

# The Definitive Guide to Zero Trust AI Security

Protecting Your MCP, Models, Agents,  
Datasets, and Business Systems



## PREFACE

### Who Is This Guide For?

This guide is designed for anyone building, deploying, or managing AI, from personal projects to large-scale enterprise systems. Whether you're an:



#### **Enterprise AI/ML Engineer, DevOps, or Security Professional**

Responsible for securing mission-critical AI deployments, ensuring compliance, and protecting sensitive corporate data.



#### **AI Startup Founder or Engineer**

Aiming to build security directly into your AI product from day one, ensuring trust and agility as you scale.

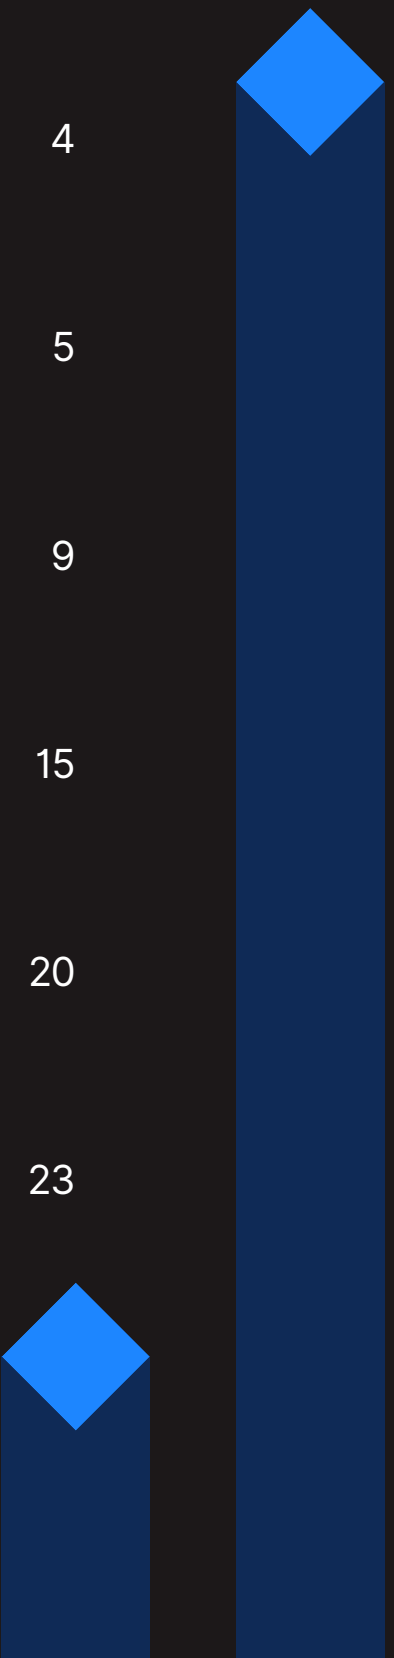


#### **AI Hobbyist or Independent Developer**

Looking to secure your personal AI projects, experiments, or home lab without wrestling with complex network configurations.

**If you're using or planning to use AI with external tools, data, or internal systems – this guide is for you.**

TABLE OF CONTENTS

Foreword: The New Frontier of Risk in the Age of AI .....	4	
Chapter 1: The AI Attack Surface - Understanding Critical Vulnerabilities .....	5	
Chapter 2: The Ideal AI Security Solution: Layering Mitigation and Protection .....	9	
Chapter 3: NoPorts: Your Blueprint for Invisible AI Security .....	15	
Chapter 4: Putting Invisible AI Security Into Practice: Getting Started with NoPorts .....	20	
Next Steps .....	23	

## FOREWORD


# The New Frontier of Risk in the Age of AI

Artificial intelligence is rapidly transforming industries, unlocking unprecedented innovation and efficiency. But as AI becomes deeply integrated into our digital lives and core business operations, it introduces a complex new landscape of security risks that traditional methods simply can't handle.

For too long, security has meant drawing lines around a perimeter, like a digital fence. But AI, with its dynamic models, autonomous agents, and fluid data flows, often makes that fence much harder, if not impossible, to define. Or worse, that fence becomes a roadblock that delays or prohibits full deployment. The stakes are incredibly high. A compromised AI system can lead to severe data breaches, intellectual property theft, service disruptions, and even financial fraud.

To harness AI's full potential safely, we must adopt a fundamentally new approach to security: zero trust, built on the concept of an invisible infrastructure. This guide will show you how to protect your AI assets, big or small, by eliminating their digital visibility and securing them from the inside out.

### Zero Trust Core Principles<sup>1</sup>

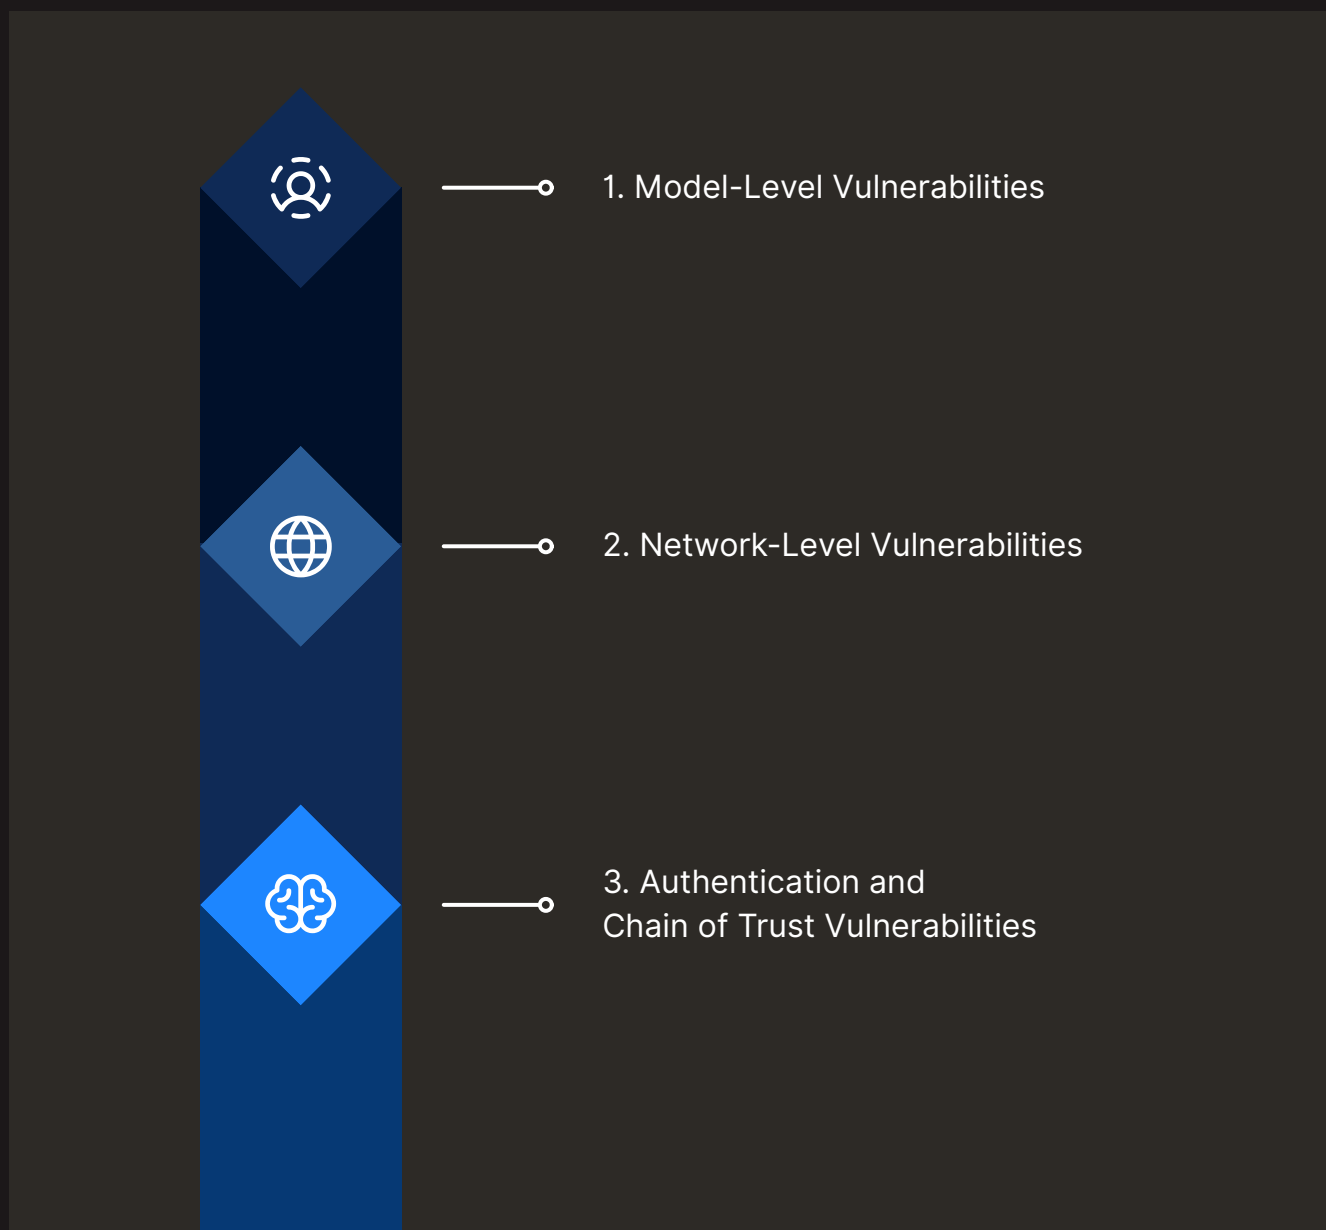


<sup>1</sup>Source: Gartner, 2024

## CHAPTER ONE

# The AI Attack Surface - Understanding Critical Vulnerabilities

The integration of AI, especially with the rise of autonomous agents and large language models (LLMs), fundamentally changes how we think about digital security. We're moving from a world of human-centric access to one where AI models and agents directly interact with, and even control, critical systems. This paradigm shift creates a vast and often unseen attack surface that can be categorized into distinct areas:



### 1. Model-Level Vulnerabilities

Model-level vulnerabilities arise from the AI model's internal logic, its training data, or how it processes information and generates responses. These are inherent risks within the AI itself, regardless of its network connectivity. While these risks exist, the Model Context Protocol (MCP) emerges as a critical strategy to mitigate many of them, as we'll explore in Chapter 2. Without proper mitigation, these vulnerabilities can lead to:

#### Bad Advice & Distorted Decisions

An AI model might be designed to provide advice or make decisions, but if it's flawed, poorly trained, or subtly manipulated (e.g., via a prompt injection that leverages its internal processing), it can lead to incorrect or harmful outcomes. This could manifest as a miscalculated game move in a personal project, or a critical financial misstatement for an enterprise, potentially leading to liability for the company.

#### Prompt Manipulation and Injection

This is a subtle yet powerful form of attack where malicious instructions or hidden prompts are engineered to exploit the model's processing capabilities. Without proper internal safeguards (like robust input validation or output sanitization at the model's interaction layer), these can override legitimate AI instructions, leading to unintended or harmful AI behavior, or cause the model to leak sensitive data or generate responses that expose your company to liability.

#### Sensitive Data Leakage & Exfiltration (Model-Driven)

An AI model, especially if over-privileged or designed without strict data minimization, could inadvertently reveal sensitive information embedded in its training data, or could be prompted to exfiltrate data it has access to. This leads to intellectual property theft, privacy breaches, and regulatory non-compliance, affecting anyone from an individual researcher to a large enterprise.

#### Factual Inaccuracies

A model might subtly misrepresent information or generate responses that are not in your best interest due to internal biases or successful manipulation, leading to a profound erosion of trust and potential operational risks.

# THE AI ATTACK SURFACE - UNDERSTANDING CRITICAL VULNERABILITIES

## 2. Network-Level Vulnerabilities

Network-level vulnerabilities are the foundational security weaknesses that arise from how AI components (including MCP servers, models, agents, and data sources) are exposed and communicate across networks. These vulnerabilities create direct entry points for attackers.


### Digital Invisibility is Crucial

Imagine a personal project server or a startup's new service with an open back door visible from the internet. It's an obvious invitation for trouble. Similarly, if your MCP server or other AI endpoints have open, listening ports, they're digitally visible to attackers. This makes them vulnerable to constant scans, probing, and direct assaults, regardless of whether it's an enterprise data center or a home lab.

### The Risk of a Compromised MCP Server (Due to Network Exposure)

MCP acts like a "digital nervous system" connecting AI models to tools and data, a compromised MCP server (reached via an open network port) becomes an incredibly dangerous pivot point. Attackers can then:

- **Execute Malicious Actions** - Instruct the AI to perform unauthorized transactions, delete intellectual property, or manipulate crucial business processes.
- **Feed Bad Data** - Inject manipulated information directly into the AI's context.
- **Act as a Backdoor for Lateral Movement** - Use the compromised MCP as a beachhead to move deeper into your network and access other sensitive systems (e.g., CRM, ERP, internal file servers).
- **Trigger DoS Attacks** - Force the AI into infinite call loops or over-consume resources, effectively shutting down your AI services.



### Unencrypted Traffic

Data in transit between AI components, MCP servers, and connected tools is vulnerable if not properly encrypted. This allows for eavesdropping, data tampering, and replay attacks.

### 3. Authentication and Chain of Trust Vulnerabilities

Authentication and chain of trust vulnerabilities stem from insufficient identity verification and weak access controls, allowing unauthorized entities to connect to, or impersonate, legitimate AI components.

#### Authentication & Access Control Gaps

Many AI component deployments lack robust authentication and granular access controls. This means unauthorized people or malicious actors can easily gain access without proper identity verification, assuming the identity of a legitimate service, agent, or person.

#### Misconfigured Access / "Operator Error" Vulnerability

Even with good intentions, the incorrect setup of permissions can expose sensitive systems. For example, a private MCP server linked to a publicly accessible AI model creates a dangerous gap. This vulnerability arises when an operator (or an automated process configured by an operator) grants excessive or inappropriate access, allowing unintended interaction between public-facing AI and sensitive internal resources. This isn't necessarily a "hack," but a security flaw stemming from misconfigured trust.

#### Excessive Privileges

Often, AI components (including MCP servers) operate with overly broad network permissions or access rights. If compromised due to weak authentication or network access, this "excessive privilege" creates a wide "blast radius" for an attacker, allowing them to cause much more damage than anticipated.

#### Unverified Connections / Spoofing

Without a strong chain of trust, it's difficult to verify that the AI model talking to the MCP, or the agent accessing a business system, is truly the one it claims to be. This enables impersonation and unauthorized interactions across your AI ecosystem.

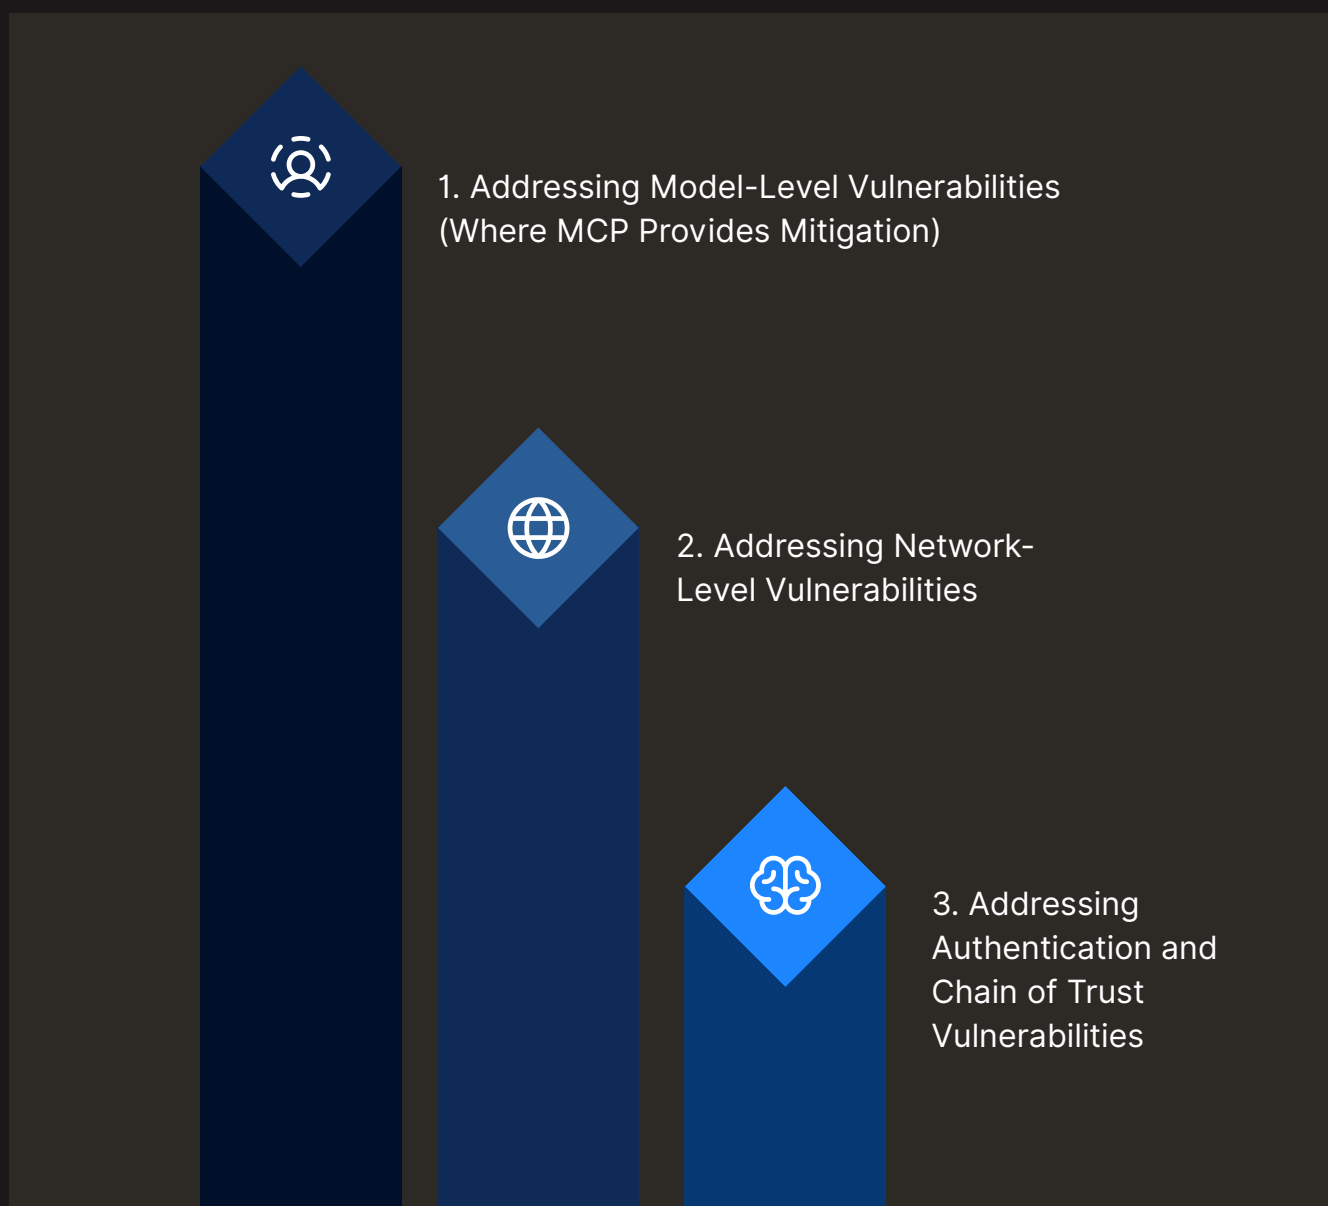


## CHAPTER TWO

# The Ideal AI Security Solution: Layering Mitigation and Protection

Given the multifaceted nature of AI security risks, an ideal solution requires a layered approach. It's not just about guarding the perimeter; it's about building robustness into the AI's interaction layer (MCP) and securing the very foundation of connectivity (network and authentication).

Here are the essential characteristics of a robust AI security solution that directly addresses the vulnerabilities outlined in Chapter 1:



## CHAPTER TWO

### 2.1 Addressing Model-Level Vulnerabilities

At the heart of many AI deployments lies the Model Context Protocol (MCP). Think of MCP as the "digital nervous system" that connects your AI models to external tools, data sources, and internal systems. It's how your AI retrieves information and issues commands. MCP emerges as a critical strategy to mitigate many model-level vulnerabilities, enhancing an LLM's reliability and ethical performance:

#### ✓ Providing Richer Context and Sources of Truth

By feeding an LLM a dependable source of factual information via MCP, the model doesn't have to "predict" its next words as often. Instead, the model focuses on reformatting the information that it has already been provided via the MCP service. This drastically reduces inaccuracies beyond what is possible with just a model alone, thereby reducing liability risks from incorrect responses.

#### ✓ Defining Clear Actions and Enabling Human Oversight

MCP can delineate specific, permissible actions for an LLM. This enables crucial human intervention: the LLM proposes an action, and a human can approve or deny it. This safeguard prevents the LLM from executing autonomous actions that could harm the business, mitigating risks of unintended or malicious actions.

#### ✓ Structured and Predictable Outputs

Unlike LLMs, which often produce unpredictable natural language outputs, MCP services ensure that the tools they expose have clearly defined inputs and outputs. This makes AI-driven actions more auditable, accountable, and, in some cases, reversible, which is vital for managing liability and ensuring trustworthy operations.

#### ✓ Meticulous Input and Output Handling

- **Rigorous Input Validation** - Any data or instructions entering the MCP server (whether from an AI model, another agent, or a person) must be strictly checked for safety, expected format, and adherence to defined policies. This is vital to prevent hidden malicious commands, prompt injections, or malformed data that could exploit the model's logic.
- **Safe Responses (Output Sanitization)** - Similarly, any data leaving the MCP server (back to the AI or other systems) should be validated and sanitized to prevent sensitive information leaks or unexpected responses that could cause issues for the AI or downstream systems.

## CHAPTER TWO

### 2.2 Addressing Network-Level Vulnerabilities

These characteristics are crucial for eliminating the external attack surface and ensuring that AI components, including MCP servers, cannot be directly targeted via network exposure. An ideal solution here focuses on invisibility, secure communication, and fine-grained network segmentation at the network layer:

#### ✓ Network Invisibility / Elimination of Open Ports

This is perhaps the most critical characteristic. An ideal solution completely eliminates the need for open, listening ports on AI components, including MCP servers, models, and agents. If attackers can't "find" your AI systems on the internet or even within your internal network with port scans, they can't initiate direct attacks. This concept of "digital invisibility" is a proactive defense that drastically shrinks the attack surface for everyone, from a hobbyist's cloud instance to an enterprise data center.

#### ✓ True Micro, or Better Yet, Picosegmentation

Beyond simply hiding endpoints, an ideal solution enables microsegmentation – creating isolated, secure zones for individual workloads or applications – and even picosegmentation, extending this isolation down to individual processes or functions. This ensures that even if one component is compromised, the breach cannot easily spread laterally across the network. This level of granular control is crucial for containing potential threats and protecting adjacent systems.

#### ✓ Solid Communication and Data Protection

- **End-to-End Encryption**

All communication, not just over the public internet, but internally between AI components, MCP servers, and connected tools, must be fully encrypted. This prevents eavesdropping, data tampering, and ensures the integrity of instructions and responses, protecting your creative prompts or your company's proprietary data.

- **Data Minimization in Transit**

While encryption is key, an ideal solution also facilitates transmitting only the bare minimum of data necessary for a given task, further reducing the potential exposure of sensitive information.

### 2.3 Addressing Authentication and Chain of Trust Vulnerabilities

These characteristics ensure that only authorized and verified entities can interact with AI components, establishing a robust security posture from the point of access. An ideal solution provides granular control and continuous verification:

#### ✓ Strict Identity and Access Control

##### • Cryptographically Proven Chain of Trust

The solution must enforce Zero Trust by ensuring that the identity of every entity in the communication chain—human, AI model, or agent—is cryptographically proven and verified at every step. This goes beyond simple passwords or IP addresses, establishing an unbreakable chain of trust. This is vital for scenarios where, for instance, Bob's agent acts on behalf of an LLM with a specific request for an MCP server, ensuring that every link in that delegation is verifiable and authorized.

##### • "Need to Know" Access (Least Privilege)

Once an entity's identity is verified, the AI component (e.g., MCP server, AI model, agent) should only be granted access to the specific datasets it's allowed to access and the absolute minimum permissions required for its specific task. This principle prevents privilege escalation and limits the damage an attacker can do if a compromise occurs, safeguarding everything from your personal data to sensitive corporate assets. Crucially, this granular control helps prevent "operator error" misconfigurations, ensuring that even well-intentioned operators cannot inadvertently create security gaps by granting inappropriate access.

##### • Edge-Based Policy Enforcement for True Zero Trust

Security policies and access controls should be enforced at the edge, on the device, or within the workload itself, rather than relying on centralized servers or network choke points. Centralized policy enforcement, common in most security models, introduces a single point of failure which fundamentally compromises true Zero Trust. This distributed approach provides resilience, low latency, and ensures that policies are effective even if the central management plane is temporarily unavailable or compromised.

#### ✓ Vigilant Monitoring and Alerting

##### • Comprehensive Audit Trail

Every single action, request, and response that passes through the AI security solution should be meticulously logged and recorded. This granular audit trail is essential for forensic analysis, troubleshooting, and maintaining accountability, whether for debugging a personal project or meeting enterprise compliance standards.

##### • Anomaly Detection & Immediate Alerts

The system should continuously analyze these logs for suspicious patterns, unusual access attempts, or abnormal behavior. Immediate, automated alerts are crucial to detect and respond to potential attacks in real-time, helping you catch issues early.

#### ✓ Scalability and Ease of Deployment

The security solution should seamlessly integrate with your existing cloud and on-premise AI infrastructure without introducing significant complexity, performance bottlenecks, or requiring extensive network reconfigurations. It should be as easy to set up for a single personal project as it is for a complex enterprise environment.

#### ✓ Compliance and Audit Readiness

The solution should inherently support and provide verifiable evidence for compliance with relevant data protection regulations (e.g., GDPR, HIPAA, CCPA). This is crucial for enterprises and startups aiming for growth and trust. For hobbyists, it provides peace of mind that their personal data handling is secure.

## CHAPTER TWO SUMMARY

# Addressing the AI Attack Surface - Understanding Critical Vulnerabilities

Critical Vulnerabilities	Risks and Consequences	Solutions
Model-Level Vulnerabilities (Where MCP Provides Mitigation)	<ul style="list-style-type: none"><li>• Bad Advice &amp; Distorted Decisions</li><li>• Prompt Manipulation and Injection</li><li>• Sensitive Data Leakage &amp; Exfiltration (Model-Driven)</li><li>• Factual Inaccuracies</li></ul>	<ul style="list-style-type: none"><li>• Providing Richer Context and Sources of Truth</li><li>• Defining Clear Actions and Enabling Human Oversight</li><li>• Structured and Predictable Outputs</li><li>• Meticulous Input and Output Handling<ul style="list-style-type: none"><li>◦ Rigorous Input Validation</li><li>◦ Safe Responses</li></ul></li></ul>
Network-Level Vulnerabilities	<ul style="list-style-type: none"><li>• Digital Invisibility is Crucial</li><li>• The Risk of a Compromised MCP Server (Due to Network Exposure)<ul style="list-style-type: none"><li>◦ Execute Malicious Actions</li><li>◦ Feed Bad Data</li><li>◦ Act as a Backdoor for Lateral Movement</li></ul></li><li>• Unencrypted Traffic</li></ul>	<ul style="list-style-type: none"><li>• Network Invisibility / Elimination of Open Ports</li><li>• True Micro, or better yet, Picosegmentation</li><li>• Solid Communication and Data Protection<ul style="list-style-type: none"><li>◦ End-to-End Encryption</li><li>◦ Data Minimization in Transit</li></ul></li></ul>
Addressing Authentication and Chain of Trust Vulnerabilities	<ul style="list-style-type: none"><li>• Authentication &amp; Access Control Gaps</li><li>• Misconfigured Access / "operator Error" Vulnerability</li><li>• Excessive Privileges</li><li>• Unverified Connections / Spoofing</li></ul>	<ul style="list-style-type: none"><li>• Strict Identity and Access Control<ul style="list-style-type: none"><li>◦ Cryptographically Proven Chain of Trust</li><li>◦ "Need to Know" Access (Least Privilege)</li><li>◦ Edge-Based Policy Enforcement for True Zero Trust</li></ul></li><li>• Vigilant Monitoring and Alerting<ul style="list-style-type: none"><li>◦ Comprehensive Audit Trail</li><li>◦ Anomaly Detection &amp; Immediate Alerts</li></ul></li><li>• Scalability and Ease of Deployment</li><li>• Compliance and Audit Readiness</li></ul>

## CHAPTER THREE

# NoPorts: Your Blueprint for Invisible AI Security

NoPorts is specifically designed to provide the foundational security that enables your entire AI ecosystem to operate safely. It directly provides solutions for network-level vulnerabilities and authentication and chain of trust vulnerabilities, which are the critical attack vectors that could compromise even well-designed MCP implementations. By eliminating open ports and building security around cryptographically verified identities, NoPorts creates a truly "invisible infrastructure" for your AI and the rest of your network, regardless of its scale.

NoPorts' primary role is to ensure that your MCPs, models, and agents can only be accessed by authorized entities over a secure, invisible network, thus enabling the MCP to fulfill its role in mitigating model-level risks without external interference.

*What all unauthorized entities see:*



## CHAPTER THREE

### 3.1 How NoPorts Delivers Invisible AI Security

NoPorts fundamentally changes how your AI components and business systems connect. Instead of relying on vulnerable open ports, NoPorts establishes outbound-only connections that are cryptographically secure and identity-aware. This unique approach means:

A Zero Trust Model  
is Baked In

Digital Invisibility and Pico-  
segmentation for All Endpoints

Automated End-to-  
End Encryption

Granular, Identity-Based Access Control &  
Edge-Based Policy Enforcement

Simplified, Secure Deployment



#### **A Zero Trust Model is Baked In**

NoPorts enforces a "never trust, always verify" policy at the most fundamental level. Every connection request, whether from an AI model, an agent, or a human, is verified against its unique cryptographic identity before any access is granted. This applies whether you're securing a personal LLM or a critical enterprise AI.



#### **Digital Invisibility and Pico-segmentation for All Endpoints**

NoPorts makes your MCP servers, AI models, agents, and other internal services undetectable from the internet and even lateral movement attempts within your internal network. Since there are no open, listening ports, attackers cannot port scan, probe, or directly assault your AI instances. If they can't see it, they can't attack it. This is a game-changer for hobbyists avoiding complex firewall rules, startups needing agile security, and enterprises requiring robust protection. Additionally, by establishing identity-based, outbound-only connections, NoPorts inherently creates micro-segmentation down to individual processes (pico-segmentation). This ensures that each component can only connect to precisely what it needs, drastically limiting lateral movement and containing potential breaches.



## CHAPTER THREE

### 3.1 How NoPorts Delivers Invisible AI Security



#### **Automated End-to-End Encryption**

All communications facilitated by NoPorts are automatically encrypted end-to-end. This means data in transit, whether from an AI agent to an MCP, or from an MCP to a business system, is always protected from eavesdropping and tampering. No manual certs or VPNs needed.



#### **Granular, Identity-Based Access Control & Edge-Based Policy Enforcement**

With NoPorts, you define precisely which specific AI models, agents, or even human developers (based on their cryptographic identity, not just an IP address) can connect to which MCP server, and what actions they can perform. This enforces least privilege by default, drastically limiting potential damage from a compromise, protecting your personal projects or your company's most sensitive data. Unlike other solutions that rely on centralized policy servers, NoPorts enforces these rules directly at the edge, on the device or workload itself. This decentralized enforcement ensures true Zero Trust, removing single points of failure and providing immediate, resilient security decisions where they are needed most. This also prevents the "operator error" of exposing private MCP servers to public models by ensuring access is only granted to explicitly authorized entities.



#### **Simplified, Secure Deployment**

NoPorts streamlines the secure deployment and management of your AI infrastructure. You don't need complex firewall rules, VPNs, or intricate network segmentation schemes that are often brittle and costly to maintain. NoPorts integrates seamlessly with your existing cloud or on-premise environments, making security accessible for everyone from a solo developer to a large enterprise.

## CHAPTER THREE

### 3.2 How NoPorts Secures Your Critical AI Components At Any Scale

NoPorts directly addresses the network-level vulnerabilities and authentication and chain of trust vulnerabilities that plague AI deployments, thereby laying the essential groundwork for enabling MCP to fulfill its role in mitigating model-level risks:

#### Securing MCP with NoPorts

##### Unassailable Protection

NoPorts ensures your MCP servers are never exposed with open ports, making them impervious to common internet-based attacks like port scanning and direct network assaults. This fundamentally eliminates the most common entry points for attackers to then bypass MCP's internal safeguards and exploit model-level vulnerabilities.

##### End to End Encrypted Data Streams

All data and instructions flowing to and from your MCP servers are automatically encrypted, safeguarding sensitive AI prompts, responses, and contextual information.


##### Precise Access & Trust

Only authorized AI agents or applications with cryptographically verified identities can initiate connections to specific MCP instances. This prevents unauthorized access, ensuring only legitimate AI components are feeding or receiving information, and allowing the MCP to apply its own rigorous input/output handling safely.

##### Simplified Management

Whether you're a hobbyist running a few small instances or an enterprise managing hundreds, NoPorts simplifies the secure setup and management of MCP servers, reducing the burden on DevOps and security teams and accelerating AI deployment.

*MCP Data flow*



## CHAPTER THREE

# 3.2 How NoPorts Secures Your Critical AI Components (At Any Scale)

### Protecting AI Models and Inference Endpoints

Securely expose your AI inference APIs to authorized consumers (other AI systems, applications, or humans) without ever opening a port. NoPorts protects against model theft, unauthorized access, and ensures the integrity of inference requests and responses, critical for protecting your intellectual property, no matter the scale.

### Hardening AI Agents

Provide secure, identity-based access for autonomous AI agents to the tools, data sources, and resources they need to operate. Control agent-to-agent (A2A) communication securely, ensuring agents only perform authorized actions with verifiable cryptographic identities. This is vital for managing anything from a personal automation script to a fleet of enterprise AI bots.

### Safeguarding AI Datasets

Secure, audited access to your sensitive training and inference data stores. NoPorts ensures that only authorized AI components and personnel can access specific datasets, preventing unauthorized data exfiltration and maintaining data privacy, whether it's your personal research or your company's confidential customer data.

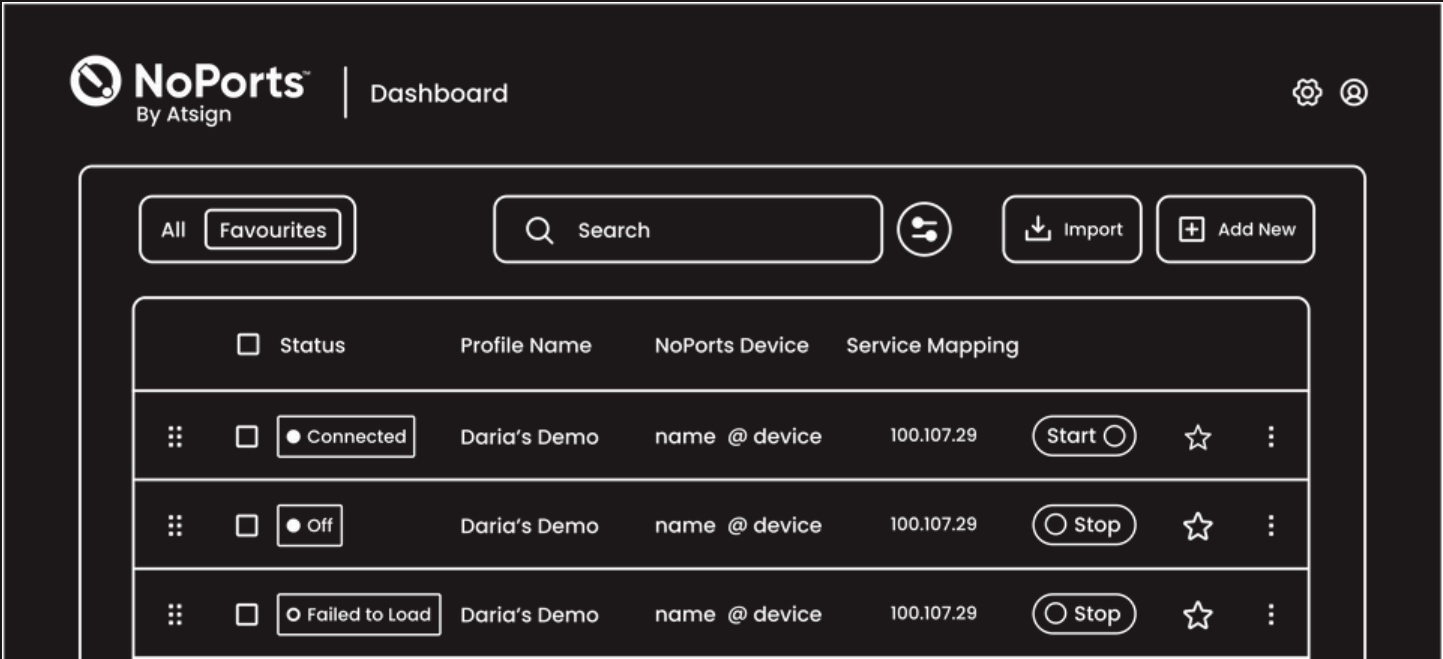
### Integrating with Business Systems

NoPorts enables secure, direct connections from your AI components (models, agents, MCP) to internal business APIs, databases, and other applications. This eliminates the need to open firewalls or expose sensitive internal services, dramatically reducing the attack surface of your critical personal or corporate applications.

# CHAPTER FOUR

## Putting Invisible AI Security into Practice: Getting Started with NoPorts

Implementing zero trust AI security doesn't have to be a daunting task. NoPorts is designed for rapid deployment and seamless integration, allowing you to quickly secure your AI infrastructure, no matter your current setup.



NoPorts Desktop Preview

CHAPTER FOUR

4.1 A Practical, Phased Approach to Getting Started



1. Identify Your AI Assets

Start by mapping out your most sensitive MCP servers, AI models, AI agents, core datasets, and the systems they interact with. Understand their current access patterns and identify potential vulnerabilities.



2. Pilot NoPorts (Easy Setup)

Begin with a small, AI component or an MCP server in a test environment. Experience firsthand how NoPorts eliminates open ports and establishes secure connections. Our intuitive setup makes this step accessible for hobbyists and quick for startups.



3. Expand Gradually

Once validated, progressively extend NoPorts to more critical AI infrastructure, validating each step. Our comprehensive documentation and responsive support team are here to guide you, whether you're building out a complex enterprise architecture or refining a personal project.

## CHAPTER FOUR

### 4.2 Beyond Security - Operational Benefits

While security is the primary driver, implementing NoPorts delivers significant operational advantages for everyone:



#### **Faster AI Deployment Cycles**

Eliminate complex network configurations, tedious firewall requests, and cumbersome VPN setups. Secure connections are established instantly based on identity, letting you get your AI models and agents running faster.



#### **Reduced Network Complexity**

Simplify your network architecture by removing the need for intricate segmentation rules and exposed ports. Focus on building amazing AI, not on managing network headaches.



#### **Simplified Access Management**

Manage access to your AI services based on clear, cryptographic identities, rather than fluctuating IP addresses or confusing network locations. This simplifies collaboration and auditing.



#### **Enhanced Auditability & Compliance**

NoPorts provides a clear, verifiable audit trail of all connections, helping you debug your AI's interactions or demonstrate adherence to industry compliance requirements.

# The Future of Secure AI is Invisible, and it's Here with NoPorts!

The promise of AI is immense, but its potential cannot be fully realized without robust security. Relying on outdated security models that expose your valuable AI infrastructure through open ports is no longer an option. The future of AI security is invisible infrastructure – a proactive, zero trust approach that eliminates the attack surface entirely, for every one and every scale.

NoPorts offers a unique, comprehensive, and easy-to-implement solution to these challenges. It empowers you to build, deploy, and scale your AI initiatives with confidence, knowing that your MCP servers, models, agents, datasets, and business systems are protected by an unassailable layer of security.

Don't let vulnerable open ports expose your critical AI assets. Embrace the power of invisible infrastructure today!

## Take the Next Step: Experience NoPorts with a Free Trial!

You now have a comprehensive understanding of the threats to your AI infrastructure and the characteristics of an ideal solution. Ready to put invisible security to work for your organization or your personal projects?

NoPorts eliminates open ports, making your AI infrastructure invisible to attackers while providing granular, identity-based access.

Start your **FREE 30-day trial** of NoPorts now at:

<https://my.noports.com/no-ports-invite/30dayfreetrial>

**Looking for specific instructions on  
securing your MCP servers with NoPorts?**

Our detailed documentation provides step-by-step guidance on how to install and configure NoPorts components for MCP: <https://docs.noports.com/use-cases/mcp>

Start building your invisible  
infrastructure today!

