

Towards a Health-Care Reference Architecture

Paul C. Brown, Ph.D.
Principal Software Architect
TIBCO Software Inc.
paul.brown@healthcarereferencearchitecture.org

John Kelly, MHA
Senior Technology Consultant
Harvard Pilgrim HealthCare
john.kelly@healthcarereferencearchitecture.org

David Querusio
Manager, Harvard Pilgrim HealthCare Integration Architecture
Perot Systems
david.querusio@healthcarereferencearchitecture.org

1 Abstract

The structure of the health care industry is in flux. In seeking more cost-effective solutions, enterprises are constantly experimenting with different business models. They are trying various combinations of roles, shuffling the deck and dealing different mixtures of risk holder, care provider, plan manager, and plan administrator. The health plans themselves are also becoming increasingly complex.

Comprehensive plans frequently end up being aggregates of smaller plans, each covering a different category of benefits that carries with it a different business model for providing and administering the care. From both a business and information systems perspective, the result is chaos – expensive chaos.

Part of the problem is that there is little standardization in the interactions between various roles involved in health care. The existing HIPAA standards focus basically on the dialog between health care providers and payers, leaving the interactions between other roles to be negotiated and customized for each business arrangement. This drives up the cost of providing health care and presents a barrier to the evolution of the health care system as a whole.

Fortunately, there is some natural order in this chaos – order that can be leveraged to manage these costs while at the same time preserving needed flexibility. The order is found in the relatively constant roles that are played within each health care business process, despite the endless organizational variations. What changes from situation to situation is the grouping of these roles into business units.

Standardizing these roles and the interactions between them will enable different and changing business models without entailing the added cost of negotiating new inter-company interactions and implementing their supporting system interfaces. Standardization will also make it possible to track and report the status of transactions as they traverse the various roles, setting the stage for improved response times and a corresponding improvement in the quality of care. This paper illustrates this approach using simplified examples drawn from the financial aspects of health care. Its purpose is to share the concept as a step towards determining the level of interest in the health care community for carrying this work forward into the development of a full-scale health care reference architecture.

2 Motivation

The structure of the health care industry is in flux. In seeking more cost-effective solutions, enterprises are constantly experimenting with different business models. They are trying various combinations of roles, shuffling the deck and dealing different mixtures of risk holder, care provider, plan manager, and plan administrator. At one extreme, health maintenance organizations (HMOs) cover the entire spectrum of health care from providing services through bearing financial risk. At the other extreme, the roles are totally fragmented. Providers, provider networks, processing intermediaries, health care plan actuaries, plan administrators, claim adjudicators, and risk holders each become independent business entities.

Complicating the situation is the fact that the health care plans themselves are becoming increasingly complex. Comprehensive plans frequently end up being aggregates of smaller plans, each covering a different category of benefits that carries with it different business model for providing and administering the care. Vision, dental, medical and psychiatric services are often completely separated, from the provider networks through to the claims processing. Some employers obtain these plans from different parties and live with the administrative fragmentation, while others obtain a unified plan from a single company that tries to present a unified front-end to an otherwise disjoint processes. From both a business and information systems perspective, the result is chaos – expensive chaos.

Part of the problem is that there is little standardization in the interactions between various roles involved in health care. The existing HIPAA standards focus basically on the dialog between health care providers and payers, leaving the interactions between other roles to be negotiated and customized for each business arrangement. HIPAA was intended to “support the electronic exchange of administrative and financial health care transactions primarily between health care providers and plans.”¹ But even where HIPAA specified interfaces are in use, they are almost always extended by additional agreements between the parties – agreements that define both the data and business rules surrounding the interactions. In practice, nearly every interface between every pair of parties today needs to be individually specified and implemented. This drives up the cost of providing health care and presents a barrier to the evolution of the health care system as a whole.

Fortunately, there is some natural order in this chaos – order that can be leveraged to manage these costs while at the same time preserving needed flexibility. The order is found in the relatively constant roles that are played within each health care business process, despite the endless organizational variations. In every instance, care is given and claims are prepared, submitted, routed, accepted, adjudicated, and paid. What changes from business model to business model are the role assignments - which business entity plays each role. While occasionally some new roles do emerge (generally these are intermediaries assuming portions of administrative responsibilities), these emerging roles generally result from subdividing an existing role.

This relative stability of roles and the interactions between them presents a significant opportunity to bring some stability to the chaos. Standardizing the roles and standardizing the interactions between them will enable different and changing business models without entailing the added cost of negotiating new inter-company interactions and implementing their supporting system interfaces. Standardizing the roles will further allow systems to be developed that can provide good support for individual roles. This will reduce the administrative cost of providing health care services while preserving the flexibility to evolve the related business structures.

Standardizing the roles makes it possible to implement a standard approach for tracking and reporting health care transactions as they progress from role to role and business to business. Standardized interfaces for querying the transaction status and reporting status changes will provide the transaction visibility required to manage overall transaction response times. Management will bring attention to

¹ Health Insurance Portability and Accountability Act of 1996 (P.L. 104-191 - known as HIPAA)

transactions that are languishing so that they can be driven towards a quicker resolution. This, in turn, will lead to a corresponding improvement in the quality of care being provided.

This paper illustrates this approach using simplified examples drawn from the financial aspects of health care. It examines the roles involved in the settling of a health care claim and illustrates the variety of ways in which these roles can be played by different parties. The interactions between these roles are examined with an eye toward standardizing the interfaces. The claim is that these interfaces are appropriate regardless of whether the interactions themselves are internal within the enterprise or external between enterprises. We also show how the abstracting of certain functions as services can further enhance the flexibility of these processes, and how the monitoring of process execution can reduce administrative costs.

The purpose of this paper is to share this concept and solicit feedback. We are interested in determining the level of interest in the health care community in carrying this work forward to develop a full-scale health-care reference architecture. Towards this end, we invite your comments – and your participation. We are interested in knowing whether you would like to invest some of your time in moving this effort forward, and welcome your suggestions as to how, exactly, to do this.

3 Challenges

3.1 The Ever-Changing Relationships between Roles and Business Entities

The roles played by different parties in the health care world are evolving. **Error! Reference source not found.** shows the roles typically associated with a provider and an insurer with respect to providing a service and settling the resulting health care claim. While historically there have been only two participants in this process, in today's world there are often many more. Providers utilize services to prepare their claims, and groups of providers form networks that submit claims on their behalf. Complicating matters even further, the claims preparers work for many service providers, and claims submitters submit claims on behalf of many providers through many claims preparers.

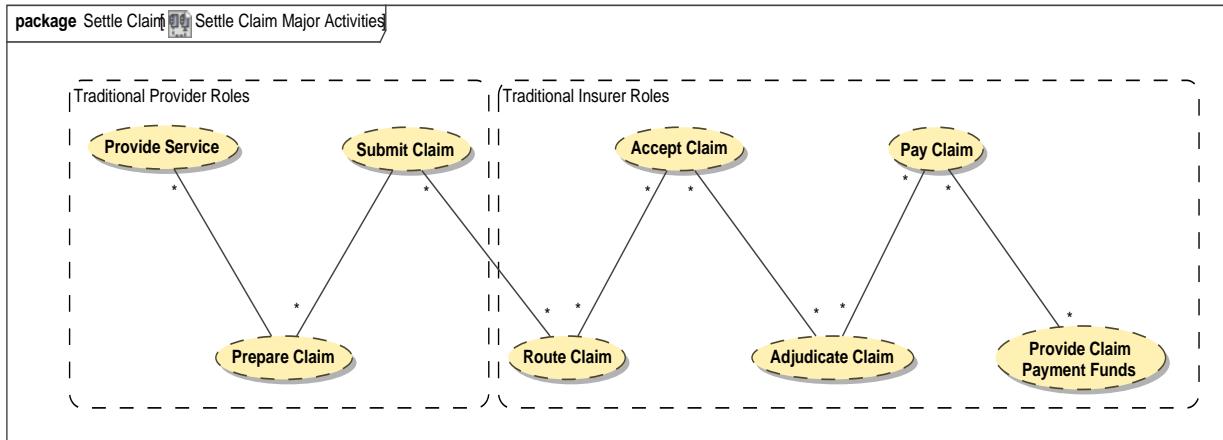


Figure 1: Traditional Provider-Insurer Roles

On the insurer side, similar splits of responsibility can be found. Some insurance “nameplates” have a single point of entry for submitting claims, then route the claims to different claims acceptors and adjudicators. Other nameplates outsource some or all of these responsibilities. Some insurers write checks against accounts in traditional banks, while others have, themselves, become banks. Even the risk holder, the entity providing the payment funds, might be an insurance company or it might be an employer. In a shared-risk situation, there may even be more than one entity providing the funds, or a reinsurer involved.

3.2 Lack of Standard Interfaces Between Roles

In examining the interactions between these roles, we find that there is only one standardized interface – the one corresponding to the traditional boundary between the provider and the insurer (Figure 2). But even here the interface is not truly standardized. The HIPAA transaction specifications are not complete either with respect to the mechanics of communications or the information content. The parties involved must negotiate the mechanics of transport and the security surrounding the interactions. In addition, the parties almost always have their own rules regarding the use of this interface and the data values carried in the messages. As a result, it is not unusual for a party to have a 100-page companion document for each HIPAA interface specifying its specific usage rules.

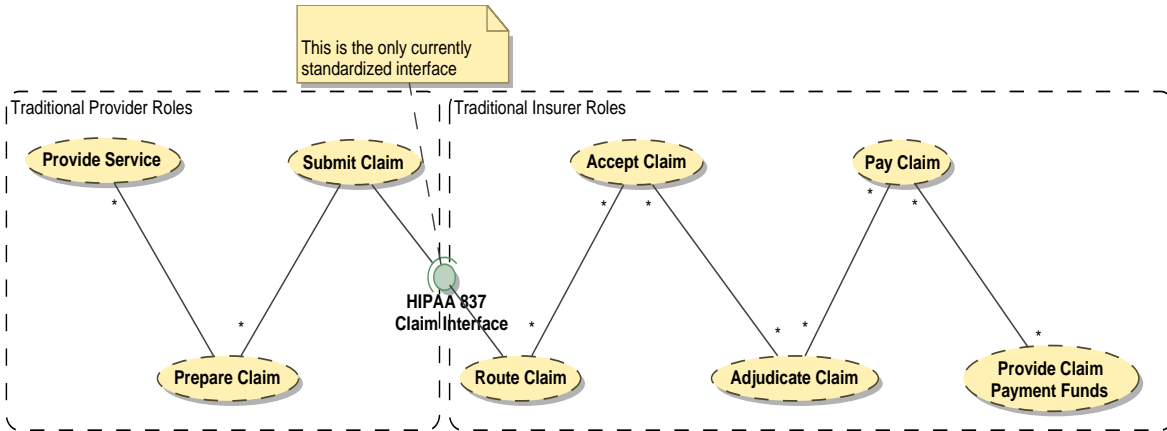


Figure 2: Interfaces

As far as the other role interactions are concerned, each party must develop its own rules and standards for these interfaces. In terms of their information structure, many of the data structures passed back and forth are loosely based on existing legacy system interfaces. Note that while these data structures are custom, they largely contain the HIPAA-required information augmented with data to support the dialog between the parties.

One common feature of the data structure extensions at both the HIPAA interfaces and the other interfaces is that much of the additional information is not related to the medical and insurance aspects of the claim – it is related to the roles, the parties playing them, and the routing of claims between them. If a `claims acceptor` receives a claim from a `claims router` and forwards it to the `claims adjudicator`, how is the `claims adjudicator` to know where to send the response? The identities of the intermediaries involved in processing the claim must somehow be known so that the response can be appropriately routed. Many of the rules regarding the use of interfaces simply establish conventions that indicate, either directly or inferentially, who the various parties are.

3.3 Replicated Functionality with Inconsistent Business Rules

Another cost driver in the health care community is the recurring need for the same logical function at various points both within a given business process and across different business processes. Figure 3, for example, shows a number of places where member eligibility can be checked and services priced in today's business processes. In some cases this function is also provided as a stand-alone business process, as when a HIPAA 270 Eligibility Check is submitted, but even here this same function is also required when a service is authorized and when a claim is adjudicated. However, despite the similarity in functional need, the actual implementation is often replicated differently and in a channel-specific manner. This replication leads to replicated business rules, partial and inconsistent implementations, and increased maintenance costs.

Towards a Health Care Reference Architecture

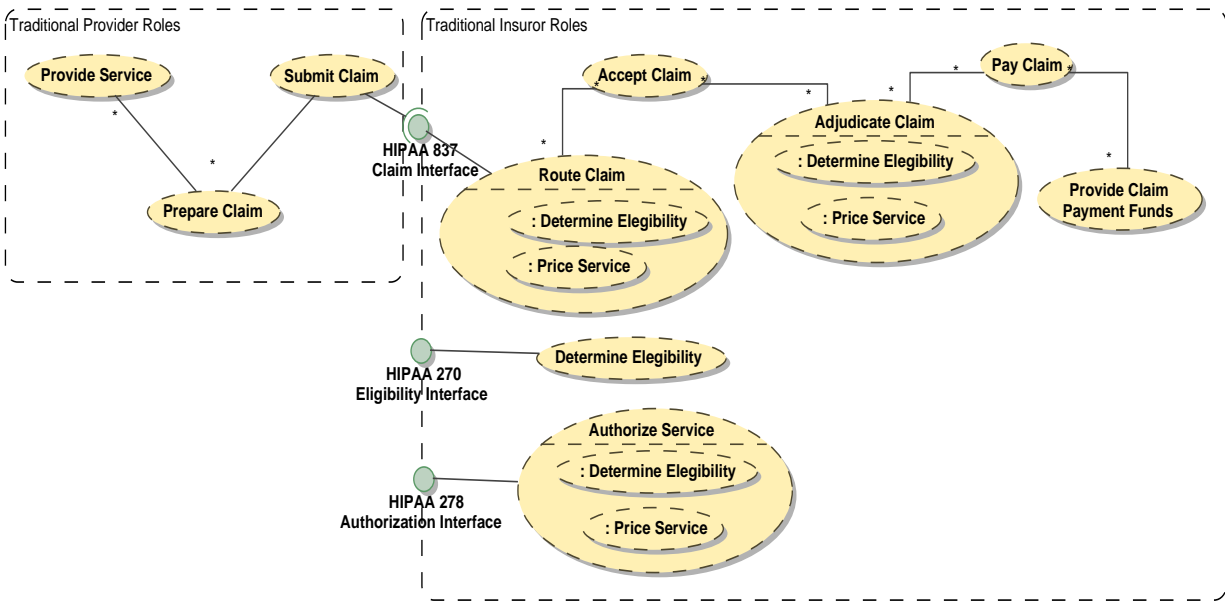


Figure 3: Uses of Determine Eligibility and Price Service Functionality

Things get even more complicated when roles are split between different business parties. When this occurs, the same function may be required by more than one party. This requires either that the parties replicate the function (along with its reference data and business rules) or that one of the parties provide an interface for the other to access the function. Neither option is attractive and both options require extra work.

The function of determining eligibility is a good case in point. An eligibility determination requires a significant amount of information about the health care policy, its rules regarding benefits, and the status of the member with respect to that policy. When a health benefits company outsources the acceptance and adjudication of claims, it is faced with the challenge of communicating this complex information in a form that can be effectively used by the other party. The other party is challenged to implement the business rules accurately.

The complexity of communicating the information and business rules is not to be underestimated. There are no standards for representing policies, benefits, and limitations, let alone standards for interpreting this information to answer eligibility questions (Figure 4). The parties involved in the business processes requiring eligibility checks are further constrained by their existing systems: their ability to represent and work with complex eligibility information and rules is defined by the systems they happen to have in house.

Towards a Health Care Reference Architecture

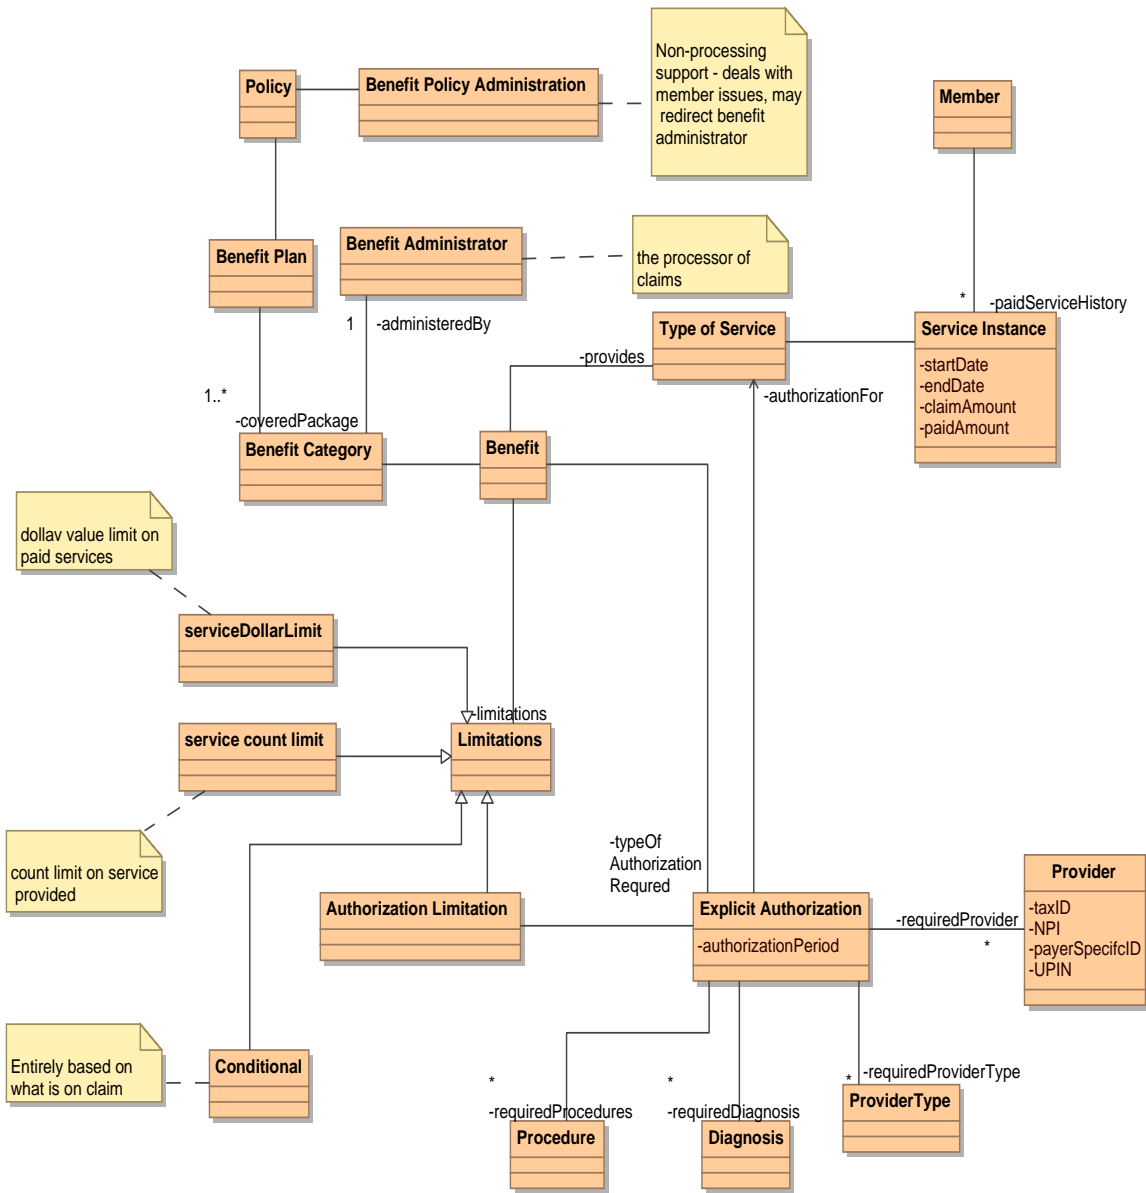


Figure 4: Sketch of Information Typically Required for Eligibility Checks

Ideally, eligibility checks are performed in the initial claim routing, before the claim is even sent for adjudication. This is particularly beneficial when there are different downstream channels for different categories of benefits. The centralized eligibility checks relieve the claims adjudicator of this responsibility and simplify the resulting interactions between the parties. Similar arguments apply to the pricing of services.

Standing in the way of this ideal is an unfortunate reality: when there are different downstream channels for different categories of benefits, the downstream parties often impose their own business rules and define their own datasets and logic to support these rules. These constraints are often driven by the limitations of the systems being used by these parties. The consequence is often different (and apparently arbitrary) rules governing different categories of benefits.

These differing downstream constraints make it difficult to implement upstream eligibility checks. The result is a limited upstream capability, often as simple as just determining whether the member was covered by the policy at the time of service. The final eligibility determination is not made until the claim is adjudicated. Discovering an eligibility problem at this late stage delays the processing of the claim and requires a more complicated feedback path. It is hardly an optimal solution for any of the parties involved.

3.4 Monolithic Health-Care Applications Span Multiple Roles

Another burden on the health care industry is created by some of the commercially available health-care applications. These applications often make inappropriate assumptions about which roles are being supported by the software. These commercial off-the-shelf (COTS) products often assume that the entire claims settlement process - claim routing, acceptance, and adjudication – is being managed and executed by one party and one piece of software. This makes it difficult for a company that wishes to outsource aspects of its operations. The systems do not support integration with other systems, nor are they capable of managing part, but not all, of the business process. Even when interfaces are available, they are proprietary and their use requires custom development on the part of the business partner. In many cases, the desired interface is actually an internal interaction within the software and is not accessible at all. For example, while many COTS health-care products determine eligibility and price services, few provide the interfaces needed to invoke these functions independently of actual claim processing.

4 The Health Care Reference Architecture Concept

Despite the ongoing experimentation with business structures, the roles being played in the health care business processes continue to remain relatively stable. This role stability can be leveraged to modularize the business processes by standardizing the roles and the interfaces between them.

4.1 Standardized Roles with Standardized Interfaces

The settlement of health care claims provides good examples of the role and interface standardization we are referring to (Figure 5). On the provider side, every provider must submit claims. This is a process which, in the end, is ultimately a business-to-business interaction between the provider and the insurer. This process requires that the claim be prepared and submitted, but there are many options for playing these roles. They may be performed by the provider directly, by one or more contracted parties, or by the provider and contracted parties. On the insurer side, the routing, claim acceptance, claim adjudication, claim payment, and funds provisioning may also be performed by various combinations of parties.

Towards a Health Care Reference Architecture

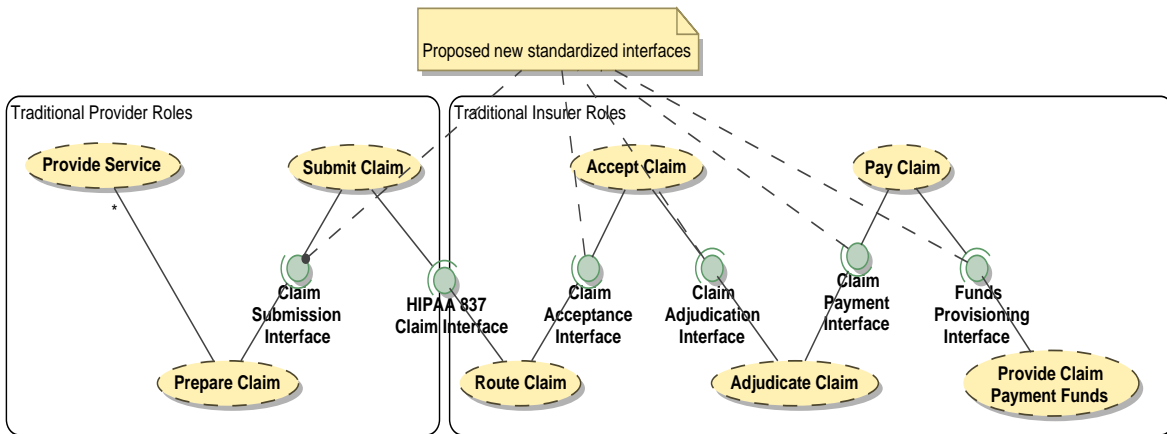


Figure 5: Standardized Roles and Related Interfaces

Standardizing the roles and the responsibilities that go along with them affords an opportunity to standardize interfaces between them. The challenge in standardizing these interfaces lies not so much in defining the operations that need to be provided but rather in standardizing the data that needs to be exchanged. The key to successfully standardizing the data lies in making the standard data structures extensible in a carefully controlled manner.

On the surface, standardization and extensibility seem to be incompatible. The key to making extensibility work is to recognize that the information being exchanged, for the most part, has an inherent organization that reflects its role and use in the world of health care. This overall organization remains the same regardless of who is playing the roles involved. The trick is to understand this organization and standardize its structure, and then define how the structure can be extended without modifying this basic organization.

Consider the concepts related to a health care claim as specified in a HIPAA 837 claim submission. Figure 6 shows the structure of this information. This HIPAA model requires information about providers to be included, but constrains the specific roles that providers can play. The same is true for other entities such as the *Submitter* and *Payer*. This model has no provisions for extensibility with respect to other roles in the claims submission process.

Towards a Health Care Reference Architecture

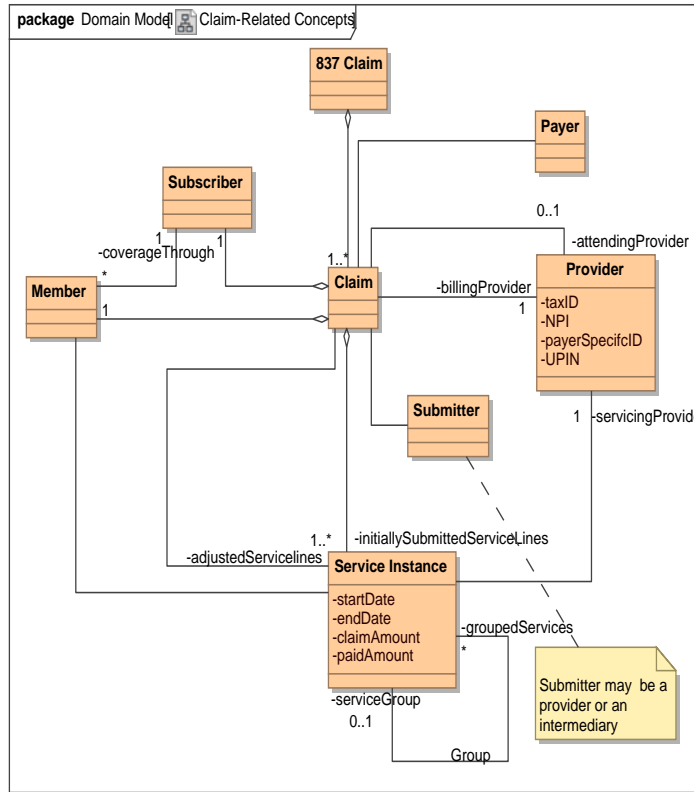


Figure 6: Claim-Related Concepts As Represented in a HIPAA 837

Towards a Health Care Reference Architecture

Now consider the more general model of this same information shown in Figure 7. This model still allows the HIPAA-specific roles for providers and other parties to be specified, but it allows other roles to be specified as well. In this example, the role of `claim preparer` has been added as an administrative role type. To support variability in the downstream processing of the claim, the administrative roles of `claim router`, `claim acceptor`, and `claim adjudicator` could be added as well.

Abstracting roles and role types is one way in which extensibility can be provided. Another is to make the set of attributes for a given concept open-ended. The core data structure specifies the standard fields, but allows other fields to be added as well. Support for this type of extensibility is inherent in the SOAP and supporting XML standards. The extensibility can be achieved through the appropriate use of XML “any” constructs, and the use of these extensions can be constrained in a well-defined and verifiable way through the use of XML namespaces.

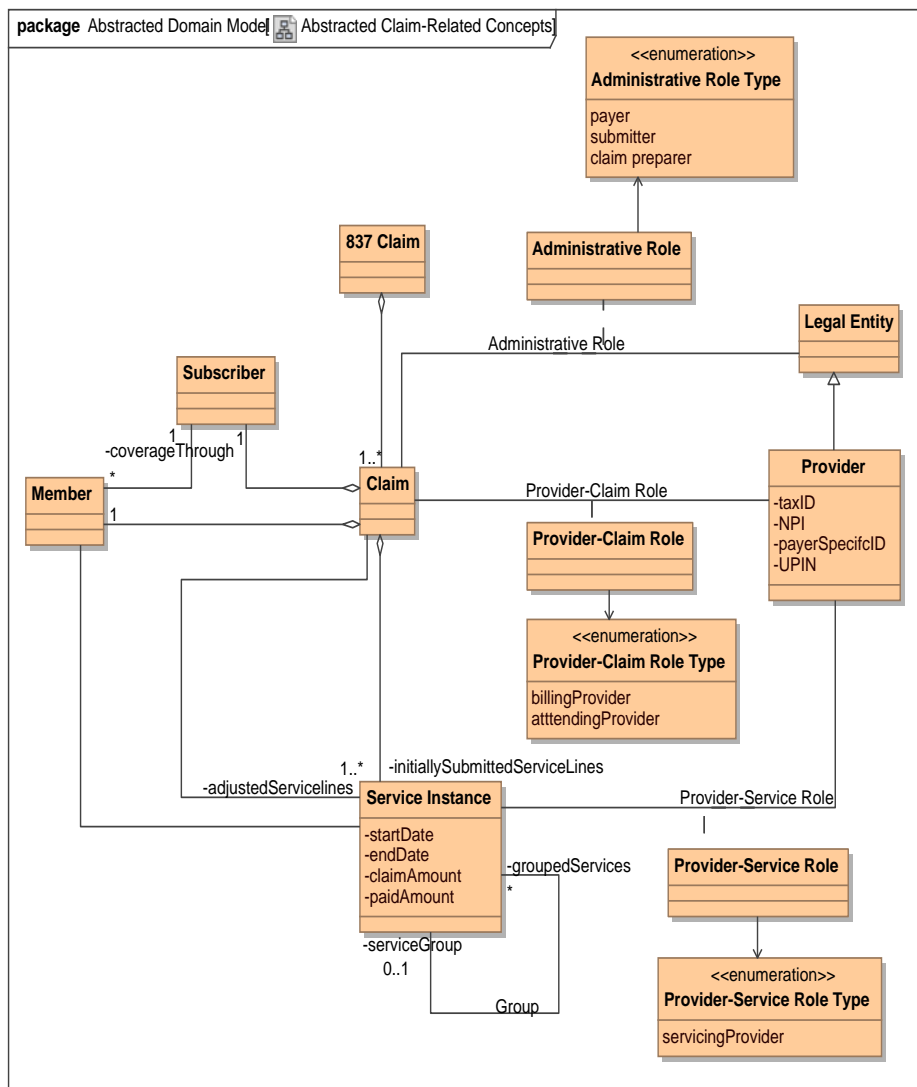


Figure 7: Abstracted Model of Claim-Related Concepts

4.2 Abstracting Common Functionality as Services

As we observed earlier, there are a number of functions that are likely appear at various places both within a given health care process and in different processes. Determining eligibility and pricing services are the two examples that were mentioned earlier. If these functions are provided as stand-alone services, then further flexibility in the implementation of these health care business processes can be achieved by standardizing these service interfaces and using these services (Figure 8). Once this is done, any of the participants in the process can be a provider of these services without modifying the business process at all. Another degree of flexibility has been obtained.

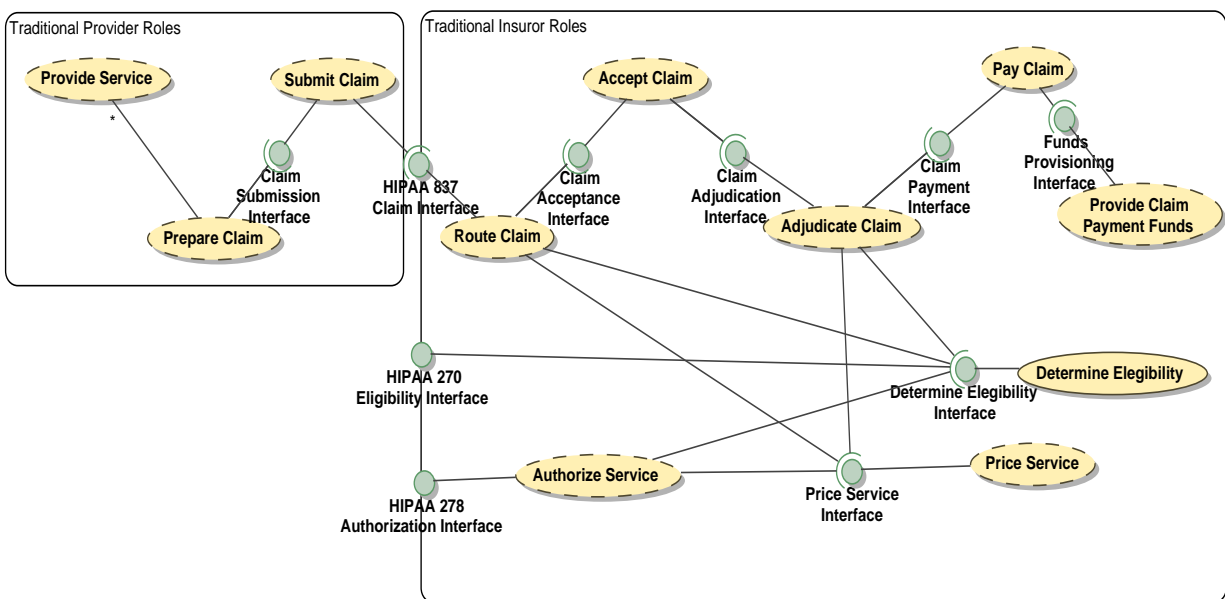


Figure 8: Use of Price Service and Determine Eligibility Services

4.3 Process Status Monitoring

Perhaps one of the biggest challenges in today's health care business processes results from the distribution of the work among multiple parties. The processing of a claim now becomes a distributed and possibly long-running multi-party transaction. This distribution introduces complexity in determining status of the process and detecting breakdowns in the process. The problem is that the actual status information is distributed among the different parties. Furthermore, due to the routing of the claim between the parties it may not even be apparent which parties have the required status information.

Tracking transactions across multiple parties is very difficult in today's health care business processes. Individual parties typically do not provide visibility into their internal processes, and no single party knows the status of an in-progress transaction. Questions about transaction status require a two-stage

investigation, first to determine which party currently has responsibility for the transaction, and second to determine that party's actual status. Tracking is expensive and time consuming in today's processes.

A solution to this problem is to introduce a new role, that of a *process monitor*, for each major process. Figure 9 shows an example process monitor for the claims settlement process. This role, which will likely be played by one of the parties already involved in the process (the claims router, in particular), provides two SOAP service interfaces. The *Claim Status Reporting Interface* provides a means by which responsibility assignments can be reported and participants can report processing status changes. The *Claim Status Query Interface* provides a means for interested parties to obtain the status of a claim.

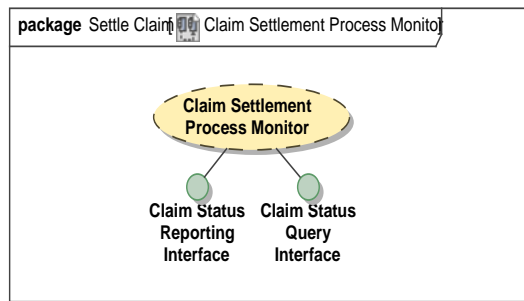


Figure 9: Claim Settlement Process Monitor

The status that is tracked and reported by this monitor is a milestone-level status that is abstracted from the processing details, thus allowing the status itself to be standardized. Figure 10 presents a simplified example of such a milestone life cycle for the processing of a claim. Each milestone represents a possible answer to the "What's the status?" question.

Towards a Health Care Reference Architecture

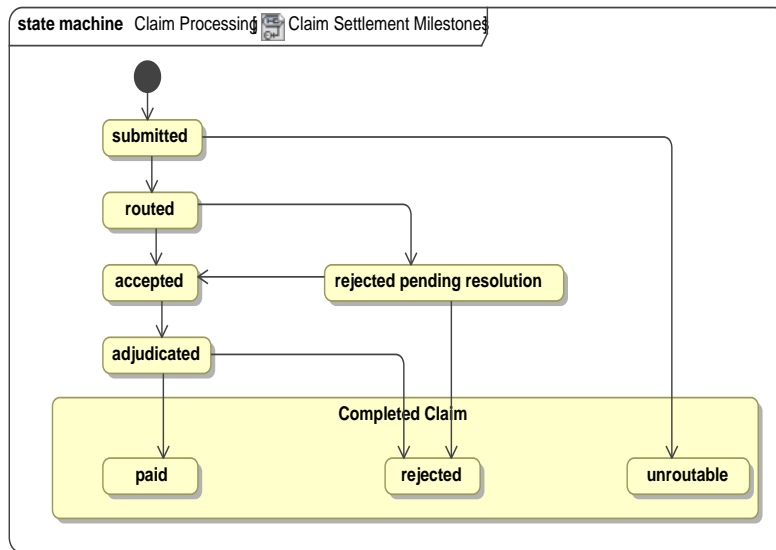


Figure 10: Simplified Claim Processing Milestones

The process monitor also keeps track of which participants have been assigned responsibility for executing specific processing roles. A query for the status of a claim would return a data structure similar to that shown in Figure 11.

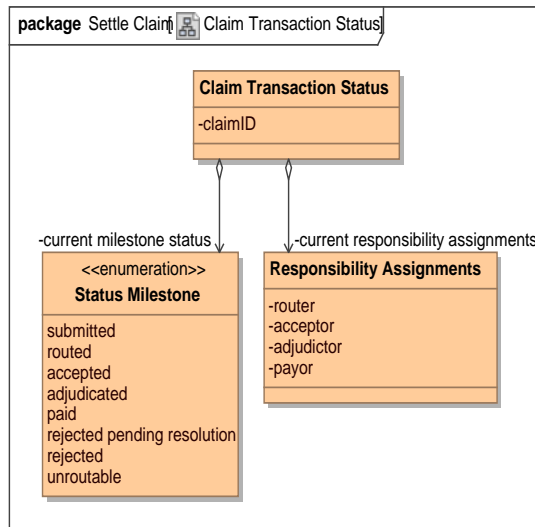


Figure 11: Example Claim Transaction Status

An alternative to central status tracking is to have each of the parties support a status query interface. With minor additions to the `Status Milestone` enumeration, individual parties could implement the `Claim Status` interface to provide external visibility of the claim transaction status. This would provide the ability for other participants in the process (the claims router, in particular) to determine the status of a claim in a standardized and automated manner.

5 Reference Architecture Concept Summary

The intent of the health care reference architecture is to standardize the partitioning of health care processing into standardized roles without constraining how those roles are actually implemented. The completed fragment of the architecture incorporating all the ideas we have been discussing is shown in Figure 12.

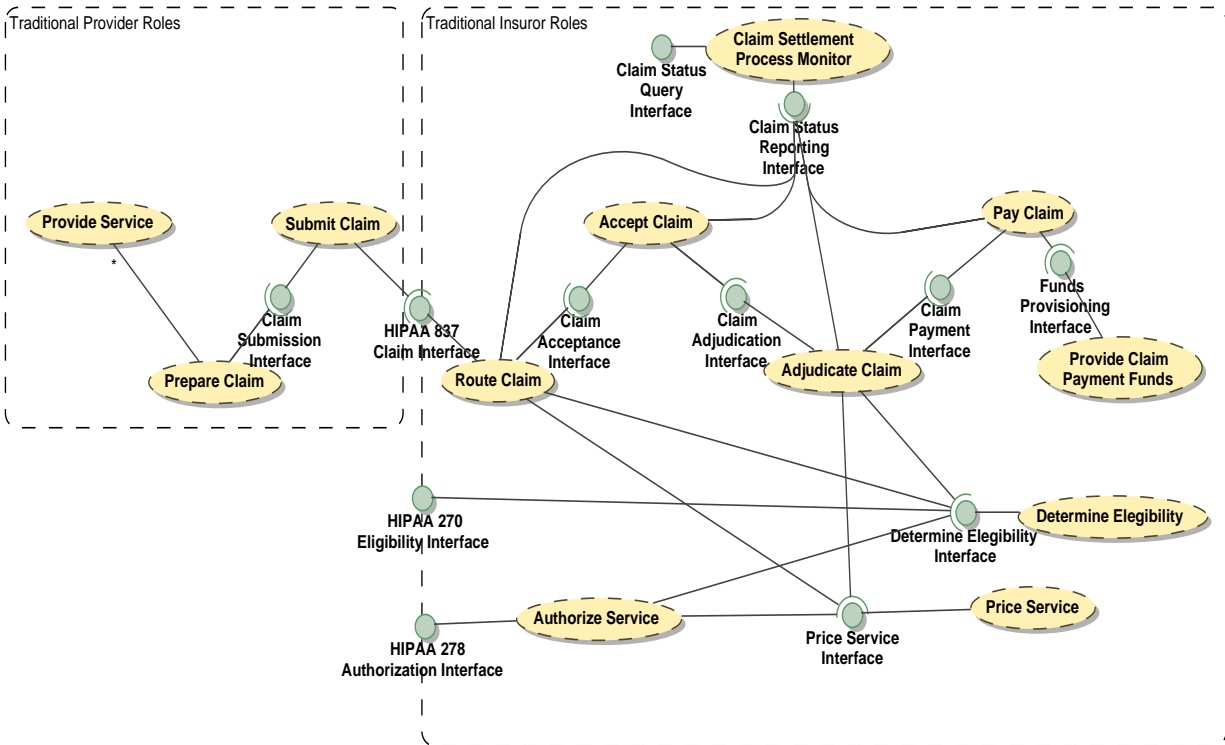


Figure 12: Completed Reference Architecture Fragment

Figure 13 shows an example in which two parties are cooperatively providing a complete health care plan. The “Nameplate” company provides full-service coverage for all benefits except vision care, which is handled by a separate company. All claims submissions are presented to the “Nameplate” company, which routes vision claims to the Vision company and routes the rest to its internal processing. Note the use of the Funds Provisioning Interface by the Vision company to replenish the account from which vision claims are paid. A full reference architecture will define the interfaces required to support all aspects of day-to-day inter-company operations.

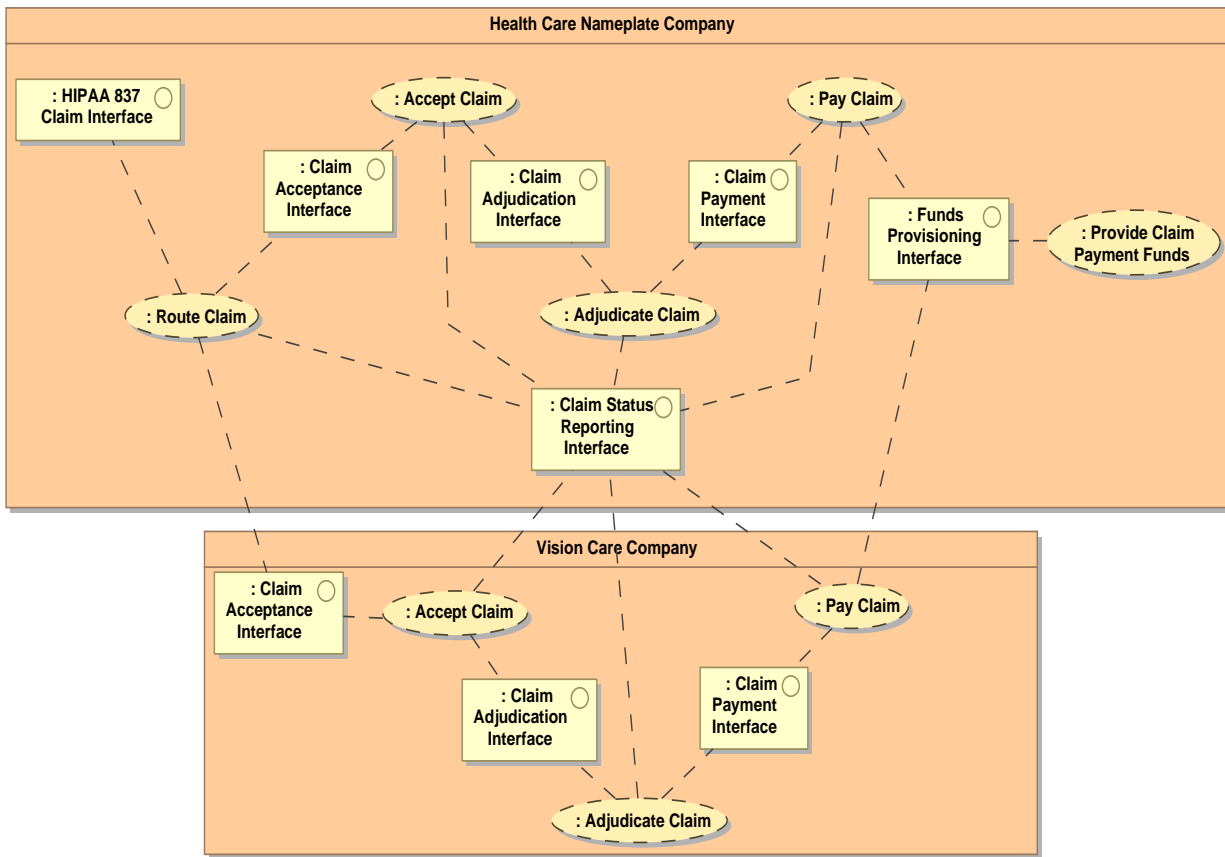


Figure 13: Example of Two Cooperating Parties

5.1 Flexibility in Business Roles

This style of reference architecture standardizes the roles in a business process and standardizes the interactions between those roles without constraining in any way the parties that play the roles. Such an architecture makes it possible for an insurance provider to retain all of the roles for medical care and, at the same time, outsource all but the claim routing role for vision and dental care. Furthermore, subsequent changes in role assignments (either outsourcing or bringing roles back in-house) no longer require the development of new interfaces.

5.2 Reduced Administrative and Development Costs

Adopting standardized interfaces between roles will significantly reduce development costs as it will no longer be necessary to develop custom interfaces for each partner relationship. Extensions, where required, can be accommodated within the standardized data structures and interfaces while being managed and controlled through the use of XML namespaces.

The use of standardized interfaces will reduce administrative costs in two areas. First, the use of SOAP and XML technology will enable more errors to be caught on the originating side of the interface, thus

reducing the number of partner interactions required to resolve problems. Second, the process monitoring function centralizes the monitoring of the distributed process. This simplifies obtaining the status of the work in progress and enables the timely detection of breakdowns in the process.

5.3 Opportunities for Role-Specific COTS Products and Services

The standardization of roles and interfaces also presents a significant opportunity for both software vendors and service providers. Software vendors can build applications to these interfaces with assurance that they will interoperate with other elements of the health care process. Service providers can implement specific roles with the confidence that their services will be easy for other participants to leverage.

6 Moving Forward

The three of us and our respective companies are very interested in your feedback regarding this concept for a health care reference architecture. Such an architecture can only succeed if it becomes a consensus-driven collaboration of participants in the health care community. Please share your thoughts with us at www.healthcarereferencearchitecture.org and let us know whether you would be willing to join this effort.