

AI DOCUMENT PROCESSOR – UK HEALTHCARE PLAYBOOK

Document-to-FHIR Automation for NHS Organisations

From document to approved, interoperable data – across email, SharePoint and Copilot.

1. Overview

NHS teams receive critical information in many different ways: documents emailed into shared mailboxes, files saved into SharePoint, and increasingly, ad-hoc requests via digital assistants and Copilot experiences.

This solution turns those documents into structured, FHIR-compliant data automatically, using Microsoft Azure.

No matter how the process starts, the workflow:

1. Extracts key clinical and administrative data using Azure Document Intelligence
2. Generates a FHIR Bundle following NHS interoperability standards
3. Stores that bundle securely in Azure Blob Storage
4. Transforms the bundle into a clean HTML summary using a FHIR-to-HTML function
5. Sends a review and approval email, so staff can validate and approve the data in one click

The result is a consistent, standards-driven pipeline that reduces manual work, improves data quality, and prepares information for integration with existing and future NHS systems.

2. Who This Is For

This solution is designed for:

- NHS Trusts and ICSs looking to standardise document-driven workflows
- Clinical and operational teams who rely on documents arriving via email or SharePoint
- Digital and IT teams who want FHIR-ready data without building everything from scratch
- Innovation and transformation teams exploring Copilot-enabled, user-initiated processing

Typical use cases include:

- Referrals and triage
- Community and outpatient assessment forms
- Discharge and clinic letters
- Diagnostics/Pathology reports

- Administrative and back-office forms

3. Three Entry Points into the Workflow

The solution supports three flexible ways to start the pipeline, so it can fit naturally into existing NHS processes.

3.1 Email: When an email arrives in an NHS inbox

- A patient referral, form, or letter is sent to a monitored NHS mailbox.
- The workflow automatically detects new messages, extracts the attachment(s), and begins processing.
- No extra steps or portals required – users keep using email as they do today.

3.2 SharePoint: When a file is saved in a document library

- A document is uploaded or saved to a designated SharePoint document library.
- The workflow is triggered as soon as the file appears or is modified.
- Ideal for teams already using SharePoint as a central repository for clinical or admin documents.

3.3 Copilot: When a user initiates processing through a Copilot agent

- A user asks a Copilot agent (e.g. “Process this document and create a FHIR summary”).
- Copilot passes the file or reference into the same Azure workflow.
- Perfect for ad-hoc requests, pilots, and interactive scenarios where staff want to trigger processing on demand.

All three entry points feed into the same core pipeline, ensuring consistent extraction, FHIR generation, and approval.

4. What the Solution Does

Key Capabilities

- **Multiple triggers:** email, SharePoint, and Copilot-initiated
- **AI-powered document understanding** with Azure Document Intelligence
- **FHIR Bundle generation** (R4) to NHS interoperability standards
- **Secure storage** of structured data in Azure Blob Storage
- **FHIR-to-HTML rendering** for human-readable summaries
- **Email-based approval workflow** with clear Approve/Reject options

5. How It Works – End-to-End Journey

Regardless of the entry point, all paths converge into the same processing pipeline.

Step 1: Ingestion – Email, SharePoint, or Copilot

- **Email inbox**
A Logic App monitors a configured NHS mailbox. When a new email arrives with an attachment (e.g. PDF, form, letter), it automatically retrieves the file.
- **SharePoint library**
Another Logic App listens for new or updated files in a specific SharePoint document library and collects the uploaded document.
- **Copilot agent**
A Copilot agent (or similar digital assistant) allows a user to select or upload a document and then calls the Azure workflow with that file as input.

 All three methods output a document file that is passed into the AI extraction stage.

Step 2: AI Extraction with Azure Document Intelligence

The document is sent securely to Azure Document Intelligence, which:

- Reads printed and handwritten text
- Extracts key-value pairs, tables, and important fields
- Produces a structured JSON representation of the document
- Adds confidence scores for every extracted field

These confidence scores make it clear which values are highly reliable and which may need human review.

Step 3: FHIR Bundle Creation

An Azure Function receives the extracted JSON and:

- Maps fields into appropriate FHIR R4 resources such as:
 - Patient
 - Observation
 - MedicationRequest
 - Additional resources depending on the use case
- Packages them into a single FHIR Bundle

This is where unstructured document content becomes interoperable, standards-based health data.

Step 4: Secure Storage in Azure Blob

The generated FHIR Bundle is saved as JSON in Azure Blob Storage.

Benefits:

- Central, secure storage of structured clinical data

- Clear audit trail from: **Original email/SharePoint/Copilot request → Extracted fields → FHIR Bundle**
- A natural integration point for analytics, downstream systems, or APIs

Step 5: FHIR-to-HTML Transformation

A second Logic App (triggered when a new bundle is created in Blob) then:

1. Retrieves the FHIR Bundle JSON
2. Calls a FHIR-to-HTML Azure Function
3. Receives a clinician-friendly HTML summary

The HTML summary typically includes:

- Patient demographics
- Key clinical observations or measurements
- Medications or requests
- Any other relevant structured data

This makes FHIR understandable to end users without exposing raw JSON or technical details.

Step 6: Approval Email Workflow

The same Logic App then sends an approval email via Microsoft 365 / Outlook:

- The email body contains the HTML summary
- The approver can view all key information at a glance
- Built-in Approve / Reject options allow for a quick decision

Approved data can be:

- Logged for audit
- Passed to other systems
- Used to update downstream workflows or records

Rejected items can be flagged for review, correction, or follow-up.

6. Why This Matters for the NHS

Interoperability by Design

- Data is structured as FHIR R4, aligning with NHS interoperability standards and Shared Care Record initiatives.
- It can be integrated with EPRs, ICS data platforms, and other systems with minimal rework.

Safety, Governance, and Transparency

- Every step is auditable, from initial entry point (email, SharePoint, Copilot) to FHIR Bundle and approval outcome.
- Confidence scores support risk-based review of extracted fields.
- Azure's security and compliance features support NHS information governance expectations.

Efficiency and Reduced Manual Work

- Less manual data entry from PDFs, emails, and forms.
- Faster turnaround from document receipt to usable, structured data.
- Staff interact through familiar tools: email, SharePoint, and modern Copilot experiences.

7. Technical Building Blocks (High-Level)

The solution is built entirely on Microsoft cloud technologies:

- **Microsoft 365**
 - NHS mailbox for email ingestion
 - SharePoint Online as a document entry point
- **Azure**
 - **Logic Apps** – workflow orchestration and triggers for email, SharePoint, and Blob
 - **Azure Document Intelligence** – AI-powered document extraction
 - **Azure Functions** – FHIR Bundle builder and FHIR-to-HTML renderer
 - **Azure Blob Storage** – secure storage of FHIR Bundles
- **Copilot / Agents**
 - Custom Copilot or conversational agent to trigger processing on demand

All components are designed to be modular, scalable, and configurable to local NHS requirements.

8. Implementation Approach

A typical implementation follows these stages:

- 1. Discovery & Design**
 - Identify document types and target pathways
 - Define FHIR mapping and approval requirements
- 2. Model & Workflow Configuration**
 - Configure Document Intelligence models for the chosen documents
 - Set up Logic Apps for email, SharePoint and Blob triggers
 - Build and configure the FHIR builder and FHIR-to-HTML functions
- 3. Pilot: Email / SharePoint / Copilot**

- Start with one entry point (e.g. email or SharePoint)
- Optionally add Copilot for on-demand processing

4. Validation & Governance

- Test end-to-end accuracy
- Validate with clinical/operational leads
- Tune thresholds and approval rules

5. Rollout & Extension

- Extend to more document types, teams, and pathways
- Integrate with downstream systems (EPRs, data platforms, etc.)

9. Key Benefits at a Glance

- Three flexible entry points: **Email, SharePoint, Copilot**
- Automated conversion from document to FHIR-compliant data
- Secure, Azure-based storage and processing
- Clinician-friendly HTML summaries instead of raw technical formats
- Built-in approval workflow with full audit trail
- Reusable pattern for multiple NHS pathways and organisations

Frequently Asked Questions (FAQs)

Document-to-FHIR Automation Workflow

1. What problem does this solution solve for NHS teams?

NHS teams receive large volumes of documents — often as PDFs or scans — that require manual review and re-keying into clinical systems.

This solution uses automation and AI to transform those documents into structured, FHIR-compliant data, reducing administrative burden and improving data quality.

2. What are the three ways a workflow can start?

The solution supports three flexible entry points to match existing NHS processes:

1. Email Trigger

When a document arrives in an NHS mailbox, the workflow extracts the attachment and begins processing automatically.

2. SharePoint Upload Trigger

When a file is saved or uploaded into a designated SharePoint document library, the workflow triggers instantly.

3. Copilot Agent Trigger

A user can manually initiate document processing through a Copilot-based conversational interface, ideal for ad-hoc or on-demand extraction.

All three entry methods feed into the same AI-powered pipeline and return consistent, FHIR-ready results.

3. What types of documents can the system handle?

Azure Document Intelligence can process:

- PDFs
- Scanned documents
- Forms and templates
- Typed and handwritten text
- Multi-page clinical or admin documents

Custom models can also be trained for NHS-specific forms or local templates.

4. How does the AI extract information from documents?

The solution uses Azure Document Intelligence to read the document and extract:

- Key-value pairs
- Tables
- Entities
- Clinical or administrative fields
- Handwritten text

Each extracted field includes a confidence score, providing transparency and supporting safe human review.

5. How is the data converted into FHIR?

An Azure Function converts the extracted information into a structured FHIR R4 Bundle, using appropriate resources such as:

- *Patient*
- *Observation*
- *MedicationRequest*
- Other resources depending on the document type

This ensures interoperability with NHS systems and shared care record platforms.

6. Where is the generated FHIR Bundle stored?

Every FHIR Bundle is stored securely in Azure Blob Storage within NHS-compliant Azure regions.

This provides an audit trail from:

Original document → Extracted fields → FHIR Bundle → Approval decision

7. How does the approval process work?

After the FHIR Bundle is created:

1. It is transformed into a clear, readable HTML summary via a FHIR-to-HTML Azure Function.
2. An email is sent to the designated reviewer or approver.
3. The approver can Approve or Reject the extracted data in one click.

This supports governance, validation, and clinical safety reviews.

8. Does the workflow support NHS interoperability standards?

Yes — the solution is built around FHIR R4, the mandated interoperability standard for NHS England.

This enables seamless integration with current and future systems such as:

- EPRs
- Shared Care Records
- ICS data platforms

9. Is the solution secure and compliant with NHS requirements?

Absolutely. The solution:

- Runs entirely within Azure, following NHS cloud security principles
- Supports DTAC expectations (cyber, data protection, interoperability)
- Uses private storage accounts and secure access patterns
- Provides full auditability and traceability

All data remains within NHS-approved cloud regions.

10. Can this workflow be adapted to different NHS use cases?

Yes — the pattern is reusable across:

- Referral management
- Discharge summaries
- Assessment forms
- Community care documentation
- Diagnostics and pathology documents
- Administrative workflows

The underlying extraction and FHIR mapping logic can be customised for any document template or pathway.

11. How can organisations extend or customise the solution?

Organisations can:

- Add new fields to the extraction model
- Create additional FHIR mappings
- Add new approval steps or routing logic
- Integrate with APIs, EPR systems, or shared care records
- Add additional Copilot agent commands for user-initiated processing

The architecture is modular and designed for flexibility.