

FHIR WHITEPAPER

Why we Decided to Invest in a FHIR Response Plan

FHIR is an exciting shift in the interoperability landscape. Many vendors, developers and clients are waiting in anticipation but are unsure of when to make the investment in implementing FHIR into their IT solutions. EMR vendors have different release offerings (resources and verbs) and road maps (timeline and functionality), many focused exclusively on satisfying MU3 Requirements. To jump into this as a replacement for existing processes can seem risky, and a hybrid solution is not very desirable. This is why we decided to invest in developing FHIR Ops API— a single connection that implements FHIR for the end client, yet can fill the gaps for functionality that the EHR doesn't expose.

What's Wrong with Integration Today — Why Do We Need FHIR?

With a full array of stable FHIR resources still a few years off, application developers have struggled to find a consistent way to integrate with the various EMRs. Integrations typically use one of two integration strategies — proprietary APIs or direct database calls — with each integration requiring deep understanding of the EMR.

All major EMR vendors now supply proprietary APIs that allow developers to push or pull data from the EMR. These APIs have varying degrees of maturity and functionality, but the consistent benefit to using the vendor-supplied API is that it requires little expertise from the developer to deploy. The EMR vendor curates a list of functionalities they think developers might want, and manage updating the API to work with past, current, and future versions of their software. That said, there are some drawbacks. Most notably, EMR vendors are pushing to find ways to monetize the APIs they provide with new transaction fees cropping up.

In some cases the transaction fees are inconsequential, but for data-intensive applications these charges can be prohibitive.

IN ONE EXAMPLE, THE TRANSACTION FEES FOR USE OF AN EMR API WERE MORE THAN 5 TIMES THE SELLING COST OF THE SOFTWARE ITSELF.

Another option is to integrate applications by making calls directly to the underlying EMR database. When done well, this option can be cost effective and very performant. The problem here is that this approach brings the highest overhead. Domain knowledge doesn't transfer between EMRs and while you provide native functionality in the EMR, you have to reinvent the wheel if you move to another EMR. Moreover, application developers must be constantly aware of the impact of new EMR versions and detail all changes to ensure that upgrades don't break the integration functionality. In many cases, this type of integration isn't even an option — as access to the database is not granted. Regardless, it's far from an ideal solution as software is now highly coupled to the EMR (and usually a specific version of the EMR).

Once the issues of how to access EMR data from a single pipeline are tackled, the next challenge of creating EMR agnostic applications arises. Imagine finally getting through your first EMR integration only to find out that when you want to sell your application into other EMR markets, the entire integration piece needs to be rewritten because there's no standardization between EMR APIs or databases.

Why Isn't FHIR Ubiquitous Yet?

As with any new standard, there is a lag between the definition of the standard itself and its adoption and implementation. This natural lag occurs for several reasons.

1 Organizations and industries need to assess the benefits of the standard and decide what issues and use cases can be effectively solved through its adoption. Bleeding edge adopters struggle to work through all the technical issues that have not yet been resolved. This process frequently uncovers gray areas that must be dealt with, and causes delays in development, longer project times and higher costs. On the other hand, being overly conservative and waiting until the entire industry has picked up the new standard can leave organizations at a competitive disadvantage and diminish the benefit derived by the standard.

2 Once a need has been identified, technical and business staff must be trained. With a new standard, this training is frequently self-directed. As the standard matures and its adoption is more prevalent, training material and instructor-led courses become more readily available.

Technical and business staff then use training to implement a solution that solves a use case in their organization. Just as with all things in life, one never does anything perfectly the first time. Adopting a standard is no different, and after the first implementation it is important to do a “lessons learned” analysis so any mistakes can be avoided in the future. This post-project assessment should be iterative and done after every implementation.

FHIR is a Total Game Changer — But It'll Take Pragmatism and Discernment to Really Reap the Rewards

At Keena Sized Solutions, we faced the same issues as other organizations in the industry. People want their data, in disparate systems, to be as complete as possible. We wanted to build our applications and have them magically work on all the major EMR Platforms without having to write separate integrations. Turns out, we weren't the only ones with this great idea — thus was the genesis of FHIR Ops. We've taken the core spec that FHIR builds on and extended it to encapsulate the many current and legacy data exchange technologies in play today. We're putting it through the paces by integrating it with our own product lines, and have started to enjoy the benefit of what many consider to be the holy grail in EMR: **unified and ubiquitous data access.**


Take a Dive Down the Rabbit Hole, and Ahead of the Curve

We've taken a hard look at FHIR and realized that it is at least 2-3 years away from being useful. Many vendors are currently focused on supporting read-only operations for a minimal set of resources, while others haven't begun any sort of meaningful implementation. FHIR Ops takes the best parts of the FHIR spec, and implements a data access redundancy strategy based on current options for a specific EMR (and version of that EMR). This granular control gets wrapped up in an easy-to-consume API that is already positioned where FHIR will be in 2-3 years.

Some supported integration methods are EMR FHIR interfaces, proprietary EMR APIs, direct database calls, XDR, XDS and HL7 messages.

FHIR and EMR Agnostic Integration, Futureproofed

When a FHIR Ops request is made, an algorithm determines which routes are supported and based off anticipated performance, cost of the call, access to data, and availability of a service — the method of access is determined and the data returned. For example, when querying demographics, calling the EMR vendor's FHIR API is frequently the easiest, most effective method for getting the data.



However, if the EMR vendor's FHIR endpoint goes down, FHIR Ops automatically adjusts and fails over to another method for retrieving the data. If you need to work with scanned images, the FHIR resources needed are non-existent. FHIR Ops will gracefully figure this out and access the needed data via another method.

FHIR resources for many common clinical use cases may not be available yet, but that hasn't slowed us down. The definitions are there, the data is there and we are bridging the gap. All of the underlying technologies that make up SMART on FHIR have been used in other industries for years. JSON, oAuth2, and RESTful services are all tried and true technologies outside of healthcare. FHIR Ops provides a framework that runs on all the SMART on FHIR protocols. When FHIR resources are available, FHIR Ops returns them.

When FHIR resources are available, FHIR Ops returns them. When they're not, FHIR Ops creates them, and gives you the option to return a full FHIR object or a normalized dataset using the SMART on FHIR protocols.


Using intelligent data source routing algorithms helps **future proof your integration**. As more FHIR resources become stable, they'll be prioritized over using each EMR vendor's proprietary APIs or direct database calls.

Oh – You're Allergic to Upgrades? So Are We

One of the biggest headaches with writing an EMR agnostic application is handling the different datasets that each EMR returns. In one EMR you get the Orange flavored patient. In another you have Banana flavor and yet another the flaveur du jour is Carburetor. Then come the upgrades, promising more of everything. We've all been there... data variation = huge overhead for everyone involved: analysts, developers, patients, providers, managers — everyone. It only further complicates an already complex question of how to improve a patient's health. The goal is to keep the data stable while still progressing on the innovation and services front. The world isn't going to stop changing, and FHIR Ops takes this philosophy to heart. Variations within EMR datasets are normalized and returned with the same structure regardless of where it originated. The value of this can't be underestimated. In essence, FHIR Ops allows you to write one integration point that can send and receive data with most major EMRs without changing a single line of code. Nonetheless, each EMR has bits of data that aren't uniformly available in other EMRs. Rather than muddy the widely accepted schema, **FHIR Ops passes any non-conformant data without loss in a metadata key/value pair object ensuring that no data available in any EMR gets trimmed.**

Another amazing feature with FHIR Ops is its ability to provide data shaping capabilities. With most APIs, developers request data and get whatever dataset the API passes back. If you only want a list of patients with their names and dates of birth, you must run a patient search and return a bloated list of patients that may contain 100 fields you are not going to use. Not only does the search take longer, but the bandwidth required to support the application is 50 times higher. Plus whittling that list down is expensive after the fact. Through parameterized requests, FHIR Ops allows you to specify the desired fields or shape of the data returned in a search. This design significantly reduces memory, bandwidth, and processing for applications built on FHIR Ops.

Finally, the development is made easy by providing a single endpoint for developers. Regardless of the EMR or hospital system you are connecting to, there is a single endpoint to interact with. By securely transmitting a different set of parameters, applications can integrate with multiple hospital systems on multiple EMRs.



End of the Firewall: Decoupling Data from Devices, Supercharging Mobile

FHIR heralds the EMR industry's acceptance of one critical fact: all devices are inherently insecure. The proverbial model of building a castle and securing it with a giant gateway + moat no longer covers everyone's ankles. In today's de-centralized and fragmented mobile marketplace, data mobility has pushed new access and device requirements. People have adopted these high-tech devices and are the true drivers for industries having to tackle how to provide services to the smart device-fueled population that we've become.

In the SMART on FHIR model of security we start with the assumption that the client device (whether used by a patient, provider, analyst or security guru) is already compromised. We weave authentication (who you are) and authorization (what you can do) into the very fabric of data exchange for each piece of data that is both requested and returned. In this way, we eliminate a vast majority of attack vectors and help further secure your PHI. This functionality also distributes the load on servers, as more computing tasks can be pushed to the client. In the real world, 10,000 smart phones using their local resources to render an app will always provide a better experience than a 10,000 core mainframe serving the same app to those remote clients. Using levers like 2-factor Authentication, device registration, caching strategies and timeouts, we can further reduce areas of exposure while still providing lightning-fast data connectivity and data integrity in client-side apps.

IN MANY WAYS, FHIR IS THE EMR'S ANSWER — AND ACKNOWLEDGMENT THAT PHI NEEDS TO EXIST ON THE OTHER SIDE OF THE FIREWALL.

People want to see and interact with their healthcare-related data on their smartphones and tablets. SMART on FHIR is a standardized vehicle that allows us to architect solutions that do just that. As an industry, we can reach the busy lives of our customers and provide value through services like Online Scheduling, Care History Access and Text Messaging to keep patients and providers connected no matter where they are.


EXAMPLES
OF WHAT
WE ARE MAKING
POSSIBLE WITH
FHIR OPS...NOW



Analysts checking on jobs, conversion statuses and other critical internal workflows. The data doesn't have to be directly tied to the EMR — the SMART on FHIR approach allows us to bolt any type of data to the same pipeline and reap the benefits.



Providers quickly checking a patient's chart from their smart device. There is no need to jump through 3 VPNs, 2 reverse proxies and run a IE9 just to see a result or update.



Parents scheduling doctors appointments for their kid's cold from their little league games.



Patients messaging their provider from anywhere in real time with chat/IM capability.



Providers initiating proactive interactions and communications with their patients — like tracking blood pressure via text messaging after a heart attack.

Come see our live demo in the HIMSS Interoperability showcase!

Where To Go From Here

FHIR is a new technology. The standard is maturing, resources are being defined (and refined), and the major vendors are at various stages of implementation.

That's why our company of 27 brilliant and agile thinkers has created FHIR Ops. A product that wraps all of the current infrastructure — HL7, ITI Profiles (XDS, XDR, PIX, etc.), Native API and more — with the security features, standard of data access and data exchange that FHIR promises. FHIR as a spec will continue to mature and when it's fully baked, when we've stepped far beyond MU3 requirements, Byte Sized Solutions will be there to help you ride that future wave of interoperability. We'll be there to help you delight your patients and providers alike, with secure and unprecedented ease of access to the healthcare data that underpins their daily lives.

AUTHORS

If you are interested in a more technical discussion or scoping of your next project, we'd love to chat.



STEVE BUTTITTA
President

sbuttitta@keena.com



JONATHAN MAPLE
Director Innovation
and Development

jmaple@keena.com