

SOFTWARE KNOWLEDGE BASE EXTRACTION & MANAGEMENT FOR ENTERPRISE

Boosting Development and Operation Efficiency
with System Knowledge Base

Table of Contents

- 01** **About Flezi Metis**
- 02** **Our Service Offering**
- 03** **The Pricing Model**
- 04** **Flezi Metis's Components**
- 05** **Success Stories**



Matters Faced by Most Enterprises




ACCURACY MATTERS

AI output is not always reliable



SIZE MATTERS

Real-world software systems are large and complex



DOMAIN MATTERS

AI lacks understanding of business processes and domain knowledge



LEGACY MATTERS

Existing models have limited understanding & support for legacy programming languages



PRIVACY MATTERS

Enterprise software systems, source code, and data need strong protection



COST MATTERS


Reasoning AI models and agents are costly and slowly



DIVERSITY MATTERS

Multiple technologies and languages across legacy systems drive higher costs when solved separately

AI Cost Crisis: ROI Problem



CONTEXT WINDOWS AS A SILENT TAX

Prompt Inflation Tax: "Context Stuffing" – including entire repositories or extensive dependency graphs in every prompt – can inflate input token cost by over **360%** without a proportional increase in quality or correctness.

360%+ INFLATION

UNBOUNDED EXECUTION RISK

65% of IT Leaders report "bill shock" from usage-based models. Autonomous agents silently retry tasks or "rehydrate" context by repeating re-sending large files.

COST EXCEED ESTIMATES BY 30% to 50%


THE 40% BUILD VS BUY THRESHOLD

Scalability Inflection Point: Hosted SaaS solutions typically become **40% more expensive** than self-managed infrastructure once an organization exceeds some million (reference: 1.2 millions) tokens processed monthly.

1.2M TOKEN/MONTH THRESHOLD

The Scaling Paradox of AI in Large Codebases

As codebase size increases, productivity gain from AI decreases



Context Window Limitations

- Performance gains decrease with larger context windows

Signal:Noise Ratio

- Large codebases have more noise that can mislead the model

Complexity

- Dependencies
- Domain-specific logic

Stanford Study of 100K Developers and 600 Companies:

<https://www.linkedin.com/pulse/does-ai-actually-boost-developer-productivity-striking-%C3%A7elebi-tcp8f/>

AI Agents Alone Aren't Enough

The problem isn't AI — it's AI without context. Flezi Metis gives your agents the knowledge to perform.

TODAY: SOFTWARE AI AGENTS ARE STRUGGLING

Huge Token Cost


Token costs balloon from context stuffing.

Slow

Low quality results lead to significantly more time for verification

Low Quality

Garbage in, garbage out. Loops make agents hallucinate



WITH FLEZI METIS — The Ultimate Software Knowledge Base



Faster

1st-pass

Correct code on the first pass, not the third, no more multiple agentic loops



Cheaper

20-55% less

Lower token cost.



Deterministic Answer

Accuracy+

Full paths exploration of code base & grounded context



Higher quality

Impact Tested

Impact analysis with deterministic network of knowledge.

AI-Driven Software Engineering Architecture With Metis



Market Problems

CONTEXT BLINDNESS


AI Agents waste **up to 90%** of their time hunting for dependencies and service relationships.

AI RELIABILITY GAP

Without grounded facts, AI "hallucinates" code that violates your systems.

DOCUMENTATION FRAGMENTATION

Scattered, outdated wikis create a mess that stalls onboarding and scales debt.



Advancing the Future of Knowledge Base Extraction & Management

- **Metis** is an advanced AI solution that empowers enterprises to develop, operate, maintain, and upgrade large, complex, legacy software systems.
- **Metis** builds a knowledge base from source code and project documents, enabling human engineers and AI agents to quickly grasp system design and business logic for tasks like design, coding, testing, and debugging.
- **Metis** provides context for AI agents as well as extracts the scope of changes and impact with high precision to cut costs in development and quality management.

✦ Flexible Cost & Infrastructure Control

Supports on-premise/private cloud deployment, optimizes TCO, easy scaling, avoids dependency on third-party providers.

✦ Full SDLC Support for Large Systems


Builds deep knowledge base for analysis, code generation, maintenance, automated testing, modernization, reducing errors and boosting team productivity.

✦ High Security & Privacy

Keeps all data internal, compliant with enterprise standards, access controls, ideal for closed environments.

✦ Customizable & Scalable

Adds domain/project-specific knowledge, continuously adapts, optimizes AI agents for organizational contexts.



CONTEXT BLINDNESS

AI Agents waste **up to 90%** of their time hunting for dependencies and service relationships.



SYSTEM MAPPING

Metis builds a structured graph of your codebase – every dependency, pattern, and relationship.

AI RELIABILITY GAP

Without grounded facts, AI "hallucinates" code that violates your systems.



ACCURATE REASONING

With grounded context from your actual systems, Metis enables AI agents to reason from facts.

DOCUMENTATION FRAGMENTATION

Scattered, outdated wikis create a mess that stalls onboarding and scales debt.



SINGLE SOURCE OF TRUTH

Metis replaces fragmented docs with unified graph for quick code access.

↓78%
Bug Fixing Time

4x
Faster Code Analysis

↓55%
Token Costs

LESS
Document Dependency

LANGUAGES SUPPORTED



1

DETERMINISTIC KNOWLEDGE MAP

100% local graph data, replacing AI "guesswork" with absolute facts and clear relationships.

2

ZERO HALLUCINATION

Provides deep, reliable insights to handle large-scale, mission-critical systems.

3


MASSIVE COST SAVINGS

Slashes token costs and boosts data security thanks to on-premises or private cloud deployment.

4

CUSTOMIZED INTELLIGENCE

Customizes for your project, injects deep domain smarts into the knowledge base, supercharging AI agents.



BEST FOR

AMS


KTLO

Legacy Modernization

Migration


New Development

Secure, Agentic AI with Enterprise Software Knowledge Base



Where Metis Fits in the Agentic Stack


From AI agents down to your ground-truth data — Metis grounds the entire stack.




Your Data is Always in Your Hand

Principles:

- Metis's core is 100% Local
- You choose the LLM



Internal Project Onboarding Pattern



Recommended Settings:

- *Large ODC: Local Server for optimized cost*
- *Project with short time usage of Metis (e.g. Reverse Engineering phase in Migration): rent a cloud machine with GPU*
- *If using Metis for rapid development task (such as SDD): Host Metis on developer machine (32GB RAM recommendation) and use external LLM (e.g. LLM Key)*

Measurement Methodology

A scalable, telemetry-driven framework for proving highly specialized tool value

Phase 1: Baseline

Establish the "Rate Card" using a controlled user group on representative tasks.

Complexity Tier (File hops)	Avg. Time Saved
Tier 1: within single file (Simple Tracing)	(example) 5 mins
Tier 2: within single repo (API/Module Impact)	(example) 25 mins
Tier 3: multiple repo (Arch. Migration)	(example) 1.5 hours

Token Cost Baseline: Compare Metis token size vs. standard AI Agent.

Phase 2: Realized Value

Track actual live usage automatically via backend telemetry. Multiply actual usage by the Phase 1 Rate Card.

1. Complexity-Adjusted Time Saved

Automatically map query depth to Phase 1 Tiers.
(e.g., 40 Tier 2 queries logged = 1,000 mins saved)

2. Token Avoidance

Track tokens NOT consumed vs. traditional AI agents. (Tokens Avoided)

THE EXECUTIVE "NOISE DISCOUNT" & FINAL ROI

To account for "curiosity queries" and non-task browsing, apply a flat, pre-agreed discount (e.g., 15%) to all telemetry. This neutralizes skepticism and makes the remaining metrics bulletproof.

$$([\Sigma \text{ Time Saved } \times \text{ Labor Cost }] + [\Sigma \text{ Tokens Saved } \times \text{ Token Price }]) \times 0.85 = \text{Project ROI}$$

Adoption Project Commitment Guideline


We do this step **after** the project specific parser is matured, or we **accept** the current cases covered by our parser

Phase 1: Baseline Rate Card (1 month)

- **Choose** 10-20 tasks to establish the baseline Rate Card, split equally between 3 tiers.
- **Time Saving Acceptance Criteria:**
 - 10% time saving on average of Tier 1
 - 20% time saving on average of Tier 2
 - 30 % time saving on average of Tier 3
- **Token Cost Saving Acceptance Criteria:**
 - 20% Token saving on average of Tier 1
 - 30% Token saving on average of Tier 2
 - 40 % Token saving on average of Tier 3

Phase 2: Realized Value (Continuously)

- Choose the **Noise Discount value**: usually 15%, could change depending on feedback of the project team.
- Continuously measure **Project ROI**
- **Metis team** is responsible for Metis related AMS support
 - **Roll out**: 1 Month extensive support to fix bugs or enhance experience, or to enhance use cases.
 - **Stable**: Once stable, we will switch to regular AMS
- **Project team** is responsible for adoption of Metis for its own members.



Use Cases Enabled with Metis

 **AMS**  **Modernization**  **Migration**  **New Development**  **Reverse Engineering**

Code Understanding

Grounded Truth that you can trust, quick and efficient way to analyze large code bases



Impact Analysis

Analyze a complete static code trace of possible impacts when changing code



Bug Fixing

Provide LLM the bug trace by enabling graph reasoning, enable fast, accurate AI bug fixing



Light Self-healing

With non-deterministic nature, clear structured code issues can be fixed live



AI Code Generation

Provide LLM the context of your code, no need for agents to overwhelm context windows with run-time analysis



Reduce Documentation Dependency

LLM treats code graph as the primary source of truth instead of possibly outdated documents



Documentation Generation

Enable Prompt-to-Detail-Design or many other documents with grounded truth



Flezi Metis for AMS: Cost, Quality & Control

1. COST OPTIMIZATION

↓ **55%** Token Costs

↓ **78%** Bug Fixing Time

4x Faster Code Analysis

- AI agents use only ~20% of tokens vs. running without Metis
- Automates root cause analysis, impact tracing, test generation & documentation
- New engineers onboard in hours, not weeks on legacy systems

2. ZERO DEFECTS

~70% Debug accuracy on 1M LOC

~70% Spec Query Accuracy

Zero Hallucination

- Deterministic knowledge map - 100% local graph, replaces AI guesswork with facts
- Full dependency tracing before every change - know what breaks before it breaks in prod

3. VENDOR FREEDOM

100% Data Stays on Your Infra

MCP Open Standard Integration

Any Foundation Model

- Works with your own LLM: Claude Code, Github Copilot, etc.
- On-premise or private cloud: zero data egress, full air-gap support
- No dependency on public AI as token prices surge

CodeWiki Transition Accelerator Flezi MS/AMS

Automated Code Documentation: From Raw Source Code to Production-Ready Wiki

Unlike manual documentation, Metis analyzes millions of lines of source code across multi-repositories using Ontology Graph KB and Agentic AI to generate a 99% accurate, production-ready Code Wiki — in a matter of days, not months.

AGENTIC CODE ANALYSIS

Processes millions of LOC across multi-repositories, automatically extracting code logic, data flow, and system architecture with precision.

ONTOLOGY GRAPH KB

Builds a structured knowledge graph linking classes, functions, dependencies and business logic into a unified, query-able knowledge base.

SME VERIFICATION GATE

Human-in-the-loop quality gate checks for Accuracy, Completeness, and Clarity, ensuring every Wiki entry is trusted and production-ready.

INSTANT KNOWLEDGE ACCESS

Wiki is searchable and immediately ready for new developers, maintenance teams, AI agents, and system migration teams.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

↓ **70%**

Onboarding Time Reduction
New developers understand the system in days instead of months

10x

Cost vs. Agentic AI Tools
10x cheaper for documentation generation at equivalent scale.

99%

Wiki Accuracy
Ensuring every entry is accurate, complete, and trusted across the team.



Long-Term Knowledge Health
Eliminates tribal knowledge dependency and documentation fragility.

WHO NEED THIS?



New Developers
Quick Onboarding



Maintenance Team
Faster Bug Fixing



Migration Team
Safer Migration

Organizations invest months in knowledge transfer during system transitions. **Critical documentation gaps create fragility, slow onboarding, and amplify migration risk** — even when the source code is available.

AI-Native SDLC Approach

AI IN EVERY PHASE

AI is not just a coding tool — it participates in analysis, planning, architecture, and review as a core team member.

CENTRALIZED KNOWLEDGE BASE

All project docs — structured, indexed, and accessible as a single source of truth. AI generates and consumes, humans review and curate.

HUMAN OVERSIGHT & CRITICAL THINKING

AI generates, humans evaluate. Every decision requires human judgment — AI accelerates, never replaces thinking.

AI-ASSISTED SDLC

Human drives, AI assists as a tool (autocomplete, suggestions)

AI-NATIVE SDLC

AI is a core team member across the **full SDLC**

The Transformation Pyramid




Structuring AI-Native Services with Metis

Our Vision: Revolutionize the full Software Life Cycle

Our Metis maps the Logical "DNA" of massive codebases (Java, .Net, C/C++, JS, TS).

Removing the "Complexity Wall" of AI, moving FPT from selling "**Man-months**" to selling "**System Certainty**".



FLEZi Metis

Service Catalog

ROOT CAUSE NAVIGATION

Bug fixing via graph-based root cause tracing

CHANGE IMPACT FORECAST

Visualizing the "blast radius" of code changes to prevent regressions.

LEGACY INSIGHT MAPS

Instant code-to-document/design mapping for legacy systems.

PRE-MIGRATION CLEANUP

Automated detection and removal of Dead Code.

TRUTH VERIFICATION AUDIT

Identifying inconsistencies between code and business requirements/docs.

CORE SYSTEM RENEWAL

Automated modernization of Java, .NET, and C++ architectural cores.

Root Cause Navigation Flezi MS/AMS

Intelligent Bug Diagnosis & Automated Root Cause Analysis

Unlike generic AI, Metis constructs a mathematically accurate, visual map of your software “DNA” (Java, .NET, C/C++, JS/TS). It traces execution paths to pinpoint root causes instantly.

GRAPH-BASED TRACING

Maps the failure path across entire repos.

75% Faster MTTR

CHANGE IMPACT FORECAST

Visualizes the “Blast Radius” of proposed fixes.

Zero-Regression Guarantee

LOGIC-TO-DOC VALIDATION

Cross-references code against business docs.

Audit-Ready Compliance

CROSS-LANGUAGE DEPTH

Supports C++ pointers & strict .NET/TS typing.

Enterprise-Grade Scaling

“Organizations are **losing 60–70% of engineering time** on root cause investigation, with **high MTTR** driven by complex system dependencies. Critical knowledge remains locked in **tribal expertise**, creating fragility, while fixes often introduce **unintended regressions** across interconnected systems.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

- ↑50%

Operational Margin Expansion
+50% teams handle higher ticket volumes with same headcount.
- 100%

SLA Compliance
Eliminates risk of financial penalties for Critical tickets.
- 40%

Lower LLM Token Costs
Graph-based context injection slashes AI overhead.
- ↑

Improved Long-Term Health
Identifies true root causes, not “band-aid” patches.

FLEXIBLE DEPLOYMENT

- 

On-Demand
Crisis Response
- 

Flezi MS
Integration
- 

“Clean-State”
Audit

Change Impact Forecast Flezi MS/AMS Migration

Eliminate Regression Risk through Graph-Based Dependency Analysis

Metis constructs a multi-dimensional graph of architecture to visualize 'Blast Radius' across all modules.

Change with Confidence. Modernize without Fear.

'BLAST RADIUS' VISUALIZATION

Downstream dependencies identified across the stack.

90% Reduction in Regressions

DEAD CODE ISOLATION

Safely identifies orphaned code for removal.

Smaller Migration Footprint

SURGICAL TEST SCOPING

Points exactly to test cases requiring execution.

50% Reduction in QA Effort

CROSS-REPO IMPACT

Traces dependencies across external APIs & Microservices.

Safe Enterprise Scaling

Legacy systems often suffer from a **"fear of the unknown,"** where massive, 1M+ line codebases harbor **hidden dependencies** that manual reviews simply cannot detect. This lack of visibility creates a **"paralysis of production,"** forcing organizations into expensive, brute-force QA cycles and delaying critical framework upgrades. Without a clear map of the **ripple effects**, every code change becomes a high-stakes gamble against system stability.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

↑ 50% **Accelerated Release Velocity**
Assessment time reduced from weeks to minutes.

↑ 60% **Cost Efficiency**
Eliminates brute-force regression testing in favor of surgical validation.

↑ 75% **Risk Mitigation (Zero-Downtime)**
Zero regressions during Java 8 to 21 modernization.

↑ 5% **Technical Debt Transparency**
Leadership makes data-driven refactoring decisions.

"SAFETY-FIRST" DEPLOYMENT

Migration Guardrail (Modernize-X)

Maintenance Safety-Net (Flezi MS)

Compliance & Audit (Truth Verification)

Legacy Insight Maps

ODC-X

Migration

Instant Architectural Discovery & Knowledge Base Generation

Metis acts as an “X-Ray” for your software estate, bypassing text and extracting ‘ground truth’ directly from source code. Navigable multi-dimensional map.

Turn Legacy Code into Searchable Intelligence. Overnight.

INSTANT DESIGN MAPPING

Generates UML, diagrams, and maps from code.

90% Faster Discovery

LOGIC EXTRACTION

Identifies hidden business rules and API endpoints.

Instant Knowledge Equity

TRIBAL KNOWLEDGE CAPTURE

Codifies system behavior before senior staff departures.

Zero-Risk Transition

REPO-WIDE SEARCH

Natural language queries find ‘where X logic is handled’.

Reduction in Knowledge Transfer

Legacy systems often operate as a “**Black Box**,” where critical business logic is trapped in **outdated documentation** or restricted to the “**Tribal Knowledge**” of a few senior engineers. This lack of transparency creates a massive **transition bottleneck**, often requiring up to six months of manual discovery that stalls ODC onboarding and high-priority migrations.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

40%

Faster Time-to-Productivity
Protects revenue realization.

150+

Manual Discovery Hours Saved
Replaced with 2-week automated sprint.

100%

Accurate Scoping
Eliminates “Discovery Surprises” during migration.



Audit & Compliance Ready
Critical for Banking, Healthcare, Automotive regulatory reviews.

SERVICE INTEGRATION



ODCX Transition
Accelerator



Modernize-X
Pre-Migration Audit



Post
Merger/Acquisition
Due Diligence

Pre-Migration Cleanup

Flezi MS/AMS

Migration

Asset Optimization through Automated Dead-Code Detection & Scope Reduction

Metis uses mathematically certain reachability analysis to identify and strip away orphaned logic with 100% accuracy.

ORPHANED LOGIC DETECTION

Unreferenced functions, variables, modules across entire repo.

20% Reduction in Migration Scope

DEPENDENCY PRUNING

Finds/removes redundant libraries & 'zombie' dependencies.

Reduced Security Vulnerabilities

COMPLEXITY COMPRESSION

Consolidates redundant logic paths.

Lower Total Cost of Ownership

PRE-MIGRATION CLEANUP

Strips 'Bloat' before moving to Cloud / Microservices.

Direct Savings on Cloud Infrastructure

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

↑ **20%** **Lower Migration Price**
Direct Cost Savings
Delivery velocity increase.

↑ **60%** **Faster Execution, Lower Memory Footprints**
Smaller binaries, leaner microservices.

↑ **20%** **Enhanced ODCX Productivity**
Focused on active, relevant code;
spend less time learning

↑ **50%** **SLA Reliability**
Flezi MS teams find real bugs faster;
Eliminating Dead Code 'Noise'.

3 PHASES OF SERVICE DEPLOYMENT



Weekly "Health Check" Audit



Modernize-X Pre-Flight



Asset Optimization (Flezi MS)

Organizations are currently paying a "**Migration Tax**" by spending resources to move and test non-business-value code, while developers lose 20% of their productivity navigating "**Dead Code**." This accumulation of "**Legacy Bloat**" unnecessarily expands the **security attack surface** and inflates cloud consumption costs for compute and storage.

Truth Verification Audit

ODC-X

Migration

Synchronizing Technical Reality with Business Intent

Metis performs a multi-dimensional cross-reference between natural language requirements and the actual source code to extract the "Ground Truth". By flagging every deviation, it ensures your cloud version matches your business intent 1:1.

LOGIC PARITY ANALYSIS

Automatically flags contradictions between code & specs.

Zero-Audit Failures

UNDOCUMENTED LOGIC DISCOVERY

Identifies "Side-Effect" logic missing from specs.

100% Logic Transparency

AUTOMATED DOC-SYNC

Updates legacy documentation to reflect current technical reality.

Instant Knowledge Equity

REQUIREMENT TRACEABILITY

Maps every business rule to the exact line of code enforcing it.

Audit-Ready Readiness

Organizations often fall into the "**Shadow Logic Trap**," where static documentation fails to keep pace with evolving code, leading to dangerous **compliance gaps** in regulated industries like Banking and Healthcare. This discrepancy creates "**Broken-on-Arrival**" **deployments** during cloud migrations because the move is based on outdated specifications rather than technical reality.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.



Hidden Logic Flaw Detection

Identifies hidden flaws before regulatory fines or production outages.

↑ 50%

Faster UAT

Automated logic parity proof slashes manual verification time.



Modernization Confidence

Cloud version matches legacy business logic 1:1.



Truth-Verified Knowledge Base

New teams learn accurately, eliminate guessing on legacy logic.

3 AUDIT STREAMS OF SERVICE DEPLOYMENT



Regulatory Compliance Sprint



Modernize-X Pre-Flight



Post-Release Validation (Flezi MS)

Core System Renewal

Automated Architectural Restructuring with Mathematical Certainty

Metis utilizes Mathematically Certain Reachability analysis to identify and isolate legacy logic cores with absolute precision. This approach allows organizations to accelerate multi-year monolith moves into manageable sprints.

SURGICAL TARGET IDENTIFICATION

Locates God Classes & cyclic dependencies.

100% Fact-Based Scoping

CHANGE IMPACT FORECAST

Real-time 'Blast Radius' Visualization of changes.

90% Reduction in Regressions

AUTOMATED LOGIC EXTRACTION

Safely separates business from infrastructure logic.

50% Faster Modernization

COMPLEXITY COMPRESSION

Flattens convoluted hierarchies & redundancy.

Lower TCO

Massive codebases are often frozen by "**Ripple Effect Fear**," where the risk of breaking undocumented logic leads to **Technical Debt Paralysis** and 40% wasted maintenance effort. This creates an **Unreliable Scoping** environment, turning every modernization attempt into a high-risk "guessing game" that threatens operational stability.

Measurable Outcomes

These outcomes are for standard projects. They may be adjusted depending on project's assessment.

↑ 20%

Profit Gain
EBIT Expansion (AMS), moves to Outcome-Based pricing.

30%

Lower Maintenance Costs
Technical Debt Reduction



Modernization Velocity
Accelerates multi-year monolith moves into multi-month sprints.



Risk Elimination
Risk Surgical validation over 'test everything' chaos.

SERVICE DEPLOYMENT



ODCX Transition Accelerator



Flezi MS Premium Upsell



Modernize-X Monolith-to-Microservices

Metis.CodeIQ – Transforming Code into Knowledge Base

- **Metis.CodeIQ** is the core code intelligence engine within **Metis**, designed to revolutionize how AI tools interact with and understand complex codebases.
- **Metis.CodeIQ** is not just an engine; it's a strategic asset for any organization looking to leverage AI for more efficient, accurate, and scalable software development.
- By transforming raw source code into a rich, structured knowledge graph, **Metis.CodeIQ** provides **unparalleled context**, enabling **superior AI-powered development workflows**.

From Code to Knowledge Graph

CodeIQ views source code as a network of interconnected elements. It extracts these relationships and constraints, converting them into a dynamic knowledge graph.

Empowering AI with Context

This graph provides deep code context, crucial for AI tools to perform precise code Q&A, automated generation, and accurate bug fixing, delivering better results.

Flexible & Scalable Architecture

Built with microservices and a clean architecture, CodeIQ supports both cloud and on-premise deployments, handling large repositories and thousands of concurrent users.

What Makes CodeIQ Different?



PRECISE BUG FIXING

Our graph-based context leads to superior accuracy in identifying and fixing bugs, outperforming generic AI chatbots.



COST OPTIMIZATION

By providing rich context directly, CodeIQ reduces token consumption significantly, using a single request per query.



SCALABLE & OPEN

Designed for horizontal scaling and continuous upgrades, it integrates seamlessly with existing tools and workflows.



ON-PREMISE READY

Meets stringent security requirements, enabling secure deployment in on-premise environments.

Metis.CodeIQ: System Architecture 1/2

- **Metis.CodeIQ** is designed to handle large repositories (2M-10M lines of code) and supports multi-tenant environments with data isolation. It's built to sustain approximately 1,000 concurrent users (CCU) or more.
- **Metis. CodeIQ** offers two primary modes of operation to maximize flexibility and utility for developers.

FULL FLOW MODE

CodeIQ directly receives user prompts, processes tasks, and returns final results. For example, converting a C++ module to Java.

INTEGRATED MODE

CodeIQ acts as a knowledge base, providing business and source code data to integrated tools like CodeVista, TestVista, AgentVista, and other third-party tools.

- **Metis. CodeIQ**'s architecture is designed for **scalability, security, and continuous evolution**, ensuring it meets the demands of modern software development.

Metis.TextIQ – Transforming Documents into Knowledge Base

- **Metis.TextIQ** is the core document intelligence engine within Metis, designed to revolutionize how AI tools interact with and understand enterprise documents.
- **Metis.TextIQ** transforms thousands of enterprise documents into a centralized, AI-searchable knowledge base, with **zero hallucination, full traceability, and enterprise-grade security**.

Knowledge Engine

Ingest thousands of documents (Excel, Word, PDF, PPTX) with hybrid AI extraction. Hierarchical chunking, semantic indexing, and AI-powered tagging with human-in-the-loop quality control.

MULTI-FORMAT

COMPLEX EXCEL

HITL TAGGING

AI Conversation Q&A

Ask questions in natural language. Agentic RAG delivers streaming answers with inline source citations, grounded in your approved knowledge base. Zero hallucination by design.

AGENTIC RAG

MULTI-TURN

CITATIONS

Living Documentation

Auto-generate PRD, HLD, LLD, Test Plans, API Specs, and more. Detect documentation drift when knowledge changes. Truth-grounded synthesis, every sentence traceable to source.

8 TEMPLATES

DRIFT DETECTION


MULTI EXPORT

What Makes TextIQ Different?

Capability	Code Mind	Notion / Confluence
On-premise / Air-Gap	☒	☐
Complex Excel Extraction	☒	☐
Zero Hallucination Design	☒	☐
HITL Smart Tagging	☒	☐
AI Consistency Checker	☒	☐
Dual-hierarchy RBAC	☒	☐
Living Doc Drift Detection	☒	☐

Metis.TextIQ: Built for Enterprise Scale

- **Dual-Hierarchy RBAC:** Org + Project hierarchies with privacy-by-default data isolation
- **Trilingual (EN/JA/VI):** Native tokenization, domain terms, and terminology databases
- **Japanese Excel Mastery:** Merged cells, color codes, multi-sheet — pixel-perfect extraction
- **IDE Integration (MCP):** Query KB from VS Code, Cursor, JetBrains — no context switching
- **Zero Hallucination:** AI reports “I don’t know” rather than fabricate — enterprise trust
- **100,000+ User Scale:** Multi-tenant architecture, 100+ tenants, 500+ concurrent users



- **7-Layer Defense-in-Depth:** Every request passes through multiple independent security barriers. 5 of 7 layers fail closed - no request passes without explicit validation.



Success Story: Enhancing AMS Efficiency with Metis

US-based Private Global Conglomerate

A highly diversified, private global conglomerate that functions as a data and infrastructure engine, offering services in automotive and communications



\$23B+
Revenue



Global
HQ: Atlanta, USA



50,000+
Employees

BUSINESS CHALLENGE

The Client's Project team is maintaining 29 repositories with over 1 million lines of code. The team is maturing in AI capabilities. Their developers currently spend 40%-60% of time searching for information in the large codebase. Under urgent pressure to boost development efficiency with AI, the Project team needs an automated AI-powered solution to provide context and reduce code search time.

SOLUTION

FPT has integrated **Metis** into the client's KTLO team, targeting a significant reduction in code search time within a large codebase, including 1M lines of code, 29 repositories:

- **Code Search Feature:** Deployed across 9 Metis.CodeIQ Knowledge Base Platform services (Knowledge Graph, Context Retrieval, Search Vector Engine, etc.).
- **Indexed Repositories:** Indexed 29 projects, 2000+ code files, 1M LOG, generating 82K nodes in the knowledge graph.
- **Issue Resolution:** Successfully resolved 50% issues, demonstrating immediate impact.

RESULTS & IMPACT

40%

REDUCED SEARCH EFFORT
Targeting a 40% reduction in code search effort.

30%

CYCLE TIME REDUCTION
Expected decrease in development team cycle time.

TECHNOLOGY HIGHLIGHTS

- **Cloud & On-Premise:** Runs seamlessly in both cloud and on-premise environments.
- **Microservice Design:** Scalable, robust architecture built on microservices.
- **Flexible Integration:** Works standalone or integrates with chatbots and tools.
- **Large-scale Handling:** Supports large code repositories and multi-tenant isolation.
- **High Concurrency:** Designed for 1,000+ concurrent users.
- **Dual Operation Modes:**
 - Full Flow Mode: Directly processes user prompts and tasks.
 - Integrated Mode: Acts as a knowledge base for integrated tools.

Success Story: Powering Agentic AMS at Scale with Metis

Japan-based Office Technology Manufacturer

A leading multinational manufacturer of multi-function printers and imaging solutions, maintaining a large-scale embedded software portfolio that powers its global product line across enterprise and consumer markets.



\$7B+
Revenue



Global
HQ: Tokyo, Japan



34,000+
Employees

BUSINESS CHALLENGE

The Client's printer software had grown so large and interconnected that no single person could fully understand the system - **a codebase exceeding 1 billion lines of embedded C/C++, with individual files running over 10,000 lines each**. Team size grew, but productivity didn't, and quality still varied release to release. Existing AI tools sped up individual work, but the most damaging defects came from cross-team intersections, putting the engagement at risk. The Client needs an AI-powered solution to unify system knowledge and operate reliably at hyperscale.

SOLUTION

FPT deployed Metis as an **agentic AMS platform operating directly across the client's entire codebase**, an AI agent layer acting as a digital engineering teammate throughout the development lifecycle:

- **Unified Knowledge Base:** Consolidating requirements, wikis, design docs, dependency graphs, and bug-fix history into a single, continuously updated living knowledge graph.
- **Problems Traced in Seconds:** Engineers see exactly where an issue starts and what a change will affect, before it becomes a costly failure.
- **Multi-Agent Software Toolchain:** Metis, CodeVista, TestVista, and AgentVista, working as a coordinated team of specialized agents to handle the full task chain: code search, bug fixing, testing, and review.
- **Valued Proven, then Scaled:** Expanded from one team with proved ROI to successful adoption across all engineering teams.

RESULTS & IMPACT

2x

PRODUCTIVITY INCREASE

Achieved a double delivery capacity without increasing headcount (Traditional methods had never achieved at this scale)

95%

FAITHFULNESS SCORE

Independently evaluated AI knowledge extraction accuracy, confirming the reliability of underlying knowledge base

TECHNOLOGY HIGHLIGHTS

- **True Multi-Agent Architecture:** Specialized AI agents (code search, bug-fix, test, review) coordinated through a unified Knowledge Base, not a single standalone chatbot
- **Model-Agnostic Inference Layer:** Runs flexibly across OpenAI, Gemini, Claude, Grok, LLaMA, and other leading models
- **Enterprise Toolchain Integration:** Connects directly via MCP with GitHub, Jira, Trello, Notion, and Confluence, embedded straight into existing workflows
- **Purpose-Built for Embedded Complexity:** Uniquely engineered to handle billion-line embedded C/C++ codebases across multi-team, cross-dependency environments

Success Story: Transforming AMS Delivery with Metis

Korea-based Global IT Services & Digital Solutions Provider

A leading Korean IT services group delivering application management services, from web applications, batch processing, BI dashboards, and database operations, for enterprise clients across finance, manufacturing, and public sector domains.



\$4B+
Revenue



Global
HQ: Seoul, South Korea



7,000+
Employees

BUSINESS CHALLENGE

The Client's AMS delivery spans web, batch, BI, and database operations - a workload still overwhelmingly manual. Existing AI tools, including Copilot, offered no meaningful support for most L1-L3 tasks: routine operations like data export, DB updates, and test execution stayed entirely manual, while even AI-assisted work like incident troubleshooting never rose above "support" level.

Delivery pressure kept mounting with no unified AI layer across the AMS lifecycle, leaving productivity gains capped and operational risk concentrated in manual, repetitive work.

SOLUTION

FPT developed a **Unified AI Managed Service** that embeds AI into the Client's daily AMS operations, powered by Metis.CodeIQ, Metis.TextIQ, and AgentBox on a shared enterprise knowledge base:

- **Unified Enterprise Knowledge:** Consolidates documents, tickets, source code, and collaboration data into a single, governed knowledge base that AI can securely access and reason over.
- **AI-powered Engineering Assistance:** Provides contextual Q&A, requirement generation, RCA, translation, and code/test support using verified and cited enterprise knowledge.
- **End-to-end Workflow Automation:** Integrates AI directly into AMS processes from ticketing, approvals, escalations, and reporting.
- **Continuous Learning and Optimization** through human feedback and learning loops.

Up to **20 AI agents** are deployed across 11 AMS task categories, enabling solutions such as Engineer Copilot, Self-Service Knowledge Portal, Automated Helpdesk Response, and Executive Dashboards, transforming AMS into an AI-assisted delivery model.

RESULTS & IMPACT

90%

Automation Rate

5x

Faster MTTR
(Mean Time to Resolution)

70%

Ticket Deflection

80%

Faster Assessment

TECHNOLOGY HIGHLIGHTS

- **Three-Engine AI Core:** Metis.CodeIQ (code intelligence), Metis.TextIQ (document/knowledge intelligence), AgentBox (workflow agent orchestration) working across a shared enterprise knowledge base
- **Flexible Deployment:** Offered as On-prem/Private LLM, Enterprise LLM (Customer-hosted), or Enterprise LLM (FPT-hosted), letting the client match security and speed-to-pilot needs to the right model
- **Model-Agnostic Inference:** Open-source models (Qwen3, Llama-3.3) for on-prem isolation, or enterprise cloud models (Azure OpenAI, Gemini, Claude) for faster time-to-value
- **Built-In Governance, Not Bolted On:** RBAC/IAM access control, DLP/PII masking, full audit trail, source-cited answers with hallucination prevention, and human approval gates
- **Compliance-Ready:** Mapped against SOC 2, GDPR, and HIPAA across the full stack



THANK YOU!

CONTACT US

FPT Tower, 10 Pham Van Bach Street, Cau Giay District, Hanoi City, Vietnam

Tel.: +84 (24) 3 768 9048

Fax: +84 (24) 3 768 9049