FHIR

Shaping the Future of Health Data Exchange

What Is FHIR®?

Fast Interoperability Healthcare Resources, or **FHIR**, is a standard for exchanging healthcare data that has been steadily gaining traction since its initial draft release in 2014. Two recent advancements will drastically increase the adoption of FHIR.
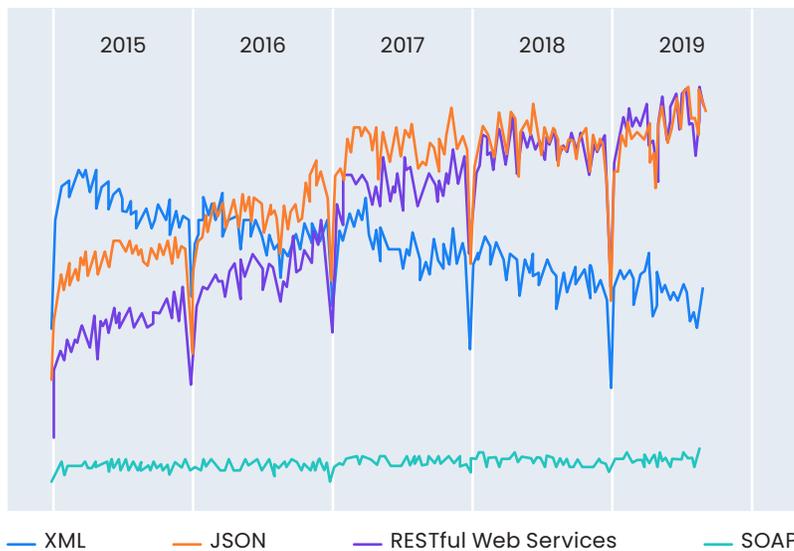
In 2019, the international standards body HL7 International released **FHIR R4**, the first normative version — meaning future changes to the standard will be backward compatible. Then in March 2020, two divisions under the U.S. Department of Health and Human Services — the Centers for Medicare and Medicaid Services (CMS), and the Office of the National Coordinator for Health IT (ONC) — published two rules requiring the use of FHIR APIs for data exchange between providers, payers, and patients.

But what is FHIR? Why was it created, and what are the implications for data exchange in the healthcare industry?

HL7 FHIR® began with an open-ended question: What would health data exchange look like if we started from scratch using modern approaches? To answer this question, HL7 turned to other industries for ideas. Recent interoperability successes pointed strongly to the use of RESTful based APIs.

HL7 FHIR combines the best features of HL7 V2, HL7 V3, and CDA, while leveraging the latest web service technologies. FHIR is based on modular components called "resources," and these resources can be combined to solve clinical and administrative problems in a practical way.

The FHIR standard is based on the following simple five key points:

1. Faster to learn and implement

2. Lower cost

3. Scales well from simple to complex

4. Flexible

5. Free

The impact of many of these key points will be seen throughout the following sections.

Figure 1: **Technology Trends in Data Exchange**

XML —— JSON —— RESTful Web Services —— SOAP

Comparison to Existing Standards

HL7 V2

HL7 V2 is a well-established standard that works well within institutions to connect applications. However, it is a legacy standard with unique syntaxes, custom tools, and a hefty learning curve for those entering the healthcare IT industry. The design of the standard is also limiting to modern devices and apps that are trying to leverage available patient data. Privacy and security are also difficult to implement. These limitations have become a barrier to patient engagement and making patient data available in the most convenient formats.

Table 1: HL7 V2 AND FHIR

Similarities	Differences
Built around re-usable "chunks" of data	Modern tools/skills
Strong forward/backward compatiblity rules	Lighter-weight specification
Extensiblity Mechanism	Each chunk (resources) is independently addressable
	Human readability is required
	Extensions are discoverable
	Instances easy to read

Comparison to Existing Standards

HL7 V3

HL7 V3 was meant to be the successor of HL7 V2. It leveraged modern standards technologies available at the time while being based on a reference model but ended up being overly complex to implement with a steep learning curve. It also had no backwards compatibility with HL7 V2, which made switching to the new standard all the more complex given how embedded HL7 V2 is within the U.S. healthcare system.

Table 2: **HL7 V3 AND FHIR**

Similarities	Differences
Based on HL7 RIM, vocabulary, and ISO data t	Similar models and syntax (reference model hidden)
Supports XML syntax	Friendly naming
	Extensibility with discovery
	Easy inter-version wire compatibility
	JSON syntax supported

Comparison to Existing Standards

HL7 CDA

HL7 CDA was incorporated into the Meaningful Use criteria, and thus has broad penetration in the United States. However, there remains many drawbacks with CDA that are being exposed by the requirements of Meaningful Use:

○ Due to the same complexity as HL7 V3, the learning curve is still steep

○ Interoperability beyond a human-to-human level is still challenging

○ CDA documents do not fit well in many workflows

○ Extensibility is difficult

HL7 is developing a FHIR Mapping Language for Consolidated CDA (C-CDA) and HL7. This effort aims to provide automated transformation between C-CDA and FHIR.

Table 3: **HL7 CDA AND FHIR**

Similarities	Differences
Support profiling for specific use-cases	Can use out of the box no templates required
Human readability is required	Not restricted to just documents
Tooling available	Implementer tooling generated with spec
Includes same similarities as with HL7 V3	Includes same differences as with HL7 V3

HL7 FHIR—Differentiators

HL7 FHIR differentiates itself from past standards technologies by focusing on the following:

- ⊘ Focus on those who are implementing the standard
- ⊘ Ensure common scenarios are easily supported
- ⊘ Leverage cross-industry web technologies
- ⊘ Require human readability as base level of interoperability, same as CDA
- ⊘ Make the standard freely available
- ⊘ Deliver flexibility by supporting multiple paradigms and architectures
- ⊘ Provide for modern governance of data

Focus on the Implementers

HL7 FHIR was written with implementers as the target audience in mind. Of high importance was the use of reference implementations, which were incorporated into the design of FHIR from the beginning. Robust testing was deemed important as well for implementers, which led to the availability of public test servers.

Starter APIs have been made available for a variety of programming methods including C#, Java, JavaScript, Objective C, and Delphi. And to allow implementers to test with others along the way, Connect-a-thons were established very early on to verify specification approaches and continue at HL7 workgroup meetings.

Support for Commonly Used Workflows

The content in the core specification is meant to cover the top 80% of use cases. This is similar to the focus of the HL7 V2 standard, only the extensions will be cleaner and easier. The idea is to focus on the real needs of the industry but allow for all the special use cases as well. HL7 V3 was commonly criticized for trying to account for too many use cases in the core standard, thus leading to an overly complex standard that was difficult to achieve.

Figure 2: **HL7 FHIR Timeline**

First Draft		First DSTU*	Second DSTU*		Third STU#	First Norm+ R4		~Second Norm+
2012	**2013**	**2014**	**2015**	**2016**	**2017**	**2018**	**2019**	**2020**
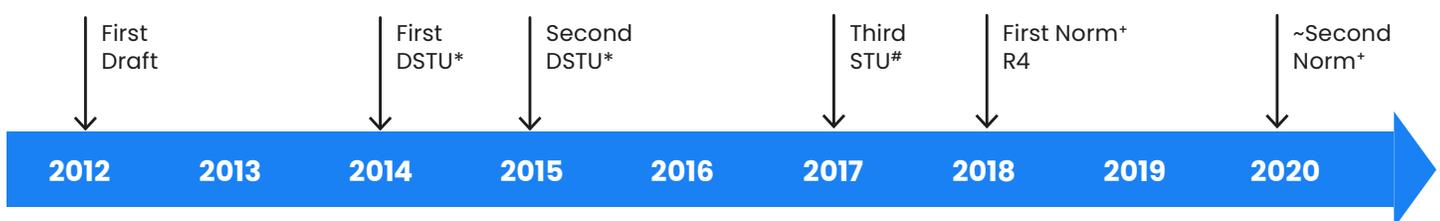
* Draft Standard for Trial Use
# Standard for Trial Use
+ Normative Edition

Information courtesy of HL7 international.

HL7 FHIR—Differentiators

Embrace Web Technologies

The use of RESTful web services as an API has been on the rise over the last decade across all industries. RESTful web services is embraced by organizations such as Facebook, Twitter, and Amazon as their primary API. In addition, related technologies such as XML, JSON, and Oauth, are also common when dealing with encoding and authorization. These technologies have well supported tools and a large talent pool of IT resources. Thus, healthcare will not be locked in to unique industry standards, but it can embrace what is used across all industries.

Human Readability

The concept of human readability being the minimum level of interoperability was introduced with the CDA standard. The idea is that if none of the structured data is able to be imported into the receiving system, the data could be viewed in a standard web browser. HL7 FHIR continues with this concept to ensure that human readability will always be an option.

Paradigms for Packaging the Payload

HL7 FHIR plans to support four interoperability paradigms, which are distinct ways of utilizing FHIR to best accommodate varying workflows. The four paradigms and when they might be used are:

- REST Small, light-weight exchanges with low coupling between systems

- Messages Communicate multiple resources in a single exchange

- Documents Focus is on persistence when data spans multiple resources

- Services Use a custom service when capabilities of other paradigms don't fit requirement

Regardless of the paradigm, the content is still based on resources. The resources are just bundled into different data sets depending on the paradigm. The diagram below details some of the key differences among the paradigms.

Table 4: HL7 CDA AND FHIR

REST	Documents	Messages	Services
Simple, out-of-the-box interoperability	Similar to CDA A bundle of resources	Similar to HL7 V2 and V3 messaging	Custom resources packaging and intelligence
Leverage HTTP: GET, POST, etc	A bundle of resources	A bundle of resources	Individual resources or bundles
Pre-defined operations	Root is a "Composition" resource	Allows request/response behaviour	Untra-complex or ultrasimple workflows
Create, Read, Update, Delete	Acts like a CDA header	Event-driven (e.g., send lab order, receive result)	Use HTTP or other protocols
Works best in environments where control resides on client side and trust relationship exists	One context	Can be asynchronous.	Only constraints is that FHIR resources are passed from one system to another
	Can be signed, authenticated, etc.		

HL7 FHIR—Differentiators

Resources Defined

Resources are small, logically discrete units of exchange. Resources define behavior and meaning, have a known identity and location, are the smallest possible unit of transaction, and provide meaningful data that is of interest to healthcare. The plan is to limit resources to about 150 in total. They are sometimes compared to an HL7 V2 segment. The resources can be extended and adapted to provide a more manageable solution to the healthcare demand for optionality and customization.

Modern Governance

Governance to maintain the privacy and security of patient data will be of utmost importance to the success of FHIR. FHIR includes modern capabilities that should allow policies to be constructed in a way that are protective but not constraining. All exchange of production data is expected to be secured using TLS/SSL. Authentication can be achieved in a number of ways, but OAuth is recommended for web-centric use. FHIR defines a Security Label infrastructure to support access control management, and FHIR also defines provenance and security event resources suitable for tracking the origins, authorship, history, status and access of resources.

Table 5: **Resource Examples**

Examples	Non-examples
Administrative • Patient • Organization • Location	Gender • Too Small Electronic
Clinical • Allergy • Family History • Care Plan	Health Record • Too Big
Infrastructure • Document Composition • Message Header • Capability Statement	Blood Pressure • Too Specific

Figure 3: **FHIR Emerging Standards**

XML	REST	JSON
Encoding Data	Communication	Encoding Data
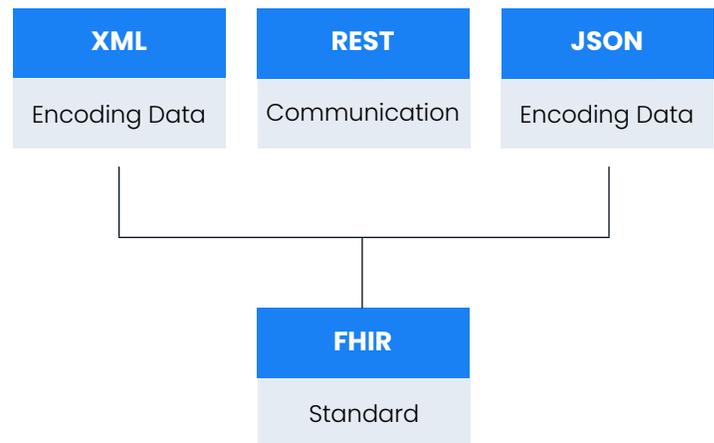
FHIR
Standard

Where Might HL7 FHIR be Used?

HL7 FHIR has the capability to be used in a variety of workflows, from small mobile devices to large multi-facility hospital information systems. FHIR not only enables new workflows, such as those related to patient engagement, but also more traditional communications between applications.

Common workflows include:

○ Traditional application-to-application interoperability within the four walls

○ External connectivity

- HIEs / ACOs
- National exchanges
- Social web
- Mobile applications
- Home health devices

The timeline diagram in Figure 4 depicts potential FHIR workflows:

○ Initially, there will likely be a need for broker applications, such as an interface engine, to act as middleware between applications that still use traditional standards and those with the new FHIR-based standard

○ Patient engagement will be a major target for FHIR. Apps will be able to leverage the lightweight.

○ REST standard to provide patients timely data and alerts.

○ FHIR will also be utilized as a link into patient data repositories, where applications all throughout the provider facility can gain access to patient data only when the patient shows up in their department.

Figure 4: FHIR Workflows
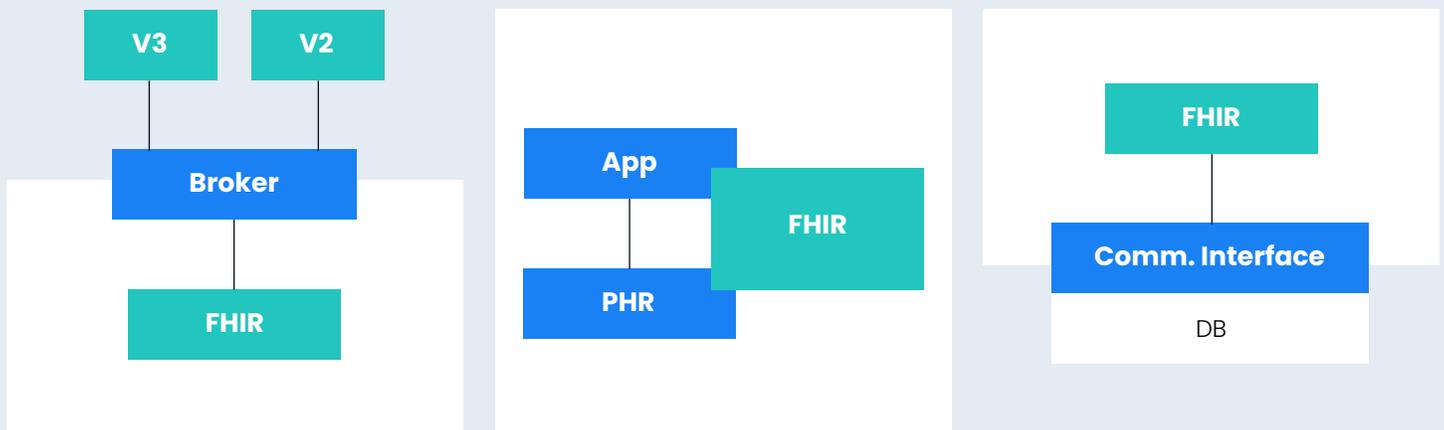
Where Might HL7 FHIR be Used?

The diagram in Figure 5 expands on the third scenario in Figure 4, patient data repository access. Many departmental applications throughout the provider facility will query the FHIR database repository for the data they need, when they need it, allowing for very efficient data transfers and an accurate single source of truth for all data.

Figure 5: **Patient Data Repository Access**
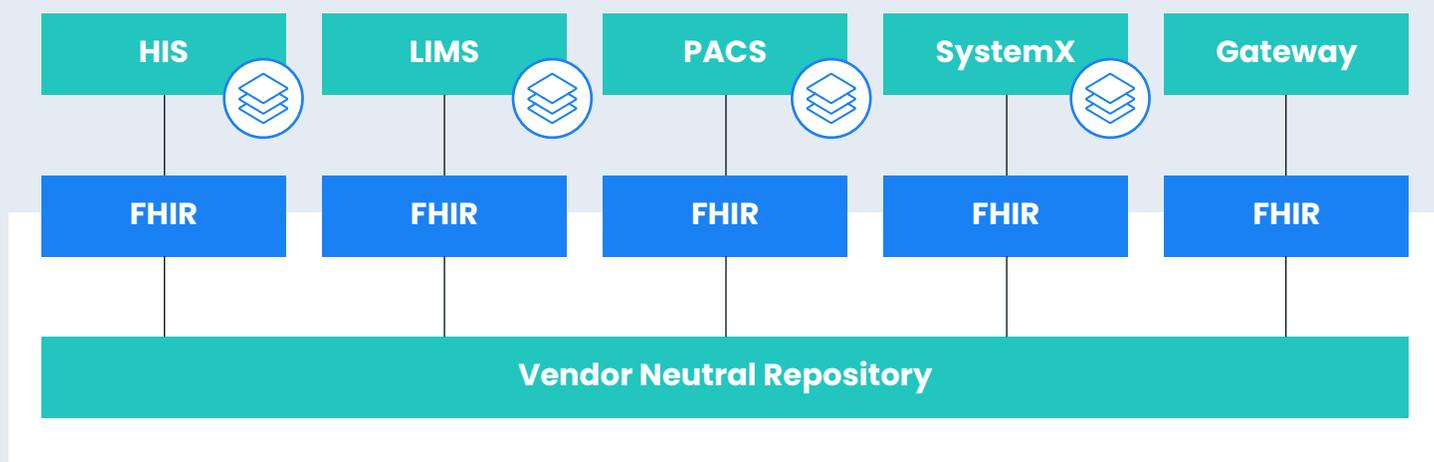
Advances in FHIR

The DaVinci Project

DaVinci is a U.S.-based project that brings together payers and providers of healthcare, along with HL7 International, with a common goal to improve healthcare and reduce waste by moving to value-based care. The health IT experts leading the DaVinci project are working together to accelerate the adoption of FHIR as the standard for integrating data exchange across communities. The project is particularly interesting in light of the new CMS interoperability and patient access rule, which requires CMS-regulated payers to be able to transfer claims from one insurance company to another based on the member's direction. It also requires CMS regulated payers to make provider directory information publicly available via a FHIR based Provider Directory API.

SMART

The HL7 FHIR standard can be used for a variety of workflows, as illustrated in the previous section. However, the key step in enabling the FHIR technology is to have major vendors open up their clinical data in a FHIR repository so other applications can leverage it to solve creative workflow problems. In addition to having the data available, there is an emerging technology that would take this even one step further. The technology is called SMART, and it allows more robust functionality and creative usage of the data directly within the EHR itself. SMART enables third party functionality to be directly embedded within the EHR application itself. In a nutshell, SMART enables third party plug-in apps to run natively inside any compliant EHR or data warehouse that supports the SMART standard. The SMART project started in 2010 with a four year, $15 million grant from the ONC. The idea was to build an app platform for healthcare allowing support apps to be chosen by clinicians. At the HIMSS15 Interoperability Showcase, several vendors demonstrated the progress they are making by incorporating SMART. Since then app development has proliferated, and there are currently 65 apps available on the SMART web site. All the apps use the same underlying set of platform specifications. The SMART specs provide means for health care organizations or developers to access discrete clinical data —such as medications, problems, lab results, immunizations and patient demographics. SMART opens up innovation in healthcare like has not been possible previously. If someone has a bright idea to enhance workflow and creates a SMART app, the healthcare provider doesn't need to wait for an EHR vendor to adopt the idea. Assuming they have a SMART compliant EHR, they can immediately take advantage of the bright idea through the use of the app.

The Argonaut Project

Argonaut is a supporting movement of FHIR in the U.S. led by volunteers. Individuals within HL7 were focused on keeping the FHIR project on target with its self-proclaimed deadlines for adoption. This focus was especially important considering the Open API requirements for Meaningful Use Stage 3. Several of the major EHR vendor organizations also wanted to do what they could to help push development and adoption in the U.S. So HL7, along with these vendor organizations, set out to do two things through the Argonaut Project: First, they funded the work to build U.S. specific profiles, and second, they combine resources to help market and drive adoption of the standard.

The organizations are working with HL7 and the ONC to shed light on the best way to use FHIR in the U.S. considering there are some parts of FHIR that are very international in perspective. For instance, because there is no consistency around the world with how medication is encoded, FHIR doesn't make any rules about how medication is encoded. But in the context of the U.S., RxNorm is the coding system used for medication. HL7 is working through many similar types of issues with the Argonaut members. This will hopefully help drive a cleaner more consistent standard for the U.S., along with a speedier adoption due to the buy-in of the larger vendor organizations that are participating.

Where Does FHIR Go From Here?

Where Will FHIR Go from Here?

FHIR R4 is a significant milestone in the advancement of modern interoperability standards. The promise of R4's normative resources is the catalyst of a stable baseline for FHIR and the fuel to drive broad participation in healthcare's interoperability ecosystem. Because of its many advantages, healthcare organizations around the world are beginning to use FHIR in a variety of contexts, including mobile phone apps, cloud communications, EHR-based data sharing, server communication.

And with the two rules released by the CMS and ONC in 2020, it is now a requirement in the United States for payers, providers, and health IT vendors to facilitate patient access to information through FHIR APIs. Healthcare organizations, even those that have achieved interoperability within their organization, will have to learn to implement new interfaces based on FHIR.

How Lyniate Can Help with FHIR

Lyniate is solely focused on healthcare interoperability, offering the No. 1 and No. 2 integration engines in the market according to KLAS — Corepoint and Rhapsody. Both solutions support all healthcare standards, including FHIR, making them ideal for handling healthcare data integrations, either on premises or in the cloud.

In addition, Lyniate offers extensive training on HL7 and FHIR. Please visit Lyniate.com and select the Education tab to learn more.

Find out more at www.lyniate.com