PUBLIC TECHNICAL REVIEW PACKAGE

LLM Acceleration and Web AI Development

A detailed public review document for custom LLM acceleration, router, streaming API, web service path, and developed AI system layer.

Core in one sentence The AI work is acceleration and Web AI system development, not simple model operation.	**Most important point** The key public point is that the LLM layer is developed infrastructure, not prompt-only model access.

LLM infra	Service path	Web AI	Live
beyond prompt-only access	selected workload evidence	API / router / UI	Web AI layer readiness

What this PDF is meant to prove

- The AI story is framed as AI system development, not simple model operation.
- LLM acceleration includes cache, draft path, router, streaming API, and service behavior.
- Web AI layer connects local model performance work to a deployed public system.
- Performance evidence remains tied to model, workload, baseline, hardware, and prompt shape.

Acceleration layer Develops faster local LLM behavior through cache, draft path, routing, and service logic.	**Service layer** Develops HTTP behavior, streaming responses, route choice, and request handling.
Web AI layer Develops API, UI, deployment checks, and public review packaging.	**Boundary** Performance results are workload-labeled, not universal speed guarantees.

Inside this document

Fast presentation line

Formysec develops AI security infrastructure by connecting security modules, custom execution foundations, LLM acceleration, and Web AI system development into one reviewable portfolio.

1. Development Role

This document explains why the AI part is development work. The point is not that a model exists. The point is that model behavior is developed into an acceleration layer, service path, and Web AI system.

- Cache and reuse reduce repeated local inference cost.
- Draft path and speculative behavior improve selected workload flow.
- Router and request classification turn model calls into controlled service behavior.
- Streaming API and UI make the AI system accessible through the web layer.
- Deployment checks and PDFs make the public package reviewable.

Before A local model can be slow, manual, and shell-bound.	**After** The model path becomes a developed service layer with web integration.
Advantage The portfolio can show AI infrastructure development rather than AI content operation.	**Limit** Performance depends on selected conditions and workload shape.

Presenter Notes

Plain-English angle This page explains 1. development role in simple terms: what was developed, why it matters, and how it should be reviewed.	**Technical angle** Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.
Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	1. Development Role describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

2. LLM and Web AI Component Register

This register gives a detailed but public-safe explanation of the AI development layer.

Component	Existing role addressed	Formysec difference	Evidence and boundary
Prefix cache / KV reuse	Repeated inference cost.	Develops reuse behavior for repeated or related prompts.	Performance depends on prompt shape and model.
Draft path	Single-path local model generation.	Develops a faster candidate path for selected workloads.	Workload-labeled support evidence only.
Router	Manual model route choice.	Develops service logic for route selection and response path control.	Public text stays architecture-level.
Streaming API	Batch-only or shell-only output.	Develops web-service behavior for live responses.	Performance depends on model and environment.
Web AI UI	Disconnected demo surface.	Develops a public system layer around the AI service path.	Site proves packaging, not external certification.
Deployment checks	Unverified local page.	Live checks confirm links, PDFs, copy lock, and deployed state.	Operational check only.

Presenter Notes

Plain-English angle This page explains 2. llm and web ai component register in simple terms: what was developed, why it matters, and how it should be reviewed.	**Technical angle** Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.
Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	2. LLM and Web AI Component Register describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

3. Existing Alternative Comparison

This is the clean way to explain why the AI layer is more than a normal chat interface.

Existing role	Typical limitation	Developed difference	Public advantage
Prompt-only AI page	Looks like simple model access.	Adds acceleration, router, API, UI, and deployment evidence.	Shows AI system development.
Local shell inference	Hard to present as a product or service.	Adds web-connected service behavior.	Readable for reviewers and visitors.
Single-path inference	No route choice or workload adaptation.	Adds routing and cache-aware service direction.	Stronger infrastructure story.
Unlabeled performance statement	Can be easy to overstate.	Labels performance evidence by selected baseline and workload.	More credible performance presentation.

Presenter Notes

Plain-English angle This page explains 3. existing alternative comparison in simple terms: what was developed, why it matters, and how it should be reviewed.	**Technical angle** Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.
Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	3. Existing Alternative Comparison describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

4. Performance Story

The performance story is strongest when it is stated through external comparison with boundaries.

Comparison anchor LLM infrastructure is compared against prompt-only and single-path model access.	**Service evidence** Cache, routing, and streaming behavior provide the public service-layer difference.
Main advantage Turns local model execution into a faster developed service path.	**Main boundary** Results depend on model, hardware, prompt length, baseline, and workload shape.

- Lead with the development work: acceleration layer, router, streaming API, Web AI service path.
- Then keep detailed speed evidence as scoped support material, not headline wording.
- State that broader model-by-model tables would further strengthen the statement.
- Do not frame the work as ordinary model access.

Presenter Notes

Plain-English angle This page explains 4. performance story in simple terms: what was developed, why it matters, and how it should be reviewed.	**Technical angle** Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.
Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	4. Performance Story describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

5. Readiness and Evidence Plan

The AI layer is the most public-facing part of the stack, so its evidence must be clear.

Area	Current public state	Why it matters	Next strengthening step
Acceleration story	Clear in site and PDF.	Shows built AI performance work.	Add more workload tables.
Service path	Router, API, streaming, and UI language present.	Shows system development.	Add a public demo flow when ready.
Web deployment	Live site and PDFs deployed.	Shows reviewable public layer.	Add versioned release notes.
Boundary control	Performance records are condition-labeled.	Prevents overstatement.	Add more external comparison descriptions.

Presenter Notes

Plain-English angle This page explains 5. readiness and evidence plan in simple terms: what was developed, why it matters, and how it should be reviewed.	**Technical angle** Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.
Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	5. Readiness and Evidence Plan describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

Plain-English Presenter Guide

This page gives the short explanation before a deeper technical review. It is designed for non-specialists and for fast presentation flow.

Review step	What it explains	What evidence follows	Do not mix with
Thesis	Formysec develops AI security infrastructure from modules to Web AI.	Master portfolio and topology.	Internal benchmark details.
Enterprise lens	A small research lab produced a broad, reviewable security and AI infrastructure package.	Enterprise review table and evidence ledger.	Universal deployment promises.
Architecture	Security base, RMEP modules, foundation, LLM acceleration, and Web AI delivery.	Stack map and layer descriptions.	Baseline comparison.
External comparison	Comparison evidence for module role, validation workflow, consistency, and LLM infrastructure.	Security-module comparison, validator-workflow comparison, LLM-infrastructure comparison, and implementation-check summaries.	General overview language.
Track details	Each field gets its own role, value, evidence style, and boundary.	Field PDFs and component pages.	One overloaded table.
Comparison	What the work may replace, complement, or improve.	Baseline comparison table.	Private construction details.
Evidence and scope	What is public, what was checked, and what remains private.	Evidence PDF, public scans, live checks.	Sensitive internal design notes.

Reviewer Q&A;

Reviewer question	Short answer	Evidence to point at	Boundary
What is the portfolio?	A technical review package for AI security infrastructure built from modules, foundation work, LLM acceleration, and Web AI.	Portfolio structure page and master PDF.	Private internals excluded.
Why should a company care?	It shows a lean lab with broad execution: security modules, AI acceleration, Web AI delivery, protocol builds, and evidence control.	Enterprise review lens, evidence ledger, and live site.	Not a final product approval statement.
Where is the evidence?	It is organized in the external comparison register so evidence does not get mixed with overview text.	External comparison table and evidence PDF.	Records are condition-labeled.
How is RMEP framed?	RMEP is structure-first module development; execution paths carry the structure rather than replace it.	RMEP PDF and result register.	Not a final-standard statement.
Is this ordinary model access?	No. The public position is development of acceleration code, routing, API behavior, Web AI layer, and evidence packaging.	LLM acceleration PDF and Web AI section.	Performance varies by workload and baseline.

Presenter Notes

Plain-English angle
This page explains plain-english presenter guide in simple terms: what was developed, why it matters, and how it should be reviewed.

Technical angle
Focus on component role, existing-role comparison, advantage, limit, evidence type, and public boundary.

Strong answer Start with the developed part, compare it with a familiar external role, then state the boundary so the statement stays credible.	**Review caution** Keep protected construction details outside the public explanation and point back to the published evidence package.

Review Questions

Reviewer question	Public answer
What was developed?	Plain-English Presenter Guide describes a developed part of the Formysec stack and connects it to the larger AI security infrastructure story.
How is it different?	The page compares the developed role against an existing technical role, then explains the advantage and the current evidence.
What is the boundary?	The public answer stays with component role, external comparison, readiness state, and public evidence while protected construction details remain excluded.

Public Boundary and Presentation Notes

This final page keeps the package strong and objective during presentation.

- The document presents component roles, external comparison labels, evidence scope, and readiness state.
- Private technical internals, protected implementation details, and sensitive construction notes stay outside this public package.
- Performance evidence remains tied to workload, environment, review method, and external comparison baseline.
- The language is intentionally objective: strong enough for a technical presentation, careful enough for review.

Recommended speaking frame

- Say that Formysec develops AI security infrastructure across security modules, execution foundations, LLM acceleration, and Web AI system layers.
- Say records are review records with stated boundaries, not universal production guarantees.
- Say that the portfolio is designed to be understandable without exposing protected internal construction.
- When asked for deeper internals, redirect to public component roles, comparison evidence, and review roadmap.