Quick Reference Guide

FHIR Fundamentals for Architects

Topic: FHIR Origins, Architectural Frameworks & Interoperability Components

History & Evolution of FHIR

FHIR (Fast Healthcare Interoperability Resources) is an interoperability standard developed by the global standards development organisation HL7 (Health Level Seven International). HL7 has also developed HL7 Version 2 (V2), Version 3 (V3), and the Clinical Document Architecture (CDA) standards.

The Evolution of FHIR

Scroll down to view the timeline of the development of FHIR and HL7 Standards, including relevant milestones within the Australian context.

Year	Milestone	Description	Scope
1987	HL7 releases Version 2 (V2)	Widely adopted for clinical messaging. Flexible but lacked semantic clarity.	Global
1995	HL7 releases Version 3 (V3)	Introduced structured, semantically rich messaging. Too complex for wide-spread adoption.	Global
2005	HL7 releases V3 CDA (Clinical Document Architecture)	XML-based document format. Used for specific use cases but limited broader interoperability.	Global
2010	Interoperability challenges persist	V2, V3 and CDA fail to meet modern data exchange demands.	Global
2011	Birth of FHIR	HL7's "Fresh Look" initiative leads to Grahame Grieve proposing FHIR.	Global
2012	HL7 releases the Draft Standard for Trial Use 0 (DSTU 0)	Shared by HL7 for community review, DSTU 0 was a working draft (a proof of concept) to test and demonstrate FHIR's potential	Global
2014	HL7 releases DSTU 1	First official release for real-world testing.	
2015	HL7 releases DSTU 2	Expanded and refined FHIR resources, introduced the FHIR Maturity Model (FMM), and enhanced terminology services, profiles and implementation guides. It set the foundation for FHIR Release 4 (R4).	Global

2014–2016	Early FHIR exploration by HL7 Australia & ADHA	FHIR pilot projects begin in Australia. ADHA explores FHIR's potential.	Australia
2016	HL7 Australia FHIR Connectathon	First national Connectathon to promote developer engagement.	AU Australia
2017	HL7 releases FHIR DSTU 3	Adds support for clinical/admin workflows and implementation tooling.	Global
2017	FHIR-based API for MHR previewed	ADHA explores FHIR interfaces for provider/consumer access to My Health Record (MHR).	Australia
2018	FHIR adoption expands	UK, Canada, and Australia begin national implementations.	Global, Australia
2018	Secure Messaging Framework released	ADHA framework includes FHIR as a long-term messaging solution.	Australia
2019	HL7 releases FHIR R4	First normative content (Patient, Observation, REST API), stable foundation for production use.	Global
2019	FHIR Implementation Guides for Australia (drafts) published	Localised profiles (e.g., medications, observations) released by ADHA & CSIRO.	Australia
2020	FHIR mandated in U.S. regulations	21st Century Cures Act Final Rule requires FHIR-based APIs.	Global
2020	COVID-19 accelerates digital interoperability	FHIR used in rapid digital health responses (e.g. eReferrals).	Australia
2021	SMART on FHIR V2 & Bulk Data API	Support for large-scale, third-party data access and population health.	Global
2021	National Digital Health Strategy update includes FHIR	Australian strategy names FHIR as key interoperability enabler.	Australia
2021	AU Core FHIR IG v1.0 released	The Commonwealth Scientific and Industrial Research Organisation’s (CSIRO) Australian eHealth Research Centre (AEHRC) publishes the Australian Core FHIR implementation guide.	Australia
2022	FHIR R5 Ballot released	Community review version introduces new resources and features.	Global

2022	e-Referral and Provider Directory IGs published	Australia publishes national FHIR-based standards to support digital services.	Australia
2023	FHIR R5 final release	Adds canonical extensions, improves search, supports advanced use cases.	Global
2023	*My Health Record* FHIR APIs in production	FHIR-based MHR APIs released for specific data views (e.g. patient summaries).	Australia
2023	ADHA roadmap for FHIR-based interoperability	National vision shifts toward FHIR-centric architecture.	Australia
2024 (Ongoing)	FHIR R6 Development	Normative expansion, real-time interoperability, global alignment.	Global
2024 (Ongoing)	FHIR Training & Adoption Programs	ADHA, HL7 Australia, and universities expand national capacity-building.	Australia
2026 (Expected)	FHIR R6 Final Release	Fully normative release expected, no breaking changes. Major stability milestone.	Global

4 Levels of Interoperability

There are 4 levels of interoperability that are supported by FHIR:

Interoperability level	**How FHIR contributes to interoperability**
Foundational The ability to send and receive (exchange) data	Provides a standard API interface for retrieving and submitting data.
Structural The data exchanged follows a format or syntax with consistent message structures.	Defines consistent formats through “resources” like Patient, Observation, MedicationRequest.
Semantic The meaning of the data is preserved across systems.	Leverages terminologies like SNOMED CT and LOINC within FHIR resources to ensure shared meaning.
Organisational Legal, governance, and workflows support trusted exchange.	Enables integration with national digital health infrastructure like My Health Record via FHIR AU Base and ADHA frameworks.

Enterprise Architecture (EA) Layers

EA layers provide a structured framework to describe how an organisation's systems, processes, and technology align with its strategic goals. Each layer builds upon the other to provide a cohesive and scalable architecture that supports interoperability, innovation, and strategic transformation.

5 Enterprise Architecture (EA) Layers

EA Layers	Focus	How FHIR aligns with EA
Business Architecture	Organisational strategy, processes, functions, and roles.	Supports clinical workflows and care coordination by enabling data sharing across departments and organisations.
Information (or Data) Architecture	Structure and flow of data within the organisation.	Structures healthcare data using defined resources and datatypes to ensure consistency and semantic interoperability
Application Architecture	Software systems and how they interact to support business processes.	Defines how software systems (e.g. EHRs, registries) interact using standard APIs and resources.
Technology (or Infrastructure) Architecture	Hardware, networks, and platforms used to run applications and store data.	Cloud hosting environments, on-premises servers, databases, API gateways, security tools. Can be implemented using standard web technologies (REST, JSON, OAuth2), making it compatible with modern IT infrastructure.
Security Architecture	Mechanisms for protecting information, ensuring privacy, and managing access control.	Identity management, encryption, audit trails, consent models. Through Implementation Guides, Profiles, and Capability Statements, FHIR provides governance tools that align with the compliance and standardisation goals of EA frameworks.

FHIR Architectural Principles

FHIR's architectural principles focus on making it efficient, scalable, and easy to implement in the real world. These principles create a flexible and robust foundation for building interoperable digital health systems.

Reuse and Composability

FHIR resources are designed according to the 80/20 rule, prioritising the 20% of features that meet 80% of interoperability needs. Resources are created to cover broad, common requirements, avoiding duplication and redundancy. When specific needs arise, they can be addressed through extensions and profiles, which allow for tailored adaptations.

Resources are also highly composable, often referencing one another, enabling complex structures to be built from smaller, reusable building blocks.

Scalability.

By following the REST architectural style, FHIR ensures that all transactions are stateless, meaning no session information is stored between requests.

This reduces memory demands and eliminates the need for ‘sticky sessions’ (where a user must always be directed to the same server to maintain state). Instead, requests can be handled by any available server, enabling horizontal scalability, i.e. the ability to add more servers to a system to increase capacity and distribute workload efficiently.

Performance

FHIR resources are lightweight and network-friendly, designed for efficient data exchange. Optimised formats exist to enhance performance in complex, multi-system transactions, although most implementations find standard JSON and XML formats sufficient.

Usability

FHIR resources are designed for both technical and non-technical users. Even without knowledge of XML or JSON syntax, the contents of a resource can be opened in a browser or text editor and still be readily understood.

Data Fidelity

FHIR is strongly typed and integrates mechanisms for linking and validating clinical terminologies. XML and JSON resources can be validated both syntactically and against business rules, ensuring data integrity and supporting semantic interoperability.

Implementability

FHIR was created to be developer-friendly and widely adopted. It uses familiar industry standards and common web technologies, making it straightforward to learn, implement, and deploy across diverse communities and systems.

Enterprise Architecture (EA) Layers

FHIR (Fast Healthcare Interoperability Resources) is an interoperability standard developed by the global standards development.

Topic: FHIR Information

FHIR Modules

FHIR Modules represent the organisational structure of the FHIR specification, dividing it into different functional areas of healthcare and system architecture. Each module groups related resources and capabilities, making it easier for implementers to navigate the specification and focus on relevant components for their use case.

There are **6 levels of FHIR resource modules**, grouped according to function.

Level 1: Foundation Resources

Foundation resources are the core, reusable building blocks that support the structure, interoperability, and governance of all other FHIR resources.

They are often used for infrastructural tasks such as providing shared definitions, technical frameworks, and common patterns that other resources depend on. While they do not usually represent direct clinical data, and are not always referenced by other resources, they play a critical role in enabling consistent and reliable data exchange.

Foundation resources are used to describe system capabilities, constrain resources, and define how FHIR is implemented.

Foundation Resources	Description	Examples
Security	Covers the principles, documentation, and resources required to ensure data security, privacy, user consent, and integrity across FHIR-based systems.	Consent, Provenance, AuditEvent
Conformance	Describes how systems implement FHIR, define constraints on resources, and	CapabilityStatement, ImplementationGuide

	specify operational capabilities. They are essential for ensuring interoperability by documenting what a system supports and how resources should be used or tailored.	OperationDefinition, SearchParameter
Terminology	Supports the use of standard clinical terminologies and vocabularies, including resources for managing code systems, value sets, and concept mappings.	CodeSystem, ValueSet, ConceptMap, NamingSystem
Documents	Supports the creation, management, and exchange of clinical documents and collections of healthcare information. These resources help package multiple related resources together into a coherent, shareable document format that can be exchanged between systems.	Bundle, Composition, DocumentReference
Other	A catch-all category for content that doesn't neatly fit under the main specification models like Security, Conformance, Terminology, or Documents.	n/a

---

Level 2: Base Resources

Base resources represent the key participants, objects, processes, and administrative controls involved in healthcare.

Base Resource Categories	Description	Examples
Individuals	These represents the people involved in healthcare.	Patient, Practitioner, RelatedPerson
Entities	These represent organisations, things or locations that play a role in healthcare.	Hospital, clinic, device
Workflow	These represent processes and tasks that manage healthcare activities, for example appointments, procedures, or care plans	Appointment, AppointmentResponse, CarePlan

Management	These relate to administrative and governance-related elements such as consent, coverage, or authorisation.	Consent, PractitionerRole

Level 3: Clinical Resources

Clinical resources represent the core healthcare information directly related to patient care, diagnosis, treatment, and clinical activities. These resources capture clinical observations, interventions, and care delivery details that are essential for managing health and wellbeing.

Clinical Resource Categories	Description	Examples
Clinical	Patient-related clinical information such as conditions, observations, and problems.	Condition, Observation, AllergyIntolerance
Diagnostic	Diagnostic tests, reports, and results.	DiagnosticReport, ImagingStudy, Specimen.
Medication	Medications, prescriptions, and administration.	Medication, MedicationRequest, MedicationAdministration.
Care Provision	Plans, care teams, and care delivery.	CarePlan, CareTeam, Encounter
Request and Response	Clinical requests and their outcomes.	ServiceRequest, Procedure, ReferralRequest

Level 4: Financial Resources

Financial resources cover the administrative and monetary aspects of healthcare, including billing, payments, and supporting information. These resources help manage the financial transactions and agreements involved in delivering healthcare services.

Financial Resource Categories	Description	Examples
Support	Information that supports financial processes but are not directly financial transactions.	EligibilityRequest, EligibilityResponse

Billing	Invoicing, claims, and billing details.	Claim, Invoice
Payment	Relates to payment transactions and financial reconciliations.	PaymentNotice, PaymentResponse
General	General financial management and administrative aspects.	Account, ChargeItem, ChargeItemDefinition

Note: A more detailed description about the Financial Module can be found here

---

Level 5: Specialised Resources

Specialised resources are designed to address specific healthcare domains or advanced use cases that go beyond basic clinical and administrative information. These resources support areas such as public health, research, clinical decision support, and quality measurement.

Specialised Resource Categories	Description	Examples
Public Health and Research	Facilitates data collection, reporting, and management for public health surveillance and clinical research.	ResearchStudy, ResearchSubject
Definitional Artifacts	Defines terminologies, protocols, and decision rules used in healthcare.	ActivityDefinition, PlanDefinition, Measure
Clinical Decision Support	Supports automated or assisted clinical decision making.	GuidanceResponse, Library
Quality Reporting	Focused on quality measurement, reporting, and improvement initiatives.	Measure, MeasureReport, Evidence

---

Level 6: Other

Level 6 does not include resources. Instead, it builds upon the composition framework established by the first five models.

Topic: FHIR Conformance

FHIR Resources Used for Conformance

The FHIR Conformance Model is the practical toolkit that brings together all the resources used to declare, constrain, and validate FHIR implementations. These resources work together to show what a system can do, how data should be structured, which roles actors (a system, device, or person) play, and how interoperability is achieved.

StructureDefinition

A resource that defines the structure and rules for FHIR data elements, resources, types, or profiles. It describes what data is allowed, how it is organised, and how it can be extended or constrained.

ActorDefinition

A relatively new FHIR resource that defines the role and responsibilities of an actor (a system, device, or person) in an interoperability scenario. It specifies what part that actor plays in workflows, which interactions it is expected to support, and which CapabilityStatements may be associated with it.

Examples:
ActorDefinition examples include roles like FHIR Client, FHIR Server, Immunisation Submitter, Prescription Dispenser, Patient Access App, or Laboratory System. Each ActorDefinition specifies who does what in an interoperability workflow, linking the role to its required capabilities.

ImplementationGuide

Bundles conformance artifacts for a specific use case or jurisdiction.

SearchParameter

A formal definition of how you can search for resources on a FHIR server. It tells the server what fields or properties can be searched on, what type of data to expect, and how the search should work.

It enables clients to query for specific subsets of resources.
Each SearchParameter defines:

- Name – The search keyword (e.g., name, date, code).
- Base Resource(s) – The resource types the search applies to.
- Type – The kind of search (e.g., string, date, token, reference).

- Expression – A FHIRPath statement pointing to the data element in the resource.

Standard vs Custom
- Standard SearchParameters are defined by the FHIR specification.
- Custom SearchParameters can be created for implementation-specific needs.

OperationDefinition

A formal description of a custom or extended operation that a FHIR server can perform, beyond the standard RESTful CRUD interactions. It defines what the operation does, what inputs it expects, and what it returns, acting as a contract for a specialised function.

Key aspects:
- Allows FHIR implementations to add capabilities for tasks that do not fit neatly into standard REST operations (i.e. Create, Read, Update, Delete, Search).
- HL7 publishes standard operations and organisations can define their own custom operations.
- OperationDefinition specifies parameters (names, types, cardinality, whether required), response format and content.
- Machine-readable – Clients can use the OperationDefinition resource to understand how to call an operation without prior manual configuration.
- Invokes a special FHIR operation that retrieves all available information related to a specific patient, returning a complete set of associated records in a single response.

CompartmentDefinition

A logical grouping of resources associated with a specific entity, such as a patient, practitioner, or device, and the rules governing access to them.
It specifies which resources belong to the compartment and the conditions under which they can be accessed in relation to the compartment's owner. This structure helps servers organise and manage data access for security, privacy, and performance purposes.

Examples of compartments in FHIR:
- Patient – all resources related to one patient (Observations, Conditions, Encounters, etc.).
- Practitioner – all resources linked to a clinician.
- Encounter – all resources tied to a particular hospital visit.
- Device – all resources linked to a specific device.

MessageDefinition

A resource that defines the structure, content, and rules for a specific type of message used in FHIR messaging.

It acts like a blueprint that tells systems the following:

- What event the message is about (e.g., patient admission, lab result available).
- Which resources must be included in the message.
- How those resources should be organised.
- Any constraints or special rules for the message.

Example:
Used in workflows where asynchronous, event-driven communication is needed, such as notifying a GP that a patient's discharge summary is ready.

GraphDefinition

A resource that describes a set of related resources and how they are linked, beginning from a single focus resource. It serves as a map that specifies all the connected resources to include and the relationships between them, providing a clear guide for retrieving or presenting interconnected information in a consistent and predictable way.

Terminology Artifacts

Formal resources that define, manage, and share the codes, value sets, and mappings used to ensure consistent meaning in clinical and administrative data. They are a key part of semantic interoperability because they standardise the way terms are represented and understood across systems. Includes: ValueSet, CodeSystem, and ConceptMap for coded data.

FHIR Conformance Model: Implementation Guide (IG)

An Implementation Guide (IG) is a formal FHIR artefact that packages together one or more profiles along with supporting documentation, examples, and conformance requirements.

It functions as both a machine-readable and human-readable blueprint, guiding implementers on how to build, validate, and integrate systems using those profiles.

Key Features of Implementation Guides

Comprehensive Packaging

Implementation Guides bring together all key artefacts (profiles, documentation, examples, and conformance rules) into a single, structured resource.

Dual Readability

They provide both human-readable guidance for implementers and machine-readable formats for automated processing.

Defined Usage Rules

These guides specify how FHIR resources, API features, and terminologies should be applied in practice.

Context-Specific Application

Implementation Guides tailor FHIR for particular use cases, such as patient data exchange or clinical workflows.

Structured Content

Each guide includes narrative explanations, technical specifications, and real-world examples to support understanding and adoption.

Standardised Format

Following a consistent, tool-friendly structure ensures that Implementation Guides can be reused and processed automatically.

Support for Interoperability

They promote consistent and accurate data exchange across systems, organisations, and jurisdictions.

AU Base and AU Core Implementation Guides

In the previous Topic, you learned that AU Base and AU Core are Australian FHIR Implementation Guides that localise the international standard by defining profiles, extensions, and terminology.

AU Base provides the foundational adaptations while AU Core establishes the minimum national requirements for consistent interoperability.

Key aspects of AU Base and Core Implementation Guides

Aspect	AU Base Implementation Guide	AU Core Implementation Guide
Purpose	Provides foundational FHIR profiles and extensions adapted for Australian contexts.	Defines minimum data expectations and standardised profiles for national interoperability.
Scope	Covers generic localisation needs applicable across all Australian healthcare domains.	Builds on AU Base to ensure national consistency in common healthcare data exchange scenarios.
Content	Profiles for common resources (e.g., Patient, Practitioner, Organisation). Extensions for Australian identifiers (e.g., Medicare Number, IHI). Australian-specific address/date/time formats. Terminology bindings using Australian vocabularies (e.g., SNOMED CT-AU).	Profiles defining mandatory elements and terminology bindings. Guidance on search parameters and must-support elements. Examples showing compliant resource instances.
Use	Acts as a foundation layer for other IGs, not tied to one workflow.	Used for real-world interoperability projects such as My Health Record integration, secure messaging, and clinical data sharing.
Relationship	Base layer for all Australian FHIR implementations.	Builds on AU Base to deliver specific interoperability requirements.

Topic: FHIR Usage

The FHIR usage model describes how components (e.g. resources, profiles, and operations) are applied in real-world systems. The usage model provides guidance on interactions, packaging, workflows, and best practices to ensure consistent and interoperable implementations.

RESTful API

Application Programming Interfaces (APIs) are essential for the rules, definitions, and protocols that allow different software applications to communicate with each other.

HL7 FHIR is a modern specification ideally suited for API-based data exchange, supporting interoperability across diverse healthcare technologies and improving access to clinical data. In the context of FHIR, a RESTful (Representational State Transfer) API serves to implement these rules in a real system.

Characteristics of RESTful API

Resource-oriented design	Each piece of healthcare data (such as a patient record or observation) is represented as a FHIR resource, identified by a unique URI and accessed through REST endpoints.
Lightweight communication	REST avoids complex middleware, enabling faster and more efficient data exchange compared to older methods like Simple Object Access Protocol (SOAP).
Facilitates scalability and flexibility	REST APIs can be used by web apps, mobile apps, IoT devices, and other systems, making FHIR adaptable to many healthcare scenarios.

RESTful API Core Operations in FHIR

Operation	**Purpose**	**HTTP Command**	**Example**
Create	Adds a new resource to the server and returns its ID.	POST [base]/[type]?_format=[mime-type]	Create a new patient: POST https:// example.org/fhir/Patient
Read	Retrieves an existing resource by its ID.	GET [base]/[type]/[id]?_format=[mime-type]	Get patient ID 123: GET https:// example.org/fhir/Patient/123
Update	Modifies an existing resource and creates a new version.	PUT [base]/[type]/[id]?_format=[mime-type]	Update patient ID 123: PUT https:// example.org/fhir/Patient/123
Delete	Removes a resource (data remains on server for audit/history).	DELETE [base]/[type]/[id]	Delete patient ID 123: DELETE https:// example.org/fhir/Patient/123
Search	Finds resources that match search criteria.	GET [base]/[type]?[parameters]&_format=[mime-type]	Search patients born in 1980: GET https://example.org/fhir/Patient?birthdate=1980
History	Retrieves all versions of a resource or all resources of a type.	GET [base]/[type]/[id]/_history?[parameters]&_format=[mime-type]	View history of patient ID 123: GET https://example.org/fhir/Patient/123/_history

Bundles

Bundle types	Description	Examples/Uses
Document bundles	Structured clinical documents	Discharge summaries, referral letters
Message bundles	Event-driven communication between systems	Lab result notifications
Transaction and batch bundles	Multiple operations sent in a single request, with response bundles mapping results back to each original entry.	Creating a patient and their associated clinical records at the same time Updating several related resources that must always be in sync Migrating linked data from one system to another without partial completion
Collection bundles	General-purpose collections of resources without transactional rules	Exporting or importing datasets for research or auditing Packaging records for bulk archival Grouping unrelated resources for distribution to another system in a non-transactional way
Search/history bundles	Returned as the result of a search query or containing historical versions of a resource	Retrieving lists of patients, encounters, or observations based on filters Paginating through large result sets Returning search matches in a standard, structured format

FHIR services

FHIR services go beyond simply storing and retrieving healthcare data. They embody the functional capabilities that power secure, interoperable, and workflow-driven integration across health systems.

This table below explains how FHIR Services extend the capabilities of basic FHIR to support real-world healthcare operations and enterprise-level interoperability. It demonstrates the evolution from a data-exchange model to a service-oriented model, showing how FHIR Services enable more robust, secure, and interoperable healthcare ecosystems.

Aspect	Basic FHIR	FHIR Services
Primary Focus	Data storage and retrieval CRUD operations on resources	Functional capabilities Business processes and workflows
Operations	Standard REST only GET, POST, PUT, DELETE Simple resource interactions	Extended operations $validate, $everything, $expand Custom business operations
Architecture	Single pattern RESTful API only Point-to-point integration	Multiple patterns REST, Messaging, Documents Event-driven, async workflows
Security	Basic authentication Simple access control Limited audit capabilities	Enterprise security OAuth2, SMART on FHIR, RBAC Full audit, consent management
Transactions	Single resource One resource at a time No rollback support	Multi-resource Coordinated updates Transactional integrity
Business Logic	Minimal Data-centric operations No workflow orchestration	Business-aligned Register Patient, Validate Identity SOA principles, service contracts
Scalability	Limited Vertical scaling Monolithic deployment	Highly scalable Modular, microservices Distributed, reusable components
Error Handling	Basic Standard HTTP codes Simple error messages	Sophisticated Agreed patterns across vendors Detailed Operation Outcome

Topic: FHIR Terminology

The table below shows Building Blocks and Terminology Services supported by FHIR.

FHIR Resources (building blocks)	FHIR Terminology Services
CodeSystem	$expand
ValueSet	$validate-code
ConceptMap	$translate

Different terminologies and code systems serve different purposes within clinical care, research, administration, and public health.

The following table shows commonly used code systems.

Code systems	Description
ICD 10 Revision 10 of the World Health Organisation International Classification of Diseases (ICD)	Used worldwide to code and classify diseases, conditions, and causes of death.
SNOMED CT Systematised Nomenclature of Medicine – Clinical Terms	A comprehensive, standardised clinical terminology used worldwide to record and share healthcare information with consistent meaning
LOINC Logical Observation Identifiers Names and Codes	Standard for identifying laboratory tests, clinical measurements, and observations.
ATC Anatomical Therapeutic Chemical Classification System	Maintained by the WHO, used for classifying medicines by therapeutic use and chemical properties.
AMT Australian Medicines Terminology	Specific to Australia, standardises naming and identification of medicines.
UCUM Unified Code for Units of Measure	Unified Code for Units of Measure UCUM is used in Australia, especially in pathology settings when no local unit is defined, ensuring consistent and unambiguous representation of measurement units across clinical systems.

Terminology Service Operations

Operation	Purpose
$validate-code	Checks if a code is valid within a specific code system or value set
$expand	Returns all codes in a value set (e.g., all SNOMED CT codes for body sites)
$lookup	Retrieves details about a specific code (e.g., display name, definition)
$translate	Converts a code from one code system to another (e.g., SNOMED CT → ICD-10)
$subsumes	Determines if one code is a broader/narrower concept than another

Topic: Integration Engines and Facades

The following are façade examples in healthcare systems.

FHIR API facade for legacy systems

A façade exposes a modern FHIR API while internally interacting with a legacy HL7 v2 system. It translates incoming FHIR requests into HL7 messages and vice versa, enabling modern applications to communicate with older infrastructure.

Security enforcement façade

A facade layer that intercepts API requests to enforce authentication, authorisation, and data masking before passing them to the backend systems. This helps centralise and standardise security controls.

Data transformation façade

A façade that transforms incoming data formats (e.g., XML, CSV, proprietary formats) into FHIR-compliant resources, allowing diverse systems to integrate with a FHIR-based platform without needing to natively support FHIR.

Consent management façade

A facade that checks patient consent records before allowing access to sensitive health data, ensuring compliance with privacy regulations and organisational policies.

Workflow orchestration façade

A facade that uses the FHIR Task resource to coordinate multi-step workflows across systems, such as lab ordering, results delivery, and clinical review, abstracting the complexity of underlying service interactions.

Integration engine or façade?

Both the FHIR façade and FHIR repository approaches act as FHIR-enabled server patterns, but they differ fundamentally in how and where data is stored, transformed, and accessed.

The below presents a comparison of features/aspects to assist architects in selecting the most suitable approach.

Feature/Aspect	FHIR façade	FHIR repository
Where data is stored	Remains in the original source systems	Copied into a dedicated FHIR Clinical Data Repository
Data conversion to FHIR	Done on-the-fly at query time	Done on ingestion when data is loaded into the repository
Query source	Queries run against the translation layer connected to the source systems	Queries run directly against the repository
Latency	Real-time mapping may add delay	Faster response for repeated queries (data is already in FHIR)
Data freshness	Always current (pulled directly from the source)	May lag depending on data load/update frequency
Implementation complexity	Requires robust mapping/transformation logic for real-time queries	Requires building and maintaining a data ingestion pipeline
Infrastructure load	Load occurs on source systems during queries	Load occurs on repository during queries; ingestion load is separate
Best suited for	When real-time, up-to-date data is critical and duplicating data is not desired	When fast query responses, historical data retention, and reduced load on source systems are priorities

Topic: Case Study: PROMS Integration Pilot using SMART on FHIR

Pilot Phases

Phase 1
A commercial app integrated into the EHR.

Phase 2
A newly developed app, using lessons from Phase 1.

The following graphic shows the delivery timeline for each phase.

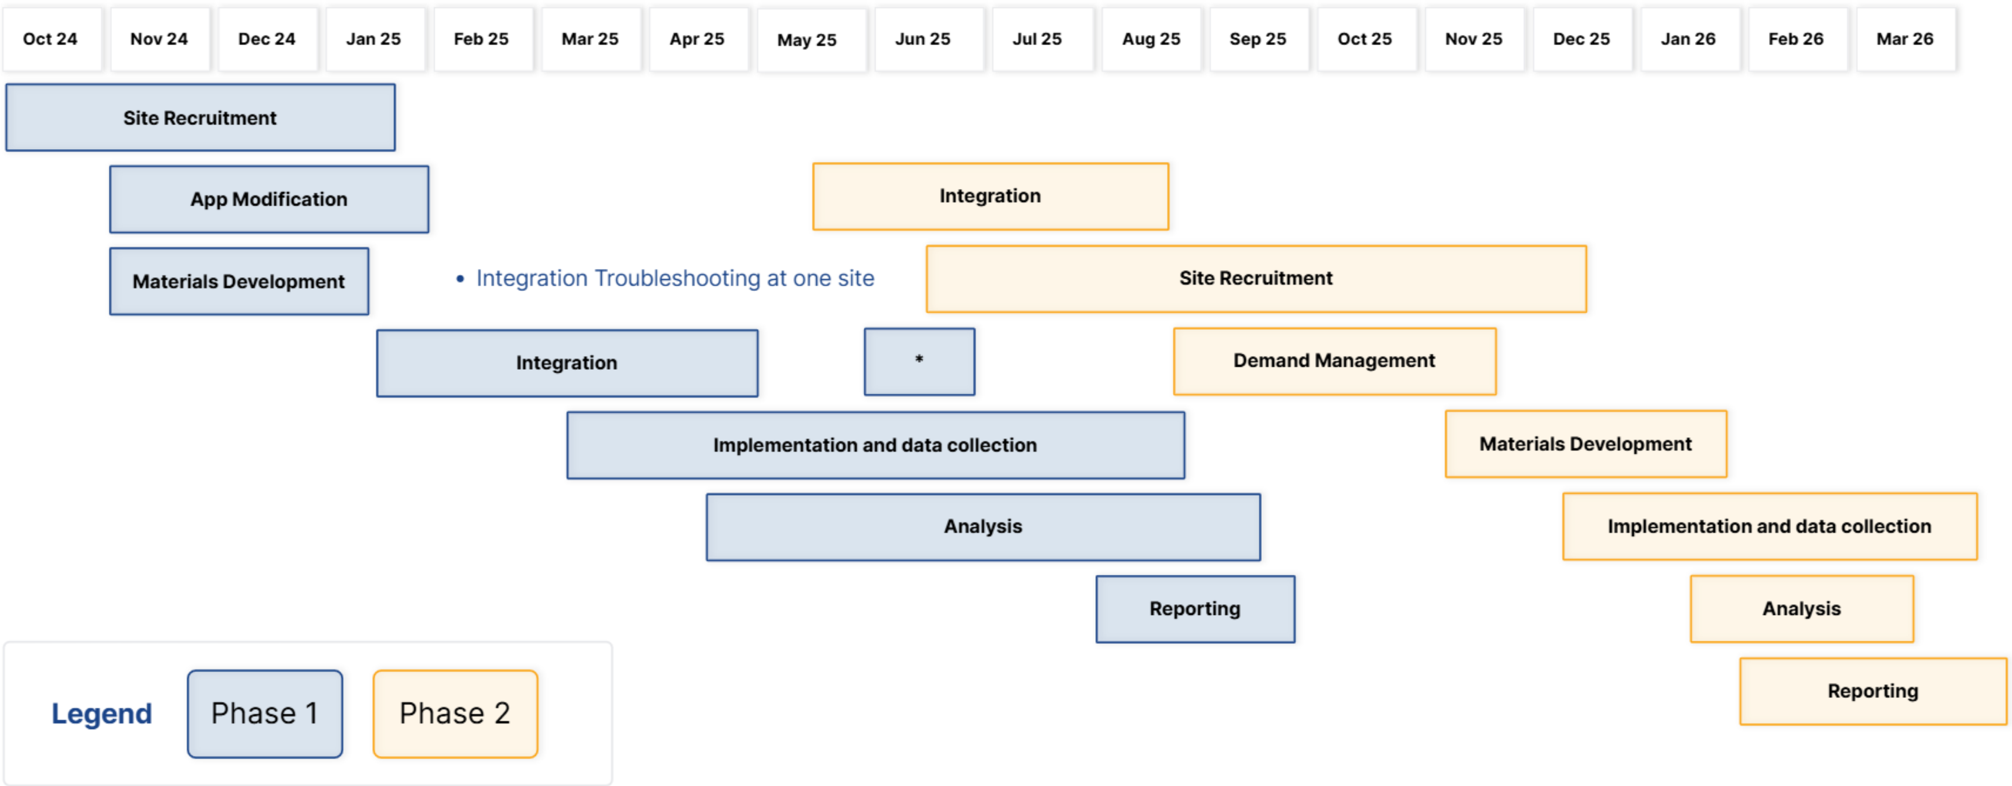
Oct 24
Nov 24
Dec 24
Jan 25
Feb 25
Mar 25
Apr 25
May 25
Jun 25
Jul 25
Aug 25
Sep 25
Oct 25
Nov 25
Dec 25
Jan 26
Feb 26
Mar 26
Site Recruitment
App Modification
Materials Development
• Integration Troubleshooting at one site
Integration
*
Implementation and data collection
Analysis
Reporting
Integration
Site Recruitment
Demand Management
Materials Development
Implementation and data collection
Analysis
Reporting
Legend
Phase 1
Phase 2