

QLCODER: A QUERY SYNTHESIZER FOR STATIC ANALYSIS OF SECURITY VULNERABILITIES

Claire Wang

University of Pennsylvania

cdwang@seas.upenn.edu

Ziyang Li

John Hopkins University

ziyang@cs.jhu.edu

Saikat Dutta

Cornell University

saikatd@cornell.edu

Mayur Naik

University of Pennsylvania

mhnaik@upenn.edu

ABSTRACT

Static analysis tools provide a powerful means to detect security vulnerabilities by specifying *queries* that encode vulnerable code patterns. However, writing such queries is challenging and requires diverse expertise in security and program analysis. To address this challenge, we present *QLCoder* – an agentic framework that automatically synthesizes queries in CodeQL, a powerful static analysis engine, directly from a given CVE metadata. QLCoder embeds an LLM in a synthesis loop with execution feedback, while constraining its reasoning using a custom MCP interface that allows structured interaction with a Language Server Protocol (for syntax guidance) and a RAG database (for semantic retrieval of queries and documentation). This approach allows QLCoder to generate syntactically and semantically valid security queries. We evaluate QLCoder on 176 existing CVEs across 111 Java projects. Building upon the Claude Code agent framework, QLCoder synthesizes correct queries that detect the CVE in the vulnerable but not in the patched versions for 53.4% of CVEs. In comparison, using only Claude Code synthesizes 10% correct queries. Our generated queries achieve an F1 score of 0.7. In comparison, the general query suites in IRIS (a recent LLM-assisted static analyzer) and CodeQL only achieve F1 scores of 0.048 and 0.073, highlighting the benefit of QLCoder’s specialized synthesized queries.

1 INTRODUCTION

Security vulnerabilities continue to grow at an unprecedented rate, with over 40,000 Common Vulnerabilities and Exposures (CVEs) reported in 2024, 28,961 CVEs reported in 2023, and 25,059 CVEs reported in 2022 (CVE, 2025). Static analysis, a technique to analyze programs without executing them, is a common way of detecting vulnerabilities. Static analysis tools such as CodeQL (GitHub, 2025a), Semgrep (Semgrep, 2023), and Infer (Meta, 2025) are widely used in industry. They provide domain-specific languages that allow specifying vulnerability patterns as queries. Such queries can be executed over structured representations of code, such as abstract syntax trees, to detect potential security vulnerabilities.

Despite their widespread use, existing query suites of static analysis tools are severely limited in coverage of vulnerabilities and precision. Extending them is difficult even for experts, as it requires knowledge of unfamiliar query languages, program analysis concepts, and security expertise. Incorrect queries can produce false alarms or miss bugs, limiting the effectiveness of static analysis. Correct queries can enable reliable detection of real vulnerabilities, supporting diverse use-cases such as regression testing, variant analysis, and patch validation, among others (Figure 1).

Meanwhile, CVE databases (MITRE, 2025; NIST, 2025; GitHub, 2025b) provide rich information about security vulnerabilities, including natural language descriptions of vulnerability patterns and records of buggy and patched versions of the affected software repositories. This resource remains largely untapped in the automated construction of static analysis queries. Recent advances in LLMs, particularly in code understanding and generation, open up the possibility of leveraging this informa-

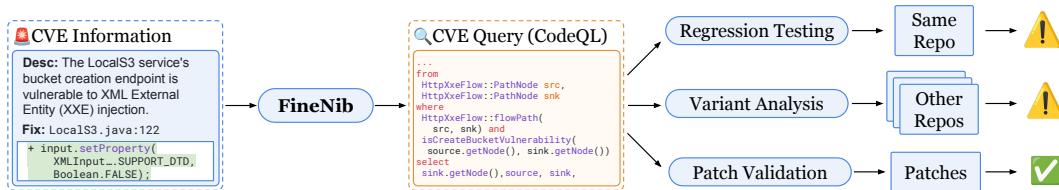


Figure 1: A CodeQL query capturing a vulnerability pattern is synthesized by QLCoder from an existing CVE and subsequently reused for regression testing, variant analysis, or patch validation.

tion to automatically synthesize queries from CVE descriptions, thereby bridging the gap between vulnerability reports and practical detection tools.

Synthesizing such queries poses significant challenges. The syntax of static analysis query languages is low-resource, richly expressive, and evolves continually. A typical query, such as the one in Figure 2(b) specifying a global dataflow pattern leaves ample room for errors in describing predicates for sources, sinks, sanitizers, and taint propagation steps. Even if the generated syntax is correct, success is measured by whether the query can identify at least one execution path traversing the bug location in the vulnerable version while producing no matches in the patched version. Achieving this requires understanding the CVE context at the level of abstract syntax trees, such as code differences that introduce a sanitizer to prevent a flow from a source to a sink. Complicating matters further, reasoning about the code changes alone is often insufficient: sources, sinks, and taint propagation steps may reside in parts of the codebase far from the modified functions or files, and the vulnerability itself may involve non-trivial dataflow chains across these components. Thus, a correct query must not only integrate information from multiple locations across the program but also capture the intricate propagation patterns to accurately characterize the vulnerability.

In this paper, we present **QLCoder**—an agentic framework that synthesizes queries in CodeQL, a powerful static analysis engine, directly from a given CVE metadata. We select CodeQL because it has the richest query language, which allows capturing complex inter-procedural vulnerability patterns. QLCoder addresses the above challenges by embedding an LLM in a structured synthesis loop that incorporates execution feedback to verify query correctness and allows interactive reasoning using a custom MCP (Model Context Protocol) interface. The MCP interface constrains the model’s reasoning using a Language Server Protocol (for syntax guidance) and a vector database of CodeQL queries and documentation (for semantic guidance). By combining these capabilities, QLCoder avoids common pitfalls of naive LLM-based approaches, such as producing ill-formed queries, hallucinating deprecated constructs, or missing subtle vulnerability patterns, and instead produces queries that are both syntactically correct and semantically precise.

We evaluate QLCoder on CWE-Bench-Java (Li et al., 2025b), which comprises 176 CVEs across 111 Java projects. These CVEs span 42 different Common Weakness Enumeration (CWE) categories and the projects range in size from 0.01 to 1.5 MLOC. To account for model training cut-offs, we include 65 CVEs reported during 2025 and target a recent CodeQL version 2.22.2 (July 2025). Using the Claude Code agent framework, QLCoder achieves query compilation and success rates of 100% and 53.4%, compared to 19% and 0% for our best agentic baseline, Gemini CLI. Further, our generated queries have an F1 score of 0.7 for detecting true positive vulnerabilities, compared to 0.048 for IRIS (Li et al., 2025b), a recent LLM-assisted static analyzer, and 0.073 for CodeQL.

We summarize our main contributions:

- **Agentic Framework for CVE-to-Query Synthesis.** We present QLCoder, an agentic framework that translates CVE descriptions into executable CodeQL queries, bridging the gap between vulnerability reports and static analysis. QLCoder introduces a novel integration of execution-guided synthesis, semantic retrieval, and structured reasoning for vulnerability query generation.
- **Evaluation on Real-World Repositories and CVEs.** We evaluate QLCoder on 176 CVEs in Java projects, covering 42 vulnerability types (CWEs) from CWE-Bench-Java. Each project involves complex inter-procedural vulnerabilities spanning multiple files. We show how QLCoder can successfully identify sources, sinks, sanitizers, and taint propagation steps, and refine queries to ensure they raise alarms on vulnerable versions while remaining silent on patched versions.
- **Comparison with Baselines.** We compare QLCoder against state-of-the-art agent frameworks and show that QLCoder achieves substantially higher compilation, success, and F1 scores. We

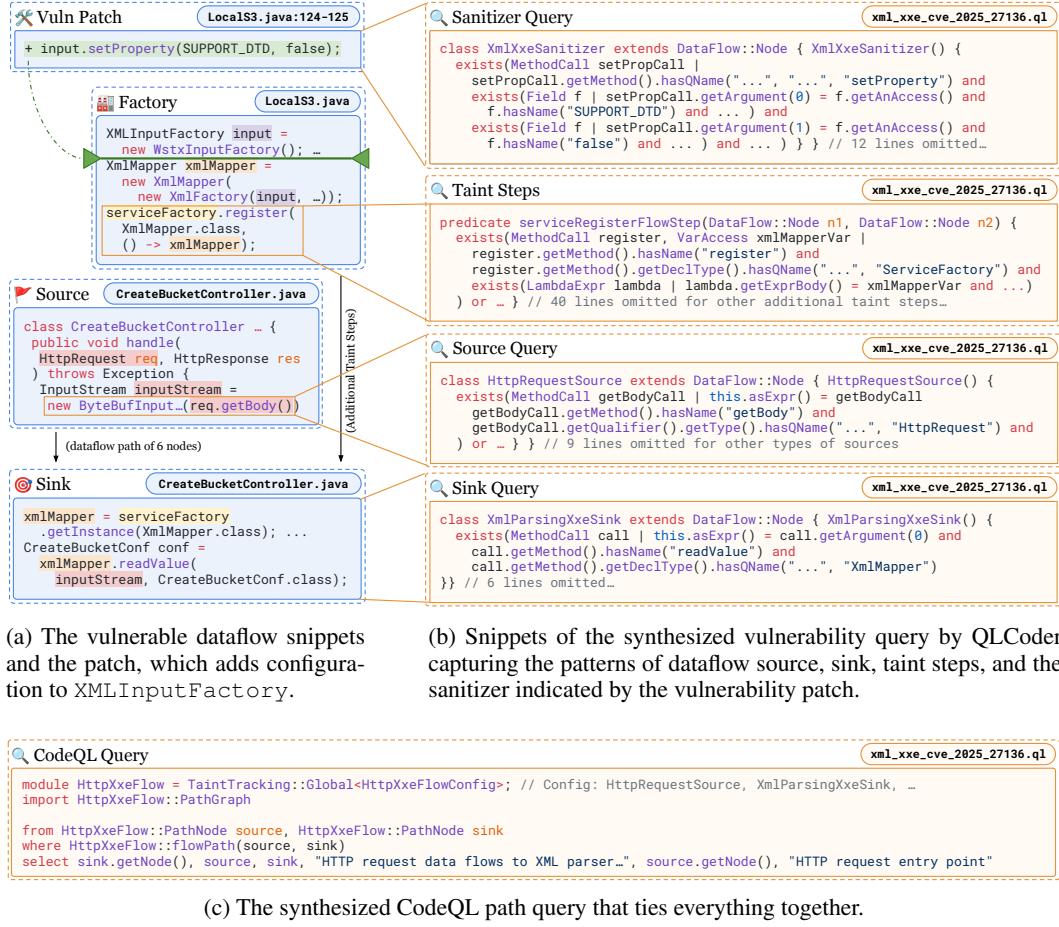


Figure 2: Illustration of vulnerability CVE-2025-27136 in repository Robothy/local-s3 which exhibits an XML External Entity Injection weakness (CWE-611). When the `XmlMapper` is not configured to disable *Document Type Definition* (DTD), the function `readValue` may declare additional entities, allowing hackers to inject malicious behavior.

also compare QLCoder’s synthesized queries with state-of-the-art static analysis frameworks and show that our queries are more precise and have higher recall.

2 ILLUSTRATIVE EXAMPLE

We illustrate the challenges of vulnerability query synthesis using CVE-2025-27136, an XML External Entity Injection (XXE) bug found in the repository Robothy/local-s3. Figure 2 depicts the vulnerability snippets, the patch, and the synthesized CodeQL query generated by QLCoder.

Vulnerability context. The vulnerability arises when the `XmlMapper` object is used to parse user-provided XML data (Figure 2a). In the vulnerable code, `XmlMapper.readValue` is called on the HTTP request body without disabling support for *Document Type Definitions* (DTDs). As a result, an attacker can inject malicious external entity declarations into the input stream, enabling server-side request forgery (SSRF) attacks, allowing for access to resources that should not be accessible from external networks, effectively leaking sensitive information. The patch mitigates the issue by configuring the underlying `XMLInputFactory` with the property `SUPPORT_DTD=false`.

Synthesizing the query. The CodeQL query that can effectively capture the vulnerability pattern needs to incorporate 1) sources such as `HttpRequest.getBody` calls where untrusted malicious information enters the program, 2) sinks such as invocations of `XmlMapper.readValue`, where the XXE vulnerability is manifested, 3) additional taint steps related to how the `XmlMapper` is

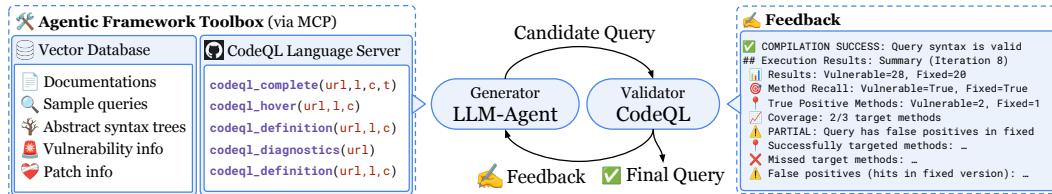


Figure 3: Overall pipeline of QLCoder’s iterative synthesis loop between an agentic query generator and a CodeQL-based validator. The generator uses a vector database and our CodeQL Language Server as tools while the validator produces compilation, execution, and coverage feedback.

constructed and configured, involving non-trivial interprocedural flows spanning multiple files, and 4) sanitizers such as calls to `setProperty(SUPPORT_DTD, false)`, so that we know that no alarm should be reported after the vulnerability has been fixed.

In general, the synthesized query must connect all these components to be able to detect the bug in the vulnerable program, while not reporting the same alarm after the vulnerability has been fixed. Figure 2b shows all the components of the CodeQL query (simplified), capturing their individual syntactic patterns. Lastly, Figure 2c connects all these components into a coherent path query by using CodeQL’s `TaintTracking::Global<.>::PathGraph` and the SQL-like `from-where-select` query, which returns the exact path from source to sink.

Challenges and solutions. Vulnerability query synthesis must overcome several tightly-coupled challenges. We hereby state the challenges and explain how QLCoder addresses them.

- *Rich expressiveness and fragility of syntactic patterns.* CodeQL is powerful but syntactically intricate: small mistakes in predicate names, qualifiers, or AST navigation often produce syntactically valid yet semantically useless queries. QLCoder mitigates this fragility through its Language Server Protocol (LSP) interface for syntax guidance and RAG database for semantic retrieval of existing CodeQL queries and documentation. These structured interactions guide predicate selection and AST navigation during synthesis, reducing off-by-name and version-mismatch errors.
- *Inter-procedural taint propagation across a large codebase.* Sources, sinks, and sanitizers typically live in different modules or files and are connected by nontrivial inter-procedural flows (lambdas, factory patterns, etc.). While CodeQL provides robust inter-procedural analysis for many common patterns, gaps in dataflow still require bridging via additional taint propagation steps. Through its custom MCP interface, QLCoder performs structured reasoning to discover candidate program points, synthesize custom taint-step predicates (e.g., service registration), and compose them into a CodeQL path query that tracks data across file and component boundaries.
- *Semantic precision: alarm on the vulnerable version, silence on the patched version.* A useful vulnerability query must not only parse correctly but also be discriminative. QLCoder enforces this semantic requirement directly during synthesis. Via an iterative refinement loop, the successful criteria states that in the fixed program, there should be no alarm being raised about the vulnerability. This incentivizes the agent to synthesize sanitizer predicates (e.g., the `setProperty` call) and use them to constrain the path query so that sanitizer presence suppresses the alarm. The resulting query thus captures the exact behavior difference, producing alarms on the vulnerable snapshot and not on the patched snapshot.

Together, these capabilities let QLCoder synthesize a semantically precise CodeQL query that can be reused for regression testing, variant analysis, or patch validation. We now elaborate on the detailed design and implementation of QLCoder.

3 QLCODER

At a high level, QLCoder operates inside a repository-aware iterative refinement loop (Figure 3). In each iteration, the agent proposes a candidate CodeQL query, a CodeQL-based validator executes and scores it on both the vulnerable and patched versions of the repository, and the agent uses the validation feedback to propose targeted repairs. The loop terminates successfully when the validator accepts a query, or fails after a fixed iteration budget. In this section, we elaborate the major design components that make the loop effective.

3.1 PROBLEM STATEMENT

The task of vulnerability detection is generally framed as a taint analysis task, where the goal of a *query* is to find dataflow paths from a *source* (e.g., an API endpoint accepting user input) to a *sink* (e.g., a database write) that lack proper *sanitization* (e.g., filtering malicious data).

We formalize the *Vulnerability Query Synthesis* problem as follows. Assume as input a vulnerable project version P_{vuln} , its fixed version P_{fixed} , and a textual CVE description (commonly available in open vulnerability reports). Let us assume we have inter-procedural dataflow program graphs for each code version: $G_{\text{vuln}} = (V_{\text{vuln}}, E_{\text{vuln}})$ and $G_{\text{fixed}} = (V_{\text{fixed}}, E_{\text{fixed}})$. Let ΔP denote the source-level patch between P_{vuln} and P_{fixed} . We represent the patch in the dataflow-graph domain as a patch subgraph $\Delta G = (\Delta V, \Delta E)$, where ΔV is the set of graph nodes that correspond to the modified program snippets.

A *vulnerability path query* Q evaluated on a graph G returns a set of dataflow paths, denoted as $\Pi = \llbracket Q \rrbracket(G)$. We write each path $\pi \in \Pi$ as $\pi = \langle v_1, \dots, v_k \rangle$, where each $v_i \in V$ is a node in the dataflow graph G . Consecutive nodes (v_i, v_{i+1}) should be either connected by an existing edge in E , or an *additional taint step* specified in the query Q , to compensate for missing edges via dataflow graph construction. Specifically, we call v_1 the *source* of path π and v_k the *sink* of π .

Synthesis task. We aim to synthesize a query Q from the vulnerability report satisfying the following requirements:

1. **Well-formedness.** Q is syntactically valid (based on the latest CodeQL syntax) and can be executed on the target CodeQL infrastructure (e.g., dataflow graphs) without runtime errors.
2. **Vulnerability detection.** Q generates at least one path π in the vulnerable version that traverses the patched region:

$$\exists \pi \in \llbracket Q \rrbracket(G_{\text{vuln}}) \text{ such that } \pi \cap \Delta V \neq \emptyset.$$
3. **Fix discrimination.** Q does not report the vulnerability in the fixed version. Concretely, no path reported on the fixed version should traverse the patched locations:

$$\forall \pi \in \llbracket Q \rrbracket(G_{\text{fixed}}), \text{ we have } \pi \cap \Delta V = \emptyset.$$

In other words, the synthesized query must be executable, must witness the vulnerability in the vulnerable version via a path that uses code touched by the fix, and must not attribute the same (patched) behavior in the fixed version. When only the well-formedness condition is satisfied, we say that the query Q is valid (denoted as $\text{valid}(Q)$); when all the conditions are satisfied, the query Q is successful (denoted as $\text{success}(Q; P_{\text{vuln}}, P_{\text{fixed}})$). Note that these criteria may admit potentially false positive paths in both versions. It might be possible to consider additional constraints regarding precision, but it might further complicate synthesis. In practice, we find most queries synthesized by QLCoder already have high precision.

3.2 DESIGN OF QLCODER

Concretely, QLCoder proceeds in an iterative refinement loop indexed by $i = [0, 1, \dots]$. Via prompting, the LLM agent-based synthesizer first proposes an initial candidate query Q_0 . For each iteration i , the validator evaluates Q_i and produces a feedback report. We consider synthesis successful at iteration i iff $\text{success}(Q_i; P_{\text{vuln}}, P_{\text{fixed}})$ holds; in that case the loop terminates and Q_i is returned. Otherwise, the synthesizer analyzes the feedback and the previous candidate Q_i , and produces the next query candidate Q_{i+1} . The loop stops successfully when $\text{success}(\cdot)$ is achieved or fails once i reaches the pre-configured limit N (in our implementation $N = 10$). The remainder of the design focuses on two aspects: 1) how the agentic synthesizer performs synthesis, and 2) how the validator generates and communicates feedback. We elaborate on both below.

Agentic synthesizer. In each iteration i , the LLM-based agentic synthesizer runs an inner *conversation loop* of up to M turns. In each turn, the agent either performs internal reasoning or issues a tool call by emitting a JSON-formatted action. When a tool call succeeds, the tool returns a JSON-formatted response that is appended to the conversation history. Conversation histories are kept local to the current refinement iteration (i.e., not carried over between iterations) to keep context compact and relevant. In practice, we set $M = 50$, i.e., the agent may interact with tools up to 50 times before generating a candidate query for validation.

Two design choices are critical for the effectiveness of this loop: 1) the *initial prompt* that initializes and constrains the agent’s behavior, and 2) the *toolbox* of callable tools, each exposed by a custom

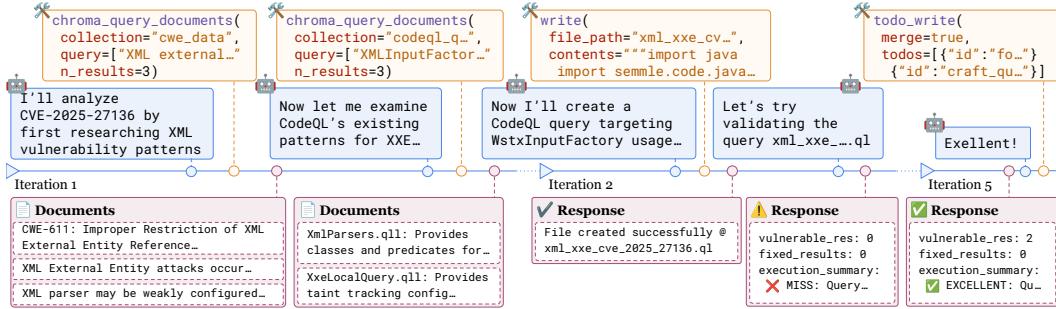


Figure 4: Illustration of example traces of conversation during the synthesis of the query in the motivating example (Figure 2). LLM-agent may think, invoke tools that are available in the toolbox, and receive responses from the MCP servers.

Model Context Protocol (MCP) server. We refer to the combined problem of designing these items as *Context Engineering* (discussed in Section 3.3).

CodeQL-Based Validator. The validator compiles and executes each candidate query against the vulnerable and fixed versions and returns a concise, structured feedback report that is used to drive refinement (Figure 3). The vulnerable and patched versions are automatically derived from patch commit hash, which is in the CVE metadata. The report contains: (i) CodeQL compilation results, (ii) execution counts (matches on vulnerable and fixed graphs), (iii) recall and coverage statistics, (iv) concrete counterexample traces and hit locations, and (v) a prioritized set of next-step recommendations (e.g., add qualifiers, synthesize sanitizer checks, or expand taint steps) that are programmatically generated via a template.

3.3 CONTEXT ENGINEERING FOR AGENTIC SYNTHESIZER

The primary goal of context engineering is to expose the LLM-based agent to the most *precise* amount of information: enough for the agent to make progress, but not so much that the LLM is confused or the cost explodes. As illustrated in Figure 3, QLCoder relies on two primary MCP servers to provide demand-driven, structured information to the agent: a retrieval-augmented vector database and a CodeQL Language Server interface. We show example traces of conversation loop in Figure 4 and describe the available tools below.

Initial prompt. Each refinement iteration begins with an *initial prompt* that kicks-off the agentic conversation loop. The initial prompt in the first iteration contains a query skeleton for reference (See § D.1 for an example). In subsequent iterations, the prompt contains a summary of the synthesis goal and constraints, the previous candidate query Q_{i-1} , and the validator feedback report. Concretely, the initial prompt emphasizes: (i) the success predicate (see $\text{Success}(\cdot)$), (ii) concrete counterexamples from previous feedback, and (iii) an explicit list of callable tools and their purpose.

Vector database. We use a retrieval-augmented vector database (ChromaDB MCP server in our implementation) to store large reference corpora without polluting the LLM prompt. The database is pre-populated with (i) vulnerability analysis notes and diffs, (ii) Common Weaknesses Enumeration (CWE) definitions, (iii) same-version CodeQL API documentation, (iv) curated CodeQL sample queries, and (v) small abstract syntax tree (AST) snippets extracted from the target repository. During a conversation loop, the agent issues compact retrieval queries (e.g., to fetch example CodeQL queries related to the CWE) and receives ranked documents or snippets on demand.

In practice, we may populate our RAG database with tens of thousands of documents. Even with this large corpus, we observe that the LLM-agent reliably retrieves exactly the kinds of artifacts it needs: CodeQL sample queries that inspire overall query structure, small AST snippets that suggest the precise syntactic navigation, and vulnerability writeups or diff excerpts that help discriminate buggy from patched behavior. These demand-driven lookups let the agent gather high-quality information without loading the main prompt with large reference corpora.

CodeQL language server. We expose the CodeQL Language Server (GitHub, 2025f) through a MCP server that the agent can call for precise syntax-aware guidance. Importantly, we developed our own CodeQL Language Server client and MCP server that ensures syntactic validity (especially

for the given CodeQL version) during query generation. The LLM agent’s MCP client makes the tool call which is received by the CodeQL MCP server. The MCP server forwards tool calls, such as `complete(file, loc, char)`, `diagnostics(file)`, and `definition(file, loc, char)`, to the underlying CodeQL process and returns JSON-serializable responses. Tools such as `completion` help the agent fill query templates and discover correct API or AST names, while `diagnostics` reveal compile or linter errors (e.g., unknown predicate names) that guide mutation. Appendix C shows the full specification and example request and response schemas. We discuss several alternative designs that we considered but found ineffective in practice. Allowing the agent unrestricted access to compile-and-run CodeQL via MCP led to severe performance degradation: compilation and full execution are expensive operations that the LLM soon overused, so we instead expose only lightweight diagnostics during the conversation and defer full compile-and-run to the end of each iteration. Permitting web search for vulnerability patterns or snippets similarly proved problematic. It is both costly and easy for the agent to rely on web lookups, which quickly pollutes the working context and degrades synthesis quality. Equipping the agent with an extensive set of heterogeneous tools led to confusion and poor tool-selection behavior; in contrast, a small, well-scoped toolbox yields more reliable actions. Finally, retaining full conversation histories across refinement iterations induced context rot and ballooning prompt sizes, so we keep histories local to each iteration. Overall, our current design is a pragmatic trade-off that balances cost, responsiveness, and synthesis effectiveness.

4 EVALUATION

We aim to answer the following research questions through our empirical evaluation:

- **RQ 1:** For how many CVEs can QLCoder successfully generate queries?
- **RQ 2:** How useful is each component of QLCoder?
- **RQ 3:** How does the choice of base agent framework affect QLCoder’s effectiveness?
- **RQ 4:** How scalable is QLCoder?

4.1 EXPERIMENTAL SETUP

We develop QLCoder on top of the Claude Code framework (Anthropic, 2025) and use Claude Sonnet 4 for all our experiments. For agent baselines, we select Codex with GPT-5 (minimal reasoning) and Gemini CLI with Gemini 2.5 Flash. For each CVE and agent baseline, we use a maximum of 10 iterations ($N = 10$). For static analysis baselines, we select IRIS Li et al. (2025b) and CodeQL (version 2.22.2) query suites. Experiments were run on machines with the following specifications: for the gpt-oss agent baseline we used a machine with For all other experiments we used an Intel Xeon Gold 6248 2.50GHz CPU, four GeForce RTX 2080 Ti GPUs, and 750GB RAM.

Dataset. We used CWE-Bench-Java (Li et al., 2025b) and its latest update, which added new CVEs from 2025. We were able to successfully build and use 111 (out of 120) Java CVEs evaluated in IRIS (Li et al., 2025b), and 65 (out of 91) 2025 CVEs. Each sample in CWE-Bench-Java comes with the CVE metadata and fix commit information associated with the bug.

4.2 EVALUATION METRICS

Besides $\text{valid}(Q)$ and $\text{success}(Q; P_{\text{vuln}}, P_{\text{fixed}})$ from Section 3.1, we use the following terms and metrics when evaluating QLCoder and baselines on the problem of vulnerability query synthesis:

$$\begin{aligned} \text{Rec}(Q) &= \mathbb{1}[\exists \pi \in \llbracket Q \rrbracket(G_{\text{vuln}}), \pi \cap \Delta V \neq \emptyset], \quad \text{Prec}(Q) = \frac{|\{\pi \in \llbracket Q \rrbracket(G_{\text{vuln}}) \mid \pi \cap \Delta V \neq \emptyset\}|}{|\llbracket Q \rrbracket(G_{\text{vuln}})|}, \\ \text{F1}(Q) &= 2 \cdot \frac{\text{Prec}(Q) \cdot \text{Rec}(Q)}{\text{Prec}(Q) + \text{Rec}(Q)}. \end{aligned}$$

4.3 RQ1: QLCODER EFFECTIVENESS

QLCoder vs. state-of-the-art QL. Table 1 shows QLCoder’s overall query synthesis success rate by CWE. Table 2 shows the notable increase in precision of QLCoder over CodeQL and IRIS. QLCoder is able to successfully synthesize 53.4% of the CVEs. For half the queries QLCoder correctly synthesizes CodeQL, detects the CVE, and does not report false positives on the fixed version of the CVE’s repository. The lack of true positive recall is why CodeQL and IRIS have significantly lower precision. CodeQL’s queries are broad, categorized by CWE queries. IRIS

Table 1: QLCoder Query Success by CWE Type

CWE Type	Total CVEs	# Success	Success (%)	Avg Precision
CWE-022 (Path Traversal)	48	31	64.6	0.75
CWE-079 (Cross-Site Scripting)	36	18	50.0	0.621
CWE-094 (Code Injection)	20	12	60.0	0.606
CWE-078 (OS Command Injection)	12	7	58.3	0.628
CWE-502 (Deserialization)	6	4	66.7	0.853
CWE-611 (XXE)	5	3	60.0	0.657
Other CWEs (≤ 4 CVEs)	49	19	38.8	0.504
Total	176	94	53.4	0.631

Table 2: Recall Performance Comparison Across Methods (Shared CVEs: 130)

Method	Recall Rate (%)	Avg Precision	Avg F1 Score
CodeQL	20.0	0.055	0.073
IRIS	35.4	0.031	0.048
QLCoder	80.0	0.672	0.700

generates all of the predicates for potential sources and sinks with CodeQL, and does not generate sanitizer or taint step predicates.

Finally, Figure 5 shows that CodeQL, IRIS, and QLCoder have significantly higher vulnerability recall rates compared to Snyk and SpotBugs. Thus highlighting CodeQL’s superior performance compared to other static analysis tools.

Impact of training cut-off. We also want to take note that Claude Sonnet 4’s training cut-off is March 2025. Table 3 shows that QLCoder performs consistently regardless of CVEs before or after the cut-off period. The CodeQL version, 2.22.2, was released in July 2025. New versions of CodeQL often include analysis improvements and new QL packs (GitHub, 2025c).

Table 3: Tool Performance Before vs After Training Cutoff

CVE Period	Total CVEs	# Recall	Success (%)	Avg Precision	Avg F1 Score
Pre-2025 (2011-2024)	111	64	57.7	0.676	0.702
2025+ (Post-cutoff)	65	30	46.2	0.555	0.583
Overall	176	94	53.4	0.631	0.658

4.4 RQ2: ABLATION STUDIES

For ablations, we chose 20 CVEs and ran QLCoder with one of the QLCoder components removed (Table 4). The ablation with no tools refers to only running Claude Code with the iterative feedback system. The high recall rate when removing access to the AST cache while lowered recall rates without the LSP server or documentation access show that the LSP and documentation lookup impact the synthesis performance more. We also include QLCoder’s performance on the same set of CVEs, and point out its significantly higher query success rate and precision score. Claude Code without tools scored a high recall rate, yet failed to synthesize queries without false positives when executed on the fixed version.

4.5 RQ3: STATE OF THE ART AGENT COMPARISON

We used Gemini CLI with Gemini 2.5 Flash . We used Codex with GPT-5 . We evaluated their performance on 20 CVEs. We achieved an increase in compilation success for both agents compared to using the agents without QLCoder in Table 5.

5 RELATED WORK

LLMs and vulnerability detection. LLMs have been used extensively for vulnerability detection and repair using techniques such as fine-tuning and prompt engineering (Zhou et al., 2024). LLMs have also been combined with existing program analysis tools for vulnerability detection. The combination of LLMs can be used for vulnerability analysis like IRIS’s (Li et al., 2025b) source and sink identification, however IRIS depends on a limited set of CWE templates derived from CodeQL’s CWE queries. IRIS also only the LLM for identifying sources and sinks. KNighter synthesizes CSA checkers given a fix commit of a C repository (Yang et al., 2025), however the checkers are written in

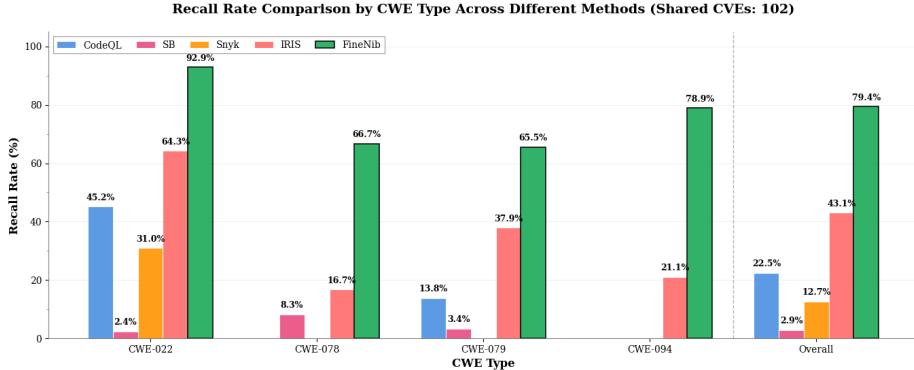


Figure 5: Recall Rate Comparison by CWE Type Across Different Methods (102 CVEs).

Table 4: Ablation Study (out of 20 CVEs)

Variant	% Successful	Recall Rate	Avg Precision	Avg F1 Score
QLCoder	55%	80%	0.67	0.69
w/o LSP	25% (-30%)	55% (-25%)	0.32	0.36
w/o Doc/Ref	20% (-35%)	55% (-25%)	0.32	0.36
w/o AST	25% (-30%)	80% ($\pm 0\%$)	0.41	0.47
	10% (-45%)	55% (-25%)	0.33	0.36

C which has more available training data. MocQ’s uses an LLM to derive a subset DSL of CodeQL and Joern, and then provides a feedback loop to the LLM though prompting via API calls is used rather than an agent with tools and MocQ uses significantly higher iterations, with a max threshold of 1,000 iterations per vulnerability experiment. (Li et al., 2025a).

LLM agents and tool usage. SWE-agent pioneered the idea of autonomous LLM agents using tools for software engineering tasks Yang et al. (2024). LSPAI Go et al. (2025), an IDE plugin, uses LSP servers to guide LLM-generated unit tests. Hazel, a live program sketching environment, uses a language server (Blinn et al., 2024) to assist code completions synthesized by LLMs. The Hazel Language Server provides the typing context of a program hole to be filled.

Low resource LLM code generation. SPEAC uses ASTs combined with constraint solving to repair LLM-generated code for low resource programming languages (Mora et al., 2024). SPEAC converts a buggy program into an AST and uses a solver to find the minimum set of AST nodes to replace, to satisfy language constraints. MultiPL-T generates datasets for low resource languages by translating high resource language code to the target language and validates translations with LLM generated unit tests (Cassano et al., 2024).

6 CONCLUSION AND LIMITATIONS

We present QLCoder, an agentic framework for synthesizing syntactically correct and precise CodeQL queries given known vulnerability patterns. We will also open source our CodeQL LSP MCP server and QLCoder. In future work, we plan to explore efficient ways to synthesize, and to combine our synthesized queries with dynamic analysis tools.

Table 5: LLM-agent baselines’ compilation and success rates on 20 CVEs from 2025.

Agent Baselines	Language Model	w/o QLCoder		with QLCoder	
		Compilation	Successful	Compilation	Successful
Gemini CLI	Gemini 2.5 Flash	19%		24% (+5%)	(=)
	Gemini 2.5 Pro			(+40%)	(=)
Codex	GPT-5,	0%		24% (+24%)	(+5%)
			0%	(+55%)	(+20%)
				(+10%)	(+5%)
				(+5%)	(+25%)

Limitations. We omit CVEs where the vulnerability involves non-Java code such as configuration files or other languages. QLCoder can be used with exploit generation to find vulnerabilities that are realized during dynamic execution. For supporting other languages that can be queried by CodeQL, the vector database can be filled with references, documentation, and example queries in other CodeQL supported languages. We also note that Claude Sonnet 4’s official training cut-off is March 2025, however the 2025 CVEs evaluated were reported between January to August 2025.

ACKNOWLEDGMENTS

We thank the reviewers for feedback that improved this paper. This research was supported by NSF award CCF #2313010 and the NSF CSGrad4US Fellowship.

REFERENCES

Anthropic. Claude code: Agentic coding assistant, 2025. URL <https://github.com/anthropics/clause-code>. Version 1.0. Released February 2025. Accessed September 2025.

Andrew Blinn, Xiang Li, June Hyung Kim, and Cyrus Omar. Statically Contextualizing Large Language Models with Typed Holes. *Artifact for Statically Contextualizing Large Language Models with Typed Holes*, 8(OOPSLA2):288:468–288:498, October 2024. doi: 10.1145/3689728. URL <https://dl.acm.org/doi/10.1145/3689728>.

Federico Cassano, John Gouwar, Francesca Lucchetti, Claire Schlesinger, Anders Freeman, Carolyn Jane Anderson, Molly Q Feldman, Michael Greenberg, Abhinav Jangda, and Arjun Guha. Knowledge transfer from high-resource to low-resource programming languages for code ILms, 2024. URL <https://arxiv.org/abs/2308.09895>.

CVE. Common vulnerabilities and exposures (cve), 2025. URL <https://www.cve.org/about/Metrics>.

GitHub. Codeql, 2025a. URL <https://codeql.github.com/>.

GitHub. Github security advisory database, 2025b. URL <https://github.com/advisories>.

GitHub. Codeql 2.23.0 (2025-09-04) — codeql, 2025c. URL <https://codeql.github.com/docs/codeql-overview/codeql-changelog/codeql-cli-2.23.0/>.

GitHub. codeql/java/ql/lib/semmle/code/java/security at main · github/codeql, 2025d. URL <https://github.com/github/codeql/tree/main/java/ql/lib/semmle/code/java/security>.

GitHub. codeql/java/ql/src/security/cwe at main · github/codeql, 2025e. URL <https://github.com/github/codeql/tree/main/java/ql/src/Security/CWE>.

GitHub. execute language-server, 2025f. URL <https://docs-internal.github.com/en/code-security/codeql-cli/codeql-cli-manual/execute-language-server>.

Gwiwhan Go, Chijin Zhou, Quan Zhang, Yu Jiang, and Zhao Wei. LSPA: An IDE Plugin for LLM-Powered Multi-Language Unit Test Generation with Language Server Protocol. In *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering*, pp. 144–149. ACM, June 2025. ISBN 979-8-4007-1276-0. doi: 10.1145/3696630.3728540. URL <https://dl.acm.org/doi/10.1145/3696630.3728540>.

Penghui Li, Songchen Yao, Josef Sarfati Korich, Changhua Luo, Jianjia Yu, Yinzhi Cao, and Junfeng Yang. Automated static vulnerability detection via a holistic neuro-symbolic approach, 2025a. URL <https://arxiv.org/abs/2504.16057>.

Z. Li, S. Dutta, and M. Naik. IRIS: LLM-Assisted Static Analysis for Detecting Security Vulnerabilities, Apr. 2025b. arXiv:2405.17238.

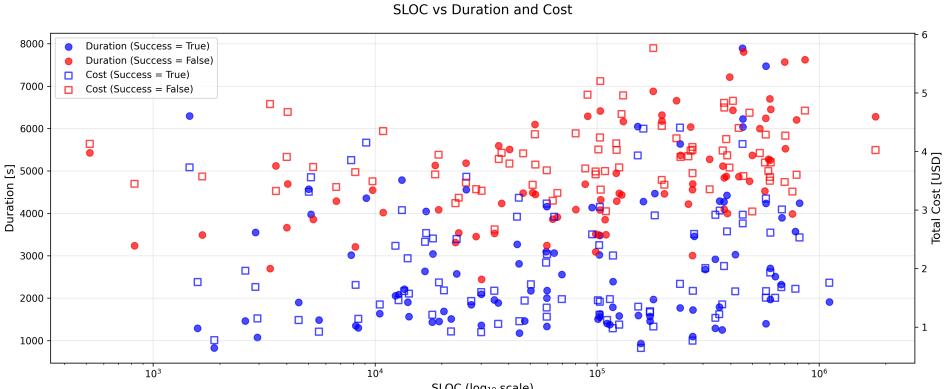


Figure 6: SLOC, Synthesis Duration, and Synthesis Cost Comparison by CVE

Meta. Infer. <https://fbinfer.com/>, 2025.

MITRE. Common vulnerabilities and exposures (cve), 2025. URL <https://www.cve.org/>.

Federico Mora, Justin Wong, Haley Lepe, Sahil Bhatia, Karim Elmaaroufi, George Varghese, Joseph E Gonzalez, Elizabeth Polgreen, and Sanjit A Seshia. Synthetic programming elicitation for text-to-code in very low-resource programming and formal languages. *Advances in Neural Information Processing Systems*, 37:105151–105170, 2024.

NIST. National vulnerability database (nvd), 2025. URL <https://nvd.nist.gov/>.

Semgrep. Semgrep. the semgrep platform. <https://semgrep.dev/>, 2023.

C. Yang, Z. Zhao, Z. Xie, H. Li, and L. Zhang. KNighter: Transforming Static Analysis with LLM-Synthesized Checkers, Apr. 2025. arXiv:2503.09002.

John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering, 2024. URL <https://arxiv.org/abs/2405.15793>.

Xin Zhou, Sicong Cao, Xiaobing Sun, and David Lo. Large language model for vulnerability detection and repair: Literature review and the road ahead, 2024. URL <https://arxiv.org/abs/2404.02525>.

A EVALUATION DETAILS

Table 6 is a more detailed breakdown of the successful query synthesis rate by CWE.

```

1 [model_providers.ollama]
2 name = "Ollama"
3 base_url = "http://localhost:11434/v1"
4 [profiles.gpt-oss-20b]
5 model_provider = "ollama"
6 model = "gpt-oss:20b"
7 wire_api = "chat"

```

A.1 EVALUATION LIMITATIONS

Codex CLI. Claude Code and Gemini CLI allow users to configure how many max turns an agent can take in a context window. As of 9/23/2025, Codex CLI does not offer this configuration. Thus we were not able to force Codex to always take up to 50 max turns each context window.



Figure 7:

IRIS. The original IRIS evaluation consists of 120 Java projects from CWE-Bench-Java. Many of these projects are old with deprecated dependencies, thus we were only able to build and use 112 of the projects with CodeQL 2.22.2. As of 9/23/2025, IRIS supports 11 CWEs and out of the 65 CVEs from 2025, we were able to use 24 of them with IRIS. When running some of the IRIS queries, the amount of sources and sink predicates in the query led to out of memory errors. This impacted 9 out of the 24 queries, thus we treat those as queries with 0 results and false recall.

B
1 Turn 5 - Assistant
2 Tool: WebSearch
3 Query: CVE-2025-27136 vulnerability details XML external entity XXE
4 Tokens: Input: 5 | Cache Read: 29,596 | Cache Create: 267 | Output: 67

Table 6: QLCoder Query Success by CWE Type

CWE Type	Total CVEs	# Success	Success (%)	Avg Precision
CWE-022 (Path Traversal)	48	31	64.6%	0.750
CWE-079 (Cross-Site Scripting)	36	18	50.0%	0.621
CWE-094 (Code Injection)	20	12	60.0%	0.606
CWE-078 (OS Command Injection)	12	7	58.3%	0.628
CWE-502 (Deserialization)	6	4	66.7%	0.853
CWE-611 (XXE)	5	3	60.0%	0.657
CWE-287 (Authentication)	4	1	25.0%	0.875
CWE-200 (Information Exposure)	3	0	0.0%	0.667
CWE-400 (Resource Consumption)	3	1	33.3%	0.556
CWE-532 (Information Exposure)	3	3	100.0%	0.686
CWE-770 (Resource Exhaustion)	3	1	33.3%	0.444
CWE-020 (Improper Input Validation)	2	2	100.0%	0.650
CWE-089 (SQL Injection)	2	2	100.0%	1.000
CWE-1333 (ReDoS)	2	0	0.0%	0.000
CWE-284 (Access Control)	2	0	0.0%	0.500
CWE-862 (Authorization)	2	0	0.0%	0.000
CWE-918 (SSRF)	2	1	50.0%	0.500
CWE-023 (Relative Path Traversal)	1	1	100.0%	1.000
CWE-044 (Path Equivalence)	1	1	100.0%	0.667
CWE-083 (Improper Neutralization)	1	1	100.0%	0.052
CWE-1325 (Improperly Controlled Memory)	1	0	0.0%	0.000
CWE-164 (Foreign Code)	1	0	0.0%	0.000
CWE-178 (Case Sensitivity)	1	0	0.0%	1.000
CWE-190 (Integer Overflow)	1	0	0.0%	0.000
CWE-264 (Permissions)	1	0	0.0%	0.000
CWE-267 (Privilege Defined)	1	0	0.0%	0.000
CWE-276 (Incorrect Permissions)	1	0	0.0%	1.000
CWE-285 (Improper Authorization)	1	1	100.0%	1.000
CWE-288 (Authentication Bypass)	1	0	0.0%	0.000
CWE-290 (Authentication Bypass)	1	1	100.0%	1.000
CWE-297 (Improper Certificate)	1	1	100.0%	1.000
CWE-312 (Cleartext Storage)	1	0	0.0%	0.000
CWE-327 (Cryptographic Issues)	1	0	0.0%	0.000
CWE-346 (Origin Validation)	1	0	0.0%	0.200
CWE-352 (CSRF)	1	1	100.0%	0.941
CWE-426 (Untrusted Search Path)	1	0	0.0%	0.000
CWE-835 (Infinite Loop)	1	0	0.0%	0.000
CWE-863 (Authorization)	1	1	100.0%	1.000
Total	176	94	53.4%	0.631

```

5 Result: Web search results for query: "CVE-2025-27136 vulnerability
6 details XML external entity XXE"
7
8 I'll search for information about the CVE-2025-27136 vulnerability and its
9 relation to XML external entity (XXE) issues.
10
11 Links: [{"title": "Locals3 CreateBucketConfiguration Endpoint XML External Entity
12 (XXE) Injection . CVE-2025-27136 . GitHub Advisory Database . GitHub",
13 "url": "https://github.com/advisories/GHSA-g6wm-2v64-wq36"}, {"title": "CVE-2025-27136 Impact, Exploitability, and Mitigation Steps | Wiz",
14 "url": "https://www.wiz.io/vulnerability-database/cve/cve-2025-27136"}, {"title": "NVD - CVE-2025-27136",
15 "url": "https://nvd.nist.gov/vuln/detail/CVE-2025-27136"}, {"title": "XML External Entity (XXE) Processing | OWASP Foundation",
16 "url": "https://owasp.org/www-community/vulnerabilities/XML_External_Entity_(XXE)_Processing"}, {"title": "What is XXE (XML External Entity) | Examples & Prevention | Imperva",
17 "url": "https://www.imperva.com/learn/application-security/xxe-xml-external-entity/"}, {"title": "CVE - Search Results",
18 "url": "https://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=xxe"}, {"title": "What is XXE (XML external entity) injection? Tutorial & Examples | Web Security Academy", "url": "https://portswigger.net/web-security/xxe"}, {"title": "XXE Complete Guide: Impact, Examples, and Prevention | HackerOne",
19 "url": "https://www.hackerone.com/knowledge-center/xxe-complete-guide-impact-examples-and-prevention"}, {"title": "XML External Entity (XXE)", "url": "https://www.invicti.com/learn/xml-external-entity-xxe"}, {"title": "XML External Entity Prevention - OWASP Cheat Sheet Series", "url": "https://cheatsheetseries.owasp.org/cheatsheets/XML_External_Entity_Prevention_Cheat_Sheet.html"}]
20
21 Turn 6 - Assistant
22 Tool: WebSearch
23 Query: XML external entity XXE CodeQL detection patterns Java XMLInputFactory
24 Tokens: Input: 4 | Cache Read: 30,998 | Cache Create: 267 | Output: 64
25 Result: Web search results for query: "XML external entity XXE CodeQL detection patterns Java
26 XMLInputFactory"
27
28 I'll search for information about XML external entity (XXE) CodeQL detection patterns for Java
29 XMLInputFactory.
30
31 Links: [{"title": "XML External Entity Prevention - OWASP Cheat Sheet Series",
32 "url": "https://cheatsheetseries.owasp.org/cheatsheets/XML_External_Entity_Prevention_Cheat_Sheet.html"}, {"title": "Resolving XML external entity in user-controlled data - CodeQL query help documentation", "url": "https://codeql.github.com/codeql-query-help/java/java-xxe/"}, {"title": "Java false positive: XXE via XMLInputFactory . Issue #7199 . github/codeql", "url": "https://github.com/github/codeql/issues/7199"}, {"title": "XML Entity Expansion in Java | SecureFlag Security Knowledge Base", "url": "https://knowledge-base.secureflag.com/vulnerabilities/xml_injection/xml_entity_expansion_java.html"}, {"title": "Detecting, Fixing, and Defending Against XXE Attacks in Python and Java | by Rohit kumar | Medium", "url": "https://rohitcoder.medium.com/comprehensive-guide-detecting-fixing-and-defending-against-xxe-attacks-in-python-and-java-e78691b4b918"}, {"title": "Prevent XML External Entity Vulnerabilities for Java - xxe", "url": "https://semgrep.dev/docs/cheat-sheets/java-xxe"}, {"title": "how to fix 'Disable XML external entity (XXE) processing' vulnerabilities in java - Stack Overflow", "url": "https://stackoverflow.com/questions/56777287/how-to-fix-disable-xml-external-entity-xxe-processing-vulnerabilities-in-jav"}, {"title": "XXE vulnerability detection | Moderne Docs", "url": "https://docs.moderne.io/openrewrite-advanced-program-analysis/security/xxe/"}, {"title": "How to prevent XML external entity attack (XXE attack) - Mkyong.com", "url": "https://mkyong.com/java/how-to-prevent-xml-external-entity-attack-xxe-attack/"}, {"title": "xml - Java XXE vulnerability - Information Security Stack Exchange", "url": "https://security.stackexchange.com/questions/260956/java-xxe-vulnerability"}]
33
34
35 Turn 7 - Assistant
36 Tool: WebSearch
37 Query: CodeQL XXE taint tracking Java XML parsing sources sinks sanitizers
38 Tokens: Input: 3 | Cache Read: 31,265 | Cache Create: 1,354 | Output: 66
39 Result: Web search results for query: "CodeQL XXE taint tracking Java XML parsing sources
40 sinks sanitizers"
41
42 I'll search for information about CodeQL XXE taint tracking in Java XML parsing, including
43 sources, sinks, and sanitizers.
44
45 Links: [{"title": "Practical Introduction to CodeQL :: jorgectf - blog", "url": "https://jorgectf.github.io/blog/post/practical-codeql-introduction/"}, {"title": "Resolving XML external entity in user-controlled data - CodeQL query help documentation", "url": "https://codeql.github.com/codeql-query-help/java/java-xxe/"}, {"title": "Using flow state for precise data flow analysis - CodeQL", "url": "https://codeql.github.com/docs/codeql-language-guides/using-flow-labels-for-precise-data-flow-analysis/"}]
46
47
48
49
50 Turn 8 - Assistant
51 Tool: WebSearch
52 Query: CodeQL XXE taint tracking Java XML parsing sources sinks sanitizers
53 Tokens: Input: 3 | Cache Read: 31,265 | Cache Create: 1,354 | Output: 66
54 Result: Web search results for query: "CodeQL XXE taint tracking Java XML parsing sources
55 sinks sanitizers"
56
57 I'll search for information about CodeQL XXE taint tracking in Java XML parsing, including
58 sources, sinks, and sanitizers.

```

```

59 {"title": "CodeQL 2.16.4 (2024-03-11) CodeQL", "url": "https://codeql.github.com/docs/codeql-overview/codeql-changelog/codeql-cli-2.16.4/"}, {"title": "Analyzing data flow in Java and Kotlin CodeQL", "url": "https://codeql.github.com/docs/codeql-language-guides/analyzing-data-flow-in-java/"}, {"title": "RealworldCTF 2024 Protected-by-Java-SE Writeup", "url": "https://intrigus.org/research/2025/05/07/realworld-ctf-2024-using-codeql-to-find-bugs-in-codeql/"}, {"title": "New dataflow API for writing custom CodeQL queries - GitHub Changelog", "url": "https://github.blog/changelog/2023-08-14-new-dataflow-api-for-writing-custom-codeql-queries/"}, {"title": "TaintTracking - CodeQL - GitHub", "url": "https://codeql.github.com/codeql-standard-libraries/javascript/semmle/javascript/dataflow/TaintTracking.qll/module.TaintTracking.html"}, {"title": "How codeql deal with java's interface invoke in Global Taint Tracking .github/codeql . Discussion #7363", "url": "https://github.com/github/codeql/discussions/7363"}, {"title": "CodeQL 2.11.0 (2022-09-28) - CodeQL", "url": "https://codeql.github.com/docs/codeql-overview/codeql-changelog/codeql-cli-2.11.0/"}]

```

Listing 1:

```

1 Turn 4 - Assistant
2 Tool: chroma_get_documents(collection_name:nist_cve_2025_brief, where:
3 {"cve_id": "CVE-2025-27136"})
4 Tokens: Input: 6 | Cache Read: 28,074 | Cache Create: 340 | Output: 105
5
6 Turn 6 - Assistant
7 Tool: chroma_query_documents(collection_name: cwe_data, query_texts:
8 ['CWE-611', 'XML External Entity', 'XXE injection'], n_results: 3)
9 Tokens: Input: 5 | Cache Read: 29,031 | Cache Create: 340 | Output: 121

```

Listing 2:

C CODEQL LANGUAGE SERVER VIA MCP

The following are the MCP tool specifications and example usage for our custom CodeQL LSP client, wrapped as an MCP server.

C.1 TOOL SPECIFICATIONS

codeql_complete Provides code completions at a specific position in a CodeQL file. Supports pagination for large completion lists and trigger character-based completions.

Inputs:

- **file_uri** (string): The URI of the CodeQL file
- **line** (number): Line number (0-based)
- **character** (number): Character position in the line (0-based)
- **trigger_character** (string, optional): Optional trigger character (e.g., ":", "::")
- **limit** (number, optional): Maximum number of completion items to return (default: 50)
- **offset** (number, optional): Starting position for pagination (default: 0)

Returns: CompletionList with pagination metadata containing completion items, each with label, kind, documentation, and text edit information.

Example usage:

```

{
  "tool": "codeql_complete",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "line": 5,
    "character": 12,
    "trigger_character": ".",
    "limit": 25
  }
}

```

codeql_hover Retrieves hover information (documentation, type information) at a specific position. Provides rich markdown documentation for CodeQL predicates, classes, and modules.

Inputs:

- `file_uri` (string): The URI of the CodeQL file
- `line` (number): Line number (0-based)
- `character` (number): Character position in the line (0-based)

Returns: Hover | null containing documentation content in markdown or plain text format, with optional range highlighting.

Example usage:

```
{
  "tool": "codeql_hover",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "line": 8,
    "character": 15
  }
}
```

codeql_definition Navigates to the definition location for a symbol at a specific position. Supports both single definitions and multiple definition locations.

Inputs:

- `file_uri` (string): The URI of the CodeQL file
- `line` (number): Line number (0-based)
- `character` (number): Character position in the line (0-based)

Returns: Location | Location[] | null containing URI and range information for definition locations.

Example usage:

```
{
  "tool": "codeql_definition",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "line": 12,
    "character": 8
  }
}
```

codeql_references Finds all references to a symbol at a specific position across the workspace. Includes both usage references and declaration references.

Inputs:

- `file_uri` (string): The URI of the CodeQL file
- `line` (number): Line number (0-based)
- `character` (number): Character position in the line (0-based)

Returns: Location[] | null containing an array of all reference locations with URI and range information.

Example usage:

```
{
  "tool": "codeql_references",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "line": 6,
    "character": 20
  }
}
```

codeql_diagnostics Retrieves diagnostics (errors, warnings, information messages) for a CodeQL file. Provides real-time syntax and semantic analysis results.

Inputs:

- `file_uri` (string): The URI of the CodeQL file

Returns: `Diagnostic[]` containing an array of diagnostic objects with severity, message, range, and optional related information.

Example usage:

```
{
  "tool": "codeql_diagnostics",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql"
  }
}
```

codeql_format Formats a CodeQL file or a specific selection within the file according to CodeQL style guidelines.

Inputs:

- `file_uri` (string): The URI of the CodeQL file
- `range` (Range, optional): Optional range to format with start and end positions

Returns: `TextEdit[]` containing an array of text edits that describe the formatting changes to be applied.

Example usage:

```
{
  "tool": "codeql_format",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "range": {
      "start": { "line": 10, "character": 0 },
      "end": { "line": 25, "character": 0 }
    }
  }
}
```

codeql_update_file Updates the content of an open CodeQL file in the language server. This allows for dynamic content modification and analysis of unsaved changes.

Inputs:

- `file_uri` (string): The URI of the CodeQL file
- `content` (string): The new complete content of the file

Returns: `string` containing a success confirmation message.

Example usage:

```
{
  "tool": "codeql_update_file",
  "arguments": {
    "file_uri": "file:///workspace/security-query.ql",
    "content": "import cpp\n\nfrom Function f\nwhere f.hasName(\"strcpy\")\nselect f"
  }
}
```

D CODEQL QUERIES

D.1 CODEQL QUERY STRUCTURE TEMPLATE

The template below is given to the LLM agent at the start of the iterative query synthesis task. The prompt instructs the LLM to use the AST nodes, along with the CodeQL LSP and CodeQL references in the vector database, to fill in this template. The prompt also takes note to find similar queries related to the given CVE’s vulnerability.

```

1 /**
2  * @name [Vulnerability Name based on analysis]
3  * @description [Description derived from the vulnerability pattern]
4  * @problem.severity error
5  * @security-severity [score based on severity]
6  * @precision high
7  * @tags security
8  * @kind path-problem
9  * @id [unique-id]
10 */
11 import java
12 import semmle.code.java.frameworks.Networking
13 import semmle.code.java.dataflow.DataFlow
14 import semmle.code.java.dataflow.FlowSources
15 import semmle.code.java.dataflow.TaintTracking
16 private import semmle.code.java.dataflow.ExternalFlow
17
18 class Source extends DataFlow::Node {
19     Source() {
20         exists([AST node type from analysis] |
21             /* Fill based on AST patterns for sources identified in Phase 1 & 2 */
22             *
23             and this.asExpr() = [appropriate mapping]
24         )
25     }
26 }
27 class Sink extends DataFlow::Node {
28     Sink() {
29         exists([AST node type] |
30             /* Fill based on AST patterns for sinks */
31             and this.asExpr() = [appropriate mapping]
32         ) or
33         exists([Alternative AST pattern] |
34             /* Additional sink patterns from analysis */
35             and [appropriate condition]
36         )
37     }
38 }
39 class Sanitizer extends DataFlow::Node {
40     Sanitizer() {
41         exists([AST node type for sanitizers] |
42             /* Fill based on sanitizer patterns from Phase 1 & 2 */
43         )
44     }
45 }
```

```

46 }
47
48 module MyPathConfig implements DataFlow::ConfigSig {
49   predicate isSource(DataFlow::Node source) {
50     source instanceof Source
51   }
52
53   predicate isSink(DataFlow::Node sink) {
54     sink instanceof Sink
55   }
56
57   predicate isBarrier(DataFlow::Node sanitizer) {
58     sanitizer instanceof Sanitizer
59   }
60
61   predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {{
62     /* Fill based on additional taint steps from analysis */
63   }}
64 }
65
66 module MyPathFlow = TaintTracking::Global<MyPathConfig>;
67 import MyPathFlow::PathGraph
68
69 from
70   MyPathFlow::PathNode source,
71   MyPathFlow::PathNode sink
72 where
73   MyPathFlow::flowPath(source, sink)
74 select
75   sink.getNode(),
76   source,
77   sink,
78   "[Alert_message_based_on_vulnerability]",
79   source.getNode(),
80   "[source_description]"

```

D.2 ITERATIVE FEEDBACK EXAMPLE

The following is an excerpt from the prompt used in a new context window that instructs the agent to improve the last iteration's query. The feedback is derived from automatically running synthesized queries against the fixed and vulnerable versions of the CVE's repository, and comparing the query's results against the changed methods in the CVE fix commit diff. The feedback also includes any successful method hits, false positive method hits, and which fixed methods have not been detected by the query.

```

1 ## Objective
2 Refine the CodeQL query based on previous iteration feedback to improve
   vulnerability detection.
3
4 ## Previous Iteration Feedback
5 ## Iteration 2 Results
6
7 ## Previous Query (Iteration 2)
8 ```ql
9 /**
10  * @name AntiSamy XSS Bypass in Style Tags
11  * @description Detects potential XSS vulnerabilities where AntiSamy HTML
   sanitization fails to properly filter malicious style tags with
   event handlers
12 * @problem.severity error
13 * @security-severity 6.1
14 * @precision high
15 * @tags security
16 * @kind path-problem

```

```

17  * @id java/antisamy-style-xss-bypass
18  */
19 import java
20 import semmle.code.java.dataflow.DataFlow
21 import semmle.code.java.dataflow.TaintTracking
22
23 class MaliciousHtmlInput extends DataFlow::Node {
24   MaliciousHtmlInput() {
25     exists(StringLiteral sl |
26       this.asExpr() = sl and
27       (
28         sl.getValue().matches("%<style%onload%>%") or
29         sl.getValue().matches("%<style%onclick%>%") or
30         sl.getValue().matches("%<style%onerror%>%") or
31         sl.getValue().matches("%<style%onmouseover%>%") or
32         sl.getValue().matches("%style%alert%") or
33         sl.getValue().matches("%<style%javascript%") or
34         sl.getValue().matches("%<style%expression%")
35       )
36     ) or
37     exists(Parameter p |
38       this.asParameter() = p and
39       (
40         p.getName().toLowerCase().matches("%html%") or
41         p.getName().toLowerCase().matches("%input%") or
42         p.getName().toLowerCase().matches("%content%")
43       )
44     )
45   }
46 }
47
48 class AntiSamyProcessing extends DataFlow::Node {
49   AntiSamyProcessing() {
50     exists(MethodCall mc |
51       this.asExpr() = mc and
52       mc.getMethod().hasName("scan") and
53       mc.getMethod().getDeclaringType()(RefType)
54       .hasQualifiedName("org.owasp.validator.html", "AntiSamy")
55     ) or
56     exists(MethodCall mc |
57       this.asExpr() = mc and
58       mc.getMethod().hasName("getCleanHTML") and
59       mc.getQualifier().getType()(RefType)
60       .hasQualifiedName("org.owasp.validator.html", "CleanResults")
61     )
62   }
63 }
64
65 class PotentialXssOutput extends DataFlow::Node {
66   PotentialXssOutput() {
67     exists(MethodCall mc |
68       this.asExpr() = mc and
69       (
70         mc.getMethod().hasName("contains") and
71         exists(StringLiteral sl |
72           mc.getAnArgument() = sl and
73           (
74             sl.getValue().matches("%alert%") or
75             sl.getValue().matches("%script%") or
76             sl.getValue().matches("%javascript%") or
77             sl.getValue().matches("%onload%") or
78             sl.getValue().matches("%onclick%")
79           )
80         )
81       )
82     )
83   }
84 }

```

```

82     ) or
83     exists(MethodCall mc |
84         this.asExpr() = mc and
85         mc.getMethod().hasName("getCleanHTML") and
86         mc.getQualifier().getType().(RefType)
87         .hasQualifiedName("org.owasp.validator.html", "CleanResults")
88     )
89   }
90 }
91
92 class ImproperSanitizationBarrier extends DataFlow::Node {
93   ImproperSanitizationBarrier() {
94     exists(MethodCall mc |
95         this.asExpr() = mc and
96         mc.getMethod().hasName("assertTrue") and
97         exists(LogNotExpr lne |
98             mc.getAnArgument() = lne and
99             lne.getExpr().(MethodCall).getMethod().hasName("contains")
100        )
101      )
102   }
103 }
104
105 module AntiSamyXssConfig implements DataFlow::ConfigSig {
106   predicate isSource(DataFlow::Node source) {
107     source instanceof MaliciousHtmlInput
108   }
109
110   predicate isSink(DataFlow::Node sink) {
111     sink instanceof PotentialXssOutput
112   }
113
114   predicate isBarrier(DataFlow::Node sanitizer) {
115     sanitizer instanceof ImproperSanitizationBarrier
116   }
117
118   predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {
119     exists(MethodCall mc |
120       n1.asExpr() = mc.getQualifier() and
121       n2.asExpr() = mc and
122       mc.getMethod().hasName("scan") and
123       mc.getMethod().getDeclaringType().(RefType)
124       .hasQualifiedName("org.owasp.validator.html", "AntiSamy")
125     ) or
126     exists(MethodCall mc |
127       n1.asExpr() = mc.getQualifier() and
128       n2.asExpr() = mc and
129       mc.getMethod().hasName("getCleanHTML")
130     )
131   }
132 }
133
134 module AntiSamyXssFlow = TaintTracking::Global<AntiSamyXssConfig>;
135 import AntiSamyXssFlow::PathGraph
136
137 from
138   AntiSamyXssFlow::PathNode source,
139   AntiSamyXssFlow::PathNode sink
140 where
141   AntiSamyXssFlow::flowPath(source, sink)
142 select
143   sink.getNode(),
144   source,
145   sink,

```

```

146  "Potential_XSS_vulnerability:_HTML_input_with_malicious_style_tags_may_
     bypass_AntiSamy_sanitization",
147  source.getNode(),
148  "malicious_HTML_input"
149  ``
150
151  ## Compilation Results
152 COMPILATION SUCCESS: Query syntax is valid
153
154  ## Execution Results
155  ## Query Evaluation Summary (Iteration 2)
156 Results: Vulnerable=8, Fixed=8
157 Method Recall: Vulnerable=True, Fixed=True
158 True Positive Methods: Vulnerable=2, Fixed=2
159 Coverage: 1/1 target methods
160 PARTIAL: Query hits targets but has false positives in fixed version
161 Method location format is path/to	hit/file.java:[Class of hit method]:[  

     Hit method]
162 Successfully targeted methods:
163  -src/main/java/org/owasp/validator/html/scan/MagicSAXFilter.java:  

     MagicSAXFilter:startElement
164 False positives (hits in fixed version):
165  - src/main/java/org/owasp/validator/html/scan/MagicSAXFilter.java:  

     MagicSAXFilter:startElement
166
167  ## Detailed Evaluation Analysis
168
169  **Method Coverage**: 1/1 target methods
170  **File Coverage**: 1/1 target files
171  **Successfully targeted files**:
172  - MagicSAXFilter.java
173
174  **Successfully targeted methods**:
175  - startElement
176
177  **False positives (hit in fixed version)**:
178  - startElement
179
180  **Evaluation Summary**:
181  - Vulnerable DB File Recall: True
182  - Fixed DB File Recall: True
183  - Vulnerable DB Method Recall: True
184  - Fixed DB Method Recall: True
185  - Total Query Results: Vulnerable=8, Fixed=8
186  - Code Flow Paths: Vulnerable=8, Fixed=8
187
188  ## Next Steps
189 We want vulnerable DB method recall and we don't want fixed DB method
     recall!
190  **Priority**: Reduce false positives by adding more specific conditions
     to avoid hitting the methods listed above.

```

D.3 QLCODER SYNTHESIZED QUERY EXAMPLES

Below are examples of successful synthesized queries. The queries successfully find > 0 true positive method hits on the vulnerable CodeQL database of the CVE's source code, and no false positive method hits on the fixed version's CodeQL database. For reference queries to compare with, CWE queries can be found on the official CodeQL repository (GitHub, 2025e;d)

CVE-2025-27136, CWE-611 - Improper Restriction of XML External Entity Reference

```

1  /**
2  * @name XML External Entity vulnerability in WstxInputFactory without
     secure configuration

```

```

3  * @description WstxInputFactory used in XmlFactory without disabling DTD
4  * support and external entities allows XXE attacks
5  * @problem.severity error
6  * @security-severity 9.1
7  * @precision high
8  * @tags security
9  * @kind path-problem
10 * @id java/wstxinputfactory-xxe
11 */
12 import java
13 import semmle.code.java.dataflow.DataFlow
14 import semmle.code.java.dataflow.FlowSources
15 import semmle.code.java.dataflow.TaintTracking
16 private import semmle.code.java.dataflow.ExternalFlow
17
18 class WstxInputFactoryCreation extends DataFlow:::Node {
19     WstxInputFactoryCreation() {
20         exists(MethodCall mc |
21             mc.getMethod().hasQualifiedName("com.ctc.wstx.stax", "WstxInputFactory", "newInstance") or
22             mc.getMethod().hasQualifiedName("com.ctc.wstx.stax", "WstxInputFactory", "newFactory")
23             |
24             this.asExpr() = mc
25         ) or
26         exists(ClassInstanceExpr cie |
27             cie.getConstructedType().hasQualifiedName("com.ctc.wstx.stax", "WstxInputFactory") and
28             this.asExpr() = cie
29         ) or
30         // Include variable access to WstxInputFactory instances (like "input
31         " parameter)
32         exists(Variable v, VarAccess va |
33             v.getType().(RefType).hasQualifiedName("com.ctc.wstx.stax", "WstxInputFactory") and
34             va.getVariable() = v and
35             this.asExpr() = va
36         )
37     }
38 }
39 class UnsafeXmlFactoryUsage extends DataFlow:::Node {
40     UnsafeXmlFactoryUsage() {
41         exists(ClassInstanceExpr xmlFactoryCall |
42             // XmlFactory constructor with WstxInputFactory parameter
43             xmlFactoryCall.getConstructedType()
44             .hasQualifiedName("com.fasterxml.jackson.dataformat.xml", "XmlFactory") and
45             xmlFactoryCall.getArgument(0) = this.asExpr()
46         ) or
47         exists(ClassInstanceExpr xmlMapperCall, ClassInstanceExpr
48             xmlFactoryCall |
49                 // XmlMapper constructor using XmlFactory with WstxInputFactory
50                 xmlMapperCall.getConstructedType()
51                 .hasQualifiedName("com.fasterxml.jackson.dataformat.xml", "XmlMapper") and
52                 xmlFactoryCall.getConstructedType()
53                 .hasQualifiedName("com.fasterxml.jackson.dataformat.xml", "XmlFactory") and
54                 xmlMapperCall.getArgument(0) = xmlFactoryCall and
55                 xmlFactoryCall.getArgument(0) = this.asExpr()
56         )
57     }
58 }

```

```

58
59 class WstxInputFactorySanitizer extends DataFlow::Node {
60     WstxInputFactorySanitizer() {
61         exists(MethodCall setPropertyCall, VarAccess factoryVar |
62             setPropertyCall.getMethod().hasQualifiedName("javax.xml.stream", "XMLInputFactory", "setProperty") and
63             setPropertyCall.getQualifier() = factoryVar and
64             (
65                 // DTD support disabled
66                 (exists(Field f |
67                     setPropertyCall.getArgument(0) = f.getAnAccess() and
68                     f.hasName("SUPPORT_DTD") and
69                     f.getDeclaringType().hasQualifiedName("javax.xml.stream", "XMLInputFactory")
70                     ) and
71                     exists(Field f |
72                         setPropertyCall.getArgument(1) = f.getAnAccess() and
73                         f.hasName("FALSE") and
74                         f.getDeclaringType().hasQualifiedName("java.lang", "Boolean")
75                     ) or
76                     // External entities disabled
77                     (exists(Field f |
78                         setPropertyCall.getArgument(0) = f.getAnAccess() and
79                         f.hasName("IS_SUPPORTING_EXTERNAL_ENTITIES") and
80                         f.getDeclaringType().hasQualifiedName("javax.xml.stream", "XMLInputFactory")
81                         ) and
82                         exists(Field f |
83                             setPropertyCall.getArgument(1) = f.getAnAccess() and
84                             f.hasName("FALSE") and
85                             f.getDeclaringType().hasQualifiedName("java.lang", "Boolean")
86                         ) )
87                     ) and
88                     this.asExpr() = factoryVar
89                 )
90             )
91     }
92
93 module WstxInputFactoryFlowConfig implements DataFlow::ConfigSig {
94     predicate isSource(DataFlow::Node source) {
95         source instanceof WstxInputFactoryCreation
96     }
97
98     predicate isSink(DataFlow::Node sink) {
99         sink instanceof UnsafeXmlFactoryUsage
100    }
101
102    predicate isBarrier(DataFlow::Node sanitizer) {
103        sanitizer instanceof WstxInputFactorySanitizer
104    }
105
106    predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {
107        // WstxInputFactory passed to XmlFactory constructor
108        exists(ClassInstanceExpr xmlFactoryCall |
109            xmlFactoryCall.getConstructedType()
110            .hasQualifiedName("com.fasterxml.jackson.dataformat.xml", "XmlFactory") and
111            xmlFactoryCall.getArgument(0) = n1.asExpr() and
112            n2.asExpr() = xmlFactoryCall
113        ) or
114        // XmlFactory passed to XmlMapper constructor
115        exists(ClassInstanceExpr xmlMapperCall |
116            xmlMapperCall.getConstructedType()
117            .hasQualifiedName("com.fasterxml.jackson.dataformat.xml", "XmlMapper") and

```

```

118     xmlMapperCall.getArgument(0) = n1.asExpr() and
119     n2.asExpr() = xmlMapperCall
120   )
121 }
122 }
123
124 module WstxInputFactoryFlow = TaintTracking::Global<
125   WstxInputFactoryFlowConfig>;
126 import WstxInputFactoryFlow::PathGraph
127
128 from
129   WstxInputFactoryFlow::PathNode source,
130   WstxInputFactoryFlow::PathNode sink
131 where
132   WstxInputFactoryFlow::flowPath(source, sink)
133 select
134   sink.getNode(),
135   source,
136   sink,
137   "WstxInputFactory_used_without_secure_configuration_flows_to_XML_parser
138   ,_allowing_XXE_attacks",
139   source.getNode(),
140   "WstxInputFactory_usage"

```

CVE-2025-0851, CWE-22 - Path Traversal

```

1 /**
2  * @name Archive path traversal vulnerability (ZipSlip) - CVE-2025-0851
3  * @description Archive entries with path traversal sequences can write
4  * files outside the intended extraction directory
5  * @problem.severity error
6  * @security-severity 9.8
7  * @precision high
8  * @tags security
9  * @kind path-problem
10 * @id java/archive-path-traversal-cve-2025-0851
11 */
12
13 import java
14 import semmle.code.java.dataflow.DataFlow
15 import semmle.code.java.dataflow.TaintTracking
16
17 /**
18  * Sources: Archive entry names from ZipEntry.getName() and
19  * TarArchiveEntry.getName()
20 */
21 class ArchiveEntryNameSource extends DataFlow::Node {
22   ArchiveEntryNameSource() {
23     exists(MethodCall mc |
24       mc.getMethod().getName() = "getName" and
25       mc.getMethod().getDeclaringType().hasQualifiedName("java.util.zip", "ZipEntry") or
26       mc.getMethod().getDeclaringType().hasQualifiedName("org.apache.
27       commons.compress.archivers.tar", "TarArchiveEntry")
28     ) and
29     this.asExpr() = mc
30   }
31 }
32
33 /**
34  * Sinks: Path resolution operations that lead to file creation
35 */
36 class PathCreationSink extends DataFlow::Node {
37   PathCreationSink() {

```

```

35  // Arguments to Path.resolve() calls
36  exists(MethodCall resolveCall |
37      resolveCall.getMethod().getName() = "resolve" and
38      resolveCall.getMethod().getDeclaringType()
39      .hasQualifiedName("java.nio.file", "Path") and
40      this.asExpr() = resolveCall.getAnArgument()
41  )
42  or
43  // Arguments to file creation operations
44  exists(MethodCall fileOp |
45  (
46      fileOp.getMethod().getName() = "createDirectories" or
47      fileOp.getMethod().getName() = "newOutputStream" or
48      fileOp.getMethod().getName() = "write" or
49      fileOp.getMethod().getName() = "copy"
50  ) and
51      fileOp.getMethod().getDeclaringType().hasQualifiedName("java.nio.
52      file", "Files") and
53      this.asExpr() = fileOp.getAnArgument()
54  )
55 }
56
57 /**
58 * Sanitizers: Proper validation that prevents path traversal
59 */
60 class PathTraversalSanitizer extends DataFlow::Node {
61     PathTraversalSanitizer() {
62         // The validateArchiveEntry method call that properly validates paths
63         // This blocks flow after the validation call is made
64         exists(MethodCall validateCall |
65             validateCall.getMethod().getName() = "validateArchiveEntry" and
66             (
67                 // Any variable assigned from validateArchiveEntry call result
68                 exists(Variable v |
69                     this.asExpr() = v.getAnAccess() and
70                     exists(AssignExpr assign |
71                         assign.getDest() = v.getAnAccess() and
72                         assign.getRhs() = validateCall
73                     )
74                 )
75                 or
76                 // Variables passed through validateArchiveEntry calls
77                 this.asExpr() = validateCall.getAnArgument() and
78                 exists(ExprStmt stmt | stmt.getExpr() = validateCall)
79             )
80         )
81         or
82         // Proper "..." validation with exception throwing (complete pattern)
83         exists(MethodCall containsCall, IfStmt ifStmt, ThrowStmt throwStmt |
84             containsCall.getMethod().getName() = "contains" and
85             containsCall.getAnArgument().(StringLiteral).getValue() = "..." and
86             ifStmt.getCondition().getAChildExpr*() = containsCall and
87             ifStmt.getThen().getAChild*() = throwStmt and
88             this.asExpr() = containsCall.getQualifier()
89         )
90         or
91         // Path normalization combined with startsWith validation
92         exists(MethodCall normalizeCall, MethodCall startsWithCall |
93             normalizeCall.getMethod().getName() = "normalize" and
94             normalizeCall.getMethod().getDeclaringType()
95             .hasQualifiedName("java.nio.file", "Path") and
96             startsWithCall.getMethod().getName() = "startsWith" and
97             startsWithCall.getMethod().getDeclaringType().
98             hasQualifiedName("java.nio.file", "Path") and

```

```

99      DataFlow::localFlow(DataFlow::exprNode(normalizeCall),
100      DataFlow::exprNode(startsWithCall.getQualifier()) and
101      this.asExpr() = normalizeCall.getQualifier()
102  )
103 }
105
106 /**
107 * Additional predicate to detect validation barriers at method level
108 */
109 predicate hasValidationCall(Callable method) {
110   exists(MethodCall validateCall |
111     validateCall.getEnclosingCallable() = method and
112     validateCall.getMethod().getName() = "validateArchiveEntry"
113   )
114 }
115
116 module PathTraversalConfig implements DataFlow::ConfigSig {
117   predicate isSource(DataFlow::Node source) {
118     source instanceof ArchiveEntryNameSource
119   }
120
121   predicate isSink(DataFlow::Node sink) {
122     sink instanceof PathCreationSink
123   }
124
125   predicate isBarrier(DataFlow::Node sanitizer) {
126     sanitizer instanceof PathTraversalSanitizer
127   }
128
129   predicate isBarrierIn(DataFlow::Node node) {
130     // Barrier at method entry if method contains validateArchiveEntry
131     // call
132     node instanceof DataFlow::ParameterNode and
133     hasValidationCall(node.getEnclosingCallable())
134   }
135
136   predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {
137     // Flow through variable assignments and declarations
138     exists(LocalVariableDeclExpr decl |
139       decl.getInit() = n1.asExpr() and
140       n2.asExpr() = decl.getVariable().getAnAccess()
141     )
142     or
143     exists(AssignExpr assign |
144       assign.getRhs() = n1.asExpr() and
145       n2.asExpr() = assign.getDest()
146     )
147     or
148     // Flow through string manipulation methods that preserve taint
149     exists(MethodCall mc |
150       mc.getAnArgument() = n1.asExpr() and
151       n2.asExpr() = mc and
152       (
153         mc.getMethod().getName() = "removeLeadingFileSeparator" or
154         mc.getMethod().getName() = "trim" or
155         mc.getMethod().getName() = "toString" or
156         mc.getMethod().getName() = "substring"
157       )
158     )
159     or
160     // Flow through Path operations
161     exists(MethodCall pathOp |
162       pathOp.getAnArgument() = n1.asExpr() and
163       n2.asExpr() = pathOp and

```

```

163     pathOp.getMethod().getName() = "resolve" and
164     pathOp.getMethod().getDeclaringType()
165     .hasQualifiedName("java.nio.file", "Path")
166   )
167 }
168 }
169
170 module PathTraversalFlow = TaintTracking::Global<PathTraversalConfig>;
171
172 import PathTraversalFlow::PathGraph
173
174 from PathTraversalFlow::PathNode source, PathTraversalFlow::PathNode sink
175 where
176   PathTraversalFlow::flowPath(source, sink) and
177   // Focus on the specific vulnerable files and methods
178   (
179     source.getNode().getEnclosingCallable().getDeclaringType()
180     .hasName("TarUtils") or
181     source.getNode().getEnclosingCallable().getDeclaringType()
182     .hasName("ZipUtils")
183   ) and
184   (
185     source.getNode().getEnclosingCallable().getName() = "untar" or
186     source.getNode().getEnclosingCallable().getName() = "unzip"
187   ) and
188   // Only report flows where validation is NOT properly done
189   not hasValidationCall(source.getNode().getEnclosingCallable())
190 select sink.getNode(), source, sink,
191   "Archive_entry_name_from_${@_flows_to_file_system_operation_without_
192   proper_path_traversal_validation, allowing_ZipSlip_attack.},
192   source.getNode(), "archive_entry_name"

```

CCVE-2025-27528, CWE-502 - Deserialization of Untrusted Data

```

1 /**
2  * @name MySQL JDBC URL parameter injection vulnerability
3  * @description Detects MySQL JDBC URLs with dangerous bracket parameters
4  *             that bypass inadequate filtering in vulnerable code
5  * @problem.severity error
6  * @security-severity 8.8
7  * @precision high
8  * @tags security
9  * @kind path-problem
10 * @id java/mysql-jdbc-url-injection
11 */
12 import java
13 import semmle.code.java.dataflow.DataFlow
14 import semmle.code.java.dataflow.TaintTracking
15
16 class MySQLDangerousBracketUrlSource extends DataFlow::Node {
17   MySQLDangerousBracketUrlSource() {
18     // String literals with dangerous MySQL parameters in bracket
19     // notation
20     exists(StringLiteral lit |
21       lit.getValue().matches("*mysql*") and
22       lit.getValue().matches("*[*]*") and
23       (
24         lit.getValue().matches("*allowLoadLocalInfile*") or
25         lit.getValue().matches("*allowUrlInLocalInfile*") or
26         lit.getValue().matches("*autoDeserialize*") or
27         lit.getValue().matches("*allowPublicKeyRetrieval*") or
28         lit.getValue().matches("*serverTimezone*") or
29         lit.getValue().matches("*user*") or
30         lit.getValue().matches("*password*")

```

```

30         ) and
31         this.asExpr() = lit
32     )
33     or
34     // Parameters to filterSensitive method that may contain dangerous
35     bracket content
36     exists(Method m, Parameter p |
37         m.hasName("filterSensitive") and
38         m.getDeclaringType().getName() = "MySQLSensitiveUrlUtils" and
39         p = m.getParameter() and
40         this.asParameter() = p
41     )
42   }
43
44 class VulnerableCodePatternSink extends DataFlow::Node {
45   VulnerableCodePatternSink() {
46     // The vulnerability: calls to filterSensitive in vulnerable code
47     // patterns
48     exists(Method m, MethodCall filterCall |
49         m.hasName("filterSensitive") and
50         m.getDeclaringType().getName() = "MySQLSensitiveUrlUtils" and
51         filterCall.getMethod() = m and
52         this.asExpr() = filterCall and
53         // Key vulnerability condition: this code exists where
54         filterSensitiveKeyByBracket method is NOT available
55         // In the vulnerable version, filterSensitiveKeyByBracket doesn't
56         exist
57         not exists(Method bracketMethod |
58             bracketMethod.hasName("filterSensitiveKeyByBracket") and
59             bracketMethod.getDeclaringType().getName() =
60             "MySQLSensitiveUrlUtils" and
61             bracketMethod.getDeclaringType() = m.getDeclaringType()
62         )
63     )
64     or
65     // Additional sink: method calls that use the result of inadequate
66     // filtering
67     exists(MethodCall mc, MethodCall filterCall |
68         filterCall.getMethod().hasName("filterSensitive") and
69         filterCall.getMethod().getDeclaringType().getName() =
70         "MySQLSensitiveUrlUtils" and
71         DataFlow::localFlow(DataFlow::exprNode(filterCall), DataFlow::
72         exprNode(mc.getArgument(_))) and
73         this.asExpr() = mc and
74         // Only vulnerable if no proper bracket filtering exists in the
75         same class
76         not exists(Method bracketMethod |
77             bracketMethod.hasName("filterSensitiveKeyByBracket") and
78             bracketMethod.getDeclaringType().getName() =
79             "MySQLSensitiveUrlUtils" and
80             bracketMethod.getDeclaringType() = filterCall.getMethod().
81             getDeclaringType()
82         )
83     )
84   }
85 }
86
87 class ProperBracketFilteringSanitizer extends DataFlow::Node {
88   ProperBracketFilteringSanitizer() {
89     // The proper bracket-based sanitization method (present only in
90     // fixed version)
91     exists(MethodCall mc |
92         mc.getMethod().hasName("filterSensitiveKeyByBracket") and

```

```

82     mc.getMethod().getDeclaringType().getName() = "MySQLSensitiveUrlUtils" and
83     this.asExpr() = mc
84   )
85 }
86 }
87
88 module MySQLJDBCUrlInjectionConfig implements DataFlow::ConfigSig {
89   predicate isSource(DataFlow::Node source) {
90     source instanceof MySQLDangerousBracketUrlSource
91   }
92
93   predicate isSink(DataFlow::Node sink) {
94     sink instanceof VulnerableCodePatternSink
95   }
96
97   predicate isBarrier(DataFlow::Node sanitizer) {
98     sanitizer instanceof ProperBracketFilteringSanitizer
99   }
100
101 predicate isAdditionalFlowStep(DataFlow::Node n1, DataFlow::Node n2) {
102   // Flow through string concatenation operations
103   exists(AddExpr addExpr |
104     n1.asExpr() = addExpr.getLeftOperand() and
105     n2.asExpr() = addExpr
106   )
107   or
108   exists(AddExpr addExpr |
109     n1.asExpr() = addExpr.getRightOperand() and
110     n2.asExpr() = addExpr
111   )
112   or
113   // Flow through variable assignments
114   exists(Assignment assign |
115     n1.asExpr() = assign.getSource() and
116     n2.asExpr() = assign.getDest()
117   )
118   or
119   // Flow through return statements
120   exists(ReturnStmt ret |
121     n1.asExpr() = ret.getResult() and
122     n2.asParameter() = ret.getEnclosingCallable().getAParameter()
123   )
124 }
125 }
126
127 module MySQLJDBCUrlInjectionFlow = TaintTracking::Global<
128   MySQLJDBCUrlInjectionConfig>;
129
129 import MySQLJDBCUrlInjectionFlow::PathGraph
130
131 from MySQLJDBCUrlInjectionFlow::PathNode source,
132   MySQLJDBCUrlInjectionFlow::PathNode sink
133 where MySQLJDBCUrlInjectionFlow::flowPath(source, sink)
134 select sink.getNode(), source, sink,
134   "MySQL_JDBC_URL_with_dangerous_bracket_parameters_flows_to_vulnerable_
134   filtering_logic_at_$@_that_lacks_proper_bracket-based_sanitization",
135   source.getNode(), "dangerous_URL_source"

```

D.4 AST EXTRACTION QUERY

Given a fix diff, QLCoder automatically parses the changed methods and files, and inserts them into an AST pretty printing query template. Below is an example of the AST extraction query used for CVE-2014-7816.

```

1 /**
2  * @name Expressions and statements for CVE-2014-7816 changed code areas
3  * @description Extract expressions and statements from vulnerability fix
4  * areas
5  * @id java/expr-stmt-diff-CVE_2014_7816
6  * @kind problem
7  * @problem.severity recommendation
8  */
9 import java
10
11 from Element e, Location l
12 where
13     l = e.getLocation() and ((l.getFile().getBaseName() = "
14         PathSeparatorHandler.java"
15     and l.getStartLine() >= 1 and l.getEndLine() <= 100) or
16
17     (l.getFile().getBaseName() = "URLDecodingHandler.java"
18     and l.getStartLine() >= 17 and l.getEndLine() <= 128)
19
20     or (l.getFile().getBaseName() = "ResourceHandler.java"
21     and l.getStartLine() >= 158 and l.getEndLine() <= 172)
22
23     or (l.getFile().getBaseName() = "io.undertow.server.handlers.builder.
24         HandlerBuilder" and l.getStartLine() >= 17
25     and l.getEndLine() <= 29)
26
27     or (l.getFile().getBaseName() = "DefaultServlet.java"
28     and l.getStartLine() >= 39
29     and l.getEndLine() <= 150)
30
31     or (l.getFile().getBaseName() = "ServletPathMatches.java"
32     and l.getStartLine() >= 32
33     and l.getEndLine() <= 140))
34 select e,
35     e.toString() as element,
36     e.getAPrimaryQlClass() as elementType,
37     l.getFile().getBaseName() as file,
38     l.getStartLine() as startLine,
39     l.getEndLine() as endLine,
40     l.getStartColumn() as startColumn,
41     l.getEndColumn() as endColumn

```