



2021 *FAST* Action Plan

March 2021



Version History

Version	Date	Description of Change
Version 1	March 16th, 2021	Published Draft



Table of Contents

Executive Summary	4
Background	5
What is <i>FAST</i>? Why Was it Established?	6
What’s Happening Now — Regulatory Drivers	7
The <i>FAST</i> Structure	8
The <i>FAST</i> Process: Transparent, Iterative, Collaborative	9
<i>FAST</i> Action Plan Overview: <i>FAST</i> Infrastructure Architecture and the Path to Implementation Through Standards Development, Guidance, and Tools	12
2021 <i>FAST</i> Action Plan Individual Solutions Summaries and Path	14
Solution 1: A US Wide Solution for FHIR Endpoint Discovery.....	15
Solution 2: A US Wide Methodology for Supporting Multiple Production Versions of FHIR.....	19
Solution 3: US-Wide Scaling Requirements for FHIR RESTful Exchange Intermediaries.....	23
Solution 4: Standards-Based Approaches for Individual Identity Management.....	27
Solution 5: US-Wide Model(s) for Scalable Security Solutions.....	31
Solution 6: Hybrid/Intermediary Exchange (Exchange with or without Intermediaries).....	38
Solution 7: A FHIR Testing and Certification Platform.....	42
Pilot Testing and <i>FAST</i> Solutions Validation Strategy	46
Conclusion	48



Executive Summary

The Health Level Seven (HL7) Fast Healthcare Interoperability Resources (FHIR) standard has provided the healthcare industry with a clear path to sharing data effectively through Application Programming Interfaces (APIs). As FHIR is adopted, healthcare organizations will be able to exchange data more seamlessly and patients will have access to their information, increasing transparency. However, our industry needs to be able to use APIs at scale. The scale challenge is significant: the number of potential connections between thousands of provider organizations, payers, public health agencies, labs, intermediaries, 3rd party applications, and patients requires enabling dynamic use of FHIR based APIs. How will this deployment work? How will our industry create an environment where FHIR solutions can be implemented at scale and support efficient health information flow between multiple exchange partners. Without scalability, the industry won't reach the full potential of interoperability.

The FHIR at Scale Taskforce (*FAST*) initiative was launched to address such industry wide FHIR scalability challenges. In a step-by-step approach, the taskforce has identified infrastructure-related challenges and barriers, developed and proposed practical solutions to address identified challenges and barriers, and ultimately, after thoroughly vetting those proposed solutions in an open, transparent and inclusive way, is now identifying and driving those solutions into effective practice through tangible outputs (such as implementation guides and other openly accessible resources).

This document speaks to the *FAST* story—its goals, scope, structure, and process—while also summarizing *FAST*'s work, describing those barriers to FHIR scalability and the core capabilities required to overcome such barriers.

Most importantly, the *FAST* Action Plan:

- Summarizes each recommended *FAST* solution
- Describes individual solution paths to implementation
- Communicates how the industry can get involved to help bring these solutions to implementation readiness and industry adoption

As the *FAST* solutions make their way toward final implementable artifacts through standards, guides, and tools, this Action Plan will remain a living draft and interactive document—until the path forward has been fully solidified. The *FAST* project will continue to work closely with all stakeholders in the industry to support the ability for FHIR to efficiently run at scale.

Note: this document has linked sections for easy reading. Use the home and arrow navigation links above.

***FAST* ACCOMPLISHMENTS & KEY ARTIFACTS**

[2020 End of Year Report](#)

[2020 Mid-Year Report](#)

[2019 Annual Report](#)

[FAST 101 Presentation](#)

[2020 SME Sessions](#)

[2020 *FAST* Workshop](#)



Background

FHIR is an HL7 standard that defines data formats and other elements for secure and efficient exchange of electronic health related information.


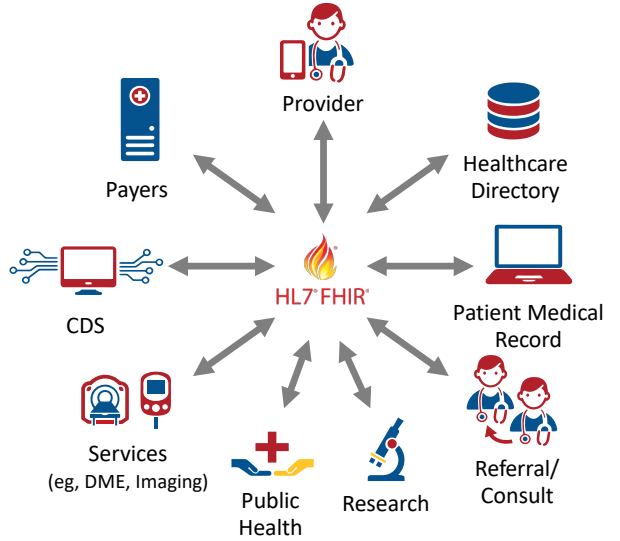
The use of the FHIR standard for health information exchange has many benefits, such as supporting REST software architecture, employing other web technologies (eg, XML, OAuth, HTTP protocol), and improving interoperability. Because FHIR contains more discrete data elements, it's easier to transmit and integrate the necessary pieces of information, improved secure access to health data—making it easier to access, interact, and transmit health data.

Recognizing the potential and benefits of the FHIR standard, the industry found itself at an inflection point regarding FHIR adoption. As of late 2018, nearly 51% of health IT developers appeared to be using a version of FHIR combined with OAuth 2.0.¹ Private companies, government institutions (eg, CMS Blue Button), and stakeholders participating in market-driven collaboratives (eg, HL7 FHIR Accelerator projects), are all developing FHIR-based solutions to many essential industry functional use cases, in an effort to move toward value-based care (VBC). However, with these growing number of FHIR-enabled solutions aiming to solve functional use cases, the industry uncovered important scalability challenges and the need for infrastructural approaches to support such functional solutions to truly run at scale, moving the industry away from the limitations of the point-to-point exchange.

FHIR V4 release as well as HL7 FHIR Accelerator Programs have also increased FHIR-based solutions' adoption. Various accelerator initiatives eg, Da Vinci, Gravity, etc. came together and built on functional use cases aiming to facilitate data integration to improve areas such as patient outcomes, provider workflow, social determinants of health, and the transition to VBC. While solving real-world business problems (ie, functional use cases), solutions coming out of the current FHIR Accelerators and other multi-stakeholder collaboratives

often do not address infrastructural barriers to FHIR scalability. These infrastructure barriers need to be solved in order to evolve from point-to-point data exchange to achieve dynamic and efficient health information flow between multiple exchange partners.

Our industry approached the Office of the National Coordinator (ONC) to facilitate a cross-stakeholder project to address infrastructure barriers to FHIR scale to solve these scalability challenges. The result was the FHIR at Scale Taskforce (*FAST*).




1. Posnack S, Barker W. Heat Wave: The U.S. is Poised to Catch FHIR in 2019. <https://www.healthit.gov/buzz-blog/interoperability/heat-wave-the-u-s-is-poised-to-catch-fhir-in-2019>. Published October 1, 2018. Accessed March 4, 2021.



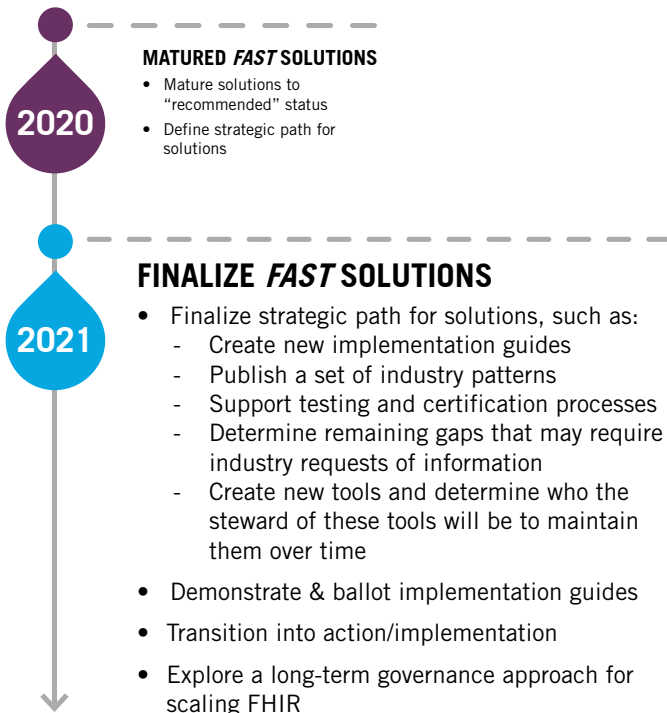
What Is FAST? Why Was It Established?

FAST was launched in late 2017 in response to an industry-recognized need to address shared FHIR scalability challenges. Acting as a neutral convener, the ONC facilitated the engagement of a highly representative group of volunteer healthcare industry stakeholders and health information technology experts, who worked together to identify these FHIR scalability infrastructural challenges, define the core capabilities that cut across those challenges, and ultimately develop solutions to effectively address them.



The FAST Initiative started by establishing a vision of a scalable ecosystem infrastructure and a project charter. In 2018, FAST transitioned to perform a landscape analysis, exploring various industry functional use cases in need of scalability solutions and defining the policy and technical barriers to be addressed. From there, FAST has moved toward defining, proposing, and maturing the taskforce scalability solutions. By engaging the industry and continuously gathering feedback, FAST was able to ensure solutions are relevant and aligned with the industry needs.

Since its inception, FAST has greatly broadened its stakeholder participation and industry engagement. Based on the applicability of the FAST approaches to many areas across the industry, and the positive impact unlocking data can have on patients' outcomes and the reduction of provider burden with the use of health IT, FAST experienced industry-wide interest and participation. The FAST Ecosystem Tiger Team has developed over 15 ecosystem use cases and the core capabilities required for scale, while the remaining 6 FAST Tiger Teams analyzed these infrastructural core capabilities needed to scale FHIR solutions, and developed solutions in the respective areas of focus for each of the Tiger Teams. In addition, by establishing and engaging with the FAST Technical Learning Community (TLC) and other FHIR initiatives, FAST continued to inform the broader public on the taskforce goals and approaches.



FAST progress has been continuously shared at industry conferences, workgroup meetings, and through yearly reports in an effort to be as transparent as possible and to gain broad industry perspectives and feedback on the solutions being developed. In order to create comprehensive and successful solutions that meet the needs of all stakeholders, an open and transparent engagement strategy was identified as critical from the start. This commitment to being inclusive of all potential stakeholders helps ensure the industry as a whole will be willing and able to one day adopt and deploy any of the solutions being put forth by FAST.

As a result, 2020 was a critical year for the FAST Initiative; the FAST-proposed solutions were more deeply vetted and validated with key subject matter experts (SMEs) from across the care continuum and outside the healthcare ecosystem. This proved to be valuable as it provided the FAST Tiger Teams with highly engaging forums for an innovative review of their proposed technical approaches that included learning from these challenges being addressed by other industries, such as finance, tech, and cloud computing.

Today, the FAST-recommended solutions are on a path to standards-based implementation guides, industry guidance, and/or asset building; they will be ready for implementation as detailed later in this FAST Action Plan.



What's Happening Now — Regulatory Drivers

Recent HHS regulations and Federal laws have set the stage for a significant expansion of the use of FHIR to address interoperability with regard to the exchange of healthcare information. The most significant regulations are:

1. [ONC 21st Century Cures Act](#): Interoperability, Information Blocking, and the ONC Health IT, Certification Program
 - Affects all Certified EHRs
2. [CMS Interoperability and Patient Access](#)
 - Affects all Medicare Advantage, Medicaid FFS, Medicaid Managed Care, CHIP FFS, CHIP Managed Care, and QHP in the Federal Exchange
 - Affects Condition of Participation Hospitals and Services (ADT)
3. [CMS Reducing Provider and Patient Burden](#) (publication pending)
 - Medicaid FFS, Medicaid Managed Care, CHIP FFS, CHIP Managed Care, and QHP in the Federal Exchange

While the following regulations do not specifically mention FHIR, the ability to effectively address the specific requirements related to Transparency of costs for specific health care service can only be effectively met by utilizing FHIR APIs and consumer-based applications where the consumer is directly involved.

4. [Hospital Transparency](#)
 - Affects all hospitals
5. [Transparency in Coverage](#)
 - Affects all group health plans and health insurance issuers offering group or individual health insurance coverage
6. Consolidated Appropriations Act – [No Surprises Act](#) Section
 - Affects all health care providers and health care facilities
 - Affects all group health plans and health insurance issuers offering group or individual health insurance coverage

Regulation	EHRs	Payers	Patients	Providers	Hospitals
ONC 21st Century Cures Act	All certified EHRs	None	300M	All using certified EHRs	All using certified EHRs
CMS Interoperability And Patient Access	None	All MA, MC, CHIP, QHP/FFE	125M	None	All CPO for ADT
CMS Reducing Burden	None	All MC, CHIP, QHP/FFE	100M	None	None
Hospital Transparency	None	None	None	None	All
Transparency in Coverage	None	All	160 M	None	None
No Surprises Act	All	All Commercial	330 M 160 M	All	All

If FHIR at scale to enable effective operations across organizations nationally was *FAST's* original call to action, the drive to adopt due to regulation is a loud and clear driver that FHIR and API adoption is here to stay, and scalability will be needed quickly in order to ramp adoption in an organized and efficient way.

Numbers are estimated based on scope of coverage



The FAST Structure

The FAST initiative adopted a taskforce Tiger Team model to harness the knowledge and expertise of various SMEs across industries, as well as to capture broader industry perspective to ensure the FAST solutions align with the needs of the industry at large.

Seven Tiger Teams, staffed entirely by volunteer SMEs, were established to focus on solving the FAST identified technical and regulatory/policy barriers, as well as identifying core capabilities essential to FHIR-based solutions scalability. Six priority areas to support FHIR to scale were identified to address:

1. Identity Resolution and Matching
2. Endpoint Directory, FHIR Version and Scale
3. Scalable Security
4. Exchange Process With and Without Intermediaries
5. Scalable FHIR Testing and Certification/ Validation Platform
6. FAST Solutions Pilots Testing

There are 2 FAST Chief Architects who ensure the taskforce is focused on the industry's most relevant infrastructural challenges and solutions, while identifying the linkages across solutions and developing an overall architectural framework that support all solutions working in harmony. The FAST Coordinating Committee and Executive Steering Committee, a mix of public-private, significant, top industry leaders provide continuous support and member resources to the taskforce.

The FAST TLC and SME Panels are external, but critical components of the FAST structure; they provide the input and feedback necessary to ensure FAST's work aligns with industry direction.

FAST Organization & Community Engagement

TIGER TEAMS IDENTIFY:

- Ecosystem Use Cases
- Technical Barriers
- Regulatory and Policy Barriers
- Core Capabilities
- Gap Analysis of Industry Models and Functional Use Cases

EXECUTIVE STEERING COMMITTEE (public-private mix)

COORDINATING COMMITTEE (public-private mix)

SEVEN TIGER TEAMS

- Ecosystem Use Case
- Identity
- Security
- Directory, Versions and Scale
- Exchange Process
- Testing & Certification
- Pilots



SUBJECT MATTER EXPERTS (SME) PANELS



TECHNICAL LEARNING COMMUNITY (TLC)



INFORMATION SHARING WITH TLC THROUGH:

- Website
- Periodic webinars
- Newsletters
- TLC Meetings
- LinkedIn Group



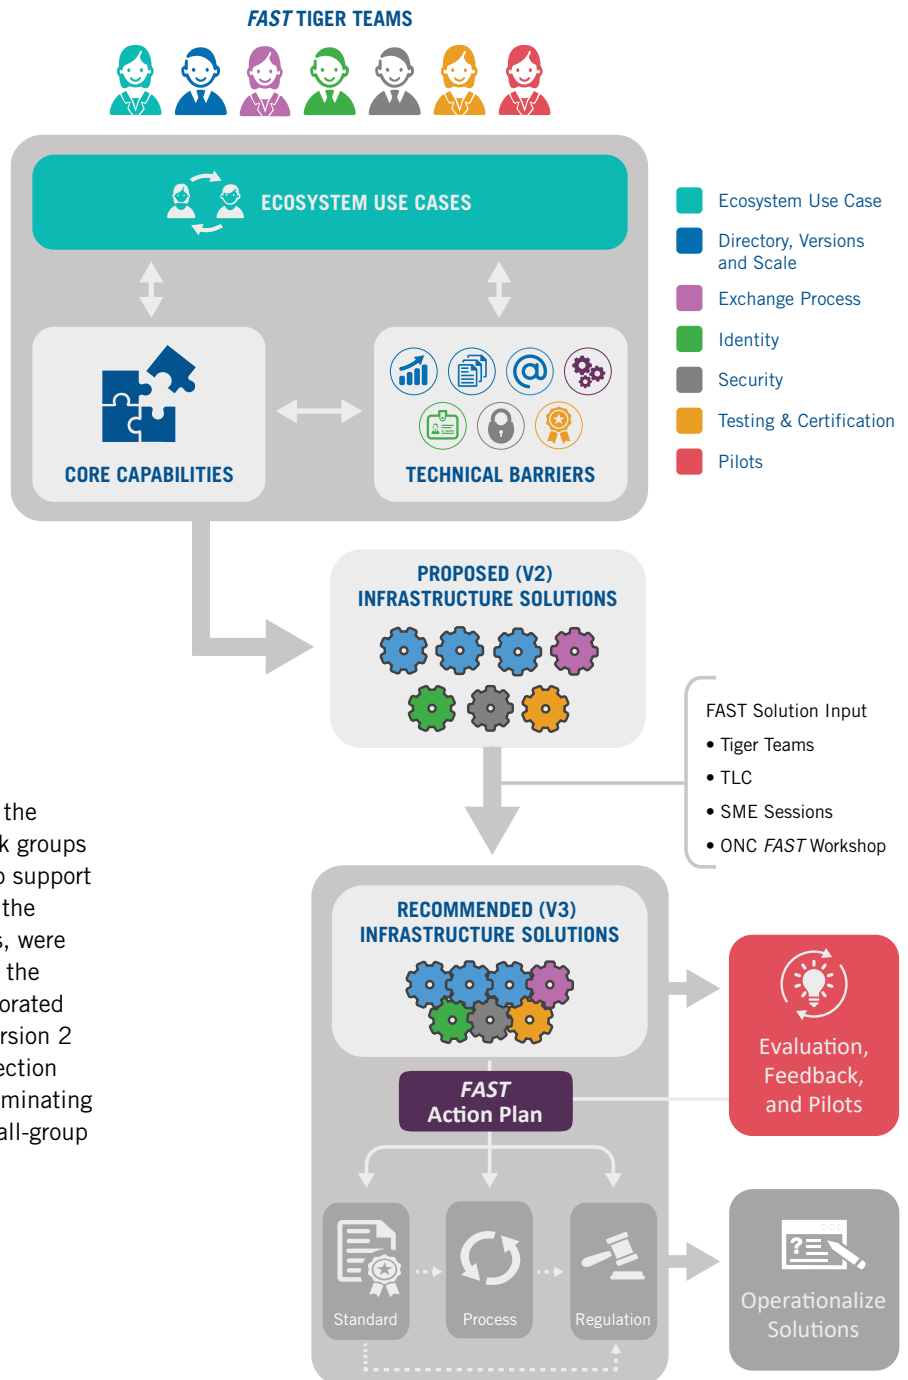
The FAST Process: Transparent, Iterative, Collaborative

Aiming to develop and propose approaches that address industry-wide applicable infrastructural FHIR scalability challenges, the FAST Initiative embarked on a multi-year iterative, transparent, and collaborative process, leading to additional standards work, processes, or other proposed measures.

The FAST Ecosystem Use Case Tiger Team was tasked to analyze a representative list of functional use cases developed by FHIR Accelerators and other HL7 Work Groups—from the perspective of the ecosystem scalability needs and challenges. Their analysis resulted in the identification of common ecosystem infrastructure challenges that most implementers might encounter, and the core capabilities required to overcome them. As the ecosystem use cases and core capabilities were being defined by the Ecosystem Tiger Team, the remaining 6 Tiger Teams, each focused on a specific scalability barrier, explored the technical and policy barriers reflected in these use cases—in line with the core capabilities as well as the potential solutions. While each Tiger Team was assigned a specific focus area, ongoing collaboration and knowledge sharing across Tiger Teams and with the industry were built into the FAST Process, through the members’ participation across Tiger Teams and cross-organizational representation by the Task Force members.

Another critical component of the FAST Process is the systematic socialization of the FAST work with work groups and SMEs, gathering feedback from the industry to support iterative solution improvement. Version 1 drafts of the solutions, developed mainly within the Tiger Teams, were shared early in the process with the FAST TLC and the feedback collected from those sessions was incorporated into the FAST-proposed solutions leading to the Version 2 set of solutions drafts. This iterative feedback collection and solutions advancement process continued, culminating with 2 important milestones in 2020: targeted small-group SME sessions and a public ONC FAST Workshop.

FAST Solution Process and Where We Are Now





SME Panel Session Highlights

In the first half of 2020, *FAST* prepared and convened small panels of SMEs, in a series of 13 working sessions. SMEs were invited to participate based on their expertise in each area and came from various industry backgrounds, including healthcare, finance, and software development.

Discussions explored how each solution could be implemented most effectively, uncovered any remaining gaps and barriers that need to be addressed, and gleaned other insights that could strengthen *FAST* solutions—as well as beginning to explore potential paths forward.

The level of subject matter expertise represented in each of these sessions was staggering and their feedback was critical to maturing the solutions into Version 3/recommended status. Following these small SME working panels to further gather industry-wide feedback, *FAST* held a full-day ONC-hosted Workshop in September 2020. The Workshop helped further validate the *FAST* solutions to FHIR scalability challenges, gain broader-industry consensus and confirm the appropriateness of the proposed path forward for each of the *FAST* solutions, gaining further industry acknowledgment, consensus, and support. The full day of events included 2 plenaries, 3 tracks with 3–4 concurrent sessions in each track to explore:

- Solution interdependencies
- Potential drivers for solution adoption. based on stakeholder nuances
- How each of the solutions could develop into standard, process, or policy
- The potential roles of the intermediaries
- Pilot testing
- Potential governance process

The Workshop enjoyed broad representation across the industry and robust, live interaction, discussions, and engagement from all the participants, panelists, and other attendees. The full agenda and access to the content covered is available through the [2020 *FAST* Workshop dedicated page](#) and the video is available on the HealthIT.gov page [here](#).

As a direct result of taskforce members' dedication to focused taskforce work through 2020, 3 HL7 Project Scope Statements (PSSs) were submitted and approved to date. The PSS approval for those 3 solutions initiates the development of HL7 FHIR Implementation Guides based on the following *FAST*-recommended solutions:

- Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)
- US-Wide Solution for FHIR Endpoint Discovery and
- US-Wide Model(s) for Scalable Security Solutions

SME PANEL SESSION HIGHLIGHTS

- 149 SMEs were nominated
- 84 SMEs were invited to participate
- 315+ slides were developed
 - 13 proposed solutions
- SMEs, in total, contributed 362 hours of their expertise
 - 5 SME Prep Sessions
 - 8 SME Sessions
- 1,262 contractor hours between Jan-June were used
- The Chief Architects and Tiger team volunteer members ramped up their work spending countless hours preparing for and attending SME Sessions

FAST WORKSHOP PARTICIPATION

- Over 130 SME Panelists were invited to participate
- Over 65 organizations were represented on SME Panels
- 10 interactive breakout sessions were conducted plus opening and closing plenaries
- Chief Architects and Tiger Team volunteers ramped up their work spending additional time preparing for, attending, and facilitating workshop sessions
- Open to general public



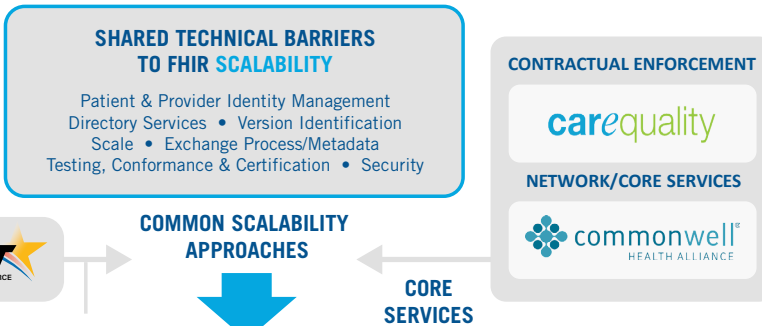
Paving the Way Towards FHIR “at Scale”

HL7® FHIR® ACCELERATOR




In addition to the internal *FAST* Process, the *FAST* Initiative recognized the importance of an ongoing broader-industry collaboration process and established an external industry-engagement strategy. The purpose of this strategy was to increase awareness about *FAST*'s work and leverage forums to seek and gather feedback from a variety of outside-industry stakeholders throughout the *FAST* solutions development process. In these presentations and collaborative sessions, *FAST* sought to help stakeholders understand how their functional use cases related to, and could be supported by, the *FAST*-proposed infrastructural solutions use cases. Additionally, *FAST* wanted to advance the proposed solutions and pave the way toward FHIR at scale through alignment with other health IT advancements and by meeting any of their scalability needs.

FUNCTIONAL USE CASES



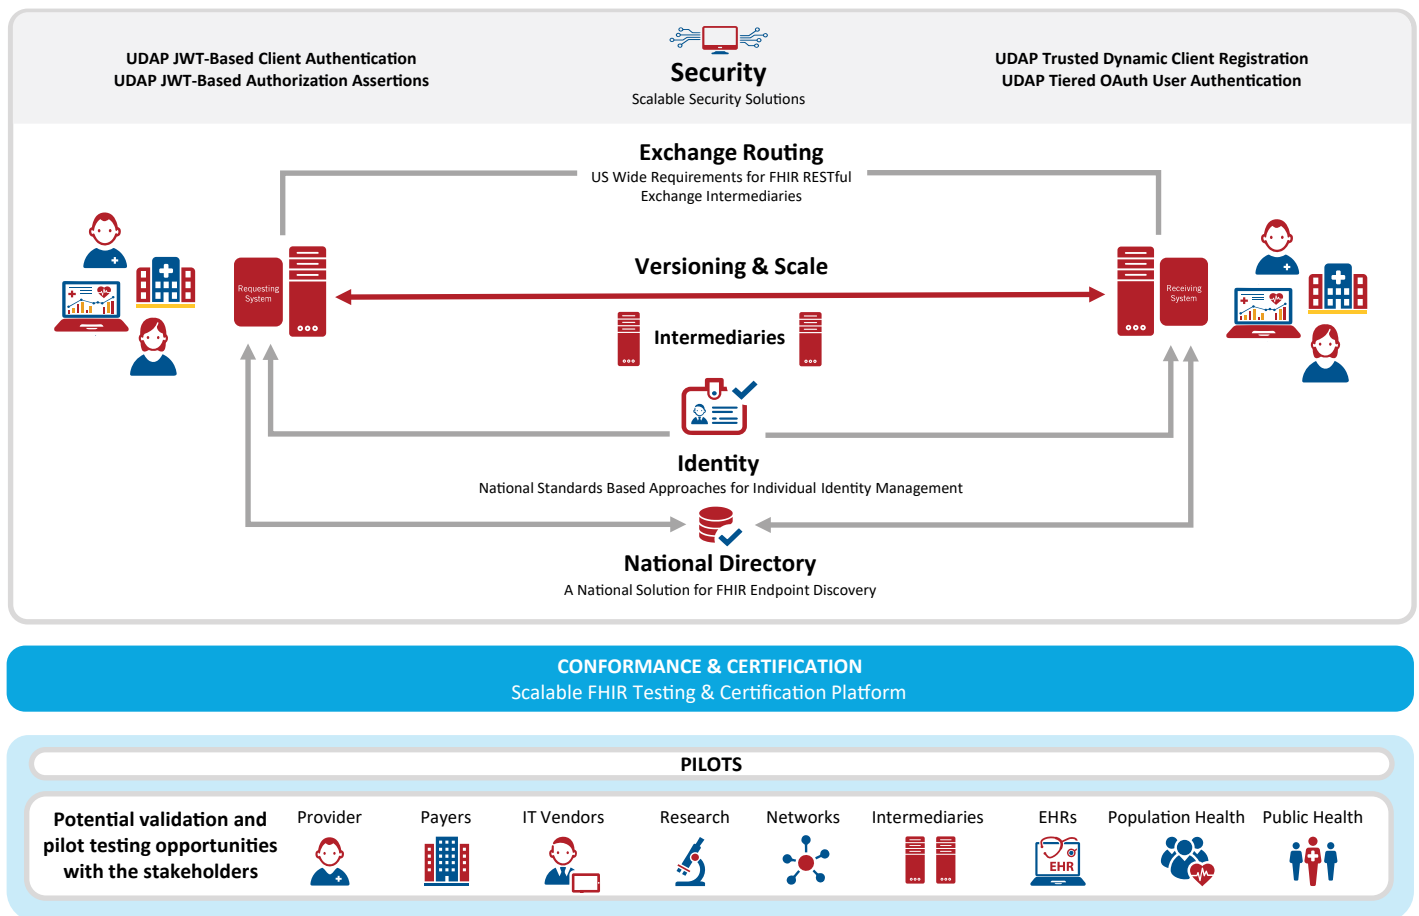
INFRASTRUCTURE USE CASES





FAST Action Plan Overview: FAST Infrastructure Architecture and the Path to Implementation Through Standards Development, Guidance, and Tools

The overall FAST conceptual infrastructure architecture and the individual FAST solutions include a battery of capabilities that can be implemented individually or ideally in concert with each other for a more comprehensive and effective approach to supporting FHIR at scale. For example, a FHIR transaction can leverage FAST solution proposals in sequence: find the API endpoint, ensure client registration and authorization, collaborate on effective patient match, and route the FHIR transaction, if necessary, through partners. The solution areas described below work together.



The diagram above represents, at a high level, that the FAST solutions are architected to provide stand-alone capabilities but also come together to form a holistic approach to FHIR at scale. Each solution can be implemented individually to support a gradual roll-out over time as the solutions become available in the standards processes. Once all solutions are available, synergistic effect is achieved.

The goal of the FAST Action Plan, depicted on the following page, is to clearly describe scope and status of each solution and whether those solutions should take the form of standard, tool or guide, or a combination. This Action Plan will also describe ways for stakeholders to get engaged and provide additional feedback or participate in Connectathons and pilot testing. As such, this is a living document and will be updated as the paths for more FAST solutions are established.

The graphic on the next page shows each solution, how it will manifest and how stakeholders can get involved and influence the path forward whether through joining HL7 work groups or participating in Connectathons.



SOLUTIONS	ARTIFACTS	INDUSTRY ENGAGEMENT	
US Wide Model(s) for Scalable SECURITY Solutions	HL7 IG: FAST: Scalable Registration, Authentication, and Authorization for FHIR Ecosystem Participants	Implementation Guide Development & Balloting JOIN the HL7 Sponsoring Workgroup: Security Targeting Sept 2021 STU Ballot	2021 Connections Look for us at the HL7 FHIR September Connection Sept 13-14, 2021 Join us at the HL7 FHIR May Connection May 17-19, 2021
A US Wide Solution FOR FHIR ENDPOINT DISCOVERY	BUILD ASSET: National Source of Truth Directory	Support Naming the Developer Organization Contact the FAST Directory, Version & Scale Tiger Team Leads	
	HL7 IG: Attestation and Validation	JOIN the HL7 Sponsoring Workgroup: Patient Administration Targeting Sept 2021 STU Ballots and Ballot for Comment (Attestation and Validation)	
	HL7 IG: Endpoint Query		
HL7 IG: Update VHDir for US Realm			
An HL7 FHIR Standard Based Solution for EXCHANGE with or without Intermediaries	HL7 IG: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)	Implementation Guide Development & Balloting JOIN the HL7 Sponsoring Workgroup: FHIR Infrastructure Targeting September 2021 STU Ballot	
US Wide SCALING Requirements for FHIR RESTful Exchange Intermediaries	Finalizing the Solution Path Forward FAST solution documentation	SME Session to be Scheduled/ Look for updates/results from SME Session on the FAST project page	Follow the FAST Technical Learning Community on LinkedIn for updates
A US Wide Methodology for Supporting Multiple Production VERSIONS of FHIR	Finalizing the Solution Path Forward FAST solution documentation	SME Session to be Scheduled Look for updates/results from SME Session on the FAST project page	
FHIR TESTING & CERTIFICATION Platform	BUILD ASSET: Testing & Certification Tooling Finalizing the Solution Path Forward Exploring potential for RFI	Support RFI Process Respond to RFI Contact the Testing & Certification Tiger Team Leads	
National Standards Based Approaches for Individual IDENTITY Management	Proposed Guidance: IAL2 Identity Verification Best Practices in Healthcare to Establish Digital Identity	Consider overlaps with other industry initiatives Socialize recommended solutions with relevant industry stakeholders Contact the Identity Tiger Team Leads	
	Proposed HL7 IG: Improving identity assurance and patient match quality through interoperable Digital Identity and Patient Matching capabilities	Implementation Guide Development & Balloting JOIN the HL7 Sponsoring Workgroup: Patient Administration Targeting 2022 STU Ballot	

Pilots, Testing, Validation & Industry Stakeholder Engagement

Provider



Payers



IT Vendors



Research



Networks



Intermediaries



EHRs



Population Health



Public Health





2021 *FAST* Action Plan

Individual Solutions Summaries and Path



Solution 1: A US Wide Solution FOR FHIR ENDPOINT DISCOVERY

In today's health care landscape, it's difficult to find and discover electronic endpoints³ for various entities, like payers, clinical care providers or other health services providers, because there isn't one central reliable source of truth. It may require "tribal" knowledge to determine where an endpoint is defined and, in many cases, there is no resource to find an endpoint for a specific provider, organization, or service. Furthermore, once an endpoint has been located, it's hard to understand the endpoint's associated capabilities for data exchange and whether the endpoint information is current and accurate.

Adding to the complexity of discovering and working with an endpoint is the diversity of endpoints that are needed (ie, specific provider endpoint, regional health insurer or other). Currently, there isn't an industry wide agreed upon infrastructure and "rules of the road" to ensure that metadata associated with an endpoint (eg, FHIR version, trust framework) is accurate and, when necessary, access is restricted (eg, the endpoint of a women's shelter or emergency response team).

The *FAST* National Healthcare Directory solution seeks to establish the technical framework and the infrastructure capabilities for a scalable national directory of validated healthcare related individual and organization electronic endpoints, where entities contribute, verify, and keep up to date their demographic information, relationships, and endpoint connection information. To accomplish this, the recommended scalable infrastructure proposes one 'source of truth', that will scale and be made available to many federated⁴ healthcare directories, which could then be queried for ongoing discovery and connection to endpoints. The relevant demographic, relationships and endpoints will be validated and maintained for accuracy in one place, the national healthcare directory, to facilitate access and appropriate exchange of health information for care delivery, public health reporting, prior authorization, reporting quality data, payment, and a multitude of other workflows in support of healthcare delivery and patient outcomes.

3. Endpoints are locations that can be connected to for the delivery or retrieval of information (eg, URL of a server or service)

4. Federated directories will maintain a copy of all or part of the validated directory information to provide discovery of and/or access to endpoints to the applications/population they serve.



@ Solution 1: A US Wide Solution FOR FHIR ENDPOINT DISCOVERY

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

The industry lacks a generally available method to find healthcare related individuals, organizations, and FHIR endpoints and their associated capabilities and attributes, as well as a common process for maintaining the information and validating its accuracy



SOLUTION

One national source for validated directory information that is available to federated directories for enhancement and integration into specific workflows



CURRENT SOLUTION

[FAST Endpoint directory proposed solution document \(Version 3\)](#)



OUT OF SCOPE

Manual / portal access, creation of a trust framework, non-FHIR related endpoints, application certification process



IN SCOPE

Healthcare validated directory including individual and entity demographics and relationships to determine appropriate endpoint(s), computable endpoint information such as versions supported, metadata for routing, trust framework, implementation guides and certification status

Federated access by HIEs, state directories, EHRs

FHIR standard implementation guides for use of the directory:

- Update the Validated Healthcare Directory (VHDir) for the US Realm
- Endpoint query
- Populating the directory - attestation and validation



STATUS

Incorporated feedback from industry stakeholders to finalize the V3 solution document. Read the full SME session report [here](#). Review slides from the FAST Workshop [here](#).

Submitted and obtained approval for HL7 Project Scope Statement (PSS) to initiate development of 3 HL7 Implementation Guides, targeted for Sept 2021 STU Ballot:

1. Develop the standard for the national healthcare directory including FHIR endpoints
2. Develop the standard for populating the directory and validating the information contained in the directory
3. Developing the standard for querying by the federated directories structure for endpoint discovery



OPEN ITEMS

- Define the minimum valuable product (MVP) and outline incremental steps/roadmap to build a validated national healthcare directory
- Develop implementation guides to support the ability to contribute to, validate, distribute, and query directory information,
- Test initial proposed solutions at HL7 FHIR Connectathons
- Complete ballot cycle



BENEFIT

- Improved discoverability of healthcare related individuals, organizations, their relationships, and endpoints
- Standardized directory information exchange
- Standardized federated directory query method(s)
- National approach for reliable endpoint data
- Decrease provider burden in providing and verifying information
- National utility on which to populate, discover and distribute endpoint capabilities further extending innovation and efficiency initiatives
- Validated information accessible to existing directory implementations



Solution 1: A US Wide Solution FOR FHIR ENDPOINT DISCOVERY

Foundational work in this space exists and would be leveraged. The [Validated Healthcare Directory](#), which established an international standard for exchanging directory information across a federated environment, will serve as the base for the FAST National Healthcare Directory. However, this work needs to be updated and localized to the United States. There are some resources that offer partial information about an endpoint, but there isn't one repository or "source of truth" of validated demographics, relationships and endpoints. Currently, providers must register at multiple societies and healthcare organizations, which duplicates efforts; it's creating a process rife with errors and out of date information, reducing the ability for endpoints to be reliably discoverable. Having one source of truth would decrease the burden on providers and owners of endpoint information who should be able to verify details about their endpoints without having to do so in multiple places. Federated access is likewise important, ensuring scalability, that the validated information can be efficiently and appropriately funneled, repurposed and accessible in multiple places, supporting existing networks and directory implementations. To further support scalability, automation of as many of these processes as possible is imperative.

The FAST National Healthcare Directory recommended solution includes 3 elements for the path forward:

1. Develop the standard for exchanging information from a national healthcare directory to a federated access architecture
2. Develop the standard for populating the directory and validating the information contained in the directory, and
3. Developing the standard for querying by the federated directory structure for endpoints

FAST Endpoint Directory — Architecture and Workflow

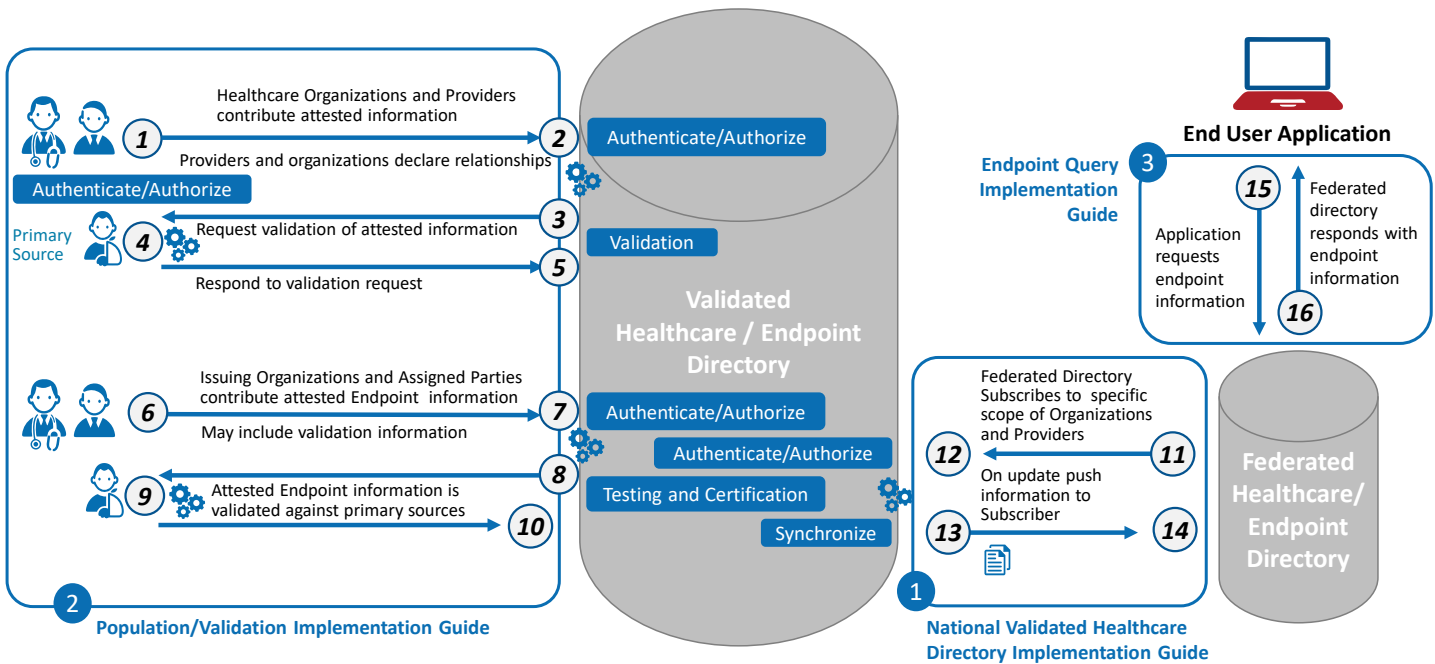
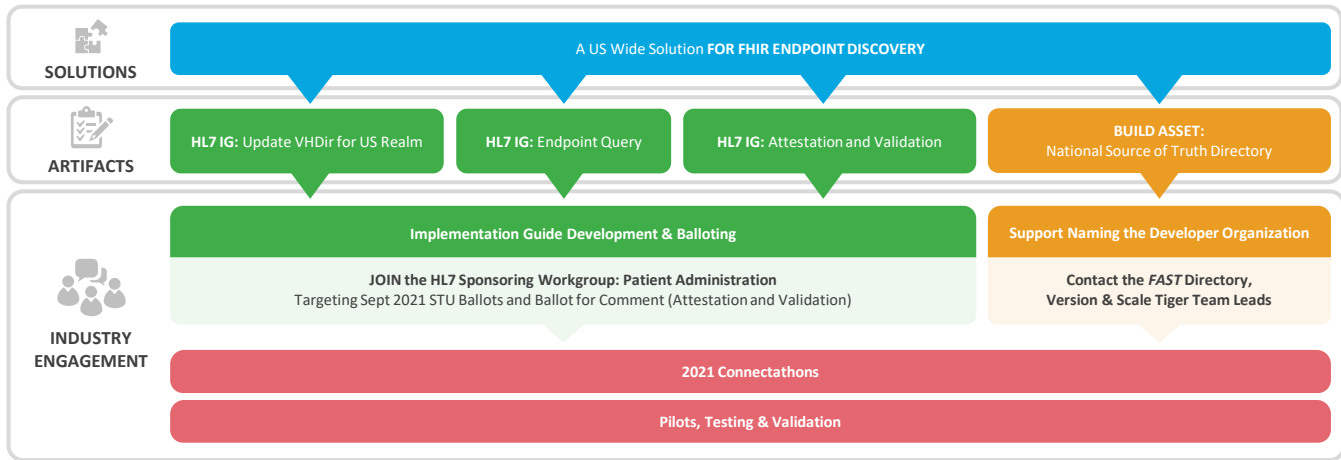


Figure 1



Solution 1: A US Wide Solution FOR FHIR ENDPOINT DISCOVERY

Next Steps and Industry Impact



The next steps for this solution are to:

- Create standards by following the HL7 Implementation Guide development process for 3 distinct Implementation guides.
 - Update VDir for US Realm
 - Endpoint Query
 - Attestation and Validation
- Identify organization to oversee/build a national directory
- Build a national directory using newly established standards to be the “source of truth”
- Issue best practices/operational processes for usage of the new “source of truth” directory



This solution offers these key benefits:

- Reducing burden by:
 - Offering endpoint information, including endpoint attributes, in one national resource for endpoint validation and distributions
 - Offering one national resource for updating and maintaining endpoint data for an organization
 - Reducing the need to create unique point to point solutions
 - Automating discovery of endpoints based on specific search criteria identifiable by originating entity
- Support adherence to the CMS proposed rule on reducing provider and patient burden by improving prior authorization processes and promoting patient’s Electronic Access to Health Information
- Ameliorate challenges related to endpoint identification issues raised by the CMS rule on Interoperability and patient access
- Streamline and support VBC workflows



This solution will have the biggest impact

on: payers, providers, health systems, EHRs, public health, application developers and community-based support services



Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Joining HL7 Patient Administration Workgroup to provide input and ballot implementation guides
- Participating in the HL7 Connectathons in May and September to demonstrate and refine the solutions
- Joining discussions and provide input on which convening organization should oversee and/or develop the national directory as the source of truth.





Solution 2: A US Wide Methodology for Supporting Multiple Production VERSIONS of FHIR

As the FHIR standard usage becomes ubiquitous across the nation, the healthcare ecosystem will need to support multiple versions of FHIR resources and their associated APIs. Unlike the classic world of healthcare interoperability in which versions rarely changed (and when it did, everyone changed at once), modern API patterns support multiple versions at the same time. Regardless of what FHIR versions are included in regulation, or the version adoption rules that stem from ONC and CMS rules, the industry is going to wind up with multiple versions of FHIR and multiple versions of resources, extensions, profiles, value sets, and implementation guides.

In today's environment, most FHIR endpoints only support one version of FHIR, and there are multiple incompatible versions of FHIR in production (eg, DSTU2, STU3, R4). These versions are not fully backward or forward compatible, which means that incompatible or "breaking" changes may exist between versions (ie, changes in one FHIR version that would cause a system or application using a different FHIR version to fail) except where FHIR resources are "normative" (ie, content is considered stable and compatible between versions). While mapping solutions exist for some FHIR resources to convert from one version to another, their capability, quality, and completeness vary from resource to resource. In addition, there is no definitive source for transforms.

Initiators of transactions must be able to determine the version(s) available at the endpoint from which they are requesting information or know that the endpoint can handle multiple versions (or be version-agnostic) so that they can request and receive data effectively (ie, "speak the same language"). Some organizations may be able to maintain data in one format and dynamically convert to the necessary FHIR version on demand, but many organizations store data in the format (ie, version) in which it was received. In order for those organizations to make data available to other entities that require different versions of FHIR, they will likely have to manage multiple representations and corresponding mappings of the same information for different FHIR versions. This model does not scale well to support a large number of exchange partners and multiple versions of FHIR artifacts.

The *FAST* team has proposed a solution for how to handle version control more effectively to handle these scenarios, potential tooling to address them, and recommendations for what organizations can do to mitigate these issues.

In an ideal future state, relevant FHIR artifacts would be normative and any variation between FHIR releases would be focused on new functionality or edge cases. New FHIR versions would be backward compatible for all normative content, and all FHIR artifacts (eg, resources, profiles, bundles) would provide version information as part of any exchange. There would also need to be policies and tooling in place to support migration to new "floor" versions of FHIR as they evolve (ie, the minimum standard implementers must meet per regulation), such as a two-year window to sunset an old version, identification of any incompatible changes between new and old versions, and HL7 tools to translate between them.

In the interim, progress can be made toward these goals with improvements in resource version identification, capability statements (ie, documentation of the functionality supported for specific FHIR versions), and tooling such as authoritative mappings across versions. Since multiple versions of FHIR are currently in production, it will be important for organizations to be able to identify what version of FHIR their exchange partners are using so they may communicate using the same version or translate to the version supported by their exchange partner if needed. When an organization performs a directory lookup for an endpoint, all directory entries should include information regarding the FHIR version(s) supported. This requirement to support endpoint version is being accounted for in the *FAST* Endpoint Directory solution and will be incorporated into the HL7 Implementation Guide(s) being developed for the exchange of directory information. In addition, as the FHIR standard continues to evolve, organizations may support different functions at different times, making the capability statement an essential component to determine current endpoint support for specific versions and functionality. All endpoints will need to support the capability statement query and the FHIR \$versions operation that returns the supported version(s). While the FHIR Capability Statement resource is normative, there are elements included within it that are not, and so it will also need to be updated to ensure that it does not change in significant ways from FHIR release to release.

As new versions of FHIR get released with incompatible changes between versions, HL7 tools will be needed to handle translation or mapping between old and new FHIR versions. Implementers will need these translation mappings tools to reconcile any differences as they communicate with exchange partners using different FHIR versions.



Solution 2: A US Wide Methodology for Supporting Multiple Production VERSIONS of FHIR

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

There are multiple incompatible versions of FHIR in production (eg, DSTU2, STU3, R4) with breaking changes between them. Until FHIR becomes “normative” (ie, content is stable between versions), the industry will need a way to manage different representations of the same information in different FHIR versions.



SOLUTION

- Resource version identification:
 - standard requirements for resources, profiles, and bundles
 - directory metadata for endpoints
 - capability statements
 - testing and validation for conformance to directory metadata and capability statements
- Support for multiple versions:
 - ability to identify endpoint version
 - ability to identify FHIR artifact version
 - ability to translate versions (at least from prior to new version)
 - Ability to document translations where appropriate (eg, provenance)



OPEN ITEMS

- Collaborate with HL7 FHIR leadership to ensure alignment with FHIR standard release plans and impact on extensions, profiles and Implementation Guides
- Identify impact on current and future ONC and CMS regulations
- Determine ability to translate non-normative resources
- Determine how version management works over time in response to new data portability requirements.
- Need to consider proliferation of Implementation Guides and Profiles that provide for different solutions to the same fundamental use case
- Consider incompatible profile constraints on underlying resource in ways that do not permit reuse by other implementation guides (eg, US Core constraints that do not support specific IGs requirements)



IN SCOPE

Managing multiple versions of FHIR and FHIR artifacts such as implementation guides, identification of supported version for a specific endpoint, transform/ translation service considerations, and the ability to appropriately manage exchange of information across multiple versions of FHIR.



OUT OF SCOPE

Specifying a single version of FHIR, requiring forward/backward compatibility for non-normative resources, addressing support for multiple versions in a single exchange.



BENEFIT

- Clarity for implementers on version management methodology
- Efficiency through establishment of authoritative source for transforms/mappings across versions
- Improved reliability of FHIR message processing



STATUS

Finalizing V2 solution documentation to obtain SME input



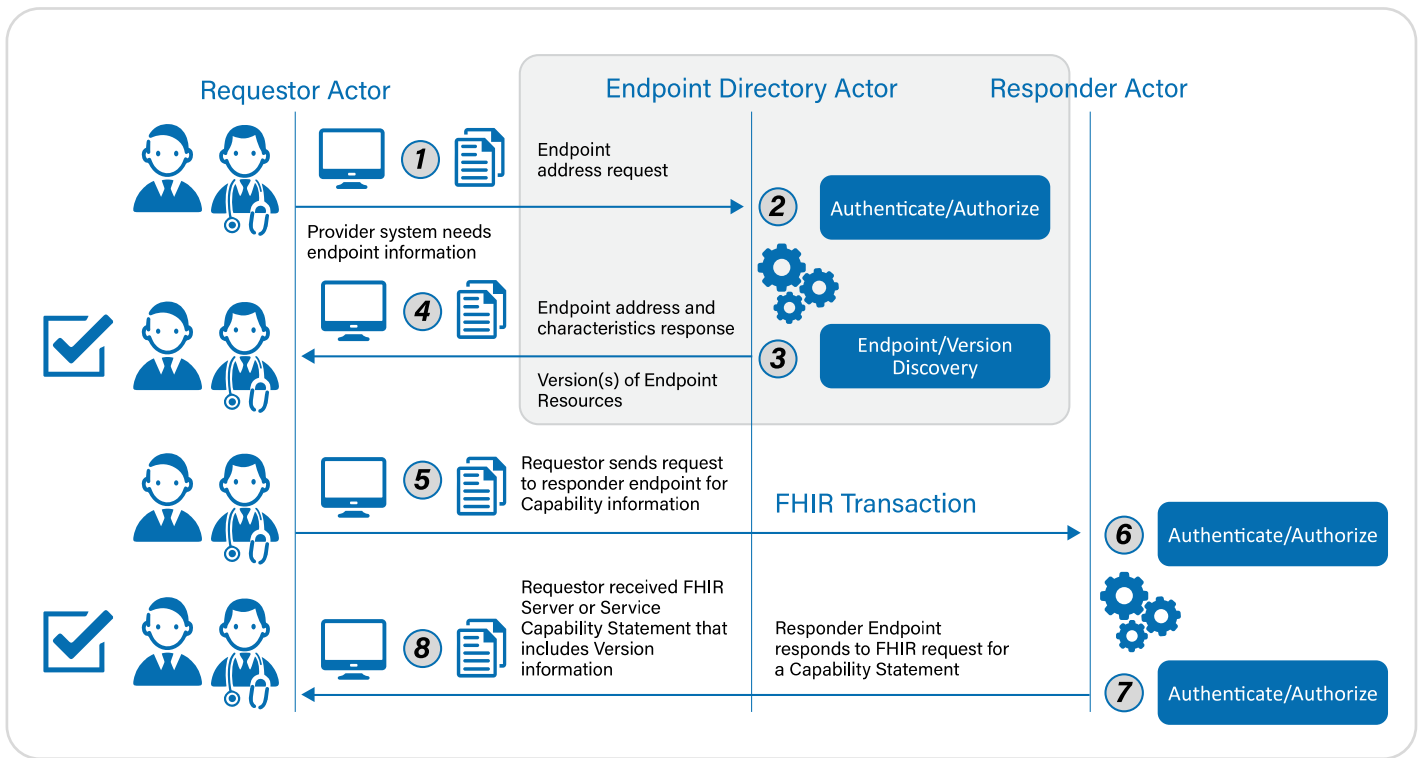
CURRENT SOLUTION

[Versioning Solution Document \(Draft V2\)](#)



Solution 2: A US Wide Methodology for Supporting Multiple Production VERSIONS of FHIR

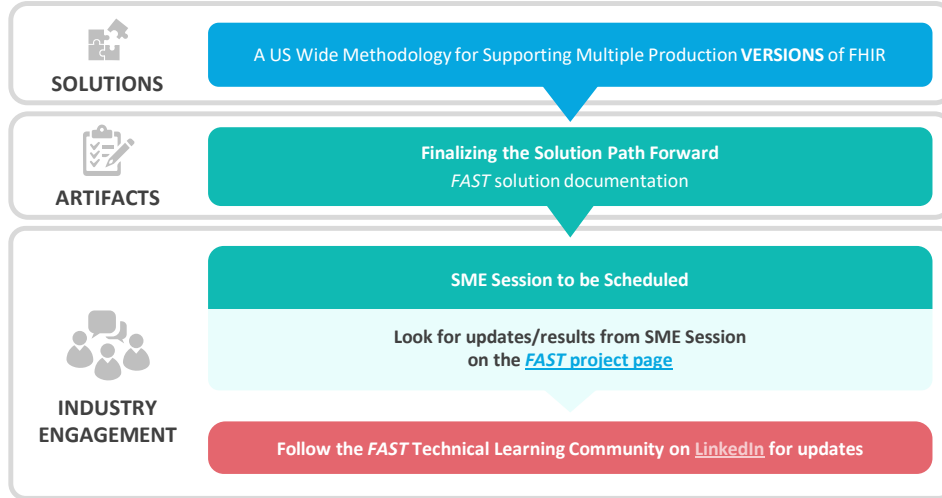
In the following diagram an entity, referred to as a requestor actor, needs to obtain information from another entity, known as the responder actor. The requestor actor initiates a request for an endpoint from the endpoint directory if not already known. The endpoint response contains metadata indicating the version(s) of FHIR supported by the endpoint and any authentication and authorization requirements. After completing the necessary authentication and authorization steps, the requestor actor requests and receives the capability statement from the responder actor which includes details on the version(s) of FHIR and Implementation Guides supported.





Solution 2: A US Wide Methodology for Supporting Multiple Production VERSIONS of FHIR

Next Steps and Industry Impact



The next steps for this solution are to:

Obtain SME input on the proposed solution



Next Steps

This solution will have the biggest impact on:

Payers, providers, health systems, EHRs, public health, application vendors, and intermediaries



Impacts

This solution offers these key benefits:

- Support adherence to the CMS proposed rules on reducing provider and patient burden by improving prior authorization processes and promoting patient's Electronic Access to Health Information
- Streamline and support value-based care workflows
- Increase scalability and improve interoperability between exchange partners
- Support for vendor product development and provides a venue for addressing FHIR version control issues



Key Benefits

Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Join the [FAST Technical Learning Community](#) on LinkedIn and stay tuned for updates coming out of the SME Session
- Visit the [FAST project page](#) for more information



Stakeholders



Solution 3: US-Wide SCALING Requirements for FHIR RESTful Exchange Intermediaries

As FHIR scales and the number of participants exchanging data grows, intermediaries (eg, clearinghouses, HIEs, or any entity that participates in the exchange of a FHIR-based transaction other than the ultimate Requestor and Responder) will wind up sitting in the middle of transactions for a variety of purposes, such as routing, value-added services (eg, endpoint resolution, patient matching, record location, version translation, error handling), and acting on behalf of organizations to become their endpoint. These intermediary implementations need to be complementary to, and supportive of, the end-to-end transaction requirements for FHIR. In order to achieve interoperability, rules of the road need to be defined for their participation in the healthcare ecosystem.

The existence of hybrid exchange models (eg, spoke/hub, direct connections/point-to-point, and regionally interconnected spoke/hub) create challenges in adopting standards for scaling FHIR and implementing consistent approaches for authentication, endpoint detection, standards for matching, and end-to-end performance. Additionally, providers and payers increasingly need to support real-time transactions embedded in the clinical workflow, which in many cases require a response prior to providers proceeding with diagnosis or treatment. Many intermediary models do not support these end-to-end synchronous real-time transactions, leading to lack of predictability, reliability, and consistent response times.

To address these challenges, the *FAST* team is proposing a solution that will support a hybrid environment moving forward, including point-to-point, gateway, and intermediary models. Minimum availability and performance requirements need to be established for any scale architecture including one or more intermediaries, with the requirement to support synchronous transactions and maintain “state” across intermediaries (ie, systems remember preceding events that occurred across multiple systems or actors). Intermediaries will need to support all FHIR workflow operations and be capable of handling volume, response time, and routing to all available endpoints. Consistent support of metadata is also needed for routing through multiple intermediaries. Finally, testing will validate conformance with these exchange and Service Level Agreement (SLA) requirements.

The goal is for ubiquitous access to permitted endpoints regardless of the architecture (ie, point-to-point, one intermediary, or multiple intermediaries). The performance reliability and availability characteristics should be substantially the same for intermediary connections and point-to-point connections and must be acceptable for real-time information exchange where there is a provider/patient waiting for the response before clinical workflow can continue. The ultimate application user, system, and API endpoint should see the rest of the world as a consistent set of endpoints, regardless of whether there is one or more intermediaries sitting in the middle of the exchange.



Solution 3: US-Wide SCALING Requirements for FHIR RESTful Exchange Intermediaries

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

Intermediaries will be needed to scale FHIR across a growing number of healthcare ecosystem participants in anticipation of future healthcare needs and increased transactions embedded in clinical workflows. Yet there are no rules of the road for intermediaries to meet minimum requirements for performance, availability, response times, or specific FHIR operations in support of synchronous end-to-end real-time data exchange



SOLUTION

- Ability to utilize intermediaries to reduce the complexity of connecting with a large number of endpoints and managing authentication and authorization with each
- Intermediaries will have predictable performance required to meet real-time exchanges
 - FHIR-enabled intermediaries
 - Standard for intermediary performance (SLAs)
 - Testing for conformance with exchange and SLAs
- Intermediaries may provide value-added services to assist in patient matching, version translation, etc.



OPEN ITEMS

- Determine scope of requirements
- Explore standards for intermediary-to-intermediary exchanges
- Clarify exchange services that must be supported
- Detail the specific availability and performance requirements
- Determine the best method for establishing requirements



IN SCOPE

- Interoperability models with, point-to-point, single and multiple intermediaries
- Issues related to RESTful FHIR exchanges/messages and related technologies (like CDS Hooks) via intermediaries
- Related authentication and authorization models (OAuth, OpenID, UDAP)
- Planning for future volume increase
- Establishing SLA and Performance requirements for intermediaries and endpoints
- Establishing functionality of endpoints and the method of declaration



BENEFIT

- Predictable end-end performance for FHIR RESTful exchanges
- Availability of FHIR endpoints regardless of the specific exchange architecture



OUT OF SCOPE

- Identification, security, directory, versioning, metadata, certification or piloting
- Ownership models
- Trust frameworks
- Legal agreements
- Non-RESTful exchange methods (eg, Direct)
- Technical Implementation



STATUS

Finalizing V2 solution documentation to obtain SME input



CURRENT SOLUTION

[FAST Scaling requirements for FHIR RESTful Exchange Intermediaries solution document \(V3 in progress\)](#)



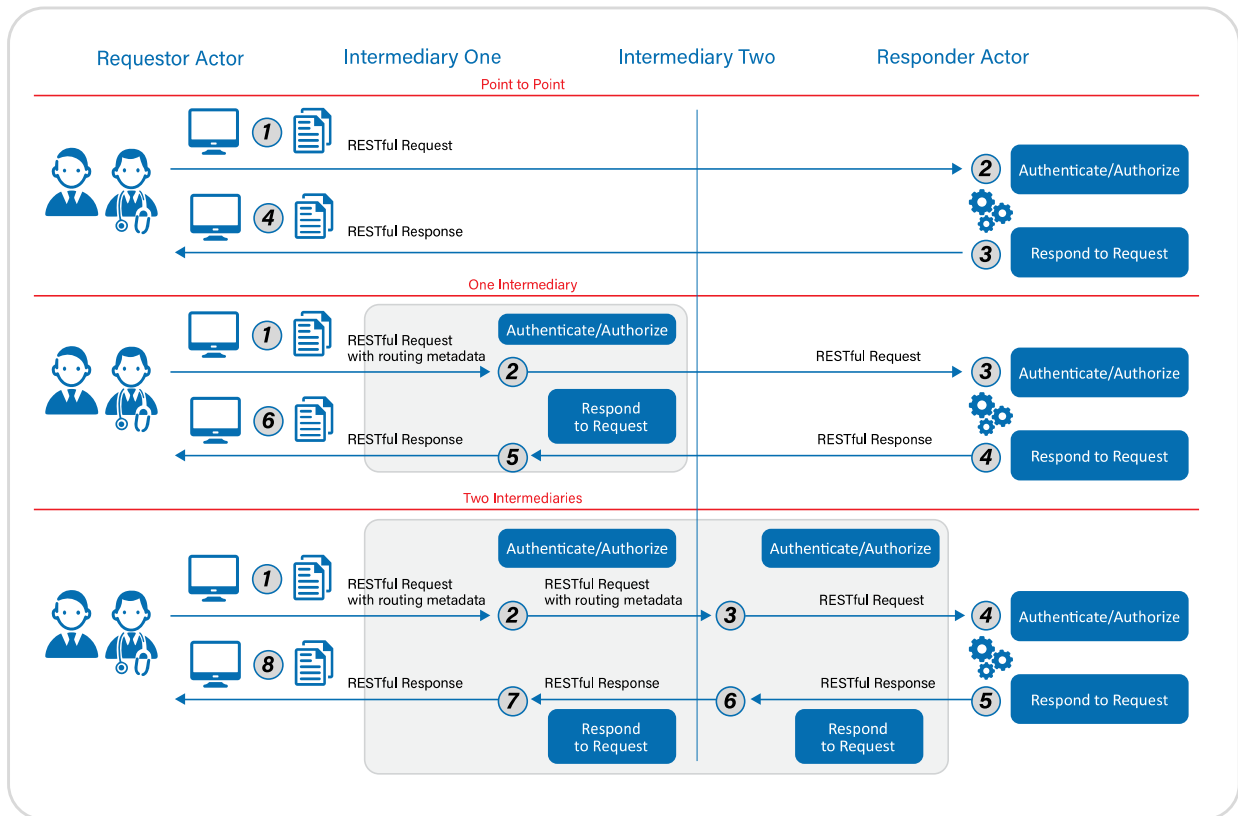
Solution 3: US-Wide SCALING Requirements for FHIR RESTful Exchange Intermediaries

Assuming the endpoint has already been determined via directory query, the following diagram represents 3 different exchange models that need to be supported in a mixed model environment with full connectivity.

Point-to-point: The Requestor uses the endpoint directory information to connect to the Responder endpoint and send the request. The Responder authenticates the Requestor, processes the request, and returns the results. The Requestor receives the response back from the Responder.

One intermediary: The Requestor uses the endpoint directory information to connect to the Intermediary endpoint and send the request including routing information. The Intermediary uses the routing information to connect to the Responder endpoint and sends the request. The Responder authenticates the Requestor/Intermediary, processes the request, and returns the results to the Intermediary. The Intermediary then returns the response to the Requestor and the Requestor receives the response from the intermediary.

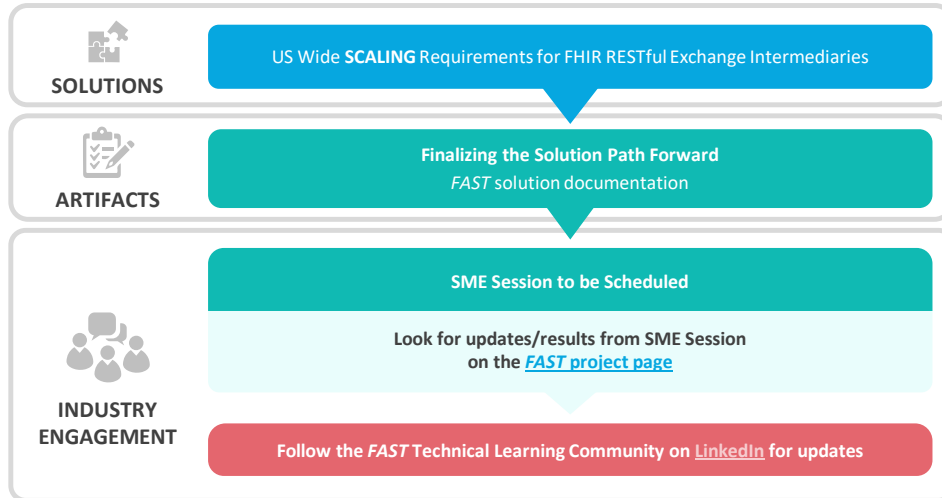
Two (or more) intermediaries: The Requestor uses the endpoint directory information to connect to Intermediary One's endpoint and send the request including routing information. Intermediary One uses endpoint directory information to connect to Intermediary Two's endpoint supporting the Responder and forwards the request including routing information. Intermediary Two uses the routing information to connect to the Responder endpoint and sends the request to the Responder. The Responder authenticates the Requestor/Intermediary, processes the request, and returns the results to Intermediary Two. Intermediary Two sends the response to Intermediary One and Intermediary One returns the response to the Requestor. The Requestor receives the response from Intermediary One.





Solution 3: US-Wide SCALING Requirements for FHIR RESTful Exchange Intermediaries

Next Steps and Industry Impact



The next steps for this solution are to:

Obtain SME input on the proposed solution



Next Steps

This solution will have the biggest impact on:

Payers, providers, health systems, EHRs, public health, application vendors, and intermediaries



Impacts

This solution offers these key benefits:

- Support adherence to the CMS proposed rules on reducing provider and patient burden by improving prior authorization processes and promoting patient's Electronic Access to Health Information
- Streamline and support value-based care workflows
- Increase scalability and improve interoperability between exchange partners



Key Benefits

Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Join the [FAST Technical Learning Community](#) on LinkedIn and stay tuned for updates coming out of the SME Session
- Visit the [FAST project page](#) for more information



Stakeholders



Solution 4: Standards-Based Approaches for Individual IDENTITY Management

Patient matching and identity management have been a challenge for many years in healthcare and are major barriers to achieving interoperability. There are several barriers to accurate and consistent patient matching across the healthcare ecosystem. For example, different organizations may rely on different types of patient identifiers such as medical record numbers or insurance IDs, but these identifiers often have limited value outside the walls of that organization. Mapping various identifiers across organizations is not a scalable approach. Another barrier to accurate patient matching and identity management is data quality and accuracy.

Patient record matching is an imprecise science, and while much has improved over the years to decrease the rate of false positive and false negative matches, more can be done to improve the accuracy and reliability of any method. Complementing patient matching techniques is a construct known as digital identity. NIST defines digital identity as “the unique representation of a subject engaged in an online transaction.” Digital identity is still in its infancy in healthcare, and no national strategy exists to date that would support widespread adoption of digital identity technologies and approaches in a coordinated and interoperable way.

Patient matching and digital identity are distinct but complementary processes.

- Patient matching provides a degree of certainty that 2 or more records should be affiliated with one another as they are likely to belong to the same person
 - Patient matching does not confirm the identity of the individual
- Digital identity, comprised of identity assurance and authentication, provides confidence that the individual is who they claim to be and they are the individual engaged in the transaction
 - Use of digital identity reduces complete reliance on traditional matching processes when querying patients across disparate systems to find other instances of records for this same individual

As APIs make data more fluid and regulation specifically requires capabilities that include patient participation, there is an opportunity to address some of these challenges and bring confidence to how the identity required by regulation can be best put to use. The Identity Tiger Team approached their solutions development with these assumptions:

- Different stakeholders have different technical and operational needs and capacities, so a one-size-fits-all approach for the industry as a whole would not be effective
- Because the United States still does not have a single, universal unique patient identifier or a national unified approach to identity and matching, healthcare will need to leverage multiple identity services and approaches to consistently identify or match patients while also managing appropriate access by only authorized requesters



Solution 4: Standards-Based Approaches for Individual IDENTITY Management

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

The industry currently employs a range of patient matching and identity management processes with inconsistencies and limited scalability beyond the walls of an organization



SOLUTION

Establish a set of patient matching and identity management patterns and best practices that the industry can adopt to reduce the variations that exist today and provide a bridge to incremental advancement and innovation



STATUS

Incorporating feedback from industry stakeholders and planning next steps for packaging up these solutions in a way they can be adopted by the community. Read the full SME session report here.



CURRENT SOLUTION

[FAST Identity proposed solution document \(version 3 in progress\)](#)



IN SCOPE

Establish a set of patient matching and identity management patterns and best practices that the industry can adopt to reduce the variations that exist today and provide a bridge to new approaches in the future

- Patient matching during payer/provider interactions: Mediated Patient Matching
- Patient-directed workflows focusing on identity management: Networked and (later) Distributed Identity Mgmt.



OUT OF SCOPE

Contractual arrangements, grammar for communicating consents, level of identity verification for certain attributes, and for access grants/delegation. (Security and directory considerations are addressed by other FAST solutions.)



OPEN ITEMS

- Pursue provider and payer identity management
- Apply proposed solutions to use cases
- Capture patient matching recommendations and best practices
- Explore preliminary steps needed to adopt solution 4, Distributed Identity Management
- Consider how regulation/policy might address challenges that can't be solved by the market
- Address consent and privacy



BENEFIT

- Improve patient safety in clinical data exchanges across organizations
- Fewer credentials for patients and organizational requesters and responders to each manage for themselves and each other
- Improved match accuracy
- Improved confidence about conditions when results will be returned
- Establish clear expectations of safe harbor for requesters and responders, when matching errors occur
- Stronger identity assurance as individual credentials assume a broader role in health IT
- Framework that can deliver a longitudinal patient record, supporting best possible care

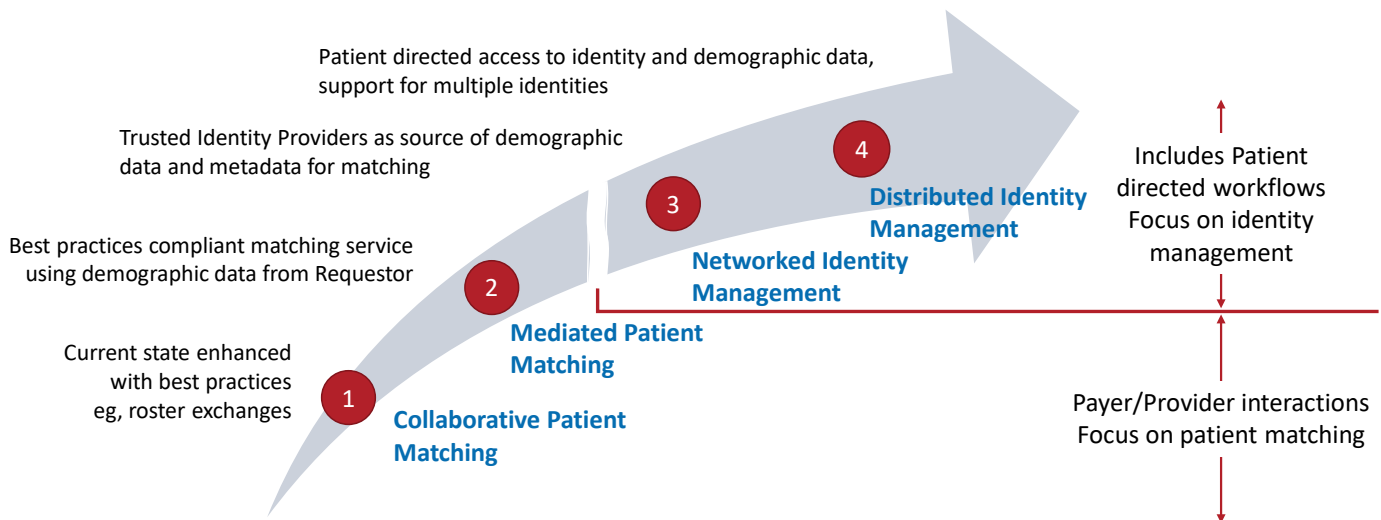


Solution 4: Standards-Based Approaches for Individual IDENTITY Management

The FAST Identity Tiger Team defined a spectrum of solution options ranging from lower complexity and less scalable to higher complexity and more scalable, with the understanding that different stakeholders have different technical and operational needs and capacities. These options are meant to offer flexibility while also solving key barriers to consistent and accurate patient matching and identity management.

The solutions options in the figure below span from solutions 1 and 2 that focus on patient matching and primarily support payer/provider interactions, while solutions 3 and 4 includes, but are not limited to patient-directed workflows and focus on identity management. The progressive path from solutions 1–2 to 3–4 also depicts increasing complexity and technical maturity needed for implementations at scale.

Solution Options: Lower to Higher Complexity



Collaborative Patient Matching relies on use of common or agreed upon patient identifiers known to both requesting and responding entities. This solution pattern closely mimics how certain entities across the industry match patient records today, enhanced with recommendations and best practices to help ensure the success of any implementation that allows this pattern.

This solution applies to scenarios where 2 parties have agreed to use a common identifier or mutually known identifiers. The FAST Identity team recommendations and best practices may be applied to address gaps that currently exist within this type of exchange.

The Requestor and Responder actors can be either a provider or a payer as the solution applies to provider-to-provider, provider-to-payer, and payer-to-payer transactions.

The patient's identity in this scenario is established by the Requester and Responder actors as part of their onboarding process that includes appropriate attribute validation. One of the reasons why collaborative approaches are not scalable is because it becomes impossible to share and manage ever-changing patients lists beyond a 1:1 or 1:few exchange relationship.



Solution 4: Standards-Based Approaches for Individual IDENTITY Management

Mediated Patient Matching represents both an incremental advancement from collaborative patient matching along the technical complexity and maturity pathway, as well as a point of entry along the pathway for those systems with the operational and technical capacity to begin efforts toward a long-term strategy to support accurate matching and digital identity. Mediated patient matching goes one step further from common approaches to suggest a matching service using demographic data at the time of each request. This service would be *FAST* best practices compliant and would match patients for end users at the point of exchange.

This solution covers patient matching in near real time, leveraging FHIR transactions to and from the matching service. The Requestor and Responder actor pairs may be represented by provider/provider, provider/payer, and payer/payer exchange scenarios. When 2 entities exchange data, the Requestor is responsible for sending the minimum required patient demographic data to be used for matching, and the Responder is accountable for matching identities of the patient(s) involved using the matching service.

Networked Identity Management is a leap in technology requirements. It builds on Mediated Patient Matching concepts but recommends moving from relying entirely on demographics-based matching to a digital identity model. This solution would use a network of trusted OpenID Connect providers, similar to the identifiers within the patient application access models within Cures Update and CMS rules the industry is in the process of deploying. In the scope of this solution parties would rely on the OpenID Connect provider's digital certificate and an OpenID identifier assigned to each user as part of an onboarding process that includes identity proofing and establishes their real-world identity and the validity of demographics made available for patient or provider matching. In this scenario responders would validate the trustworthiness of the associated identity provider via its digital certificate and use the recorded patient or provider OpenID identifier and other verified user profile data to match on the identifier or, if the identifier is not yet known in their system, fall back to a demographics-based matching process or other interrogation of the OpenID provider or the holder of the identifier.

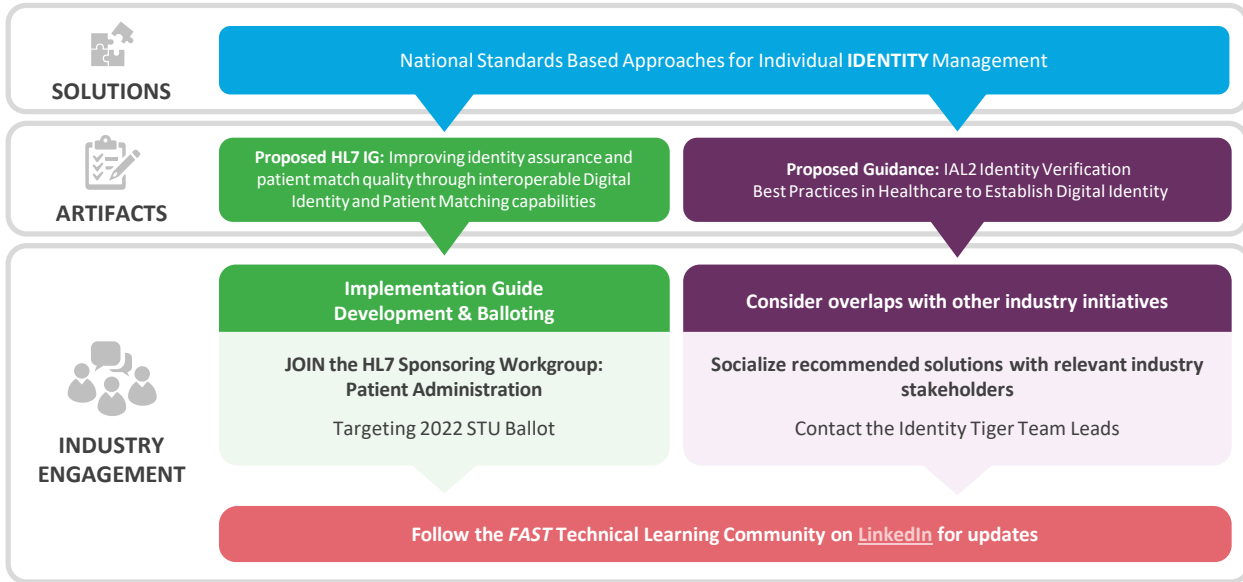
Distributed Identity Management is a solution that would give the patient the responsibility and complete control of their identity and assignment of permissions for entities to access their health information. Trusted third-party identity services would assign and manage patient identities and associated identifiers. During a patient information request, the Requester relies on the identity service to validate identity claims made by the patient or Requester, for use by the Responder in making its authentication or authorization decision. The solution also positions the industry to explore the concept of multiple digital identities with different verifiable claims for the same patient for potential use as consent profiles. The mechanisms for digital identity management are still being defined by the industry and processes for using them in a privacy-preserving, trusted/federated, and secure healthcare context are only just emerging.

These 4 solutions allow for incremental improvements in patient matching and identity management as technology matures and adoption increases.



Solution 4: Standards-Based Approaches for Individual IDENTITY Management

Next Steps and Industry Impact



The next steps for this solution are to:

- Create and publish IAL2 Identity Verification Best Practices in Healthcare to Establish Digital Identity
- Develop HL7 Implementation Guide: FHIR \$match Query for Cross-organizational (B2B) Exchange



This solution will have the biggest impact on:

payers, providers, health systems, EHRs, public health, app vendors



Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Join sponsored HL7 Work Group once identified
- Provide input and ballot of implementation guide
- Provide input on proposed guidance
- Help identify any overlaps with other industry initiatives
- Participate in HL7 Connectathons
- Participate in pilots



Stakeholders

This solution offers these key benefits:

- Supports improved patient safety through more accurate data exchange across organizations
- Reducing burden through
 - Consistent and accurate patient matching
 - Leveraging standards to minimize the need for unique solutions between exchange partners
- Support adherence to the CMS-proposed rule on reducing provider and patient burden by improving prior authorization processes and promoting patient's Electronic Access to Health Information
- Streamline and support VBC workflows
- Increase scalability and improve interoperability between exchange partners



Key Benefits



Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

Security in healthcare faces unique challenges due to the increasing adoption of the FHIR standard for improved healthcare interoperability, and the growing interest in implementing FHIR-based APIs to facilitate access to health data for clinical care and health services delivery. Security solutions needed to be re-evaluated with regards to their scalability, and actions taken to address any identified gaps. The *FAST* Scalable Security solutions are intended to address potential current limitations of FHIR-based information exchange to appropriately authenticate users (confirm they are who they say they are) and authorize users (confirm they have required permissions) to see the data requested.

As the FHIR ecosystem grows and the number of deployed servers and clients multiplies, several aspects of the registration, authentication, and authorization processes that must occur before FHIR resources can be exchanged are potential bottlenecks. To facilitate the effective scaling of the ecosystem, automated approaches for application registration and more robust mechanisms to reliably identify participants and manage credentials are needed. For larger ecosystems with numerous requestors and responders, a distributed system of authoritative information can be leveraged using digital certificates to enable a scalable dynamic solution for client (ie, FHIR client/Requestor) registration.

The registration problem alone is exemplified by statistics shared from a scenario analysis that considered manual registration of 60 current Blue Button API clients across 907 US Health Plans. Extrapolating from the CMS experience, the time required by human facilitated administrative activities at the Health Plans to register 60 applications (eg, review meetings, actual generation and sharing of API credentials with app developer, etc.) was estimated at 73 person-years.⁵ This estimate only addresses one type of registration interaction (payer/consumer), with additional registration effort required for all provider/consumer, provider/provider, provider/payer, and payer/payer pairings. The healthcare dollars expended in non-value, manual-related client registration activities can be recaptured many times over as automated registration using digital solutions is adopted across the FHIR API ecosystem nationwide.

The ONC FHIR at Scale Taskforce's Security Tiger Team was formed in late 2018 to investigate these issues and identify potential solutions. The Tiger Team identified the components below as building blocks to be used by implementers to address the issues above and enhance the overall scalability of the FHIR ecosystem:

- The Unified Data Access Profiles (UDAP)^{6,7} for dynamic client registration (which empowers trust community actors and client apps to get OAuth credentials from one another in a scalable fashion)
- Client authentication (which allows network actors to identify software components to one another)
- Client authorization (which allows the relying party/Responder to make decisions on what data to release)
- Tiered OAuth (which allows for reusability of user credentials across participants in the network where appropriate)

The recommended solutions leverage and build upon work already done in the industry and extends it to achieve the intended outcome in a FHIR-enabled API healthcare environment that scales nationwide.

5. Seib A, Scrimshire M. "Making it easier for Patients and Data Holders," EHNAC AHIP Webinar, 2020.

6. UDAP Implementation Guide for Registration and Authorization of Business-to-Business Health Apps. [UDAP.org. https://www.udap.org/udap-ig-b2b-health-apps.html](https://www.udap.org/udap-ig-b2b-health-apps.html). Published November 9, 2020. Accessed March 4, 2021.

7. UDAP Implementation Guide for Registration and Authorization of Consumer Facing Health Apps. [UDAP.org. https://www.udap.org/udap-ig-consumer-facing-health-apps.html](https://www.udap.org/udap-ig-consumer-facing-health-apps.html). Published November 8, 2020. Accessed March 4, 2021.



Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

Today, we have limitations on our ability to ensure, in a scalable way, that the Requestor of information using a FHIR-based information exchange is appropriately authenticated and has the authorization to see the data requested. Current registration processes are manual and too time-consuming to support expected growth



SOLUTION

Leverage existing credentials and authorizations and best practice standards to establish common security processes that facilitate automated exchange and reuse existing infrastructure where possible



IN SCOPE

- Trusted dynamic client registration using UDAP
- JWT-based client authentication and authorization



OUT OF SCOPE

Directory for endpoint discovery, trust policy governance, requirements for a specific architecture, patient/provider, or provider/patient



STATUS

Incorporating feedback from industry stakeholders. Read the full SME session report [here](#)



OPEN ITEMS

Cross-solution overlaps, explore standard authorization metadata requirements, recommendations related to privacy



CURRENT SOLUTION

[FAST Security proposed solution document \(version 3 draft\)](#)



BENEFIT

Increased confidence and consistency and reduced administrative burden

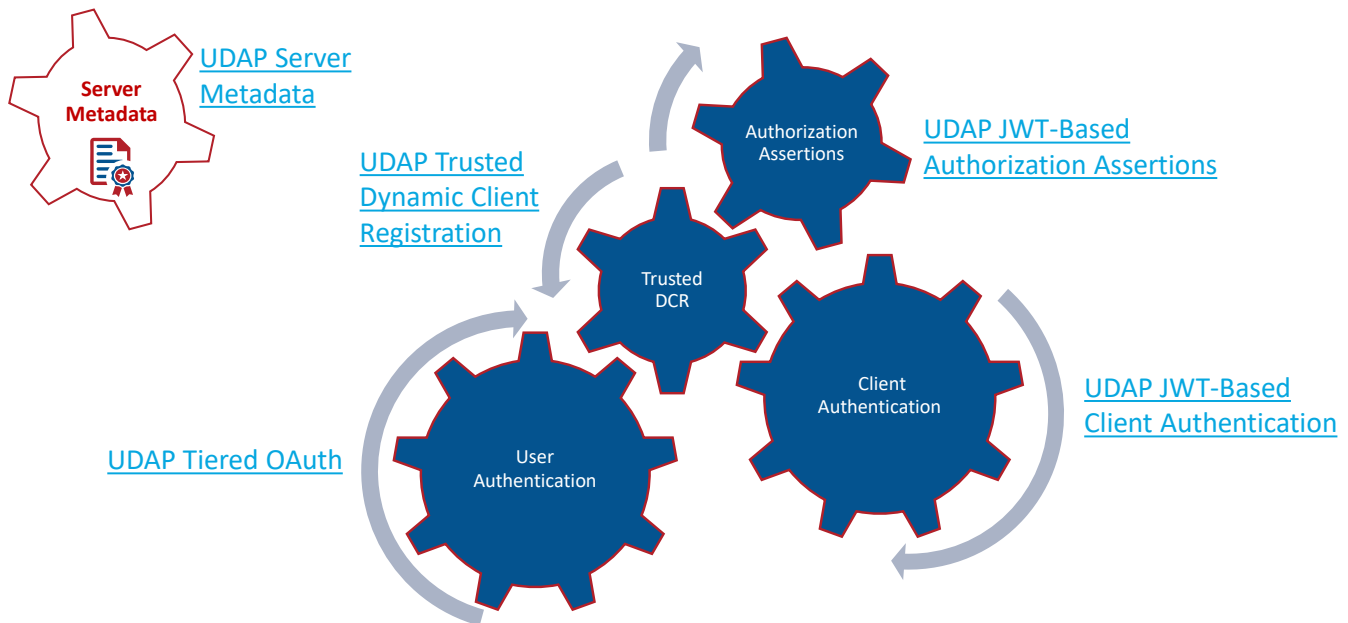


Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

The Trusted Ecosystem proposed by the *FAST* Security Tiger Team would automate existing manual process leveraging existing trust frameworks. These proposed solutions are not about building new trust frameworks, but instead are about networking existing trust frameworks into a larger trust community. These solutions leverage existing standards that have proven effective, avoiding a need to rip and replace. The intention is to leverage existing trust frameworks and standards to better support healthcare IT use cases—including those that apply to FHIR APIs on a nationwide scale. The infrastructure is intended to be robust and reusable, so as new use cases arise no additional development would be required because key features already built into the framework can be enabled. These solutions propose reusing authorization and authentication across the trust community to minimize time required to exchange data with a new endpoint and still be confident that the requester of data is who they say they are and have the appropriate permissions to see the information they are requesting. The solutions also enable requests to carry additional information—such as organization affiliations, certifications, and additional context elements—that data owners can use to make more informed authorization decisions.

The proposed solution components would allow the FHIR community to implement:

- Integration of existing public key infrastructure mechanisms with registration, authentication, and authorization processes to establish robust trust networks with reusable credentials to identify actors
- Trusted dynamic client registration
- Client app submissions of self-assertions, third-party certifications, or other endorsements to servers and vice-versa
- Client app assertions of additional information for a given session so that resource holders can more finely scope access tokens, including information related to consent or purpose of use
- Increase security and assurance in identity of all actors by using asymmetric cryptographic methods for authentication, including specific protocols to support network-wide revocation of credentials
- Dynamic federation of user credentials to facilitate reuse of credentials and single sign-on



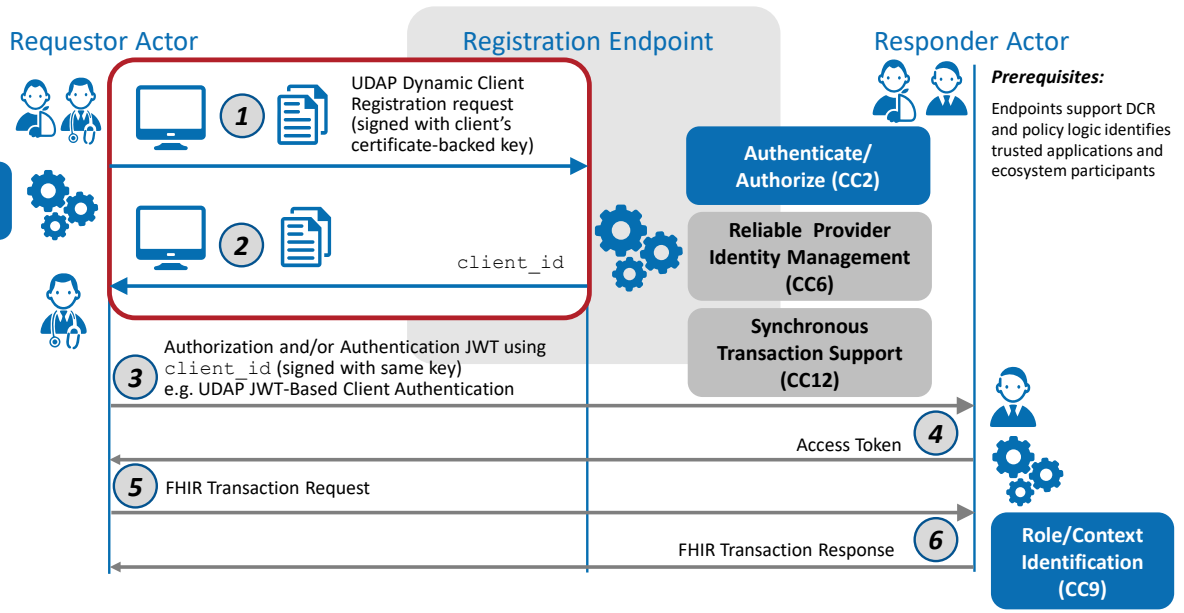


Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

Prerequisites:

- Client app obtains a trusted digital certificate and end user at requestor organization is authorized to use the app
- Client app may examine UDAP Server Metadata before attempting a formal request

Data Provenance (CC5)



■ In Scope
■ Out of Scope

Complexity Rating

Medium: Extends existing OAuth 2.0 specifications; reuses existing technologies to validate JWTs and X.509 Certificates

ID	Description	Notes
1	Client app requests registration with endpoint	UDAP client registration signed with client's certificate-backed key
2	Endpoint validates registration and returns client_id	Registration request may encompass certifications and endorsements. Used to communicate information about an app from trusted Endorsers (ie, App Certifiers); see UDAP Certifications and Endorsements



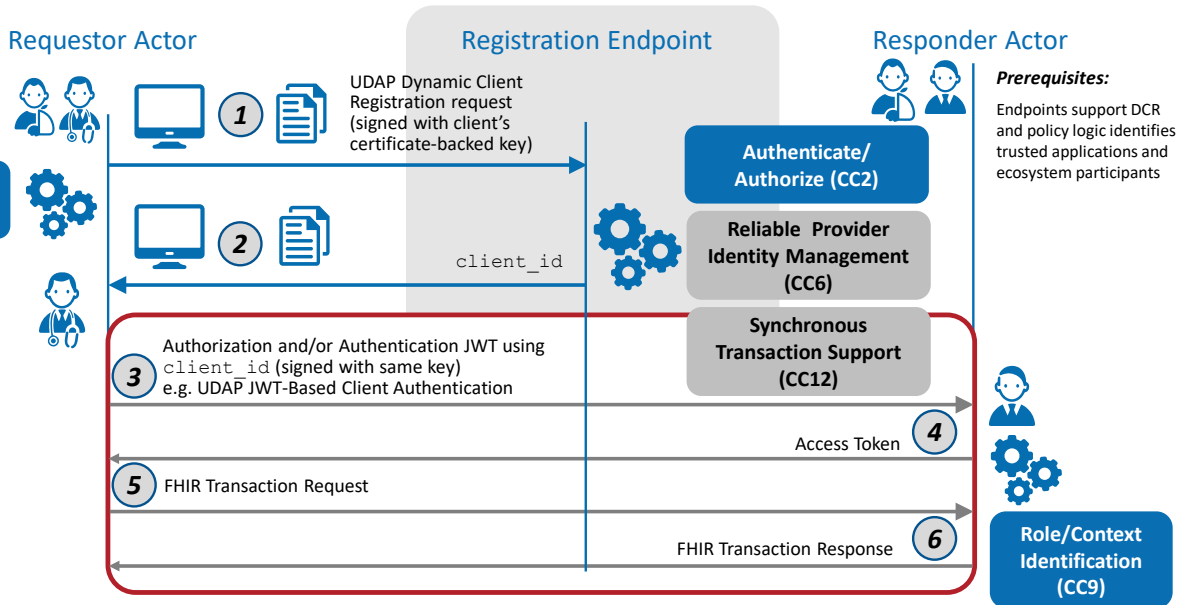
Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

Prerequisites:

- Client app obtains a trusted digital certificate and end user at requestor organization is authorized to use the app
- Client app may examine UDAP Server Metadata before attempting a formal request

Data Provenance (CC5)

In Scope
Out of Scope



Complexity Rating

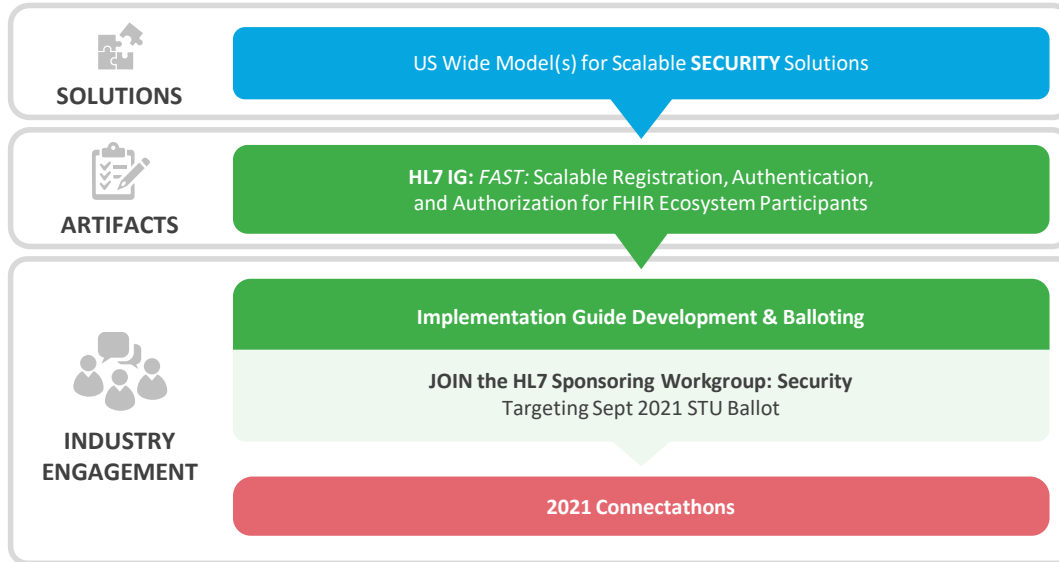
Medium: Extends existing OAuth 2.0 specifications; reuses existing technologies to validate JWTs and X.509 Certificates

ID	Description	Notes
3	Authenticate and/or Authorize using client ID	UDAP JWT-based Client Authentication
4	Return Access Token to Client	
5	FHIR Transaction Request	Steps 5 and 6 in the flow diagram are typical capabilities already in place and supported by a FHIR endpoint (ie, Responder) such as...
6	FHIR Transaction Response	(5) – Requestor includes access token in the FHIR transaction request (6) – FHIR endpoint (Responder) receives access token then executes local policy logic regarding the request; returns requested information (or denies access)
6	FHIR Transaction Response	



Solution 5: US-Wide Model(s) for Scalable SECURITY Solutions

Next Steps and Industry Impact



The next steps for this solution are to:

Develop HL7 Implementation Guide: *FAST*: Scalable Registration, Authentication, and Authorization for FHIR Ecosystem Participants



This solution offers these key benefits:

- Increased confidence in request authorization decisions
- Leveraging already existing trust frameworks makes adoption more approachable



This solution will have the biggest impact on:

payers, providers, health systems, EHRs, public health, app vendors



Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Joining the HL7 Security Work Group
- Providing input and ballot of implementation guide
- Participating in HL7 Connectathons
- Participating in pilots





Solution 6: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)

In today's environment, FHIR integration is typically point-to-point without the need for routing information. In this model, it is usually known who the Requestor and Responder are without the need for multi-hop routing data. As FHIR scales and the number of actors grow, a hybrid environment will exist where dynamic point-to-point as well as intermediary or multi-intermediary models (eg, clearinghouses, HIEs, national networks, and other exchanges) are being used, and routing information will become more important.

The Exchange Tiger Team recognized that for the foreseeable future, the healthcare industry will be made up of a complex ecosystem of a variety of actors with a need to exchange data. The healthcare ecosystem is operating in a hybrid environment, meaning transactions may consist of dynamic point-to-point and intermediary models. The team understood that whatever solution was developed would need to provide reliable data exchange regardless of whether the transaction is synchronous or asynchronous, what identifiers are used or how many "hops" the data may make before reaching its final destination.

As the need for integration between different actors in healthcare has grown, transaction routing across one or more intermediaries is recognized as a key aspect of reliable exchange. One example scenario is when a payer uses a clearinghouse intermediary as their gateway for transactions. There are both technical and business operational value-adds in this intermediary model. This model was born in the world of the original X12 transaction set and it is expected to continue in the evolving RESTful FHIR API integration model. Other networks including HIEs and national networks have emerged as brokering intermediaries (document access/exchange, e-prescribing, etc.) that may also engage in FHIR-based interoperability. In the model described above, the interaction originator will know the final destination but will not need to be concerned with whether intermediaries are involved in the message routing. The intermediary, or intermediaries, will need to have origination and routing information available during the life cycle of the transaction to ensure appropriate delivery. The solution that will be described later in this section leverages the use of RESTful headers to carry the routing information along with a transaction, to ensure reliable exchange regardless of whether the transaction is dynamic point-to-point or via intermediaries.



Solution 6: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

FHIR information exchange is typically performed point to point between 2 trusted system endpoints. Because healthcare participants may also wish to leverage intermediaries in FHIR exchanges, a solution for conveying routing information is needed



SOLUTION

Employ RESTful header parameters to send originator and destination information for use by exchange intermediaries



IN SCOPE

Exchange using intermediaries in addition to point-to-point connections
Method for exchanging routing information within HTTP REST headers



OUT OF SCOPE

- Value set defining exchange identifiers
- Capturing provenance information from exchange through multiple intermediary “hops”
- Security considerations are being addressed by the FAST Security Tiger Team through a separate implementation guide



STATUS

Incorporated feedback from industry stakeholders. Read the full SME session report [here](#). Connectathon testing and implementation guide development through the HL7 balloting process.



OPEN ITEMS

- Development of HL7 Exchange with or without Intermediaries Implementation Guide
- Continue to test solution at upcoming HL7 FHIR Connectathons
- Complete ballot cycle



CURRENT SOLUTION

[FAST Exchange solution document \(version 3\)](#)



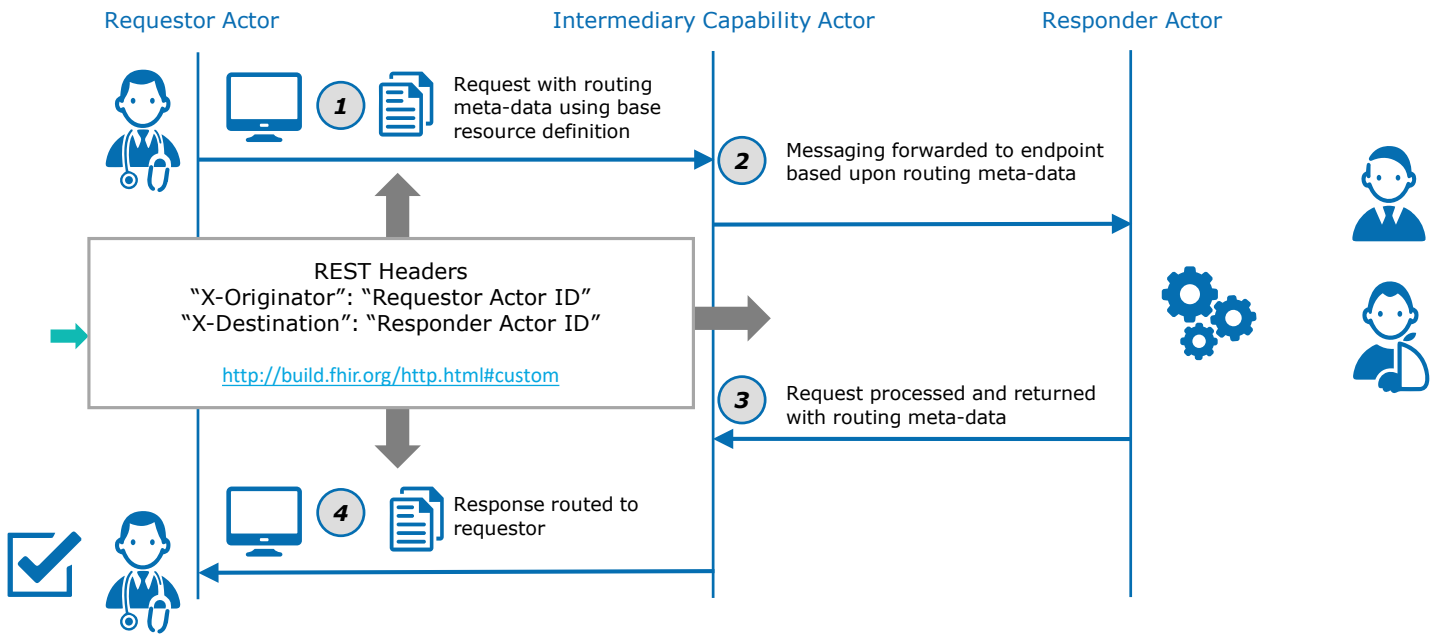
BENEFIT

- Consistent and accurate routing
- Common pattern, used for many years in healthcare and other industries
- Lightweight
- Even when there is no FHIR resource being exchanged (eg, searches and matches), routing information is still available
- Universally usable, regardless of FHIR transaction—it’s resource agnostic



Solution 6: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)

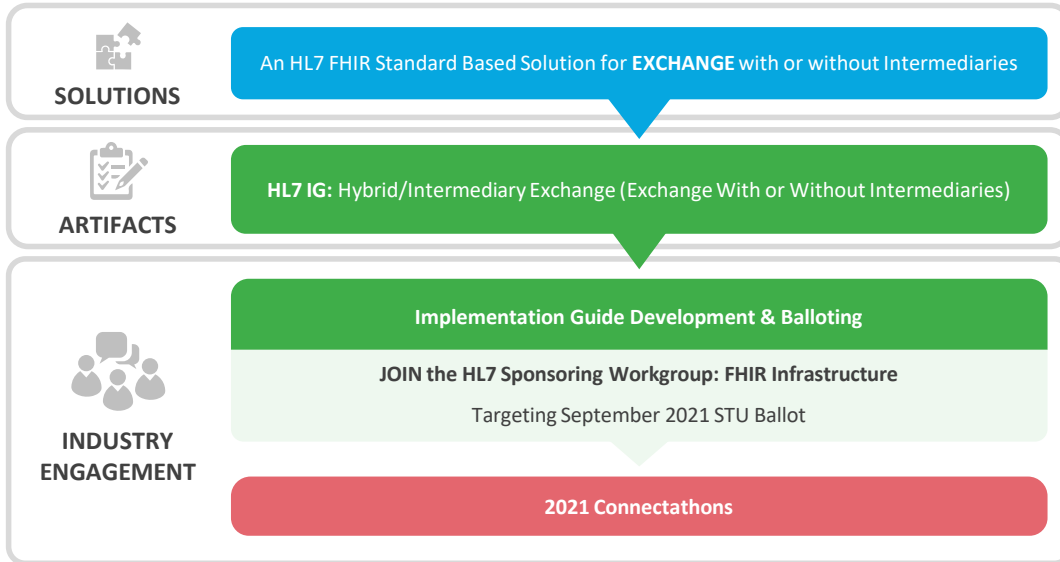
The illustration below includes an intermediary since the solution must operate in a hybrid model of point-to-point exchange or exchanging through intermediaries. The solution is depicted in a relatively simple transaction. The Requestor actor could be a payer, provider, or other stakeholder. The originator of the transaction in step 1 is sending the transaction, with the standardized routing information within the HTTP header, through an intermediary. The intermediary in step 2 can determine where the transaction should be forwarded based on the routing information included in the header and sends the transaction on to its destination. The recipient of the transaction processes the request and then can route the transaction back to the originator because the routing information included in the header is standardized. One or more intermediaries may be involved in routing the transaction, but the originator need only be concerned with identifying the final destination, and the Responder with identifying the originator to return the response.





Solution 6: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)

Next Steps and Industry Impact



The next steps for this solution are to:

Develop HL7 Implementation Guide: Hybrid/Intermediary Exchange (Exchange With or Without Intermediaries)



This solution will have the biggest impact on:

payers, providers, health systems, EHRs, public health, app vendors



Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Join sponsored HL7 FHIR Infrastructure Work Group
- Provide input and ballot of implementation guide
- Participate in HL7 Connectathons
- Participate in pilots



This solution offers these key benefits:

- Reducing burden through
 - Enabling organizations to delegate routing details to an intermediary reducing costs and complexity
 - Leveraging standards to minimize the need for unique solutions between exchange partners
- Enables intermediaries to offer value-add services in conjunction with FHIR routing
- Support adherence to the CMS-proposed rule on reducing provider and patient burden by improving prior authorization processes and promoting patient's Electronic Access to Health Information
- Streamline and support VBC workflows
- Increase scalability and improve interoperability between exchange partners





Solution 7: A FHIR TESTING AND CERTIFICATION Platform

Generally testing and certification in health IT helps provide a baseline assurance that technology will perform clinical care and data exchange functions in accordance with interoperability standards and user-centered design. FHIR is a named standard in federal regulations and a testing and certification program to validate that FHIR solutions can be implemented at scale is necessary to ensure interoperability. There are a number of documents including FHIR Implementation Guides that have been developed by *FAST* team members. Structuring *FAST* guidance into FHIR Implementation Guides will facilitate actionable and measurable guidance for the industry to adopt FHIR at scale. There are several challenges to creating a testing and certification program around these requirements. The *FAST* Tiger Team identified key technical and regulatory barriers that need to be addressed to consistently and reliably validate the baseline conformance for FHIR solutions to be scalable. These barriers include:

- Lack of documentation outlining requirements that will be tested against
- No definition of baseline FHIR conformance
- Limited number of tools able to test the multifaceted aspects of a solution simply using the instructions in the implementation guide as the requirements
- No culture/tools for multifaceted iterative solutions testing throughout the production cycle
- Lack of certification governance body to grant FHIR certification or scalability readiness

The *FAST* Testing and Certification Tiger Team focused on process and specifications for testing and certification of the requirements in the areas of identity, security, endpoint discovery, scaling, and exchange in addition to other FHIR readiness criteria, which include:

1. Endpoint discovery
2. Authentication
3. Authorization
4. Resource version identification
5. Reliable patient identity management
6. Data provenance
7. Reliable provider identity management
8. Event/Message/Topic Subscription/Publication
9. Guaranteed message delivery
10. Role/Context identification
11. Readiness credential
12. Standard-based endpoint access
13. Synchronous-transaction support
14. Asynchronous-transaction support
15. Reliable-payer identification

This work was accomplished in alignment with the other *FAST* Tiger Teams and industry stakeholders to identify the potential assertions that should define the baseline. Some consistent validation challenges that are addressed by the Scalable FHIR Testing and Certification Platform are:

1. Standardizing on FHIR version 4.0.1
2. Addressing stakeholder maturity levels
3. Automating complex workflow testing
4. Inclusion of backward compatibility as new versions of FHIR are released



Solution 7: A FHIR TESTING AND CERTIFICATION Platform

Quick Reference Summary of Barrier and Solution With Links to Resources



BARRIER

FHIR testing capabilities and an associated accreditation/certification are needed to support reliable, trusted exchange between healthcare participants. It must be a process in which specification/requirements that are well established and broadly shared can be absolutely confirmed



SOLUTION

Testing platform supporting the base FHIR Specification and FAST Readiness Criteria:

- Automated tool to perform testing and validation
- ONC FHIR Testing and Certification Program



IN SCOPE

- Establish scope of assertions for validation testing and certification to FAST Readiness and Assessment Criteria, including the base FHIR specification
- Draft Request for Information (RFI) to gather information on the best approach to testing platform/tool development and maintenance



OUT OF SCOPE

- HL7 FHIR Validation Engine RFP development to select entity to provide services
- Validate ease of establishing connections, conformance to non-blocking requirements, conformance to HIPAA patient privacy



STATUS

Incorporating feedback from industry stakeholders. Read the full SME session report [here](#).

Exploring a potential RFI to:

- Establish the scope of FAST Readiness Criteria assertions that need validation
- Identify existing tools/platforms in the industry that could meet the proposed FAST testing, validation, and certification requirements
- Obtain input on the feasibility of a certification model that would act as validation to industry players without being exclusionary



BENEFIT

- Testing improves consistency in performance of data sharing
- Validation tool decreases the burden of achieving interoperability with FHIR
- Certification provides assurance that products support interoperability
- Certification provides a benchmark and support of industry standards and implementation guides (IGs) to enable integration/interoperability with exchange partners
- Certification offers a level of stakeholder trust to exchange partners



OPEN ITEMS

- Assertion Table Completion: to document solution guidance type, clarify test method, and who needs to test
- Explore and finalize an RFI



CURRENT SOLUTION

[FAST Testing and Certification solution document \(version 3 in progress\)](#)

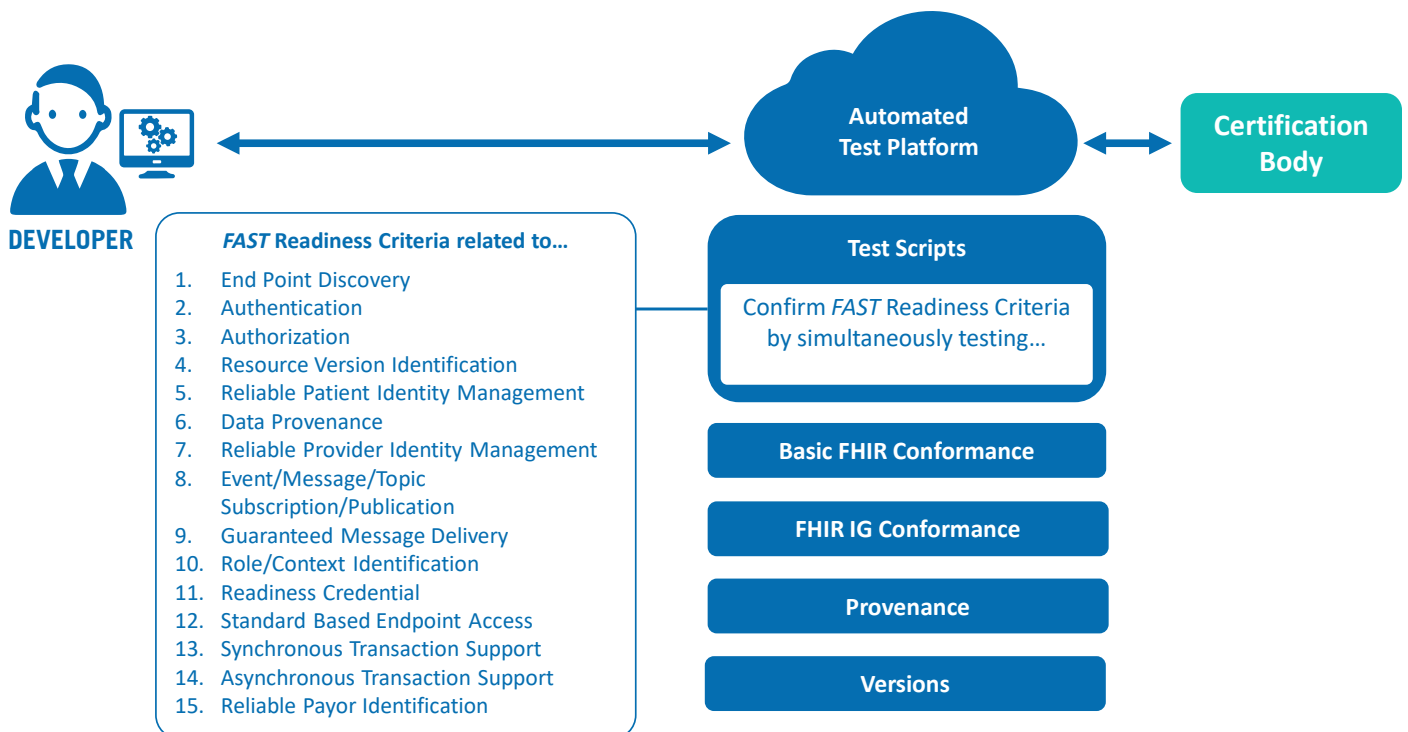


Solution 7: A FHIR TESTING AND CERTIFICATION Platform

The *FAST* Testing and Certification Tiger Team reviewed existing certification models and tools, such as FHIR tooling platforms Inferno and Touchstone. This review helped inform and refine the approach on how to test and certify scalable FHIR-enabled solutions. There are challenges to the development of a comprehensive testing and certification solution including the fact that FHIR is continually evolving, lack of minimum conformance, and the need to support validation of several layers, versions, and companion standards. Another factor considered was how to create a testing and certification program that would be easy to use and affordable to encourage adoption vs becoming a barrier to interoperability through FHIR implementation.

An important matter addressed by a FHIR Testing and Certification Platform is the need to test throughout the development and implementation of a solution. Waiting to test late in the software development cycle (ie, at certification time) reveals bugs that are exponentially more expensive to remediate than if they had been discovered and addressed prior to building on top of the bug. It is important that developers of FHIR-based systems are empowered with test scripts and automated tooling that enable them to test frequently to iteratively meet the requirements of FHIR Implementation Guides as they first begin to develop software, to avoid costly redevelopment.

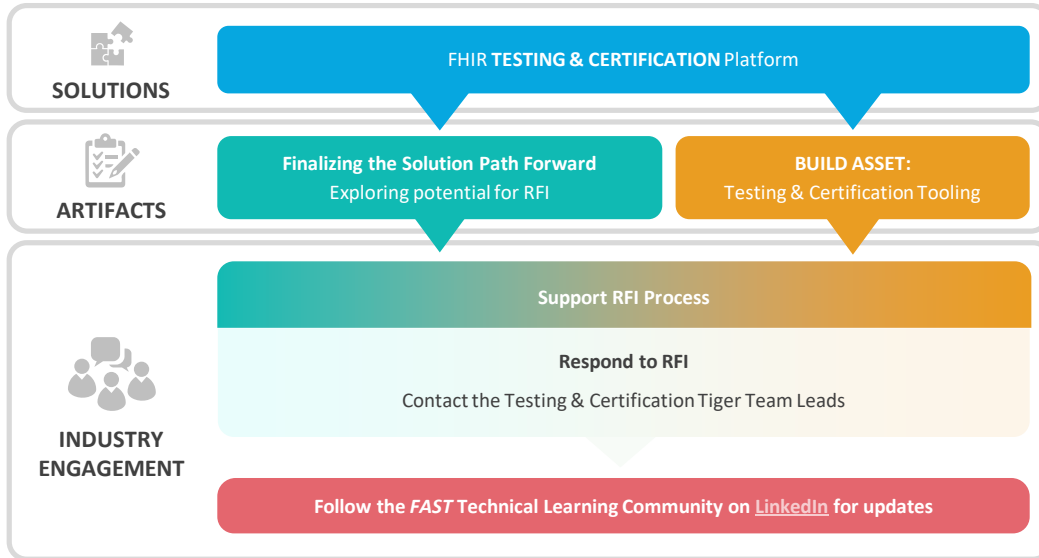
The *FAST* FHIR Testing and Certification Platform would provide 1 automated tool to allow developers to proactively test their solution early and often, ensure compliance and serve as the validation tool to obtain certification. This approach allows developers to iteratively test and improve their solution against the established prerequisites for base FHIR, as well as *FAST* readiness criteria. An automated tool leveraging the existing HL7 FHIR validation engine and allowing simultaneous testing of multiple FHIR scalability criteria would help reduce costs and improve efficiency. The proposed automated tool should include an easy-to-understand dashboard that would identify areas that need to be addressed before passing. This solution will provide an ongoing validation mechanism as solution developers need to adapt to new versions of FHIR, changes to infrastructure, and the needs of new stakeholders over time.





Solution 7: A FHIR TESTING AND CERTIFICATION Platform

Next Steps and Industry Impact



The next steps for this solution are to:

- Finalize the solution path forward: exploring industry RFI
- Identify potential governing bodies to manage and maintain this solution
- Develop certification program
- Build Testing and Certification Tool



This solution will have the biggest impact on: payers, providers, health systems, EHRs, public health, app vendors



Stakeholders of all stripes and types can help shape how this solution crystallizes for industry use by:

- Join sponsored HL7 Conformance Infrastructure Work Group to ensure consistency across IGs
- Participate in HL7 Connectathons
- Participate in pilots



This solution offers these key benefits:

- Reducing burden through
 - Providing a tool to allow testing early and often through application development thereby more easily achieving interoperability
- Improve patient experience and engagement through more consistency with how APIs are developed
- Assure data exchange and standardization between EHR systems and other exchange endpoints
- Support adherence to the CMS-proposed rule on reducing provider and patient burden by improving prior authorization processes and promoting patient's Electronic Access to Health Information
- Streamline and support VBC workflows
- Increase scalability and improve interoperability between exchange partners






Pilot Testing and FAST Solutions Validation Strategy

FAST's goal is to provide infrastructure patterns through the recommended solutions previously defined, so that FHIR solutions can scale across the US healthcare system. Since FAST is in the process of identifying entities who are best suited to further develop and implement the recommended solutions longer term, the FAST team needs to define a strategy for building solution piloting, testing and validation into this approach. The team has identified 4 overlapping phases of solutions testing and industry validation with successive levels of sophistication, that will leverage near-term efforts to inform longer-term initiatives:

1. Validation/technical testing through Connectathons
2. Pilot Toolkit/Sandbox introduced for FHIR Accelerators and other FHIR initiatives
3. ONC-convened demonstration projects (subject to approval and funding)
4. Coordinated multi-stakeholder pilot(s) (subject to approval and funding)



Phase 1: Validation/Technical Testing Through Connectathons

As FAST finalizes recommended solutions, testing becomes more of a reality. In the near term, HL7 FHIR Connectathons allow ready and willing ecosystem participants to validate individual FAST solutions and provide feedback to inform further development. FAST participation in Connectathons raises awareness of this work with organizations throughout the industry and provides real-world experience in testing these solutions that can inform later testing phases and initiatives, as well as implementation guide development. Iterative Connectathon testing and feedback is already taking place and will continue throughout subsequent Phases 2–4 described below.



Phase 2: Pilot Toolkit/Sandbox Introduced for FHIR Accelerators and Other FHIR Initiatives

The *FAST* team recognizes that in order for entities to test the *FAST* solutions they'll need easy access to tools so that they don't need to spend effort discovering how to test each solution or hunt for the information they need. Taking learnings from Phase 1 Connectathon Testing, the *FAST* team will layer on additional implementer support by developing a toolkit that includes materials that help implementers understand what the *FAST* core capabilities are and how to use them, in the form of process/supporting documentation, specifications/implementation guides, and supplemental guides. In addition, a sandbox will be developed so that entities can test that they can interact with the *FAST* core capabilities as well as some available functional use cases (eg, Da Vinci, etc.). The team is currently working through sandbox requirements, but the environment could potentially include synthetic data, APIs for directory services, or simulate intermediary routing. The Pilot Toolkit and Sandbox will continue to be available throughout Connectathon testing and Phases 3 and 4 described below.

Phase 3: ONC-Convened Demonstration Projects (subject to approval and funding)

Once solutions have been individually tested through Connectathons and tested in a sandbox, then the next step is to start defining some projects to test solutions through industry partnerships incorporating the lessons learned from earlier testing phases. These implementations will be industry-led and funded, where ONC would recruit organizations based upon some basic requirements for participation and provide materials as guidance to help them define their proposed demonstration projects. For example, demonstration projects could focus on payer-to-provider, payer-to-payer, or payer-to-application exchange with multiple entities representing each type of stakeholder in support of CMS' Patient Access API requirements.

Phase 4: Coordinated Multi-Stakeholder Pilot(s) (subject to approval and funding)

Lessons learned from previous phases will be incorporated into Phase 4 to define coordinated pilots where multiple healthcare actors in the industry integrate with each other to exchange information and are using multiple *FAST* solutions. This phase represents the most complex level of testing, which will require program level coordination, support, and evaluation. Potential funding may come through joint stakeholder efforts or industry in-kind support.



Conclusion

Over the course of the past 3 years, *FAST* volunteers have dedicated a great deal of their time and expertise to develop and advance the *FAST* set of infrastructure scalability solutions, providing a vision of the infrastructure architecture needed to support FHIR solutions at scale, but the work isn't done yet. As these *FAST* solutions embark upon their paths to be implementation-ready, industry stakeholders still have a role to play, whether providing input during the implementation guide development process, providing suggestions on which organizations would be a good fit to be the long-term stewards of *FAST* solutions or by participating in pilots.

The most important step, of course, is for these *FAST* solutions to eventually be adopted and put into practice. Each of these solutions solves for specific challenges that have been faced as FHIR adoption has grown. These solutions, while they can stand alone, work better together in an integrated architecture for more seamless data flows. Adopting as an integrated architecture, while not required, may also aid in meeting requirements of interoperability rules. Solving these problems will lead to a variety of benefits including:

- Burden reduction through less re-work, avoiding the need for unique point-to-point connections and having a central location to maintain endpoint details
- Increased confidence and consistency in patient matching
- Providing one source of truth for FHIR endpoints
- One tool for iterative testing and certification

To learn more about *FAST* solutions and how you can get involved, please contact Stephen Konya, ONC Lead for the *FAST* Initiative: at ONC (Stephen.Konya@hhs.gov).

Acknowledgments

Thank you to all the *FAST* volunteers, taskforce staff, and SMEs who lent your expertise and perspective over the years. We look forward to the exciting next phase of *FAST* work.