



Eco-system for disease specific clinical workflow
and data integration

*Building the Foundations for Standardized
Clinical AI: The SYMPHONY Contributions*



Project number:	ITEA 21026
Written by:	Yusuf Sayita
Date:	2026-03-01

ITEA Roadmap challenge:
Smart Health

Executive Summary

Healthcare today faces a fundamental Interoperability Crisis, often visible as a lack of meaningful “connectivity” between systems. Hospitals generate huge amounts of data every day – imaging, lab results, genomics, and patient-reported outcomes – but this data is locked in silos. Systems technically connect, but they do not truly interoperate. Formats, semantics, and workflows differ so much that data cannot move or be reused in a useful way. This makes AI adoption very difficult because AI requires structured, standardized, and machine-interpretable data.

Symphony’s solution is the Open Data Backbone. Instead of creating new standards, Symphony brings together existing global standards like HL7 FHIR, OpenEHR, DICOM, and IHE profiles into one consistent architecture. The backbone is built on the Six Interface Model, which separates life cycle layers and makes data flow clearer and traceable across the hospital IT ecosystem. This supports compliance, traceability, and human-in-the-loop control, preparing healthcare providers for regulations such as GDPR and the upcoming EU AI Act.

Key Result: Symphony validated this standards-based approach in four real clinical use cases:

- **Prostate Cancer** – AI-guided guideline automation using structured and unified data.
- **Abdominal Aortic Aneurysm** – Chained AI for biomechanical risk assessment using DICOM SR and Segmentation.
- **Atrial Fibrillation** – Patient engagement with FHIR-based mobile apps and chatbot.
- **Multiple Sclerosis** – Multi-layer workflow integrating imaging AI, mobile vision tests, and predictive models.

Symphony shows that **interoperability, not just connectivity**, is the foundation for safe, explainable, and scalable clinical AI. It is more than a technical implementation — it is a blueprint for future-ready hospitals.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	2
1 INTRODUCTION & BACKGROUND	5
1.1 The Interoperability Crisis in Healthcare	5
1.2 Why AI Needs Standardisation	5
1.3 Symphony’s Mandate: Evolution, Not Revolution.....	5
1.4 Architectural Principle: Life Cycle Decoupling	5
1.5 Compliance and Trust – Preparing for EU AI Act.....	6
2 THE SYMPHONY BLUEPRINT – THE OPEN DATA BACKBONE	7
2.1 Why Architecture Matters	7
2.2 The Open Data Backbone	7
2.3 Philosophy: Evolution, Not Revolution.....	8
3 THE STANDARDS ENGINE – TECHNICAL HOW-TO	9
3.1 The Hybrid Model: FHIR + OpenEHR.....	9
3.2 The AI Engine: DICOM SR and Segmentation	10
3.3 Workflow Orchestration: IHE Profiles	10
3.4 Feedback Loop – Closing the Gap	11
4 CLINICAL IMPACT – THE USE CASES.....	12
4.1 Use Case 1: Prostate Cancer (PC).....	12
4.2 Use Case 2: Abdominal Aortic Aneurysm (AA)	12
4.3 Use Case 3: Atrial Fibrillation (AF)	13
4.4 Use Case 4: Multiple Sclerosis (MS)	14
5 CHALLENGES AND LESSONS LEARNED	17
5.1 Vendor Readiness	17
5.2 Lack of UPS-RS Support in PACS	17
5.3 Mapping Complexity Between FHIR and OpenEHR	17
5.4 Standardisation Gaps.....	18

5.5 Data Privacy and Compliance..... 18

5.6 Limited Support for DICOM Segmentation Objects in PACS..... 18

5.7 Summary 19

6 GUIDANCE FOR FUTURE IMPLEMENTERS 20

6.1 Start Small and Grow 20

6.2 Demand Standards in Contracts 20

6.3 Validate PACS Capabilities Early 20

6.4 Prepare for Mapping Complexity..... 20

6.5 Expect Gaps in Standards 21

6.6 Make Compliance Part of Design 21

6.7 Build for User Experience 21

6.8 Plan for Marketplace Integration 21

6.9 Final Advice 22

7 CONCLUSION 22

1 Introduction & Background

1.1 The Interoperability Crisis in Healthcare

Healthcare systems in Europe produce huge amounts of data every day. This includes MRI and CT images, genomic sequences, and data from wearable devices. But the main problem is not the lack of data, it is the lack of **interoperable data** flow. Information is locked inside proprietary systems, often stored in unstructured formats or non-standard schemas. This creates isolated data silos, sometimes called a “**digital archipelago**”, where valuable patient information cannot be shared or reused effectively.

This became very visible during the **COVID-19 pandemic**, when hospitals struggled to exchange critical patient data even within the same country. The inability to share data quickly showed how weak the legacy IT infrastructures are. And because modern healthcare is moving from **volume-based care to value-based care**, which depends a lot on longitudinal patient information, the lack of interoperability keeps clinical workflows fragmented and inefficient.

1.2 Why AI Needs Standardisation

Artificial Intelligence (AI) is expected to transform healthcare by supporting diagnosis, predicting disease progression, and automating guideline-based decisions. But AI depends on **high-quality, structured, and interoperable data**. Without standardized data exchange, AI cannot be trained on diverse datasets and cannot be integrated into clinical workflows.

Symphony recognized that the biggest barrier for AI adoption is not the algorithm itself, but the infrastructure that connects data sources. This is why the project focused on building a **standardized interoperability backbone** rather than creating new proprietary formats.

1.3 Symphony’s Mandate: Evolution, Not Revolution

The **ITEA4 Symphony project** was designed to solve this interoperability challenge. Instead of inventing new standards, Symphony combines widely accepted standards:

- **HL7 FHIR** for data exchange
- **OpenEHR** for semantic persistence
- **DICOM** for imaging and AI results
- **IHE Profiles** for workflow orchestration

This approach is documented in **Deliverable D7.3/D7.5: Standardisation & Dissemination Plan**, which explains the strategy for adopting interoperability standards. The project’s philosophy is clear:

“Use what works, integrate it smartly, and make it scalable.”

Symphony is an international consortium with hospitals (Karolinska, Leiden UMC), technology partners (Philips), and SMEs. Its goal is to create an **Open Data Backbone** that supports disease-specific workflows and AI integration.

1.4 Architectural Principle: Life Cycle Decoupling

One of the key ideas in Symphony, described in **Deliverable D2.1: Reference Architecture for Open Ecosystem**, is **Life Cycle Decoupling**. In traditional hospital IT systems, data storage, workflow logic, and user interface are tightly coupled. This makes upgrades risky and expensive. Symphony separates these layers so they can evolve independently.

For example:

-
- A hospital can upgrade its AI storage backend without changing the clinician dashboard.
 - Applications interact through standardized interfaces, not direct database queries.

This principle is enforced through the **Six Interface Model** defined in D2.1, which organizes data flow between sources, stores, and applications using standards like FHIR, OpenEHR, and DICOM.

1.5 Compliance and Trust – Preparing for EU AI Act

In healthcare, trust is not only about technology, but also about legal and ethical compliance. Symphony was designed with GDPR in mind, but now we also look at the **EU AI Act**, which is becoming very important for high-risk AI systems in Europe.

The EU AI Act asks for **traceability**, **transparency**, and **human-in-the-loop control**.

Symphony architecture already supports these requirements by design:

- The **Six Interface Model** makes data flow clear and documented. Every transaction between applications, data sources, and stores can be traced.
- **IHE ATNA logging** provides audit trails for all data access and actions. This is not only good IT practice, but it is a strong foundation for accountability.
- Workflow orchestration with **UPS-RS** and validation steps ensures clinicians stay in control of AI decisions. AI results are never “black box”; they are validated before use.

This means Symphony is not just about interoperability. It is an infrastructure that helps hospitals prepare for EU AI Act compliance. By combining standards, logging, and human validation, Symphony makes AI safer and more transparent for clinical use.

2 The Symphony Blueprint – The Open Data Backbone

2.1 Why Architecture Matters

In healthcare IT, many systems are monolithic. This means data storage, workflow logic, and user interface are all tightly connected. If you change one part, everything can break. This is why hospitals often keep old systems for years, because upgrades are risky and expensive.

Symphony project wants to change this. The idea is **Life Cycle Decoupling**. This means each layer of the system – data sources, data stores, applications, and workflow orchestration – should be independent. If you upgrade your AI engine, you should not need to change your EHR or dashboard.

2.2 The Open Data Backbone

The backbone is not one big system. It is a set of **standardized interfaces** that connect different components. This makes the system flexible and scalable. Hospitals can add new AI tools or replace storage without breaking workflows.

Symphony defines **six interface categories** in **Deliverable D2.1: Reference Architecture for Open Ecosystem**. These interfaces are the “rules of the road” for data exchange:

1. **Interface I – Applications → Data Sources**
Apps can pull data from EHR or imaging systems using HL7 v2, HL7 FHIR, and DICOM Query/Retrieve.
2. **Interface II – Data Sources → Applications**
When new data is available (like a CT scan), it can be pushed to apps using DICOM Send or FHIR Subscription.
3. **Interface III – Workflow Management ↔ Applications**
Orchestration of tasks like “Start Analysis” using IHE AIW-I and REST APIs.
4. **Interface IV – Data Stores → Applications**
Query longitudinal data using OpenEHR and FHIR.
5. **Interface V – Applications → Data Stores**
Persist AI results using DICOM SR, FHIR Observations, and OpenEHR Compositions.
6. **Interface VI – Data Sources → Data Stores**
Archive raw data for long-term storage using DICOM Storage and HL7 MDM.

Handling Different DICOM Data Models

One practical question we got during the project was whether Symphony also controls the **different DICOM data models**, for example classic MR versus Enhanced MR, or CT versus Enhanced CT. Symphony does not force one specific DICOM object model. Instead, the idea is that the AI components and workflow manager must be able to handle these variations in a standards-based way. So, the orchestration layer does not try to “normalize” all DICOM flavors itself, but it works with whatever the modality or PACS provides, if it is valid DICOM. In many cases the AI component will include a small preprocessing step to convert Enhanced MR into something the model can use. This flexibility was important because hospitals do not always have the same modalities or the same acquisition protocols. Symphony focuses on the **interoperable interfaces**, not on prescribing one imaging data model.

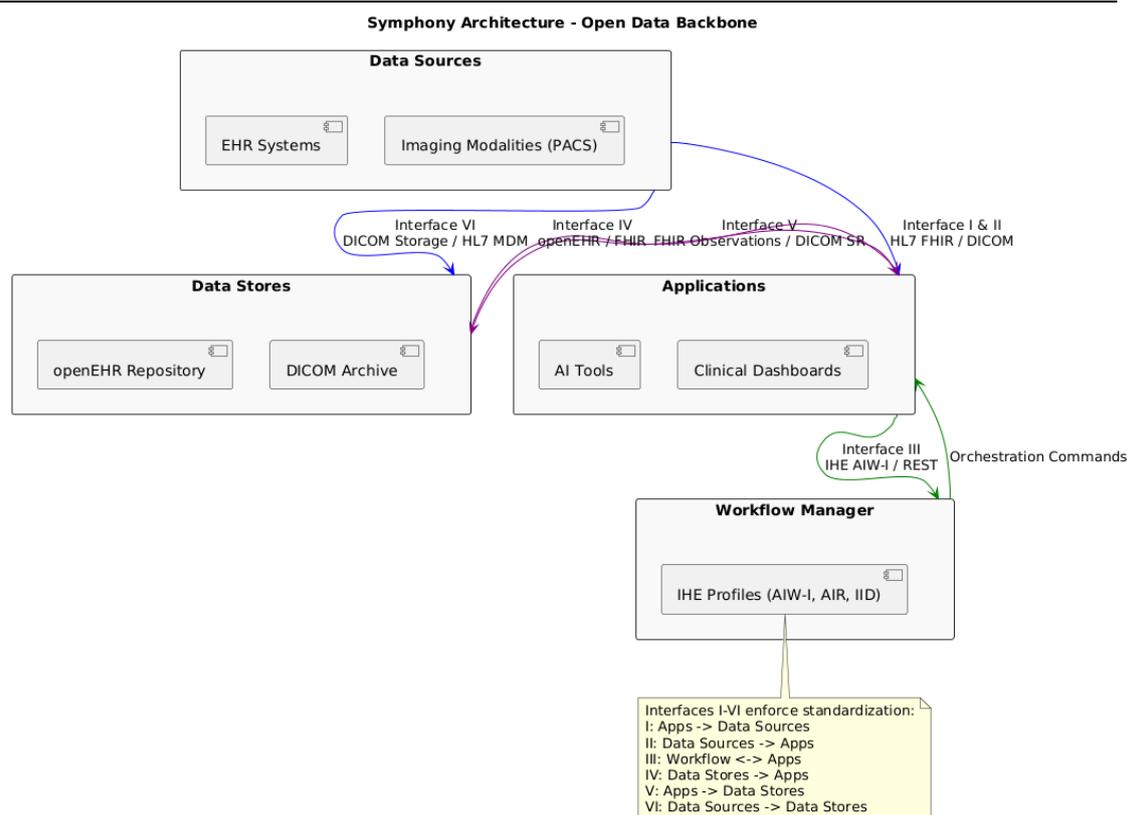


Figure 1 Symphony Architecture - Open Data Backbone

2.3 Philosophy: Evolution, Not Revolution

Symphony made a very clear decision: **reuse existing standards instead of creating new ones**. There are several reasons for this:

- 1. Avoid Adding “Standardisation Noise”**
 Healthcare already has many standards. If we create new ones, hospitals and vendors will face more complexity. This slows down adoption and increases cost.
- 2. Leverage Global Adoption**
 Standards like HL7 FHIR, OpenEHR, and DICOM are already widely used. Hospitals and vendors know them. By reusing these, Symphony reduces the learning curve and integration effort.
- 3. Ensure Interoperability Across Borders**
 Healthcare is international. If we use global standards, systems in different countries can communicate. This is very important for projects like Symphony that involve partners from The Netherlands, Sweden, Spain, Canada, and Türkiye.
- 4. Focus on Integration, Not Reinvention**
 The real challenge is not inventing new formats, but making existing ones work together. Symphony invests in building the **Open Data Backbone** and mapping between standards (like FHIR and OpenEHR), instead of creating new protocols.
- 5. Compliance and Trust**
 Existing standards already have compliance frameworks (GDPR, MDR). For example, IHE ATNA provides audit trails for security. Reusing these standards makes regulatory approval easier.

This philosophy is documented in **Deliverable D7.3/D7.5: Standardisation & Dissemination Plan** and **D5.3: Final Report for Standards and Open API**.

3 The Standards Engine – Technical How-To

Symphony project is not only about architecture, but also about **how to make standards work together in real workflows**. This chapter explains the technical approach for interoperability and AI integration.

3.1 The Hybrid Model: FHIR + OpenEHR

One big question was: **Which standard should we use for clinical data?**

FHIR is very popular for data exchange because it uses modern web technologies (REST, JSON). It is easy for developers and good for mobile apps. But FHIR has a problem: it is very flexible. Two systems can speak FHIR but still not understand each other because profiles are different.

OpenEHR solves this problem because it focuses on **semantic meaning**. It uses archetypes defined by clinicians, so data keeps its meaning for long time. This is very important for diseases like Multiple Sclerosis and Prostate Cancer where data are stored for decades.

So, Symphony decided to use **both**:

- **FHIR for the wire**: For communication between apps and lightweight clients.
- **OpenEHR for the store**: For backend persistence and complex queries.

This is described in **Deliverable D5.3: Final Report for Standards and Open API**.

3.1.1 Challenge

Mapping between FHIR and OpenEHR is hard. FHIR is graph-based, OpenEHR is tree-based. Symphony had to build transformation engines to convert FHIR resources into OpenEHR compositions. This was a big lesson learned.

Symphony Hybrid Model - FHIR and openEHR

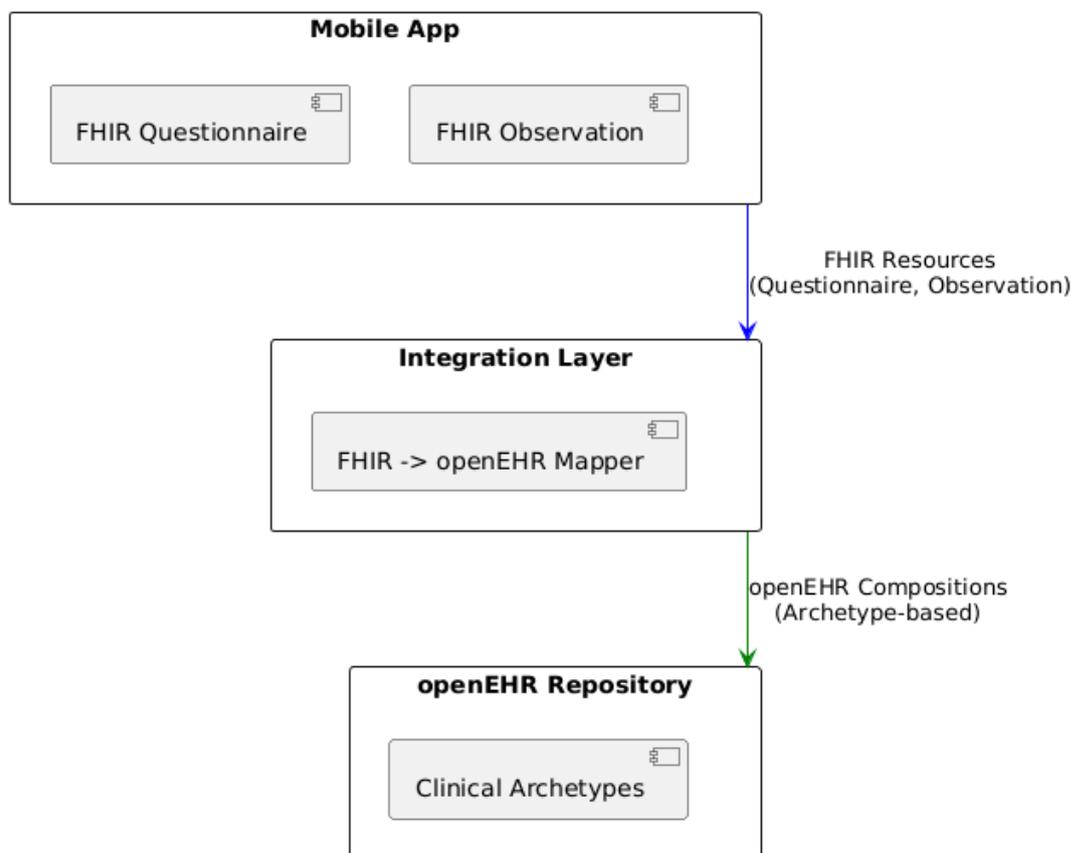


Figure 2 Symphony Hybrid Model - FHIR and OpenEHR

3.2 The AI Engine: DICOM SR and Segmentation

When AI analyses medical images, it produces results like lesion measurements, classifications, or segmentation masks. If these results are stored as simple images (JPEG with a red circle), they are **dead data**. Computers cannot read them, and workflows cannot automate decisions.

Symphony uses **DICOM Structured Reporting (SR)** and **DICOM Segmentation Objects** because:

- **Structured Reporting (SR)** allows AI to store results as machine-readable data:
 - Example: Lesion Diameter = 12 mm, SNOMED Code = Malignant.
 - These values can trigger alerts or guideline recommendations automatically.
- **Segmentation Objects** store pixel-level masks in a standardized way:
 - Example: Tumour region highlighted in CT scan.
 - This enables visualization and quantitative analysis later.

This approach is documented in **Deliverable D5.3: Final Report for Standards and Open API**.

Benefits

- **Interoperability:** Any PACS or viewer that supports DICOM SR can display AI results.
- **Automation:** Workflow engines can read SR data and trigger next steps.
- **Compliance:** DICOM SR supports audit trails and metadata for GDPR.

3.3 Workflow Orchestration: IHE Profiles

AI integration is not only about data formats. It is about **orchestration** – how tasks are triggered, monitored, and completed without breaking clinical workflow.

Symphony uses **IHE Profiles** for this:

- **IHE AIW-I (AI Workflow for Imaging):**
 - Defines actors and transactions for requesting and performing AI analysis.
 - Example: Workflow Manager creates a **DICOM UPS-RS (Unified Procedure Step)** for AI task.
- **IHE AIR (AI Results):**
 - Standardizes how AI results (DICOM SR, Segmentation) are stored and retrieved.
 - Ensures results are linked to original study and patient context.
- **IHE IID (Invoke Image Display):**
 - Allows one-click opening of images in PACS from Symphony dashboard.
 - Improves usability and reduces clicks for clinicians.

Challenges

- Many open-source PACS do not support UPS-RS yet.
- AIR profile is still evolving; lacks native support for interactive AI corrections (e.g., radiologist adjusting bounding box).

Addressing Interaction Gaps

During Symphony adoption, we saw that current IHE AIR profile does not include user interaction steps for approving or correcting AI results. This was a practical gap for workflows where clinicians need to validate or adjust AI findings. Now this missing part is also recognized by standard makers and is being addressed in the upcoming **IHE AIRAI (AI Result Approval)** profile. This shows that Symphony experience is aligned with how standards are evolving, even if Symphony is not directly contributing to the new profile.

3.4 Feedback Loop – Closing the Gap

One missing link in many AI workflows is the feedback loop. AI models can drift over time and need retraining with real clinical corrections. Symphony addresses this with the **Feedback Loop API**, described in Deliverable D2.5. When a clinician adjusts an AI result (for example resizing a segmentation mask or correcting a wrong lesion), this correction is stored again in the repository as a new **Ground Truth annotation**. This makes the AI pipeline learn from real hospital practice instead of only the initial training dataset.

Before any corrected result is sent back to an AI vendor or an external training environment, the Privacy Gateway applies the same **DICOM de-identification profiles** that are already used for the normal inference workflows. This means the corrected segmentation or updated SR is stripped of direct identifiers (name, date of birth, study UID if required) and only keeps the technical structures that the model needs. The Feedback Loop API only exposes de-identified data to external AI systems, while the hospital PACS keeps the full clinical version. This way, the model developers get valid ground-truth data for retraining, but the hospital stays compliant with GDPR and the internal data-sharing rules. The roundtrip is possible because Symphony separates *identity* from *content* and uses standard DICOM metadata to keep the technical alignment intact.

This mechanism is important because it links daily clinical practice with the continuous evolution of AI models, and it does this in a standardized and privacy-preserving way.

4 Clinical Impact – The Use Cases

Symphony project is not only theory. It was tested in **four real clinical scenarios** to show how standardisation works in practice. Each use case demonstrates different challenges and how standards helped to solve them.

4.1 Use Case 1: Prostate Cancer (PC)

4.1.1 Clinical Problem

Prostate cancer care faces a common challenge: the inability to utilize AI effectively due to fragmented and non-standardized data.

4.1.2 Standardisation Solution:

- **Guideline Automation:** Symphony automated **EAU guidelines** for prostate cancer.
- **Data Integration:** During the Symphony project we have used OpenEHR archetypes and templates for standardisation of data from all different source systems, including pathology, radiology, urology findings as well as patient-reported data. Using a Clinical Data Repository (CDR) for storing all source data in a standardised format we can utilize AI in a scalable manner, and especially algorithms by third parties.

4.2 Use Case 2: Abdominal Aortic Aneurysm (AA)

4.2.1 Clinical Problem

The decision to operate on an aneurysm is very critical and usually based on the maximum diameter of the aorta. But this is not always accurate because rupture risk depends on more complex biomechanical factors.

4.2.2 Innovation

Symphony introduced **chained AI inference** for better risk assessment:

- **Step 1:** First AI algorithm segments the aorta from CT angiography images and stores the segmentation as **DICOM Surface Segmentation Object**.
- **Step 2:** Second AI algorithm uses this segmentation to compute **biomechanical wall stress analysis** and other advanced metrics.
- **Step 3:** Results from the second AI are stored as **DICOM Structured Reporting (SR)** for clinical interpretation.

The **hospital PACS acts as the single data lake** between these AI algorithms. This is possible because:

- All intermediate and final results are stored in **interoperable DICOM objects**.
- PACS is a legacy system, but because Symphony uses standards, it can serve as the backbone for AI orchestration without replacing existing infrastructure.

4.2.3 Why This Matters

- **Interoperability enables chaining:** Without standardized outputs (DICOM SR and Segmentation), the second AI could not reuse the first AI's results.
- **Vendor-neutral integration:** Hospitals do not need new proprietary systems; PACS remains the central hub.

- **Automation and scalability:** Workflow manager orchestrates tasks using IHE AIW-I and ensures results are linked to the correct patient and study.

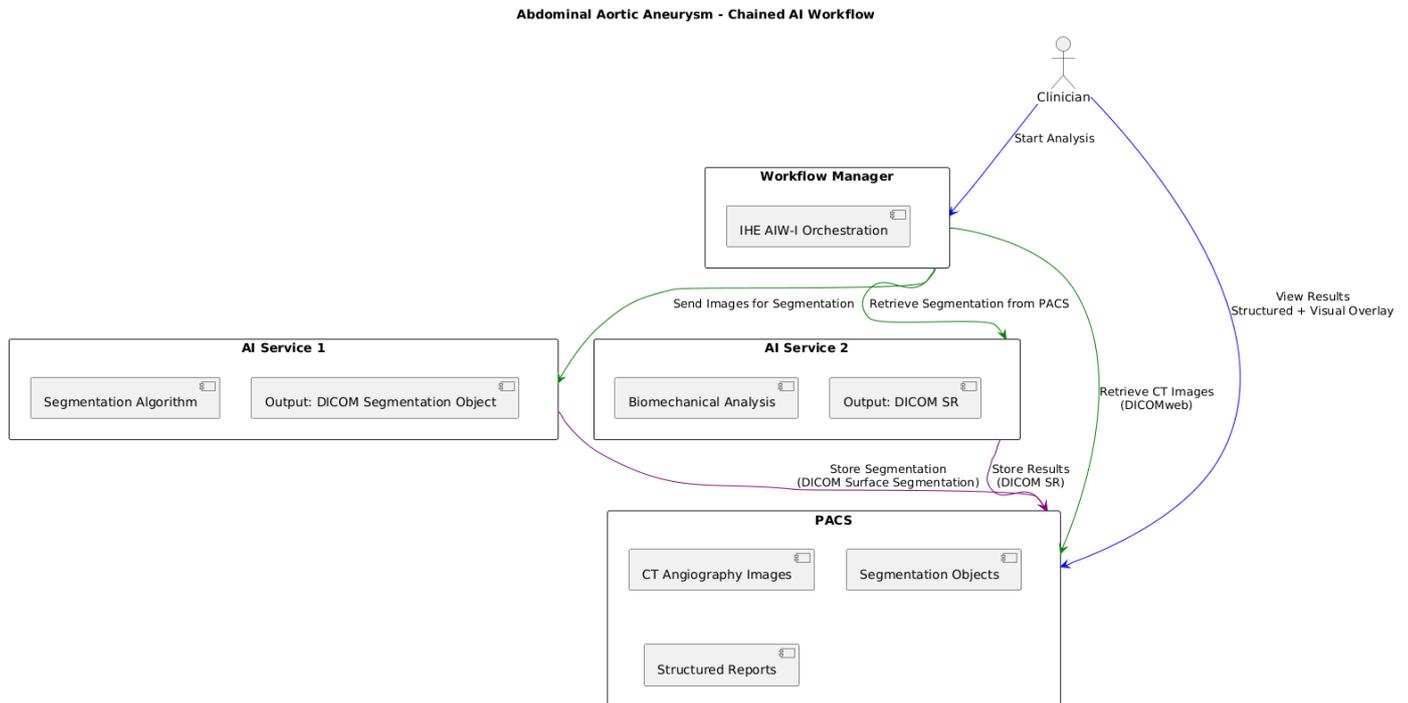


Figure 3 Abdominal Aortic Aneurysm - Chained AI Workflow

4.3 Use Case 3: Atrial Fibrillation (AF)

4.3.1 Clinical Problem

Patients often fail to follow medication plans.

4.3.2 Standardisation Solution:

- **Patient Engagement:** A chatbot powered by LLM interacts with patients to gather symptoms.
- **Hybrid Model:** Patient-Reported Outcome Measures (PROM) captured via mobile app using **FHIR questionnaires**, stored in **OpenEHR repository**.
- **Workflow Integration:** Data flows through standardized APIs.

Atrial Fibrillation Workflow - Patient Engagement

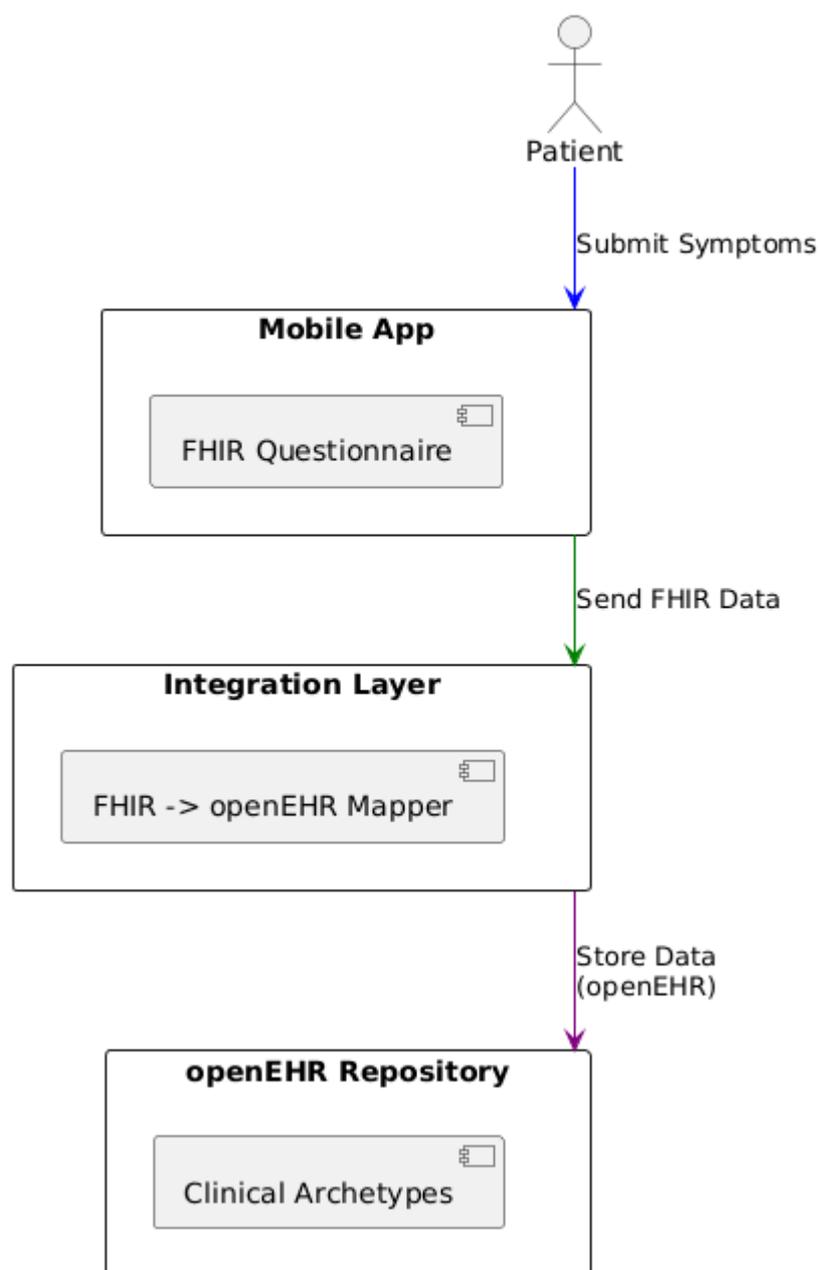


Figure 4 Atrial Fibrillation Workflow - Patient Engagement

4.4 Use Case 4: Multiple Sclerosis (MS)

4.4.1 Clinical Problem

MS is a chronic, disabling disease of the brain and spinal cord. Monitoring progression is complex because lesions evolve over time, and linked conditions like optic neuritis and sarcopenia require additional tests. Traditional workflows are slow—interpreting MRI can take 15 minutes to an hour per patient—and lack coordination between specialties.

4.4.2 Symphony Innovation

The MS use case in Symphony is a **multi-layered workflow combining imaging AI, mobile health apps, and predictive analytics**:

4.4.2.1 AI-Based Lesion Segmentation and Structured Reporting

- MRI images are processed by an AI algorithm that performs **lesion segmentation**.
- Outputs are stored as **DICOM Surface Segmentation Objects (SSO)** and **DICOM Structured Reports (SR)**.
- These results are visualized in an AI viewer (Philips and ARD Group) for clinician validation. Clinicians can **accept or reject** the AI findings.
- Once validated, results are stored in the **hospital PACS**, acting as the central data lake.

4.4.2.2 Chained AI Inference for Severity Analysis

- After validation, a second AI algorithm retrieves the segmentation and SR from PACS.
- It performs **severity scoring** and progression analysis (e.g., comparing lesion load between two time points).
- Results are encoded again in **DICOM SR** for interoperability and stored back in PACS.

4.4.2.3 Mobile Applications for Linked Conditions

- **Optic Neuritis Detection**: Mobile apps perform vision tests such as:
 - Contrast sensitivity
 - Visual acuity
 - Critical flicker fusion
- These results are captured using **FHIR resources** and stored in the clinical data repository.
- Additional apps monitor sarcopenia using motion capture and IMU sensors, integrated via FHIR.

4.4.2.4 Predictive Models

- Karolinska Institute develops ML models (Logistic Regression, GBM, XGBoost) to predict disability progression using:
 - Demographics
 - EDSS scores
 - Lesion features from MRI
 - Vision test results
- These models help clinicians identify high-risk patients early.

4.4.3 Standards and Interoperability

- **DICOM SR & SSO**: For AI imaging results and segmentation.
- **IHE AIW-I**: Orchestrates AI workflows. Note: As the hospital PACS lacked native **UPS-RS** support, Symphony implemented a custom UPS middleware to manage these tasks.
- **IHE AIR**: Standardizes AI result encoding and display.
 - While clinicians can accept or reject AI findings, Symphony identified this as a gap in current IHE AIR profiles and implemented a design that anticipates the future IHE AIRAI standard.
- **IHE IID**: Enables one-click image display for clinicians.
- **HL7 FHIR**: Used for vision test results and questionnaires. Note: Due to a lack of standard profiles for specialized eye tests, the project developed custom FHIR extensions to ensure data structure.

- **OpenEHR:** For storing longitudinal clinical data.

4.4.4 Key Features

- **Explainable AI:** Segmentation results displayed in 3D viewers (ARD Group) with transparency and slice-by-slice navigation.
- **Digital Twin:** Combines imaging, vision tests, and sensor data for holistic monitoring.
- **Interdisciplinary Collaboration:** Neurologists, radiologists, ophthalmologists, and physiotherapists share data via a privacy-preserving gateway.
- **Continuous Monitoring:** Mobile apps enable early detection of optic neuritis and sarcopenia.

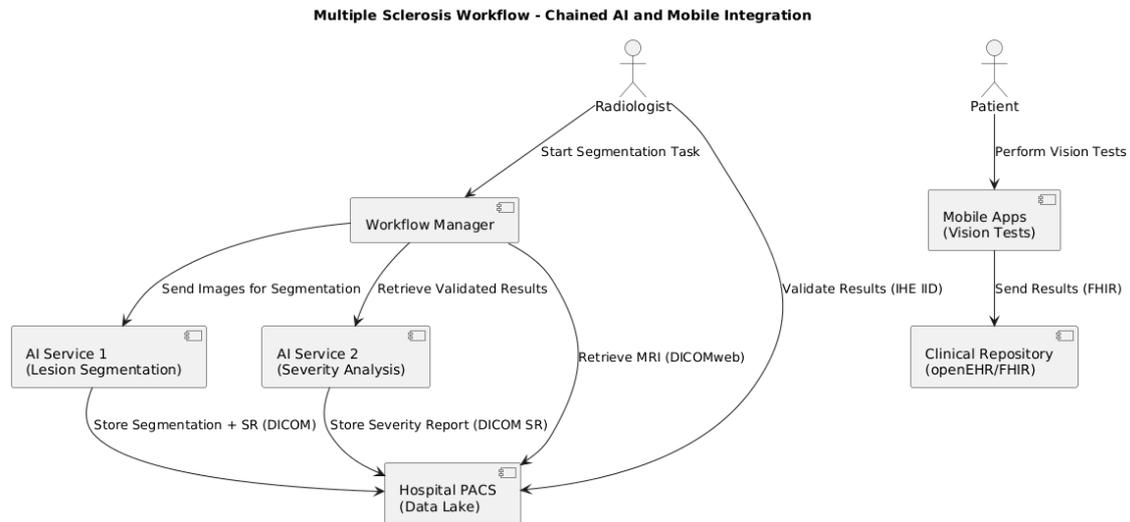


Figure 5 Multiple Sclerosis Workflow - Chained AI and Mobile Integration

5 Challenges and Lessons Learned

Symphony was not only a technical project, but it was also a big integration effort across hospitals, vendors, and research partners. We learned many things the hard way. Below are the main challenges and what we learned from them.

5.1 Vendor Readiness

Challenge:

Many AI vendors and hospital IT teams lacked familiarity with healthcare interoperability standards such as **DICOM SR**, **IHE AIW-I**, or **UPS-RS**. Most were accustomed to using custom APIs or transmitting results in non-standard formats like JPEG or PDF.

Example:

In the **Multiple Sclerosis** and **Aortic Aneurysm** use cases, the initial AI vendors delivered results as secondary capture images (screenshots). This made the data "dead" and useless for automation because the workflow manager could not extract or read lesion measurements. The Symphony team had to provide technical training to help the vendors implement **DICOM SR**.

Reference:

Deliverable **D5.3** describes this issue and how we created guidelines for vendors to adopt IHE profiles.

Lesson Learned:

Hospitals should integrate compliance requirements directly into their contracts and RFPs, explicitly stating that "AI must support DICOM SR and UPS-RS". It is also critical to start projects with technical training sessions for vendors to align expectations early.

5.2 Lack of UPS-RS Support in PACS

Challenge:

IHE AIW-I profile depends on **DICOM Unified Procedure Step – RESTful Services (UPS-RS)** for orchestration. But most PACS systems, especially open-source ones, do not support UPS-RS.

Example:

In the **MS use case**, we needed to create UPS tasks for lesion segmentation and severity analysis. The hospital PACS could not handle UPS-RS, so we built a custom UPS server as a middleware. This added complexity and delayed integration.

Reference:

Deliverable **D4.3** explains how we implemented orchestration using custom UPS services.

Lesson Learned:

Check PACS capabilities early. If UPS-RS is missing, plan for custom orchestration components.

5.3 Mapping Complexity Between FHIR and OpenEHR

Challenge:

We decided to use **FHIR for exchange** and **OpenEHR for persistence**. This is good for interoperability, but mapping between them is hard because:

- FHIR is graph-based.
- OpenEHR uses tree-based archetypes.

Example:

In the **Atrial Fibrillation use case**, patient-reported outcome measures (PROMs) were captured via FHIR questionnaires in a mobile app. These needed to be stored in

OpenEHR for longitudinal analysis. We had to build transformation engines to convert FHIR Observations into OpenEHR Compositions.

Reference:

Deliverable **D2.1** and **D5.3** describe the hybrid model and mapping challenges.

Lesson Learned:

If you use hybrid model, plan for mapping effort and test early. It is not trivial.

5.4 Standardisation Gaps

Challenge:

Some clinical tests do not have official FHIR profiles. For example, **vision tests for MS** (contrast sensitivity, flicker fusion) were missing.

Example:

In the **MS use case**, mobile apps collected vision test results. We had to create custom FHIR extensions because standard profiles were not available.

Reference:

Deliverable **D6.4** mentions these extensions and how they were implemented.

Lesson Learned:

Be ready to extend standards for niche clinical data. Document these extensions well for future reuse.

Future Standards – Closing the Gap

During Symphony adoption, we noticed that interactive user validation of AI results was missing in current standards. This gap is now also recognized by standard makers and is being addressed in the upcoming IHE AIRAI (AI Result Approval) profile. This shows that Symphony's practical experience is aligned with the evolution of standards, even if the project is not directly contributing to the profile development.

5.5 Data Privacy and Compliance

Challenge:

GDPR compliance was a big topic. We used **IHE ATNA** for audit trails and **DICOM de-identification profiles** for imaging. But making sure all partners follow same rules was challenging.

Example:

In the **AA use case**, biomechanical analysis required sharing CT images with external AI services. We had to implement de-identification pipelines before sending data.

Reference:

Deliverable **D1.3** and **D5.2** describe privacy tools and compliance strategies.

Lesson Learned:

Legal and ethical requirements should be part of technical design from start. Do not treat them as an afterthought.

Note: GDPR was a key requirement during Symphony design, but now the **EU AI Act** is also relevant. Symphony's approach with audit trails (IHE ATNA), traceability, and human-in-the-loop workflows aligns well with the Act's principles for high-risk AI systems.

5.6 Limited Support for DICOM Segmentation Objects in PACS

Challenge:

One practical limitation we encountered is that **not all PACS systems support the DICOM Surface Segmentation Object (SSO)**. Some hospitals still run older PACS versions that can store only classic DICOM images and basic Structured Reports, but they cannot ingest or visualize segmentation objects. This became visible especially in

the Aortic Aneurysm (AA) and Multiple Sclerosis (MS) use cases, where segmentation was a key part of the chained AI inference workflow.

Example:

In the AA use case, the first AI algorithm outputs a full 3D segmentation of the aorta as a DICOM SSO. But the hospital PACS in some settings could not store or index this object type. Without segmentation support, the second AI algorithm could not retrieve the result directly from the PACS for its biomechanical analysis.

How Symphony Managed This:

Instead of forcing the hospital to upgrade or replace their PACS, Symphony used a flexible design:

- If the PACS **supports** SSO, the segmentation is stored directly in the PACS.
- If the PACS **does not support** SSO, the segmentation is stored in an **external DICOM-compatible storage** (e.g., a DICOMweb-compliant object store) that behaves as an additional data endpoint in the Open Data Backbone.
- The Workflow Manager does not depend on PACS-native visualization; it only needs stable references, so the AI viewers can load segmentation objects from whichever store supports them.

Lesson Learned:

AI workflows should not assume full DICOM support in legacy PACS installations. A flexible architecture with alternative DICOMweb storage endpoints makes chained AI workflows possible even in hospitals with older imaging infrastructure. This avoids expensive PACS replacements and keeps the workflow interoperable.

5.7 Summary

Symphony showed that interoperability is not only about standards, but also about **people, processes, and readiness**. The main advice for future projects:

- Train vendors early.
- Validate PACS capabilities.
- Plan for mapping complexity.
- Prepare for gaps in standards.
- Integrate compliance from day one.

6 Guidance for Future Implementers

Symphony project was a big learning experience. We tried to combine standards, AI, and clinical workflows in real hospitals. It was not easy, but now we know what works and what does not. This chapter gives practical advice for anyone who wants to build similar systems. These are not just theory; they come from real problems we solved.

6.1 Start Small and Grow

Do not try to implement everything at once. Begin with one standard and one workflow. For example:

- Start with a **FHIR server** for basic data exchange.
- Then add **IHE AIW-I** for AI orchestration.
- Later integrate **OpenEHR** for persistence.

Example from Symphony:

In the **Atrial Fibrillation use case**, we first implemented FHIR questionnaires for patient-reported outcomes. Only after that worked, we added OpenEHR storage and AI chatbot integration.

Reference:

Deliverable **D6.3** shows this incremental approach.

6.2 Demand Standards in Contracts

Hospitals should put interoperability requirements in their RFPs. Write clearly:

- AI must support **DICOM SR** for results.
- Workflow must use **IHE AIW-I** and **UPS-RS**.
- Data exchange must use **FHIR** or **OpenEHR**.

Example:

In the **Prostate Cancer use case**, the first vendor delivered JPEG screenshots. We had to rework everything. If the contract had required DICOM SR, this problem would not happen.

Reference:

Deliverable **D5.3** explains why vendor compliance is critical.

6.3 Validate PACS Capabilities Early

Do not assume your PACS supports everything. Many PACS do not have UPS-RS or DICOMweb fully implemented.

- Test PACS for **UPS-RS** before planning AI orchestration.
- If missing, plan for middleware.

Example:

In the **MS use case**, we discovered late that PACS could not handle UPS tasks. We built a custom UPS server, but this delayed integration.

Reference:

Deliverable **D4.3** describes this workaround.

6.4 Prepare for Mapping Complexity

If you use hybrid model (FHIR + OpenEHR), mapping is hard.

- FHIR is flexible and graph based.
- OpenEHR uses strict archetypes.

Plan for:

- Transformation engines.
- Validation rules.
- Testing with real clinical data.

Example:

In the **AF use case**, mapping patient reported outcome measures (PROM) from FHIR questionnaires to OpenEHR compositions took weeks.

Reference:

Deliverable **D2.1** and **D5.3** show mapping strategies.

6.5 Expect Gaps in Standards

Not all clinical tests have profiles. For example:

- Vision tests for MS (contrast sensitivity, flicker fusion).
- Sarcopenia monitoring.

You will need custom FHIR extensions. Document them well for reuse.

Example:

In the **MS use case**, we created FHIR extensions for eye tests.

Reference:

Deliverable **D6.4** explains these extensions.

6.6 Make Compliance Part of Design

GDPR and MDR are not optional.

- Use **IHE ATNA** for audit trails.
- Apply **DICOM de-identification profiles** for imaging.
- Secure APIs for data exchange.

Example:

In the **AA use case**, we implemented de-identification before sending CT images to external AI services.

Reference:

Deliverable **D1.3** and **D5.2** describe compliance tools.

6.7 Build for User Experience

Technical integration is useless if clinicians hate it.

- Use **IHE IID** for one-click image display.
- Keep click-path short.
- Provide clear visualization for AI results.

Example:

In the **MS use case**, radiologists could open PACS with AI overlays in one click. This was key for adoption.

Reference:

Deliverable **D6.4** shows the UI design.

6.8 Plan for Marketplace Integration

If you want scalability, design for plug-and-play AI.

- Use **Application Registration API**.
- Define manifests for input/output.
- Automate wiring using standards.

Example:

In the **AA use case**, biomechanical analysis algorithm was integrated using manifest-based approach.

Reference:

Deliverable **D2.5** explains marketplace architecture.

6.9 Final Advice

- Start small but think big.
- Standards are your friend, but they need effort.
- Compliance is not a checkbox; it is a design principle.
- User experience decides success, not just technology.

7 Conclusion

When we started Symphony project, the goal was very clear but also very hard: make AI useful in real clinical workflows by solving the interoperability problem. After almost four years of work, we can say the main idea was correct – the problem is not the algorithm, the problem is the infrastructure.

What We Learned

- AI cannot scale without **standardized data exchange**. Hospitals have many systems, many vendors, and without standards everything becomes custom integration.
- Using existing standards like **HL7 FHIR, OpenEHR, DICOM, and IHE profiles** was the right decision. It was not easy, but it avoided creating new “standardisation noise”.
- Interoperability is not only technical. It is also about **vendor readiness, legal compliance, and user experience**.

What We Achieved

- We built an **Open Data Backbone** with six interface model described in **D2.1**.
- We implemented **hybrid model** (FHIR for wire, OpenEHR for store) and solved mapping challenges.
- We standardized AI workflows using **IHE AIW-I, AIR, and IID**, and stored AI results in **DICOM SR** and **Segmentation Objects**. Symphony also introduced the concept of a Feedback Loop API (see D2.5), enabling clinician corrections to be stored as new ground truth annotations for AI retraining. This closes the loop for continuous improvement.
- We validated everything in **four real clinical use cases**: Prostate Cancer, Aortic Aneurysm, Atrial Fibrillation, and Multiple Sclerosis.
- We designed an **Open Algorithm Marketplace** for plug-and-play AI integration (**D2.5**).

Why This Matters

Healthcare is moving to value-based care and precision medicine. AI can help, but only if it is integrated in a safe and scalable way. Symphony shows that this is possible when you combine:

- Standards
- Compliance
- User-centred design

Final Message

Infrastructure is destiny. Build the right architecture, and AI will follow—secure, compliant, and ready for future regulations. Symphony is not the end; it is a blueprint for the future. We hope this white paper helps hospitals, vendors, and policymakers to understand that interoperability is the foundation for digital health.