (2025); SUSE (2025). To reduce residual exposure, **SECCOMP** can be combined with AppArmor or SELinux type enforcement to constrain file access and process capabilities, and **no_new_privs** can be set so processes cannot gain additional privilege

through execution or setuid transitions Linux Kernel Project (2012). These controls significantly reduce the kernel level attack surface even when arbitrary code executes.

Even with strong preventive controls in place, runtime behavior must still be monitored, which makes **runtime detection** an essential component of this defense model. Host based sensors or eBPF backed tools such as Falco can observe system calls, file activity, and network flows and surface policy violating behaviors including unexpected process launches, unauthorized writes in system directories, or connections to non allowlisted endpoints Sysdig (2023). NIST recommends maintaining a runtime threat profile for each container that tracks processes, network activity, and filesystem changes so deviations are surfaced quickly Souppaya et al. (2017). In Kubernetes environments, these detections can trigger containment actions such as tightening NetworkPolicies, blocking metadata service access, scaling affected pods to zero, and revoking just in time credentials before subsequent invocations MCP Security Working Group (2025a). Through this continuous feedback loop, runtime monitoring reinforces isolation and capability reduction, ensuring that prevention, containment, and response operate as a coherent defense in depth system.

Secret hygiene represents a parallel control plane concern because every host, client, and server in an MCP deployment carries a non human identity, so a leaked token on any one component can rapidly expand into system wide access. Empirical analyses of public MCP server repositories have already identified secrets in approximately five percent of projects, with API keys and bearer tokens left in configuration files, examples, and container images Ferry (2025). To counter this sprawl, operators should adopt a secret broker pattern in which servers and tools obtain short lived, least privilege credentials from a vault at call time rather than relying on hard coded keys. Credentials should rotate automatically and be segmented across distinct wallets or accounts that enforce spending limits for high risk tools MCP Security Working Group (2025b). The orchestration layer should treat secrets as out of band data so that prompts, tool descriptions, logs, and traces carry only opaque handles that MCP servers resolve through authenticated APIs, ensuring that credentials function as scoped and ephemeral capabilities rather than standing privileges MCP Security Working Group (2025b). Beyond credential handling, security must also be enforced at the deployment layer, where platform level controls determine the effective permissions and isolation boundaries of agents and tools. At the platform level, **policy scoped deployment** ensures that agents and tools operate only within environments aligned with their declared permissions. This alignment begins with isolating each MCP component in its own pod and enforcing a strict security context so that execution privileges are narrowly bounded from the outset. In practice, this includes running workloads as non root users, dropping unnecessary Linux capabilities, applying an enforced seccomp profile rather than leaving it unconfined, and mounting the root filesystem as read only. These execution constraints must then be reinforced at the network boundary. NetworkPolicies should restrict connectivity so that pods can reach only the MCP servers, data stores, and control endpoints required for their roles Kubernetes (2024a). By narrowing ingress and egress paths in this way, the platform reduces the opportunity for lateral movement or unintended data exposure. Beyond execution and network isolation, access to orchestration and cloud resources must also reflect the same minimal scope. Fine grained RBAC should ensure that each component receives only the cloud and API permissions strictly necessary for its function rather than broad cluster wide roles Kubernetes (2024b). When these layers are applied together, the declared policy scope becomes enforceable in practice rather than merely descriptive. Finally, any activity that falls outside this declared scope, such as unexpected filesystem access or egress to non allowlisted destinations, should be surfaced through audit logs and treated as a policy violation that triggers automated containment. Together, these controls constrain runtime behavior and ensure that deviations from declared scope are observable and enforceable.

Finally, metadata sanitization is necessary because tool metadata becomes part of the model context and can therefore function as a channel for adversarial instructions. Microsoft security guidance documents a form of tool poisoning in which malicious instructions are embedded in tool names, descriptions, or examples and then surfaced to the model during discovery or selection, steering it toward unintended actions or silent data exfiltration even when the underlying implementation is benign Young & Delimarsky (2025). This risk is not limited to MCP specific tooling. Research on instruction backdoors shows that natural language instructions can encode persistent triggers that reliably alter model behavior once introduced into context Xu et al. (2023). Taken together, these findings highlight that metadata itself must be treated as untrusted input. To reduce this attack surface, MCP hosts should enforce strict schemas and sanitization

Table 6: MCP controls: verifiable checks and runtime signals (part 2).

Control Family	Verifiable controls & how to check	Operational signals & metrics
Cryptographic Signing & Registry	• <i>Signed artifacts</i>: Enforce cosign verify in CI/load; unsigned → rejected. • <i>Provenance checks</i>: SLSA attestations verified pre-deploy. • <i>Server identity</i>: Valid/pinned TLS certs; expired/mismatch → refused. • <i>Allow-listed tools</i>: Unknown tool ID/hash → rejected. • <i>Change control</i>: Tool definition drift → block until approved.	• Sig/provenance failures; registry drift (hash/version). • Cert expiry horizon; cert validation error rate. • Registry access anomalies (bulk pulls, odd hours).
Fine-Grained Authorization & Policy	• <i>Minimal scopes</i>: Inspect tokens/consents; out-of-scope call → denied. • <i>RBAC/ABAC tests</i>: Simulate unauthorized principal → deny. • <i>Allow/deny lists</i>: Non-approved endpoint/domain → blocked. • <i>User consent</i>: High-privilege tools require approval; no approval → cannot run. • <i>No “god” tokens</i>: Audience-bound, service-specific; passthrough → denied.	• Deny/allow counts by rule/tool; surge flags. • Scope-expansion attempts vs baseline. • Policy drift detections; unauthorized edits. • Consent prompts/grants per period (anomalies flagged).
Monitoring & Logging	• <i>Structured audit logs</i>: Immutable/WORM or hash-chained; field presence checks. • <i>Distributed tracing</i>: Correlate edge→server→tool for a sample request. • <i>Alerting pipeline tests</i>: Inject synthetic anomalies; ensure alerts fire. • <i>Retention & time sync</i>: Enforce retention/NTP; simulate log tamper/clock skew → alert.	• Behavioral baselines/anomaly scores (per-tool freq, prompt/output size, fan-out). • Security event volumes: auth failures, policy denials, WAF/HIDS alerts. • MTTD/MTTR; alert acknowledgement/closure SLAs. • Log/trace coverage: % requests with full trace; drop rate. • Exfil indicators: unusual egress volume/destinations/entropy shifts.

on tool metadata and prompts before they are exposed to the model. In practice, this includes constraining lengths, whitelisting allowed characters, stripping markup, and validating structured fields against defined schemas, with validation failures treated as hard errors OWASP (2023). While sanitization alone cannot eliminate prompt based and metadata based attacks, it establishes a defensive baseline in which all content entering the agent context is assumed untrusted and filtered as early as possible.

4.4 Cryptographic Signing & Registry

A core pillar of MCP security is ensuring the authenticity and integrity of communicating components, including tools, servers, and clients, through cryptographic verification. Without strong signing and validation mechanisms, MCP ecosystems remain vulnerable to tool poisoning, server spoofing, and broader supply chain compromise. This is especially important in MCP because agents routinely rely on externally provided tool definitions, manifests, and server metadata during ordinary workflows. As a result, trust cannot rest on names, endpoints, or transport alone, but must instead be grounded in verifiable provenance that allows clients and operators to confirm who produced an artifact and whether it has been altered since publication Microsoft (2025b).

This foundation begins with **cryptographic identities and digital signing of artifacts**. In a hardened MCP deployment, tools and servers do not appear as anonymous binaries but as components with verifiable cryptographic identities that clients can authenticate before accepting code or configuration Microsoft (2025b). By signing tool definitions, manifests, and server binaries with trusted keys, for example through X.509 certificates or Sigstore based attestations, providers make authenticity and integrity directly verifiable at the point of installation and update. Clients should validate these signatures as a mandatory admission step and treat any missing, invalid, or untrusted signature as a hard failure. Framed this way, signing is not merely a packaging convenience but a direct control against tool poisoning, server spoofing, and unauthorized deployment. Once artifact identity is established, **trusted tool registries** provide the operational mechanism for carrying that trust into routine discovery and deployment workflows. Tool discovery in MCP should therefore begin from a registry that the organization explicitly trusts rather than from arbitrary public catalogs. This registry should serve as the supply chain anchor for MCP servers and tools by recording expected versions, hashes, and cryptographic signatures, and by storing these records in tamper evident forms such as immutable ledgers or signed Merkle tree based transparency logs so that attempts to rewrite release history become detectable Sigstore (2022). Under this model, only registry listed artifacts should be treated as eligible for discovery or execution, sharply limiting the ability of unvetted tools to enter day to day workflows. Building on registry based control planes proposed in prior work, administrators should govern onboarding, enforce fine grained admission policies, and apply automated trust scoring so that squatted or malicious tools are identified and blocked before they are surfaced to agents Narajala et al. (2025).

These registry controls must be complemented by **server authentication** so that trust in artifacts is matched by trust in the endpoints that actually serve them. MCP hosts should treat every server connection as untrusted until it is authenticated and should rely on TLS certificates so that clients connect only to endpoints whose identities they can verify Paktiti (2025). On top of these transport level checks, server binaries, scripts, and container images should also be signed and verified before deployment or update in line with secure software development guidance Souppaya et al. (2022b). In service to service settings, hosts should further enforce mutual TLS with strict certificate validation or pinning and reject self signed or mismatched identities so that only authenticated MCP servers participate in the ecosystem. This matters because transport authentication confirms the endpoint being reached, whereas artifact signing confirms the integrity of the code and configuration that endpoint is actually running.

Trust must also persist after deployment, which makes **runtime signature verification** an essential continuation of these earlier checks. MCP clients should not assume that installation time verification is sufficient because a registered tool can later be modified while retaining the same name, endpoint, or apparent version. Whenever a tool definition changes, even if the version number does not, the client should treat the update as untrusted and require explicit revalidation or a fresh approval workflow before execution proceeds. This kind of continuous verification is particularly important against rug pull style attacks in which a provider silently alters a tool after users have already granted consent Bhatt et al. (2025). By reattaching trust to a specific signed state rather than to a static tool identity, runtime verification helps ensure that post approval drift does not quietly reenter trusted workflows.

4.5 Fine-Grained Auth & Policy

MCP implementations should adopt least privilege so that each tool operates only with the scopes or permissions it actually needs. Because MCP relies on OAuth like tokens to carry these permissions, those tokens must remain tightly scoped and closely aligned with the intended function of each tool Lodderstedt et al. (2025). A mail reading tool that only needs read access to inbox messages, for example, should never receive credentials that also permit mailbox deletion. In practice, this means that authorization must be attached to the specific action being performed rather than granted broadly at the level of the overall agent or session. Role based and attribute based access control can help enforce these boundaries, while policy frameworks such as Open Policy Agent and AWS Verified Permissions make it possible to attach contextual rules to each tool or token OPA (2025); AWS (2025). The MCP authorization specification builds on OAuth 2.x and therefore allows implementers to encode permissions directly in scopes and JWT claims. Prior work further recommends including an explicit list of required OAuth scopes in each tool’s signed definition so that the client requests exactly those scopes and the user approves them at consent time Bhatt et al. (2025).

Together, these mechanisms turn least privilege from a general principle into an enforceable authorization boundary for day to day MCP operation.

Scoped API tokens provide the most direct way to enforce this boundary during execution. If a single bearer token can drive many MCP tools and servers, then one compromise is enough to unlock large parts of the deployment and enable lateral movement across workflows. To reduce this risk, the host should treat each tool invocation as a request for a narrowly defined capability rather than as an extension of a broad standing session. Before the call is issued, the authorization layer can mint an OAuth access token or API key whose audience is limited to one MCP server and whose scopes match only the specific operation that tool is permitted to perform, such as read only access to a constrained file namespace or a single portfolio quote endpoint. The server must then enforce this contract by validating audience and scope on every request and rejecting calls that exceed the policy attached to the tool. Applied consistently, this pattern replaces long lived multi purpose credentials with short lived fine grained tokens, which means that even if one token is exfiltrated it cannot be replayed to invoke other tools, reach unrelated servers, or silently expand agent privileges in downstream workflows Lodderstedt et al. (2025); Paktiti (2025). This execution level scoping becomes stronger when it is backed by **RBAC and policy engines** that govern authorization as a dedicated control layer across the deployment. In MCP environments, principals such as users, agents, and tool servers should be associated with roles or attributes, and every request should be evaluated against explicit policy before any tool or resource is invoked Ferraiolo et al. (2001). Building on this model, a policy engine such as Open Policy Agent or Cedar can receive the request context and determine whether the invoking principal is permitted to perform the requested operation on the target resource Open Policy Agent Contributors (2025); Cutler et al. (2024). This makes access control declarative, auditable, and easier to evolve over time. It also allows the authorization layer to incorporate contextual signals such as time, source location, or recent activity, which is important because the same action may be acceptable in one context and unsafe in another. In this way, policy engines extend least privilege beyond static role assignment and support context aware authorization that remains aligned with runtime conditions.

The same principle also applies earlier in the workflow through **allow and deny lists**, which constrain what tools and servers can enter the trusted operating set in the first place. In stronger MCP deployments, tool and server selection should be treated as an allowlist decision rather than as an open ended discovery process driven by whatever appears in public registries Cyber Sierra Knowledge Team (2025). A practical approach is to maintain a curated set of MCP servers, tools, and domains for each environment, pinned by origin, registry, and stable identifiers such as image digests or UUIDs, so that agents tend to invoke only components that have already been reviewed Cyber Sierra Knowledge Team (2025). When a new server or tool is proposed, it can pass through a lightweight intake process in which security and legal stakeholders review provenance, licensing, and data handling before it is added to the allowlist, while anything unvetted remains denied by default. At runtime, the host should reinforce this stance by preferring approved registries and rejecting discovery or fetch requests that target unapproved sources or suspicious look alike names. This reduces the chance that agents quietly bind to opportunistic or high risk endpoints during ordinary operation Enkrypt AI (2025). Even after a tool or server has passed these admission checks, trust should not be treated as permanent or unconditional. **User approval workflows** provide an additional safeguard by making privilege expansion visible at the point where new capabilities are about to be exercised. When an MCP host encounters a tool for the first time or detects a new version, it should pause automatic use and present a clear approval request to the user or administrator. That request should explain what data the tool can read or write, which external systems it can contact, and what irreversible actions it can trigger, so that the approver can compare the requested privileges with the intended use. If the requested scope is broader than necessary, the workflow should highlight that mismatch and either constrain the permissions or escalate the decision to a higher trust tier. In high assurance deployments, these approval flows align with MCP guidance that emphasizes explicit user consent for tool and resource access and are reinforced by recommendations for human in the loop approval during sensitive or high impact operations Humanloop (2025); Wiz Research (2025). By introducing a review checkpoint when capabilities change or expand, approval workflows help ensure that least privilege remains enforceable throughout operation rather than only at initial admission.

4.6 Monitoring and Logging

Robust monitoring and logging provide the autonomous detection and forensic layer for MCP deployments by turning each authentication event, tool invocation, and data transfer into security evidence that can be queried and correlated. When telemetry is treated as a first class signal rather than as a by product of execution, hosts and operators gain the ability to detect misconfigurations and attacks in near real time while also preserving the traces needed to reconstruct how threats emerge and propagate across tools, servers, and agents. This visibility begins with **comprehensive event capture**. Instead of recording only final outcomes, deployments should treat every major control point as a source of evidence. Authentication attempts, tool invocations, prompt characteristics such as size and selected tools, configuration edits, policy updates, and security control decisions should each produce structured traces. These records should then flow into immutable append only or WORM storage so that incident responders retain evidence they can trust when reconstructing what occurred Kent & Souppaya (2006). Once this telemetry stream is established, operators can route it through a centralized and correlated pipeline such as a SIEM that aggregates feeds from platforms like Splunk or Elastic, allowing analysts to follow an attacker across components and piece together a coherent incident timeline Rouse (2023). And, as monitoring moves closer to execution, **runtime syscall and container telemetry** would extend this evidence layer into the underlying infrastructure. In Kubernetes based MCP deployments, hosts and servers should not be treated as opaque runtime units because the underlying nodes can observe each workload through a stream of low level signals. Kernel level and container aware sensors such as eBPF based probes and Falco daemons can monitor system calls in real time and enrich each event with container, pod, and namespace context Di Pietro & Lombardi (2021); Sysdig (2023). Building on this visibility, Falco applies rule driven policies to characterize how an MCP component behaves over time and raises alerts when observed patterns depart from established baselines, such as an unexpected shell spawn, a write into a protected directory, an attempted privilege escalation, or an outbound socket to an unusual destination Di Pietro & Lombardi (2021); Sysdig (2023). In Kubernetes environments, this style of monitoring aligns naturally with existing isolation boundaries because pod scoping and namespaces provide clearer behavioral units, making early signs of compromise easier to distinguish from routine system noise GCP (2023).

This infrastructure level telemetry becomes even more valuable when connected to **identity aware request tracing** at the protocol boundary. In MCP deployments, the first policy decision about whether an agent may reach a given server often occurs at an identity aware proxy or gateway rather than inside the tool itself. At that boundary, each JSON RPC request and connection attempt can be logged as a normalized event that captures principal identity claims, token status and scope, source IP and device context, the targeted MCP server, tool or resource, the invoked method, observed latency, and the final allow or deny outcome Pomerium (2023). As these records accumulate, they form an identity centered trace of MCP activity that makes tool use, context access, and configuration changes directly observable. They also align naturally with compliance requirements around access monitoring and change traceability, which strengthens the evidentiary basis for frameworks such as SOC 2 and HIPAA in MCP based systems. Once request level and runtime telemetry are in place, **behavioral baselining and anomaly analytics** help distinguish ordinary activity from suspicious deviations. In MCP deployments, behavioral models can characterize routine operation by learning statistical or machine learning baselines over signals such as per tool call frequency, typical prompt length and structure, response size, inter request timing, and the chains of tools and servers that recur in normal workflows ScaleSec (2025); Mooney (2025). With such a baseline established, departures from it become meaningful indicators of possible misuse. A sudden spike in invocations of a previously quiet tool, a sequence of abnormally long prompts directed at a sensitive server, or an unexpected fan out across many downstream resources can all indicate prompt injection, unauthorized tool chaining, or covert enumeration. These patterns can then trigger targeted investigation or automated containment before the behavior develops into a broader incident.

A closely related extension of this analysis is **data exfiltration pattern analysis**, which focuses on how anomalous outbound behavior may reveal covert leakage and therefore connects naturally to the broader problem of telemetry integrity. MCP operators should correlate outbound request volume, destination cardinality, payload size variance, and time of day deviations to identify possible exfiltration channels. Outlier egress flows become especially concerning when they target newly observed domains or exhibit abrupt entropy shifts in transferred data, since such patterns may indicate staging, covert encoding, or

unauthorized movement of sensitive information. Detecting these anomalies early allows the monitoring layer to trigger containment workflows before large scale leakage is completed. For these signals to remain operationally useful, however, the underlying telemetry must also be collected and preserved in a trustworthy form. Deployments therefore need **sanitization and integrity controls** that protect both evidentiary value and operational safety by scrubbing or tokenizing credentials, secrets, and personally identifiable information before logs are aggregated, while still retaining structural features such as field names, lengths, and categorical tags that are necessary for investigation. At the same time, each log record should be cryptographically hashed, or linked through chained hashing, so that tampering becomes detectable if any record is altered after collection. Trust can be strengthened further by periodically anchoring hash roots to an external trusted timestamping service or ledger. Supported by accurate time synchronization and retention policies aligned with incident response requirements and regulatory minima, these controls ensure that telemetry remains both safe to store and reliable enough to support forensic reconstruction and accountability.




Figure 14: End-to-end threat lifecycle for the RADE attack, illustrating the dual-axis taxonomy in action. Stage 1 shows the attack entry via a poisoned file upload. Stage 2 traces propagation through the MCP file server, hidden instruction execution, cross-tool chaining, and credential exfiltration. Stage 3 applies the dual-axis threat taxonomy introduced in Section 3 to classify the attack by what (Retrieval-based Poisoning), where (Host, Client, Tools, Server, Data Layer), and when (Invocation and Execution phases), yielding the resulting threat class Prompt, Input, and Context Injection. Stage 4 maps this classification to the Section 4 control plane, distinguishing four direct controls (Secure Gateway, Zero Trust, Fine-Grained Authorization, Monitoring and Logging) from two indirect controls (Containerization, Cryptographic Signing) per Table 4.

Taken together, these control families show that MCP security depends on coordinated defenses across the full request lifecycle, not on any single safeguard. Secure gateways reduce exposure at the boundary. Zero trust, cryptographic signing, and fine grained authorization restrict which identities, tools, and artifacts can be trusted and under what conditions. Containerization and sandboxing help contain compromise and reduce blast radius when prevention fails. Monitoring and logging preserve the visibility needed to detect misuse, investigate incidents, and support recovery. Figure 14 makes this concrete by tracing the RADE attack path from poisoned input to cross tool propagation, and by showing how layered controls can disrupt that path at multiple stages.

5 Structured Evaluation of Verifiable Checks

After introducing the control families, we now test whether they can be implemented as concrete checks in a real MCP deployment. The point of this evaluation is not to build a large benchmark or to claim broad empirical coverage. Instead, it examines a narrower question that follows directly from the framework developed in the earlier sections. If the taxonomy is useful operationally, then checks derived from it should be able to detect representative MCP failures, place those failures at the correct component and lifecycle phase, and leave behind evidence that an operator could inspect later. In that sense, this experiment connects the descriptive taxonomy in Section 3 to the control-oriented view developed in Section 4.

Table 7: Benchmark scenarios used to evaluate executable MCP controls

ID	Scenario	Viol.	Component	Lifecycle	Expected evidence
S1	Over-broad file scope exposed through <code>read_file</code>	Yes	Tools	Creation / Registration	policy violation
S2	Blocklisted outbound fetch denied by egress policy	Yes	Tools	Invocation / Execution	policy violation
S3	Poisoned tool description rejected during discovery	Yes	Tools	Discovery	metadata violation
S4	Unapproved tool hash rejected at admission	Yes	Tools	Creation / Registration	hash mismatch
S5	Out-of-scope token denied for a privileged call	Yes	Auth. Infra.	Invocation / Execution	authorization denial
S6	Unauthorized principal denied for a privileged call	Yes	Auth. Infra.	Invocation / Execution	authorization denial
S7	Post-approval drift detected through hash divergence	Yes	Tools	Update / Maintenance	hash mismatch
S8	A real egress denial occurs but the required denial trace is missing	Yes	Server	Update / Maintenance	observability gap
S9	Approved tool with valid scope and metadata	No	Tools	Creation / Registration	clean execution
S10	Authorized HTTP call with a valid token and complete evidence	No	Auth. Infra.	Invocation / Execution	clean execution
S11	A post-approval integrity failure yields a hash mismatch although a naive phase mapping favors creation and registration	Yes	Tools	Update / Maintenance	hash mismatch
S12	An over-broad scope failure surfaces as an authz denial although the root cause lies in tool configuration	Yes	Tools	Creation / Registration	authorization denial

The testbed is intentionally small, but it exercises the main interactions that matter for the taxonomy. It contains one host, one MCP client, and three servers that expose different kinds of risk. The first server provides a directory-scoped file tool over `stdio`. The second provides a network-facing fetch tool over HTTP with an egress allow-list. The third provides a privileged secret retrieval tool behind token and scope checks. Checks derived from the framework are then applied at the points where the system can actually intervene. Scope checks reject requests when file access, destination rules, or token privileges exceed policy. Integrity checks compare approved and observed tool hashes, which makes it possible to catch both unapproved admission and later drift. Discovery checks block poisoned or policy-violating tool descriptions before invocation. Observability checks verify that selected denials leave behind the trace or audit artifact an operator would expect to review afterward. The benchmark therefore does more than record whether a

Table 8: Final benchmark results

Metric	Dual-axis	Lifecycle-only	Component-only
Violation detection rate	100%	100%	100%
False positive rate	0%	0%	0%
Localization accuracy	100%	50%	80%
Evidence completeness	100%	100%	100%
Localized / detected	10/10	5/10	8/10

request succeeds or fails. It also asks where the failure belongs, when the relevant control should have fired, and what evidence should remain once the event is over.

The benchmark contains 12 scenarios in total, including 10 violations and 2 benign cases as shown in Table 7. The violating cases were chosen to cover the kinds of failures the framework is meant to distinguish, including over-broad scope, blocked outbound fetches, poisoned metadata, integrity failures at admission, authorization denials, post-approval drift, and an observability-gap case. Two additional cases were designed to separate single-axis from dual-axis reasoning more clearly. In one of them, a post-approval integrity problem produces a hash-mismatch signal that a naive mapping associates with creation and registration even though the correct phase is update and maintenance. In the other, an over-broad scope problem appears as an authorization denial even though the root cause lies in tool configuration rather than in the authorization infrastructure. Each scenario is annotated with a ground-truth failure mode, component label, lifecycle label, expected enforcement outcome, and expected evidence artifact. The two benign cases play an equally important role because they confirm that valid registration and valid authorized execution do not trigger false alarms.

The value of the dual-axis view becomes clearer when it is compared with two reduced baselines. One baseline assigns only a lifecycle phase using naive signal-based inference from runtime output. The other assigns only a component using the same kind of runtime signal. By contrast, the proposed method uses the taxonomy to trace the failure back to its source and assigns both a component and a lifecycle phase. This difference matters because the visible symptom and the true enforcement point often do not coincide. A denial at runtime may reflect a scope problem that should have been constrained earlier. A hash mismatch may be observed only after the system is already running. An observability failure may appear only when a real denial occurs but the expected trace is missing. For that reason, performance is evaluated along four dimensions, whether a method flags real violations, whether it avoids false alarms on benign cases, whether it localizes failures correctly, and whether the expected audit artifact is actually produced. For the dual-axis method, correct localization requires both the right component and the right lifecycle phase. For the reduced baselines, correctness is defined only on the single axis each one models.

Table 8 reports the final results. All three methods detect all 10 violating scenarios and raise no false positives on the 2 benign cases, so the main difference does not lie in simple detection. It appears in localization. The dual-axis method localizes all 10 violating cases correctly, while the lifecycle-only baseline localizes 5 of 10 and the component-only baseline localizes 8 of 10. The asymmetry is important. Component-only reasoning still retains part of the failure location, but lifecycle-only reasoning degrades more sharply when the visible signal does not match the phase in which the relevant control should have fired. This is clearest in S1 and S12, where the observed denial signal is associated naively with invocation and execution even though the root cause belongs to creation and registration. S7 and S11 expose a different kind of mismatch. In those cases, the hash-mismatch signal is mapped naively to creation and registration, but the true phase is update and maintenance because the tool had previously been approved and then drifted. S8 reveals a related but slightly different failure. There, the runtime denial is real, but the benchmark then checks for the required server-side denial trace and finds that it is missing, which makes the failure an observability problem rooted in the server and its maintenance state rather than only a policy-enforcement event.

S12 also shows why component-only reasoning is insufficient. The system returns an authorization denial, but the root cause lies in tool configuration and registration rather than in the authorization infrastructure itself. The dual-axis method resolves these ambiguous cases by using failure-mode annotations from the scenario taxonomy to trace each failure back to its source, whereas the single-axis baselines rely only on the raw response signal. These are precisely the situations in which a single-axis view can still detect that something went wrong while losing track of where responsibility lies or when the control should have fired. Evidence completeness is 100% in the final run because each violating case produces the expected artifact. Taken together, these findings show that the benchmark primarily serves to validate whether the framework can localize failures correctly and preserve actionable evidence.

The benchmark remains intentionally compact and should therefore be interpreted as a structured validation exercise rather than as a comprehensive suite. Even with that limitation, it shows that the framework has operational value beyond description alone. Checks derived from the taxonomy can be implemented in a real MCP stack, they can detect representative failures reliably, and they can place those failures more accurately than reduced single-axis views while preserving concrete audit evidence for debugging and review. Larger multi-host and multi-agent settings may reveal further interactions, but the central result already holds in this smaller deployment. The additional structure in the taxonomy does not merely reorganize known risks. It improves how failures are localized and how the resulting evidence can be used.

6 Emerging Frameworks and Protocol Extensions from MCP Security Research

Having mapped where and when MCP specific threats emerge through our dual axis taxonomy in Section 3, attached concrete controls and verification hooks to each failure mode in Section 4, and evaluated those checks in a compact benchmark in Section 5, we now turn to how these ideas are beginning to take shape in real systems. At this point, the question is no longer only which threats exist or which controls are desirable in principle, but how recent MCP security proposals are translating them into practical mechanisms. Seen in this way, emerging frameworks can be read as early attempts to operationalize parts of the control plane developed in the previous sections, with each emphasizing a different point in the MCP lifecycle and a different portion of the attack surface. Rather than exhaustively surveying every prototype, we focus on a small set of representative efforts that show how runtime guards, Zero Trust patterns, protocol-aware intermediaries, and security-conscious interfaces are beginning to crystallize within the MCP ecosystem.

One prominent direction is represented by **Model Contextual Integrity Protocol**, proposed by Jing et al. (2025), which augments MCP with structured tracking and a dedicated guard model along the tool invocation path. MCIP introduces a Model Contextual Integrity log format that records each MCP interaction as a trajectory of information flows, giving tool calls and their surrounding context a structured representation that can be audited and reused as labeled telemetry. On top of this telemetry, the MCIP Guardian model classifies trajectories and individual tool invocations as safe or as one of several interaction phase risk types, allowing the system to block or flag unsafe calls even when they are syntactically valid. Through the lens of our taxonomy, this framework most directly addresses Prompt Input & Context Injection threats such as Poisoned Prompt Templates, RADE, Format Embedded Poisoning, and related interaction phase misuse of tools on the client and server sides. At the same time, it provides a concrete realization of the Monitoring & Logging and Fine Grained Authorization & Policy control families discussed earlier, since the MCI logs instantiate the structured telemetry assumed by those controls and the Guardian model acts as a runtime decision point for individual tool calls.

A related but more infrastructure centric direction appears in the **MCP Gateway Architecture** described by Brett (2025). In this design, a gateway sits between model hosts and self hosted MCP servers as a reverse proxy, terminating TLS, centralizing OAuth 2.1 authentication, and funnelling all JSON RPC traffic through a single inspection and routing point with secure tunnels to backend servers. This shifts security enforcement toward the network and protocol boundary, which makes the gateway a natural embodiment of the Secure Gateway / Firewall family developed in Section 4. In particular, it aligns closely with the Application Layer Gateways and Validation and Host based Firewalls and OS Hardening strategies, since the gateway becomes the protocol aware choke point through which requests are authenticated, inspected, rate limited, and routed before they ever reach an MCP server. Because OAuth enforcement and scoped tokens are attached at this

boundary, the design also supports the Zero Trust controls of Per Request Identity Verification and Just In Time Least Privilege Access. As a result, the gateway pattern directly mitigates several Identity & Trust Management threats, including Server Name Collisions and Impersonation, Mutual Authentication and Encryption Gaps, Lack of Verified Server Identity or PKI, and Implicit Trust in Unauthenticated Servers, while also reducing part of the blast radius for Prompt Input & Context Injection and Credential Leakage and Data Exposure by ensuring that cross boundary traffic is constrained and logged at a single ingress point. This emphasis on the gateway as a controlled choke point is pushed further by **MCP Guardian**, introduced by Kumar et al. (2025a), which presents a security first middleware layer that proxies traffic between MCP clients and servers rather than simply relaying it. In this design, all MCP requests flow through the guardian, which enforces authenticated sessions with scoped credentials, applies strict per token and per tool rate limiting, and records end to end logs and traces for every call. At the same time, WAF style inspection examines JSON RPC payloads for malformed structure, oversized bodies, and common injection patterns before they reach a server. Compared with the broader gateway architecture discussed above, MCP Guardian places even greater weight on active security enforcement at runtime, and therefore serves as an especially direct realization of the Secure Gateway / Firewall, Zero Trust Architecture, and Monitoring & Logging control families. It concretely instantiates the Application Layer Gateways and Validation, Per Request Identity Verification, Just In Time Least Privilege Access, Comprehensive Event Capture, and Identity Aware Request Tracing strategies described earlier in the paper. Through this layered enforcement model, MCP Guardian hardens several categories in our taxonomy at once, including Identity & Trust Management failures such as Server Name Collisions and Impersonation and Implicit Trust in Unauthenticated Servers, as well as Malicious or Poisoned Tool Servers and Prompt Input & Context Injection threats whose payloads can be filtered at the edge. It also improves visibility into Credential Leakage & Data Exposure by routing sensitive flows through a single observable middleware layer rather than through multiple ad hoc endpoints.

While these gateway oriented systems concentrate primarily on protocol boundaries and live request enforcement, **Enterprise Grade Security**, proposed by Narajala & Habler (2025), broadens the focus to a defense in depth blueprint for hardening MCP deployments as a whole. Their framework begins from a systematic threat model and adapts NIST and OWASP guidance to the MCP stack, arguing for MCP aware application gateways, continuous identity verification, host based EDR and HIDS, segmentation, and rich telemetry for intrusion detection. This broader framing is important because it connects several control families that are often treated in isolation and instead presents them as mutually reinforcing parts of one deployment model. Under our taxonomy, this blueprint directly addresses Authorization & Access Control Gaps such as Decentralized and Inconsistent Policy Enforcement, No Context Aware Revalidation, Uneven RBAC Support, Binary Permissions, and Confused Deputy or Role Leakage, while also covering Supply Chain & Code Integrity risks such as Trojanized or Compromised Installers and Unverified Package Sources and Automatic Invocation. From the perspective of Section 4, it operationalizes the Secure Gateway / Firewall, Zero Trust Architecture, Cryptographic Signing & Registry, Fine Grained Authorization & Policy, and Monitoring & Logging families through segmented deployments, hardened identity paths, MCP aware gateways, and SIEM backed intrusion detection over MCP traffic.

A different form of operationalization appears in **MCP Safety Scanner**, developed by Radosevich & Halloran (2025), which approaches MCP security through empirical auditing rather than through an inline enforcement layer. Their work shows that LLMs connected to standard MCP servers can be steered into malicious code execution, remote access control, and credential theft during routine workflows. Building on the `McpSafetyScanner` codebase, they instantiate an agent based auditor that discovers a server’s tools, crafts adversarial test flows, and returns a structured report of vulnerabilities and suggested remediations that operators can run before deployment. This shifts attention from runtime mediation to pre deployment and recurring assessment, which makes the framework especially relevant to the interaction between Prompt Input & Context Injection, Malicious or Poisoned Tool Servers, and Credential Leakage and Data Exposure during Invocation / Execution. In terms of our mitigation strategies, it points toward MCP specific scanning as a practical extension of the Cryptographic Signing & Registry and Monitoring & Logging families, where recurring audits and adversarial checks complement registry hygiene, version pinning, and runtime telemetry by catching high impact misconfigurations before they become cross workflow incidents.

This auditing oriented perspective is developed further by **MCP Scan**, introduced by McGinley-Stempel (2025), which treats tool manifests and installed MCP configurations as a supply chain that should be inspected rather than blindly trusted. The `mcp-scan` CLI inspects local MCP configuration files for clients such as Claude, Cursor, and Windsurf, enumerates all registered tools, and runs static checks over their descriptions and metadata to flag known prompt injection and tool poisoning patterns. It also detects cross origin tool shadowing, in which one server silently overrides another’s tools, and identifies rug pull scenarios by comparing current tool hashes against previously known values. In practical terms, it plays a role similar to `npm audit` for MCP and can already be embedded into CI pipelines so that changes to tool manifests or newly added servers are scanned automatically before deployment. Through our taxonomy, this framework strengthens Prompt Input & Context Injection, Supply Chain & Code Integrity, and Malicious or Poisoned Tool Servers by turning poisoned metadata, cross origin escalation, and silent post approval redefinition into concrete findings that can block a release. From the perspective of Section 4, it gives operators a practical way to realize the Cryptographic Signing & Registry family in everyday workflows, particularly Trusted Tool Registries, allow listed tools, and change control, because its hash based checks and drift reports mirror the guidance to reject unknown tool identities and block execution on unapproved definition changes even though the current implementation remains primarily static and rule based.

Taken together, these frameworks show that the dual axis taxonomy and control plane developed in this paper already map onto concrete mechanisms such as structured context logs, guard models, protocol aware gateways, middleware enforcement layers, and supply chain scanners. At the same time, the proposals surveyed here tend to concentrate their effort on a relatively narrow band of risks centered on prompt injection, malicious or misconfigured servers, and unsafe tool use during live interaction. Many also define their own policy schemas, telemetry formats, or trust assumptions rather than converging on shared standards, and only partially address lifecycle wide hardening, multi tenant deployments, and multi agent coordination. As a result, real world MCP deployments will still need to assemble these point solutions into more standardized patterns for authentication, authorization, evaluation, and continuous assurance so that defenses can be composed, compared, and audited across hosts and servers. The next section therefore moves beyond individual systems and turns to broader research directions, with particular attention to common security schemas, shared benchmarking infrastructure, and more robust self monitoring behavior in MCP based ecosystems.

7 Limitations and Future Directions

This work has important limitations that directly shape the research directions that follow. First, the benchmark is deliberately limited in scope and is intended to show that the framework can be instantiated, enforced, and evaluated in practice, rather than to offer broad empirical coverage of MCP deployments. Second, the taxonomy is built from a mixed evidence base that includes archival research, specifications, incident reports, and practitioner sources. That breadth is appropriate for a fast-moving ecosystem, but it also means that some categories are supported more by emerging operational evidence than by mature longitudinal study. Third, although the paper shows that dual-axis mapping improves localization, control placement, and evidence collection, it does not yet provide a fully operationalized framework for deployment-wide scoring or assurance. Taken together, these limitations point to the broader question of what a more mature MCP security ecosystem should look like. Addressing that question requires moving beyond ad hoc controls and one-off prototypes toward shared mechanisms that can be embedded in specifications, implementations, evaluation practice, and governance. The rest of this section therefore outlines research directions that connect protocol design, benchmarking, human oversight, and ecosystem trust, so that MCP security can develop as a coordinated discipline rather than remain a loose collection of isolated tools.

7.1 Standardized Authentication & Authorization Schemas

There is a clear need for official authentication and authorization frameworks in the MCP specification. At present, **authentication and authorization are explicitly optional for MCP implementations**, as noted in the official draft Model Context Protocol (2025l). In practice this means a client can talk to several MCP servers in the same workspace, each with its own bespoke model of identity, credentials, and permissions. One server might use API keys, another might rely on environment variables, and a third might

skip authentication entirely, so access control behavior becomes hard to reason about or audit at system level. A more coherent future is one where MCP defines an OAuth 2.0 or PKI-style handshake so that clients can verify server identity (for example, servers sign their manifest and clients verify this against a certificate authority or trusted registry), and where there is a built-in notion of user or agent identity and roles so that access control can be expressed per user and per role. Standard schemas would then let an enterprise state, in one place, that “LLM agents of type X can only access tools A, B, C with read-only permissions,” rather than encoding that policy in ad hoc configuration spread across many hosts.

Within this picture, research is needed into lightweight authentication methods that fit dynamic, tool-centric workflows. One plausible direction is a token exchange at session start, followed by JSON Web Tokens attached to each JSON-RPC call, carrying user, agent, and session context Microsoft (2025). This naturally connects to Zero Trust principles, where every call must be authenticated and authorized, not just the initial connection, and where servers treat each tool invocation as a fresh decision rather than assuming that once a channel is open it remains trustworthy.

In addition, **mutual authentication**, where both the client and server verify one another, would help prevent rogue clients or servers from silently joining the ecosystem. An MCP server offering access to sensitive internal systems might require **client-side certificates or API keys** issued only to approved AI applications, and refuse connections from anything else. The OAuth 2.0 ecosystem already supports this pattern with mutual-TLS (mTLS), which binds client credentials to certificates in a way that resists credential theft and replay Lodderstedt et al. (2020), and similar patterns could be adopted or profiled for MCP-aware stacks.

On the authorization side, standardized schemas could allow tools to declare their scopes in a way that closely mirrors OAuth scopes. A tool might declare that it can only read from a particular class of resources, or that it requires a specific organizational role such as `OrgAdmin` to perform write operations. Today, many of these constraints live out-of-band in documentation, environment configuration, or human convention. Encoding them directly in MCP manifests and protocol messages would make them machine-checkable and auditable, at the cost of additional complexity in server and client implementations. This direction is reflected in SMCP, which extends MCP with unified identity management, robust mutual authentication, ongoing security context propagation, fine-grained policy enforcement, and comprehensive audit logging Hou et al. (2026). Together, these mechanisms show how authentication and authorization patterns could be embedded directly into the protocol itself.

It is plausible that researchers and industry groups, such as OWASP’s Agentic Security Initiative OWASP Foundation (2025a), will push these patterns toward a more formal standard. A “MCP 2.0” security model could look much closer to OAuth-style connector security, with client identifiers, secrets, permission scopes, and first-class support for mechanisms such as mTLS or Decentralized Identifiers (DIDs) for trust establishment Sporny et al. (2022). In that world, MCP servers and hosts would speak a shared language of identities, credentials, and scopes, making it far easier to reason about which agents can call which tools, under what conditions, and with what blast radius if something goes wrong.

7.2 Security Benchmarking and Evaluation Frameworks

While tools like **MCPSafetyScanner** Halloran (2025) and early benchmarks such as **MCPBench** Modelscope Team (2025) are a useful first wave, the community still lacks *standardized benchmarking suites for MCP security*. A mature ecosystem would have a shared corpus of attack scenarios, including prompt injection chains, RADE bypass attempts, and over-privileged tool calls, that any new MCP implementation or LLM-based agent can be run against in a repeatable way. Much as **CVE benchmarks** help characterize software vulnerabilities across vendors MITRE Corporation (1999), an “MCP Security Benchmark” could drive convergence on comparable risk scores and structured evaluation reports, making it easier to understand how well a given stack withstands known classes of abuse.

In parallel, there is a need for **evaluation frameworks that focus on agent behavior under adversarial conditions**, not only on server-side vulnerabilities. These frameworks would spell out concrete behaviors to measure, such as whether agents seek user confirmation when instructions are ambiguous or high impact, how they respond when tool outputs conflict, and whether they leak sensitive data when juggling multiple contexts

or tenants. Emerging standards like **NIST AI 600-1** National Institute of Standards and Technology (2024) and community-developed guidance such as **OWASP’s Top 10 for LLMs** OWASP Foundation (2023b) could naturally evolve MCP-specific test modules, so that security evaluations cover both the protocol surface and the decision-making patterns of the agents that drive it.

One concrete direction is an **MCP CTF (Capture The Flag)** ecosystem, where researchers and practitioners interact with deliberately vulnerable MCP servers in a controlled environment. In such a setting, teams would be challenged to coax agents into leaking synthetic secrets, performing unintended actions, or chaining tools in ways that violate policy, while organizers track which exploit patterns are most successful. Gamified platforms like *Hack The Box’s MCP Server* and the *MCP-Kali-Server* project are already experimenting with this style of hands-on training and evaluation Hack The Box (2025); Nahya (2025), and a more formalized CTF series could feed directly into benchmark suites and best practices.

Finally, **community-driven red-teaming and incident sharing** will be essential if MCP deployments are to harden quickly in the face of real-world attacks. Organizations that deploy MCP-capable agents can contribute by sharing anonymized incident reports, red-team findings, and mitigations, turning isolated failures into shared learning. Experiences from efforts such as Microsoft’s internal AI red team and public AI security contests at DEFCON Johnson & Mehrotra (2023); Newman (2023) suggest that this kind of collaborative threat discovery surfaces edge cases that formal analysis alone might miss. Over time, such a feedback loop could converge into a widely adopted **MCP Hardening Guide**, analogous to secure configuration baselines for operating systems or cloud platforms, giving operators a concrete checklist for deploying and maintaining MCP systems with security in mind from day one.

7.3 Human-in-the-Loop UX and Control Design

Today, user consent UIs around AI tools are often rudimentary, even though the underlying actions can be complex and high impact. This creates a research and design opportunity to build more intuitive, graded, and informative consent interfaces for AI-driven tools Nakao et al. (2022); Mandel et al. (2019). Instead of a generic prompt such as “Allow tool to run? (Yes/No)”, future interfaces could present richer, structured context, for example “This tool will delete five files. Estimated impact: High. Admin approval required.” or require multi-factor authentication before performing sensitive actions. Giving users simple ways to express standing preferences (for example, “Never allow AI to send email externally” or “Ask me if more than 10 records will be affected”) can bridge the gap between today’s coarse consent and fully automatic operation. There is substantial Human–Computer Interaction (HCI) work to be done on how to surface an AI’s intentions so that users can make informed decisions, including natural language summaries of proposed actions such as “I found three files that contain your SSN and propose to upload them to an encrypted vault. Do I have your permission?” Cappuccio et al. (2025); Chromik & Butz (2021).

A second pillar is explicit oversight mechanisms such as kill switches and supervisory controls Orseau & Armstrong (2016); Hadfield-Menell et al. (2017). Users and operators may need a visible “panic button” that immediately halts all tool usage when something appears to be going wrong, as well as an oversight mode where a human or a separate AI guardian monitors the agent’s calls in real time and can intervene or veto. Recent work has proposed **Agent Governors**, dedicated processes that continuously watch an AI system’s actions and decisions against a policy budget Xiang et al. (2024). In an MCP setting, this might correspond to an intermediary component that inspects each tool invocation and requires dual approval for certain categories of tools, similar to how some safety-critical operations in other domains require two independent operators to confirm an action before it proceeds.

7.4 Governance, Certification, and Ecosystem Trust

For MCP to be viable at scale (especially in enterprises and regulated industries), there needs to be an ecosystem of trust and certification for third-party tools. This could take inspiration from mobile app stores or browser extensions, e.g., an MCP Server Marketplace where submissions undergo security review, automated scanning, and signing Apple Inc. (2021); Iqbal et al. (2023). Furthermore, we might see the emergence of MCP security certifications, similar to how software can get a SOC2 AICPA (2023) or FedRAMP Program Management Office (2024) certification. For instance, an MCP connector handling healthcare data might

need to certify HIPAA compliance. This implies not just technical measures, but documentation, audit logs, and perhaps even third-party pen-tests.

Another governance angle is community-driven threat intel sharing for MCP. As MCP threats evolve, a central database of known malicious servers or attack signatures (like an “MCP threat feed”) could help automatically guard systems (similar to anti-virus definitions or threat intelligence feeds used in network security) MITRE Corporation (2023); Google (2024). For example, if a certain package name on PyPI is identified as malware, that info should propagate quickly to all MCP security tools Phylum Research (2022).

From a standards perspective, groups like the **Frontier Model Forum**, which includes OpenAI, Anthropic, Google, and Microsoft, are already publishing shared safety recommendations and could play a similar role in defining **best practices for MCP usage**. Building on that momentum, **the Model Context Protocol Authorization draft**, maintained by Anthropic, sketches concrete flows for securing MCP interactions with API tokens and OAuth-style credentials Model Context Protocol (2025l). The next step is to turn these proposals into a broadly adopted baseline, especially for sensitive surfaces such as registries and authentication, which will require coordinated governance and consensus building by AI safety coalitions and open-spec bodies rather than isolated vendor efforts. In parallel, **regulators** are already stepping in. Under the EU AI Act, high-risk AI systems must log key lifecycle events, including automated decisions, for audit and compliance purposes EUA (2024). If future breaches or misuse are traced to MCP-powered AI actions, for example sending wire transfers or altering health records, similar **logging and traceability mandates** are likely to be extended to finance, healthcare, and government sectors, creating a strong external push for MCP deployments to adopt robust, standardized security and audit practices.

7.5 Continual Model Robustness and Alignment

Finally, a more research-oriented direction is to improve the LLMs themselves to be more security-conscious when using tools. Currently, LLMs have a limited understanding of security context; they will readily follow instructions Li et al. (2024); Zhao et al. (2025); Ge et al. (2025). Future research might train models with “security awareness,” so the AI can flag its own concerns: “This request would result in sending sensitive data externally, which seems unsafe. I should ask for confirmation.” Early work in **Constitutional AI** instructs LLMs to critique their own actions according to a constitution Bai et al. (2022), and **RLHF** has been shown to be effective in training models to obey complex human preferences Christiano et al. (2017). Furthermore, Recent efforts to extend Constitutional AI into tool-using agents, such as the Agentic Superego framework Watson et al. (2025), illustrate how internal ethics checks can be embedded into MCP workflows, allowing AI systems to proactively surface security-relevant explanations and seek user confirmation before executing sensitive actions.

One emerging concept is **Prompt Flow Integrity (PFI)**, a security framework aimed at ensuring that instructions received by an AI agent have not been tampered with and do not result in unintended privilege escalation Kim et al. (2025). PFI involves isolating prompt-processing stages and enforcing guardrails around tool-use permissions. To support this, **differential testing techniques**, where multiple agent instances cross-check each other’s decisions-have been proposed to detect inconsistencies or malicious deviations Rao et al. (2025). Additionally, recent work shows that **instrumenting the model’s reasoning process** (e.g., via trace analysis or sabotage detection) can proactively catch when an LLM is veering into unsafe behavior Kutasov et al. (2025). Related work is beginning to push this idea closer to MCP tool use itself, with MCPShield proposing an adaptive trust-calibration layer that evaluates risk before and after tool execution so that unsafe or suspicious actions can be constrained and reevaluated during runtime Zhou et al. (2026). Together, these strategies can enhance the robustness of tool-using agents by enforcing internal consistency, highlighting anomalies, and providing early intervention signals within an MCP workflow.

Another future direction is **context-aware access control**, where the AI’s own understanding of context labels (confidential, public, etc.) influences its tool choices (for instance, refuse to use a public posting tool if the content is labeled confidential). This requires the AI to accurately classify data sensitivity, which is nontrivial, but not impossible with metadata and training.

Lastly, integrating AI explainability into MCP workflows is vital because when a model can articulate why it is invoking a tool with specific data, users or oversight systems can better judge whether that action is

appropriate. Bilal et al. (2025) show that LLMs can produce rich, natural-language explanations of their own reasoning, and work on layered prompting in multi-agent settings demonstrates how to surface action-level justifications, which in turn creates a foundation for real-time human-in-the-loop validation before tool execution Faisal & Tunkel (2025).

8 Conclusion

Taken together, our findings show that MCP security is governed less by isolated vulnerabilities and more by structured cascades that traverse components and lifecycle phases. By constructing a failure mode oriented taxonomy and mapping each threat to specific components and phases, we turn scattered incidents into a single structured space. Once threats are viewed in this joint space of failure mode, component, and phase, patterns that previously appeared unrelated become instances of the same underlying weakness. Concrete MCP scenarios make this structure visible. A creation time misconfiguration that widens a tool scope, an interaction time prompt manipulation that steers an agent toward that tool, and a maintenance time drift that leaves stale credentials in place may appear as distinct issues, yet they are different manifestations of shared failures in identity, isolation, or integrity rather than separate classes of bugs. This perspective keeps the taxonomy compact while still reflecting the range of behaviors observed in real MCP deployments.

Building on this view, the taxonomy provides a scaffold for a mitigation oriented control plane that spans design, deployment, and operation. For each failure mode, it identifies the components and phases where relatively small changes in configuration, authentication, or isolation can have the greatest security effect. At those control points, the taxonomy is paired with concrete mechanisms and verification artifacts, so that abstract failure modes become specific checks and enforcement actions rather than only descriptive categories. Section 5 shows that this control oriented view can also be realized in a compact benchmark, where representative MCP failures are detected through verifiable checks, localized to the relevant component and lifecycle phase, and tied to explicit audit artifacts. Read together, these elements make scoped capabilities, hardened gateway patterns, sandboxed execution environments, and structured telemetry appear not as isolated defenses, but as recurring control patterns that can be adapted across deployments while preserving their intended security effect. Because many threats span multiple component and phase cells, no single control is sufficient on its own. Effective enforcement must therefore be layered across configuration validation, pre deployment checks, and runtime monitoring, which helps constrain the blast radius of individual failures, slows the propagation of cross phase cascades, and leaves operators with structured evidence for diagnosis and remediation.

This map is intended to give both researchers and practitioners a shared structure for reasoning about MCP threats and controls. For researchers, it clarifies which regions of the MCP threat space remain weakly understood and which combinations of failure mode, component, and phase are not yet covered by existing frameworks and benchmarks. For system architects and operators, it supports systematic reasoning about where to place controls, which forms of evidence to collect, and how to prioritize remediation work when incidents occur. Taken together, these roles mean that the taxonomy and control plane currently operate as a static reference that guides design decisions, structured reviews, and incident analysis. Building on this, a natural next step is to embed the same structure into the infrastructure of MCP deployments so that the same concepts inform automated detection and enforcement. One line of work is the development of telemetry driven analytics that align anomalies in MCP traffic with specific cells in the taxonomy and provide early warning of emerging attacks. A second line of work is the design of integrity enforcement mechanisms that compare expected agent and tool behavior with observed executions and emit explicit violations when workflows deviate from policy. A third line of work is the construction of scalable policy engines that express MCP specific controls in a uniform language and drive configuration, deployment, and runtime enforcement from the same rule base, particularly as recent work explores shared-context collaboration, large-scale tool routing, and governance-aware enterprise data access Jayanti & Han (2026); Yao et al. (2026); Tonnarelli et al. (2026). Pursued together, these directions can move MCP from a loosely structured collection of components toward a security aware environment where threats are mapped, monitored, and constrained using the structure laid out in this survey.

References

- Regulation (eu) 2024/1689 — ai act, 2024. URL https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ%3AL_202401689. Published in the Official Journal of the European Union.
- Cve-2024-33899 – national vulnerability database. <https://nvd.nist.gov/vuln/detail/CVE-2024-33899>, 2024a.
- Cve-2024-36052 – national vulnerability database. <https://nvd.nist.gov/vuln/detail/CVE-2024-36052>, 2024b.
- mcp-get: Discover and manage model context protocol servers. <https://github.com/michaellatman/mcp-get>, 2025a.
- mcp-installer: Automated mcp server deployment tool. <https://github.com/anaisbetts/mcp-installer>, 2025b.
- Arash Ahmadi, Sarah S. Sharif, and Yaser M. Banad. Mcp bridge: A lightweight, llm-agnostic restful proxy for model context protocol servers. *arXiv preprint arXiv:2504.08999*, 2025.
- AICPA. Soc 2® - soc for service organizations: Trust services criteria (2017 tsc with revised points of focus), 2023. URL <https://www.aicpa.org/topic/audit-assurance/audit-and-assurance-greater-than-soc-2>. Trust Services Criteria with revised Points of Focus, published by AICPA on 30 Sep 2023.
- Fatih Kadir Akın. Introducing MCP tools: A command-line inspector for model context protocol servers. <https://blog.fka.dev/blog/2025-03-26-introducing-mcp-tools-cli/>, 2025. Accessed 25 July 2025.
- Yoshihiro Ando. Zero-trust architecture for mcp-based ai agents: A unified cli approach. *TechRxiv*, feb 2026. doi: 10.36227/techrxiv.177155647.75363542/v1.
- Anthropic. Introducing the model context protocol. <https://www.anthropic.com/news/model-context-protocol>, 2024a. Accessed: 2025-11-21.
- Anthropic. Many-shot jailbreaking. Technical report, 2024b. URL <https://www.anthropic.com/research/many-shot-jailbreaking>. Demonstrates long-context “many-shot” jailbreaks via repeated prompting.
- Anthropic. Three sketches of asl-4 safety case components, 2024c. URL <https://alignment.anthropic.com/2024/safety-cases/>. Discusses data-poisoning based backdoors and challenges in ruling them out.
- Anthropic. Model context protocol connector specification. <https://docs.anthropic.com/en/docs/agents-and-tools/mcp-connector>, 2025.
- Anthropic. Connect claude code to tools via mcp, 2025. URL <https://code.claude.com/docs/en/mcp>. Anthropic documentation warning that third-party MCP servers may fetch untrusted content and expose users to prompt injection risk.
- Anthropic. Mitigate jailbreaks and prompt injections, 2025a. URL <https://docs.anthropic.com/en/docs/test-and-evaluate/strengthen-guardrails/mitigate-jailbreaks>. Official Anthropic guidance on adversarial prompts/prompt injection and guardrails.
- Anthropic. How we built our multi-agent research system, 06 2025b. URL <https://www.anthropic.com/engineering/built-multi-agent-research-system>. Reports coordination challenges, emergent behaviors, deployment/hand-off considerations in production multi-agent systems.
- Anthropic. Tools, 2025c. URL <https://modelcontextprotocol.io/docs/concepts/tools>. Model Context Protocol documentation (accessed July 24, 2025).
- Apple Inc. Building a trusted ecosystem for millions of apps, 2021. URL https://www.apple.com/privacy/docs/Building_a_Trusted_Ecosystem_for_Millions_of_Apps.pdf. White paper released June 23, 2021.

- Astrix Research Team. State of MCP server security 2025. <https://astrix.security/learn/blog/state-of-mcp-server-security-2025/>, 2025. Accessed: 2025-11-25.
- AWS. Aws verified permissions, 2025. URL <https://aws.amazon.com/verified-permissions/>.
- AWS Compute Blog. Anatomy of cve-2019-5736: A runc container escape. <https://aws.amazon.com/blogs/compute/anatomy-of-cve-2019-5736-a-runc-container-escape/>, 2019. Online; accessed 24 July 2025.
- Backslash Security. Top 10 risks of using mcp servers in ides, 06 2025. URL <https://www.backslash.security/blog/top-risks-mcp-servers-ide>. Industry analysis: misconfigurations and untrusted servers can expand capabilities and exfiltrate data.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. Constitutional ai: Harmlessness from ai feedback, 2022. arXiv preprint. <https://arxiv.org/abs/2212.08073>.
- Chaithanya Bandi, Ben Hertzberg, Geobio Boo, Tejas Polakam, Jeff Da, Sami Hassaan, Manasi Sharma, Andrew Park, Ernesto Hernandez, Dan Rambado, Ivan Salazar, Rafael Cruz, Chetan Rane, Ben Levin, Brad Kenstler, and Bing Liu. Mcp-atlas: A large-scale benchmark for tool-use competency with real mcp servers. *arXiv preprint arXiv:2602.00933*, 2026.
- Frédéric Barthelet. Better MCP tools/call error responses: Help your AI recover gracefully. <https://dev.to/alpic/better-mcp-tools-call-error-responses-help-your-ai-recover-gracefully-15c7>, July 2025. DEV Community blog, accessed 23 November 2025.
- Akhil Behl. Navigating mcp security: Key considerations and mitigation strategies for the enterprise, May 2025. URL <https://fractal.ai/blog/navigating-mcp-security-key-considerations-and-mitigation-strategies-for-enterprises>.
- Carolina Beretta, Gus Carlock, and Andrew Pease. Mcp tools: Attack vectors and defense recommendations for autonomous agents. <https://www.elastic.co/security-labs/mcp-tools-attack-defense-recommendations>, September 2025. Elastic Security Labs blog.
- Luca Beurer-Kellner and Marc Fischer. Whatsapp mcp exploited: Exfiltrating your message history via mcp. <https://invariantlabs.ai/blog/whatsapp-mcp-exploited>, 2025.
- Manish Bhatt, Vineeth Sai Narajala, and Idan Habler. Etdi: Mitigating tool squatting and rug pull attacks in model context protocol (mcp) by using oauth-enhanced tool definitions and policy-based access control, June 2025. URL <https://arxiv.org/abs/2506.01333>. arXiv:2506.01333v1.
- Ahsan Bilal, David Ebert, and Beiyu Lin. Llms for explainable ai: A comprehensive survey. 16(3):32, 2025. Article 32. Also available as arXiv:2504.00125. <https://arxiv.org/abs/2504.00125>.
- Nitin Birur. We scanned 1,000 MCP servers. ~33% had critical vulnerabilities. <https://www.enkrypta.com/blog/we-scanned-1-000-mcp-servers-33-had-critical-vulnerabilities>, October 2025. Enkrypt AI blog, Product Updates.
- Ivo Brett. Simplified and secure mcp gateways for enterprise ai integration, April 2025. v1 (submitted 28 Apr 2025). <https://arxiv.org/abs/2504.19997>.

- Eleonora Cappuccio, Andrea Esposito, Francesco Greco, Giuseppe Desolda, Rosa Lanzilotti, and Salvatore Rinzivillo. Explanation user interfaces: A systematic literature review, 2025. URL <https://arxiv.org/abs/2505.20085>. arXiv preprint; published May 26, 2025.
- Ramaswamy Chandramouli. Guide to a secure enterprise network landscape, 2022. URL <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-215.pdf>. Final release.
- Ramaswamy Chandramouli and Zack Butcher. Building secure microservices-based applications using service-mesh architecture, 2022. Final publication. <https://csrc.nist.gov/pubs/sp/800/204/a/final>.
- Edward Y. Chang and Longling Geng. Sagallm: Context management, validation, and transaction guarantees for multi-agent LLM planning. *Proceedings of the VLDB Endowment*, 18(12):4874–4886, 2025. doi: 10.14778/3750601.3750611. URL <https://www.vldb.org/pvldb/vol18/p4874-chang.pdf>.
- Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*, pp. 4299–4307, 2017. URL <https://papers.nips.cc/paper/7017-deep-reinforcement-learning-from-human-preferences>. NeurIPS 2017, accepted as part of NIPS 30.
- Michael Chromik and Andreas Butz. Human–xai interaction: A review and design principles for explanation user interfaces. In *Human–Computer Interaction – INTERACT 2021*, volume 12934 of *Lecture Notes in Computer Science*, pp. 619–640. Springer, 2021. doi: 10.1007/978-3-030-85616-8_36. URL https://doi.org/10.1007/978-3-030-85616-8_36.
- Cloudflare. What is the model context protocol (mcp)?, 2025. URL <https://www.cloudflare.com/learning/ai/what-is-model-context-protocol-mcp/>. Cloudflare Learning Center. Accessed: March 15, 2026.
- Cloudflare Learning Center. Understanding mutual tls (mtls). <https://developers.cloudflare.com/ssl/mutual-tls>, 2024.
- Omri Cohen and Sarit Yerushalmi. Python’s colorama typosquatting meets ‘fade stealer’ malware. *Imperva Blog*, 2024. <https://www.imperva.com/blog/pythons-colorama-typosquatting-meets-fade-stealer-malware/>.
- Elena Cross. The “s” in mcp stands for security. <https://dev.to/elenacross/the-s-in-mcp-stands-for-security-1gge>, April 2025. Blog post, DEV Community.
- CSF Tools. Nist sp 800-53 rev. 4 control si-10: Information input validation, 2025. URL <https://csf.tools/reference/nist-sp-800-53/r4/si/si-10/>. Accessed 19 Jul 2025.
- Curity. Design MCP authorization to securely expose APIs. <https://curity.io/resources/learn/design-mcp-authorization-apis/>, 2025. Accessed: 2025-11-25.
- Joseph W. Cutler, Craig Disselkoen, Aaron Eline, Shaobo He, Kyle Headley, Michael Hicks, Kesha Hietala, Eleftherios Ioannidis, John Kastner, Anwar Mamat, Darin McAdams, Matt McCutchen, Neha Rungta, Emina Torlak, and Andrew M. Wells. Cedar: A new language for expressive, fast, safe, and analyzable authorization (extended version). <https://arxiv.org/abs/2403.04651>, 2024. Accessed Dec. 1, 2025.
- Cyber Sierra Knowledge Team. MCP security risks: A CISO’s guide to mitigation. <https://cybersierra.co/blog/mcp-security-risks-guide/>, August 2025. Accessed Nov. 30, 2025.
- CyberArk Threat Research. Is your ai safe? threat analysis of mcp, 2025. <https://www.cyberark.com/resources/threat-research-blog/is-your-ai-safe-threat-analysis-of-mcp-model-context-protocol>.
- Data Science Dojo Staff. The state of mcp security in 2025: Key risks, attack vectors, and case studies, September 2025. URL <https://datasciencedojo.com/blog/mcp-security-risks-and-challenges/>. Blog article.

- R. Di Pietro and F. Lombardi. Falco: Kernel-space threat detection with ebpf. *Journal of Cybersecurity*, 2021.
- Docker. Seccomp security profiles for docker, 2025. Default Docker seccomp profile and blocked syscalls. <https://docs.docker.com/engine/security/seccomp/>.
- Vardhan Dongre, Ryan A. Rossi, Viet Dac Lai, David Seunghyun Yoon, Dilek Hakkani-Tür, and Trung Bui. Drift no more? context equilibria in multi-turn LLM interactions. <https://arxiv.org/abs/2510.07777>, 2025.
- dts-solution. Securing the model context protocol (mcp) in enterprise ai deployments. <https://www.dts-solution.com/securing-the-model-context-protocol-mcp-in-enterprise-ai-deployments/>, 2025.
- Siddharth Dushantha. Ansi escape injection vulnerability in winrar, 2024. URL <https://sdushantha.medium.com/ansi-escape-injection-vulnerability-in-winrar-a2cbfac4b983>.
- Abul Ehtesham, Aditi Singh, Gaurav Kumar Gupta, and Saket Kumar. A survey of agent interoperability protocols: Model context protocol (mcp), agent communication protocol (acp), agent-to-agent protocol (a2a), and agent network protocol (anp). 2025. URL <https://arxiv.org/abs/2505.02279>. Preprint.
- Elastic Security Labs. Mcp tools: Attack vectors and defense recommendations for autonomous agents. <https://www.elastic.co/security-labs/mcp-tools-attack-defense-recommendations>, 2025.
- Emergent Mind. Model context protocols (mcps) - security variability. <https://www.emergentmind.com/topics/model-context-protocols-mcps>, 2025. Online; accessed 24 July 2025.
- Enkrypt AI. MCP security vulnerabilities: Attacks, detection, and prevention. <https://www.enkryptai.com/blog/mcp-security-vulnerabilities-attacks-detection-and-prevention>, October 2025. Accessed Nov. 30, 2025.
- Eriksen, Charlie. Ratatouille: A malicious recipe hidden in rand-user-agent. *Aikido Blog*, 2025. <https://www.aikido.dev/blog/catching-a-rat-remote-access-trojan-rand-user-agent-supply-chain-compromise>.
- Bashir Faisal and David Tunkel. Explainable ai in multi-agent systems: Advancing transparency with layered prompting. In *Explainable AI — INTERACT 2025*, pp. 619–640. Springer, 2025. Proceedings paper (details via ResearchGate). https://www.researchgate.net/publication/388835453_Explainable_AI_in_Multi-Agent_Systems_Advancing_Transparency_with_Layered_Prompting.
- Junfeng Fang, Zijun Yao, Ruipeng Wang, Haokai Ma, Xiang Wang, and Tat-Seng Chua. We Should Identify and Mitigate Third-Party Safety Risks in MCP-Powered Agent Systems. *arXiv preprint arXiv:2506.13666*, June 2025. doi: 10.48550/arXiv.2506.13666. URL <https://arxiv.org/abs/2506.13666>.
- FedRAMP Program Management Office. Fedramp agency authorization playbook, version 4.0, 2024. URL https://www.fedramp.gov/assets/resources/documents/Agency_Authorization_Playbook.pdf. Playbook (Version 4.0 published to align with OMB Memo M-24-15).
- Xiang Fei, Xiawu Zheng, and Hao Feng. Mcp-zero: Proactive toolchain construction for llm agents from scratch. *arXiv preprint arXiv:2506.01056*, 2025. URL <https://arxiv.org/abs/2506.01056>.
- Robert Fenstermacher. How to secure mcp: Threats and defenses. <https://stytych.com/blog/mcp-security>, 2025. Accessed 24 July 2025.
- Mohamed Amine Ferrag, Norbert Tihanyi, Djallel Hamouda, Leandros Maglaras, and Merouane Debbah. From prompt injections to protocol exploits: Threats in LLM-powered AI agents workflows. *arXiv preprint arXiv:2506.23260*, June 2025. doi: 10.48550/arXiv.2506.23260. URL <https://arxiv.org/abs/2506.23260>.

- David F. Ferraiolo, Ravi Sandhu, Serban Gavrilă, D. Richard Kuhn, and Ramaswamy Chandramouli. Proposed NIST standard for role-based access control. *ACM Transactions on Information and System Security*, 4(3):224–274, 2001. doi: 10.1145/501978.501980.
- Gaëtan Ferry. A look into the secrets of MCP: The new secret leak source. <https://blog.gitguardian.com/a-look-into-the-secrets-of-mcp/>, 2025. GitGuardian Blog, accessed Nov. 25, 2025.
- Field, Michael. What is mcp? diving deep into the future of remote ai context. <https://konghq.com/blog/learning-center/what-is-mcp>, 2025. Online; Kong Inc. Learning Center; accessed 24 July 2025.
- Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM CCS*, 2015. URL <https://rist.tech.cornell.edu/papers/mi-ccs.pdf>.
- Muhammad Liman Gambo and Ahmad Almulhem. Zero trust architecture: A systematic literature review. *arXiv preprint arXiv:2503.11659*, March 2025. doi: 10.48550/arXiv.2503.11659. URL <https://arxiv.org/abs/2503.11659>.
- GCP. Using falco with gke autopilot, 2023. URL <https://cloud.google.com/architecture/falco>.
- Yubin Ge, Neeraja Kirtane, Hao Peng, and Dilek Hakkani-Tür. Llms are vulnerable to malicious prompts disguised as scientific language, 2025. URL <https://arxiv.org/abs/2501.14073>. arXiv preprint (16 pages, includes figures and tables).
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92, 2021. doi: 10.1145/3458723. URL <https://dl.acm.org/doi/10.1145/3458723>.
- General Analysis. Supabase MCP can leak your entire SQL database. <https://www.generalanalysis.com/blog/supabase-mcp-blog>, June 2025. Accessed: 2025-11-22.
- GitHub, Inc. Configure mcp server access for your organization or enterprise, 09 2025. URL <https://docs.github.com/en/copilot/how-tos/administer-copilot/configure-mcp-server-access>. Each registry entry points to a server manifest describing tools, resources, and prompts. Accessed 2025-10-12.
- Google. Osv - open source vulnerabilities: A distributed vulnerability database for open source, 2024. URL <https://osv.dev/>. Website launched by 2024, actively reviewed in early 2025.
- Abhishek Goswami. Agentic jwt: A secure delegation protocol for autonomous ai agents. *arXiv preprint arXiv:2509.13597*, 2025. URL <https://arxiv.org/abs/2509.13597>.
- Yongjian Guo, Puzhuo Liu, Wanlun Ma, Zehang Deng, Xiaogang Zhu, Peng Di, Xi Xiao, and Sheng Wen. Systematic analysis of MCP security, 2025. URL <https://arxiv.org/abs/2508.12538>. arXiv preprint.
- Hack The Box. Introducing the mcp server for ctf competitions at hack the box, June 2025. Blog post (“Multiple Authors”), published 16 Jun 2025. Accessed 19 Jul 2025. <https://www.hackthebox.com/blog/model-context-protocol>.
- Dylan Hadfield-Menell, Anca Dragan, Pieter Abbeel, and Stuart Russell. The off-switch game. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, pp. 220–227, 2017. doi: 10.24963/ijcai.2017/32. URL <https://doi.org/10.24963/ijcai.2017/32>.
- B. Halloran et al. Mcp safety training: Learning to refuse falsely benign mcp exploits using improved preference alignment. *arXiv preprint arXiv:2505.23634*, 2025. <https://arxiv.org/abs/2505.23634>.
- John Halloran. Mcpsafetyscanner: A safety auditor for model context protocol (mcp) servers, 2025. URL <https://github.com/johnhalloran321/mcpSafetyScanner>. Open-source on GitHub with activity as of April 29, 2025 (initial commit date not publicly disclosed).

- Jared Hanson and Model Context Protocol maintainers. Authentication - discussion #64 (model context protocol). <https://github.com/modelcontextprotocol/modelcontextprotocol/discussions/64>, 2024. GitHub discussion post, opened 25Nov2024. Accessed 24Jul2025.
- D. Hardt, T. Lodderstedt, and D. Fett. Oauth 2.1 authorization framework (ietf draft). <https://datatracker.ietf.org/doc/html/draft-ietf-oauth-v2-1>, 2024.
- Mohammed Mehedi Hasan, Hao Li, Emad Fallahzadeh, Gopi Krishnan Rajbahadur, Bram Adams, and Ahmed E. Hassan. Model context protocol (mcp) at first glance: Studying the security and maintainability of MCP servers, 2025. URL <https://arxiv.org/abs/2506.13538>. arXiv preprint, v4.
- Mohammed Mehedi Hasan, Hao Li, Gopi Krishnan Rajbahadur, Bram Adams, and Ahmed E. Hassan. Model context protocol (mcp) tool descriptions are smelly! towards improving ai agent efficiency with augmented mcp tool descriptions. *arXiv preprint arXiv:2602.14878*, 2026.
- Hitachi Vantara. Transport layer security (tls) overview. <https://knowledge.hitachivantara.com>, 2023.
- HiveTrail. Securing mcp server authentication: From oauth 2.1 challenges to production-ready solutions. <https://hivetrail.com/blog/securing-mcp-server-authentication/>, 2025. Online; accessed 24 July 2025.
- Keith Hoodlet. Deceiving users with ansi terminal codes in mcp, April 2025. URL <https://blog.trailofbits.com/2025/04/29/deceiving-users-with-ansi-terminal-codes-in-mcp/>.
- Kashish Hora. Error handling in custom MCP servers. <https://mcpicat.io/guides/error-handling-custom-mcp-servers/>, 2025. MCPcat guide, accessed 23 November 2025.
- Xinyi Hou, Yanjie Zhao, Shenao Wang, and Haoyu Wang. Model Context Protocol (MCP): Landscape, Security Threats, and Future Research Directions. *arXiv preprint arXiv:2503.23278*, March 2025. doi: 10.48550/arXiv.2503.23278. URL <https://arxiv.org/abs/2503.23278>.
- Xinyi Hou, Shenao Wang, Yifan Zhang, Ziluo Xue, Yanjie Zhao, Cai Fu, and Haoyu Wang. Smcp: Secure model context protocol. *arXiv preprint arXiv:2602.01129*, 2026.
- Humanloop. Model context protocol (mcp) explained. <https://humanloop.com/blog/mcp>, April 2025. Accessed Dec. 1, 2025.
- Invariant Labs. Mcp security notification: Tool poisoning attacks, 2025. <https://invariantlabs.ai/blog/mcp-security-notification-tool-poisoning-attacks>.
- Umar Iqbal, Tadayoshi Kohno, and Franziska Roesner. Llm platform security: Applying a systematic evaluation framework to openai’s chatgpt plugins, 2023. URL <https://arxiv.org/abs/2309.10254>. arXiv preprint; v2 issued 2024-07-26.
- istio.io. Istio security (authentication, authorization, and mutual tls) concepts, 2025. URL <https://istio.io/latest/docs/concepts/security/>. Living documentation covering identity, mutual TLS, and authorization features for service-to-service security.
- Ithena. Mcp audit trails: Meeting enterprise compliance when ai agents access your data. <https://www.ithena.one/blog/mcp-audit-trails>, 2025. Accessed 2025-10-12.
- Meenakshi Amulya Jayanti and X.Y. Han. Enhancing model context protocol (mcp) with context-aware server collaboration. *arXiv preprint arXiv:2601.11595*, 2026. doi: 10.48550/arXiv.2601.11595.
- Huihao Jing, Haoran Li, Wenbin Hu, Qi Hu, Heli Xu, Tianshu Chu, Peizhao Hu, and Yangqiu Song. Mcip: Protecting mcp safety via model contextual integrity protocol, June 2025. v4 (last revised 4 Jun 2025; first submitted 20 May 2025). <https://arxiv.org/abs/2505.14590>.
- Khari Johnson and Dhruv Mehrotra. This showdown between humans and chatbots could keep you safe from bad ai, 2023. URL <https://www.wired.com/story/showdown-humans-chatbots-defcon/>. Published August 14, 2023 at 19:51 (site local time).

- Ashur Kanoon. Securing MCP server communications: Beyond authentication. <https://aembit.io/blog/securing-mcp-server-communications-best-practices/>, 2025. Aembit Blog, accessed 2025-11-22.
- K. Kent and M. Souppaya. Guide to computer security log management. Technical report, NIST SP 800-92, 2006.
- Juhee Kim, Woohyuk Choi, and Byoungyoung Lee. Prompt flow integrity to prevent privilege escalation in llm agents, 2025. arXiv preprint (version 2, revised 21 Apr 2025; original submission 17 Mar 2025). DOI: 10.48550/arXiv.2503.15547. <https://arxiv.org/abs/2503.15547>.
- Richard Kissel, Andrew Regenscheid, Matthew Scholl, and Kevin Stine. Nist special publication 800-88 revision 1: Guidelines for media sanitization, 2014. <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-88r1.pdf>.
- Knostic Team. How model context protocol (mcp) servers communicate, 2025. URL <https://www.knostic.ai/blog/model-context-protocol-communication>. Accessed: 2025-11-23.
- Kubernetes. Pod security standards, 2022. URL <https://kubernetes.io/docs/concepts/security/pod-security-standards/>. Section: Running as Non-root (accessed with deep link fragment).
- Kubernetes. Network policies, 2024a. URL <https://kubernetes.io/docs/concepts/services-networking/network-policies/>. Official Kubernetes Documentation.
- Kubernetes. Role based access control good practices, 2024b. URL <https://kubernetes.io/docs/concepts/security/rbac-good-practices/>. Official Kubernetes Documentation.
- Kubernetes. Volumes. <https://kubernetes.io/docs/concepts/storage/volumes/>, 2025a. Warns that hostPath presents many security risks and should be avoided. Accessed 2025-11-25.
- Kubernetes. Configure a security context for a pod or container. <https://kubernetes.io/docs/tasks/configure-pod-container/security-context/>, 2025b. Covers non-root, capability limits, read-only root. Accessed 2025-11-25.
- Sonu Kumar, Anubhav Girdhar, Ritesh Patil, and Divyansh Tripathi. Mcp guardian: A security-first layer for safeguarding mcp-based ai system, May 2025a. v2 (last revised 19 May 2025; first submitted 17 Apr 2025). <https://arxiv.org/abs/2504.12757>.
- Sonu Kumar, Anubhav Girdhar, Ritesh Patil, and Divyansh Tripathi. MCP guardian: A security-first layer for safeguarding MCP-based ai system. *arXiv preprint arXiv:2504.12757*, April 2025b. doi: 10.48550/arXiv.2504.12757. URL <https://arxiv.org/abs/2504.12757>.
- Jonathan Kutasov, Yuqi Sun, Paul Colognese, Teun van der Weij, Linda Petrini, Chen Bo Calvin Zhang, John Hughes, Xiang Deng, Henry Sleight, Tyler Tracy, Buck Shlegeris, and Joe Benton. Shade-arena: Evaluating sabotage and monitoring in llm agents, 2025. arXiv preprint (v2, revised 8 Jul 2025; original submission: 17 Jun 2025). DOI: 10.48550/arXiv.2506.15740. <https://arxiv.org/abs/2506.15740>.
- Melvin Lammerts. The ai protocol under siege: Mcp server vulnerabilities expose critical threats, July 2025. URL <https://hadrian.io/blog/the-ai-protocol-under-siege-mcp-server-vulnerabilities-expose-critical-threats>. Hadrian blog post.
- B. Laurie, D. Mehta, A. Raghunathan, B. Ray, A. Roy, and D. Sculley. An introduction to google’s approach for secure ai agents, 2024. URL <https://research.google/pubs/an-introduction-to-googles-approach-for-secure-ai-agents/>. Google Research.
- Ruiqi Li, Zhiqiang Wang, Yunhao Yao, and Xiang-Yang Li. Mcp-ityp: An automated framework for implicit tool poisoning in mcp. *arXiv preprint arXiv:2601.07395*, 2026a.
- Xiaofan Li and Xing Gao. Toward understanding security issues in the model context protocol ecosystem. *arXiv preprint arXiv:2510.16558*, 2025. doi: 10.48550/arXiv.2510.16558. URL <https://arxiv.org/abs/2510.16558>.

- Zekun Li, Baolin Peng, Pengcheng He, and Xifeng Yan. Evaluating the instruction-following robustness of large language models to prompt injection. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 557–568, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.33. URL <https://aclanthology.org/2024.emnlp-main.33/>. EMNLP 2024.
- Zhihao Li, Kun Li, Boyang Ma, Minghui Xu, Yue Zhang, and Xiuzhen Cheng. We Urgently Need Privilege Management in MCP: A measurement of API usage in MCP ecosystems. *arXiv preprint arXiv:2507.06250*, July 2025. doi: 10.48550/arXiv.2507.06250. URL <https://arxiv.org/abs/2507.06250>.
- Zhihao Li, Boyang Ma, Xuelong Dai, Minghui Xu, Yue Zhang, Biwei Yan, and Kun Li. Don’t believe everything you read: Understanding and measuring mcp behavior under misleading tool descriptions. *arXiv preprint arXiv:2602.03580*, 2026b.
- Linux Kernel Project. no_new_privs: Linux kernel userspace api documentation, 2012. URL https://docs.kernel.org/userspace-api/no_new_privs.html.
- Torsten Lodderstedt, Sébastien Dronia, and Marius Scurtescu. OAuth 2.0 Token Revocation. <https://www.rfc-editor.org/rfc/rfc7009>, August 2013. Proposed Standard.
- Torsten Lodderstedt, Brian Campbell, and Daniel Fett. OAuth 2.0 mutual-tls client authentication and certificate-bound access tokens. Request for Comments RFC 8705, Internet Engineering Task Force (IETF), January 2020. Standards Track. <https://datatracker.ietf.org/doc/html/rfc8705>.
- Torsten Lodderstedt, John Bradley, Andrey Labunets, and Daniel Fett. Best current practice for oauth 2.0 security. RFC 9700 (BCP 240), January 2025. URL <https://www.rfc-editor.org/rfc/rfc9700>. Updates RFC 6749, RFC 6750, RFC 6819.
- Avi Lumelsky. Critical rce vulnerability in anthropic mcp inspector - cve-2025-49596, June 2025. URL <https://www.oligo.security/blog/critical-rce-vulnerability-in-anthropic-mcp-inspector-cve-2025-49596>. Blog post; CVSS 9.4; MCP Inspector prior to v0.14.1.
- Elias Lumer, Anmol Gulati, Vamse Kumar Subbiah, Pradeep Honaganahalli Basavaraju, and James A. Burke. Scalemcp: Dynamic and auto-synchronizing model context protocol tools for llm agents. *arXiv preprint arXiv:2505.06416*, 2025. URL <https://arxiv.org/abs/2505.06416>.
- Jessica Lyons. Anthropic won’t fix a bug in its sqlite mcp server, leaving thousands of forks vulnerable. https://www.theregister.com/2025/06/25/anthropic_sql_injection_flaw_unfixed/, June 2025. The Register, accessed 25 July 2025.
- Lori MacVittie. Mcp actually does introduce new risks—are you ready?, 2025. URL <https://www.f5.com/company/blog/mcp-actually-does-introduce-new-risks-are-you-ready>.
- Narek Maloyan and Dmitry Namiot. Breaking the protocol: Security analysis of the model context protocol specification and prompt injection vulnerabilities in tool-integrated llm agents. *arXiv preprint arXiv:2601.17549*, 2026.
- Travis Mandel, Jahnu Best, Randall H. Tanaka, Hiram Temple, Chansen Haili, Kayla Schlectinger, and Roy Szeto. Let’s keep it safe: Designing user interfaces that allow everyone to contribute to ai safety, 2019. URL <https://arxiv.org/abs/1907.04446>. arXiv preprint; published July 9, 2019.
- Matt. What are the critical MCP security risks every developer must know? <https://prefactor.tech/blog/mcp-security-risks-developer-guide>, July 2025. Prefactor blog, accessed 23 November 2025.
- Reed McGinley-Stempel. npm-audit for mcp security: A deep-dive on mcp-scan. <https://stytych.com/blog/mcp-scan/>, May 2025. Blog article (published 2 May 2025). <https://stytych.com/blog/mcp-scan/>.
- Reed McGinley-Stempel. Model context protocol (mcp): A comprehensive introduction for developers, 03 2025. URL <https://stytych.com/blog/model-context-protocol-introduction>. Accessed 2025-10-12.

- MCP Security Working Group. Runtime isolation for mcp servers. <https://modelcontextprotocol-security.io/build/runtime-isolation/>, 2025a. Accessed: 2025-11-25.
- MCP Security Working Group. Secrets & credential management. <https://modelcontextprotocol-security.io/hardening/secrets-management.html>, 2025b. Model Context Protocol Security Hardening Guide, accessed Nov. 25, 2025.
- MCP Working Group. Model context protocol specification v2025-06-18, 2025a. <https://modelcontextprotocol.io/specification/2025-06-18>.
- MCP Working Group. Authorization – security considerations: Confused deputy problem, 2025b. URL <https://modelcontextprotocol.io/specification/draft/basic/authorization>.
- Darren Meyer. Pypi supply chain attack uncovered: Colorama and colorizr name confusion. *Checkmarx Zero*, 2025. <https://checkmarx.com/zero-post/python-pypi-supply-chain-attack-colorama>.
- Microsoft. Oauth 2.0 client credentials grant flow (microsoft learn documentation), 2025. Describes client credential flow and JWT usage (living document). Accessed 19 Jul 2025. <https://learn.microsoft.com/en-us/azure/active-directory/develop/v2-oauth2-client-creds-grant-flow>.
- Microsoft. Microsoft extends zero trust to secure the agentic workforce, 2025a. URL <https://www.microsoft.com/en-us/security/blog/2025/05/19/microsoft-extends-zero-trust-to-secure-the-agentic-workforce/>. Microsoft Security Blog.
- Microsoft. Plug, play, and prey: The security risks of the model context protocol, 2025b. URL <https://techcommunity.microsoft.com/blog/microsoftdefendercloudblog/plug-play-and-prey-the-security-risks-of-the-model-context-protocol/4410829>.
- Microsoft. Use mcp servers in vs code, October 2025. URL <https://code.visualstudio.com/docs/copilot/customization/mcp-servers>. Shows adding or switching MCP servers via mcp.json without client code changes.
- MITRE Corporation. Cve - common vulnerabilities and exposures, 1999. URL <https://cve.mitre.org/cve/>. CVE program publicly launched September 1999; maintained by MITRE.
- MITRE Corporation. Mitre adversarial threat landscape for artificial-intelligence systems (atlas) fact sheet, 2023. URL https://atlas.mitre.org/pdf-files/MITRE_ATLAS_Fact_Sheet.pdf. Fact sheet referencing early 2023 releases; exact publication date not provided.
- Model Context Protocol. Architecture overview, 2025a. URL <https://modelcontextprotocol.io/docs/learn/architecture>. Documentation page describing MCP participants, layers (data and transport), lifecycle, and primitives.
- Model Context Protocol. Base protocol - lifecycle, 06 2025b. URL <https://modelcontextprotocol.io/specification/2025-06-18/basic/lifecycle>. Protocol Revision: 2025-06-18. Accessed 2025-10-12.
- Model Context Protocol. Schema reference - implementation, 06 2025c. URL <https://modelcontextprotocol.io/specification/2025-06-18/schema>. Defines implementation metadata fields {name, version, title?}. Accessed 2025-10-12.
- Model Context Protocol. Server features - prompts, 06 2025d. URL <https://modelcontextprotocol.io/specification/2025-06-18/server/prompts>. Servers expose discoverable prompt templates; clients can retrieve and parameterize them. Accessed 2025-10-12.
- Model Context Protocol. Server features - resources, 06 2025e. URL <https://modelcontextprotocol.io/specification/2025-06-18/server/resources>. Resources are uniquely identified by URI; include files, database schemas, app-specific info. Accessed 2025-10-12.

- Model Context Protocol. Model context protocol (mcp) - architecture specification. <https://modelcontextprotocol.io/specification/2025-06-18/architecture>, 2025f. Describes the client–host–server design and role separation.
- Model Context Protocol. Understanding mcp clients, June 2025g. URL <https://modelcontextprotocol.io/docs/learn/client-concepts>. Explains that each client handles one direct communication with one server.
- Model Context Protocol. Connect to local mcp servers, June 2025h. URL <https://modelcontextprotocol.io/docs/develop/connect-local-servers>. Claude Desktop config via `claude_desktop_config.json` to add servers; no client code changes.
- Model Context Protocol. Mcp lifecycle (initialization, operation, shutdown). <https://modelcontextprotocol.io/specification/2025-03-26/basic/lifecycle>, 2025i.
- Model Context Protocol. Security best practices – session hijacking, 2025j. URL https://modelcontextprotocol.io/specification/draft/basic/security_best_practices. Accessed 25 July 2025.
- Model Context Protocol. Model context protocol: Authorization, 2025k. URL <https://modelcontextprotocol.io/docs/tutorials/security/authorization>. Authorization specification and guidance.
- Model Context Protocol. Authorization - mcp basic specification, version 2025-06-18, 2025l. URL <https://modelcontextprotocol.io/specification/2025-06-18/basic/authorization>. Token audience binding; prohibition on token passthrough; confused deputy guidance.
- Model Context Protocol. Understanding mcp servers, 2025m. URL <https://modelcontextprotocol.io/docs/learn/server-concepts>.
- Model Context Protocol. Sampling (client features) - mcp specification, version 2025-06-18, 2025n. URL <https://modelcontextprotocol.io/specification/2025-06-18/client/sampling>.
- Model Context Protocol. Model context protocol: Security best practices, 2025o. URL https://modelcontextprotocol.io/specification/2025-06-18/basic/security_best_practices. Security best practices and mitigations.
- Model Context Protocol. Tools (server features) - security considerations, version 2025-06-18, 2025p. URL <https://modelcontextprotocol.io/specification/2025-06-18/server/tools>. Validating inputs, access controls, rate limiting, sanitizing outputs; tool list change notifications.
- Model Context Protocol. Model context protocol: Transports, 2025q. URL <https://modelcontextprotocol.io/docs/concepts/transports>. Documentation page.
- Model Context Protocol Security. Preference manipulation. <https://modelcontextprotocol-security.io/https/tool-poisoning/preference-manipulation/>, 2025. MCP Security TTPs: Tool Poisoning & Metadata Attacks.
- Model Context Protocol Security Working Group. Dependency vulnerabilities, 2025. URL <https://modelcontextprotocol-security.io/https/supply-chain/dependency-vulnerabilities/>. Model Context Protocol Security TTPs.
- Modelscope Team. Mcpbench: An evaluation framework for mcp servers, 2025. URL <https://github.com/modelscope/MCPBench>. Open-sourced on GitHub (supports evaluation of MCP Servers on accuracy, latency, token usage).
- Mallory Mooney. How to monitor MCP server activity for security risks. <https://www.datadoghq.com/blog/mcp-detection-rules/>, November 2025. Datadog blog, accessed Dec. 1, 2025.
- M. Morley. Json-rpc 2.0 specification. <https://www.jsonrpc.org/specification>, 2010. Accessed: 2025-10-11.

- Sumeet Ramesh Motwani, Mikhail Baranchuk, Martin Strohmeier, Vijay Bolina, Philip H. S. Torr, Lewis Hammond, and Christian Schroeder de Witt. Secret collusion among generative ai agents: Multi-agent deception via steganography. *arXiv*, 2024. URL <https://arxiv.org/abs/2402.07510>. v4 last revised 2025-04-14.
- Muayad Ali. Mcp security considerations for safer agent-tool interoperability, 2025. <https://writer.com/engineering/mcp-security-considerations/>.
- Mustafa Naamnih and Jonathan Ginzburg. Threat research: Hundreds of mcp servers vulnerable to abuse, June 2025. URL <https://www.backslash.security/blog/hundreds-of-mcp-servers-vulnerable-to-abuse>. Backslash Research Team blog post.
- Yousof Nahya. Mcp-kali-server. GitHub repository, April 2025. Initial release 8 Apr 2025. Accessed 19 Jul 2025. <https://github.com/Wh0am123/MCP-Kali-Server>.
- Yuri Nakao, Lorenzo Strappelli, Simone Stumpf, Aisha Naseer, Daniele Regoli, and Giulia Del Gamba. Towards responsible ai: A design space exploration of human-centered artificial intelligence user interfaces to investigate fairness. *International Journal of Human-Computer Interaction*, 39(9):1762–1788, May 2022. doi: 10.1080/10447318.2022.2067936. URL <https://doi.org/10.1080/10447318.2022.2067936>.
- V. Narajala and I. Habler. Enterprise-grade security for the model context protocol (mcp): Frameworks and mitigation strategies. *arXiv preprint arXiv:2504.08623*, 2025. URL <https://arxiv.org/abs/2504.08623>.
- Vineeth Sai Narajala, Ken Huang, and Idan Habler. Securing GenAI Multi-Agent Systems Against Tool Squatting: A Zero Trust registry-based approach. *arXiv preprint arXiv:2504.19951*, April 2025. doi: 10.48550/arXiv.2504.19951. URL <https://arxiv.org/abs/2504.19951>.
- National Institute of Standards and Technology. Artificial intelligence risk management framework: Generative artificial intelligence profile. Technical Report NIST AI 600-1, National Institute of Standards and Technology, Gaithersburg, MD, July 2024. W3C-style profile companion to AI RMF 1.0. Published July 2024 (release announced 26 Jul 2024). Accessed 19 Jul 2025. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>.
- National Vulnerability Database. CVE-2019-5736, 2019. URL <https://nvd.nist.gov/vuln/detail/CVE-2019-5736>. Accessed: March 15, 2026.
- National Vulnerability Database. CVE-2020-14343, 2020a. URL <https://nvd.nist.gov/vuln/detail/CVE-2020-14343>. Accessed: March 15, 2026.
- National Vulnerability Database. CVE-2020-1747, 2020b. URL <https://nvd.nist.gov/vuln/detail/CVE-2020-1747>. Accessed: March 15, 2026.
- National Vulnerability Database. CVE-2022-1471, 2022. URL <https://nvd.nist.gov/vuln/detail/CVE-2022-1471>. Accessed: March 15, 2026.
- National Vulnerability Database. CVE-2025-49596: Model Context Protocol Inspector remote code execution vulnerability. <https://nvd.nist.gov/vuln/detail/CVE-2025-49596>, 2025a. Accessed: 2025-11-22.
- National Vulnerability Database. CVE-2025-49596: Mcp inspector remote code execution, 2025b. URL <https://nvd.nist.gov/vuln/detail/CVE-2025-49596>.
- Nate Nelson. Hundreds of mcp servers expose ai models to abuse, rce, June 2025. URL <https://www.darkreading.com/cloud-security/hundreds-mcp-servers-ai-models-abuse-rce>. News article summarizing Backslash Security findings.
- Network Intelligence. Model context protocol (mcp) security checklist: Complete protection guide, 2026. URL <https://www.networkintelligence.ai/blogs/model-context-protocol-mcp-security-checklist>. Accessed: 2025-11-23.

- Lily Hay Newman. Microsoft’s ai red team has already made the case for itself, 2023. URL <https://www.wired.com/story/microsoft-ai-red-team/>. Published August 7, 2023 at 13:51 (site local time).
- Sharma Nikhil. Securing mcp servers: A comprehensive guide to authentication and authorization, 2025. URL <https://www.infracloud.io/blogs/securing-mcp-servers/>. See “traditional MCP implementations rely on static API tokens that provide all-or-nothing access.”
- OPA. Open policy agent documentation, 2025. URL <https://www.openpolicyagent.org/docs>.
- Open Policy Agent Contributors. Open policy agent: Policy-based control for cloud native environments. <https://openpolicyagent.org/>, 2025. Accessed Dec. 1, 2025.
- OpenAI. Define tools. <https://developers.openai.com/apps-sdk/plan/tools/>, 2025. “Use the checklist below to turn your use cases into well-scoped tools *before you touch the SDK*.” Accessed 2025-10-12.
- OpenSearch Project. Remove mcp tools api. <https://docs.opensearch.org/latest/ml-commons-plugin/api/mcp-server-apis/remove-mcp-tools/>, 2025. Accessed 25 July 2025.
- Laurent Orseau and Stuart Armstrong. Safely interruptible agents. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence (UAI 2016)*, pp. 557–566, June 2016. doi: 10.5555/3020948.3021006. URL <https://doi.org/10.5555/3020948.3021006>. Presented June 2016; original version released June 1, 2016; revised (v2) October 28, 2016.
- OSTC. The mosaic effect: Why “harmless” data isnt harmless. *Medium*, 2025. <https://ostradecraft.medium.com/the-mosaic-effect-why-harmless-data-isnt-harmless-bc63bbd3d072>.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, et al. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf. Also available as arXiv:2203.02155.
- OWASP. Llm prompt injection prevention cheat sheet. https://cheatsheetseries.owasp.org/cheatsheets/LLM_Prompt_Injection_Prevention_Cheat_Sheet.html, 2023. OWASP Cheat Sheet Series, accessed Nov. 25, 2025.
- OWASP Foundation. Owasp api security top 10 (2019) — a4: Lack of resources & rate limiting, 2019. OWASP API Security Top 10 2019 Edition, A4 entry. <https://owasp.org/API-Security/editions/2019/en/0xa4-lack-of-resources-and-rate-limiting/>.
- OWASP Foundation. Zero trust architecture cheat sheet, 2021. URL https://cheatsheetseries.owasp.org/cheatsheets/Zero_Trust_Architecture_Cheat_Sheet.html. OWASP Cheat Sheet Series guidance on Zero Trust principles and implementation steps.
- OWASP Foundation. Owasp api security top 10 (2023) — a4: Unrestricted resource consumption, 2023a. OWASP API Security Top 10 2023 Edition, A4 entry. <https://owasp.org/API-Security/editions/2023/en/0xa4-unrestricted-resource-consumption/>.
- OWASP Foundation. Owasp top 10 for large language model applications, version 1.1, 2023b. URL <https://owasp.org/www-project-top-10-for-large-language-model-applications/>. Version 1.1 update announced October 18, 2023.
- OWASP Foundation. Owasp genai — agentic security initiative, 2025a. Initiative overview page. Accessed 19 Jul 2025. <https://genai.owasp.org/initiatives/agentic-security>.
- OWASP Foundation. OWASP MCP Top 10. <https://owasp.org/www-project-mcp-top-10/>, 2025b. Accessed 2025-11-29.

- OWASP Foundation. Input validation cheat sheet, 2025c. URL https://cheatsheetseries.owasp.org/cheatsheets/Input_Validation_Cheat_Sheet.html. OWASP Cheat Sheet Series; accessed 19 Jul 2025.
- OWASP Foundation. Owasp top 10: 2025 risks & mitigations for llms and genai apps, 2025d. URL <https://genai.owasp.org/llm-top-10/>.
- Maria Paktiti. The complete guide to MCP security: How to secure MCP servers & clients. <https://workos.com/blog/mcp-security-risks-best-practices>, 2025. WorkOS Blog, accessed 2025-11-22.
- Palo Alto Networks. Model context protocol (mcp): A security overview, 2025. <https://www.paloaltonetworks.com/blog/cloud-security/model-context-protocol-mcp-a-security-overview>.
- Palo Alto Networks Community. Mcp security exposed: What you need to know now, 2025. <https://live.paloaltonetworks.com/t5/community-blogs/mcp-security-exposed-what-you-need-to-know-now/ba-p/1227143>.
- Palo Alto Networks Unit 42. Ai agents are here. so are the threats., 2025. <https://unit42.paloaltonetworks.com/agent-ai-threats/>.
- Sean Park. Why a classic mcp server vulnerability can undermine your entire ai agent. https://www.trendmicro.com/en_us/research/25/f/why-a-classic-mcp-server-vulnerability-can-undermine-your-entire-ai-agent.html, June 2025. Trend Micro Research blog, accessed 25 July 2025.
- Amadeo Pellicce. [sane] local mcp. <https://blog.openint.dev/sane-local-mcp/>, 2025. Accessed 24 July 2025.
- Phylum Research. Disrupting a pypi software supply-chain threat actor, 2022. URL <https://blog.phylum.io/disrupting-a-software-supply-chain-threat-actor-building-a-botnet/>. Phylum Research blog post on tracking and disrupting the Gooberworm campaign.
- Pillar Security. The security risks of model context protocol (mcp). *Pillar Security Blog*, 2025. URL <https://www.pillar.security/blog/the-security-risks-of-model-context-protocol-mcp>.
- Ping Identity. Securing MCP servers. <https://developer.pingidentity.com/identity-for-ai/agents/idai-securing-mcp-servers.html>, 2025. Accessed: 2025-11-25.
- Pomerium. Secure access for mcp, 2023. URL <https://www.pomerium.com/blog/secure-access-for-mcp>.
- Christian Posta. Part one: MCP authorization the hard way, 2025a. URL <https://www.solo.io/blog/understanding-mcp-authorization-step-by-step-part-one>. Deep dive into MCP Authorization with step-by-step examples and source code.
- Christian Posta. Deep dive: Mcp and a2a attack vectors for ai agents, 2025b. <https://www.solo.io/blog/deep-dive-mcp-and-a2a-attack-vectors-for-ai-agents/>.
- Christian Posta. The mcp authorization spec is a mess for enterprise. <https://blog.christianposta.com/the-updated-mcp-oauth-spec-is-a-mess/>, 2025c. Accessed 24 July 2025.
- Prompt Security Team. Security for agentic ai: Unveiling mcp gateway & mcp risk assessment, May 2025a. URL <https://www.prompt.security/blog/security-for-agentic-ai-unveiling-mcp-gateway-mcp-risk-assessment>. Blog post.
- Prompt Security Team. Introducing prompt security’s ai risk assessment tool, July 2025b. URL <https://www.prompt.security/blog/introducing-prompt-securitys-ai-risk-assessment-tool>. Blog post announcing risk scoring for AI sites and MCP servers.
- Proofpoint Threat Research. Beyond credentials: Weaponizing oauth applications for persistent cloud access. *Proofpoint Threat Insight Blog*, 2025. URL <https://www.proofpoint.com/us/blog/threat-insight/beyond-credentials-weaponizing-oauth-applications-persistent-cloud-access>.

- Brandon Radosevich and John Halloran. MCP Safety Audit: LLMs with the model context protocol allow major security exploits. *arXiv preprint arXiv:2504.03767*, April 2025. doi: 10.48550/arXiv.2504.03767. URL <https://arxiv.org/abs/2504.03767>.
- Ajeet Singh Raina. Mcp horror stories: The security issues threatening ai infrastructure, July 2025. URL <https://www.docker.com/blog/mcp-security-issues-threatening-ai-infrastructure/>. Blog post summarizing ecosystem-wide MCP security findings.
- Snegha Ramnarayanan. Model context protocol security. <https://www.wiz.io/academy/model-context-protocol-security>, 2025. Wiz CloudSec Academy, accessed 2025-11-22.
- Nikitha Rao, Elizabeth Gilbert, Tahina Ramananandro, Nikhil Swamy, Claire Le Goues, and Sarah Fakhoury. Diffspec: Differential testing with llms using natural language specifications and code artifacts, 2025. arXiv preprint, version 3 (revised 6 May 2025; original submission 5 Oct 2024). DOI: 10.48550/arXiv.2410.04249. <https://arxiv.org/abs/2410.04249>.
- Rohit Raveendran. Decoding the pod termination lifecycle in kubernetes: a comprehensive guide, December 2024. URL <https://www.cncf.io/blog/2024/12/19/decoding-the-pod-termination-lifecycle-in-kubernetes-a-comprehensive-guide/>. Member post, CNCF Blog.
- Red Hat, Inc. *Security Hardening*. Red Hat Enterprise Linux Documentation, 2025. PDF available at https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/8/pdf/security_hardening/security-hardening.pdf.
- Red Hat Product Security. Model context protocol (mcp): Understanding security risks and controls, 2025. URL <https://www.redhat.com/en/blog/model-context-protocol-mcp-understanding-security-risks-and-controls>.
- Scott Rose, Oliver Borchert, Stu Mitchell, and Sean Connelly. Zero trust architecture. NIST Special Publication 800-207, National Institute of Standards and Technology, Gaithersburg, MD, August 2020. URL <https://csrc.nist.gov/pubs/sp/800/207/final>.
- R. Ross, M. McEvelley, and M. Winstead. Engineering trustworthy secure systems. NIST Special Publication (SP) 800-160, Volume 1, Revision 1, National Institute of Standards and Technology, Gaithersburg, MD, 11 2022. URL <https://doi.org/10.6028/NIST.SP.800-160v1r1>.
- Margaret Rouse. What is siem (security information and event management)?, 2023. URL <https://www.techtarget.com/searchsecurity/definition/security-information-and-event-management-SIEM>. Accessed 2025-07-19.
- Ahmad Sadeddin. Securing model context protocol (mcp) servers: Threats and best practices. <https://corgea.com/Learn/securing-model-context-protocol-%28mcp%29-servers-threats-and-best-practices>, 2025. Online; accessed 24 July 2025.
- Abu Sakib. Monitoring mcp security and agent behavior with moesif, July 2025. URL <https://www.moesif.com/blog/monitoring/model-context-protocol/Monitoring-MCP-Security-And-Agent-Behavior-With-Moesif/>. Accessed 25 July 2025.
- ScaleSec. MCP server security best practices for safer AI adoption. <https://scalesec.com/blog/mcp-server-security-best-practices>, 2025. Accessed: 2025-11-25.
- Ory Segal, Aviv Sasson, and Elad Shuster. The new ai attack surface - how cortex cloud secures mcp. <https://www.paloaltonetworks.com/blog/2025/06/cloud-security-model-context-protocol-mcp-security/>, June 2025.
- Adriana Sejfia and Max Schäfer. Practical automated detection of malicious npm packages. In *Proceedings of the 44th International Conference on Software Engineering (ICSE '22)*, pp. 1681–1692, Pittsburgh, PA, USA, May 2022. ACM. doi: 10.1145/3510003.3510104. URL <https://doi.org/10.1145/3510003.3510104>.

- Yujia Shen, Peng Zhang, and et al. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *arXiv preprint arXiv:2303.17580*, 2024.
- Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, 2017. URL <https://arxiv.org/abs/1610.05820>.
- Sigstore. Rekor: Immutable, tamper resistant transparency logs for the software supply chain. <https://docs.sigstore.dev/logging/overview/>, 2022. accessed Dec. 1, 2025.
- Aditi Singh, Abul Ehtesham, Saket Kumar, and Tala Talaei Khoei. A survey of the model context protocol (mcp): Standardizing context to enhance large language models (llms). *Preprints*, April 2025. doi: 10.20944/preprints202504.0245.v1. URL <https://www.preprints.org/manuscript/202504.0245/v1>. Preprint; not peer reviewed.
- SlowMist Team. MCP Security Checklist: A Security Guide for the AI Tool Ecosystem. <https://github.com/slowmist/MCP-Security-Checklist>, 2025. Accessed: 2025-11-23.
- SLSA Community. Slsa specification, version 1.1 (approved), 2025. URL <https://slsa.dev/spec/v1.1/>.
- Smithery. Smithery cli - command line tool for mcp server management. <https://github.com/smithery-ai/cli>, 2025.
- Snyk Security Research. Arbitrary code execution in autogen, October 2024. URL <https://security.snyk.io/vuln/SNYK-PYTHON-AUTOGEN-8161199>. Advisory ID: SNYK-PYTHON-AUTOGEN-8161199; disclosed 2024-10-01.
- SOCRadar. Mcp servers: Everything you need to know, 2025. URL <https://socradar.io/mcp-servers-everything-you-need-to-know/>. “Bonus 3: Most Common Configuration Mistakes in Real Deployments.”
- Hao Song, Yiming Shen, Wenxuan Luo, Leixin Guo, Ting Chen, Jiashui Wang, Beibei Li, Xiaosong Zhang, and Jiachi Chen. Beyond the protocol: Unveiling attack vectors in the model context protocol ecosystem. *arXiv preprint arXiv:2506.02040*, 2025. URL <https://arxiv.org/abs/2506.02040>.
- Murugiah Souppaya, John Morello, and Karen Scarfone. Application container security guide, 2017. URL <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-190.pdf>. Final release.
- Murugiah Souppaya, Karen Scarfone, and Donna Dodson. Secure software development framework (ssdf) version 1.1: Recommendations for mitigating the risk of software vulnerabilities. NIST Special Publication (SP) 800-218, National Institute of Standards and Technology, Gaithersburg, MD, 2022a. URL <https://doi.org/10.6028/NIST.SP.800-218>.
- Murugiah Souppaya, Karen Scarfone, and Donna Dodson. Secure software development framework (ssdf) version 1.1: Recommendations for mitigating the risk of software vulnerabilities. NIST Special Publication NIST Special Publication 800-218, National Institute of Standards and Technology, Gaithersburg, MD, February 2022b. URL <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-218.pdf>. Version 1.1.
- Speakeasy. What is mcp authorization? <https://www.speakeasy.com/mcp/authorization>, 2025. Accessed 24 July 2025.
- Ryan Spletzer. The many contexts of model context protocol. <https://www.spletzer.com/2025/05/the-many-contexts-of-model-context-protocol/>, May 2025. Accessed: 2025-11-22.
- Manu Sporny, Dave Longley, Markus Sabadello, Drummond Reed, Orie Steele, Christopher Allen, Jonathan Holt, and Dmitri Zagidulin. Decentralized identifiers (dids) v1.0, July 2022. W3C Recommendation 19 July 2022. <https://www.w3.org/TR/did-core/>.
- sshh.io. Everything wrong with mcp. <https://blog.sshh.io/p/everything-wrong-with-mcp>, 2025.

- SsoJet. Mcp security best practices: How to secure model context protocol. <https://ssojet.com/blog/what-are-the-best-practices-for-mcp-security>, 2025. Accessed: 2025-11-25.
- Strobes Security Labs. Model context protocol and its critical vulnerabilities, 2025. URL <https://strobess.co/blog/mcp-model-context-protocol-and-its-critical-vulnerabilities/>.
- Stytch. Mcp authentication and authorization implementation guide, 07 2025. URL <https://stytch.com/blog/MCP-authentication-and-authorization-guide/>. Accessed 2025-10-12.
- Supabase. Model context protocol (mcp). URL <https://supabase.com/docs/guides/getting-started/mcp>. Supabase MCP documentation and security risks guidance.
- SUSE. Apparmor profile components and syntax: Ptrace rules, 2025. AppArmor ptrace limitation and rule semantics. <https://documentation.suse.com/en-us/sles/15-SP7/html/SLES-all/cha-apparmor-profiles.html>.
- Sysdig. Falco: Open source runtime security, 2023. URL <https://falco.org>.
- Elham Tabassi. Artificial intelligence risk management framework (ai rmf 1.0). NIST AI Series AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD, 01 2023a. URL <https://doi.org/10.6028/NIST.AI.100-1>.
- Elham Tabassi. Artificial intelligence risk management framework (ai rmf 1.0). NIST AI Series NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD, January 2023b. URL <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10>.
- The Hacker News. Malicious pypi, npm, and ruby packages exposed in ongoing campaign. *The Hacker News*, 2025. <https://thehackernews.com/2025/06/malicious-pypi-npm-and-ruby-packages.html>.
- Marco Tonnarelli, Filippo Scaramuzza, Simon Harrer, and Linus W. Dietz. Data product mcp: Chat with your enterprise data. *arXiv preprint arXiv:2601.08687*, 2026. doi: 10.48550/arXiv.2601.08687.
- Bill Toulas. Supply chain attack hits npm package with 45,000 weekly downloads. *BleepingComputer*, 2025. <https://www.bleepingcomputer.com/news/security/supply-chain-attack-hits-npm-package-with-45-000-weekly-downloads/>.
- Nova Trauben. Resolving CVE-2022-1471 with the snakeyaml 2.0 release. <https://www.veracode.com/blog/resolving-cve-2022-1471-snakeyaml-20-release-0/>, 2023. Veracode Blog, March 3.
- TrollEye Security. Thousands of AI MCP servers exposed without any authentication controls, 2025. URL <https://www.trolleyesecurity.com/articles-news-mcp-servers-exposed-without-any-security/>. Accessed: 2025-11-22.
- TrustGraph. tg-delete-mcp-tool. <https://docs.trustgraph.ai/reference/cli/tg-delete-mcp-tool.html>, 2025. Accessed 25 July 2025.
- VeraCrypt Project. Memory dump files, 2015. <https://veracrypt.io/en/Memory%20Dump%20Files.html>.
- Nell Watson, Ahmed Amer, Evan Harris, Preeti Ravindra, and Shujun Zhang. Personalized constitutionally-aligned agentic superego: Secure ai behavior aligned to diverse human values, 2025. arXiv preprint. <https://arxiv.org/abs/2506.13774>.
- T. Wigren and Software Freedom Conservancy. Json-rpc 2.0 specification. <https://www.jsonrpc.org/specification>, 2020.
- Simon Willison. Supabase mcp can leak your entire sql database, July 2025a. URL <https://simonwillison.net/2025/Jul/6/supabase-mcp-lethal-trifecta/>. Link blog post discussing the Supabase MCP support-ticket prompt injection scenario.
- Simon Willison. Model context protocol has prompt-injection security problems. <https://simonwillison.net/2025/Apr/9/mcp-prompt-injection/>, 2025b.

- Wiz Research. Mcp and llm security research briefing. <https://www.wiz.io/blog/mcp-security-research-briefing>, April 2025. Accessed Dec. 1, 2025.
- Yulong Wu, Viktor Schlegel, and Riza Batista-Navarro. Natural context drift undermines the natural language understanding of large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, Suzhou, China, 2025. Association for Computational Linguistics. doi: 10.18653/v1/2025.findings-emnlp.65. URL <https://aclanthology.org/2025.findings-emnlp.65/>.
- Zhen Xiang, Z. Yu, C. Xiao, Y. Wu, and X. Zhang. Guardagent: Safeguard llm agents by a guard agent via knowledge-enabled reasoning, 2024. URL <https://arxiv.org/abs/2406.09187>. arXiv preprint; revised version (v2) published 2025-03-09.
- Wenpeng Xing, Zhonghao Qi, Yupeng Qin, Yilin Li, Caini Chang, Jiahui Yu, Changting Lin, Zhenzhen Xie, and Meng Han. Mcp-guard: A defense framework for model context protocol integrity in large language model applications, 2025. URL <https://arxiv.org/abs/2508.10991>. Version 2.
- Jiashu Xu, Mingyu Derek Ma, Fei Wang, Chaowei Xiao, and Muhao Chen. Instructions as backdoors: Backdoor vulnerabilities of instruction tuning for large language models. *arXiv:2305.14710*, 2023. URL <https://arxiv.org/abs/2305.14710>.
- Yifan Yang, Zhi Liu, Rohan Motwani, Kai Zhang, and Yuchen Xu. Mcpsecbench: A systematic security benchmark and playground for the model context protocol. *arXiv preprint arXiv:2508.13220*, 2025. URL <https://arxiv.org/abs/2508.13220>.
- Zhiyuan Yao, Zishan Xu, Yifu Guo, Zhiguang Han, Cheng Yang, Shuo Zhang, Weinan Zhang, Xingshan Zeng, and Weiwen Liu. Toolace-mcp: Generalizing history-aware routing from mcp tools to the agent web. *arXiv preprint arXiv:2601.08276*, 2026. doi: 10.48550/arXiv.2601.08276.
- Sarah Young and Den Delimarsky. Protecting against indirect prompt injection attacks in mcp, 2025. URL <https://devblogs.microsoft.com/blog/protecting-against-indirect-injection-attacks-mcp>. Guidance blog post on mitigating indirect (cross-domain) prompt injection and tool poisoning in Model Context Protocol.
- Chaojia Yu, Zihan Cheng, Hanwen Cui, Yishuo Gao, Zexu Luo, Yijin Wang, Hangbin Zheng, and Yong Zhao. A survey on agent workflow – status and future. *arXiv preprint arXiv:2508.01186*, 2025. doi: 10.48550/arXiv.2508.01186. URL <https://arxiv.org/abs/2508.01186>.
- Dongsen Zhang, Zekun Li, Xu Luo, Xuannan Liu, Peipei Li, and Wenjun Xu. MCP Security Bench (MSB): Benchmarking Attacks Against Model Context Protocol in LLM Agents. *arXiv preprint arXiv:2510.15994*, 2025. doi: 10.48550/arXiv.2510.15994. URL <https://arxiv.org/abs/2510.15994>.
- Xuandong Zhao, Xianjun Yang, Tianyu Pang, Chao Du, Lei Li, Yu-Xiang Wang, and William Yang Wang. Weak-to-strong jailbreaking on large language models, 2025. URL <https://openreview.net/forum?id=Nazzz5GJ4g>. Submitted to ICLR 2025; revised 5 Feb 2025; original submission 25 Sep 2024.
- Allen Zhou. Diving into the MCP authorization specification. <https://www.descope.com/blog/post/mcp-auth-spec>, 2025. Online; accessed 24 July 2025.
- Zhenhong Zhou, Yuanhe Zhang, Hongwei Cai, Moayad Aloqaily, Ouns Bouachir, Linsey Pang, Prakhar Mehrotra, Kun Wang, and Qingsong Wen. Mcpshield: A security cognition layer for adaptive trust calibration in model context protocol agents. *arXiv preprint arXiv:2602.14281*, 2026.
- Xuanjun Zong, Zhiqi Shen, Lei Wang, Yunshi Lan, and Chao Yang. Mcp-safetybench: A benchmark for safety evaluation of large language models with real-world mcp servers. *arXiv preprint arXiv:2512.15163*, 2025.

A Appendix

This appendix provides standalone definitions for each threat in the MCP threat taxonomy. Each subsection aligns with a category defined in Section 3 of the main paper. To avoid redundancy, the dual-axis mappings to components, lifecycle phases, STRIDE, and MAESTRO are not reproduced here, as they are already presented in Tables 1 through 3. This appendix is intended to serve as a reference for readers seeking precise definitions of individual threats beyond the synthesized treatment in Section 3.

A.1 Identity and Trust Management

This category addresses how MCP deployments name, authenticate, and bind clients, servers, and tools. The absence of a global naming authority and the lack of protocol-level cryptographic identity enforcement mean that clients and agents routinely rely on self-reported metadata to make trust decisions. The four threats below represent the primary failure modes through which this structural weakness is exploited.

A.1.1 Server Name Collisions and Impersonation

Adversaries register MCP server names that are typographically or visually similar to trusted ones, for example `dropboxuploader` versus `dropbox-uploader`, to redirect users or agents toward attacker-controlled endpoints. During configuration and discovery in ecosystems that rely primarily on human-readable names and descriptions, and that do not enforce verified namespaces, signed metadata, or equivalent provenance checks, clients may bind to a lookalike server instead of the intended one. As a result, names alone are not a reliable basis for distinguishing legitimate servers from impersonating ones without additional trust verification.

Evidence. Empirical analyses of MCP-style registries and adjacent plugin ecosystems show that weak submission vetting allows adversaries to publish lookalike packages and endpoints, creating realistic typosquatting and impersonation opportunities that are difficult for both human operators and automated agents to detect during routine registration and discovery Hou et al. (2025); Posta (2025b); dts-solution (2025).

Impact. Once a client or agent is bound to an impersonating server, the adversary can intercept sensitive context, inject malicious tool outputs, and alter execution paths without detection. Because the same client configuration is often reused across multiple agents and hosts, a single successful name collision can propagate as a persistent supply chain foothold across an entire MCP deployment.

A.1.2 Mutual Authentication and Encryption Gaps

MCP’s early specification defined client–server transports including `stdio` and HTTP without mandating mutual authentication or secure channel enforcement at the protocol level, leaving TLS configuration, certificate validation, and client identity verification to individual implementations rather than requiring them as baseline protocol guarantees Hanson & maintainers (2024). OAuth-based authorization was later introduced as an optional transport-level mechanism for HTTP-based deployments rather than a universal protocol requirement, so authentication and authorization strength may vary across deployments. In deployments lacking mutual authentication, strict certificate validation, or equivalent endpoint verification, an attacker who controls a network path or misconfigured gateway may be able to interpose on traffic without protocol-level detection.

Evidence. Independent security reviews report MCP endpoints exposed on the public internet with weak or absent authentication controls, or relying only on server-side TLS without client identity checks Model Context Protocol (2025l); TrollEye Security (2025). Security guidance for MCP deployments increasingly recommends mutual TLS or equivalent cryptographic endpoint verification in response to these observed gaps Ramnarayanan (2025); Paktiti (2025); Kanoon (2025).

Impact. A client may silently bind to attacker-controlled endpoints, transmit sensitive context, and act on manipulated tool outputs, enabling man-in-the-middle attacks, credential theft, and long-lived workflow hijacking across multiple hosts and orchestrated agents.

A.1.3 Lack of Verified Server Identity or PKI

MCP does not define a native mechanism for binding servers to cryptographic identities beyond generic HTTPS/TLS and OAuth infrastructure, so clients and registries may rely heavily on human-readable metadata such as names, descriptions, or repository provenance when selecting servers Anthropic (2025); Posta (2025b). Unlike protocols that enforce a standardized ecosystem-wide identity-binding mechanism for endpoint authenticity, MCP leaves identity verification largely to deployment-specific choices. As a result, environments that rely primarily on metadata-based discovery without stronger provenance checks may remain vulnerable to server impersonation and misleading server listings.

Evidence. Security analyses highlight that MCP ecosystems can expose identity-confusion risks when discovery depends on names, descriptions, or other human-readable metadata without stronger provenance verification, so effective identity assurance depends on deployment-specific controls rather than a universal protocol-enforced identity layer Anthropic (2025); Posta (2025b).

Impact. The absence of a standardized identity-binding layer makes it more difficult to reliably distinguish legitimate from adversarial servers during registration and discovery, increasing the risk of impersonation, workflow hijacking, and long-lived malicious footholds, particularly in decentralized MCP ecosystems where trust verification is inconsistent.

A.1.4 Implicit Trust in Unauthenticated Servers

Early MCP deployments often ran servers on localhost or within tightly controlled internal networks, where implicit trust and minimal authentication were treated as acceptable design assumptions Spletzer (2025). These assumptions were reasonable in their original setting, but they became embedded in tools, configurations, and client implementations that were later reused in broader network contexts without equivalent security hardening. When such tools and agents are later exposed over HTTP, cloud infrastructure, or shared internal networks, the inherited trust model can become exploitable because network-reachable services may accept tool invocations without robust client authentication or identity verification.

Evidence. CVE-2025-49596 documents a concrete instance of this pattern. In affected versions of MCP Inspector below 0.14.1, authentication was missing between the Inspector client and proxy. As a result, unauthenticated requests could launch MCP commands over `stdio`, leading to remote code execution in the context of the vulnerable process National Vulnerability Database (2025a).

Impact. Agents that bind to and invoke unverified services can effectively extend their own privileges to those services, allowing malicious or misconfigured servers to trigger privileged tool executions, expand access, and exfiltrate sensitive data while the surrounding workflow continues to appear legitimate.

A.2 Supply Chain and Code Integrity Risks

MCP’s reliance on decentralized, third-party tool ecosystems and ad hoc deployment practices means that many deployments operate without centralized vetting or consistent integrity policies for server binaries, manifests, and libraries. The five threats below represent the primary failure modes through which supply chain and code integrity breakdowns arise.

A.2.1 Trojanized or Compromised Installers

Tools such as `Smithery-CLI`, `mcp-get`, and `mcp-installer` provide streamlined workflows for discovering, installing, and configuring MCP servers for use with clients, agents, and hosts Smithery (2025); mcp (2025a;b). If an installer binary, its package source, or its execution path is trojanized, the same convenience workflow can become a delivery mechanism for a backdoored MCP server or malicious configuration. A trojanized installer can add malicious server entries or package-based servers under plausible names and persist those changes in local configuration files or installation paths that are reused by clients on subsequent runs. A single privileged run can add rogue server entries, alter local configuration, or install malicious packages that blend into later discovery and execution workflows.

Evidence. Public documentation for these installers describes discovery, installation, local client integration, and configuration-management capabilities Smithery (2025); mcp (2025a;b). These documented capabilities confirm that installer-mediated workflows can modify the same persistent configuration surface on which MCP-aware clients later rely.

Impact. Any MCP-aware client that relies on the modified configuration risks telemetry exfiltration, credential harvesting via newly installed tools, and staged payloads that execute in later updates or orchestrated workflows. Such compromise can persist across multiple agents, hosts, and projects that share the same installer-managed environment.

A.2.2 Malicious Dependencies and Backdoored Code

MCP tools distributed through ecosystems such as PyPI and npm inherit the broader software supply-chain risks of those platforms, including dependency hijacking, typosquatting, and backdoor injection into both direct and transitive dependencies. Attackers can exploit dependency chains to deliver malicious code through packages that are installed either directly or indirectly, allowing MCP servers or tools built on those packages to execute attacker-controlled code with the tool’s privileges Sejfia & Schäfer (2022). Within an MCP deployment, packaging a server or tool around such backdoored dependencies means that routine tool invocation may silently activate malicious functionality.

Evidence. Recent incidents such as the compromise of `rand-user-agent` demonstrate how seemingly innocuous packages can be repurposed into delivery vectors for remote access trojans Toulas (2025); Eriksen, Charlie (2025). Typosquatting and name-confusion cases involving `colorizr` and `colorama` similarly show how small, deceptive naming changes can steer operators toward malicious code Meyer (2025); The Hacker News (2025); Cohen & Yerushalmi (2024).

Impact. Routine MCP tool calls can become reliable opportunities for remote access, credential theft, or data exfiltration wherever the compromised package or image is reused, even when the tool appears to behave normally from the operator’s perspective.

A.2.3 Lack of Integrity Verification

MCP’s core protocol does not define a universal, mandatory mechanism for message signing or end-to-end integrity verification across transports and package types, so integrity assurance for tool calls, server artifacts, and server responses depends on transport-specific, registry-specific, and deployment-specific controls rather than a protocol-wide guarantee Strobes Security Labs (2025). An adversary who can modify a repository, package feed, or other software distribution path may be able to swap or patch MCP server code or alter its behavior without any MCP-specific integrity signal.

Evidence. Security guidance recommends signing MCP components, verifying dependency integrity, and using hardened build pipelines with SAST and software composition analysis, treating MCP servers as ordinary executable code whose integrity protections must be implemented around MCP deployments rather than being enforced by the protocol itself Red Hat Product Security (2025).

Impact. Routine installs, updates, and tool invocations can become software supply-chain or message-integrity attack vectors that propagate across clients and workflows that depend on a compromised server, especially where stronger registry, transport, or build-pipeline controls are absent.

A.2.4 Vulnerable Redeployments and Stale Manifests

Lost, delayed, or corrupted updates can cause MCP deployments to drift into inconsistent state, leaving different agents or hosts operating on conflicting context or stale tool assumptions. In environments where MCP configurations and tool definitions are reused across developer IDEs, CI systems, and shared agent hosts, incomplete propagation of updates can cause some nodes to continue operating on outdated state while others advance to newer behavior.

Evidence. Production-oriented MCP security guidance describes drift arising when updates are lost, delayed, or corrupted, causing agents to operate on conflicting information. The same guidance recommends

strict synchronization, integrity checks, and automated conflict resolution to reduce this risk MacVittie (2025).

Impact. Hosts and agents may continue invoking tools under stale assumptions or inconsistent state long after a corrective update is issued. An attacker or failure condition that exploits this inconsistency can therefore affect multiple workflows and nodes before operators recognize that different parts of the deployment are no longer operating on the same effective state.

A.2.5 Unverified Package Sources and Automatic Invocation

In MCP deployments, hosts can translate LLM-generated outputs into external tool invocations once servers are configured, making server-provided tool metadata part of normal agent workflow selection. Recent ecosystem analyses of several public MCP registries report weak or inconsistent vetting and show that a substantial number of listed servers can be hijacked Li & Gao (2025). When servers from unverified sources are installed, a compromised server can place malicious tools directly in the host’s execution path, especially in deployments that auto-execute tools or rely on permissive approval settings.

Evidence. A large-scale security scan of 1,000 public MCP servers found that about 32% exposed at least one critical vulnerability Birur (2025). This indicates that unverified server sources represent a material risk in the current MCP ecosystem.

Impact. A single compromised server can exfiltrate chat or backend data or trigger unintended actions during otherwise routine workflows, with the risk amplified when hosts do not require strong confirmation, input review, or result-validation safeguards before tool execution Model Context Protocol (2025p).

A.3 Malicious or Poisoned Tool Servers

Because MCP clients rely on server-advertised names, schemas, and tool metadata to bind calls, malicious or manipulated definitions can misbind requests to untrusted implementations, inject tainted context, and exploit name or command collisions to bypass policy. The seven threats below represent the primary failure modes through which this binding surface is exploited.

A.3.1 Poisoned Tool Metadata

Adversaries can embed hidden instructions in MCP tool descriptions or description-bearing docstrings that are surfaced to the LLM during tool discovery and selection, causing these fields to be interpreted as instructions rather than neutral metadata. Poisoned metadata can influence agent behavior by steering the model toward high-privilege actions or covert exfiltration paths. A seemingly harmless `add()` tool whose description instructs the model to read and exfiltrate `~/.ssh/id_rsa` illustrates how a host that auto-registers or insufficiently reviews such a tool can be coerced into leaking credentials during routine completions Invariant Labs (2025).

Evidence. Large-scale empirical evidence shows that description-level cues materially shape agent tool selection and that MCP does not enforce consistency between tool descriptions and underlying implementations, allowing benign-looking descriptions to mask materially different behavior Li et al. (2026b).

Impact. Hosts or agents that auto-register tools or insufficiently validate tool descriptions can be coerced into leaking SSH keys, configuration files, or API credentials from their local environment during otherwise routine workflows, turning generic tool metadata into a reusable exfiltration channel.

A.3.2 Tool Shadowing and Name Collisions

When multiple MCP servers expose tools with identical or confusingly similar names, adversaries can exploit string-based discovery and selection logic to place attacker-controlled tools alongside legitimate ones. Deliberate tool shadowing and name collisions can divert invocations intended for trusted tools toward malicious implementations that reuse similar names and descriptions, allowing attackers to hijack critical workflows while mimicking the expected interface.

Evidence. Security analysis of MCP tool selection shows that tool name collisions arise when different servers offer identical or similar tool names, and that AI models rely on tool names and descriptions when choosing which tool to invoke. This makes similarly named malicious tools a realistic way to bias selection and redirect execution, especially when hosts aggregate tools from multiple servers without stronger provenance or approval controls Beretta et al. (2025).

Impact. Data and execution can be silently redirected across workflows even though invocations still appear to be routine tool calls. Attackers may intercept arguments, model outputs, and configuration updates intended for trusted tools, then inject modified artifacts back into the same pipelines.

A.3.3 Pre-Invocation Code Injection

Malicious MCP tools can embed harmful behavior either in description-level metadata or in unsafe execution paths. Line-jumping attacks place immediate-action directives in tool description fields, allowing connected servers to influence agent behavior or trigger hidden actions before any explicit tool invocation occurs. Separately, tools that interpolate unvalidated parameters directly into shell commands, such as `os.system("notify-send " + info["msg"])`, can turn crafted inputs into command injection and remote code execution during routine invocation.

Evidence. Published MCP security analyses describe both patterns. Security guidance on line-jumping shows that malicious instructions embedded in tool description fields can influence agent behavior before any explicit tool invocation occurs Fenstermacher (2025). Separate analysis of insecure MCP server implementations shows that unvalidated shell interpolation can turn ordinary tool inputs into command injection and remote code execution during routine invocation Cross (2025).

Impact. Compromised tools can leak secrets, modify configuration, or run arbitrary commands during routine agent workflows, sometimes before the user recognizes that a sensitive action has been triggered and sometimes through seemingly normal tool execution.

A.3.4 Rug Pull via Dynamic Tool Redefinition

A tool that is initially trusted can later be redefined if its server code or tool description is updated in place, while the host continues to resolve the same registered tool. Song et al. describe this as a rug-pull pattern, in which an MCP server first appears benign and is later modified to embed malicious instructions; they also show that malicious servers uploaded to aggregation platforms can trigger harmful actions in users’ local environments after installation Song et al. (2025).

Evidence. ETDI proposes immutable, versioned, cryptographically signed tool definitions as a direct defense against rug-pull attacks. It further argues that changes to a tool’s effective behavior should require a new definition version and renewed approval, indicating that baseline MCP does not provide these protections by default Bhatt et al. (2025).

Impact. Previously benign servers can become durable backdoors, allowing attackers to hijack file access, API calls, or transaction flows across hosts and workflows that continue to trust the original registration, even though the tool’s effective behavior has changed.

A.3.5 Preference Manipulation

Attackers can steer tool selection by shaping what the agent sees during discovery through persuasive descriptions that claim greater safety or capability, suggest that a tool is the preferred default, or otherwise increase its relative salience compared with competing tools. Unlike poisoned tool metadata that embeds explicit directives, preference manipulation relies on subtle non-imperative phrasing and selection bias so that the attacker’s tool is chosen under ordinary tool-selection behavior. The manipulation operates entirely within the metadata layer exposed to the model during discovery.

Evidence. Public MCP security analyses consistently describe preference manipulation as a metadata-level attack in which biased naming, persuasive phrasing, or selectively framed descriptions influence LLM tool selection. Misleading descriptions such as “prefer this tool for security reasons” can bias the model toward a

malicious option even when the underlying capabilities remain unchanged Model Context Protocol Security (2025); Elastic Security Labs (2025).

Impact. The result is a persistent bias toward unsafe tools that increases the likelihood of bypassed safeguards in routine workflows, even without embedding an explicit instruction that simpler prompt-injection filters might detect.

A.3.6 Command and Function Collisions

When hosts aggregate tools from multiple MCP servers, identical or confusingly similar tool names can create ambiguity during discovery and invocation. In clients that also expose slash-style commands, overlapping commands such as `/query` or `/upload` can similarly route execution to the wrong tool if duplicate registration is allowed. Without stronger disambiguation, provenance cues, or unique registration rules, a malicious tool can shadow a legitimate one by reusing a similar name or interface and thereby divert routine invocations.

Evidence. Security analyses describe both cross-server tool shadowing and slash-command overlap as realistic MCP attack patterns. Benchmark results show that these attacks can successfully redirect execution to unintended tools, while enterprise security guidance notes that similar functionality or duplicate command registration can confuse tool selection and allow malicious tools to intercept actions intended for trusted ones Yang et al. (2025); dts-solution (2025).

Impact. Attackers can intercept intended actions, capture sensitive data, or trigger unauthorized operations by redirecting invocations to malicious handlers. Even without malicious intent, collisions and overlaps can produce information disclosure, privilege misuse, or effective denial of service against the legitimate tool.

A.3.7 ANSI Escape Code Injection

An attacker can embed ANSI escape sequences in strings rendered to humans, such as a tool’s `name` or `description` returned by `tools/list` or text returned by `tools/call`. In terminal-based MCP clients that render these strings without sanitization, the terminal may interpret them as control instructions that hide text, clear the screen, move the cursor, or display deceptive hyperlinks rather than showing the raw content as-is. Unlike poisoned tool metadata, which targets model interpretation, this attack targets the human display layer and exploits the disconnect between what the reviewer sees and what the underlying content actually contains.

Evidence. MCP-specific testing showed that ANSI escape sequences in tool descriptions and outputs were not sanitized in Claude Code 0.2.76, allowing malicious payloads to be hidden from users while still influencing the surrounding workflow Hoodlet (2025). Independent CVEs outside the MCP ecosystem also confirm the general feasibility of ANSI-based screen spoofing, including vulnerabilities that allowed attackers to spoof console output and even create fake file listings or hide genuine content through escape-sequence injection nvd (2024a;b); Dushantha (2024).

Impact. Staged warnings, masked prior lines, and misleading screens raise the chance of unsafe approvals or missed indicators during tool or log review, allowing attackers to obtain human authorization for actions that the operator would not approve if the display were accurate.

A.4 Prompt, Input, and Context Injection Attacks

This class of attacks targets how MCP agents assemble and reuse context at runtime. An agent’s working context is constructed from many channels including user prompts, server-supplied templates, retrieved content, local files, API responses, and intermediate tool outputs, and any of these channels can carry hidden instructions. The four threats below represent the primary failure modes through which insecure context assembly is exploited.

A.4.1 Poisoned Prompt Templates

MCP servers can expose reusable prompt templates that clients discover through `prompts/list` and retrieve through `prompts/get`. Because these prompts are provided by the server as structured templates intended

to shape model interactions, a compromised server can embed hidden instructions that are presented as trusted guidance rather than ordinary untrusted user input Model Context Protocol (2025m).

Evidence. Public security analysis notes that a malicious server can supply a subtly modified prompt designed to trick the application into revealing sensitive information or misusing tools. Because such prompts are delivered as reusable templates from the server, a single compromised or hijacked server can propagate the same malicious guidance across every host that imports and invokes them Palo Alto Networks (2025).

Impact. MCP prompt templates can become high-leverage inputs whose compromise steers assistants toward information leakage, credential exposure, or unsafe tool use across multiple hosts and workflows that rely on the affected server.

A.4.2 Retrieval-Based Poisoning (RADE)

Adversaries can embed malicious MCP instructions inside files or online content that retrieval tools ingest, so the LLM later treats them as benign context and follows them as part of an otherwise ordinary query. Retrieval-Agent DEception (RADE) shows that when a crafted file is added to a vector database, a model can retrieve attacker-written instructions and then use connected MCP tools to add SSH keys, search for secrets, or exfiltrate credentials. Follow-on work further broadens this threat model by showing that attackers may only need to post malicious content online if MCP-enabled scraping and retrieval pipelines ingest it automatically Radosevich & Halloran (2025).

Evidence. In the original RADE demonstrations, Claude used the Chroma MCP server to retrieve malicious content from a vector database and then used filesystem and Slack-related tools to add attacker SSH keys and exfiltrate API keys. Subsequent TRADE experiments showed that similar retrieval-based attacks can succeed even when the payload is hosted online rather than downloaded directly, and that existing refusal guardrails often remain weak against these falsely benign retrieval exploits Halloran et al. (2025).

Impact. A single compromised document, indexed file, or ingested web page can escalate into system-level data exfiltration or unauthorized access across workflows that connect retrieval servers to high-privilege tools, with the payload executing as a natural consequence of normal information retrieval.

A.4.3 Transitive-Trust Poisoning (Fourth-Party Injection)

Transitive trust poisoning arises when a trusted MCP server retrieves content from external systems and forwards it into the model context as trusted contextual input. If that upstream content contains hidden instructions, the server can become an unwitting relay that launders attacker-controlled text across a trust boundary the model treats as higher-confidence than ordinary user input. This risk remains even when deployments add partial mitigations, because retrieved content may still influence downstream model behavior if sanitization or review is incomplete Supabase.

Evidence. In a publicly discussed Supabase MCP scenario integrated with Cursor, the assistant operated with `service_role` privileges while reading customer-submitted ticket content. A single crafted support ticket could induce the agent to read a sensitive table such as `integration_tokens` and write the results back into the same ticket, turning retrieved upstream content into an exfiltration path Willison (2025a).

Impact. A single poisoned upstream message can become a channel for sensitive-data exfiltration and unsafe tool use in workflows that rely on the same trusted retrieval path, even when the MCP infrastructure itself has not been directly compromised.

A.4.4 Format-Embedded Poisoning (Weaponized Content Channels)

Format-embedded poisoning arises when MCP clients surface retrieved documents, web content, chat messages, or tool outputs directly into the model context, allowing attackers to hide model-interpretable instructions inside content that appears to be ordinary data. The attack relies on the fact that once such content is passed through a trusted retrieval or messaging path, the model may interpret embedded instructions as part of the task context rather than as inert payload.

Evidence. Public MCP guidance warns that servers fetching untrusted content expose clients to prompt-injection risk Anthropic (2025). In MCP-specific demonstrations, a crafted file indexed through the Chroma MCP server caused Claude to retrieve attacker-controlled instructions and then use other connected tools to add SSH keys and exfiltrate API keys Radosevich & Halloran (2025). Separate experiments showed that a single injected WhatsApp message returned through a trusted MCP integration could compromise the agent and exfiltrate sensitive contact information, even without installing a separate malicious server Beurer-Kellner & Fischer (2025).

Impact. A poisoned document, message, or retrieved page can become a portable attack payload across workflows that ingest the same content channel, turning ordinary retrieval into a path for data exfiltration and unsafe follow-on tool use.

A.5 Credential Leakage and Data Exposure

Where the preceding categories focus on trust establishment, code integrity, and unsafe context assembly, this section turns to what MCP workflows can unintentionally reveal once those mechanisms begin operating across real services. MCP servers often function as high-privilege intermediaries that hold tokens, relay content, and combine outputs from multiple tools, so exposure can arise not only from direct secret theft but also from logging, forwarding, retention, and cross-system aggregation. The threats collected here therefore capture the main ways credentials and sensitive data escape their intended scope in MCP deployments.

A.5.1 Centralized Credential Storage

MCP servers often consolidate API keys, OAuth tokens, and other service credentials in a single configuration file or secret store, so compromising that host or backing credential store can give an attacker valid bearer tokens for every connected integration. The centralized nature of the failure means that a single breach can cascade across multiple integrations rather than requiring separate compromise of each downstream service.

Evidence. Public security analysis of MCP deployments describes servers that retain long-lived OAuth tokens for email and enterprise services, and warns that a breached server can reuse those tokens to perform arbitrary tool actions and API calls that appear indistinguishable from legitimate agent activity. The same analysis also notes that MCP servers are high-value targets because they often store authentication tokens for multiple services at once Pillar Security (2025).

Impact. Once attackers control these credentials, downstream APIs may treat their traffic as ordinary MCP usage, allowing them to exfiltrate data, send or delete messages, and invoke administrative operations across multiple connected services until credentials are rotated or revoked, with blast radius proportional to the number of services consolidated behind that store.

A.5.2 Insecure Secret Management

In MCP deployments, secrets such as API keys and OAuth tokens can leak through operational choices such as hard-coded configuration files or other insecure credential-handling practices. Because MCP servers often mediate access to multiple external services at once, a single credential exposure can grant attackers broad access across connected integrations.

Evidence. Public security analysis of MCP describes how servers often store API keys and may retain access to multiple external services simultaneously, which magnifies the damage of a single leak. The same guidance recommends secret scanning for configuration files and using environment variables or dedicated secret-management solutions instead of hard-coded credentials Palo Alto Networks (2025).

Impact. Once exposed, these bearer tokens can allow attackers to issue legitimate-looking calls to downstream services on behalf of the MCP host or agent until revocation and rotation fully propagate, with blast radius increasing alongside the number of services consolidated behind the same credential store.

A.5.3 Overprivileged and Persistent Tokens

Tools and servers may request broad permission scopes for external services and rely on OAuth tokens that can remain valid beyond a single interactive session. Even when secrets are not directly exposed, the breadth of the granted scope means that a single stolen token may authorize actions across multiple connected services simultaneously.

Evidence. Public security analysis of MCP deployments describes servers that store OAuth tokens for connected services such as email, with a compromised server able to reuse those tokens to issue tool actions and API calls that are indistinguishable from legitimate agent behavior Pillar Security (2025). Separate OAuth incident research shows that malicious applications can retain cloud access even after password resets or multifactor authentication changes because the authorized application remains valid until explicitly revoked Proofpoint Threat Research (2025).

Impact. A compromised overprivileged token becomes a long-lived capability that lets an adversary operate across multiple tools and workflows on behalf of the MCP host or agent until rotation and revocation fully complete, with the scope of the token determining the full extent of actions the attacker can perform without triggering an additional authorization step.

A.5.4 Cross-Tool Data Flows

In complex MCP agent pipelines, a single hidden directive can steer data from internal tools, retrieved resources, or conversation history toward external endpoints within one multi-step interaction. The base MCP specification standardizes how hosts, clients, and servers exchange tools, resources, prompts, and sampling requests, but leaves security controls to implementors rather than defining protocol-level provenance tracking or data-flow restrictions across chained tool use.

Evidence. In a documented indirect prompt injection scenario, an assistant retrieved a compromised webpage and was then induced to URL-encode the user’s conversation history and send it as a query parameter to an attacker-controlled domain Palo Alto Networks Unit 42 (2025). Separate MCP-specific threat analysis shows a composability-chaining pattern in which a seemingly benign MCP server forwards a tool request to a second hidden server, merges the returned malicious instructions into the response, and causes the model to exfiltrate sensitive environment variables within the same session CyberArk Threat Research (2025).

Impact. Cross-tool data flows can silently move private conversations, local files, and environment-scoped secrets from internal agents and hosts to attacker endpoints. Because each individual tool call may still look legitimate in isolation, per-call monitoring alone may miss the larger exfiltration sequence.

A.5.5 Trojanized MCP Servers

A compromised or malicious MCP server can observe and modify tool traffic, including user queries, tool parameters, file contents, and credential-bearing responses, while continuing to present itself as a legitimate endpoint. Because agent workflows may rely on server responses as contextual input for subsequent reasoning and tool use, a trojanized server can occupy a central position in the agent’s decision loop and influence downstream behavior across the workflow.

Evidence. Public security analysis of MCP deployments describes servers as high-value targets because they may store OAuth tokens for connected services, allowing a breached server to reuse those tokens for arbitrary tool actions and API calls that appear legitimate to downstream services Pillar Security (2025). Separate MCP security guidance notes that a rogue server can mimic or modify the responses of genuine services, reroute messages, or tamper with outputs while the agent continues to treat the server as a trusted integration Muayad Ali (2025).

Impact. A single trojanized server can become both a credential-theft pivot and a workflow-poisoning point, enabling cascading data leakage and unauthorized actions across the workflows that depend on that server.

A.5.6 Emergent Aggregation Risks

In MCP deployments that connect multiple servers at once, agents can accumulate low-sensitivity fragments such as calendar entries, file names, and chat metadata into a shared model context even when each individual tool call appears benign. Because MCP enables a single assistant to access multiple external tools and data sources through one workflow, outputs from email, file, database, and calendar servers can be correlated in the same reasoning loop. The protocol leaves privacy and access-control safeguards to implementors rather than defining protocol-level limits on how such cross-source outputs may be combined.

Evidence. Public MCP security guidance identifies data aggregation as a distinct risk, noting that centralized access to multiple services enables correlation across sources and can expose sensitive patterns that are not visible in any one result alone Pillar Security (2025); Behl (2025). More generally, work on the mosaic effect shows that records or metadata that pose little risk in isolation can reveal identities, routines, or other sensitive attributes once aggregated OSTC (2025). Figure 10 illustrates this directly, where aggregating low-risk outputs from email, calendar, and file servers exposes sensitive project details that no single call would reveal.

Impact. Because each invocation may reveal only a small fragment, emergent leaks are difficult to detect through per-call review even when the accumulated context becomes sensitive. An attacker or over-privileged agent that can observe partial results from several tools may still reconstruct detailed organizational maps, schedules, or user profiles from the combined context.

A.5.7 Residual Data Exposure

Within MCP deployments, tool inputs, outputs, or credentials may persist in host-side storage artifacts such as logs, caches, temporary files, or crash dumps after a session ends. If these artifacts are not explicitly sanitized, sensitive values intended to be ephemeral can remain recoverable on disk long after the original interaction has completed.

Evidence. Guidance on media sanitization explains that ordinary deletion or reformatting may leave residual data that can still be reconstructed unless explicit clearing or purging is applied Kissel et al. (2014). Operating-system guidance for encryption tools similarly warns that crash dump files can contain cached passwords, encryption keys, and the contents of sensitive files from RAM VeraCrypt Project (2015).

Impact. An attacker who later gains access to the host or its retained storage artifacts may be able to recover prior tool inputs, outputs, or credentials that were expected to be transient, even without access to the original MCP session.

A.6 Authorization and Access Control Gaps

Authorization failures in MCP rarely arise from the absence of credentials alone. More often, they stem from permissions that are too coarse, too static, or too weakly revalidated as workflows evolve. Even when tokens, approvals, and annotations are present, enforcement can still break down when policy remains binary, contextual risk is ignored, consent becomes routine, or different servers apply incompatible access rules. Taken together, the threats in this section show how least privilege erodes in practice as MCP deployments become longer-lived, more automated, and more distributed.

A.6.1 Binary Permissions

In many MCP deployments, authorization remains coarse-grained, and a bearer token or equivalent approval can grant broad access to all tools exposed by a server without finer-grained controls over which specific actions may be invoked. This all-or-nothing model leaves little room for intermediate permission levels such as allowing read access while blocking writes or administrative operations.

Evidence. Practical hardening guidance describes traditional MCP implementations as relying on static API tokens that provide binary, unrestricted access to all available tools, violating least privilege Nikhil (2025). Enterprise security work correspondingly argues for fine-grained, scoped access tokens, per-request

authorization, and role- or context-aware controls, indicating that these protections are not provided by baseline MCP deployments by default Narajala & Habler (2025).

Impact. An attacker who compromises an overprivileged tool token can gain broad control over data access, state changes, and destructive operations reachable through the approved tool surface. A single leaked credential may therefore be sufficient to exercise the full capability set exposed to that token until revocation, rotation, or additional authorization checks take effect.

A.6.2 Lack of Risk-Aware Differentiation

MCP does not define a mandatory mechanism for separating high-sensitivity tool calls from low-risk ones. The specification treats `ToolAnnotations` such as `readOnlyHint` and `destructiveHint` as advisory meta-data rather than enforceable security controls, and the documentation warns that clients should never make security-critical decisions based solely on these annotations Anthropic (2025c).

Evidence. Recent MCP security frameworks argue that descriptive metadata is insufficient on its own and recommend Zero Trust style controls such as fine-grained, scoped access tokens, just-in-time access, context-aware access decisions, and per-request authorization to gate each interaction Narajala & Habler (2025).

Impact. In deployments that do not add such extra-protocol safeguards, destructive actions may be exposed through the same tool-selection and invocation flow as benign read-only operations, increasing the chance that a misrouted call, compromised agent flow, or weak approval path leads to high-impact changes Anthropic (2025c); Narajala & Habler (2025).

A.6.3 Consent Fatigue and Automation Drift

Repeated permission prompts for MCP tools can desensitize users over time, causing them to click through approvals or skim long, complex requests instead of carefully reviewing what a tool is about to do. In deployments where previously approved tools are allowed to run without renewed scrutiny, the intended human-in-the-loop safeguard can gradually weaken and become easier for attackers to bypass.

Evidence. One published MCP threat analysis describes a sampling-based attack in which a malicious server hides exfiltration instructions inside a long, innocent-looking story prompt and relies on a user who no longer reads the request or response carefully. The same analysis assumes that a shell-capable tool had already been approved for constant use under an “Always allow” setting, allowing the attacker to bypass meaningful review and extract environment variables through the generated output CyberArk Threat Research (2025).

Impact. Sensitive actions such as exfiltrating environment variables or invoking shell-capable tools can occur without fresh user scrutiny. In workflows that reuse stale approvals or permissive trust settings, the human oversight layer that MCP depends on for safety can be substantially weakened.

A.6.4 Confused Deputy and Role Leakage

In multi-server or delegated MCP deployments, an agent or intermediary server with higher privileges can be induced to execute actions on behalf of a less-privileged requester when user identity, authorization scope, or token audience are not tightly bound to each request. Where tool handoffs occur frequently and automatically, each improperly scoped delegation can become a path for privilege escalation.

Evidence. The MCP authorization specification explicitly warns that MCP servers acting as intermediaries to third-party APIs can be abused through token passthrough and failures of token audience validation, creating a confused-deputy risk when downstream actions are performed under the wrong authority MCP Working Group (2025b). Separate empirical audits of MCP-enabled agents show that large language models can be induced to chain tools and servers to perform remote access, credential theft, and other unauthorized actions within a single workflow Radosevich & Halloran (2025).

Impact. Coarse session-wide grants, shared tool contexts, and missing per-request scoping can allow nominally low-sensitivity requests to transitively exercise higher-privilege capabilities under the agent’s or server’s

effective identity, enabling silent privilege escalation, unauthorized data access, and side effects that are difficult to attribute from individual tool calls alone.

A.6.5 Decentralized and Inconsistent Policy Enforcement

Each MCP server in a deployment can independently define and enforce its own OAuth settings, tool allow-lists, logging rules, and update policies. Although an MCP host may enforce local security policies for the clients it manages, the protocol does not standardize a fleet-wide mechanism for keeping these controls aligned across independently deployed servers. As a result, a vulnerability discovered and patched on one server may remain exploitable on others if policy and update state drift across the deployment.

Evidence. Official MCP architecture documentation describes hosts as managing separate client connections to individual servers, while broader ecosystem overviews note that decentralized MCP environments introduce security variability and version-management challenges Model Context Protocol (2025f); Emergent Mind (2025). Enterprise gateway guidance correspondingly frames centralized enforcement as an added control layer for achieving consistent authentication and security policy across multiple MCP servers, implying that such consistency is not provided by baseline MCP deployments by default Field, Michael (2025).

Impact. Two servers that expose the same nominal tool may still enforce materially different protections, so compromising the weakest instance can give an adversary capabilities that users incorrectly assume are governed by a uniform security baseline across the broader deployment.

A.6.6 No Context-Aware Revalidation

Clients obtain OAuth-based access tokens for MCP servers and then reuse those tokens for subsequent tool and resource calls. The MCP authorization specification standardizes token discovery, issuance, validation, and runtime scope step-up for HTTP-based transports, but it does not define an intent-aware or continuous authorization mechanism that rebinds each call to changing workflow context after a valid bearer token has been issued MCP Working Group (2025a).

Evidence. Recent work on secure delegation for autonomous agents argues that bearer-style OAuth and JWT tokens satisfy identity checks at issuance time but remain possession-based thereafter, so any party that holds the token can continue issuing API calls until expiry. That analysis further notes that resource servers cannot distinguish authorized intent from unauthorized execution unless tokens are augmented with intent, delegation, or per-agent binding information Goswami (2025).

Impact. A token granted in a low-risk or tightly supervised context can continue to authorize sensitive MCP tool invocations under materially different runtime conditions, yielding durable access that is difficult to constrain once the workflow, agent state, or surrounding trust context has changed.

A.6.7 Uneven RBAC Support

MCP leaves detailed authorization design to individual servers. The authorization specification standardizes transport-level OAuth for HTTP-based transports, but it does not define shared scopes or an interoperable RBAC model across servers, so a role or permission enforced on one server need not have any corresponding meaning on another Model Context Protocol (2025l); Zhou (2025).

Evidence. Practitioner guidance shows that MCP servers can implement authorization using patterns such as RBAC with viewer, editor, and admin roles, ABAC based on token claims and resource attributes, or external policy engines, but presents these as implementation choices that server operators must assemble themselves Speakeasy (2025). This server-by-server flexibility can produce fragmented authorization and make least-privilege reasoning harder across multi-server deployments Zhou (2025).

Impact. Fragmented and inconsistent authorization makes it difficult to understand and audit an agent’s effective capabilities across tools and hosts, and an agent operating across multiple servers may have materially different permissions on each one without any operator-visible summary of the combined capability surface.

A.6.8 Orphaned Sessions and Credentials

Hosts and servers may shut down, restart, or be decommissioned while leaving access tokens, API keys, or other session identifiers valid if those credentials are not explicitly rotated, expired, or invalidated during teardown. In environments where MCP services are frequently restarted or redeployed, retained credentials can outlive the process or session that originally obtained them and remain usable until their normal expiry or explicit revocation.

Evidence. The MCP authorization specification warns that attackers who obtain tokens stored by the client, or cached or logged on the server, can access protected resources with requests that appear legitimate to the resource server, and it recommends short-lived access tokens to reduce the impact of leaked credentials Model Context Protocol (2025l). Practical MCP hardening guidance similarly recommends continuous validation of short-lived tokens and scope-based authorization at every interaction point, reflecting the risk of persistent access when old credentials are not actively invalidated HiveTrail (2025).

Impact. Any attacker or previously authorized client that retains those credentials may continue calling MCP servers and downstream services after restart or teardown until expiry, rotation, or revocation takes effect, turning a temporary deployment lifecycle event into a longer-lived access path.

A.7 Execution Environment and Sandbox Risks

Once an MCP workflow crosses from reasoning into execution, the security question is no longer only what the model selects, but what the runtime is actually allowed to do. Tools that execute with broad filesystem access, unrestricted network reach, shared infrastructure, or weak isolation can turn prompt-level mistakes and server-side compromise into host takeover, lateral movement, or resource exhaustion. The threats in this section therefore center on the containment boundary itself, including how execution is isolated, what privileges it inherits, and how far a failure can propagate beyond the initial tool invocation.

A.7.1 Lack of Execution Isolation

In the `stdio` transport, the client launches the MCP server as a subprocess on the same machine and communicates with it over standard input and output. MCP does not require sandboxing or resource confinement for such local servers, so absent additional controls they execute with the privileges of the client process rather than inside a mandatory containment boundary.

Evidence. Public security guidance notes that local MCP servers can execute commands with client privileges and typically run without strong isolation unless operators add containerization, sandboxing, or similar host-level controls Red Hat Product Security (2025).

Impact. Compromise of an unsandboxed local server can yield host-level consequences, including arbitrary command execution, sensitive file access, data exfiltration, and persistent effects on workflows or agents that depend on that server.

A.7.2 Sandbox Escape Vulnerabilities

Even when MCP servers run inside containers or virtual machines, they still depend on the underlying container runtime or hypervisor for isolation. A vulnerability or escape in that lower layer can therefore let a compromised server interact with host resources outside its intended sandbox, and the severity of the failure grows when multiple workflows share the same host or runtime.

Evidence. Analyses of container-escape vulnerabilities such as CVE-2019-5736 in `runc` show that crafted workloads can overwrite host binaries and obtain code execution on the host National Vulnerability Database (2019). Separately, MCP safety audits show that once tools can reach the host environment, LLM-driven agents can be induced to inject malicious commands into system files and exfiltrate API keys or other secrets from environment variables and configuration files Radosevich & Halloran (2025).

Impact. A single server compromise can escalate into a host-level breach that enables persistence and affects multiple MCP workflows, and every workflow sharing the same container runtime inherits the blast radius of the escape.

A.7.3 Root or Elevated Execution

Running tools and helper processes with root or Administrator privileges removes the operating system privilege separation that would otherwise limit the impact of a compromise. As a result, any successful exploit or prompt-driven misuse of a tool can escalate directly into host-level control without an intermediate containment boundary.

Evidence. Official MCP security guidance warns that local MCP servers run with the same privileges as the client and recommends sandboxing them with minimal default privileges Model Context Protocol (2025o). Public advisories for CVE-2025-49596 further show that MCP Inspector versions below 0.14.1 allowed unauthenticated requests to launch MCP commands over `stdio`, resulting in remote code execution in the context of the Inspector process National Vulnerability Database (2025a).

Impact. If the affected process runs with elevated privileges, a single exploit can yield broad host-level control, including arbitrary command execution, access to sensitive files and configuration, and the ability to alter local artifacts in ways that hinder detection and recovery.

A.7.4 Prompt-Based Remote Code Execution

An agent that can invoke shell or scripting tools can be driven by prompt injection to issue arbitrary operating system commands, because the model may treat injected tool requests as actionable instructions rather than as untrusted input. When shell-capable MCP tools are available, a single successful prompt injection can therefore translate directly into command execution on the host.

Evidence. An MCP safety study demonstrates a proof-of-concept attack in which an injected prompt causes the agent to append a netcat command to the user's shell run configuration file, creating a reverse shell that connects back to an attacker-controlled host whenever a new terminal is opened Radosevich & Halloran (2025).

Impact. A single successful prompt injection into an agent with shell access can enable persistent backdoor installation, credential theft from local files and environment variables, and further compromise of workflows that rely on the affected MCP tooling.

A.7.5 DoS and Resource Exhaustion

Although the MCP specification requires servers to rate limit tool invocations, it does not define a shared quota model, default thresholds, or a fleet-wide enforcement mechanism across deployments. A misconfigured or malicious agent can therefore still repeatedly write large files, flood external APIs, or spawn CPU-heavy tasks until host resources or downstream services degrade or fail if server-side limits are weak or inconsistently implemented. In shared MCP deployments, resource exhaustion triggered by one agent can degrade availability for other agents and workflows that depend on the same server or infrastructure.

Evidence. The MCP tools specification requires servers to rate limit tool invocations, while operational hardening guidance recommends implementing quotas and rate limits per user, agent, or session together with timeouts and throttling for costly operations Model Context Protocol (2025p); Sadeddin (2025).

Impact. A single faulty workflow or compromised agent can still exhaust CPU, memory, or third-party quotas and cause denial of service for other workloads when limits are weak or inconsistently enforced. Because quota design and operational thresholds are left to implementers, these protections must still be configured and maintained by each MCP server operator.

A.7.6 Vulnerable Third-Party Dependencies

MCP servers and tool plug-ins often depend on third-party parsers, HTTP clients, and SDKs, so outdated or unverified dependencies can expand the server’s attack surface and expose ordinary tool workflows to known software-supply-chain weaknesses. MCP security guidance accordingly recommends software composition analysis, dependency verification, and hardened build pipelines because vulnerable dependencies become part of the effective attack surface of the server itself Red Hat Product Security (2025).

Evidence. Widely deployed parsing libraries have repeatedly shown that attacker-controlled content can trigger serious compromise when unsafe parsing paths are enabled. For example, SnakeYAML’s `Constructor()` allowed attacker-supplied YAML deserialization to lead to remote code execution, and PyYAML versions before 5.3.1 and 5.4 were likewise vulnerable to arbitrary code execution when processing untrusted YAML through unsafe loading paths National Vulnerability Database (2022; 2020b;a).

Impact. A single vulnerable dependency can turn an otherwise routine fetch-and-parse operation into remote code execution on the MCP server host, and because these flaws may reside in transitive dependencies several layers below the application code, the resulting compromise can affect every workflow that relies on the affected tool.

A.7.7 Lateral Movement via Shared Infrastructure

In shared MCP deployments, multiple servers may run on the same hosts or virtual networks and may centralize OAuth tokens or API credentials for connected services. If one server or its credential store is compromised, an attacker can reuse those tokens and existing network reach to access other connected services or pivot through the same trust boundary Pillar Security (2025).

Evidence. Public security analysis describes MCP servers as high-value targets because they often store authentication tokens for multiple services at once. Official MCP security guidance likewise warns that if a token is accepted by multiple services without proper validation, an attacker compromising one service can use that token to access other connected services Model Context Protocol (2025o).

Impact. A single server compromise in a shared deployment can expand into lateral access across connected services, enabling unauthorized tool use, token reuse, and data exfiltration well beyond the originally breached integration.

A.8 Workflow and Configuration Risks

Even when tools are trusted and execution is contained, MCP workflows remain vulnerable if operational state stops being consistent across steps, sessions, and servers. Small mismatches in context, configuration, or teardown can accumulate until the workflow that actually runs is no longer the one users and operators intended. This section focuses on the failures of continuity and control that make that divergence possible.

A.8.1 Context Drift and State Inconsistency

Multi-step MCP workflows rely on consistent propagation and update of shared state across tools and servers, so when an early step is delayed, lost, corrupted, or applied inconsistently, later steps may continue from an outdated or conflicting view of the world. Small local inconsistencies can therefore accumulate into broader workflow failures that may only become visible after downstream tools have already acted on stale context.

Evidence. Work on context drift in LLMs shows that model behavior can diverge over time as surrounding context departs from a stable, goal-consistent reference Wu et al. (2025); Dongre et al. (2025). Related work on multi-agent planning frameworks argues that transaction guarantees, validation, and context management are necessary because current systems struggle with inconsistent shared state, context loss, and failures that require later repair Chang & Geng (2025). MCP-specific operational guidance likewise warns that lost, delayed, or corrupted updates can cause agents to operate on conflicting information MacVittie (2025).

Impact. In MCP deployments that compose models into long-running workflows, inconsistencies and contextual drift can accumulate across steps and agents, causing downstream tools to act on outdated or contradictory information and turning operational failures into broader integrity and reliability problems.

A.8.2 Configuration Drift Across Servers

Enterprise MCP deployments can accumulate many server instances with small divergences in authentication settings, logging, allowed scopes, and dependency versions, so uneven maintenance and decentralized administration may leave some nodes substantially weaker than others. As a result, a vulnerability mitigated or remediated on one server may remain exploitable on another if security posture drifts across the fleet.

Evidence. Practitioner risk assessments report that more than 13,000 MCP servers were publicly available on GitHub by May 2025 and characterize the ecosystem as frequently unauthenticated and lightly supervised Prompt Security Team (2025a). Complementary large-scale analysis of 1,899 open-source MCP servers finds that 7.2% contain general vulnerabilities and 5.5% exhibit MCP-specific tool poisoning, indicating that materially different security postures persist across publicly available MCP servers Hasan et al. (2025).

Impact. Once even a single weak or poisoned server remains inside an agent’s trust boundary, an attacker can use that outlier as an entry point to harvest secrets, misroute tool calls, or compromise workflows that implicitly treat servers across the deployment as equivalently trustworthy.

A.8.3 Transactional Integrity Gaps

HTTP-based MCP interactions send each client JSON-RPC message as a separate HTTP POST request, and although the transport can support sessions and streaming, the protocol does not define atomic multi-step transaction semantics or built-in rollback across related calls Knostic Team (2025). As a result, recovery from partial failure is left to higher-level workflow logic rather than guaranteed by the protocol itself.

Evidence. Practitioner guidance places rollback for suspicious tool-call sequences at the workflow layer rather than the protocol layer Network Intelligence (2026). Related work on multi-agent LLM planning likewise argues that current systems lack transactional safeguards and can enter inconsistent states or partial-failure conditions unless additional compensation, validation, and recovery logic is added above the base interaction layer Chang & Geng (2025).

Impact. In MCP-based enterprise automations that span multiple services, a misaligned or aborted run can leave one subsystem updated while others remain unchanged, creating integrity and policy violations that may only be detected after the workflow has completed.

A.8.4 Patch Gaps and Vulnerable Instances

MCP servers often reuse reference implementations, connectors, and community-maintained packages, so when vulnerabilities appear in widely copied code, unmaintained forks can quietly retain exploitable behavior across registries and internal deployments. The same forking model that makes MCP tools easy to customize and redistribute also increases the chance that insecure code will persist in downstream copies long after the original project is archived or no longer maintained.

Evidence. Recent reporting describes a SQL injection flaw in Anthropic’s SQLite MCP server that had been forked or copied more than 5,000 times before the repository was archived, with no official patch planned and downstream maintainers left to audit and fix their own copies Park (2025); Lyons (2025).

Impact. A single vulnerable MCP server can leave many downstream agents and workflows exposed to stored prompt injection, data exfiltration, and workflow hijacking even after the original project has been deprecated. Because remediation depends on each fork maintainer acting independently, vulnerable copies can persist across the ecosystem long after the original disclosure.

A.8.5 Session Isolation in Multi-Agent Environments

Shared MCP servers that host multiple agents or tenants can implement sessions with weak boundaries, so an attacker who can guess, replay, or intercept a session identifier may be able to impersonate another client or cause state intended for one user or agent to be delivered into a different active session. In multi-agent systems where context is handed off automatically across components, a single hijacked session can therefore contaminate downstream agent behavior if later steps trust the injected or misrouted state.

Evidence. MCP security guidance explicitly describes both session-hijack impersonation and session-hijack prompt-injection scenarios in stateful HTTP deployments. It warns that servers must not use sessions for authentication, must verify inbound requests, should bind session identifiers to user-specific information, and must not associate state with session identifiers alone, precisely because weak session handling can enable unauthorized actions and cross-user state confusion Model Context Protocol (2025j).

Impact. A compromised session can expose private tools, prompts, and retrieved state across workflows that share the affected server, leading to unauthorized actions and disclosure of user-specific context well beyond the original interaction.

A.8.6 Unsafe Fallbacks and Error Handling

When MCP tool calls fail, servers can return structured application-level error payloads inside otherwise successful `tools/call` results, and these payloads may be visible to the model and influence its next step Anthropic (2025c).

Evidence. Public guidance on MCP error handling recommends returning context-rich but sanitized error messages because models may use them to recover, retry, or request human intervention Barthelet (2025).

Impact. If client-visible error messages include internal URLs, stack traces, credentials, or other diagnostic details, an ordinary tool failure can become an information-disclosure or workflow-manipulation path that exposes sensitive system information to the model or end user and steers the agent toward unsafe recovery behavior Hora (2025).

A.8.7 Operational Visibility and Logging Gaps

Effective forensics in MCP deployments depends on consistent, high-granularity logging of tool invocations, context changes, and access events, yet many real-world implementations either log too little or do so inconsistently. These telemetry gaps make it difficult to reconstruct agent behavior or determine which tools ran and what data was accessed.

Evidence. Monitoring guidance for MCP deployments reports that the vast majority of implementations lack sufficient telemetry to audit what users and agents are doing Sakib (2025). Enterprise security work likewise identifies *Insufficient Auditability* as a distinct MCP risk, defined by inadequate logging that constrains detection and investigation of security events Narajala & Habler (2025).

Impact. Logging gaps leave operators unable to reliably reconstruct agent behavior or verify what tools ran and what data was accessed across workflows. In the absence of adequate telemetry, incident response can degrade into post hoc guesswork about what the agent actually did during the affected period.

A.8.8 Incomplete Tool De-Registration

Outdated MCP tools or endpoints that are not properly removed during synchronization can continue to appear in indexes, configuration stores, or orchestration layers after operators believe they have been retired. In large deployments where tools are frequently added, updated, and deleted, the gap between the server’s current tool list and downstream cached or indexed state can grow over time if removal is not handled correctly.

Evidence. Research on MCP-aware tool synchronization treats the server tool list as the single source of truth and includes explicit create, update, and delete operations to keep downstream tool stores aligned Lumer et al. (2025). Production systems likewise expose dedicated removal operations for MCP

tools, including OpenSearch’s Remove MCP Tools API and TrustGraph’s `tg-delete-mcp-tool` command, indicating that stale tool records must be actively deleted rather than assumed to disappear automatically OpenSearch Project (2025); TrustGraph (2025).

Impact. If removal or synchronization steps are skipped or fail, stale tool records can linger in indexes and configuration groups, causing agents or orchestration layers to continue discovering or invoking tools that operators intended to retire. This creates a persistent mismatch between operational intent and effective capability exposure, increasing the risk of misrouting, outdated behavior, and avoidable attack surface.

A.9 Emergent Threats from Model Behavior and Human Factors

A.9.1 Model Inversion Attacks.

When multiple MCP agents rely on a shared model backend, an attacker may use one agent interface to issue carefully crafted queries that infer sensitive properties of data the model has learned during training or adaptation. In enterprise MCP deployments, this creates a shared privacy boundary in which the confidentiality of information exposed through one workflow depends not only on that workflow’s own controls, but also on whether other agent interfaces can be used to probe the same underlying model.

Evidence. Prior work on model inversion shows that black-box access to model outputs can, in some settings, reveal sensitive attributes or approximate information about the data used to train the model Fredrikson et al. (2015).

Impact. In MCP environments built around shared model backends, a malicious workflow could use ordinary inference access to extract information correlated with proprietary or tenant-sensitive data, even without compromising MCP servers, tool implementations, or the surrounding infrastructure.

A.9.2 Membership Inference.

In deployments where MCP agents expose a shared enterprise model trained or adapted on internal data, an attacker can probe one agent interface to test whether particular records, users, or documents were included in the model’s training corpus. Because the attack requires only query access to normal model outputs, it may be available to any user or external actor who can interact with the deployment through ordinary channels.

Evidence. Prior work on membership inference shows that black-box access to model outputs can, in some settings, distinguish training examples from non-members with significant reliability Shokri et al. (2017).

Impact. In MCP deployments built around shared enterprise models, attackers may be able to infer whether sensitive internal records contributed to training without compromising MCP servers or tools, enabling targeted follow-on attacks against individuals, datasets, or business processes linked to the confirmed training data.

A.9.3 Pre-MCP Data Poisoning and Backdoored Models.

MCP implementations typically treat the underlying model as a trusted component, so covert behaviors introduced during pre-training, fine-tuning, or instruction tuning may be inherited by every agent that relies on it. Because such behaviors are embedded in the model itself rather than in MCP infrastructure, they can persist across server restarts, configuration changes, and software updates until the affected model is replaced, retrained, or otherwise remediated.

Evidence. Prior work shows that poisoned training or instruction-tuning pipelines can implant behaviors that remain dormant under normal use but activate in response to specific prompts, triggers, or contextual patterns Xu et al. (2023). Broader safety analyses also note that data poisoning and hidden model-level behaviors can be difficult to detect or rule out once incorporated into deployed systems Anthropic (2024c).

Impact. In MCP deployments built around a shared compromised model, ordinary conversations or tool-mediated contexts may activate hidden behaviors that leak information, bypass intended safeguards, or

systematically bias downstream tool use across multiple agents and workflows. The attack surface therefore extends beyond any single MCP component to the full model training, adaptation, and deployment pipeline.

A.9.4 Malicious Agent Operators.

Because MCP hosts translate model outputs into tool invocations and enforce the surrounding configuration and permission boundaries, operators who connect unvetted servers or grant overly broad privileges can expand an agent’s effective execution authority beyond intended policy. In such cases, abuse may not require exploitation of any protocol flaw at all. Instead, it can arise from trusted administrative control over host configuration, server selection, and authorization settings.

Evidence. Recent analyses of MCP ecosystems describe hosts as components that translate LLM-generated outputs into external tool invocations and note weak or missing output-verification and trust-validation mechanisms in current deployments Li & Gao (2025); Guo et al. (2025). These trust assumptions mean that high-impact behavior can result not only from malicious servers or prompt manipulation, but also from unsafe operator choices about what servers to connect and what permissions to grant.

Impact. A malicious or careless operator can cause routine-looking workflows to execute high-privilege actions, expose sensitive data, or alter downstream behavior across multiple agents. Because the operator acts through legitimate configuration channels, the resulting abuse may be difficult to distinguish from authorized administrative activity, and downstream servers or connected services may observe only apparently valid requests.

A.9.5 Insider Threats and Shadow MCP Servers.

Ungoverned or unauthorized MCP servers may be introduced inside organizations without central oversight, creating shadow MCP instances that access internal resources while bypassing standard monitoring, authentication, and governance controls. Because MCP servers can often be deployed and connected quickly, the organizational attack surface may expand faster than security and compliance teams can inventory, review, or contain it.

Evidence. OWASP guidance defines shadow MCP servers as unapproved or unsupervised MCP deployments operating outside formal organizational governance and warns that they can bypass centralized authentication, monitoring, and data-governance controls while creating new unmonitored endpoints OWASP Foundation (2025b). Broader ecosystem analyses further show that MCP deployments rely heavily on host-server trust relationships and that weak vetting of server onboarding can expose users and organizations to malicious or unsafe server behavior Li & Gao (2025).

Impact. Each shadow server introduces an unreviewed trust boundary that may expose sensitive data, execute internal capabilities, or provide attackers with an unmonitored foothold into enterprise systems. Because these deployments sit outside normal governance and asset-tracking processes, their misuse or compromise may remain undiscovered until after substantial operational, security, or compliance harm has occurred.

A.9.6 Agent Interface Social Engineering.

The agent interface and host create an attack surface where adversaries can issue crafted prompts or interaction sequences that steer tool selection and execution. Because these attacks operate through seemingly legitimate interactions rather than low-level protocol exploits, they can remain effective even when core MCP mechanisms are implemented as intended.

Evidence. Prior work on MCP security treats the agent interface and host as a primary attack surface where adversaries can present seemingly legitimate requests or workflows that drive high-impact tool invocations and data exfiltration under the guise of normal usage Guo et al. (2025).

Impact. Attackers can manipulate agent behavior through ordinary interaction channels, causing unsafe tool use or disclosure without exploiting a protocol flaw in MCP itself. These risks therefore cannot be addressed

through protocol hardening alone and also require stronger validation, approval design, and interaction-level safeguards.

A.9.7 Agent Collusion.

Two or more compromised agents may coordinate to bypass access controls or exfiltrate data covertly, including through encoded or steganographic content in inter-agent messages. In multi-agent MCP deployments, such coordination can combine capabilities that no single agent would be authorized to exercise alone, while remaining difficult to detect when oversight is fragmented across individual agents or workflows.

Evidence. Prior work on secret collusion among generative AI agents formalizes the risk of covert coordination through steganographic communication and presents empirical results suggesting that such behaviors warrant continued monitoring as model capabilities improve Motwani et al. (2024). MCP further expands the surface for nested and server-initiated agentic interactions through Sampling, which allows servers to request LLM generations inside other MCP features when implementations expose that capability Model Context Protocol (2025n).

Impact. Collusion-based attacks can be difficult to identify because each agent’s local behavior may appear policy-compliant when inspected in isolation, with the actual violation emerging only when cross-agent information flow, coordination patterns, or hidden communication channels are reconstructed end to end.

A.9.8 Unverified Agent Handoffs.

When one agent or MCP server delegates work to another without verifying identity, audience, or effective scope, the handoff can create privilege-escalation or confused-deputy failures. In complex multi-agent topologies where delegation occurs frequently and automatically, each unverified handoff becomes a potential escalation path that may be difficult to enumerate, monitor, or constrain exhaustively.

Evidence. The MCP *Authorization* specification explicitly requires audience validation and forbids token passthrough to prevent confused-deputy failures in delegated interactions Model Context Protocol (2025l). Broader multi-agent engineering experience likewise shows that systems with multiple cooperating agents introduce additional coordination and handoff complexity, increasing the risk that trust boundaries will be enforced inconsistently in practice Anthropic (2025b).

Impact. An unverified handoff can allow a lower-privileged agent to trigger actions through a more privileged agent or server that fails to validate the caller’s authority. The resulting escalation mirrors the confused-deputy pattern described in A.6.4, but here the failure arises across agent or delegation boundaries rather than within a single server interaction.

A.9.9 Orchestration Drift.

Over time, loosely coupled MCP workflows may diverge from their intended behavior, accumulate shadow logic, or reintroduce deprecated behaviors as tools, prompts, and coordination patterns evolve. Such orchestration drift can be difficult to detect because the affected workflow may still produce plausible outputs under normal conditions, with the divergence becoming visible only when an edge case, deployment change, or adversarial input exposes the gap between the intended and actual execution logic.

Evidence. MCP provides observability and change-management primitives such as tool list change notifications and structured logging, which help implementations track evolving capabilities and runtime behavior Model Context Protocol (2025p). Field reports and engineering experience further show that MCP and multi-agent deployments can accumulate hard-to-trace workflow logic, exhibit emergent behavior under small prompt or orchestration changes, and require careful coordination during deployment to avoid breaking active agents or silently altering behavior Backslash Security (2025); Anthropic (2025b).

Impact. Drifted workflows may silently grant capabilities beyond what operators intended, re-enable previously closed attack surfaces, or degrade policy enforcement in ways that remain invisible during routine

use. Because the divergence often accumulates gradually across prompts, tools, and orchestration logic, point-in-time reviews of configuration alone may fail to capture the actual behavior of the live system.