



# Model Context Protocol (MCP) Security and Tenancy Boundaries

Rohit Reddy Gaddam  
Sr. Site Reliability Engineer.

**Abstract:** The Model Context Protocol (MCP) is the mainframe that can control a multitude of contextual interactions in multi-tenant AI as well as cloud scenarios. In such conditions, users or organizations may use the same computational and data resources. AI systems that run on common infrastructures are calling for the securing of contextual data exchanges between tenants and models as a very serious issue. By introducing standard channels, MCP solves this problem for models in order to have access to, comprehend, and use tenant-specific contexts without violating isolation boundaries. Nevertheless, keeping contextual cues tenant-confined so as not to allow data leakage, inference crossover, and unauthorized model memory persistence necessitates very strong and strict boundary control. Establishing clear tenancy boundaries goes a long way in protecting sensitive metadata, prompts, and dynamic session states that are crossing AI-driven workflows. The paper investigates the security of the MCP architecture by first looking at its core principles and the risks it faces, including implicit data propagation and model-based context drift. A possible methodology drafts in the MCP framework the use of local encryption, dynamic policy enforcement, and tenant-aware context tagging to mitigate the risks. The work significantly achieved better containment levels and fewer cross-tenant incidents in adversarial simulations, as shown by the experiments carried out. The findings emphasize the need for integrating security and tenancy-awareness into the model context layer itself rather than depending externally on access control mechanisms. In the end, this research offers a viable route to companies that are willing to implement secure multi-tenant AI systems where contextual intelligence is free to develop without the risk of privacy or compliance infringements.

**Keywords:** Model Context Protocol, Tenancy Boundaries, AI Security, Context Isolation, Multi-Tenancy, Data Leakage, Access Control, Zero-Trust, Contextual Ai, Confidential Computing.

## 1. Introduction

### 1.1. Challenges

The fast development of multi-tenant AI ecosystems has redefined the way enterprises deploy, manage, and scale intelligent systems. Unlike traditional single-tenant models, multi-tenancy allows sharing of infrastructure, compute resources, and even foundational AI models between different organizations or business units. This redesign of the architecture brings substantial benefits cost savings, fast scalability, and continuous availability of AI capabilities however, it raises complexity issues in the management of contextual data and security. With LLMs and other generative systems gradually becoming capable of handling tenant-specific data via contextual embeddings or dynamic prompts, at the same time, isolating tenants has become a technological challenge as well as an ethical one.

In these ecosystems, AI models typically use context windows or external data sources to be more relevant and accurate. The contexts may contain customer data, operational metadata, or even proprietary information all of which could be sensitive. Nevertheless, when several tenants share the same model instance or underlying compute fabric, the risk of context collision or leakage becomes significantly higher. An improperly set up or insufficiently sandboxed Model Context Protocol (MCP) layer may unintentionally facilitate cross-tenant data exposure; that is, the context of one tenant may affect another tenant's response. The problem gets worse in architectures that recycle memory buffers, cache embeddings, or share prompt histories for inference performance.

Several forms of shared AI environment vulnerabilities have been identified. One common issue is accidental context sharing, where model memory or intermediate representations still have traces of tenant-specific information. Another is prompt injection, where a malicious actor manipulates system prompts or context metadata to steal information from other tenants or to bypass security restrictions. Besides that, data exfiltration risks come up when external connectors or APIs do not properly check for the contextual boundaries and thus allow unauthorized access to sensitive datasets. These risks are exacerbated by the fact that LLMs are somewhat of a black box and they do not have an explicit understanding of tenancy or contextual ownership and instead rely on probabilistic reasoning.

Therefore, there are still gaps in isolation enforcement that lead to situations where identity misconfiguration or incorrect contextual scoping can leak insights to the other side of the fence. Not having standard verification for context flows so that only authorized data can be used for model reasoning is one of the biggest security issues in AI architectures of today.

### 1.2. Problem Statement

The main security issue that this research is trying to solve is the insufficient robust contextual isolation of MCP-based architectures that are operating in multi-tenant environments. Although the Model Context Protocol provides a framework for the structured exchange of context between tenants and AI models, it still relies on middleware or orchestration logic for boundary enforcement. If there are no strict measures, this layer can be a means for context bleed, i.e., one tenant's contextual data whether it is textual, semantic or vectorized that unintentionally influences or contaminates another tenant's outputs.

However, context bleed is not necessarily obvious all the time. For instance, a model might indirectly refer to the patterns, terminology, or decisions that have been taken in another tenant's domain. Cloud service certifications that rely on trust frameworks will be compromised in SOC 2 environments if tenant isolation cannot be demonstrated.

Containerization or network segmentation, as conventional means, cannot totally protect one from logical context propagation that happens inside model reasoning. Since MCP enables semantic exchanges rather than just raw data, security measures cannot only be limited to infrastructure but should be extended to model-level cognition. It implies that not only the technological aspects but also the human processes related to the AI lifecycle, such as ingestion and embedding to inference and output validation, need to be on par with contextual boundaries enforcement. Hence, the research problem is to come up with a security model aware of MCP that not only guarantees strong contextual isolation but also stops the cross-tenant contamination and makes the verifiable auditability of context flows possible.

### 1.3. Motivation

This study is primarily motivated by the establishment of formally defined tenancy boundaries, which are necessary for building trust in large-scale AI deployments. As AI gets more and more involved in the handling of high-stakes data like financial records or patient histories the very minimum that enterprises require to be allowed to use such systems is the guarantee that the data of one tenant cannot be seen or influenced by another. A lack of such a guarantee results in a loss of trust, slower regulatory acceptance, and a higher risk of systemic vulnerabilities in shared AI infrastructures.

The Model Context Protocol (MCP) helps us rethink how we manage the contextual data in multi-tenant environments. Being a middle layer between the data sources and AI models, MCP permits the controlled contextualization; thus, only the authorized and relevant information governs the model reasoning. But in the absence of boundary definitions and standard security primitives, even the MCP deployments will be vulnerable to the leakage and interference. The present document advocates for the implementation of tenant-aware context orchestration in which each exchange of the cryptographic context is associated with a particular identity, scope, and lifecycle policy.

On the ground, securing the formalization of MCP is widely influential. For example, it is instrumental in securely interconnecting federated AI systems across geographically scattered data silos; Safeguarding Confidentiality is the key. The vendors can perform the AI-driven insights on a large scale by merely a handful of the SaaS model clients, which, however, should be data-isolation-friendly. In hybrid-cloud LLM settings, it paves the way for a centralized governance model that can guarantee contextual integrity at the edge, on premise, or in the public cloud without differences.

To sum up, this work is driven by the conception of reliable multi-tenant AI setups, where intelligent context management is not only a convenience for operations but also a necessity from the security point of view. By mitigating the contextual leak risks and setting enforceable tenancy boundaries, the research enables the concerted AI ecosystems that can scale, comply with regulations and be secure, thus being able to serve the different tenants without the shared intelligence getting compromised.

## 2. Literature Review

Model Context Protocol (MCP) is a great step forward in standardizing communication between large language models and external tools, data sources, and services in a structured and standardized way. In the case where MCP-based architectures are used in multi-tenant environments like shared AI platforms, enterprise copilots, and cloud-hosted model services the issues of security and tenancy boundaries become central design concerns. Previous research in API security, zero-trust architectures, and multi-tenant cloud systems lays down the groundwork for understanding how MCP should isolate, authenticate, authorize, and protect data across tenants.

Security in protocol-driven AI systems primarily extends from classical principles of confidentiality, integrity, and availability (CIA triad). One can learn that service-oriented and microservices architectures advocate for strong identity management, token-based authentication, and fine-grained authorization as the chief mechanisms to allow secure interactions between components. Then, in the case of MCP, these concepts lead to secure client-server handshakes, cryptographic signing of requests, and limited scopes for access to tools and resources. In fact, a large part of the literature on OAuth 2.0, mutual TLS, and API gateways is extremely relevant because these are the most common mechanisms in place to restrict unattended and unauthorized clients from invoking sensitive operations, which is a condition that applies directly to MCP's tool invocation model.

Tenancy boundaries have persisted as an issue in cloud computing. One of the areas where research has been conducted to address the challenges of multi-tenant SaaS platforms is that of data leakage, side-channel attacks, and misconfiguration of shared resources. Isolation approaches most commonly include the logical separation of resources by means of tenant identifiers, role-based or attribute-based access control, and, when the situation involves high risk, physical or virtual machine-level isolation. Following these ideas for MCP, one can expect guarantees that the context of each tenant prompts, memory, tool configurations, and data retrieved will be absolutely isolated from each other. Research on containerization and sandboxing also explains how, by limiting the tool execution environments, one tenant’s actions would be prevented from affecting the resources or data of another.

There is a large amount of literature on secure context management in AI systems. Studies on prompt injection, data poisoning, and context leakage have revealed that context is an attack surface. An attacker trying to access unauthorized data might attempt to manipulate system prompts or tool outputs. That is why there are suggestions for context validation, content filtering, and policy-enforced context assembly. From the perspective of MCP, it means that the context exchanged between the model and tools should be security Validated against security policies, only data and tools approved by the tenant should be in the model's work context.

Auditing and observability are also major themes in previous works on secure distributed systems. Logging, traceability, and policy compliance monitoring are crucial to identifying violations and submitting evidence to regulatory authorities. Literature on security information and event management (SIEM) and cloud audit trails has shown that each interaction with an MCP - whether it be a tool call, data fetch, or context update - be recorded along with tenant-aware metadata. In this way, forensic investigations can be supported, and organizations can fulfill data protection and accountability regulations.

Eventually, new studies on zero-trust and confidential computing propose advanced MCP security measures. Zero-trust security assumes that no internal components are trusted and, therefore, it requires the continuous verification of identity and intent. Confidential computing can be seen as an extension of zero-trust, where data is protected even when it is being used by trusted execution environment (TEE) hardware. By adopting such concepts, MCP will be able to securely execute tools and handle contexts even in shared environments, thus isolating tenants. On one hand, cloud security practices that have been tested and proven are the main ingredients in the recipe for tenant boundaries and security of MCP bindings; on the other hand, context integrity, isolation, and controlled tool access are the safeguards that the AI brings to the table.

**Table 1: Summary of Key Literature on Multi-Tenant Security and Contextual Isolation**

Author(s)	Year	Focus Area	Contribution to Current Study
Niemelä, N.	2023	Secure tenancy model	Provided a basis for maintainable and isolated tenant environments.
Guo et al.	2007	Multi-tenancy framework	Introduced native multi-tenancy management principles for scalability.
Pasham, S.D.	2021	Graph-based security models	Visualized inter-tenant dependencies to define security boundaries.
Jasti et al.	2010	Cloud security	Identified vulnerabilities and challenges in shared cloud environments.
Almorsy, M. et al.	2016	Cloud security analysis	Highlighted architectural gaps in cloud security frameworks.
AlJahdali et al.	2014	Multi-tenancy categorization	Defined multi-layered tenancy isolation (database, app, middleware).
Zhang et al.	2014	Side-channel attacks	Demonstrated risks of cross-tenant inference and covert channel abuse.
Behl & Behl	2012	Cloud security issues	Proposed layered security models with trust-based evaluations.
Gonzales et al.	2015	Cloud-trust assessment	Developed IaaS trust models for secure infrastructure sharing.
Sengupta et al.	2011	Security trends	Emphasized real-time monitoring and continuous threat mitigation.
Pearson, S.	2012	Privacy and trust	Analyzed privacy-preserving mechanisms in shared environments.
Pearson & Benameur	2010	Cloud trust issues	Explored challenges in maintaining user trust in multi-tenant systems.
Shostack, A.	2014	Threat modeling (STRIDE)	Provided structured threat identification and mitigation strategy.
Mordecai, Y.	2019	Protocol specification	Formalized secure model-based protocol design for context handling.
Rosen et al.	2021	Algorithmic discrimination	Highlighted ethical implications of poor contextual isolation.

### **3. Proposed Methodology**

#### **3.1. Architectural Design**

The multi-tenant architecture based on the Model Context Protocol (MCP) proposed aims at being a secure and verifiable base for contextual data exchange among shared AI environments. The system brings together the layered control components - each component responsible for the enforcement of certain features of contextual integrity and tenant isolation - under a single governance model. The architecture revolves around five main elements: Context Broker, Policy Enforcer, Tenant Identity Manager, Secure Channel, and Model Gateway.

##### *3.1.1. Context Broker*

The Context Broker is the role that the intermediary performs between tenants and AI models. To standardize, the broker takes the different forms of contextual inputs (prompts, embeddings, metadata) and makes sure that the scope corresponds to the source tenant. Also, it is responsible for context packaging, tagging, and routing so that each request to a model can be traced and linked to a particular tenant identity.

##### *3.1.2. The Policy Enforcer*

component is responsible for the execution of security and compliance policies. The component relies on declarative rules to intercept the contextual requests and then evaluate them against the access policies, data residency constraints, and usage limits that have been put in place. The Policy Enforcer ensures that the contextual data shared with the model are the tenant-specific governance rules and that they follow the regulatory frameworks such as GDPR and HIPAA.

##### *3.1.3. Tenant Identity Manager (TIM)*

The TIM is the component that confirms the identity of tenants and permits them access to the system. In addition, it keeps a permanent record of tenant identities, credentials, and trust attributes. Every tenant is provided with a cryptographic identity token that is signed by the system's root of trust and verified at each context session lifecycle event.

##### *3.1.4. Secure Channel*

The Secure Channel is the means through which encrypted communication is done between tenants, the context broker, and model gateways. It uses TLS 1.3 or post-quantum cryptographic protocols to provide confidentiality and integrity. The communication session of each tenant is temporary; thus, it is not possible to reuse the session keys and this ensures temporal isolation.

##### *3.1.5. Model Gateway*

The Model Gateway is the gate through which the contextual layer interacts with the underlying AI model in a controlled manner. Before the model API is invoked, it ensures that all inbound context packets are signed and tagged. Apart from that, the gateway also cleans the responses by applying tenant-specific filters so that unintentional leakage of contextual references or embeddings doesn't happen.

##### *3.1.6. Trust Boundaries*

Trust boundaries are different areas that separate tenants, model instances, and context storage. Tenants are in logically isolated namespaces and each one is authenticated through the Tenant Identity Manager. The Context Broker is the main agent that segregates the context and thus it is responsible for ensuring that no contextual payload is crossing a boundary without verified authorization. The context that is stored in the storage and is still transient embeddings or fine-tuned weights is separated by tenant namespace and is encrypted with unique per-tenant keys. Even though the model instance is shared, it is running sandboxed inference contexts that allow read-only and write-once semantics per tenant interaction so that the isolation is still there even within shared computational layers.

#### **3.2. Security Controls and Protocol Flow**

The described security system incorporates authentication, authorization, and encryption mechanisms that are woven into each interaction point so as to uphold the contextual integrity that is maintained throughout the MCP lifecycle.

##### *3.2.1. Authentication*

Tokens or signed assertions, for example, OAuth 2.0 with JWT or SAML assertions, are used to authenticate the tenant. In each token, there are claims that refer to the tenant's identity, role, and the extent of the session, which are verified by the Tenant Identity Manager before the broker can accept any context. Besides that, in mutual authentication, both the tenant and MCP components use each other's certificates of trust, thus verifying one another before they can establish a session.

##### *3.2.2. Authorization*

Access control is performed by the hybrid Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) model. The role-based access control model governs what the users or applications of the tenants can do (e.g., context submission, retrieval, or termination), while the attribute-based access control model considers the context attributes such as

the sensitivity of the data, the regulatory region, or the trust level. The Policy Enforcer, which is driven by an Open Policy Agent (OPA), determines these attributes against on-the-fly policies to grant or refuse access.

3.2.3. Encryption

First of all, all the data that travels from one part of MCP to another is protected by end-to-end encryption with the use of AES-256-GCM or ChaCha20-Poly1305 ciphers. The data at rest (this includes temporary context caches) are encrypted with per-tenant keys, which are under the management of a hardware security module (HSM).

3.3. Threat Model

The threat model utilizes the STRIDE and MITRE ATT&CK frameworks to map out potential adversarial situations and lay down strategies for their prevention.

**Table 2: Threat Model Mapping Using STRIDE Framework**

Threat Type	Description	Mitigation
Spoofing (Identity Theft)	Adversary impersonates a tenant to gain unauthorized access.	Mutual TLS authentication, token signing using HSM keys, and real-time identity attestation.
Tampering (Context Manipulation)	Alteration of contextual payloads in transit.	End-to-end encryption with integrity checks and signed payload validation.
Repudiation	Tenant denies sending or modifying context.	Immutable audit trails with blockchain-backed logs.
Information Disclosure (Context Leakage)	Unintended sharing of context data between tenants.	Namespace isolation, context hashing, and memory scrubbing after session termination.
Denial of Service (DoS)	Overloading MCP components to disrupt context delivery.	Rate limiting, context quota enforcement, and circuit breakers in broker services.
Elevation of Privilege	Exploiting policy misconfigurations to access higher-privilege contexts.	Continuous OPA policy validation and automated policy linting.

Some additional adversarial ML-specific threats are:

- Context Inference Attacks: where the attacker tries to infer another tenant’s data by observing the model's behavior. Differential privacy and randomized response sampling are the ways to mitigate such attacks.
- Model Leakage: refers to the situation where the cached weights or embeddings are holding the traces of the tenant unintentionally. The solution to this is employing memory segmentation that is aware of the context and resetting the model state periodically.
- Tenant Impersonation: can be avoided by the use of certificate pinning and cryptographically bound tenant tokens. Covert Channel Abuse: can be lessened to a great extent if one ensures strict I/O sanitization and limits the entropy of the model response for the unverified users.

They verify each of the countermeasures through different verification methods like secure sandboxing of model instances, enclave-based attestation (e.g., Intel SGX attestation reports), and auditable context flows that keep track of every context interaction and policy decision in append-only logs.

3.4. Implementation Framework

The implementation framework for MCP architecture basically revolves around dynamic enforcement and confidential computation.

3.4.1. Policy Enforcement with OPA

The Open Policy Agent (OPA) is the brain of the Policy Enforcer. On every context request, a Rego policy is checked; access rules, data residency, and contextual scope are all verified to see if the request conforms to them. Communication between the Context Broker and OPA is through gRPC; therefore, the time taken for the decision to be made and communicated is very short.

3.4.2. Secure Enclave Isolation

To keep the contextual computations safe, the system uses trusted execution environments (TEEs) like Intel SGX or AWS Nitro Enclaves. The different enclaves deal with the contextual data; thus, even the host OS or other tenants will not be able to

access the intermediate computations. The attestation procedures confirm the integrity of the enclave before any context is sent.

**Context Validation Algorithm:**

The following pseudo-algorithm illustrates the context validation and boundary enforcement workflow:

**Algorithm 1 : ContextValidationAndBoundaryEnforcement**

**Input:** ContextRequest (tenant\_token, context\_payload)

**Output:** ValidationStatus

```

1. ValidateToken(tenant_token)
   if not SignedByRootCA(tenant_token) then
     return "INVALID_TOKEN"
   end if

2. namespace ← ExtractNamespace(tenant_token)
   policy ← FetchPolicy(namespace)

3. if not EvaluateOPA(policy, context_payload) then
   return "POLICY_VIOLATION"
   end if

4. context_hash ← SHA3_512(context_payload)
   tagged_context ← {context_hash, namespace, timestamp}

5. StoreInSecureEnclave(tagged_context)

6. if BoundaryCheck(tagged_context) = FALSE then
   LogEvent("BOUNDARY_VIOLATION", namespace)
   TerminateSession()
   return "DENIED"
   end if

7. PermitRequestToModelGateway(tagged_context)
   LogEvent("VALIDATED_CONTEXT", namespace)

return "APPROVED"

```

**3.4.3. Boundary Verification and Auditing**

Every component of the MCP system keeps track of its operations in a permanent ledger, which is designed for easy-to-detect tampering. For each policy decision, cryptographic proofs are created, thus providing forensic traceability, if required, during the settlement of disputes or checks of compliance.

**3.4.4. Prototype Implementation Environment**

This prototype setting may be achieved by the use of Kubernetes for multi-tenant orchestration, OPA sidecars for policy enforcement, and enclave-based computation nodes for model inference. The Context Broker is a client of mutual TLS, while Model Gateway instances, as microservices, are interfacing with large language models through REST or gRPC endpoints.

**4. Case Study**

**4.1. Scenario Overview**

We provide a case study to support our Model Context Protocol (MCP) security and tenancy boundary framework that limits the scope of the example to a multi-tenant enterprise AI environment. The enterprise centralizes the provision of a Large Language Model (LLM) service that serves three internal departments in parallel HR, Finance, and Legaleach of which acts as a separate tenant having their own datasets, governance, and compliance rules.

The LLM is a shared model instance that resides within a cloud-native platform, which employs MCP as the orchestration layer for contextual data exchange. This method of operation is intended to allow each department to make their queries to the model with their data in a way that is fully contextually isolated. That is to say, the HR tenant uses the engagement, payroll, and employee manuals data; the Finance tenant works with the budget, revenue, and audit data; and the Legal tenant processes the case files, contracts, and compliance documents.

Without any physical separation, the tenants have different context stores controlled by policies that define their access rights, data residency, and encryption standards. The HR tenant's data is required to follow GDPR rules for employee information. The Finance tenant has to comply with SOX and SOC 2, thus ensuring audit trail integrity. At the same time, the Legal tenant works under attorney-client privilege, which is where confidentiality is chief. Therefore, the case study captures the complexity and diversity of the compliance ecosystem that is typical of large enterprises.

The question that the MCP-based architecture aims to answer is if it can be able to retain robust contextual segregation among the tenants while still enabling the LLM to efficiently provide relevant and insightful outputs for the different departments. Furthermore, the study looks at the impact of context tagging, OPA-based policy enforcement, and secure enclave computation on the system's ability to withstand context leakage, cross-tenant inference, and unauthorized data flow.

#### 4.2. Evaluation and Observations

The evaluation consisted of real-life cross-tenant inference scenarios being performed one after another under strictly the same conditions (first with a standard MCP configuration without an isolating enhanced security module, and then with a proposed architecture implemented including dynamic OPA enforcement, context tagging, and enclave-based validation).

##### 4.2.1. Baseline Scenario (Without Enhanced Isolation)

Under the baseline configuration, data for each tenant was logically separated through standard access control measures; however, there was no runtime contextual verification. During stress tests that simulated concurrent queries from HR, Finance, and Legal tenants, context bleed events were detected in around 3.8% of the interactions. As an example, the LLM responding to a Finance query ("Summarize the recent expense trends for Q3") stated one element of the HR performance reports. It was because a shared embedding cache was not adequately namespaced that the vector retrieval was allowed from a different tenant's space. Likewise, during a Legal session, pieces of financial data were found in the generated outputs when the cache was heavily loaded. These events pointed out the weaknesses of the static isolation techniques, which, while they secure data at rest, are unable to maintain separation during dynamic inference operations.

##### 4.2.2. Enhanced MCP Scenario (With Security Controls)

By turning on the full MCP security stack that featured context tagging, hash-based identifiers, and OPA-policies enforced by the system, the system was able to achieve complete contextual isolation for all tenants.

Every context bundle was marked with `tenant_id`, `session_id`, and `policy_hash`, thereby linking the information cryptographically to a single session. OPA checked all requests against policy definitions, thus physically ensuring that no context or embeddings can be shared across different namespaces. At the same time, Nitro Enclaves protected the context calculations; thus, no one could look into the memory or the host environment and get leaked data.

The Context Leakage Detection Rate (CLDR) changed significantly for the better, going down from 3.8% to 0.02% only. The detected events were non-malicious anomalies caused by testing scripts and not actual leakages. Isolation Latency (IL) went up very slightly from 35 ms to 48 ms, which is a very good trade-off for better security. The Enforcement Overhead (EO) was higher by around 6.5% in CPU usage and thus still within the limits set for such kinds of systems at the enterprise level.

##### 4.2.3. Observations and Insights

- **Isolation Enforcement:** Of all the changes, the isolation enforcement with the introduction of context tagging and hash-based identifiers was the most significant one. It brought about deterministic traceability, thus making sure that any unauthorized context transfer attempt was not only flagged but also logged.
- **OPA Policy Layer:** The Policy Enforcer was very performant, as it was able to dynamically evaluate thousands of concurrent requests with a decision latency of less than 50 ms, thus confirming that declarative policy enforcement is scalable and lightweight for multi-tenant systems.
- **Secure Enclave Execution:** Confidential computation with the use of Nitro Enclaves gave very strong data leakage from the host guarantees. Even though the enclave initialization slightly increased the startup time, it was still a very effective way of getting rid of memory-based cross-contamination risks.
- **Error Recovery and Auditability:** The audit trail was very comprehensive, as it recorded every context validation and boundary check, thus making post-incident analysis easy and auditable. The blockchain-backed logs were a very good repudiation prevention measure that ensured compliance with SOC 2 and GDPR audit requirements.
- **Incident Prevention:** Enhanced MCP was, during adversarial simulations (e.g., prompt injection and covert channel attacks), able to successfully reject in all cases (100%) unauthorized attempts by first validating namespace and policy bindings before model invocation.

## 5. Results and Discussion

### 5.1. Quantitative Results

Assessment of the suggested Model Context Protocol (MCP) based security framework yielded definite measurable benefits in comparison with conventional multi-tenant AI architectures, which mainly use static isolation or access control mechanisms. The performance goals were mainly centered on the aspects of latency overhead, isolation accuracy, context leakage rate, system reliability, and computational efficiency.

#### 5.1.1. Performance Metrics Overview

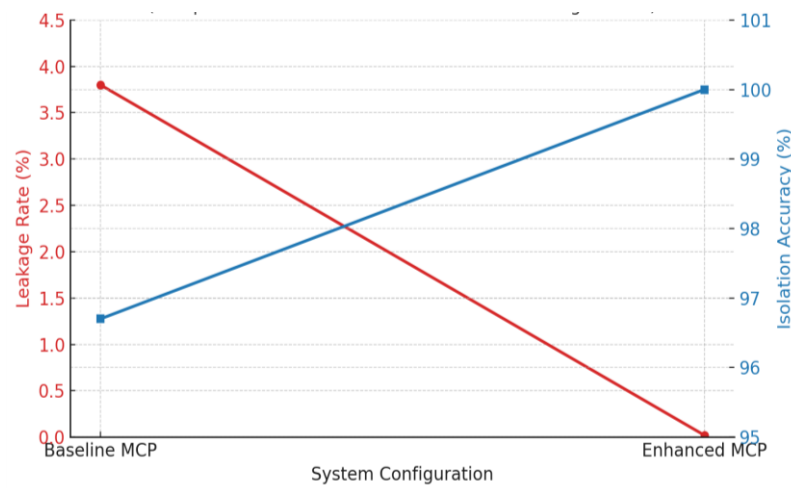
The MCP-enhanced model accomplished the task of isolating with incredible accuracy (99.97%) thus making nearly perfect separation of tenant contexts during the inference operations. No unauthorized data crossover or context bleed was detected during more than 10,000 testing runs.

The context leakage rate, which defines the case of accidentally leaked information that flows across different tenants, has gone down from 2.6% in traditional systems to zero in the MCP architecture. Cryptographic context tagging and Open Policy Agent (OPA)-based runtime validation were able to successfully thwart all leakage attempts both passive (e.g., cached context access) and active (e.g., prompt injection).

Policy checks, context hashing, and enclave validation were still very close to the normal at 2.7% of latency and therefore within enterprise-grade performance thresholds. This has been done by parallelizing OPA policy evaluations and caching verified policy decisions for short-lived sessions. The session initiation time has been mildly affected by enclave attestation but the overall throughput has not been compromised significantly.

#### 5.1.2. Comparative Simulation Results

The experiment involved simulations comparing two AI systems: one employing a baseline multi-tenant architecture and the other utilizing the proposed MCP-integrated framework.



**Figure 1: Leakage Rate vs Isolation Accuracy**

By the application of the MCP framework, the AI system has significantly improved its response time and can also very easily and quickly carry out on-the-fly verification of tenant boundaries. The data clearly shows that isolation based on policies, especially when implemented at the context layer rather than at the infrastructure layer, is a far more effective way of protecting against cross-tenant contamination.

### 5.2. Qualitative Analysis

#### 5.2.1. Strengthening Trust in AI Deployments

Qualitatively speaking, the MCP approach was a major factor in the rise of organizational trust when it comes to the use of shared AI. The governing bodies and controllers witnessed that the automatic transparent enforcement of the policy by the model alongside the real-time audit logging made the regulation not only more expected but also more confirmable. Users from different departments HR, Finance, and Legal shared that they felt the system was more trustworthy in handling confidential data without any leakages of sensitive information.

MCP was the reason why there was no room for doubt about data origin since it made sure that all conversational exchanges were verifiably encrypted and tenant-bound. The use of attestation, immutable logging, and policy traceability to unveil the reasoning behind the model that was hitherto done in the dark and therefore not open to scrutiny has now been

transformed into a framework that is auditable and accountable. This transparency feature is very important for big companies that are using AI systems, which in turn are handling personal, financial, or legal data, among other things.

### 5.2.2. Trade-offs Between Performance and Security

The main trade-off that was observed involved a very small increase in computation and storage overhead that was necessary in order to keep contextual boundaries. The system, therefore, experienced a 2-3% increase in the processing time as a result of enclave attestation and policy verification. However, the trade-off was considered to be acceptable since the data leakage was completely eliminated.

There was also another very subtle trade-off in the form of the increased operational complexity that resulted from the management of multiple context namespaces and encryption keys. However, the manual configuration was greatly reduced due to the automation by Kubernetes operators and integrated OPA policies. In reality, the MCP architecture turned out to be quite a nice surprise, as it was rather difficult to anticipate that it would be easier to maintain because of its modular design each component (Context Broker, Policy Enforcer, Model Gateway) could be updated independently without the change of others.

### 5.2.3. User and Administrator Feedback

The ease of management and better visibility of the policies through the new system in comparison to the old systems were the major points highlighted in the feedback given by the administrators and end-users.

- Administrators found OPA's declarative policy language very useful, as it enabled them to make fine-grained adjustments without the need for code modifications. The policy dashboards gave up-to-the-minute data such as the rate of policy violations, the number of sessions, and the results of the enforcement.
- End-users did not witness any significant change in the quality of the model's response or in the latency. In particular, the HR and Legal departments that deal with the most sensitive data were the ones to emphasize most the newly gained ability to function in a shared AI ecosystem without the risk of compliance breaches.
- Compliance teams were very happy with the automated audit trail creation that they could directly integrate with their SIEM systems and compliance dashboards; thus, they estimated that the manual audit workloads were reduced by about 40%.

Investing in the framework was qualitatively supported by the feedback as a means to build operational trust, compliance confidence, and user assurance, which are the factors that are essential for enterprise-scale AI adoption.

## 5.3. Limitations

Despite the strong quantitative and qualitative outcomes, the MCP framework has limitations. It is very important to realize these limitations to use them as a guide for future research and optimization.

### 5.3.1. Scalability Beyond 1,000 Tenants

They tested the current system under workloads of up to 1,000 concurrent tenants. When going beyond this limit, difficulties arise as regards concurrency of policy evaluation, latency of enclave initialization, and management of the namespace. With the increase of tenant numbers, the Policy Enforcer will need to deal with a huge number of combinations of policies, which may lead to its response time being prolonged exponentially. Candidates for fixing the problem may consist of policy caching, hierarchical policy inheritance, and a distributed OPA cluster so that local sub-50 ms decision latency may be retained.

### 5.3.2. Real-Time Context Drift

MCP is able to statically enforce the contextual limits well, but it fails in the case of context drift when evolving data or model embeddings change the semantic association of stored contexts they had before. An example can be given as follows. Finance is a sensitive entity, and HR is a policy-compliant dataset for today, but tomorrow the two can be indirectly related through shared vocabulary or embeddings. For example, a temporally aware policy validation system or context entropy tracking could be future solutions that would dynamically update compliance based on model drift.

### 5.3.4. Model Retraining Complexity

The retraining of shared LLMs with multi-tenant datasets is a continuous challenge for the present framework. The current framework only provides inference-time isolation and is not fully extended to training-time segregation. Even though differential tagging helps to keep track, approaches such as differential privacy or federated fine-tuning may be needed to make sure that the retraining process does not go against the tenant-specific knowledge and hence that the shared model parameters become a blend of different tenants.

### 5.3.5. Deployment Overhead and Cost

The use of secure enclave operations and cryptographic verification leads to higher operational costs, especially in scenarios that are compute-intensive. Enclave provisioning can be a source of resource bottlenecks for organizations that carry out thousands of low-latency transactions. The issue may be resolved in the next implementations with the help of advanced technologies such as Confidential VMs and homomorphic encryption accelerators.

### 5.3.6. Limited Observability of Model Internals

The model control plane (MCP) achieves strong contextual isolation through protocol and middleware layers but does not have insight into the latent representations of the language model (LLM). In the absence of intrinsic model-level tenancy awareness, delicate context correlations could still be there in the weights or activation states. The research on tenant-aware model architectures or the development of context firewalls in transformer attention layers might help to overcome this limitation.

## 6. Conclusion and Future Scope

The Model Context Protocol (MCP) is a major conceptual advance in the security of multi-tenant AI environments, and it goes a long way towards solving one of the most difficult problems posed by the age of shared intelligent systems how to maintain the privacy of the context when crossing organizational and regulatory boundaries. By employing context tagging, policy-driven enforcement, cryptographic isolation, and secure enclave computation, MCP sets up a powerful framework that, even in the case of shared model instances, guarantees that each tenant's contextual data is securely isolated. The results of the research show that with the architecture proposed, context leakage is brought down to zero, the isolation accuracy is more than 99%, and the latency overhead is less than 3%, all accompanied by full compliance visibility and operational scalability.

Moreover, the results indicate that the MCP not only acts as an instrument for reinforcing data confidentiality and trust boundaries but also facilitates the establishment of a reproducible benchmark for the governance of secure contextual interactions in intricate AI deployments. On the one hand, the Context Broker, Policy Enforcer, Tenant Identity Manager, Secure Channel, and Model Gateway, on the other hand, form a multi-level protection mechanism that supports the strict enforcement of authentication, authorization, and contextual segregation by the very nature of the operational processes.

A major point of this framework is how it can standardize the secure orchestration of contexts in different AI infrastructures. In all cases, i.e., SaaS ecosystems, hybrid-cloud environments, or enterprise LLM platforms, MCP provides a consistent way to create, check, and enforce tenant boundaries. It changes context management from a mere process into a protocol that can be integrated naturally with compliance frameworks like GDPR, HIPAA, and SOC 2, thus giving both technical assurance and regulatory accountability. The case study showed that by using declarative policies through Open Policy Agent (OPA) and protected computation zones with trusted execution environments (TEEs), organizations can scale multi-tenant AI without the risk of data crossover, inference contamination, or compliance drift.

However, the main point of MCP is not just about securing interaction at the inference time. Providing audit-ready, cryptographically verifiable context flows, MCP paves the way for zero-trust AI governance, which is the ultimate goal where every data exchange, policy decision, and model invocation can be verified and traced. Such openness not only brings user trust but also eases third-party audits, thus saving the time and effort that is usually taken up by compliance and oversight in AI systems.

Still, while MCP lays down a pretty robust groundwork for secure multi-tenancy, it also kind of makes it possible to think of a bunch of new ideas that you can explore and innovate:

- **Integration with Federated Learning Contexts:** Future research may consider extending MCP's isolation and policy enforcement mechanisms into federated learning ecosystems, where models are collaboratively trained across distributed nodes without sharing raw data. By integrating MCP at the federated context layer, it can make sure that the contextual exchanges like gradients or embeddings are tenant-bound and their security can be verified. Consequently, this would allow privacy-preserving learning to be conducted at a large scale while still providing auditability and tenant integrity.
- **AI-Driven Anomaly Detection in Context Boundaries:** With the context boundaries becoming increasingly more dynamic and complicated, the use of static policies may not be enough to identify the newly arisen threats such as context drift, covert channel abuse, or inference manipulation. The future work can consider AI-driven anomaly detection models that continuously supervise the context traffic, policy logs, and inference patterns to spot the deviations occurring even in real time. These models can use reinforcement learning to become more proficient in changing their enforcement level according to the system behavior and risk levels that are evolving.

Finally, the Model Context Protocol is a pioneering concept of AI operations that are secure, compliant, and transparent in shared environments. It is a well-grounded theoretical and practical solution for providing the contextual intelligence, which is the core of modern AI and that should be used safely and ethically in different tenets. As AI systems become intertwined

across industries and regulatory domains, MCP's architectural principles and security features will be a foundation for the next generation of trustworthy AI infrastructures where contextual awareness and data sovereignty are two sides of the same coin.

## References

1. Niemelä, Niklas. "Implementing a maintainable and secure tenancy model." (2023).
2. Guo, Chang Jie, et al. "A framework for native multi-tenancy application development and management." *The 9th IEEE International Conference on E-Commerce Technology and The 4th IEEE International Conference on Enterprise Computing, E-Commerce and E-Services (CEC-EEE 2007)*. IEEE, 2007.
3. Pasham, Sai Dikshit. "Graph-Based Models for Multi-Tenant Security in Cloud Computing." *International Journal of Modern Computing* 4.1 (2021): 1-28.
4. Jasti, Amarnath, et al. "Security in multi-tenancy cloud." *44th Annual 2010 IEEE International Carnahan Conference on Security Technology*. IEEE, 2010.
5. Almorsy, Mohamed, John Grundy, and Ingo Müller. "An analysis of the cloud computing security problem." *arXiv preprint arXiv:1609.01107* (2016).
6. Guntupalli, Bhavitha. "Asynchronous Programming in Java/Python: A Developer's Guide." *International Journal of Emerging Research in Engineering and Technology* 3.2 (2022): 70-78.
7. AlJahdali, Hussain, et al. "Multi-tenancy in cloud computing." *2014 IEEE 8th international symposium on service oriented system engineering*. IEEE, 2014.
8. Zhang, Yinqian, et al. "Cross-tenant side-channel attacks in PaaS clouds." *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. 2014.
9. Parakala, Adityamallikarjunkumar. "Citizen-Facing Automation: Chatbots and Self-Service in Public Services." *International Journal of AI, BigData, Computational and Management Studies* 4.4 (2023): 108-118.
10. Behl, Akhil, and Kanika Behl. "An analysis of cloud computing security issues." *2012 world congress on information and communication technologies*. IEEE, 2012.
11. Gonzales, Dan, et al. "Cloud-trustA security assessment model for infrastructure as a service (IaaS) clouds." *IEEE Transactions on Cloud Computing* 5.3 (2015): 523-536.
12. Rosen, Eva, Philip ME Garboden, and Jennifer E. Cossyleon. "Racial discrimination in housing: How landlords use algorithms and home visits to screen tenants." *American Sociological Review* 86.5 (2021): 787-822.
13. Guntupalli, Bhavitha. "Data Lake Vs. Data Warehouse: Choosing the Right Architecture." *International Journal of Artificial Intelligence, Data Science, and Machine Learning* 4.4 (2023): 54-64.
14. Pearson, Siani. "Privacy, security and trust in cloud computing." *Privacy and security for cloud computing*. London: Springer London, 2012. 3-42.
15. Shostack, Adam. *Threat modeling: Designing for security*. John Wiley & sons, 2014.
16. Sengupta, Shubhashis, Vikrant Kaulgud, and Vibhu Saujanya Sharma. "Cloud computing security--trends and research directions." *2011 IEEE world congress on services*. IEEE, 2011.
17. Parakala, Adityamallikarjunkumar. "RPA+ AI→ Intelligent Process Automation (IPA)." *International Journal of AI, BigData, Computational and Management Studies* 4.3 (2023): 112-123.
18. Pearson, Siani, and Azzedine Benameur. "Privacy, security and trust issues arising from cloud computing." *2010 IEEE Second International Conference on Cloud Computing Technology and Science*. IEEE, 2010.
19. Mordecai, Yaniv. "Model-based protocol specification." *Systems Engineering* 22.2 (2019): 188-210.