



White Paper: The Strategic Value of JProfiler in Agentic Development

Enabling AI Coding Agents to Diagnose and Resolve JVM Performance Issues
Autonomously

Executive Summary

AI coding agents are transforming software development, taking on tasks from writing code to debugging and refactoring. Yet these agents share a critical blind spot: they cannot observe what happens at runtime. Performance bottlenecks, memory leaks, slow database queries, and inefficient HTTP calls are invisible to an agent that can only read and write source code.

JProfiler's MCP server eliminates this blind spot. By exposing JProfiler's profiling capabilities through the Model Context Protocol (MCP), AI agents gain direct access to CPU hotspots, call trees, database query analysis, and HTTP connection profiling. The MCP server is specifically designed for agent consumption: data is condensed and structured, the workflow is guided so agents follow it reliably, and the tools integrate seamlessly into agentic development environments.




This white paper explains how JProfiler's MCP server turns AI coding agents from code editors into full-stack performance engineers, and why this capability represents a strategic advantage for organizations adopting agentic development.

Introduction

Agentic development is rapidly becoming a core part of how software organizations build and maintain their applications. AI coding agents can read source code, write new implementations, run tests, and iterate on feedback. For many routine tasks, they already operate with remarkable autonomy.

However, there is a fundamental gap in what these agents can do today. While they excel at working with static artifacts like source code, configuration files, and test outputs, **they are blind to runtime behavior**. Performance problems, memory leaks, and inefficient database interactions only manifest when the application is running. An agent that cannot observe runtime data is limited to guesswork when addressing these issues.

JProfiler's MCP server bridges this gap by giving AI agents direct access to profiling data through the Model Context Protocol. This white paper examines three dimensions of this capability:

-  Runtime Observability for AI Agents
-  Codebase Navigation via Call Trees
-  Agent-Optimized Data and Workflow

We will explore how each of these dimensions creates tangible value, and why equipping AI agents with profiling tools represents a strategic investment that pays for itself many times over. The discussion is intended for technical leaders and decision-makers evaluating how to maximize the return on their organization's adoption of AI-assisted development.

Agents Need Runtime Observability

The Blind Spot of AI Coding Agents

Today's AI coding agents are powerful assistants for reading, writing, and refactoring code. But when it comes to performance and memory problems, they are working with a severe handicap: they cannot see what the application actually does at runtime.

Consider a developer asked to fix a slow API endpoint. A human developer would profile the application, identify that 80% of the response time is spent in a particular database query, trace the query back to a missing index, and apply the fix. Without profiling data, an AI agent facing the same problem can only read the source code and guess. It might restructure the code, add caching where none is needed, or optimize a method that accounts for a negligible fraction of the total time.

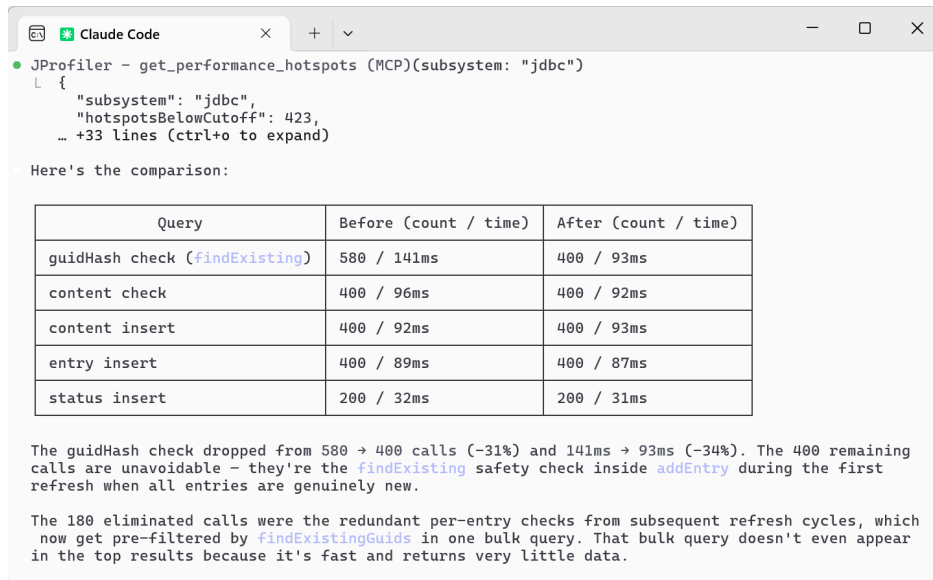
The analogy is straightforward: **asking an agent to fix a performance problem without runtime data is like asking a doctor to diagnose a patient without being allowed to examine them**. The agent needs eyes and ears into the running application.

What JProfiler's MCP Server Provides

JProfiler's MCP server gives AI agents direct access to the same profiling data that human developers rely on. Through the Model Context Protocol, agents can:

- Profile a running JVM process or load an existing snapshot for analysis
- Retrieve CPU hotspots that show exactly which methods consume the most execution time

- Analyze database query performance across JDBC, JPA, and MongoDB subsystems, identifying slow queries and their frequencies
- Inspect HTTP client calls to detect excessive or slow outgoing requests
- Drill down into back traces to understand the full chain of callers leading to each hotspot



This transforms the agent from a code-level tool into one that understands the runtime characteristics of the application. Instead of guessing which code paths matter, the agent can **see precisely where time is being spent** and direct its efforts accordingly.

Beyond CPU: Subsystem-Level Insight

Performance problems in modern Java applications rarely stem from raw CPU inefficiency alone. They arise in the interaction between the application and external systems: databases, HTTP services, message queues, and file systems. JProfiler's MCP server exposes dedicated subsystems for these areas, giving agents visibility into:

- **JDBC and JPA:** Slow SQL statements, N+1 query patterns, and connection pool exhaustion
- **HTTP Client:** Excessive outgoing calls, slow third-party services, and redundant requests
- **MongoDB:** Unoptimized queries, missing indexes, and collection scan patterns

Each subsystem provides hotspot-level data with execution counts, timing, and stack traces that tie the problem directly to the responsible code. This enables agents to identify not just that something is slow, but why it is slow, how frequently the problem occurs, and exactly where in the codebase to apply the fix.

Profiling Live Applications

A particularly powerful capability is that agents can attach to running JVM processes and profile them directly. Rather than relying on saved snapshots or log files, the agent can observe the application under its actual workload, in its actual environment. This is critical because many performance problems only manifest under specific conditions that are difficult to reproduce locally.

JProfiler's MCP server also supports attaching to JVMs running in local Docker containers. The agent can profile the application exactly as it runs in production-like conditions, without any container reconfiguration.

Navigating Complex Codebases

The Challenge of Scale

Enterprise Java applications routinely span hundreds of thousands or millions of lines of code, organized across thousands of classes and modules. When a performance issue is reported, the first and often hardest question is: where in this vast codebase should you look?

Human developers rely on experience, intuition, and profiling tools to narrow their focus. AI agents, lacking runtime context, must search through code using text patterns, dependency graphs, and static analysis. In a large codebase, this approach produces an overwhelming number of candidates. The agent may spend significant effort exploring code paths that turn out to be irrelevant to the actual performance problem.

A Performance-Weighted Map of the Codebase

JProfiler's call tree and hotspot views give agents something far more effective: a map of the codebase weighted by actual runtime behavior. Instead of searching through source code to find what might be relevant, the agent receives a ranked list of the methods that actually consume the most time.

```
jprofiler - expand_performance_hotspot (MCP)(id: 4)
└─ {"subsystem": "cpu", "wallTime": false, "method": "org.gjt.jclasslib.io.SourceDataInput.readByte
   eArray(int)", "micros": 408475, "callers": [{"method": "org.gjt.jclasslib.structures.constants.
   ConstantUtf8Info.readUTF", "micros": 359144, "percent": 87.92, "callers": [{"method": "org.gjt.jc
   ... +36 lines (ctrl+o to expand)

Here's the clear picture. The jclasslib reading time breakdown (from the total ~1353ms in
Structure.read):
```

Hotspot	Time	% of read
<code>ConstantUtf8Info.readUTF</code>	440ms	34%
└ of which <code>readByteArray</code> (inside <code>readUTF</code>)	359ms	27%
└ of which <code>decodeToString</code> (inside <code>readUTF</code>)	282ms	21%
<code>SourceDataInput.readByteArray</code> (CodeAttribute bytecode)	49ms	4%
<code>ConstantType.read</code> (dispatch + other constants)	243ms	18%
<code>AttributeContainer.readAttributes</code> + <code>create</code>	356ms	26%

```
UTF-8 decoding dominates - readUTF accounts for ~34% of total read time. Within that,
readByteArray (the byte[] allocation + copy from the source buffer) takes 359ms and decodeToString
takes 282ms. These overlap (readByteArray is called from within readUTF), but the key point is
clear: allocating the byte array and converting it to a String are the two biggest single costs.
```

This is a fundamentally different approach to codebase navigation. Rather than asking "which code could cause this problem?", the agent asks "which code **does** cause this problem?" The answer comes directly from the profiling data.

Drill-Down Navigation

Once an agent identifies a hotspot, JProfiler's MCP server enables systematic drill-down through the call hierarchy. The agent can expand any hotspot to reveal its callers, understanding the full execution path that leads to the bottleneck. This back-trace capability is essential because:

- A slow method may be called from multiple locations, but only one call site may be responsible for the majority of the time spent
- The fix often needs to be applied not at the hotspot itself but at a higher level in the call chain, such as reducing the number of calls or adding caching at the right layer

Focusing on What Matters

Package filtering allows agents to focus their analysis on application code while excluding framework internals. When an agent profiles a Spring Boot application, it does not need to wade through Spring framework methods to find the application-level code that needs optimization. The MCP server's filtering capabilities let the agent **cut directly to the relevant code paths**, reducing the search space from thousands of methods to a manageable set of actionable targets.

For organizations with large codebases, this capability alone justifies the investment. An agent that can identify the exact methods responsible for a performance regression in minutes would otherwise require hours of manual investigation by a senior developer.

Optimized For Agent Consumption

The Problem with Raw Profiling Data

Traditional profiling tools present data through rich graphical interfaces designed for human consumption. When agents are given access to raw profiling data, such as JFR recordings or the output of simple JVM profilers, the result is an overwhelming volume of method names, timings, and call relationships. Agents processing this much data lose focus on what matters, hallucinate connections between unrelated data points, or simply ignore the output entirely. Feeding large amounts of unprocessed profiling data into an AI agent is not just ineffective, it is expensive: every token of irrelevant data adds to the cost of the agent interaction without improving the outcome.

Condensed, Structured Data Presentation

JProfiler's MCP server is specifically engineered to present profiling data in a format that AI agents can process effectively. Several mechanisms work together to achieve this:

- **Intelligent cutoff filtering**

Only hotspots that account for a significant percentage of total execution time are included in results. Minor contributors are filtered out automatically, keeping the agent focused on impactful findings.

- **Maximum node counts**

Results are bounded to prevent data overload. The agent receives the most important data points without being overwhelmed by the long tail of insignificant entries.

- **Below-cutoff counters**

The agent is informed how many items were filtered out, so it knows whether the visible results represent the complete picture or whether further investigation with refined filters might be warranted.

- **Hierarchical structure**

Data is organized in a tree structure that mirrors the natural hierarchy of method calls, allowing agents to navigate from high-level summaries to specific details systematically.

Systematic Navigation

The MCP server is designed for **systematic, multi-step investigations**. The agent can retrieve the top hotspots, select the most promising one, expand it to see its callers, and continue drilling down. At each step, the data is scoped and structured so the agent maintains a clear picture of where it is in the analysis and what to explore next.

Comparison Mode for Validating Fixes

When an agent applies a performance fix, it needs to verify that the fix actually worked. JProfiler's MCP server supports loading a baseline snapshot alongside the current data, enabling direct before-and-after comparisons. The agent can see which hotspots improved, which regressed, and which are new, providing concrete evidence that a code change had the intended effect.

A Guided Workflow Agents Actually Follow

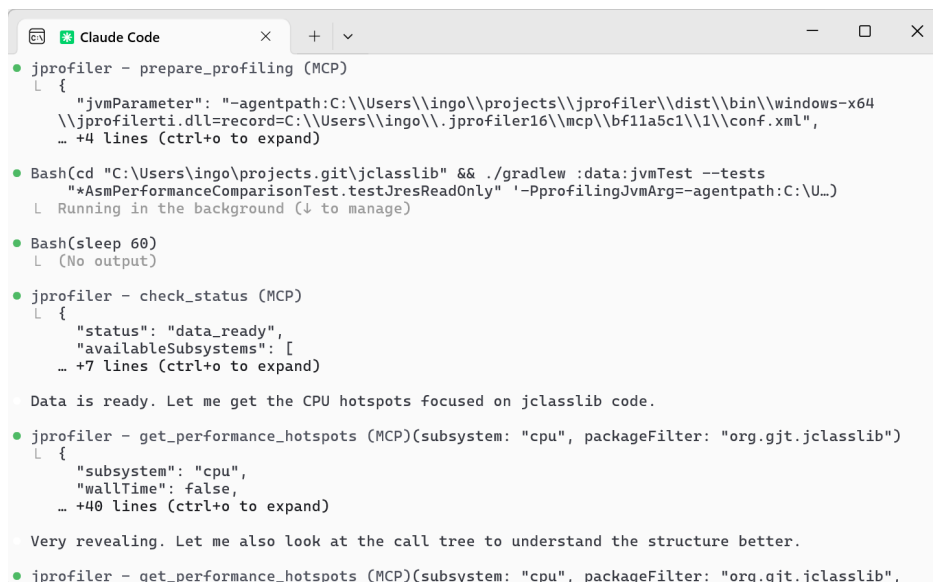
The MCP Server Adoption Problem

The Model Context Protocol has enabled a rapidly growing ecosystem of tool integrations for AI agents. However, a common challenge has emerged: many MCP servers expose flat lists of tools without guidance on how to use them. The result is that agents often call tools in the wrong order, skip essential setup steps, or fail to follow up on initial results. In practice, agents frequently ignore MCP servers whose tools they do not understand how to sequence effectively.

This is not a limitation of the agents themselves but of the tool design. **An MCP server that does not guide its users will not be used effectively**, regardless of how powerful its underlying capabilities are.

A Directed Workflow

JProfiler's MCP server addresses this with a carefully designed workflow that mirrors how expert human profiler users work: set up profiling, discover where time is being spent, then drill down into the most significant hotspots. The agent is guided through each phase, from initial setup to detailed root-cause analysis.



```
Claude Code
• jprofiler - prepare_profiling (MCP)
  L {
    "jvmParameter": "-agentpath:C:\\Users\\ingo\\projects\\jprofiler\\dist\\bin\\windows-x64
      \\jprofilerti.dll=record=C:\\Users\\ingo\\.jprofiler16\\mcp\\bf11a5c1\\1\\conf.xml",
    ... +4 lines (ctrl+o to expand)

• Bash(cd "C:\\Users\\ingo\\projects.git\\jclasslib" && ./gradlew :data:jvmTest --tests
  "*AsmPerformanceComparisonTest.testJresReadOnly" '-PprofilingJvmArg=-agentpath:C:\\U...)
  L Running in the background (↓ to manage)

• Bash(sleep 60)
  L (No output)

• jprofiler - check_status (MCP)
  L {
    "status": "data_ready",
    "availableSubsystems": [
    ... +7 lines (ctrl+o to expand)

Data is ready. Let me get the CPU hotspots focused on jclasslib code.

• jprofiler - get_performance_hotspots (MCP)(subsystem: "cpu", packageFilter: "org.gjt.jclasslib")
  L {
    "subsystem": "cpu",
    "wallTime": false,
    ... +40 lines (ctrl+o to expand)

Very revealing. Let me also look at the call tree to understand the structure better.

• jprofiler - get_performance_hotspots (MCP)(subsystem: "cpu", packageFilter: "org.gjt.jclasslib",
```

The key innovation is that every tool response tells the agent what to do next. Because these suggestions are part of the response data itself, the agent sees them as it processes the results. This is fundamentally different from static documentation that agents typically skip. In practice, agents follow the workflow reliably and consistently, producing useful profiling investigations without human intervention.

Zero-Configuration Setup

A significant barrier to tool adoption in agentic workflows is complex setup requirements. JProfiler's MCP server eliminates this barrier through automatic handling of installation and licensing:

- JProfiler is downloaded and installed automatically on first use, with no external dependencies or manual configuration steps
- License management is handled through the MCP protocol's input elicitation, allowing agents to guide users through license entry or evaluation mode without leaving the development environment

This means that from the moment a development team adds JProfiler's MCP server to their agent configuration, **profiling capabilities are immediately available** without any setup overhead.

ROI And Competitive Advantage

Developer Productivity Multiplier

Performance investigations are among the most time-consuming tasks in software development. They require deep expertise, familiarity with the codebase, and often pull senior engineers away from feature work for hours or days. When an AI agent can autonomously profile an application, identify the root cause of a performance issue, and propose a targeted fix, the impact on team productivity is substantial.

Instead of a senior developer spending half a day investigating why a service's response time increased, the agent handles the investigation and presents a diagnosis with supporting profiling data. The developer reviews the findings and approves the fix in minutes. This **shifts senior engineers from fire-fighting to high-value work** like architecture, feature development, and mentoring.

Faster Incident Response

Performance regressions in production are costly. The time between detecting a regression and identifying its root cause is often the longest and most expensive part of the incident lifecycle. AI agents equipped with profiling tools can serve as first responders: they can profile the affected application, identify the hotspot, compare it against a baseline snapshot, and produce a diagnosis before a human engineer has even begun their investigation.

Even when the agent's fix requires human review before deployment, the diagnostic work it performs reduces incident resolution time significantly. With the [average cost of IT downtime estimated at over \\$5,600 per minute^{\(1\)}](#), this acceleration alone can justify the entire tooling investment.

Shift-Left Performance Engineering

Traditionally, performance problems are discovered late: in staging, during load testing, or worst of all, in production. Research on the cost of software defects consistently shows that

⁽¹⁾ <https://www.the-future-of-commerce.com/2023/10/19/the-cost-of-downtime/>

issues found in production are an order of magnitude more expensive to resolve than those caught during development. Agentic profiling enables a shift-left approach where performance analysis happens earlier in the development cycle.

An agent can profile code changes during development or as part of a code review workflow, catching regressions before they are merged. This prevents the costly cycle of deploying, discovering, diagnosing, and fixing performance issues that characterizes traditional development workflows.

Reduced Context-Switching Costs

Research has shown⁽²⁾ that after an interruption, it takes an average of 23 minutes to return to the original task. When a senior developer is interrupted to diagnose a performance issue, the cost extends beyond the investigation time itself. The developer loses focus on their current task and needs considerable time to regain productive flow afterward.

By delegating performance investigations to AI agents, organizations **protect their developers' focus time**. The agent works in the background, and the developer is only involved at the final review stage.

Cost Perspective

The effectiveness of AI agents is directly proportional to the quality of the tools they can access. A single prevented production incident or a single optimized cloud workload typically saves more than the cost of JProfiler licenses for an entire team for multiple years. Teams that equip their agents with runtime observability will ship faster, with fewer performance regressions, and with more efficient use of cloud resources.

Conclusion

AI coding agents are capable of remarkable autonomy in reading, writing, and refactoring code. However, an agent that cannot observe runtime behavior is fundamentally limited in the problems it can solve.

JProfiler's MCP server removes this limitation. It gives AI agents the ability to profile JVM applications, navigate complex codebases through performance-weighted call trees, and identify the root causes of performance and memory problems. The data is presented in formats that agents can process effectively, and the guided workflow ensures that agents use the tools reliably and systematically.

For development organizations, this translates into concrete benefits: faster performance investigations, reduced incident response times, fewer production regressions, and more efficient use of senior engineering talent. The cost of JProfiler is minimal compared to the compounding returns of agents that can autonomously diagnose and resolve performance issues.

Equipping your AI agents with JProfiler's MCP server is a strategic decision that positions your organization at the forefront of agentic development, turning your agents from capable code assistants into **full-stack performance engineers**.

⁽²⁾ <https://dl.acm.org/doi/10.1145/1357054.1357072>