Technical Brief

STAMUS NETWORKS

Clear NDR® and Model Context Protocol (MCP)

Strategic Enablers for the AI-Powered SOC

Introduction

Modern security operations face an impossible challenge: the volume and complexity of threats have outpaced human capacity to analyze them. Security teams are drowning in alerts, struggling with analyst shortages, and racing against attackers who move faster than manual investigation processes allow.

Organizations are turning to large language models (LLMs) to help them address alert fatigue and triage bottlenecks, expertise scaling, response and investigation speed, analysis of complex data patterns, and reporting consistency.

While public LLMs like Claude, Gemini, Mistral, or GPT-4 provide powerful general capabilities, many organizations are deploying custom domain-specific LLMs or local LLM instances to achieve greater control, privacy, and optimization for their security operations. The core drivers for creating a custom LLM include data sovereignty and privacy, domain-specific optimization, cost optimization, reduced response latency, and organization-specific policies.

When combined with an MCP-capable AI client and MCP integration to tools like Clear NDR, these custom LLMs create a powerful new security architecture that addresses critical enterprise concerns

What is Model Context Protocol (MCP)

Model Context Protocol (MCP) – a new capability of Clear NDR – empowers security practitioners to tap into rich network security intelligence generated by Clear NDR using natural language commands and queries. This can dramatically accelerate threat detection and response operations.

MCP is an open standard developed by Anthropic that enables AI clients and agents to securely connect external data sources and tools and large language models. Think of it as a standardized way for AI models to "plug into" different systems and access real-time information or perform actions beyond their training data.

In summary, MCP is the protocol that connects large language models (LLMs) to the resources inside Clear NDR and other modern security data sources

Where MCP Fits in the Tech Stack

Model Context Protocol (MCP) enables seamless integration between security tools and large language models through a client-server architecture.

In this architecture, MCP-enabled security systems implement MCP servers that expose their capabilities as standardized resources (read-only data such as network flows, threat detections, and forensic evidence) and tools (actionable functions like querying events, retrieving packet captures, or updating hunting rules). Network detection and response (NDR) systems, security event and incident management (SIEM) systems, endpoint detection and response (EDR) systems are examples of security systems that may implement MCP servers.

MCP-capable AI clients – ranging from simple conversational interfaces like Open WebUI to sophisticated autonomous agent frameworks like LangGraph or CrewAI – sit between the human analyst and their choice of large language model, whether that's a cloud-based service like Claude or GPT-4, or a locally-deployed model like Llama or Mistral running on-premises for complete data sovereignty.

These AI clients connect to multiple MCP servers simultaneously using a standardized JSON-RPC protocol, allowing the underlying LLM to access and analyze security data from diverse sources. This means a security analyst using a single MCP-capable AI client can ask one question and have the LLM automatically query an NDR for network intelligence, an EDR for endpoint data, their SIEM for log correlation, IAM for identity context, and threat intelligence platforms for IOC enrichment—all without switching interfaces or manually correlating findings.

The diagram below illustrates this architecture.

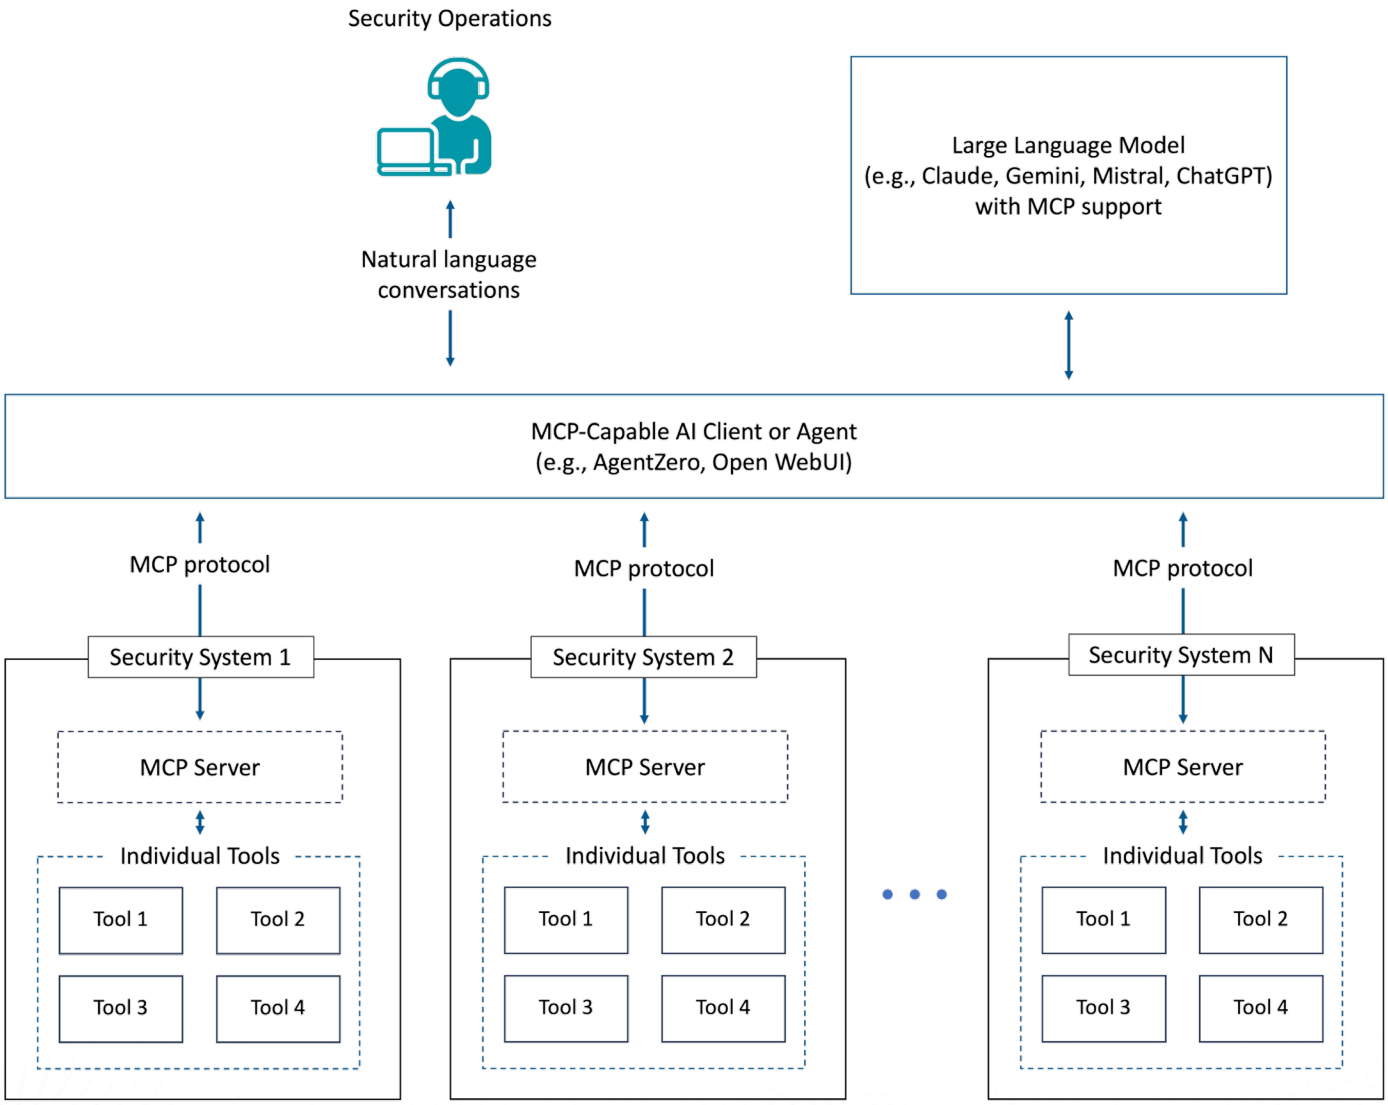

The architecture eliminates the need for custom integrations between each AI system and each security tool, instead providing a universal interface where security operations become truly unified.

Analysts can investigate threats conversationally, workflow systems can orchestrate complex incident response playbooks across multiple tools, and autonomous agents can conduct 24/7 threat hunting, all through a single, standardized protocol. This makes security operations both more efficient and more effective while giving organizations complete flexibility in their choice of AI infrastructure.

Natural Language Investigation of Lateral Movement

To illustrate how Clear NDR's MCP integration makes security operations more accessible and efficient, consider a straightforward but common scenario: a security practitioner needs to check for lateral movement threats in their network.

While Clear NDR's expert-level user interface provides powerful capabilities for experienced analysts to hunt through network data, filter threat families, examine host behaviors, and drill down into specific alerts, MCP integration offers an alternative approach that's particularly valuable for junior analysts, analysts working under time pressure, or teams wanting to democratize security expertise across their organization.

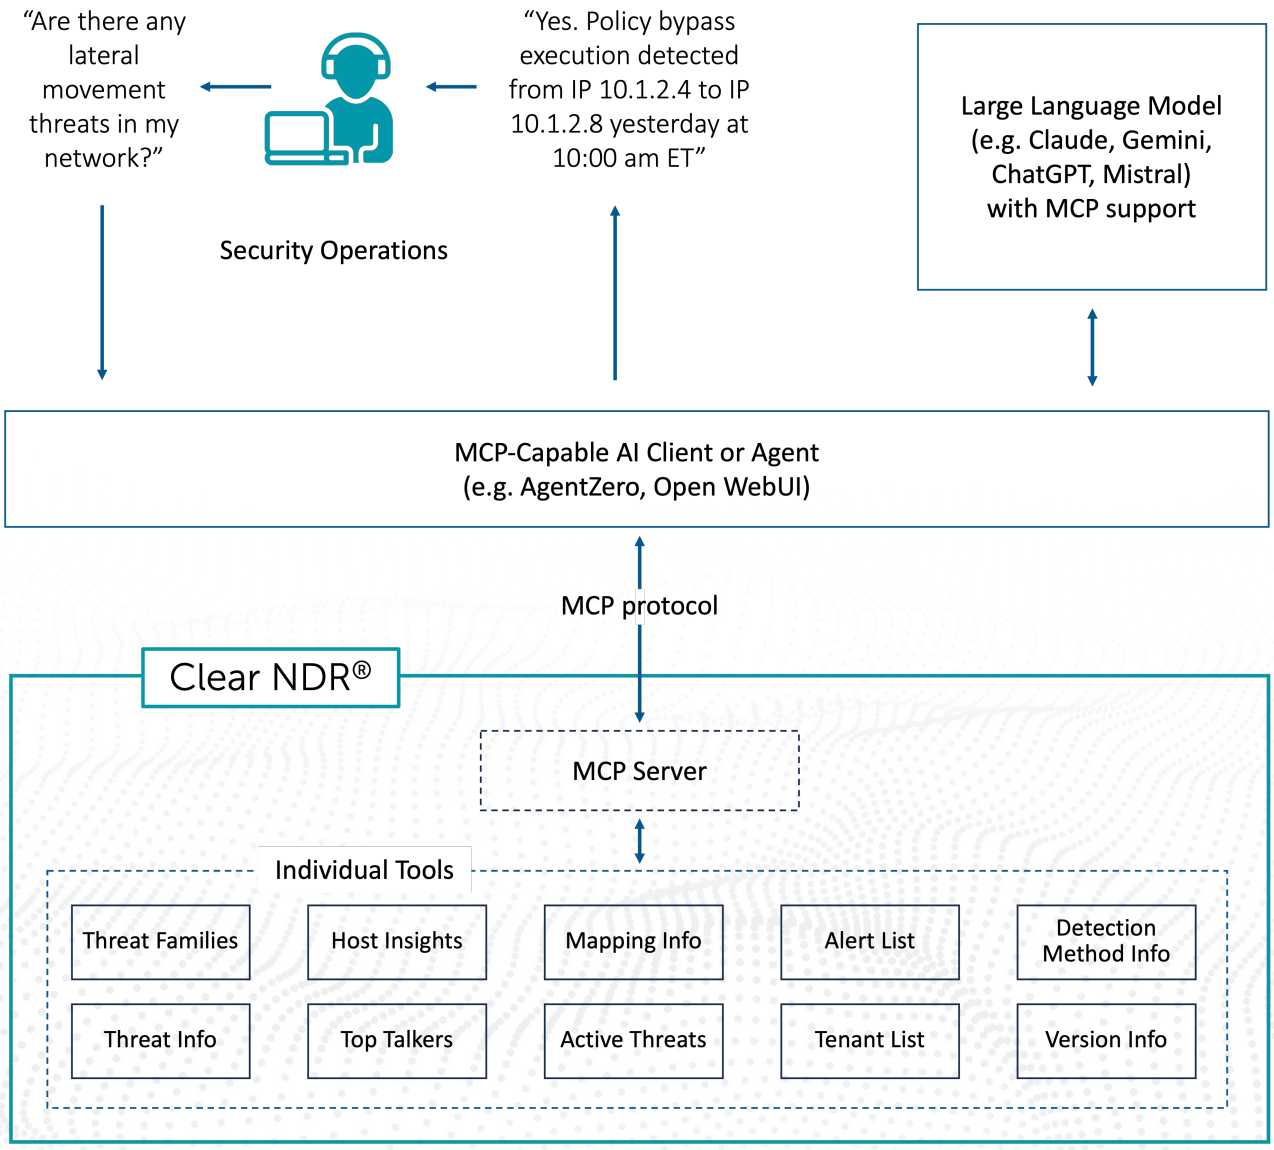

Instead of navigating through Clear NDR's interface – selecting the appropriate views, configuring filters for lateral movement indicators, reviewing threat family classifications, and manually correlating host details – the practitioner can simply use natural language to ask an MCP-capable AI client: "Are there any lateral movement threats in my network?" The AI client then orchestrates the necessary queries to Clear NDR's MCP server, intelligently calling the appropriate tools (Active Threats, Threat Info, Host Details) to gather comprehensive information, and synthesizes the findings into a clear, actionable answer.

This natural language approach doesn't replace Clear NDR's powerful interface for deep investigations; rather, it complements it by making Clear NDR's rich network intelligence instantly accessible to a broader range of users and use cases, while dramatically accelerating initial triage and threat assessment workflows.

The diagram on the previous page illustrates the specific elements – including the tools exposed by the MCP server in Clear NDR – used in this example.

Investigation Example: The Query

A security practitioner using an MCP-capable AI client asks:

Are there any lateral movement threats in my network?

How it Works

Step 1: The MCP-Capable AI Client Receives and Routes the Query

The security practitioner asks their question through the MCP-capable AI client interface. The client receives this natural language query and passes it to the large language model (LLM), which then interprets that "lateral movement threats" means checking for suspicious internal network activity, particularly commands executed between hosts that might indicate an attacker moving through the network. The LLM then sees that the Clear NDR MCP server has a tool for this information and asks the client to query the tool.

Step 2: The AI Client Queries Clear NDR's "Active Threats" Tool

Based on the LLM's analysis, the MCP-capable AI client uses the MCP protocol to ask Clear NDR's MCP server for any recent threats related to lateral movement in the last 24 hours. The AI client calls Clear NDR's "Active Threats" tool to learn the specifics. Clear NDR

provides details showing that the source IP (10.1.2.4) executed commands on the destination IP (10.1.2.8) at 10:00 AM ET via SMB protocol, bypassing security policies. Clear NDR identifies a high-severity "Policy Bypass Execution" threat that occurred yesterday.

Step 3: The AI Client Enriches with Host Context

To provide a complete picture, the AI client uses Clear NDR's "Host Details" tool to understand the systems involved. Clear NDR reveals that the source is a finance department workstation showing elevated privileges, and the destination is a critical corporate file server.

Step 4: The LLM Synthesizes Everything into a Clear Answer

The MCP-capable AI client sends the data it has received as well as the initial prompt to the LLM. The LLM then analyzes all the information retrieved through these three MCP tool calls and constructs a natural language response, which the AI client then presents to the security practitioner.

The Response

The following is an example response offered by the system (actual results will vary with the particular LLM selection):

> *Yes. A policy bypass execution was detected from IP 10.1.2.4 to IP 10.1.2.8 today at 10:00 AM ET.*
>
> *This represents a high-confidence lateral movement threat where a workstation in the finance department (10.1.2.4) executed commands on a critical corporate file server (10.1.2.8) by bypassing security policies via SMB protocol. This activity aligns with MITRE ATT&CK tactic TA0008 (Lateral Movement).*
>
> *The source host shows signs of elevated privileges, and Clear NDR has captured complete forensic evidence including packet captures and protocol transaction logs. I recommend immediate investigation of the finance workstation for potential compromise.*

This architecture eliminates the need for custom integrations between each AI system and each security tool, instead providing a universal interface where any MCP-capable AI client can immediately leverage Clear NDR's high-fidelity network intelligence alongside other MCP-enabled security tools like EDR platforms, identity providers, and threat intelligence feeds.

The Power of Multi-Tool Orchestration

While this example shows Clear NDR alone, the same MCP-capable AI client could simultaneously query multiple security systems in response to a single question:

Clear NDR MCP Server (using tools like Active Threats, Host Details, Threat Info):

- Network evidence of lateral movement
- Host behavioral analysis
- Forensic packet captures

SentinelOne EDR MCP Server:

- Process execution details on both hosts
- Memory analysis for credential theft
- Persistence mechanisms

Okta Identity MCP Server:

- Authentication events for accounts on both systems
- Privilege escalation attempts
- Unusual access patterns

In this case, the analyst asks ONE question, and the MCP-capable AI client orchestrates queries across ALL systems, correlating findings into a single, comprehensive answer.

Clear NDR's MCP Tools Enable Intelligent Investigations

Currently, Clear NDR exposes ten MCP tools that allow the LLM to construct investigation strategies tailored to each query. Each query triggers a different combination of these tools based on what the LLM determines is needed to answer the question accurately. The table below outlines the ten tools available in the initial release (Update 42 or U42)

Official	Simplified Shorthand
clear_ndr.alert_list: Retrieves a paginated list of alerts.	Alert List
clear_ndr.family_info: Retrieves information about one or more threat families.	Threat Families
clear_ndr.get_ip_details: Retrieves detailed Host Insights for a specified IP address.	Host Insights
clear_ndr.mapping_info: Provides information about common fields for Lucene queries.	Mapping Info
clear_ndr.rules: Gets rule information from the SID (signature_id field in the alert event).	Detection Method Info
clear_ndr.talkers: Retrieves a list of the most active "talkers" (hosts) on the network.	Top Talkers
clear_ndr.threat_info: Retrieves information about one or more threats.	Threat Info
clear_ndr.threat_list: Retrieves a list of threats (Declarations of Compromise) observed within your network.	Active Threats
clear_ndr.tenants: Lists available tenants.	Tenant List
clear_ndr.version: Retrieves the current version information for Clear NDR.	Version Info

The Complementary Value Proposition

Clear NDR's MCP integration doesn't replace the platform's powerful expert interface—it enhances it by providing an alternative interaction model optimized for different use cases and user profiles.

Organizations gain the flexibility to choose the right approach for each situation: natural language queries through MCP for rapid assessment and broad accessibility, or the expert interface for deep forensic analysis and advanced threat hunting.

Understanding when and how to leverage each approach enables security teams to maximize both the speed and depth of their operations.

Using Clear NDR's User Interface

Clear NDR's native user interface is purpose-built for security professionals who need granular control and comprehensive visibility into network threats. The interface enables analysts to construct sophisticated queries, examine complex attack patterns, and navigate through layers of network telemetry with precision based on their expertise and intuition.

- Powerful filtering and visualization capabilities
- Direct control over queries and investigation paths
- Deep access to all network telemetry and evidence
- Ideal for complex hunting and forensic analysis
- User profile: Experienced security analysts

Using MCP-Capable AI Client with Clear NDR

MCP integration brings Clear NDR's powerful network intelligence to security practitioners through conversational interaction that eliminates the learning curve associated with expert security tools. Complex, multi-stage queries that would require navigating through several interface screens happen transparently in response to a single natural language question, with automatic correlation across Clear NDR, EDR, identity providers, and threat intelligence platforms.

- Natural language queries accessible to any skill level
- Automatic orchestration of investigation steps
- Instant synthesis of findings from multiple tools
- Rapid triage and initial threat assessment
- User profile: All security team members, especially during rapid response

Best Practice: Use Both

The most effective security operations leverage both approaches strategically, recognizing that MCP integration and Clear NDR's expert interface serve complementary purposes rather than competing ones.

- MCP for rapid triage, initial assessment, and cross-tool correlation
- Clear NDR UI for deep investigation, advanced hunting, and detailed forensics
- The combination maximizes both speed and depth of security operations

This is how Model Context Protocol transforms security operations: not by replacing sophisticated tools like Clear NDR's user interface, but by making their rich intelligence conversationally accessible, enabling faster triage, democratizing security expertise across teams, and seamlessly correlating findings across multiple security platforms through natural language interaction.

Clear NDR and MCP: Transformative for the SOC

For organizations deploying Clear NDR, MCP integration represents a strategic inflection point. Clear NDR's strength has always been its ability to extract rich, high-fidelity network telemetry and provide transparent, evidence-based threat detection through Declarations of Compromise.

By implementing a native MCP server that exposes this network intelligence through standardized tools, Clear NDR becomes the essential network data foundation for AI-powered security operations – whether organizations use cloud-based LLMs like Claude and GPT-4, or deploy private, domain-specific models for complete data sovereignty.

This approach amplifies the AI investments organizations have already made in platforms like SentinelOne AI SIEM, Splunk, Google Chronicle, CrowdStrike Falcon Next-Gen SIEM, or custom LLM applications, while positioning Clear NDR as the network intelligence multiplier that makes any AI security system more accurate, complete, and effective.

Experience Clear NDR's MCP Integration Yourself

Ready to see how Model Context Protocol can transform your security operations? Stamus Networks invites you to experience firsthand how Clear NDR's native MCP server integration delivers conversational access to high-fidelity network intelligence while seamlessly connecting with your existing security tools.

Request a Live Demo

See Clear NDR's MCP integration in action with a personalized demonstration tailored to your environment. Our security experts will show you how natural language queries can replace complex, multi-tool investigations—whether you're using cloud-based LLMs, deploying private models for data sovereignty, or integrating with AI-powered SIEM platforms like SentinelOne or CrowdStrike Falcon.

Schedule a Complementary On-Site Evaluation

Experience the power of MCP-enabled security operations in your own environment. Stamus Networks offers complementary on-site evaluations where we'll deploy Clear NDR in your network, configure MCP integration with your chosen AI client, and demonstrate real-world threat detection and investigation workflows using your actual network telemetry. See for yourself how Clear NDR's rich network intelligence becomes the foundation for truly autonomous, AI-powered security operations

Contact us Today

Visit: www.stamus-networks.com/mcp-demo
Email: contact@stamus-networks.com

Discover why leading organizations trust Clear NDR as their network intelligence foundation for AI-powered security operations.

STAMUS® NETWORKS

5 Avenue Ingres
75016 Paris
France

450 E 96th St. Suite 500
Indianapolis, IN 46240
United States

contact@stamus-networks.com
www.stamus-networks.com

