

Information Sharing Model for Health and Human Services Agencies

Carrie Hoff, Deputy Director, Health and Human Services Agency, County of San Diego | Carrie.Hoff@sdcounty.ca.gov
Subhankar Sarkar, Associate Partner, IBM and Chief Architect, ConnectWellSD | ssarkar@us.ibm.com

Abstract— Nearly everyone involved with Health and Human Services (HHS) agencies believes that better information leads to better care. There is also the recognition that this information has become dispersed and fragmented over time. Legislative shifts, different funding sources and increased demand has led to proliferation of information systems supporting delivery of HHS programs. This paper proposes a model for information sharing across this systems portfolio, such that the mission of providing better care may still be achieved. The model manifests as an information exchange, the resources wherein can be securely accessed by a variety of stakeholders.

The model comprises of the following patterns – Identity Resolution, Composite Customer View, and Authorization and Consent – bound together in a Service Oriented Architecture (SOA). Various methods of integrating the composite customer view into business processes are also discussed. The entire model pivots on the central principle of person-centric care, wherein information is brought together from multiple systems and channels of service to holistically fulfill the client’s needs.

To assemble information about a customer from multiple sources, first an identity must be established for the customer that spans these sources. The model uses a Master Client Index (MCI) for this purpose, which includes various demographic and other attributes of the customer that define his identity. The MCI is built from the various source systems, and is in effect an index of indices.

When a user requests customer information on the exchange, the request goes through the MCI for identity resolution. Once a customer’s identity has been resolved, the system can then make federated callouts to multiple source systems to obtain relevant data. This data is then assembled by subject area, and presented to the user, as per the Composite Customer View pattern described in this paper.

This information sharing model would not be executable in the absence of an Authorization and Consent pattern. While information sharing across the agency is essential for person-centric care, such sharing must respect the relevant policy and statutory controls. These controls come from multiple sources – HIPAA, 42 CFR Part 2, WIC 827 (CA Child Welfare), WIC 5328 (Ca Mental Health), State Penal Code, Title 17 CCR (Public Health), etc. This labyrinth of regulations can be an insurmountable barrier to effective information sharing. This paper describes a robust, fine grained and flexible access

control model that enables the agency to maintain compliance, without impeding the authorized flows of information across the enterprise.

The model described in this paper was formed and implemented at the County of San Diego, as part of the Connect Well San Diego project, itself a part of the Live Well San Diego program. Variants have been implemented at the counties of Los Angeles and Sonoma, and the Ministry of Social Development in New Zealand.

I. INTRODUCTION

Certain defining characteristics pertain to HHS systems portfolios, and the domain in which they operate.

1. There’s a multitude of systems, each purporting to serve a certain group of social programs, often instituted through a specific funding source. There is a mix of home grown systems and commercial-off-the-shelf (COTS) products. Most of the data is stored in relational databases of various brands. The application software is usually client-server or web based.
2. These systems exhibit tremendous variance in size, technology and flexibility. More importantly, they are semantically disparate, in that they represent customers and service flows in very different ways.
3. A given customer, more often than not, participates in multiple programs managed by multiple systems. Case workers using these different systems are unable to form a holistic picture of the customer or his needs.
4. Many agencies establish data warehouses to bring together information from the various systems. While such warehouses provide important analytic insight at an aggregate level, they have limited relevance at the point of service for a specific customer.
5. Almost all the information in these systems is PII/PHI; authorization requirements are therefore complex and compliance requirements onerous. In order to meet these requirements, most agencies build these systems as silos, and grant a few workers access to each system. Authorization and consent policies are narrowly formulated and individually applied at the system level. Paradoxically, this adherence seldom results in the

customer's betterment, because it traps relevant information in opaque containers.

6. These systems have various types of user access controls which are not designed for information sharing. Rather, the user access controls are designed with the assumption that all users will also work in the program/department (or be a contracted service provider for the program/department) that 'owns' the system. This design assumption results in simple frameworks for segregation of access to data, which are usually based on a single attribute (user role).
7. These systems collect a wide variety of data that is supportive of effective service delivery. In addition to collecting data specific to a service, program, or treatment type, these systems collect and track other pieces of information. The recognition of the role of the social determinants of health in developing an effective plan with a customer, the challenges with accessing information across programs, and for some programs, the enrollment criteria, are logical drivers toward the diverse set of information collected. A program providing families with economic assistance may collect information about a customer's literacy level, and an alcohol and drug treatment program may capture a customer's status as a military veteran.

The County of San Diego Health and Human Services Agency (HHSA) recognized these challenges in 2008 and began forming a vision of the future. These developed into "Live Well, San Diego" a strategic blueprint for a region that is Healthy, Safe, and Thriving. To achieve wellness in its fullest sense, the County of San Diego HHSA needed an integrated view of the individuals and families receiving services, and the ability to collaborate across program silos as integrated teams with shared customers. HHSA serves approximately 1 million customers per year in a region that has a population of 4.4 million people. The scope of programs includes social services, health, behavioral health, protective services, and housing. To support this breadth and scale, the HHSA developed ConnectWellSD - an information exchange that promoted secure data sharing between County departments, and enabled a Person-centric service perspective. This objective was accomplished in 2016. Since then, the ConnectWellSD program has progressed with several projects that leverage the information exchange and make it an intrinsic part of HHSA operations.

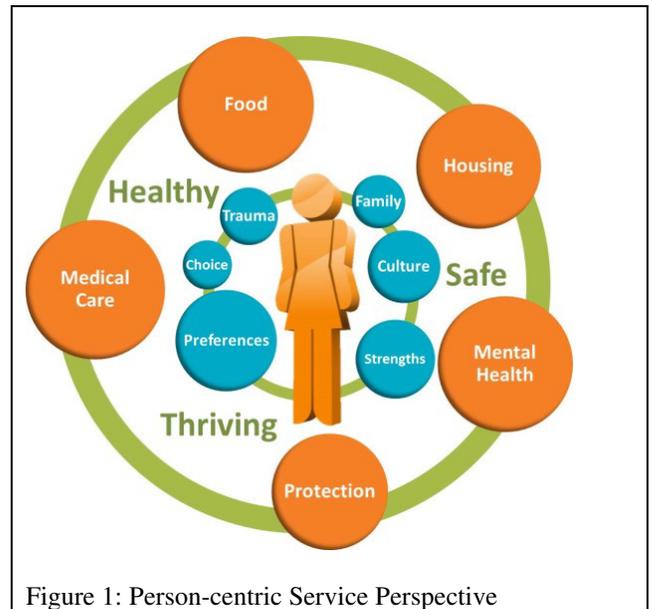


Figure 1: Person-centric Service Perspective

The information sharing model described in this paper, and realized through the ConnectWellSD information exchange, is based on the following principles:

1. Person-centric care: HHS agencies exist to serve the customer. Services should then be provided based on a holistic understanding of the client's needs, not based on what agency location he happens to walk into. HHS customers often do not have a full understanding of the social programs available for his support; also, his participation in one program may influence the results of, and suitability for, another program. It is in the case worker's brief to inform the customer of his choices, and to extend a service that meets his needs holistically than in a piecemeal fashion. Of importance is the additional tenet that the customer is a member of his or her own care team, and therefore self-service access must be considered in the information exchange.
2. Policy based authorization to information: HHS agencies need to comply with regulations e.g. HIPAA and 42 CFR Part 2. They need security and privacy controls, including capture and execution of client consent. They also need policies around what information can legally be shared as 'need to know', in the absence of explicit consent. A flexible access control model, that realizes the agency's authorization policy, is a necessary enabler. This paper posits that the customer perceives the agency as a single entity, not a composition of departments. The agency, too, must share customer information as per policies that further the customer's well being, rather than boxing information in disjoint systems.
3. Composite customer view as a foundation to better service: The composite customer view is rendered as a component, which can be integrated into business processes in several ways. Agencies may choose to integrate this composite view into their current point

of service applications, or create a new application that performs cross-cutting collaborative functions such as referrals. The composite view component may also be integrated into a data warehouse. The versatility of the model stems from this ability to support multiple integration points.

II. MODEL OVERVIEW

The information sharing model comprises of the following patterns:

1. Identity resolution: This pattern concerns itself with establishing an identity of the customer across multiple systems, based on certain demographic characteristics that are provided as input.
2. View composition: This pattern concerns itself with assimilating information about the customer from various systems, and assembling and presenting the information by subject area.
3. Authorization and consent: This pattern concerns itself with applying privacy controls on the presented data. The agency's authorization policies, and its incorporation of customer consent, determine what information is accessible.

The different components of the model are formed and integrated in an SOA. The principles of SOA – loose binding, statelessness and composability – are essential towards the building of the model. As such, SOA is not described in a separate section, but permeates the description of every component.

III. IDENTITY RESOLUTION

To assemble information about a customer from multiple sources, first an identity must be established for the customer that spans these sources. The model uses a Master Data Management (MDM) system for this purpose. The MDM serves a Master Client Index (MCI), which includes various demographic and other attributes of the customer that define his identity: name, address, phone, date of birth, SSN, driver's license #, alien registration #, ethnicity, gender, language, etc. The MCI is built from the various source systems, and is in effect an index of indices. A user of the composite view, be it a human or a system, will start with a search and lookup of this index, by entering some information about the client. MDM will then return a set of composite records along with a match probability. When the user selects one such composite, the identity of the customer across the various source systems would have been resolved.

The key to establishing the MCI is the effective cross-linking of client information across source systems. Not just deterministic techniques, but also fuzzy matching techniques, which account for anomalies such as misspelled names or transposed digits in ID #s, must be supported. Figure 2 illustrates the formation of the MCI.

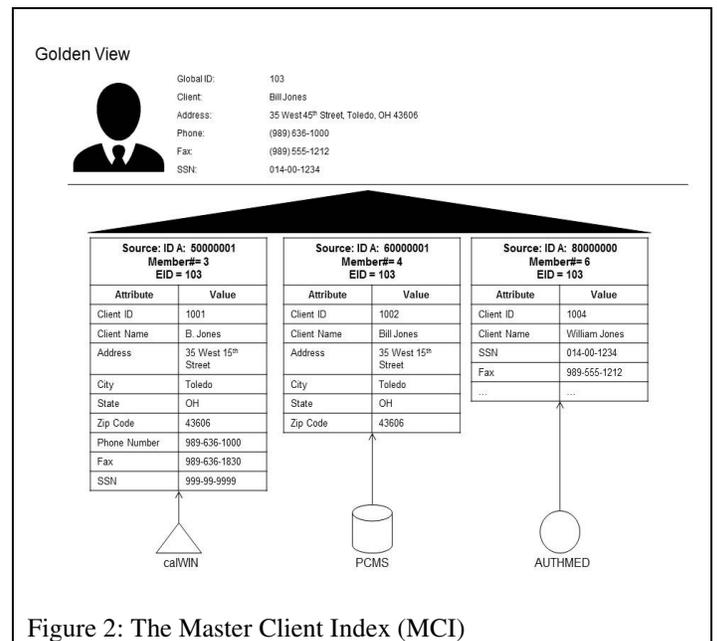


Figure 2: The Master Client Index (MCI)

The model deliberately employs the *registry style* of MDM, not the *hub style*. In the hub style, master data is maintained in an MDM hub, to which all systems must subscribe. Use of an MDM hub is an imperative in certain scenarios, e.g. the Federal Housing Administration's maintaining an Institutional Master File for approved lenders. Transactional (underwriting) systems must then subscribe to this MDM system such that only approved entities (lenders) are used. This approach is usually not suited for large and diverse Health and Human Services (HHS) portfolios; the cost of modifying every system to draw customers from the hub is prohibitive. The registry style offers a more practical and effective alternate, where source systems contribute to the MCI, and the MDM system employs matching techniques to link records across source systems. Source systems do not have to be modified to support the registry style. A Search and Lookup service is provisioned on the MDM registry, which users on the exchange use to identify customer and locate their information.

IV. VIEW COMPOSITION

This component of the model composes the 360 degree customer view, based on relevant information in the source systems. The information is arranged by subject area i.e. what it is, rather than where it came from. Although the system maintains source system traceability in the background, the origin of the data need not be presented to the user. Perhaps more importantly, the origin of the data must not be required for a user to navigate to information needed to perform their job. If data was not arranged by subject area, but instead by the system of origin, it would make it difficult for service providing staff to quickly consume and process customer information, and render navigation of the information exchange awkward at best.

With the View Composition part of this model, information in a certain subject area - medical history, say - may come from behavioral health, public health and Alcohol and Substance Abuse systems. Each subject area is presented as a resource on the information exchange. The requestor requests the resource or the collection of resources, not data in a specific system. This arrangement allows for the breakdown of inter-departmental silos, and exposes information as per the agency's authorization policy described later in this paper. Figure 3 illustrates the concept.

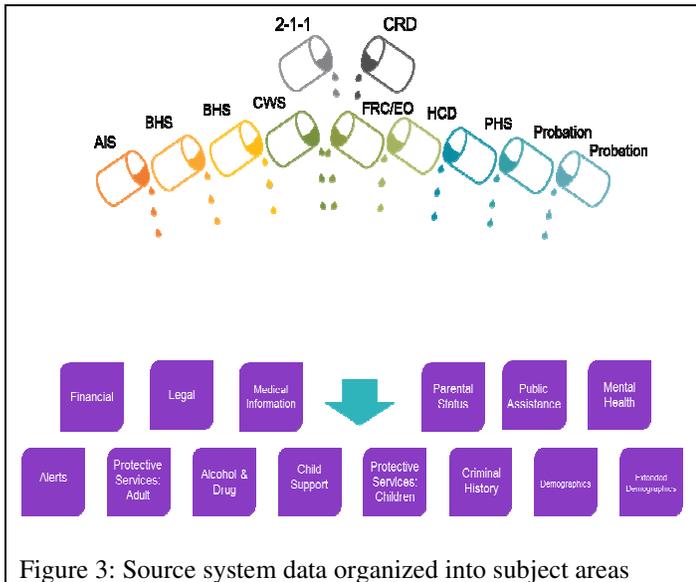


Figure 3: Source system data organized into subject areas

The composition of the 360 degree view relies on the identity resolution pattern described earlier. Each system contains certain demographic data about the customer, which goes into the MCI, and is used to resolve the customer's identity and his index of indices (source system identifiers). Once a customer's identity has been resolved, this index drives federated callouts to multiple source system. The information retrieved is then assembled by subject area, as per the metadata descriptor for the customer view, and presented as resources on the information exchange. The requestor only needs to specify the resource(s) it needs, and need not have any understanding of the back-end sources or mechanics. Figure 4 illustrates this concept.

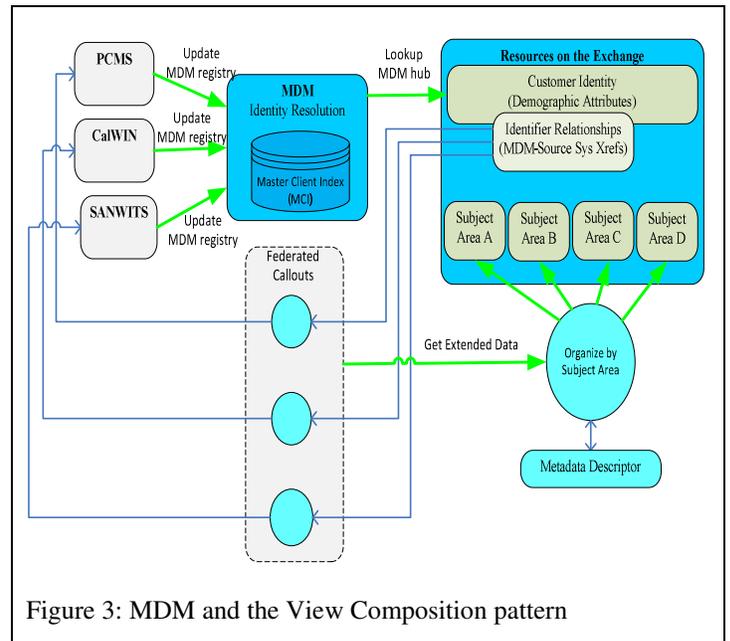


Figure 3: MDM and the View Composition pattern

V. AUTHORIZATION AND CONSENT

This information sharing model would not be executable in the absence of an Authorization and Consent pattern. While information sharing across the agency is essential for person-centric care, such sharing must respect the relevant policy and statutory controls. These controls come from multiple sources – HIPAA, 42 CFR Part 2, WIC 827 (CA Child Welfare), WIC 5328 (Ca Mental Health), State Penal Code, Title 17 CCR (Public Health), etc. This labyrinth of regulations can be an insurmountable barrier to effective information sharing. This paper describes a robust, fine grained and flexible access control model that enables the agency to maintain compliance, without impeding the authorized flows of information across the enterprise.

The information that drives the care coordination and Person-centric service model falls in several spheres. Some information, such as demographics, is freely accessible across the enterprise. Some information, such as medical information, may be subject to customer consent. In the absence of customer consent, the agency's authorization policy may still allow some level of sharing. Some information may not be subject to person consent but agency's policies may still restrict sharing. Also, there may be customers, such as offenders on probation, who may have waived their privacy rights as a condition to their release. Information collected in a system that is outside the program boundary presents another facet to the application of authorization and consent. The number of possible permutations is large, and to make it more difficult, regulations are often in flux. This paper presents an Authorization and Consent pattern that encapsulates this complexity, and enables information sharing within this agency within the bounds of compliance.

The pattern uses Attribute Based Access Control (ABAC) to secure the Object (the thing to be protected), from the Subject (the thing that is seeking to access the Object). Object and Subject come from Access Control nomenclature, and are abstract entities within this pattern. The Object can be a database table, a record within a database table, an application page, or a web service. The Subject can be a human or a system. Any attribute within the Object can be used as an access controlling attribute. To keep the model manageable, Subject side attributes are limited to 3:

- Organization – The department or office or firm the user works for, i.e. the node in the organization hierarchy to which he functionally associates.
- Role – The user level e.g. worker, supervisor, executive, which usually corresponds to his level in the organization and his authority thereof.
- Service Delivery Model – This reflects the line of work of the user and the nature of his customer engagement: Medical, Social, Transactional, Protective services, Law Enforcement and Case Administration. This list will likely be similar across HHS agencies.

Access control policies are formulated using these attributes. A Subject is granted access to an Object defined by certain attribute-value pairs, based on his Organization, Role and Service Delivery Model. The presence or absence of customer consent determines the policy to be applied. The customer consents to sharing all or certain sections of his data, but does not direct what users it can be shared with. In other words, consent associates to the Object side, not the Subject side. The customer cannot specify the Subject side as he is not expected to understand the internal organization or workflows of the agency. As far as the customer is concerned, he is interacting with the agency in its entirety. The Subject side i.e. who can access what sections of the customer’s data in the presence or absence of consent, is governed by agency policy.

It is important to note that this is far more than a Consent solution. A binary, point-to-point Consent solution that establishes trust between a consenting party and receiving party, is not enough to sustain information sharing in a large HHSA information exchange. In this pattern, Consent is conceptualized as a policy shift. Let’s imagine a continuum of access: 0 to 10, with 10 being full access to the customer’s data and 0 being no access at all. Consent nudges access closer to 10, but may not take it all the way to 10. For example, even though the customer has consented to sharing his medical data, policy may still deny law enforcement access to that data. Similarly, in the absence of consent, access may not be 0. Even if the client has not consent to sharing his probation data, say, policy may grant law enforcement access to that data.

The ABAC model is illustrated in Figure 5, which shows the Subject side and the Object side, and the influence of Consent. The transactional context may also be a factor e.g. when the

request is coming in thru a certain protected API, a specific policy set may apply.

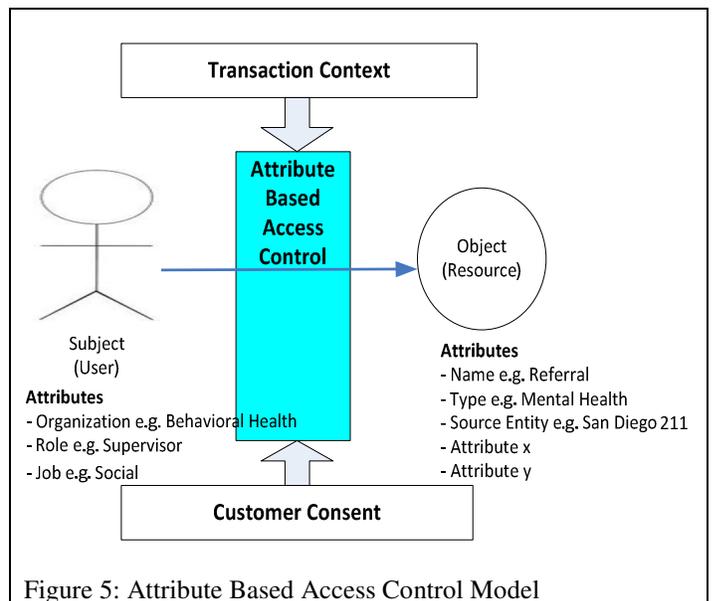


Figure 5: Attribute Based Access Control Model

While this model is comprehensive and extensible, certain issues endemic to an ABAC model need to be carefully handled. For example, a large number of Object side access controlling attributes may hinder performance. Also, rules collision is a frequent issue in ABAC systems. ConnectWellSD avoided these pitfalls through an ABAC data model that pre-empted rule conflicts, and that allowed rules to be described as a simple combination of attributes rather than detailed procedural logic.

Architecturally, the ABAC system is organized as illustrated in Figure 6. The solution comprises of the following:

- A Policy Administration Point (PAP), whereby access control policies, including application of consent, can be administered.
- A Consent Administration Point (CAP), whereby client consent can be captured and managed through its life cycle.
- A Policy Deployment Point (PDP) for privacy controls, reachable through standard APIs.
- Policy Enforcement Points (PEPs) in various applications, which intercept requests for protected data on the information exchange, and route to the PDP for adjudication.
- An audit trail for every request, that includes the source and reason for the data request, the transactional context around the request, and if the request was allowed or denied. This audit trail supports compliance reporting.

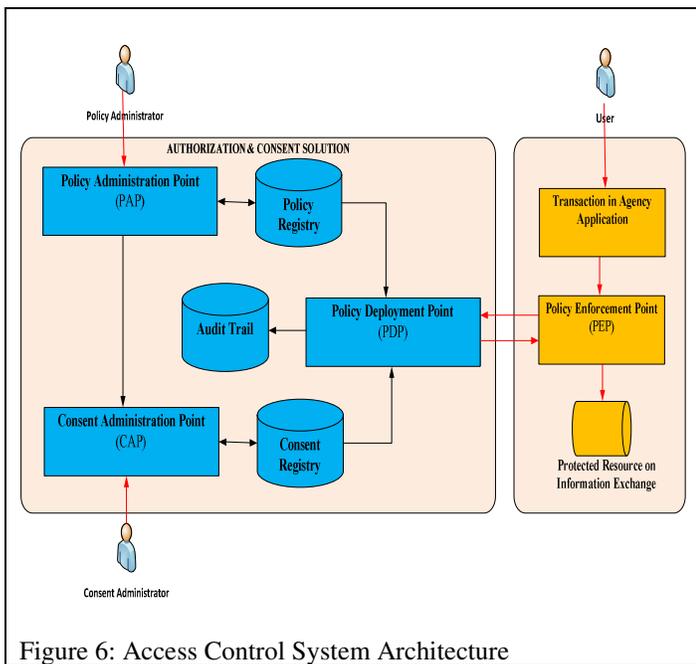


Figure 6: Access Control System Architecture

VI. IMPLEMENTATION NOTES

The ConnectWellSD information exchange now controls access to 14 categories of data acquired from 9 different source systems. Nearly 5000 users, a major part of the County’s workforce, use this system every day. The information exchange is used to drive Referrals and Collaborative Service Team (CST) activity. The system is being used to support the California Whole Person Care initiative, for which it’s ideally suited. The system is also used by County Eligibility Operations to process applications for public assistance programs.

Upcoming initiatives include opening this information exchange to local and State partners, and establishing collaboration across the larger ecosystem. One of the challenges is to arrive at an information exchange schema that can work across enterprises, and a common definition for resources on the Object side of the access control model. The County of San Diego is looking to use the National Information Exchange Model (NIEM) for this purpose.

VII. RETROSPECTIVE OBSERVATIONS

While this model is presented in a structured way, the work to arrive at this model was not always linear. Key elements of the development process that enabled the creation of this model include the following:

- Use Cases and Scenarios. Before there were any discussions of technology, the County of San Diego HHSa defined a common service delivery business process or service flow. This common set of steps was used to then select the initial collaborative transactions that the information would support. These use cases helped provide a focal point for

discussions about data sharing, user access, consent, and composition of data view. The use cases helped constrain the working discussions and also provide concrete ways in which the data would be used.

- Integrated workgroups. Defining what data should come together for an integrated view will only have utility if there is an understanding of how the data will be applied to decision-making, priority setting, and case planning. The County of San Diego HHSa established an integrated workgroup with representatives from all of the disciplines (departments/programs) that would ultimately use the system. Through working sessions where each source system was displayed screen-by-screen, representatives identified which fields would support the delivery of the use cases and how the field would be put to use. Capturing this discussion helped support decision making and design related to access control.
- Abstracting the current business organization into business process patterns.¹ The current business structure of ‘departments’ and ‘programs’ provided broad categorization but was not sufficient to define the dimensions of user access. Developing definitions of common business process patterns among groups of jobs that would apply to all ‘departments’ and ‘programs’ was necessary to define the dimensions of user access with enough depth and flexibility to work across a broad and varied user group.
- Separating questions being asked of the data. Initial discussions about access to integrated data ended in a holding pattern because of an underlying assumption that a customer look up would return information that could answer any question that a system user may pose to the data. By separating the questions and actions into granular units, progress was possible because a logical organization of objects and access began to emerge. For example, when looking a customer up in the system, only the question, ‘Is this person known?’ is being asked. Therefore, only data that answers that question is returned (example: first name, last name, address, phone number, and gender).

¹ Polya, George (1945) How to Solve It. Princeton University Press