Tilburg University

Service-oriented design and development methodology

Papazoglou, M.; van den Heuvel, W.J.A.M.

Published in:
International Journal of Web Engineering and Technology

Publication date:
2006

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Papazoglou, M., & van den Heuvel, W. J. A. M. (2006). Service-oriented design and development methodology. *International Journal of Web Engineering and Technology*, 2(4), 412-442.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Service-Oriented Design and Development Methodology

Michael P. Papazoglou*

INFOLAB, Department of Information Systems and Management,
Tilburg University, PO Box 90153, Tilburg 5000 LE, The Netherlands
E-mail: mikep@uvt.nl

*Corresponding author

Willem-Jan van den Heuvel

INFOLAB, Department of Information Systems and Management,
Tilburg University, PO Box 90153, Tilburg 5000 LE, The Netherlands
E-mail: wjheuvel@uvt.nl

Abstract: SOA is rapidly emerging as the premier integration and architectural approach in contemporary complex, heterogeneous computing environments. SOA is not simply about deploying software: it also requires that organizations evaluate their business models, come up with service-oriented analysis and design techniques, deployment and support plans, and carefully evaluate partner/customer/supplier relationships. Since SOA is based on open standards and is frequently realized using Web services, developing meaningful Web service and business process specifications is an important requirement for SOA applications that leverage Web services. Designers and developers cannot be expected to oversee a complex service-oriented development project without relying on a sound design and development methodology. This paper provides an overview of the methods and techniques used in service-oriented design and development. Aim of this paper is to examine a service development methodology from the point of view of both service producers and requesters and review the range of elements in this methodology that are available to them.

Keywords: service oriented computing; service oriented architecture; business processes; web services; design and development methodologies.

Biographical notes: Michael P. Papazoglou is a professor of computer science and director of the INFOLAB at the University of Tilburg in the Netherlands. His research interests include distributed systems, service-oriented computing and Web services, enterprise application integration, and e-Business technologies and applications. He received PhD in computer systems engineering from the University of Edinburgh.

Willem-Jan van den Heuvel is an associate professor of information Systems at the University of Tilburg in the Netherlands. His research interests include service-oriented computing, alignment of new enterprise system with legacy systems, and system evolution. He received a PhD in computer science from the University of Tilburg.

1 INTRODUCTION

SOAs provide a set of guidelines, principles and techniques in which business processes, information and enterprise assets can be effectively (re)organized and (re)deployed to support and enable strategic plans and productivity levels that are required by competitive business environments [Papazoglou 2003]. In this way, new processes and alliances need to be routinely mapped to services that can be used, modified, built or syndicated.

To achieve such business requirements, the internal architecture of an SOA needs to evolve into a multi-tier,

service-based system, often with a diversified technical implementation [Cauldwell 2001]. This diversity is the result of a very broad spectrum of business and performance requirements as well as different execution and reuse contexts. As a consequence, older software development paradigms for object-oriented and component-based development cannot be blindly applied to SOA and Web services.

Many enterprises in their early use of SOA, suppose that they can port existing components to act as Web services just by creating wrappers and leaving the underlying component untouched. Since component methodologies

focus on the interface, many developers assume that these methodologies apply equally well to service-oriented architectures. As a consequence, implementing a thin SOAP/WSDL/UDDI layer on top of existing applications or components that realize the Web services is by now widely practiced by the software industry. Yet, this is in no way sufficient to construct commercial strength enterprise applications. Unless the nature of the component makes it suitable for use as a Web service, and most are not, it takes serious thought and redesign effort to properly deliver components functionality through a Web service. While relatively simple Web services may be effectively built that way, a service-based development methodology is of critical importance to specify, construct, refine and customize highly volatile business processes from internally and externally available Web services.

In this paper we concentrate on the workings of a services design and development methodology that provides sufficient principles and guidelines to specify, construct and refine and customize highly volatile business processes choreographed from a set of internal and external Web services.

2 CHARACTERISTICS OF SERVICE DEVELOPMENT LIFE CYCLE METHODOLOGY

A Web Services Lifecycle Development [Papazoglou 2006] methodology should focus on analyzing, designing and producing an SOA in such a way that it aligns with business process interactions between trading partners in order to accomplish a common business goal, e.g., requisition and payment of a product, and stated functional and non-functional business requirements, e.g., performance, security, scalability, and so forth.

Service-oriented design and development incorporates a broad range of capabilities, technologies, tools, and skill sets, that include [Arsanjani 2004], [Brown 2005]:

- Managing the entire services lifecycle—including identifying, designing, developing, deploying, finding, applying, evolving, and maintaining services.
- Establishing a platform and programming model, which includes connecting, deploying, and managing services within a specific runtime platform.
- Adopting best practices and tools for architecting services-oriented solutions in repeatable, predictable ways that deal with changing business needs.
- Delivering high quality workable service-oriented solutions that respect QoS requirements. These solutions may be implemented on best-practices, such as tried and tested methods for implementing security, ensuring performance, compliance with standards for interoperability, and designing for change.

Fundamental to the above capabilities is that business goals and requirements should always drive downstream design, development, and testing to transform business processes into composite applications that automate and integrate the enterprise. In this way business requirements can be traced across the entire lifecycle from business goals, through software designs and code assets, to composite applications.

Service design and development is about identifying the right services, organizing them in a manageable hierarchy of composite services (smaller grained often supported larger grained), choreographing them together for supporting a business process. A business service or process can be composed of finer-grained services that need to be supported by infrastructure services and management services such as those providing technical utility such as logging, security, or authentication, and those that manage resources.

Classifying related business processes that exhibit common functional characteristics and objectives can raise the level of abstraction in an SOA. In this way, business process conglomerations can be created and organized under a service domain. Classifying business processes into logical service domains simplifies an SOA by reducing the number of business processes and services that need to be addressed. A *service domain*, also referred to as *business domain*, is a functional domain comprising a set of current and future business processes that share common capabilities and functionality and can collaborate with each other to accomplish a higher-level business objective, e.g., loans, insurance, banking, finance, manufacturing, human resources, etc. In this way a business can be partitioned into a set of disjoint domains. Such domains can be leveraged from multiple architectural reasons such as load balancing, access control, and vertical or horizontal partitioning of business logic.

Figure 1 Web Services Development Life Cycle hierarchy.

Figure -1, shows that a service domain such as distribution is subdivided into higher-level business processes such as purchasing, order management and inventory. In this figure, the order management business process, which we shall use as a running example throughout the paper, performs order volume analysis, margin analysis, sales forecasting and demand forecasting across any region, product or period. It can also provide summary and transaction detail data on order fulfilment and shipment according to item, sales representative, customer, warehouse, order type, payment term, and period. Furthermore, it can track order quantities, payments, margins on past and upcoming shipments, and cancellations for each order. The order management process in Figure 1 is shown to provide business services for creating, modifying, suspending, cancelling, querying orders, and scheduling order activities. Business services can also create, modify, and delete bulk orders and order activities, while customers be informed of the progress of an order and its order activities. Business services in the order management process are used to create and track orders for a product, a service or a resource, and is used to capture the customer-selected service details. Information that is captured as part of an order may include customer account information, product offering and quality of service details, SLA details, access information, scheduling information, and so forth.

Business services are supported by infrastructure, management and monitoring services. These services provide the infrastructure enabling the integration of services through the introduction of a reliable set of capabilities, such as intelligent routing, protocol mediation, and other transformation mechanisms, often considered as part of the Enterprise Service Bus [Chappell 2004], [Papazoglou 2006]. This layer also provides the capabilities required for enabling the development, delivery, maintenance and provisioning of services as well as capabilities that monitor, manage, and maintain QoS such as security, performance, and availability. It also provides services that monitor the health of SOA applications, giving insights into the health of systems and networks, and into the status and behaviour patterns of applications making them thus more suitable for mission-critical computing environments. Monitoring services implement all important standards implementations of WS-Management and other relevant protocols and standards such as WS-Policy and WS-Agreement.

All service domains, business processes and services are automatically populated with financial and operational functions and data available from resources such as ERP, databases, CRM and other systems, which lie at the bottom of the service lifecycle development hierarchy. However, component implementation is an issue that can seriously impact the quality of available services. Both services and their implementation components need to be designed with the appropriate level of granularity. The granularity of components should be the prime concern of the developer

responsible for providing component implementations (see section-4.3).

3 WEB SERVICES DEVELOPMENT LIFE CYCLE METHODOLOGY BASELINE

This section presents the elements of a service-oriented design and development methodology that is partly based on other successful related development models such as the Rational Unified Process ([RUP 2001], [Kruchten2004]), Component-based Development [Herzum 2000] and Business Process Modelling [Harmon2003] and concentrates on the levels of the Web Services Development Life Cycle hierarchy depicted in Figure 2.

Figure 2 Phases of the service-oriented design and development methodology.

A service-oriented design and development methodology is based on an iterative and incremental process that comprises one preparatory and eight distinct main phases that concentrate on business processes. These are planning, analysis and design (A&D), construction and testing, provisioning, deployment, execution and monitoring. These phases may be traversed iteratively (see Figure 2). This approach is one of continuous invention, discovery, and implementation with each iteration forcing the development team to drive the software development project's artefacts closer to completion in a predictable and repeatable manner. The approach considers multiple realization scenarios for business processes and Web services that take into account both technical and business concerns.

The planning phase constitutes a preparatory phase that serves to streamline and organize consequent phases in the methodology. During the planning phase, the project feasibility, goals, rules and procedures are set and requirements are gathered. The analysis phase is based on a thorough business case analysis that considers various alternatives for implementing business processes, e.g., by

wrapping enterprise systems or acquiring new services, while the design phase aims at identifying and specifying Web services and business processes in a stepwise manner. Service construction and testing involves coding Web services and business processes using the specifications that were developed during the design phase. It also involves testing the coded services and processes for functional correctness and completeness as well as for interoperability. The service provisioning phase then enforces the business model for service provisioning, that is chosen during the planning phase. This activity encompasses such as issues service metering, service rating and service billing. Once the provisioning model has been established, the Web services may be deployed and advertised in a repository system. The final phase in the methodology deals with the execution and monitoring of Web services. This phase includes the actual binding and run-time invocation of the deployed services as well as managing and monitoring their lifecycle.

The workings of this methodology will concern us for the rest of this paper and will exemplified by means of the order management process that we introduced earlier.

4 SERVICE-ORIENTED DESIGN AND DEVELOPMENT PRINCIPLES

A service-oriented design and development methodology focuses on business processes, which it considers as reusable building blocks that are independent of applications and the computing platforms on which they run. This promotes the idea of viewing enterprise solutions as federations of services connected via well-specified contracts.

The course of designing a business processes goes through the stages outlined earlier. However, in order to design useful and reliable business processes that are developed on the basis of existing or newly coded services we need to apply sound service design principles that guarantee that services are self-contained and come equipped with clearly defined boundaries and service end-points to allow for service composability. Two key principles serve as the foundation for service- and business process design: service coupling and cohesion.

4.1 Service coupling

It is important that grouping of activities in business processes is as independent as possible from other such groupings in other processes. One way of measuring service design quality is coupling, or the degree of interdependence between two business processes. The objective is to minimise coupling, that is, to make (self-contained) business processes as independent as possible by not having any knowledge of or relying on any other business processes. Low coupling between business processes indicates a well-partitioned system that avoids problems of service redundancy and duplication.

Coupling can be achieved by reducing the number of connections between services in a business process, eliminating unnecessary relationships between them, and by reducing the number of necessary relationships - if possible.. Coupling is a very broad concept, however, and for service design can be organized along the following dimensions:

1. *Representational coupling*: Business processes should not depend on specific representational or implementation details and assumptions of one another, e.g., business processes do not need to know the scripting language that was used to compose their underlying services. These concerns lead to the exploitation of interoperability and reusability for service design. Representational coupling is useful for supporting interchangeable/replaceable services and multiple service versions.
2. *Identity coupling*: Connection channels between services should be unaware of who is providing the service. It is not desirable to keep track of the targets (recipients) of service messages, especially when they are likely to change or when discovering the best service provider is not a trivial matter.
3. *Communication protocol coupling*: A sender of a message should rely only on those effects necessary to achieve effective communication. The number of messages exchanged between a sender and addressee in order to accomplish a certain goal should be minimal, given the applied communication model, e.g., one-way, request/response, and solicit/response. For example, one-way style of communication where a service end point receives a message without having to send an acknowledgement places the lowest possible demands on the service performing the operation..

4.2 Service cohesion

Cohesion is the degree of the strength of functional relatedness of operations within a service. Service aggregators should create strong, highly cohesive business processes; business processes whose services and service operations are strongly and genuinely related to one another. A business process with highly related services and related responsibilities, and which does not do a tremendous amount of computational work, has high design cohesion. The guidelines by which to increase service cohesion are as follows:

1. *Functional service cohesion*: A functionally cohesive business process should perform one and only one problem-related task and contain only services necessary for that purpose. At the same time the operations in the services of the business process must also be highly related to

one another, i.e., highly cohesive. Consider services such as get “product price”, “check product availability”, and “check creditworthiness”, in an order management business process.

2. *Communicational service cohesion*: A communicationally cohesive business process is one whose activities and services use the same input and output messages. Communicationally cohesive business processes are cleanly decoupled from other processes as their activities are hardly related to activities in other processes.
3. *Logical service cohesion*: A logically cohesive business process is one whose services all contribute to tasks of the same general category by performing a set of independent but logically similar functions (alternatives) that are tied together by means of control flows. A typical example of this is mode of payment.

Like low coupling, high cohesion is a service-oriented design and development principle to keep in mind during all stages in the methodology. High cohesion increases the clarity and ease of comprehension of the design; simplifies maintenance and future enhancements; achieves service granularity at a fairly reasonable level; and often supports low coupling. Highly related functionality supports increased reuse potential as a highly cohesive service module can be used for very specific purposes.

4.3 Service Granularity

Service granularity refers to the scope of functionality exposed by a service. Services may exhibit different levels of granularity. An implementation component can be of various granularity levels. Fine-grained component (and service) implementations provide a small amount of business-process usefulness, such as basic data access. Larger granularities are compositions of smaller grained components and possibly other artefacts, where the composition taken as a whole conforms to the enterprise component definition. The coarseness of the service operations to be exposed depends on business usage scenarios and requirements and should be at a relatively coarse-level reflecting the requirements of business processes.

A coarse-grained interface might be the complete processing for a given service, such as “SubmitPurchaseOrder”, where the message contains all of the business information needed to define a purchase order. A fine-grained interface might have separate operations for: “CreateNewPurchaseOrder”, “SetShippingAddress”, “AddItem”, and so forth. This example illustrates that fine-grained services might be services that provide basic data access or rudimentary operations. These services are of little value to business applications. Services of the most

value are coarse-grained services that are appropriately structured to meet specific business needs. These coarse-grained services can be created from one or more existing systems by defining and exposing interfaces that meet business process requirements. Web service interfaces may be invoked via messages, which may likewise be defined as coarse- or fine-grained entities.

Fine-grained messages result in increased network traffic and make handling errors more difficult. This significantly hinders cross-enterprise integration. However, internal use of Web services may be beneficial in those cases where the internal network is faster and more stable. A higher number of fine-grained services and messages might therefore be acceptable for EAI applications.

From the perspective of service-oriented design and development it is preferable to create higher-level, coarse-grained interfaces that implement a complete business process. This technique provides the client with access to a specific business service, rather than getting and setting specific data values and sending a large number of messages. Enterprises can use a single (discrete) service to accomplish a specific business task, such as billing or inventory control or they may compose several services together to create a distributed e-Business application such as customised ordering, customer support, procurement, and logistical support. These services are collaborative in nature and some of them may require transactional functionality. Enabling business to take place via limited message exchanges is the best way to design a Web services interface for complex distributed applications that make use of SOAs.

5 PHASES OF THE SERVICE-ORIENTED DESIGN AND DEVELOPMENT METHODOLOGY

The objective of the service-oriented design and development methodology is to achieve service integration as well as service interoperability.

5.1 The Planning Phase

The planning phase determines the feasibility, nature and scope of service-solutions in the context of an enterprise. A strategic task for any organization is to achieve a service technology “fit” with its current environment. The key requirement in this phase is thus to understand the business environment and to make sure that all necessary controls are incorporated into the design of a service-oriented solution. Activities in this phase include analyzing the business needs in measurable goals, reviewing of the current technology landscape, conceptualizing the requirements of the new environment and mapping those to new or available implementations. Planning also includes a financial analysis of the costs and benefits including a budget and a software development plan including tasks, deliverables, and schedule.

A business environment is usually large and complex. Business experts at the service provider's side provide a categorization and decomposition of the business environment into business areas based on the functions being served by sets of business processes. The stakeholders in this business environment view the discrete units of work done within their enterprise organized as business processes, which are in reality higher-order services that are further decomposed into simpler services.

The planning phase is very similar to that of software development methodologies including the RUP [RUP 2001], [Royce 1998], and will be not discussed any further.

5.2 The Analysis Phase

Service-oriented analysis is a phase during which the requirements of a new application are investigated. This includes reviewing business goals and objectives that drive the development of business processes. Business analysts complete an "as-is" process model to allow the various stakeholders understand the portfolio of available services and business processes. The organization designs, simulates, and analyzes potential changes to the current application portfolio for potential return-on-investment (ROI) before it commits to any changes to business processes. This analysis results in the development of the "to-be" process model that an SOA solution is intended to implement.

The analysis phase examines the existing services portfolio at the service provider's side to understand which policies and processes are already in place and which need to be introduced and implemented.

The analysis phase encourages a radical view of process (re)-design and supports the re-engineering of business processes. Its main objective is the reuse (or repurposing) of business process functionality in new composite applications. To achieve this objective the analysis phase comprises four main activities: process identification, process scoping, business gap analysis, and process realization.

5.2.1 Process Identification

Understanding how a business process works and how component functionality differs or can get adjusted between applications is an important milestone when identifying suitable business process candidates. When designing an application, developers must first analyse application functionality and develop a logical model of what an enterprise does in terms of business processes and the services the business requires from them, e.g., what is the shipping and billing addresses, what is the required delivery time, what is the delivery schedule and so on. Thus the objective of this step is to identify the services that need to be aggregated into a business process whose interface has a high viscosity.

The key factor is being able to recognize functionality that is essentially self-sufficient for the purposes of a business process. It is important when designing a business process

to identify the functionality that should be included in it and the functionality that is best incorporated into another business processes. Here, we can apply the design principles of coupling and cohesion to achieve this. For instance, an order management process has low communication protocol coupling with a material requirements process.

Figure 3 Business process identification and scoping (adapted from [El Sawy 2001]).

Process identification could start with comparing the abstract business process portfolio with standard process definitions, such as for example, RosettaNet's "Manage Purchase Order" (PIP3A4). This PIP encompasses four complementary process segments supporting the entire chain of activities from purchase order creation to tracking-and-tracing, each of which is further decomposed into multiple individual processes. A quick-scan of this PIP reveals that Segment 3A "Quote and Order Entry" and Segment 3C "Returns and Finance" may be combined into the process Order Management.

5.2.2 Process Scoping

Defining the scope of business processes helps ensure that a process does not become monolithic mimicking a complete application. Unbundling functionality into separate (sub-) business processes will prevent business processes from becoming overly large, complex and difficult to maintain. For example, designing a business processes that handles online purchasing would require the removal of packaging and shipping information and costs to different business processes. In this example the three functions are mutually exclusive and should be implemented separately. The functionality included in these business processes is not only discrete and identifiable but it is also loosely coupled to other parts of the application.

The *scope of a business process* is defined as an aggregation of aspects that include where the process starts

and ends, the typical customers (users) of the process, the inputs and outputs that the customers of the process expect to see, the external entities, e.g., suppliers or logistics providers, that the process is expected to interface with, and the different types of events that start an instance of the process.

Figure 3 illustrates how business identification and scoping interact and result either in processes that can be reused and are candidates for design or in processes that need to be redesigned and re-engineered.

5.2.3 Business Gap Analysis

Gap analysis is a technique that purposes a business process and Web services realization strategy by incrementally adding more implementation details to an abstract service/process interface. Gap analysis commences with comparing candidate service functionality with available software service implementations that may be assembled within the enclosures of a newly conceived business process. A gap analysis strategy may be developed in stages and results in a recommendation to do development work, reuse or purchase Web services. Several service realization possibilities are discussed later in this section. For the moment let us assume that there may exist software components internal to an organization that provide good match. These may include service implementations previously developed by the enterprise, externally supplied service realizations available on a subscription or pay per use basis. Service realizations may be a blend of service and service-enabled implementations. In this way, portfolios of services possibly accessible on a global scale will complement and sometimes even entirely replace monolithic applications as the new fabric of business processes

5.2.4 Process Realization Analysis

Process realization analysis is an approach that considers diverse business process realization scenarios evaluated in terms of costs, risks, benefits and return of investment in accordance with business requirements and priorities. Service providers consider the following four realization options (which may be mixed in various combinations) to develop new business processes [Brittenham 2001]:

1. Green-field development: This step involves describing how a new interface for a Web service will be created on the basis of the Web service implementation. Green field development assumes that first a service is implemented and subsequently the service interface is derived from the new Web service implementation. During this step the programming languages and models that are appropriate for implementing the new Web service are also determined.

2. Top-down development: Using this realization option a new service can be developed that conforms to an existing service interface. This type of service interface is usually part of an industry standard that can be developed by any number of service providers. Processes are usually deployed top-down from a business level process blueprint. The main benefit of the top-down service development is the consistency of the applications and integration mechanisms. It is also rather easy to evolve the service-oriented solution across the enterprise as the industry evolves. Main problems with this realization option are the costs involved in development as well as the costs of achieving consensus on a high-level SOA architecture throughout the enterprise [Graham 2005].
3. Bottom-up development: Using this option a new service interface is developed for an existing application. The existing application can be coded as a Java, C++ program, Enterprise Java Bean (EJB), etc, or could be a back end legacy application. This option usually involves creating a Web service interface from the application programming interface (API) of the application that implements the Web service. Bottom is well suited for an environment that includes several heterogeneous technologies and platforms or uses rapidly evolving technologies.
4. Meet-in-the-middle- development: This option is used when an already existing Web service interface - for which an implementation already exists - is partially mapped onto a new service or process definition. This option involves service realizations that mix service and service-enabled implementations. This approach may thus involve creating a wrapper for the existing applications that need to be service-enabled and that are to be combined with the already existing Web service interface. Meet-in-the-middle development realization strategies offer a middle ground that attempts to take advantage of some of the benefits of the other approaches while extenuating some of the most notable problems and risks.

One of the issues with the top-down, bottom-up and meet-in-the-middle development options is that they are rather ambiguous regarding which business processes an enterprise should start from and how these can be combined to form business scenarios. To address this problem, service development solutions need to target specific focal points and common practises within the enterprise such as those that are specified by its corresponding sector reference models. Reference models - a typical example of which is RosetaNet - address a common large proportion of the basic

‘plumbing’ for a specific sector, from a process, operational function, integration, and data point of view. Such a *verticalized development model* presents the ideal architecture for supporting service development. This makes certain that the development team are aware of known best practices and standard processes so that they do not reinvent the wheel. For example, developers could use RosettaNet’s standard processes PIP4B2 (“Notify of Shipment Receipt”) and PIP4C1 (“Distribute Inventory Report”) for applications in which suppliers develop replenishment plans for consignment inventory at buyers. Product receipt and inventory information is notified using PIP4B2 and PIP4C1, respectively between consignment warehouses and suppliers.

The various options for process realization analysis emphasize the separation of specification from implementation that allows Web services to be realized in different ways, e.g., top-down or meet-in-the-middle development. It then becomes important to plan effectively when deciding how to realize or provision services; we need to carefully examine the diversity of realization alternatives and make the right choice. The service realization strategy involves choosing from an increasing diversity of different options for services, in addition to service reuse, which may be mixed in various combinations. This includes reusing or repurposing already existing Web services, business processes or business process logic; developing new Web services or business processes logic from scratch; purchasing/leasing/paying per use for services; outsourcing service design and implementation regarding Web services or (parts of) business processes; and, using wrappers and/or adapters to revamp existing enterprise (COTS) components or existing (ERP/legacy) systems.

Process realization results in a business architecture represented by business processes and the set of normalized business functions extracted from the analysis of these processes. During process realization, analysis decisions are made whether to reuse a particular enterprise asset, e.g., service or business process. To determine the quality of a specific asset, quality metrics are used that evaluate its flexibility, extensibility, maintainability, and level of cohesion and coupling. The process realization analysis estimates existing and expected operational costs, integration costs, service and process customisation costs, service and process provisioning costs and architecture costs for each scenario realization scenario. Develops and project managers could use tools such as, for example, IBM Rational Portfolio Manager to gain insight into the business benefits, costs, and risks of the SOA services portfolio [Brown 2005]. Architecture costs are associated with acquiring artefacts for realizing the target architecture including servers, specialized software, training, and required network bandwidth.

Service analysis is logically succeeded by service design, during which conceptual processes and services are transformed into a set of related, platform-agnostic interfaces. Designing a service-oriented application requires developers to model and define well-documented interfaces for all major service components prior to constructing the services themselves. Service design is based on a twin-track development approach that provides two production lines: one to produce services (possibly out of pre-existing components), and another to assemble (compose) services out of reusable service constellations. This calls for a business process model that forces developers to determine how services combine and interact jointly to produce higher level services.

Service design, just like service analysis, has its own special characteristics and techniques, which we shall describe in this section. We shall first start by describing a broad set of service design concerns.

6.1 Service Design Concerns

A number of important concerns exist that influence design decisions and result in an efficient design of service interfaces, if taken seriously. These concerns bring into operation the design principles for service-enabled processes. Prime concerns include managing service granularity, designing for service reuse and designing for service composability.

6.1.1 Managing Service and Component Granularity

Identifying the appropriate level of granularity for a service or its underlying component is a difficult undertaking as granularity is very much application context dependent. In general, there are several heuristics that can be used to identify the right level of granularity for services and implementation components. These include clearly identifiable business concepts, highly usable and reusable concepts, concepts that have a high-degree of cohesion and low-degree of coupling and must be functionally cohesive. Many vertical sectors, e.g., automotive, travel industry and so on, have already started standardising business entities and processes by choosing their own levels of granularity.

6.1.2 Designing for Service Reusability

When designing services it is important to be able to design them for reuse so that they can perform a given function wherever this function is required within an enterprise. To design for service reuse one must make services more generic, abstracting away from differences in requirements between one situation and another, and attempting to use the generic service in multiple contexts where it is applicable. Designing a solution that is reusable requires keeping it as simple as possible. There are intuitive techniques that facilitate reuse that are related to design issues such as identification and granularity of services. These include

looking for common behaviour that exists in more than one place in the system and trying to generalize behaviour so that it is reusable.

When designing a service-based application, it is possible to extract common behaviour and provide it by means of a generic services so that multiple clients can use it directly. It is important, however, that when designing enterprise services that business logic is kept common and consistent across the enterprise so that generalization is not required. Nevertheless there are cases where fine-tuning, specialization or variation of business logic functionality is required. Consider for instance, discounting practices that differ depending on the type of customer being handled. In those cases it is customary to produce a generalized solution with customisation points to allow for service variations.

6.1.3 Designing for Service Composability

In order to design useful and reliable services we need to apply sound service design principles that guarantee that services are self-contained, modular and support service composability. The design principles that underlie component reusability revolve around the two well-known software design guidelines service coupling and cohesion.

6.2 Specifying Services

A service specification is a set of three specification elements, all equally important. These are [Johnston 2005]:

- *Structural specification*: This focuses defining the service types, messages, port types and operations.
- *Behavioural specification*: This entails understanding the effects and side effects of service operations and the semantics of input and output messages. If, for example we consider an order management service we might expect to see a service that lists "place order," "cancel order," and "update order," as available operations. The behavioural specification for this ordering service might then describe how one cannot update or cancel an order you did not place, or that after an order has been cancelled it cannot be updated.
- *Policy specification*: This denotes policy assertions and constraints on the service. Policy assertions may cover security, manageability, etc.

During service design, service interfaces that were identified during the analysis phase are specified based on service coupling and cohesion criteria as well as on the basis of the service design concerns that we examined in section-6.1. In case that reference models are available business processes and service interfaces can be derived on their basis. For the remainder of this section we consider a sample purchase order business process for supply chain service-oriented applications that is derived on the basis of RosettaNet's order management PIP cluster. By applying the dominant cohesion criterion, namely functional cohesion, this process

may be decomposed into several sub-processes such as "Quote and Order Entry", "Transportation and Distribution", and "Returns and Finance", which conform to RosettaNet's segments 3A, 3B and 3C respectively. The "Quote and Order Entry" sub-process allows partners to exchange price and availability information, quotes, purchase orders and order status, and enables partners to send requested orders to other partners. The "Transportation and Distribution" sub-process enables communication of shipping- and delivery-related information with the ability to make changes and handle exceptions and claims. Finally, the "Returns and Finance" sub-process provides for issuance of billing, payment and reconciliation of debits, credits and invoices between partners as well as supports product return and its financial impact. By applying functional cohesion again, the Quote and Order Entry sub-process may be decomposed into several services such "Request Quote", "Request Price and Availability", "Request Purchase Order" and "Query Order Status". These three services conform to RosettaNet's PIPs 3A1, 3A2, 3A4 and 3A5, respectively.

6.2.1 Structural and Behavioral Service Specification

In the following we will briefly examine the course of specifying an interface for a Web service in WSDL [Chinnici 2004]. Web service interface specification comprises four steps: describing the service interface, specifying operation parameters, designating the messaging and transport protocol, and finally fusing port types, bindings and actual location (a URI) of the Web-services. These steps themselves are rather trivial and already described in-depth in literature, e.g., in [Alonso 2004] a detailed approach for specifying services is outlined. The following guidelines and principles are relevant while developing a WSDL specification:

- The service interface should only contain port types (operations) that are logically related or functionally cohesive. For example, the service "Request Purchase Order" captures the operations "Purchase Order Request" and "ReceiptAcknowledgement" as they are functionally cohesive;
- Messages within a particular port type should be tightly coupled by representational coupling and communication protocol coupling. For example, the operation "Purchase Order Request" may have one input message (PurchaseOrderID) and one output message (PurchaseOrder) sharing the same communication protocol (e.g., SOAP) and representation (atomic XML Schema datatypes);
- Coupling between services should be minimized. For example, the services "Request Purchase Order" and "Query Order Status" are autonomous, having no interdependencies.

```

... ..
<portType name="CanReceive3A42_PortType">
  <!-- name of operation is same as name of message -->
  <operation name="PurchaseOrderRequest">
    <output message="tns:PurchaseOrderRequest"/>
  </operation>
  <operation name="ReceiptAcknowledgment">
    <output message="tns:ReceiptAcknowledgment"/>
  </operation>
</portType>

<portType name="CanSend3A42_PortType">
  <!-- name of operation is same as name of message -->
  <operation name="PurchaseOrderConfirmation">
    <input message="tns:PurchaseOrderConfirmation"/>
  </operation>
  <operation name="ReceiptAcknowledgment">
    <input message="tns:ReceiptAcknowledgment"/>
  </operation>
  <operation name="Exception">
    <input message="tns:Exception"/>
  </operation>
</portType>
... ..

```

Figure 4 WSDL excerpt for request purchase order.

Specifying the service interface: A WSDL specification outlines operations, messages, types and protocol information. Figure 4 shows an abridged specification for a service interfaced and operation parameters for the “Request Purchase Order” service. The WSDL example in Figure 4 illustrates that the Web service defines two `<portType>` named “CanReceive3A42_PortType” and “CanSend3A42_PortType”. The `<portType>` “CanReceive3A42_PortType” supports two `<operation>`s, which are called “PurchaseOrderRequest” and “ReceiptAcknowledgment”.

Specifying operation parameters: After having defined the operations, designers need to specify the parameters they contain. A typical operation defines a sequence containing an input message followed by an output message. When defining operation parameters (messages) it is important to decide whether simple or complex types will be used. As Web services become increasingly more complex, and message-based SOAP will be more appropriate, complex schemas need to be created. The `<operation>` “PurchaseOrderRequest” in Figure 4 is shown to contain an output message “PurchaseOrderRequest” which includes a single part named “PO-body”. This part is shown in Figure 5 to be associated with the complex type “PIP3A4PurchaseOrderRequest” that is further specified to the level of atomic (XSD) types in the compartment that is embraced with the `<wsdl:types>` tag. As Figure 5 indicates, well-factored Web services often result in a straightforward `<portType>` element where the business complexity is moved into the business data declaration.

Several graphical Web services development environments and toolkits exist today. These enable developers to rapidly create, view, and edit services using WSDL and manage issues such as correct syntax and validation, inspecting and testing Web services and accelerating many common XML development tasks encountered when developing Web service enabled applications.

6.2.2 Service Programming Style

In addition to structural and behavioural service specification the service programming style must also be specified during the service design phase. Determining the service programming style is largely a design issue as different applications impose different programming-style requirements for Web services. Consider for example an application that deals with purchase order requests, purchase order confirmations and delivery information. This application requires that request messages may contain purchase orders in the form of XML documents while response messages may contain purchase order receipts or delivery information again in the form of XML documents. This type of application uses data-oriented Web services. Moreover, there is no real urgency for a response message to follow a request immediately if at all. In contrast to this consider an application that provides businesses with up-to-the-instant credit standings. Before completing a business transaction, a business may require to check a potential customer’s credit standing. In this scenario, a request would be sent to the credit check Web service provider, e.g., a bank, processed, and a response indicating the potential customer’s credit rating would be returned in real time. This type of Web service relies on an RPC- or process-oriented programming style. In these types of applications the client invoking the Web service needs an immediate response or may even require that the Web services interact in a back-and-forth conversational way.

```

<wsdl:types>
  <xsd:complexType name="PIP3A4PurchaseOrderRequest">
    <xsd:sequence>
      <xsd:element ref="PurchaseOrder"/>
      <xsd:element ref="fromRole"/>
      <xsd:element ref="toRole"/>
      <xsd:element ref="thisDocumentGenerationDateTime"/>
      <xsd:element ref="thisDocumentIdentifier"/>
      <xsd:element ref="GlobalDocumentFunctionCode"/>
    </xsd:sequence>
  </xsd:complexType>

  <xsd:complexType name="PurchaseOrder">
    <xsd:sequence>
      <xsd:element ref="deliverTo" minOccurs="0"/>
      <xsd:element ref="comment" minOccurs="0"/>
      <xsd:element ref="packListRequirements" minOccurs="0"/>
      <xsd:element ref="ProductLineItem" maxOccurs="unbounded"/>
      <xsd:element ref="GlobalShipmentTermsCode"/>
      <xsd:element ref="RevisionNumber"/>
      <xsd:element ref="prePaymentCheckNumber" minOccurs="0"/>
      <xsd:element ref="QuoteIdentifier" minOccurs="0"/>
      <xsd:element ref="WireTransferIdentifier" minOccurs="0"/>
      <xsd:element ref="AccountDescription" minOccurs="0"/>
      <xsd:element ref="generalServicesAdministrationNumber"
        minOccurs="0"/>
      <xsd:element ref="secondaryBuyerPurchaseOrderIdentifier"
        minOccurs="0"/>
      <xsd:element ref="GlobalFinanceTermsCode"/>
      <xsd:element ref="PartnerDescription" maxOccurs="unbounded"/>
      <xsd:element ref="secondaryBuyer" minOccurs="0"/>
      <xsd:element ref="GlobalPurchaseOrderTypeCode"/>
    </xsd:sequence>
  </xsd:complexType>
</wsdl:types>

... ..
<message name="PurchaseOrderRequest">
  <part name="PO-body" type="tns:PIP3A4PurchaseOrderRequest"/>
</message>

```

Figure 5 Specifying parameters for operations in Figure 4.

The use of document-based messaging promotes loose coupling. In general, there is no reason for the Web service client to know the name of the remote methods. In document-based messaging, the Web services subsystem receives the document, invokes the appropriate methods,

and responds. Using document-based messaging results in loose coupling, since the message constructs the document, but does not indicate the steps to process that document. With document-based messaging, it is the only receiver that knows the steps to handle an incoming document. Thus, the receiver can add additional steps, delete steps, and so on, without impacting the client.

6.2.3 Service Policy Concerns

The design of service-oriented solutions, like any other complex structure, requires early architectural decisions supported by well-understood design techniques, structural patterns, and styles that go far beyond ensuring "simple" functional correctness, and deal with non-functional service concerns. These patterns address common QoS issues that include performance requirements, information regarding service reliability, scalability, and availability, transactional requirements, change management and notification, and so on. In general, non-functional service characteristics describe the broader context of a service, e.g., what business function the service accomplishes, how it fits into a broader business process as well as characteristics of the hosting environment such as whether the component provider ensures security and privacy, what kind of auditing, security and privacy policy is enforced by the component provider, what levels of quality of component are available and so on.

Typical service related non-functional concerns that are addressed by policies include security issues and authorization concerns, and policy models. For example, a service provider could specify a policy stating that a given Web service requires Kerberos tokens, digital signatures, and encryption that can be used by clients use such policy information to determine whether they can use the particular service under consideration. In another example, an authentication model may require that a client authenticates itself by presenting its encoded credentials or may require that XML signatures be generated for Web service requests.

As a wide range of services is provided across a network it is natural that services would benefit from the use of policy management models, which could determine the configuration of a network of differentiated services according to business rules or application-level policies. There are many reasons why enterprises might want to give different levels of service to different customers or why they might need different levels of priority to different business transaction models involving Web services. Therefore, it is only natural that such criteria constitute important elements of a service design methodology and are considered to be equally important to technical policies such as security or authentication.

6.3 Specifying Business Processes

Designers should be able to compose (or decompose) and relate to each other process models that are developed in different parts of the enterprise, or by partners. They also

should be in a position to incrementally refine processes and export process improvements achieved in one part of the business to other parts of the business, with adaptation as required. Industry best practices and patterns must also be taken into account when designing processes and abstract process models can act as blueprints for subsequent concrete models.

Once business processes are extracted and their boundaries are clearly demarcated, they need to be described in the abstract. This step comprises four separate tasks, one deriving the process structure, one linking it to business roles, which reflect responsibilities of the trading partners, e.g., a buyer, a seller and a shipper in the order management process, and one specifying non-functional characteristics of business processes. The first task is to choose the type of service composition. The choice is between orchestration versus choreography. If a choice for orchestration is made three tasks follow to orchestrate a process. These are defined using the Web services Business Process Execution Language (BPEL) [Andrews 2003]. In the following we will place emphasis on orchestration given that today there are several implementations for BPEL while the WS-CDL (Choreography Description Language) is still being specified.

6.3.1 Describing the Business Process Structure

The first step in the business process design is to specify the business structure and the functions of the business process. The business process structure refers to the logical flow or progression of the business process. A business process reveals how an individual process activity (<PortType>) is linked with another in order to achieve a business objective. To assemble a higher-level service (or process) by combining other Web services, the service aggregator needs to select potential services that need to be composed depending on how these services and their operations fit within the enclosures a business process and how they relate to one another. Subsequently, the service provider needs to connect the process interface to the interfaces of imported services and plug them together. Business processes can be scripted using BPEL.

The abstract description of a process encompasses the following tasks:

1. *Identify, group and describe the activities that together implement a business process:* The objective of this action is to identify the services that need to be combined in order to generate a business process and then describe the usage interface of the overall business process. The functions of a business process are expressed in terms of the activities or the services that need to be performed by a specific business process. For instance, the registration of a new customer is an activity in a sales order process. The structure of a business process describes how an individual process activity (<PortType>) is linked with one

another. To assemble a higher-level service by combining other Web services, the service designer needs to:

- a. Select the services to compose by looking at how these services and their operations within a business process relate to one another.
 - b. Connect the usage interface of the business process to the interfaces of imported services and plug them together.
2. *Describe activity dependencies, conditions or synchronisation:* A process definition can organise activities into varying structures such as hierarchical, conditional and activity dependency definitions. In a hierarchical definition processes activities have a hierarchical structure. For instance, the activity of sending an insurance policy for a shipped order can be divided into three sub-activities: compute the insurance premium, notify insurance, mail insurance premium to the customer. In process definitions that have a conditional activity structure activities are performed only if certain conditions are met. For instance, it may be company policy to send a second billing notice to a trading partner when an invoice is more than two months overdue. Activity dependency definitions signify dependencies between activities and sub-activities in a process. In any process definition, sub-activities can execute only after their parent activity has commenced. This means that sub-activities are implicitly dependent on their parent activity. In other cases there might be an explicit dependency between activities: an activity may only be able to start when another specific activity has completed. For instance, a shipment cannot be sent to a customer if the customer has not been sent an invoice.
3. *Describe the implementation of the business process:* provide a BPEL definition, that maps the operations and interfaces of imported services to those of another in order to create the usage interface of the business process (higher-level Web service).

Figure 6 illustrates an abbreviated snippet of the BPEL specification for the order management process. The first step in the process flow is the initial buyer request. The listing shows that three activities are planned in parallel. An inventory service is contacted in order to check the inventory, and a credit service is contacted in order to receive a credit check for the customer. Finally, a billing service is contacted to bill the customer. Upon receiving the responses back from the credit, inventory and billing services, the supplier would construct a message back to the buyer.

6.3.2 Describing Business Roles

The second step during business process design is to identify responsibilities associated with business process activities and the roles that are responsible for performing them. Roles may thus invoke, receive and reply to business activities. Each service provider is expected to properly fulfill the business responsibility of implementing a business activity as one or more port types of a Web service, which perform a specific role. The result of this phase actually constitutes the foundation for implementing business policies, notably role-based access control and security policies.

Figure 7 illustrates the different parties that interact within the business process in the course of processing a client's purchase order. This listing describes how the "PurchaseOrder" service is interfaced with the port types of associated service specifications including a credit check and price calculation (billing) service (not shown in this figure) to create an order management process. Each `<partnerLink>` definition in this listing is characterised by a `<partnerLinkType>`.

```
<sequence>
  <receive partner="Manufacturer" portType="lms:PurchaseOrderPortType"
    operation="Purchase" variable="PO"
    createInstance="yes" >

  </receive>
  <flow>
    <links>
      <link name="inventory-check"/>
      <link name="credit-check"/>
    </links>
    <!-- Check inventory -->
    <invoke partner="InventoryChecker"
      portType="lms:InventoryPortType"
      operation="checkInventory"
      ...
      <source linkName="inventory-check"/>
    </invoke>
    <!-- Check credit -->
    <invoke partner="CreditChecker"
      portType="lms:CreditCheckPortType"
      operation="checkCredit"
      ...
      <source linkName="credit-check"/>
    </invoke>
    <!-- Issue bill once inventory and credit checks are succesful -->
    <invoke partner="BillingService"
      portType="lms:BillingPortType" operation="billClient"
      inputVariable="billRequest" outputVariable="Invoice" >
      joinCondition="getLinkStatus("inventory-check") AND
        getLinkStatus("credit-check") />
      <target linkName="inventory-check"/>
      <target linkName="credit-check"/>
    </invoke>
  </flow>
  ...
  <reply partnerLink="Purchasing" portType="lms:purchaseOrderPT"
    operation="Purchase" variable="Invoice"/>
</sequence>
```

Figure 6 BPEL process flow for Purchase Order Process.

Developers can use automated tools to assist them with designing and developing business process. Toolsets such as IBM's WebSphere Business Modeller enables analysts to model, simulate, and analyze complex business processes quickly and effectively. This toolset can be used to model "as-is" and "to-be" business processes, allocate resources, and perform "what-if" simulations to optimize and estimate business benefits [Brown 2005]. These models can then be transformed into UML and Business Process Execution Language models to jumpstart integration activities. These descriptions can be used to orchestrate the

constructed services as part of a business process workflow. Based on the overall business process model defined in WebSphere Business Modeller and exported as a BPEL, developers bring together the overall workflow by wiring together service implementations. To achieve this they can use WebSphere Integration Developer to import, create, enact, and manage business processes described in BPEL.

Figure 7 Defining roles in BPEL.

Specialized graphical notations such as the Business Process Modelling Notation (BPMN) can also be used for modelling business processes. BPMN is an attempt at a standards based business process modelling language that can unambiguously define business logic and information requirements to the extent that the resultant models are executable [White 2004], [Owen 2004]. BPMN is intended for allowing users to express the complex semantics of business processes. The net result is that business logic and information requirements are maintained in a model that is consistent with and reflects the requirements of business. BPMN is based on BPML's process execution meta-model and can produce directly fully executable BPEL processes.

6.3.1 Non-Functional Business Process Concerns

The process design sub-phase must also deal with non-functional process design concerns including among other things performance, payment model, security model, and, transactional behaviour. In the following we will briefly mention typical business process related, non-functional concerns.

Service Level Agreements (SLAs) provide a proven vehicle for not only capturing non-functional requirements but also monitoring and enforcing them. SLAs are special legal agreements that encapsulate multiple concerns, and symmetrically fuse the perspective of service provider and client. Besides mutual commitments regarding to-be-delivered services, e.g., scalability and availability, the SLA should stipulate penalties, contingency plans for exceptional situations, and mechanisms for disaster recovery.

As an example consider security policies targeted by an SLA. Such an SLA can be used to bundle security policies to protect multi-party collaborations. Knowing that a new business process adopts a Web services security standard such as WS-Security [Atkinson 2002] is not enough information to enable successful composition. The client needs to know if the services in the business process actually require WS-Security, what kind of security tokens they are capable of processing, and which one they prefer. Moreover, the client must determine if the service should communicate using signed messages. If so, it must determine what token type must be used for the digital signatures. Finally, the client must decide on when to encrypt the messages, which algorithm to use, and how to exchange a shared key with the service. Trying to orchestrate services without understanding these technical details will inevitably lead to erroneous results. For example, the purchase order service in the order management process may indicate that it only accepts username tokens that are based signed messages using X.509 certificate that is cryptographically endorsed by a third party.

7. THE SERVICE CONSTRUCTION PHASE

The construction phase of the lifecycle methodology includes development of the Web services implementation, the definition of the service interface description and the definition of the service implementation description [Brittenham 2001]. On the provider side the implementation of a Web service can be provided by either creating a new Web service, or by transforming existing applications into Web services, or by composing new Web services from other (reusable) Web services and applications. Unlike the previous phases that focus only the provider, the service construction phase also considers service requesters. On the service client side although the service requester progresses through similar lifecycle elements as the service provider, different tasks are performed during each construction step.

This phase may involve green-field code development, however, in most cases it will consist of modifying existing (J2EE or .NET) services or constructing wrappers on top of existing legacy applications. Thus, during the construction phase the process realization scenario identified in section 5.2.4 must be implemented.

8. THE SERVICE TEST PHASE

Service testing is generally characterized as a validation exercise ascertaining that requirements have been met and that the deliverables are at an acceptable level in accordance with existing standards during the analysis, design and implementation phases of the service-oriented design and development life cycle. The result of testing is a "healthy"

service-oriented application that performs well enough to satisfy the needs of its customers.

The most interesting type of testing for service implementations is *dynamic testing* that consists of running the implementation and comparing its actual to its expected behaviour before it is deployed. If the actual behaviour differs from the expected behaviour, a defect has been found. In the context of services, dynamic testing is used to perform a variety of types of tests such as functional tests, performance and stress tests, assembly tests and interface tests.

Functional testing covers how well the system executes the functions it is expected to execute—including user commands, data manipulation, searches and business processes, and integration activities. Functional testing covers the obvious surface type of functions, as well as the back-end system operations, such as security, database transactions and how upgrades affect the system [Brown 2002].

The focus of performance testing in service-oriented environments is monitoring the system on-line response times and transaction rates under peak workload conditions. It also involves load testing, which measures the system's ability to handle varied workloads. Performance testing is related to stress testing, which looks for errors produced by low resources or competition for resources. It is also related to volume testing, which subjects the software to larger and larger amounts of data to determine its point of failure.

The objective of interface testing is to ensure that any service developed to interface with other service functions properly outside of its surrounding process. Interface testing should be performed while testing the function that is affected, e.g., an order management process calling an inventory service.

Finally, assembly testing ensures that all services function properly when assembled into business processes. It also verifies that services that interact and interoperate function properly when assembled as apart of business processes.

In addition to functional tests, performance and stress tests, assembly tests and interface tests, there are a variety of additional tests that may need to be performed during the service test phase. These include network congestion tests, security tests, installability tests, compatibility tests, usability tests, and upgrade tests. Tests have to be conducted to ensure that service security requirements such as privacy, message integrity, authentication, authorization and non-repudiation.

9. THE SERVICE PROVISIONING PHASE

As Web services become acceptable from industry, organizations realise that there are several intricate issues pertaining to the deployment aspects of revenue generating Web services. Service provisioning is central to operating revenue generating Web services between organisations. The provisioning requirements for Web services impose serious implications for the development methodology of

services. Service provisioning is a complex mixture of technical and business aspects for supporting service client activities and involves choices for service governance, service certification, service enrolment, service auditing, metering, billing and managing operations that control the behaviour of a service during its use. We provide an overview of the most salient features of service provisioning in what follows.

9.1 Service Governance

The goal of service governance is to align the business strategy and imperatives of an enterprise with its IT initiatives [Mitra 2005], [Ports 2003]. When applied to service-oriented applications service governance may involve reviews of internal development projects as well external reviews from the perspective of the service providers, using results from gap analysis.

Typical issues for internal reviews include whether the right types of services have been selected, whether all requirements for new services have been identified and so forth. To this end service governance may use as input the findings of the business gap analysis sub-phase. Other internal reviews issues also include whether the use of a particular service within an application would conform to enterprise specific or government mandated privacy rules, whether service implementation does not compromise enterprise-specific intellectual property, and so on. To achieve its stated objectives and support an enterprise's business objectives on strategic, functional, and operational levels, service governance provides a well-defined structure. It defines the rules, processes, metrics, and organizational constructs needed for effective planning, decision-making, steering, and control of the SOA engagement to meet the business requirements of an enterprise and its customers [Balzer 2004].

Two different governance models are possible. These are central governance versus distributed governance.

With *central governance*, the governing body within an enterprise has representation from each service domain as well as from independent parties that do not have direct responsibility for any of the service domains. The central governance council reviews any additions or deletions to the list of services, along with changes to existing services, before authorizing the implementation of such changes. Central governance suits an entire enterprise.

With *distributed governance* each business unit has autonomous control over how it provides the services within its own enterprise. This requires a functional service domain approach. A central governance committee can provide guidelines and standards to different teams. Distributed governance suits distributed teams better.

9.2 Service Certification

To establish that a service possesses some desired property we need to use knowledge to predict properties that an assembled application may attain. Certification depends on

compositional reasoning [Bachmann 2000], which identifies which properties of services are material for predicting or achieving some end-system properties, such as performance, safety, scalability and so on, and how to predict the “values” of end-system properties from service properties. These contractual specifications must be expressive enough to capture all of the properties imposed by frameworks that will lead to measurable end system quality attributes.

9.3 Service Metering and Rating

This process requires that service providers come up with viable business cases that address factors such as service metering, rating and billing.

Service metering model: Use of a service by a client must be metered if the service provider requires usage-based billing. Then the service provider needs to audit the service as it is used and bill for it. This could typically be done on a periodic basis and requires that a metering and accounting model for the use of the service be established. The model could allow the establishment of a service contract for each new subscriber and tacking and billing for using the subscribed hosted services. To achieve this, the service-metering model could operate on the assumption that Web services with a high degree of value are contracted via, for example, SLAs.

Service rating/billing model: Software organisations that are used to the traditional up-front license/ongoing maintenance pricing structure for software should come up with annuity-based pricing models for the Web services they provide. The pricing (rating) model could determine subscriber rates based on subscription and usage events. For example, the pricing model could calculate charges for services based on the quality and precision of the service and on individual metering events based on a service-rating scheme. The billing model associates rating details with the correct client account. It provides adequate information to allow the retrieval and payment of billing details by the client and the correct disbursement of payments to the service provider’s suppliers (who are in turn service providers offering wholesale services to the original provider).

9.4 Service Billing Strategies

Increasingly, business models for commercial Web service provisioning will become a matter of concern to service providers. From the perspective of the service provider a complex trading Web service is a commercializable software commodity. For example, a service provider may decide to offer simple services (with no quality of service guarantee) for free, while it would charge a nominal fee for use of its complex (added value) Web services. With complex trading Web services the quality of service plays a high role of importance and the service is offered for a price. These types of services are very different from the selling of shrink-wrapped software components, in that payment should be on an execution basis for the delivery of

the service, rather than on a one-off payment for an implementation of the software. For complex trading Web services, the service provider may have different charging alternatives. These may include: payment on a per use basis, payment on a subscription basis, payment on a leasing basis, lifetime services, free services:, free services with hidden value.

10. THE SERVICE DEPLOYMENT PHASE

Deployment means rolling out new processes to all the participants, including other enterprises, applications and other processes. The service-oriented development methodology promotes a separation between development and deployment activities, which are grouped into separate phases that can occur at different times, and that different individuals with different skills can perform. This yields a true separation of concerns, enabling developers to repurpose software components, services and business processes.

The tasks associated with the deployment phase of the Web service development lifecycle include the publication of the service interface and service implementation definition. Services are deployed at the service provider side according to the four service realization option that we examined in section 5.2.4.

11. THE SERVICE EXECUTION PHASE

Execution means ensuring that the new process is carried out by all participants – people, other organisations, systems and other processes. During the execution phase, Web services are fully deployed and operational. During this stage of the lifecycle, a service requester can find the service definition and invoke all defined service operations. The run-time functions include static and dynamic binding, service interactions as a function of Simple Object Access Protocol (SOAP) serialization/deserialization and messaging and interactions with back-end legacy systems (if necessary).

12. THE SERVICE MONITORING PHASE

The service monitoring phase concerns itself with service level measurement and monitoring is the continuous and closed-loop procedure of measuring, monitoring, reporting and improving the quality of service of systems and applications delivered by service-oriented solutions. Service level monitoring is a disciplined methodology for establishing acceptable levels of service that address business objectives, processes and costs.

The service monitoring phase targets continuous evaluation of service level objectives and performance. To achieve this objective service monitoring requires that a set of QoS metrics is gathered on the basis of SLAs, given that

an SLA is an understanding of expectation of service. In addition, workloads need to be monitored and the service weights for request queues might need to be readjusted. This allows a service provider to ensure that the promised performance level is being delivered, and to take appropriate actions to rectify non-compliance with an SLA such as reprioritization and reallocation of resources.

To determine whether an objective has been met SLA QoS metrics are evaluated based on measurable data about a service - e.g., response time, throughput, availability, and so on - performance during specified times, and periodic evaluations. SLAs include other observable objectives which are useful for service monitoring. These include compliance with differentiated service-level offerings, i.e., providing differentiated QoS for various types of customers (gold, silver, bronze), individualized service-level offerings, and requests policing which ensures that the number requests per customer stay within a predefined limit. All these need to be also monitored and assessed. A key aspect of defining measurable objectives is to set warning thresholds and alerts for compliance failures. This results in pre-emptively addressing issues before compliance failures occur. For instance, if the response time of a particular service is degrading then the step could be automatically routed to a backup service.

OUTLOOK

In this paper we have described an experimental methodology for service-oriented design and development. The methodology that we presented reflects an attempt in defining a foundation of design and development principles that applies equally well to Web services and business processes. The methodology takes into account a set of development models (e.g., top-down, bottom-up and hybrid), stresses reliance on reference models, and considers several service realization scenarios (including green field development, outsourcing and legacy wrapping). During service and process design, not only the functional requirements of services and processes are considered but also their non-functional characteristics, e.g., security, transactional properties and policies, are taken into account.

In contrast to traditional software development approaches, the methodology that we introduced in this article emphasizes activities revolving around service provisioning, deployment, execution and monitoring. We believe that these activities will become increasingly important in the world of services as they contribute to the concept of *adaptive service capabilities* where services and processes can continually morph themselves to respond to environmental demands and changes without compromising on operational and financial efficiencies. In this way, business processes could be analysed in detail instantaneously, discovering and selecting suitable external services, detecting problems in the service interactions, searching for possible alternative solutions, monitoring

execution step by step, upgrading and versioning themselves, and so on.

Service adaptivity is particularly useful for integrated supply chains as it implies that an integrated supply chain solution can leverage collaborative, monitoring and control abilities to manage product variability and successfully exploit the benefits of available-to-promise (ATP) capabilities. For example, consider the case where an enterprise receives a direct request from its customer order-entry service. An order promising service routes this request instantaneously to all sites that could fulfil the order. Frequently there are multiple, hierarchically ordered partners with facilities in different geographic regions. The ATP service for available and planned inventory is then checked against the date requested by the customer and the appropriate quantities. If necessary, substitute choices are offered. The ATP results are then sent to a transportation-planning business process of a logistics service provider to determine transportation time and delivery dates. The results are subsequently relayed to the order-promising service, which selects the fulfilment site and responds to the customer-order service for approval. Order acceptance is then propagated back through the system, driving the acceptances. However, if the material is not available, the order promising service can use the capable-to-promise functionality to contact a production-scheduling service and establish a date for the products promised. All these steps involve conversation between processes that span enterprises, with customized alerts set up across the network to track exceptions and provide manual intervention if necessary.

We intend to further strengthen and refine the approach by conducting several real-life case studies in different sectors, from which experience will be gained and more design concerns may be derived. In addition, we plan to develop an integrated toolset to effectively support the methodology.

REFERENCES

- [Andrews 2003] T. Andrews et al., Business Process Execution Language for Web Services, Version 1.1, 2003
- [Alonso 2004] G. Alonso and F. Casati and H. Kuno and V. Machiraju, Web Services: Concepts, Architectures and Applications, Springer, Heidelberg, 2004
- [Arsanjani 2004] A. Arsanjani "Service-oriented Modeling and Architecture", IBM developerworks, November 2004, available at: <http://www-106.ibm.com/developerworks/library/ws-soa-design1/>.
- [Atkinson 2002] Bob Atkinson et al. Web Services Security (WS-Security), International Business Machines Corporation, Microsoft Corporation, VeriSign, Inc., Version 1.0, April 2002
- [Bachmann 2000] F. Bachmann et. al. "Technical Concepts of Component-Based Software Engineering", Technical Report, Carnegie-Mellon Univ., CMU/SEI-2000-TR-008 ESC-TR-2000-007, 2nd Edition, May 2000.
- [Balzer 2004] Y. Balzer "Strong governance principles ensure a successful outcome", IBM developerworks July 2004, available at: <http://www-106.ibm.com/developerworks/library/ws-improvesoa/>.

- [Brown 2002] C. Brown, G. Cobb, R. Culbertson “Testing ...”, Prentice Hall, April 2002.
- [Brown 2005] A. Brown et. al., “SOA Development Using the IBM Rational Software Development Platform: A Practical Guide”, Rational Software, September 2005.
- [Brittenham 2001] P. Brittenham “Web-services Development Concepts”, IBM Software Group, May 2001, available at: <http://www-06.ibm.com/software/solutions/webservices/>
- [Cauldwell 2001] P. Cauldwell, et. al. “XML Web Services”, Wrox Press Ltd., 2001.
- [Chappell 2004] D. Chappell, “Enterprise Services Bus”, O’Reilly, 2004
- [Chinnici 2004] R. Chinnici and M. Gudgin and J.-J. Moreau and J. Schlimmer and S. Weerarana, Web Services Description Language (WSDL) Version 2.0, March 2004, w3c.org 4.
- [Deora 2003] V. Deora et al. A Quality of Service Management Framework based on User Expectations Proceedings of the First International Conference on Service Oriented Computing ,pp.: 104-114, Springer-Verlag, 2003
- [Harmon 2003] P. Harmon “Second Generation Business Process Methodologies”, Business Process Trends, vol. 1, no. 5, May 2003.
- [Herzum 2000] P. Herzum, O. Sims “Business component Factory”, J. Wiley & Sons Inc., 2000.
- [Johnston 2005] S. Johnston “Modelling Service-oriented Solutions”, IBM developerworks, July 2005, available at: <http://www128.ibm.com/developerworks/rational/library/johnston/>.
- [Kruchten 2004] P. Kruchten “Rational Unified Process—An Introduction”, 3rd edition, Addison-Wesley, 2004.
- [Mitra 2005] T. Mitra “A Case for SOA Governance”, IBM developerworks August 2005, available at: <http://www-106.ibm.com/developerworks/webservices/library/ws-soa-govern/index.html>.
- [Owen 2004] M. Owen, J. Raj “BPMN and Business Process Management: An Introduction to the New Business Process Modelling Standard”, Business Process Trends, March, 2004, available at: www.bptrends.com.
- [Papazoglou 2003] M.P. Papazoglou and G. Georgakopoulos, Introduction to the Special Issue about Service-Oriented Computing, CACM, October 2003, 46(10): 24-29.
- [Papazoglou 2006] M.P. Papazoglou, Principles and Foundations of Web Services: A holistic view, Addison-Wesley, to appear: 2006.
- [Ports 2003] M. Potts et. al “Web Service Manageability – Specification (WS-Manageability)”, OASIS, September 2003.
- [Royce 1998] W. Royce “Software Project Management | A Unified Framework”, Addison-Wesley, 1998.
- [RUP 2001] Rational Software Corporation “Rational Unified Process: Best Practices for Software Development Teams”, Technical Paper TP026B, Rev. 11/01, November 2001, available at <http://www.therationaledge.com>.
- [White 2004] S. A. White “Introduction to BPMN”, Business Process Trends, July, 2004, available at: www.bptrends.com
- [Zimmerman 2004] O. Zimmerman, P. Korgdahl, C. Gee “Elements of Service-oriented Analysis and Design”, IBM developerworks, June 2004, available at: <http://www-106.ibm.com/developerworks/library/ws-soad1/>.