


CONFIDENTIAL CLIENT CASE STUDY

CMMC Compliance AI Backend

Agentic RAG, NIST-grounded reasoning and structured compliance-document generation for CMMC workflows.

A US-based compliance technology company engaged Byond Boundrys to build the AI backend for a CMMC documentation platform, including agentic assistance, control-aware NIST retrieval, gap assessment, document verification and SSP, POA&M and SRM generation.

Project classification: Byond Boundrys delivered the AI backend and APIs for a US-based compliance technology company; the client owned the product, frontend, customer experience and go-to-market.



INDUSTRY Cybersecurity & Compliance	CLIENT TYPE US-based compliance technology company	PROJECT STAGE Production implementation
EVIDENCE Pilot-observed	DELIVERY SCOPE Agentic AI, NIST RAG and audit-ready document generation	

Control-aware compliance assistance The backend maps requests to relevant CMMC and NIST controls before retrieving evidence and generating a response.	Grounded and cited responses Generated guidance retains standards references so compliance professionals can review the basis of each answer.	Structured gap and document evaluation Specialised services assess gaps, policies, procedures and evidence against control-specific requirements.
--	---	---

01 / BUSINESS CONTEXT

Who needed the solution

A US-based compliance technology company building a CMMC-exclusive documentation platform engaged Byond Boundrys to architect and build the AI backend that powers their workflow. The platform itself serves defense contractors, Managed Service Providers (MSPs/MSSPs), and CMMC consultants who need to prepare audit-ready documentation aligned to CMMC and NIST 800-171. Our scope was strictly the AI/ML layer—an agentic AI compliance chatbot, the NIST-focused reasoning engine, the agentic RAG system, gap assessment intelligence, document verification AI, and auto-generation engines for System Security Plans (SSP), Plans of Action and Milestones (POA&M), and Shared Responsibility Matrix (SRM) artifacts. The client owned the product, frontend application, customer experience, and go-to-market; we delivered the AI services that make the platform work.

02 / CHALLENGE

What needed to change

The backend had to support formal CMMC assessment workflows with control-level precision, auditability and privacy boundaries.

Generic LLMs could hallucinate on specialised control language

Unsupported interpretations could undermine assessor confidence and remediation decisions.

Document-level RAG lacked control-level precision

The workflow needed retrieval aligned to specific CMMC and NIST 800-171 objectives, not broad document summaries.

Gap assessment spans 110+ Level 2 objectives

The system required structured reasoning, prioritisation and remediation guidance across many related controls.

Policies and evidence needed objective-level verification

Uploaded documents had to be scored against specific requirements rather than merely summarised.

SSP, POA&M and SRM outputs had to be assessor-ready

Generated artefacts required consistent structure, citations and traceability to validated evidence.

Sensitive CUI required a privacy-first architecture

The AI layer was designed without client-network access and without retaining sensitive CUI beyond the request lifecycle.

03 / WORKFLOW TRANSFORMATION

Before and after

Previous workflow

- Generic LLM responses were difficult to trust for specialised control language
- Document-level search missed objective-specific evidence
- Gap assessments required manual reasoning across 110+ objectives
- Policies and evidence were reviewed manually against control requirements
- SSP, POA&M and SRM artefacts required repetitive manual preparation
- Sensitive information could not be exposed through uncontrolled integrations

Structured workflow

- Classify request intent and map it to relevant controls
 - Retrieve NIST and client evidence with control-aware RAG
 - Evaluate posture, gaps and remediation needs through specialised agents
 - Verify uploaded documents against objective-level requirements
 - Generate citation-backed SSP, POA&M and SRM artefacts
 - Return structured outputs through privacy-first backend APIs for human review
-

04 / SOLUTION

How we approached it

A LangGraph-based multi-agent service layer decomposed requests into intent detection, control mapping, grounded retrieval, compliance evaluation and structured output generation.

01

Intent detection

Classify each request as gap assessment, document verification, artefact generation or remediation guidance.

02

Control mapping

Link the request and supplied evidence to relevant CMMC and NIST 800-171 controls and objectives.

03

Control-aware retrieval

Chunk standards and uploaded documentation at control boundaries, embed the content and retrieve objective-level evidence.

04

Gap assessment engine

Evaluate posture across CMMC Level 2 objectives and return prioritised, NIST-cited remediation guidance.

05

Document analyser

Verify policies, procedures and evidence against specific objectives and return actionable quality feedback.


06

Artefact generators

Produce structured SSP, POA&M and SRM outputs from validated evidence and control mappings.

05 / EXAMPLE WORKFLOW

From input to reviewable output



06 / DELIVERY SCOPE

What the delivery covered

Compliance workflow capabilities

- Real-time control-aware compliance chatbot
- Gap assessment across CMMC Level 2 objectives
- Policy, procedure and evidence verification
- NIST-cited remediation guidance
- SSP, POA&M and SRM generation

AI, retrieval and evaluation

- LangGraph multi-agent orchestration
- Control-aware chunking and retrieval
- Pinecone vector indexing and semantic retrieval
- Structured-output and citation enforcement
- Quality scoring, traceability and validation


Backend, cloud and privacy

- Node.js backend and API service layer
- Azure OpenAI reasoning and generation
- Microsoft Azure / Azure Government deployment foundation
- SOLID-based modular architecture for framework extensibility
- No client-network access and no CUI persistence beyond request scope

Scope boundary

Byond Boundrys delivered only the AI backend and APIs. The client retained ownership of the frontend, product experience and go-to-market. Quantitative impact figures are excluded from the public page.

How the system is organised



Modular service boundary

Clean APIs separate the AI backend from the client-owned frontend and product experience.

08 / BUSINESS VALUE

Value observed during validation

- | | | |
|---|--|----------------|
| 1 | Control-aware compliance assistance
The backend maps requests to relevant CMMC and NIST controls before retrieving evidence and generating a response. | Pilot-observed |
| 2 | Grounded and cited responses
Generated guidance retains standards references so compliance professionals can review the basis of each answer. | Pilot-observed |
| 3 | Structured gap and document evaluation
Specialised services assess gaps, policies, procedures and evidence against control-specific requirements. | Pilot-observed |
| 4 | SSP, POA&M and SRM generation
The API layer produces structured compliance artefacts from validated mappings, evidence and remediation inputs. | Pilot-observed |
| 5 | Privacy-first request processing
The AI layer is separated from client networks and does not persist CUI beyond the request lifecycle described in the source. | Pilot-observed |

09 / TECHNOLOGY

Technology and component roles

NO	Node.js Backend runtime and API layer exposing AI services to the client product	AZ	Azure OpenAI Reasoning, evaluation and structured compliance-document generation	LA	LangGraph Stateful orchestration of specialised compliance agents
LA	LangChain RAG and LLM application orchestration	PI	Pinecone Control-level semantic vector retrieval	TE	text-embedding-ada-002 Embeddings for standards, controls and uploaded documentation - verify current production model

MI

Microsoft Azure

Secure production deployment foundation, including Azure Government where applicable

CO

Control-aware PDF pipeline

Parsing, boundary-aware chunking and normalisation of standards and evidence

10 / RELATED WORK

Explore related case studies

INSURANCE & COMPLIANCE

InsureAI

RAG-based policy review, real-time knowledge retrieval and automated policy-gap analysis.

HEALTHCARE & CLINICAL RESEARCH

Clinical Data Harmonization

Hybrid RAG and fine-tuned small language models for explainable source-to-standard clinical data mapping.

AI SAFETY & CONTENT INTEGRITY

AI Text Detector SLM

A fine-tuned small language model for explainable AI-text detection and adversarial robustness.

Working on a similar challenge?

Discuss the workflow, data, integrations and evidence required for a practical first release.