Key Considerations Related to Policies for Interoperable, Federated Provider Directories

Version 1.0

(Last revised: March 6, 2014)

Prepared for:

This document was prepared for the U.S. Department of Health and Human Services (HHS), Office of the National Coordinator for Health Information Technology (ONC), State Health Information Exchange (HIE) Cooperative Agreement Program, through the State HIE Provider Directory Community of Practice (CoP). It was developed as a resource for ONC's State HIE Cooperative Agreement grantee community, and is not intended for public distribution.

Prepared by:

This document was created by Krysora, an information technology development and consulting company with a focus on health care and health information exchange.

DISCLAIMER

This document was created with funding support from ONC, but does not necessarily reflect the views or opinions of HHS or ONC.

Introduction

Recent months have seen advancement of technical requirements for provider directories. Efforts like the EHR/HIE Interoperability Workgroup, the Western States Consortium project, ONC's Modular Specifications, and work by ONC State HIE cooperative agreement partners have furthered the development of specifications based on Integrating the Healthcare Enterprise's (IHE's) Healthcare Provider Directory (HPD) profile. By detailing necessary aspects of provider directory interactions and addressing federated relationships, these specifications provide a framework enabling technical interoperability between directory implementations.

With the opportunities created by widespread technical interoperability and the capabilities offered by these standards, implementers and operators are moving to enable their directories to connect with others to form networks of directories. Parties who have actively piloted approaches to this have found that successfully forming such networks requires sufficient alignment of policies and scalable ways to extend those networks. Their experiences (and experiences of parties in related arenas that have faced similar challenges, such as the Direct ecosystem) demonstrate that, at best, insufficiently aligned policies delay operators from being able to participate in or form such networks; at worst, they can even prevent operators running under such policies from connecting at all.

In light of this, based on previous work performed by parties piloting provider directory policies (notably the Western States Consortium) as well as information gathered in interviews with parties active in this space – in particular DirectTrust, EHR/HIE Interoperability Workgroup, Healtheway, and the National Association for Trusted Exchange (NATE), formerly known as the Western States Consortium (WSC) – this document sets forth certain key considerations related to policies for interoperable provider directories forming networks that could contain federated relationships. These considerations are not intended to prescribe policy. Rather, they should be viewed as an informative resource for directory operators and implementers as they consider their directory policies and participating in or forming networks with other directories.

Overarching Considerations

Terminology

Assessing policy alignment can be difficult if the concepts covered by the policies in question are not clearly defined and understood. Two examples of key concepts that could be defined differently by various directory operators include:

1. The very idea of "provider directory". Different parties may have varying ideas of what actually constitutes a provider directory (e.g., some may understand a "provider directory" to store a certain set of data related to actual providers and provider organizations, while others may consider a "provider directory" to host a more expansive set of data related to providers, plans, and others).
2. What constitutes federation. Federated directory networks can be constructed in a number of ways. A directory network may be "flat", with all members directly connected to one another, or may feature more complicated structures such as hierarchical arrangements (whereby data is stored "at the edges" and queries are routed through the environment starting at a central point, which then distributes them to next-level delegates that may in turn further distribute the queries to their own delegates).

Defining these and other concepts (and keeping them current as they may evolve) as part of a policy framework can help operators better communicate the nature and coverage of their policies and facilitate assessing alignment of their policies with those of other operators.

Incremental Policy-Making

Use cases can evolve over time as the needs of users grow and change. Trying to anticipate out of the gate every potential use case a directory might serve into the future and craft associated policies may be tempting to directory operators but could hinder and delay serving the needs of today. As well, within a group wishing to connect their directories as a network, it may take time for members to be comfortable and capable of meeting network-wide policy requirements. Directory operators and networks may want to consider taking an incremental, iterative approach to policy-making, enabling policies to be focused on agreed-upon needs as those needs emerge and to be rolled out as they can provide value.

Considerations Related to Specific Provider Directory Policy Areas

Electronic Service Information (ESI) Discovery

Electronic Service Information (ESI) housed in provider directories defines how to communicate with an information exchange partner. It includes such information as electronic information addresses (e.g., Direct addresses and URLs for query services) as well as supported transport,

content, and security options. ESI discovery refers to the process by which a party queries for the ESI of a potential exchange partner, and with the inclusion of health information exchange-related measures within Meaningful Use and the push for more effective delivery of healthcare. ESI discovery is an important use case of directories that fulfills a core need of providers, who increasingly are conducting administrative and clinical transactions electronically.

Collection and Maintenance Policies

In looking at supporting ESI discovery, a provider directory must have processes by which ESI is collected and maintained for entities within it. Policies aimed at maintaining the currency and accuracy of such data will not just be a concern of those directly served by the directory but, also of those within the network to which the directory is connected. See **Quality of Data** below for more on this.

Minimum Dataset Policies

Searches for ESI for a given entity within the directory are based on demographics and other information identifying that entity. Depending on the type of entity in question (e.g., individual, organization), examples of this information could include but are not limited to name (first, last, organizational), identifiers (such as National Provider Identifier), specialty, relationships and affiliations, physical addresses, and phone and fax numbers. The **S&I Framework Provider Directory Initiative**[1] has detailed this **use case**[2], as well as **dataset considerations**[3] for potential data elements that requestors could expect to search. While the potential data are rather basic in nature and are likely to be collected and maintained by most directories for the entities they store, it's possible that a given directory may not do so for some or all entities, or that a directory may restrict access to some or all of this data, thereby making them unavailable to search. However, operators of such directories may want to consider that pilots in this space have indicated providers and other users can be less comfortable exchanging with entities they discover in directories if they feel that the supported data upon which they can search is not comprehensive enough to identify an entity as the sought after exchange partner. This can influence adoption and usage of directories, as well as factor into whether other directories will be willing to connect due to the potentially low value of such connections to their own users.

[1] http://wiki.siframework.org/Provider+Directories

[2] http://wiki.siframework.org/PD+-+Query+for+Electronic+Service+Information+including+Electronic+Address++Use+Case

[3] http://wiki.siframework.org/PD+-+Query+for+Electronic+Service+Information+including+Electronic+Address++Use+Case#x12.0 Dataset Considerations

Data Return Policies

Similarly, while a directory may allow searches for ESI using a range of demographic and other identifying information, it may have a policy prohibiting some or all of this data to be returned as part of an ESI search result, potentially only returning the actual ESI itself associated with entities[4]. While a result containing just ESI for matching entities is technically sufficient for parties to engage in exchange, in practice, a search could match multiple entities, and as above, providers and other users, in order to avoid disclosing PHI to the wrong parties, may be unwilling to engage in exchange if they feel there is insufficient information provided in the search results to identify their potential exchange partners.

Quality of Data

As a source of data, a provider directory's value to its users is directly dependent on the quality of the data it delivers. A provider directory that delivers outdated or inaccurate information may find that its users are hesitant to rely on the directory, in particular if they plan to act on this information in ways that carries or is perceived to carry risk, such as using electronic service information (ESI) to engage in health information exchange. Other directories also may be less inclined to enter into relationships and connect to such a directory, and those that already have done so may even elect to terminate their relationships over time and to disconnect the directory from their network.

Policies focused on establishing and maintaining data quality can help a directory to deliver the value its relying parties, both existing and potential, expect. Things to consider include:

- Identifying information such as demographics and affiliations are fundamental data elements. Any directory that does not place an emphasis on establishing and maintaining current and accurate identifying information will be challenged to serve core use cases, including ESI discovery.
- When looking at tackling quality over a broad range of data, taking an iterative approach and focusing on those data elements associated with higher priority and important use cases can be one way to make the process more tractable.
- Accuracy and currency of data in the directory may depend on processes and actions of data submitters. Policies covering data for which the directory itself is responsible will need to be reflected in agreements with data submitters who have responsibility for their data.
- Information associated with entities that no longer exist or are inactive (e.g., an individual who has retired, an organization that has closed its doors) can create false expectations and result in frustration when parties relying on that information attempt to perform desired

[4] This is distinct from any potential operational policies a directory may have related to refusing queries that are deemed to pose an operational or business risk to the directory, such as broad queries containing sufficient wildcards that would result in the directory responding with its entire dataset.

use cases. A directory with policies and processes to minimize the presence of such information from a directory will provide higher value to relying parties.

Permitted Purposes

Past efforts in this space have shown that parties who can contribute data to a directory, as well as the entities themselves that could appear there, will participate only if they understand and accept who can access data and how data can be used. Clear statements within a directory's policy framework defining the permitted purposes for which data can be used and by whom can provide comfort to possible participants.

When defining permitted uses, directory operators may want to note that, via requirements in state HIE Operational Plans as well as in the "Governance RFI"[5], ONC has called upon directory operators to make their directories open and as accessible as possible to foster health information exchange. While data contributors may desire comfort, putting in place too many restrictions or restrictions generally deemed unreasonable could deter or hinder users from exchanging information, damaging a key value proposition directories seek to offer.

As well, business or financial drivers could incline some directory operators to permit data to be used for purposes beyond furthering treatment and other health care-related activities (e.g., marketing). It's beyond the scope of this document to recommend particular uses that a directory should permit, but for health care-related directories, permitting data to be used in ways that go beyond health care may inhibit participation.

Considerations Specific to Directory Networks

Scalability

While one-off contracts can provide a way for directory operators to begin forming relationships with other directories, continued use to create a broad network of relationships is not scalable (doing so is an "N-Squared" problem, wherein the number of contracts required grows exponentially as the number of parties involved increases). A group of operators looking to build a broad network will need to consider other, more scalable ways to agree upon policies, bind one another to those policies, and manage their relationships, such as using a single, common, multi-party agreement that each member of the network signs once upon admittance into the group.

As well, when two provider directories decide to connect, it's possible that each may require the other to complete an onboarding process wherein compliance to policies and conformance

[5] Proposed Rule document posted May 15, 2012. *Nationwide Health Information Network: Conditions for Trusted Exchange*, http://www.regulations.gov/#!documentDetail;D=HHS-OS-2012-0006-0001

to technical standards are assessed. A group of directory operators interested in forming a network may need to consider standardizing on common onboarding criteria and procedures to avoid another form of the "N-Squared" problem, one in which a new network entrant has to pass through the onboarding process of each and every directory in the current network, and each and every directory in the current network has to pass through the new entrant's onboarding process.

Federated Relationships

When a directory network features federated relationships (whereby directory queries could be received by certain nodes on the network which in turn delegate those queries to subordinates and then aggregate and deliver a composite response), unless policies and procedures of the network are constituent components of agreements with flow down requirements to delegates, it's possible that there could be some directories effectively acting as members of the network that are not bound by the network's policies and that do not comply with them. A network may want to consider this when constructing its policies and the agreements that bind its members to them.

Directory Intermediaries

Some directory network architectures, often those with federated relationships, feature intermediaries that orchestrate directory searches among a set of directories by distributing queries and aggregating responses. In some cases, these intermediaries can be directories that answer certain queries themselves while also performing orchestration, and in others, be routers that solely act as orchestrators. Regardless, by its nature, any intermediary that is part of the network will have access to all the responses from the directories it fronts, and by design, whether a given node on the network is an intermediary and how it functions may not be obvious to other members of the network. Directory operators forming a network may want to consider whether and how intermediaries might be part of their network to ensure their policy framework applies not just to directories but also to those intermediaries.

Security

Once directories join to form a network, security becomes the shared responsibility of all members of that network. With data in each directory now available to all, security incidents involving one directory could lead to issues elsewhere like inappropriate access. To mitigate this, as part of the security policies for the network, directory operators may want to look at addressing traceability and authentication.

Traceability

Traceability refers to the ability to track a specific directory transaction using an audit trail or other recorded information. In a single directory environment, the directory itself has all the information it needs to record and make traceability straightforward, such as the source of

each search, the query itself, and whether information was returned. Should a question arise regarding one or more transactions, the directory operator can refer to the directory's logs and obtain a complete picture. Traceability within a directory network, however, is more complicated. Once issued by the source, any number of intermediaries can relay that search request, and multiple directories can return responses. For traceability to be possible, a directory network may want to consider the following.

- Full traceability will require every member of the network to maintain adequate audit logs. This means:
    - Each member will need to log all incoming and outgoing transactions, including date and time, nature of the transaction (e.g., query or response, incoming or outgoing), from whom the query or response was actually received (if the transaction is incoming), to whom the query or response is passed (if the transaction is outgoing), and any other information that could help identify a particular transaction.
    - The directory actually first receiving a query from a user will also need to log information that can tie the query to that user.
    - Any directory providing a response will also need to log if its response was negative or positive, and if positive, enough information to later identify what was conveyed.
- The network will need policies governing requests for logs and related procedures.

Authentication

Authentication is the process of verifying that someone or something is who or what it purports to be. While traceability helps with reacting to potential incidents and determining their impact after the fact through tracking, authentication is a proactive measure aimed at preventing certain incidents. In the context of a directory network, authentication can be used to establish that all the parties involved in a transaction – the user who initiated the query and the directories and intermediaries that act to answer that query – are who they say they are and are valid members of the network bound by its policies.

If looking at addressing authentication in the security policies of their network, operators may want to consider:

- Interoperability. Directories and intermediaries that support incompatible authentication mechanisms will not be able to communicate, so defining and requiring standard authentication mechanisms for transactions between directories and intermediaries within the network is key.
- Authentication of users. If users are not authenticated, then they may perform actions for which they are unauthorized. As well, it may not be possible to tie transactions to them in support of traceability.