Open Medical Inference

Open infrastructure for federated AI inference in healthcare

A project of the Medical Informatics Initiative (MII)

Interoperable AI services for German university medicine

Interoperable KI-Services für die deutsche Universitätsmedizin

An MII project building open, standards-based infrastructure to bring AI inference services into clinical and research workflows across university medicine.

Dr. Obioma Pelka · Kaushik Manjunatha · Universitätsmedizin Essen (IKIM)

16 partner sites · ~€8.5M · BMFTR / MII

9 work packages, all functional

2 standards used: HL7 FHIR + DICOM

THE PROJECT

What is OMI?

Open Medical Inference, a Medical Informatics Initiative methods platform.

THE GOAL

Make AI inference services discoverable and usable across university medicine, through one open, standards-based protocol and reference infrastructure.

A site without its own AI capacity sends pseudonymised data to a service hosted elsewhere, and gets the results back.

Funding & scale
BMFTR / MII · ~€8.5M · 16 sites

A consortium
16 sites · four MII consortia

Open
Open source · FHIR IG on Simplifier

MOTIVATION

The problem we set out to solve

AI models exist. Getting them to patients across sites does not work.

01

Two worlds

Data standards (FHIR, DICOM) and AI backends assume different things. Every integration is built from scratch.

02

The last mile

Strong models rarely reach care at other sites: no common way to discover and invoke them across the MII.

03

Governance

The EU AI Act and the MDR require traceability. Ad-hoc adapters leave no audit trail.

OUR HYPOTHESIS: the bottleneck is interoperability and governance, not the models.

OUR APPROACH

How we planned to solve it

An open specification, a reference implementation, evaluation.

Specify

An open FHIR protocol, an imaging module for the MII core dataset, and a governance framework.

Implement

Reference components anyone can deploy: client, gateway, DICOMweb adapter, service registry.

Evaluate

Roll out across sites, connect to RACOON / NUM nodes, and test it in practice.

PRINCIPLE: standards over glue, governance as infrastructure, low entry barriers. Distributed inference, not federated learning.

ARCHITECTURE

The full OMI architecture

From a requesting site to AI services hosted elsewhere.

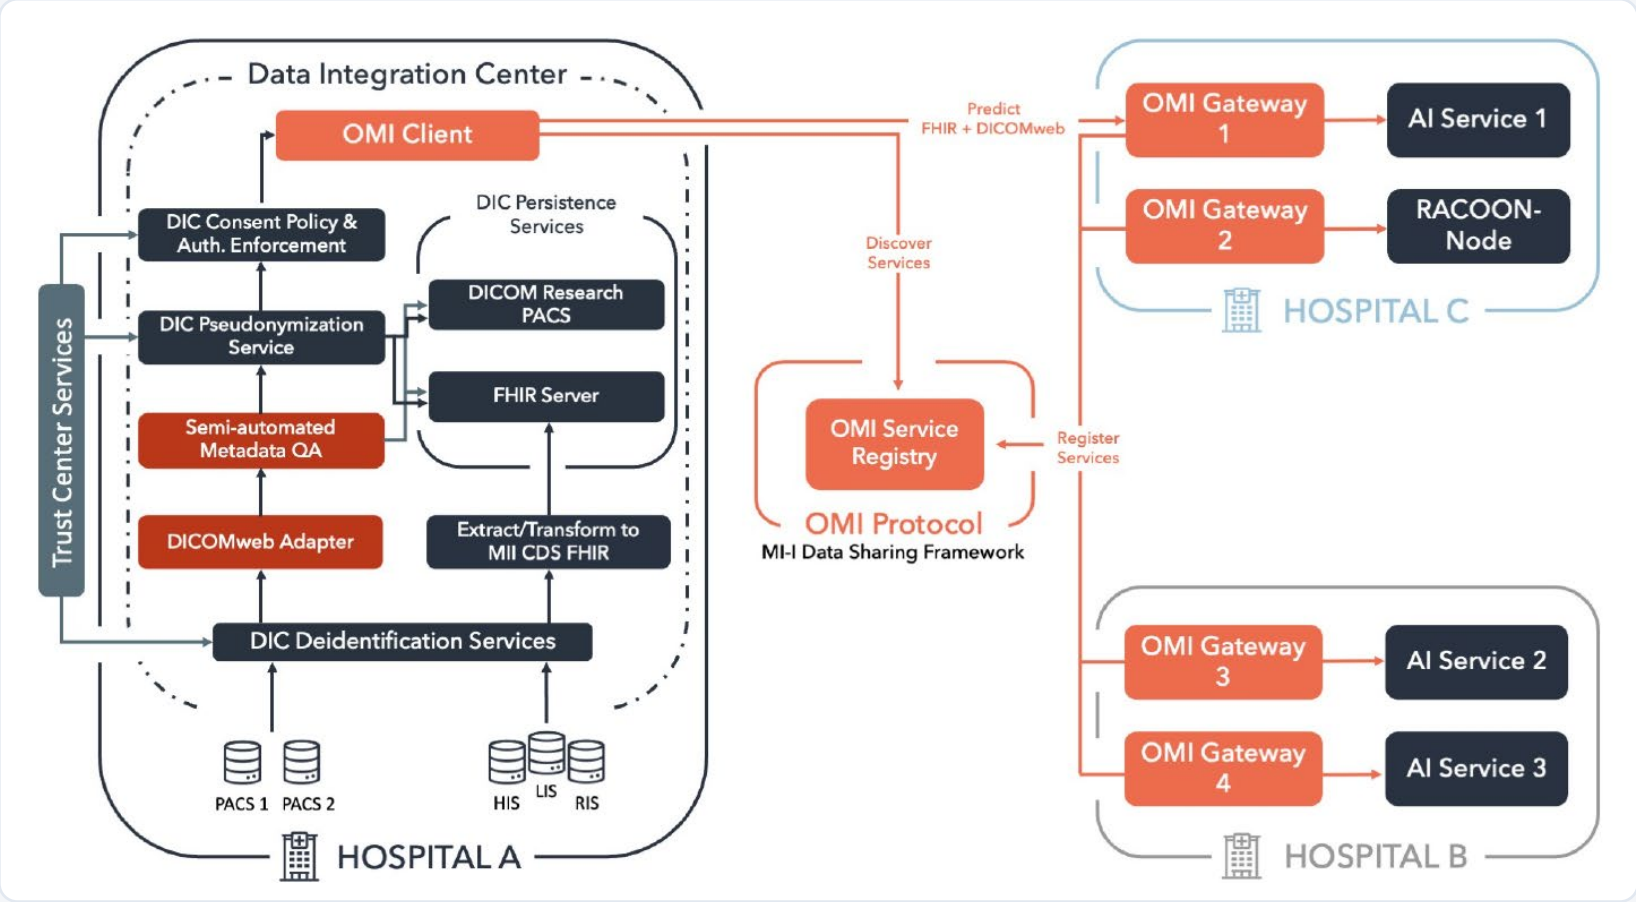

Requesting site

Data prepared and de-identified at the origin.

Protocol + registry

Discover and register services over the OMI protocol.

Inference sites

Gateways host AI services; results return.

WORK PACKAGES

Nine work packages

What each part of the consortium builds.

Specification

WP1 Protocol — open FHIR contract

WP2 Imaging KDS module

WP7 Governance & ethics

Implementation

WP3 DICOMweb adapter to PACS

WP4 Gateway: validate & route

WP5 Client: request inference

WP6 Registry: find services

Evaluation

WP8 Roll-out & evaluation

WP9 Project management

KDS: The imaging extension to the MII core dataset (WP2) is specified and released. ✓

THE GATEWAY (WP4)

Validation and audit on the inference path

Every stage validated; an audit trail by design.

Figure 2: OMI Gateway Processing Pipeline

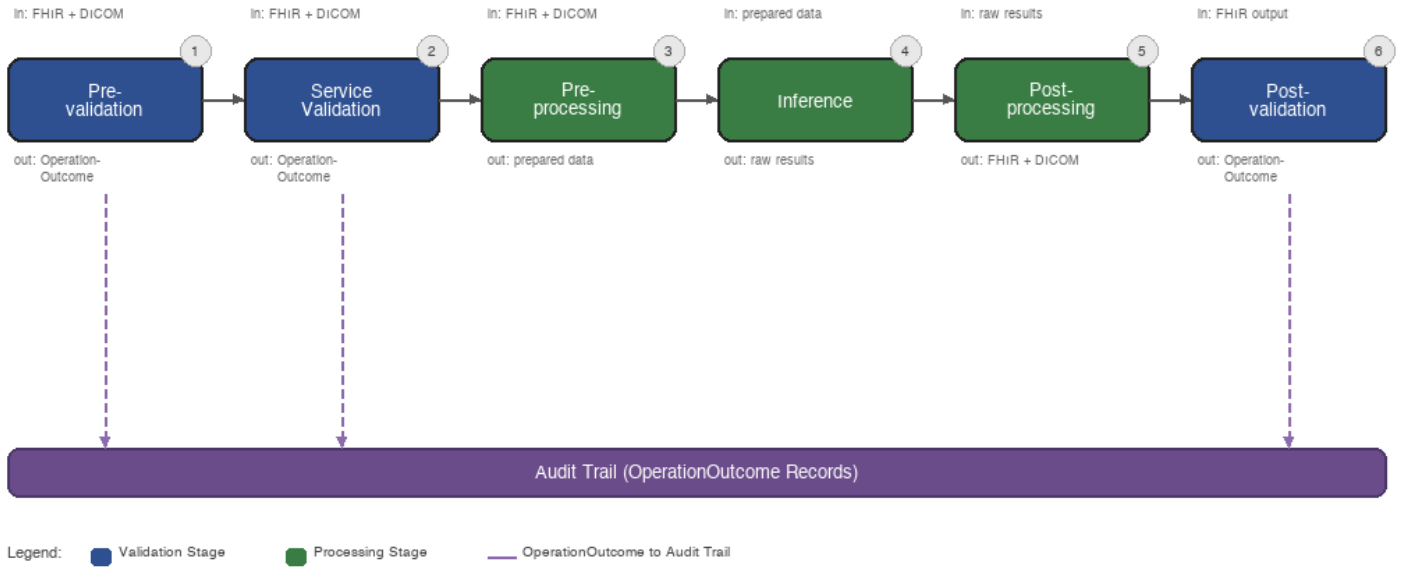

WHY IT MATTERS

Inputs and outputs checked against FHIR and DICOM at every stage.

Each stage writes an OperationOutcome: a machine-readable audit trail.

Plugins: add a service, not a rebuild.

THE IMAGING PATH (WP3)

Bridging AI services to hospital PACS

The DICOMweb adapter: the same mediation, for imaging.

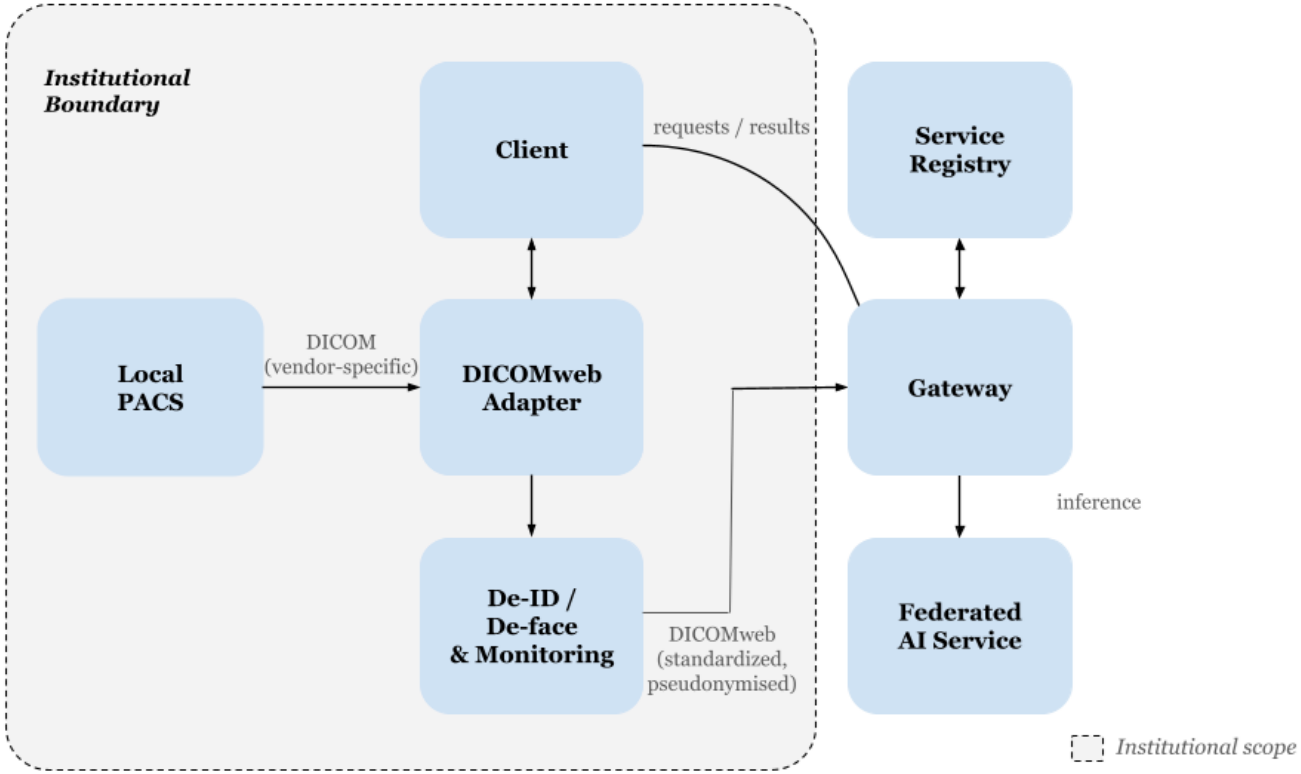

WHAT WE LEARNED

Existing tools (Orthanc) cannot proxy DICOMweb to remote PACS, so a dedicated layer is needed.

IN PRODUCTION

6,061 requests · 100% success

Two-week clinical deployment.

EXAMPLE AI SERVICE

Example: requesting BOA inference

BOA, the Body & Organ Analyzer, as an OMI service.

1. **Discover —** find BOA in the registry; read its FHIR input / output.
2. **Submit —** one REST call: FHIR + the CT series; get a job ID.
3. **Pseudonymise & validate —** de-identified at origin; inputs checked; audit written.
4. **Infer —** gateway routes to the BOA plugin; segmentation runs.
5. **Retrieve —** results as FHIR + DICOM segmentation, with audit trail.

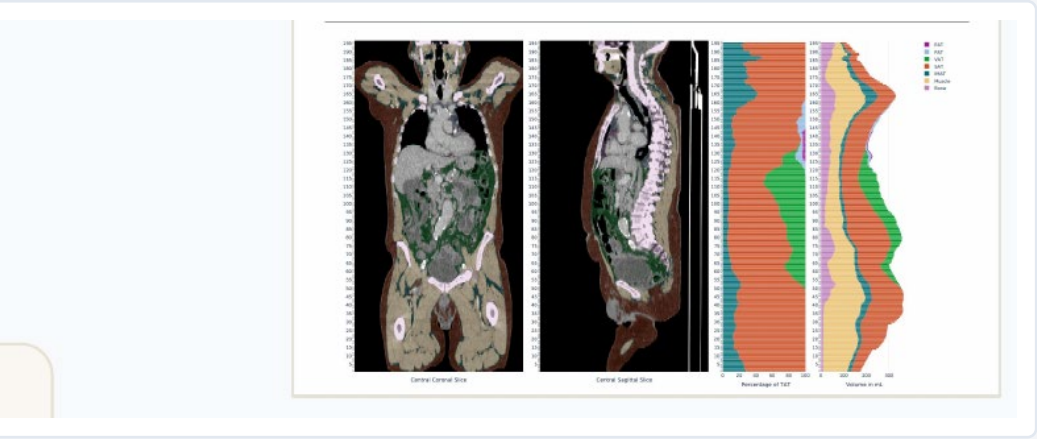

BOA report: body-composition & organ segmentation on CT

ONE CALL FROM THE RESEARCHER. OMI handles discovery, pseudonymisation, validation, transcoding, and the audit trail.

PERFORMANCE

Low, predictable gateway overhead

100 sequential requests; single host.

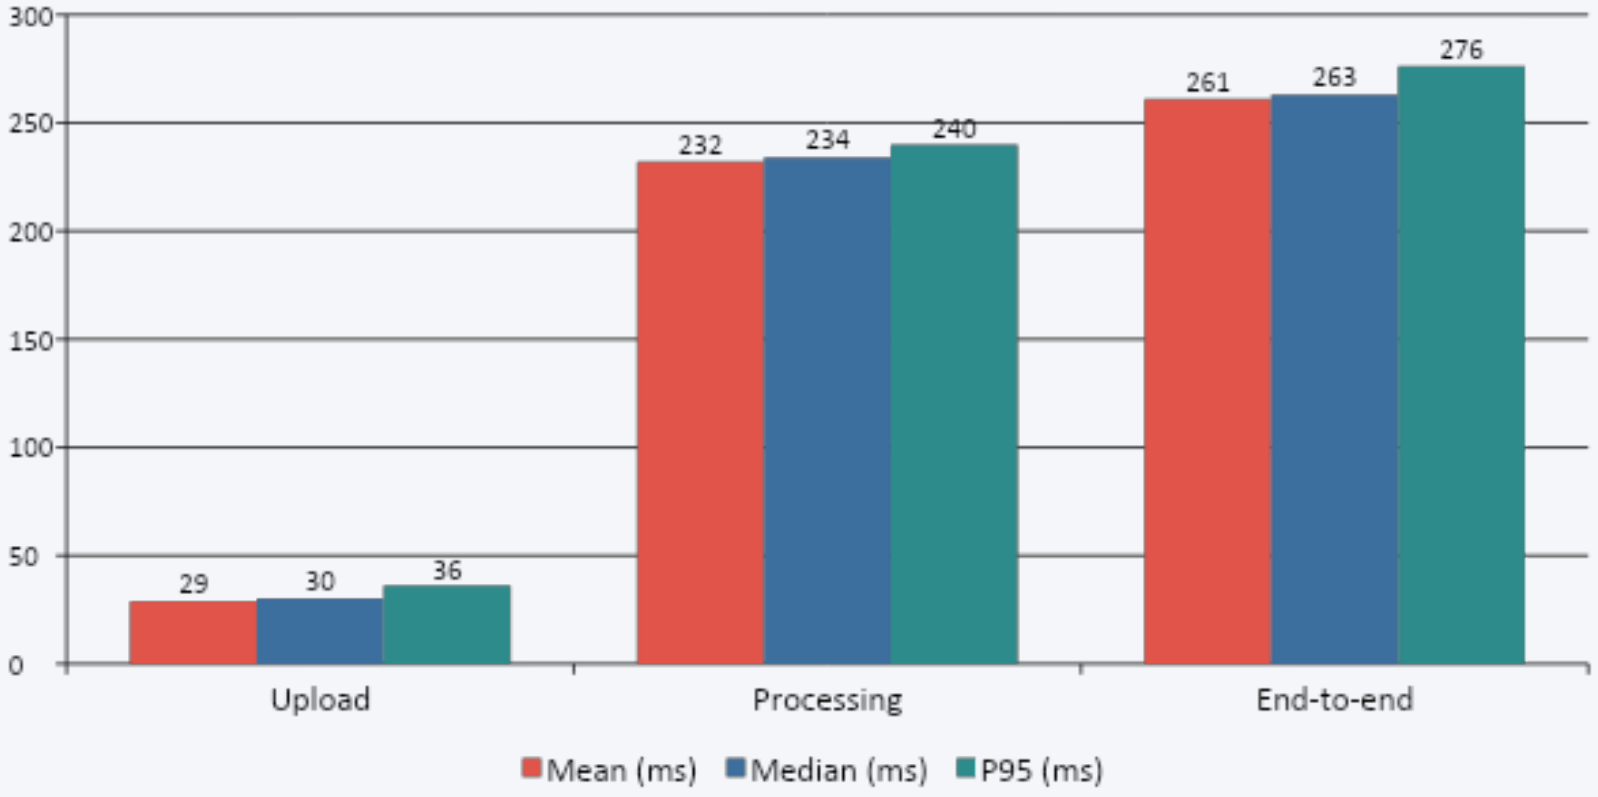

DISPERSION (SD)

Upload ± 4 ms

Processing ± 21 ms

End-to-end ± 22 ms

Small and predictable relative to model inference.

POSITIONING

How the OMI gateway compares

	OMI	Kaapana	JIP	MONAI
OMI protocol compliance	**Yes**	No	No	No
FHIR validation pipeline	**Yes**	No	No	No
OperationOutcome audit trail	**Yes**	No	No	No
Plugin-based extensibility	**Yes**	Yes	Yes	Yes
DICOMweb support	**Yes**	Yes	No	Yes
Open source	**Yes**	Yes	Yes	Yes

DIFFERENTIATORS: protocol compliance, FHIR validation, the audit trail.

STATUS

Where we are now

Halfway; all components functional.

Done

- Protocol + FHIR IG published
- KDS imaging module released
- Gateway + adapter, open source
- In production at one site
- Peer-reviewed publications

In progress

- Evaluation (WP8), from 2026
- Multi-site roll-out
- More AI services onboarded
- Defacing + in-adapter audit

Limitations

- Single-site evaluation so far
- Latency on one host only
- High concurrency untested
- Some governance still planned

RESULTS

Published results

Documented and peer-reviewed across the consortium.

RöFo 2026	**Open Medical Inference: protocols, data exchange & AI integration**	*Pelka, Sigle, Werner … Nensa*
GMDS 2026	**The OMI Gateway**	*Pelka, Eil, Manjunatha, Singh, …*
GMDS 2026	**DICOMweb adapter for distributed AI in imaging**	*Pelka, Kamzol, Manjunatha, Singh, Omeirat, Nensa*
Whitepaper 2024	**Bridging (AI-)services & clinical use: the OMI-protocol**	*Sigle, Werner, Schweizer, …*
Conf. 2026	**AI governance & ethics framework: expert evaluation**	*Apfelbacher, Christoforaki, … Ganslandt*
Conf. 2024	**A governance framework for AI in a university hospital**	*Apfelbacher, Kocman, Prokosch, Christoph*
Conf. 2024	**OMI meets RACOON: MI + radiology infrastructure**	*Prokosch, Nensa, Penzkofer*

OUTLOOK

Sustainability and use

Carrying OMI beyond the project, and how it can serve the MII.

What is planned

- Sustainability beyond funding (Verstetigung)
- Governance & AI-monitoring as a service
- Maintained open-source release + FHIR IG
- Federation across sites; MII / NUM

How it can be used

- Run a model at one site on data at another
- Validate a model across sites, one interface
- Bring research / vendor AI in, governed
- Standardised, auditable AI evaluation

TAKE-HOME

Three things to remember

01

Standards-based

Services discovered and invoked the same way across sites, over FHIR and DICOM.

02

Governable

Validation and an automatic audit trail: transparent and traceable by construction.

03

Real

Nine functional components, in production, peer-reviewed and open source.

Thank you.

To the whole OMI consortium, and everyone working toward the goal.

omi.ikim.nrw