CYBAGE
Delivering Value. Scientifically.

WHITEPAPER

FHIR Interoperability in Healthcare

Architecture, Legacy Coexistence, and Cross-Standard Interoperability

Table of contents

1. Introduction

Healthcare interoperability is entering a new phase. Fast Healthcare Interoperability Resources (FHIR) provides a standardized, API-based approach to exchanging healthcare information, while U.S. regulations are accelerating its adoption. CMS rules now extend FHIR-based exchange across patient access, provider access, payer-to-payer exchange, and prior authorization; ONC certification requires standardized FHIR APIs for certified health IT; and TEFCA is extending interoperability toward nationwide health information exchange.

As a result, FHIR is becoming part of broader healthcare interoperability strategies rather than remaining limited to individual applications or modernization initiatives. Payers, providers, health technology vendors, and intermediaries increasingly need to expose, consume, and exchange information through FHIR while continuing to support existing regulatory and business workflows.

However, healthcare is not a FHIR-native ecosystem. HL7 v2 remains deeply embedded in clinical messaging, X12 in administrative transactions, and CDA in clinical documents, while substantial data remains in relational databases, proprietary applications, and established integration platforms. FHIR adoption does not replace these systems; it must coexist with them.

The resulting architectural challenge is one of coexistence and transformation: introducing FHIR while enabling information to move between new and existing systems without losing meaning or integrity. Depending on the environment, this may involve exposing FHIR over legacy systems, transforming data into a FHIR repository, bridging established message flows with FHIR interfaces, or combining these approaches in a hybrid architecture.

This paper examines how FHIR can coexist with established healthcare standards, the architectures and exchange patterns that enable this coexistence, the challenges of translating information across different data models, and the controls required to make FHIR data reliable for downstream applications, analytics, and AI-enabled healthcare solutions.

2. FHIR Conformance to Interoperability

FHIR provides a common framework for representing and exchanging healthcare information, but implementing FHIR does not guarantee interoperability. A system can generate valid FHIR resources, expose them through a FHIR API, and conform to a required profile or Implementation Guide while still producing information that another system cannot reliably interpret.

This is partly due to FHIR's design. FHIR uses modular resources such as Patient, Encounter, Observation, Coverage, Claim, and Organization, while its base specification remains intentionally broad to support different systems, jurisdictions, and use cases. Profiles and Implementation Guides add context-specific constraints, including cardinalities, terminology bindings, extensions, supported elements, and operations.

Interoperability therefore requires alignment beyond base FHIR conformance. Trading partners must agree on the applicable profiles and Implementation Guide versions, terminology and identifier systems, and the meaning of the information exchanged. A resource may be technically valid but still contain an incorrect identifier namespace, an inappropriate terminology code, or information that was misinterpreted during transformation.

This is especially important when introducing FHIR over existing healthcare data. HL7 v2, X12, CDA, and legacy databases represent information differently from FHIR. Transformation is therefore more than syntax conversion: source concepts, relationships, identifiers, codes, and business rules must be mapped to the appropriate FHIR representation while preserving meaning and meeting the constraints of the target profile or Implementation Guide.

3. The FHIR Foundation Layer

FHIR provides a common structure for representing healthcare information while deliberately leaving room for that structure to be constrained for different jurisdictions, workflows, and contexts. The components below are the foundational structures every implementation builds on, and these components group into four tiers: structure, meaning, conformance, and exchange.

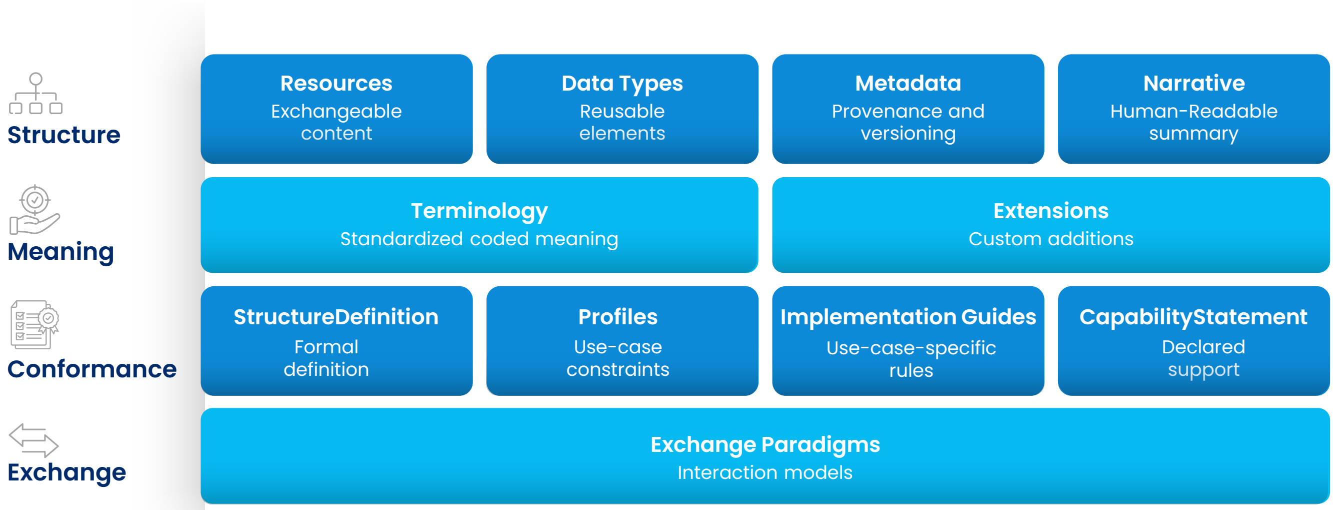

Resources

The basic building block of FHIR. All exchangeable content is defined as a resource: Patient, Encounter, Observation, Condition, Coverage, Claim, and around [145 types in R4 across](#) clinical, administrative, financial, and infrastructure domains. Each resource is independently identifiable and can reference others, making FHIR data a connected graph rather than a document. The base specification deliberately defines only the elements broadly applicable across implementations, the 80/20 philosophy, leaving use-case-specific requirements to profiles and extensions. A base R4 Patient does not require a name, identifier, or birth date. This enables global adoption, but it also means base FHIR alone does not guarantee interoperability.

Data Types

Resources are assembled from reusable data types. Primitive types hold single values such as string, boolean, and dateTime. Complex types group related elements into recurring structures such as HumanName, Identifier, CodeableConcept, and Period. A constraint applied to a data type carries the same meaning wherever that type appears.

Metadata

Every resource carries metadata: a version identifier, a last-updated timestamp, and the profiles it claims to conform to.

Narrative

Most resources carry a human-readable XHTML rendering of their content, so a clinician can read the resource even where the receiving system cannot process every element.

Extensions

Extensions represent information beyond the base resource. Each is defined by a canonical URL so its meaning stays consistent across implementations, supporting local requirements without changing the base specification.

Terminology

A coded element is only interpretable if both systems agree on the codes. A CodeSystem defines codes and their meanings, such as LOINC, SNOMED CT, or ICD-10. A ValueSet selects codes from one or more code systems for a specific context. A ConceptMap defines relationships between codes in different systems, which makes legacy code translation expressible as a FHIR artifact rather than application logic.

StructureDefinition

Every resource, data type, extension, and profile is formally defined through a StructureDefinition: what it contains, cardinality, data types, terminology bindings, and constraints. For profiles, it carries two views: the differential shows only what a profile changes; the snapshot shows the full structure after changes are applied, including inherited constraints. StructureDefinitions also define invariants, rules a resource must satisfy to be valid even when structure and types are correct.

Profiles

A profile constrains a resource for a specific use case or jurisdiction: required elements, restricted types, terminology bindings, fixed values, and Must Support obligations. Two systems can conform to the same base resource while implementing different profiles, which is why profile-level agreement, not base conformance, is what interoperability rests on.

Implementation Guides

An Implementation Guide assembles profiles, extensions, terminology, and exchange requirements into a single specification an implementation must follow. Examples include US Core, Da Vinci PAS, and CARIN Blue Button. The meaningful conformance target is a specific guide and version, not "FHIR compliant."

CapabilityStatement

A server declares what it actually supports through a CapabilityStatement: the resources it exposes, the interactions and search parameters it accepts, and the profiles it claims. This allows a client to verify conformance at runtime rather than relying on documentation.

Exchange

FHIR defines how systems move this data through a REST API and three further paradigms. These are covered in the Exchange Paradigms section.

4. FHIR Exchange Paradigms

FHIR building blocks set out how healthcare data is structured and how conformance narrows it down for real use. But interoperability also depends on how that data moves between systems. Clinical workflows, administrative processes, and integrations all have different communication needs, so one interaction pattern cannot cover them all. FHIR handles this with four exchange paradigms. Each suits a different scenario, and all four use the same underlying information model.

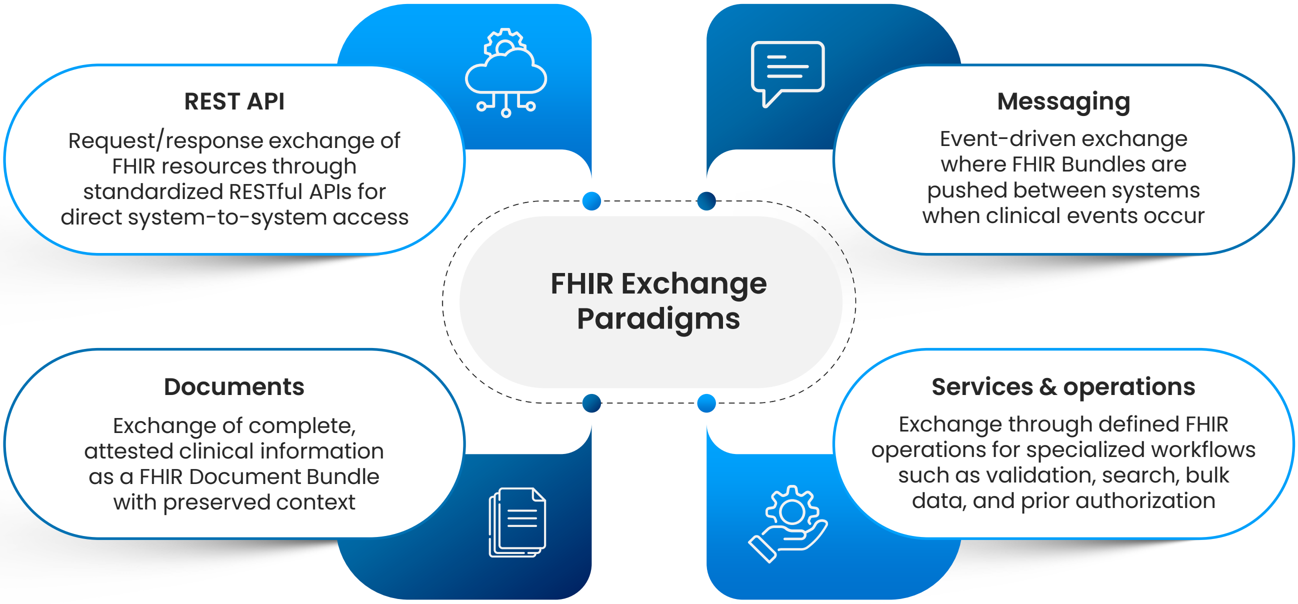

The four are meant to work together. A single implementation can expose a REST API for direct system-to-system access, push Bundles as messages when clinical events happen, exchange Document Bundles across organizational boundaries, and invoke operations for validation, bulk data, and prior authorization. Which one you use depends on the workflow, not on technical preference.

5. The FHIR Architectural Patterns

Choosing a paradigm settles how systems talk to each other. It does not settle where FHIR data sits or when it gets converted, and that question is unavoidable because FHIR arrives at an estate that already works. HL7 v2 still carries clinical messaging, where existing interfaces hold site-specific customization, and there is no business driver to replace a working feed. EDI X12 still carries administrative transactions mandated under HIPAA. CDA documents remain in use for exchanging and preserving clinical information. FHIR is introduced alongside these standards, not in place of them, so translation and coexistence are architectural concerns rather than migration steps.

Four patterns are used in practice. They differ on two things: whether the source system or a dedicated FHIR server holds the data, and whether conversion happens when data is written or when it is requested.

FHIR Facade

Data stays in the legacy system. A FHIR API layer maps it on request, so no migration is needed.

Event Driven Pipeline

When source data changes, the FHIR server updates automatically. Applications are notified rather than repeatedly checking.

FHIR Repository

Source data is converted to FHIR during ingestion and stored in a dedicated FHIR server, which applications then query directly.

Hybrid

Some resource types are stored in the FHIR server. Others stay in the legacy system, converted only when requested.

6. Use Case: Prior Authorization at the X12 and FHIR Seam

Prior authorization is a clear example of X12 and FHIR coexisting within the same workflow. CMS-0057-F requires impacted payers to offer a FHIR-based Prior Authorization API using the Da Vinci PAS Implementation Guide, while HIPAA continues to use the X12 278 for applicable prior authorization transactions. PAS therefore provides the FHIR-based provider interaction, with translation to X12 278 where required. Delegated utilization management platforms may participate across this workflow by consuming PAS Bundles, translating them into their internal authorization model, and supporting X12 278 transactions across multiple health plans.
A full regulatory analysis of CMS-0057-F, covering scope, timelines, API requirements, and enforcement posture, is available in [Cybage's dedicated blog on the Final Rule.]

7. The Challenge Landscape

While FHIR establishes a standardized foundation for healthcare interoperability, successful production implementations depend on far more than conformant APIs and valid resources. As organizations scale FHIR across clinical, administrative, and analytical workloads, they encounter architectural challenges related to version management, semantic consistency, operational performance, and governance. Addressing these challenges is essential to achieving reliable interoperability in production environments.

Standards Version Fragmentation

FHIR evolves across releases, and different versions can remain in production simultaneously. While normative content provides backward-compatibility guarantees, changes across releases, particularly in trial-use content, can still affect elements, structures, types, and bindings. Supporting multiple versions therefore requires careful version assessment and, where differences affect data exchange, explicit mapping and transformation.

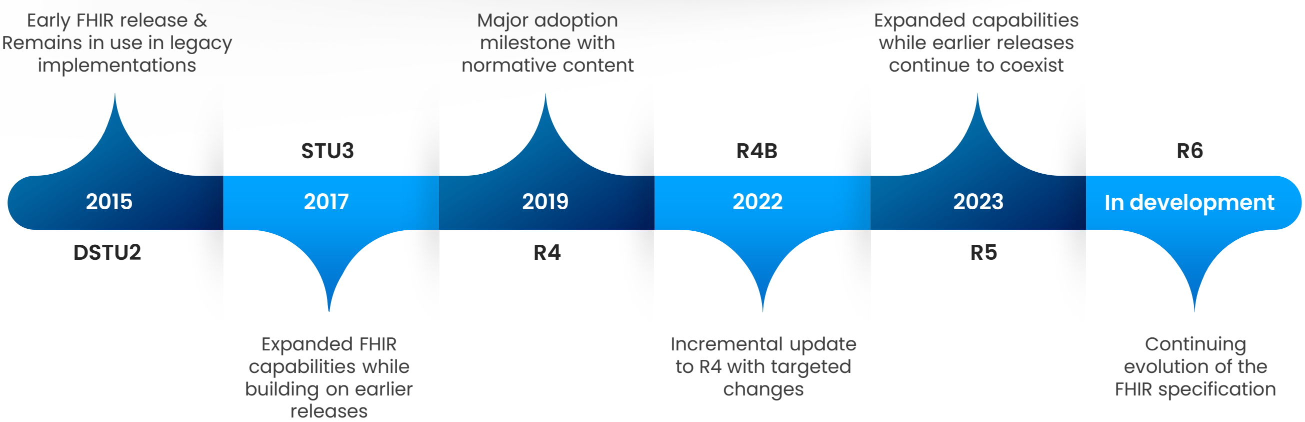

Implementation Guide Version Drift

Implementation guides also evolve through multiple versions. US Core illustrates this: successive releases tighten profiles, add Must Support requirements, and incorporate new USCDI data elements, while certification and regulatory programs may remain tied to earlier versions. This means partners using different versions can have different requirements and terminology bindings, creating interoperability issues.
The same applies across IG families. Da Vinci releases depend on specific US Core versions, while CARIN Blue Button follows its own release cycle, creating different dependencies for organizations implementing multiple guides. Version registers and reconciliation gates help ensure changes are deliberate and tested rather than silent.

Terminology Management

FHIR relies on external terminologies such as **SNOMED CT, LOINC, RxNorm, ICD, and CPT,** each with its own release cycle and licensing requirements. As concepts and value sets evolve, different systems can validate the same data against different terminology versions, creating silent semantic divergence. Version-pinning terminology, recording code-system versions where supported, and treating terminology updates as managed releases with regression testing help maintain consistency. Versions, disagreeing about whether a specific code is valid. The control is to version-pin terminology the way IGs are pinned, record code system versions inside resources where the data type allows it, and treat terminology updates as managed releases with regression testing.

Legacy Data Transformation and Mapping

FHIR implementations rarely originate from native FHIR data. Most organizations introduce FHIR by transforming information from HL7 v2, CDA, X12, proprietary databases, or canonical enterprise models. Maintaining semantic fidelity across these transformations remains one of the most complex architectural challenges. Mapping is seldom a one-time activity; source systems evolve, implementation guides change, and business rules expand over time. Sustainable interoperability therefore depends on governed mapping specifications, transformation testing, and traceability rather than individual interface implementations.

Governance and Change Management

FHIR implementations evolve continuously as regulations, implementation guides, terminology systems, and business workflows change. Without governance, organizations accumulate inconsistent profiles, duplicated extensions, incompatible mappings, and fragmented APIs that undermine interoperability. Effective governance establishes ownership of profiles, extensions, terminology assets, mapping specifications, and implementation guide adoption, ensuring changes remain coordinated across the enterprise rather than occurring independently within individual projects.

8. The Road Ahead

FHIR adoption is no longer driven solely by interoperability initiatives. Regulatory mandates, evolving implementation guides, cloud-native healthcare platforms, and AI-driven applications are expanding the role of FHIR beyond data exchange into a foundational component of modern healthcare architecture. Understanding these trends provides insight into how interoperability strategies are likely to evolve over the coming years.

Regulatory Evolution

Healthcare interoperability regulation is steadily expanding from data access toward end-to-end digital workflows. Initial regulations focused on exposing standardized FHIR APIs for patient and provider access. More recent mandates extend those requirements into operational processes such as prior authorization, payer-to-payer exchange, and provider access, reflecting a broader shift from information availability to workflow interoperability.

For solution architects, the implication is clear. FHIR should no longer be viewed as an isolated API strategy but as an integral part of regulated healthcare workflows. At the same time, existing standards such as X12 continue to remain mandatory for many administrative transactions, reinforcing that future interoperability architectures will rely on coexistence rather than replacement.

Evolution of the FHIR Specification

FHIR continues to evolve through successive releases, but adoption takes time. R4 remains widely used across U.S. healthcare, while R5 introduces new capabilities that can be adopted selectively, including through backport guides. R6 represents the next stage of the specification's evolution, with greater normative maturity expected to support broader adoption over time.

Moving between major FHIR releases is therefore not simply a version upgrade. It requires coordinated adoption, version assessment, and mapping across implementation guides, systems, and trading partners. Organizations that maintain strong transformation and mapping practices will be better positioned for future FHIR version transitions.

Global Adoption of FHIR

While regulatory adoption has accelerated most rapidly in the United States, FHIR has become the preferred interoperability standard across healthcare systems worldwide. National initiatives in Europe through the European Health Data Space (EHDS), NHS England, Australia, Canada, and several Asia-Pacific healthcare programs increasingly position FHIR as the foundation for modern digital health ecosystems.

Although implementation approaches differ across countries, the architectural direction remains consistent. FHIR is becoming the common interoperability layer connecting clinical applications, national health exchanges, patient-facing services, and digital health platforms while continuing to coexist with established national standards where operational or regulatory requirements persist.

FHIR as the Foundation for Analytics and AI

FHIR is increasingly becoming an important foundation for healthcare analytics and AI by providing a standardized way to represent and exchange clinical data across systems. This gives AI applications more consistent access to patient information from different sources, supporting use cases such as clinical decision support, predictive analytics, and AI-assisted workflows. As AI adoption grows, FHIR can also provide a standardized structure for **AI-generated healthcare data.** Clinical documentation, coding, mapping, and ambient AI applications can produce FHIR resources that are validated, linked to standard terminologies, and tracked through provenance, making AI-generated information easier to integrate into healthcare workflows.

9. Conclusion

FHIR is not replacing the standards that came before it; it is becoming part of a broader interoperability architecture where multiple standards, systems, and versions coexist. The long-term value of FHIR lies in how effectively organizations can connect these environments while preserving the meaning, quality, and integrity of the data being exchanged. Successful FHIR adoption therefore depends not only on implementing FHIR, but on the architecture, mappings, governance, and interoperability decisions that make it work alongside the systems already in place.

10. How Cybage Helps

Cybage helps healthcare organizations advance interoperability through FHIR. Our capabilities span FHIR-native solutions, integration of FHIR with established standards such as HL7, X12, and CDA, EDI-to-FHIR transformation, legacy data transformation, implementation guide and terminology mapping, and integration with existing healthcare systems, helping organizations adopt FHIR while building on their existing interoperability investments.

Ready to build a FHIR interoperability strategy that works with your existing healthcare ecosystem?

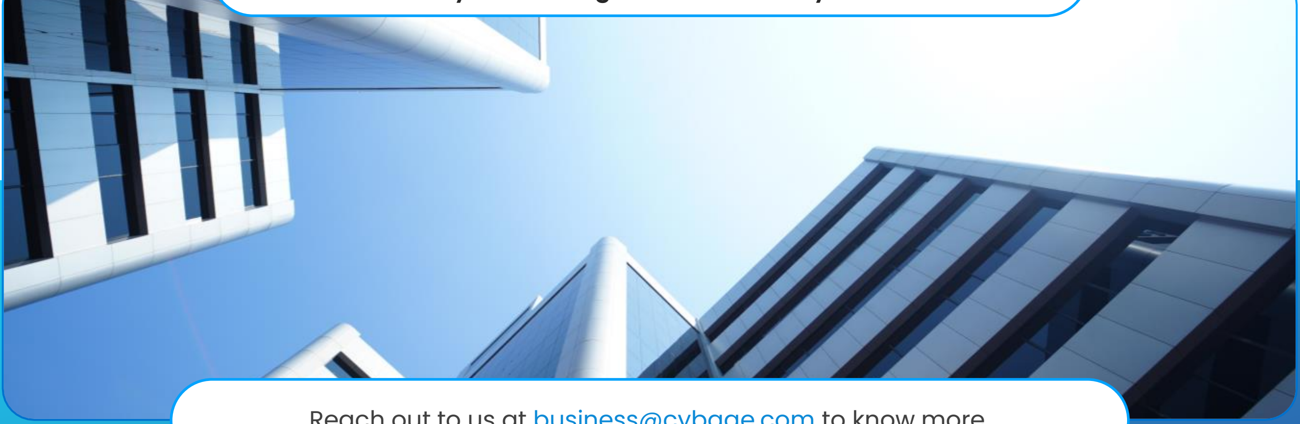

Reach out to us at business@cybage.com to know more.