

Stateless Process Enactment

Raf Haesen

Lotte De Rore

Stijn Goedertier

Monique Snoeck

Wilfried Lemahieu

Stephan Poelmans

Katholieke Universiteit Leuven
Department of Decision Sciences & Information Management
Naamsestraat 69, B-3000 Leuven, Belgium
firstName.lastName@econ.kuleuven.be

1. NAME

Stateless Process Enactment

2. ALIAS

Flowless Workflow

3. AUDIENCE

Business process designers, Process-Aware Information Systems designers

4. MOTIVATION

John from the claims department at InsuranceCo is quite dissatisfied after his exploratory study of some popular workflow engines available in the market. In his opinion, they all act like some "Big Brother" who needs complete control over all existing information systems. John considers their claim handling process that is enacted by such a workflow engine, but for which a specific activity can be executed without intervention of the engine. However, in that case, the state of the running process will not be updated in the workflow engine and as a consequence it will be erroneous. STATELESS PROCESS ENACTMENT facilitates the construction of a workflow engine, in which the state of processes is not explicitly stored.

5. CONTEXT

An average company contains many information systems that are typically designed according to a layered approach: the lowest layer contains a set of database transactions that enable the consistent retrieval and modification of business data. On top of this data layer, applications are offered that support the execution of clearly delineated activities, such as opening a new claim case, registering damage, fraud detection or indemnity payment.

Preliminary versions of these papers were workshopped at Pattern Languages of Programming (PLoP) '07 September 5-8, 2007, Monticello, IL, USA. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission. Copyright is held by the authors. ISBN: 978-1-60558-411-9.

In the topmost layer, a business process represents the coordination of a set of activities to achieve a high-level business goal. For example, a claim handling process coordinates different activities to process insurance claims. The coordination of these activities can be automated by means of a workflow engine that interprets process descriptions and manages the accompanying process instances.

Figure 1 sketches a typical service-oriented enterprise architecture, in which each service type externalizes functionalities of a particular layer. A business domain offers business concept services and activity services, while a workflow engine offers workflow services:

- *Business concept services* provide access to business data and associated business logic in a transactional way.
- *Activity services* support the execution of activities to be executed by human actors or machines. An activity service can use one or more business concept services.
- *Workflow services* coordinate activity services according to a business process.

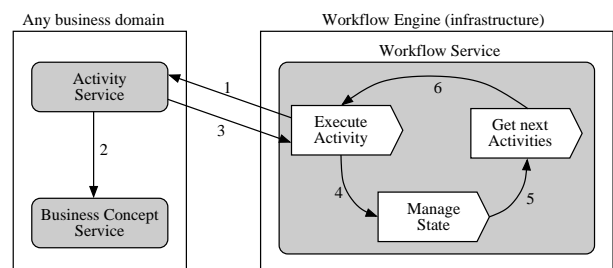


Figure 1: Service-Oriented Enterprise Architecture

Figure 1 also indicates the main steps of process enactment: the workflow engine typically guides the actor to a particular activity service (step 1). Eventually one or more business concept services will change a part of the business data (step 2). After execution of the activity, the workflow service continues the enactment of the process (step 3). After the workflow service has assessed the state of the process (step 4) the next activities can be determined (step 5) and executed (step 6).

6. PROBLEM

Figure 2 shows a typical representation of a business process, using the Business Process Modeling Notation (BPMN). Since all activities and their routing are explicitly predefined, state can be maintained in an efficient way, as it only involves updating one or more points of execution. The example represents a simplified version of a claim handling process with the sole focus on the control flow between the activities. After the case is opened, a claim handling expert needs to assess the damage of the involved object. Subsequently, the indemnity can be calculated in parallel with the investigation for a possible fraudulent case. If fraud is detected, the case is immediately closed; otherwise the indemnity is paid and finally the claim case is closed.

Most existing workflow engines explicitly store the state of running process instances as part of the workflow control data [5]. Such a workflow engine is called **stateful** since the enabling (pre-condition) and completion (post-condition) of activities is explicitly maintained in the engine for each activity of the process.

Stateful enactment boils down to the replication of the state of business concepts into the workflow engine, entailing all known problems associated with replication of data such as synchronization issues. The following two scenarios will show how inconsistencies can arise between the data in the business domains and the workflow control data (see Figure 3):

1. Experts should not be forced to follow predefined processes since the obligation to follow a predefined path of activities might be counterproductive. For example, a senior claim handler with many years of experience knows which activities he has to perform to handle a claim without the assistance of a workflow engine. As a consequence the workflow engine may not be notified about the completion of each executed activity.
2. To reduce costs, the workflow engine may be installed on a machine with low Service Level Agreements. Therefore the engine may occasionally fail. Suppose a particular activity – of a process under control of the workflow engine – has successfully been executed but flagging off that activity in the workflow engine fails because the engine is down at that moment. In other words, activity execution and process state modification form no single unit of work (UOW). If the engine is up again, it will return that same activity since for the engine it appears not to be executed yet. The repeated execution of one activity may obviously leave the system in an incorrect state.

Summarizing the problem, a workflow engine may fail or an actor may choose not to be assisted by a workflow engine. How can we ensure that the use of a workflow engine is optional?

7. FORCES

The forces can be described in terms of a set of quality characteristics that need to be considered. Furthermore it is appropriate to distinguish between forces that arise during process modeling (design time) on the one hand and process enactment (runtime) on the other hand.

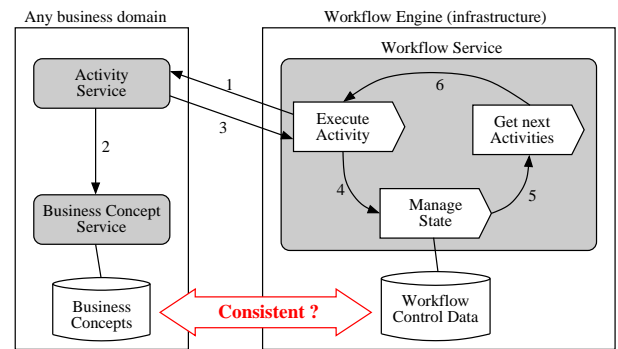


Figure 3: Architecture for stateful process enactment

- Process modeling forces

- **Executability:** The specification of a process is interpreted by the workflow engine. Therefore the specification must be complete and correct. Activities of a process model like the one in Figure 2 are identified only by their respective names. To improve executability the workflow engine should e.g. be able to automatically determine when a particular activity can be registered as completed.
- **Comprehensibility:** A process model must be easily understood by business people. For example, Figure 2 depicts a process model that can easily and rapidly be interpreted. However, in order to enable the automated execution of business processes (executability force), the corresponding process models should be as detailed as possible, which will most likely be at the expense of the understandability for the human actors.
- **Flexibility:** In order to minimally constrain a human actor, a process model should allow the execution of a ‘happy path’ and some variants of that path. However, each variant should not be explicitly elaborated, as it will lead to overly complex process models.
- **Expressivity:** A process modeling language should have a rich semantics that enables modeling all relevant requirements. The Workflow Patterns initiative¹ has identified different categories of patterns that should be supported by a workflow engine. A distinction is made between control-flow, data, resource, operational and exception handling patterns.

- Process enactment forces

- **Performance:** The introduction of workflow support should not create performance issues, especially not in the existing operational information systems. As explained in the context, activity services and business concept services are often realized by means of optimized mainframe transactions. During process enactment the workflow engine should influence these transactions to the minimal amount possible.

¹See <http://www.workflowpatterns.com>




Figure 2: Claim handling process

- **Progress monitoring and enforcement:** The workflow engine must ensure that enabled activities are effectively being executed. The problem description states that an actor may execute activities without intervention of the workflow engine. However, these activities may influence the state of process instances and therefore the activity instances that need to be executed.

8. SOLUTION

Synchronization problems between business data and process state can be avoided if the process state is not explicitly stored. Therefore, one should create a workflow engine that reconstructs the process state each time it is needed. Such a workflow engine is **stateless** since the conditions for enabling and completing activities are derived from the state of one or more business concepts. As illustrated in Figure 4, the workflow engine reconstructs the state of a process through the use of business concept services, each time the next possible activities are to be determined. As opposed to stateful process enactment, the state of a process is not explicitly stored as part of the workflow control data.

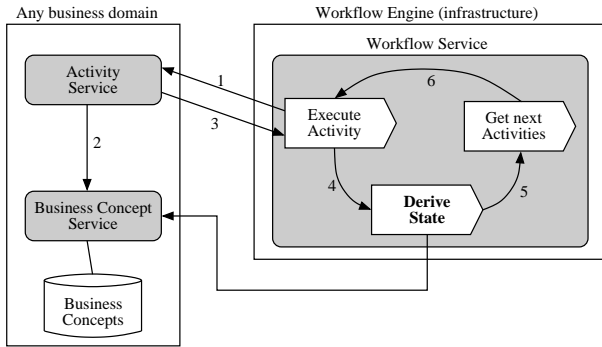


Figure 4: Architecture for stateless process enactment

Stateless process enactment can be seen as the extension of one of the key principles of the case handling paradigm, which states that an activity is completed if the associated mandatory data objects are entered [4]. As illustrated in Figure 5, both the pre-conditions and post-conditions of an activity depend on (the existence and/or state of) one or more (attributes of) business concepts. If the pre-conditions of a particular activity in the context of a process are met, that activity can be executed. On the other hand, an activity is completed if all its post-conditions are met. Note that the pre- and post-conditions can be added to the activities of an existing process model. However, as we will illustrate below, a set of activities enriched with pre- and post-conditions describes an (implicit) process model on its

own.




Figure 5: Activity pre- and post-conditions

Consider the enactment of the claim handling process in a stateless way. To do so the activities must have pre- and post-conditions that are expressed in terms of business concepts. An overview of the relevant business concepts and their attributes is depicted in Figure 6 using the Unified Modeling Language (UML). A client has one or more contracts and each contract insures one or more objects. For the sake of simplicity, we assume that each claim case is about exactly one damaged object. A claim case can be in five different states: opened, fraudulent, genuine, paid and closed. The possible indemnity of a case can be derived from the assessed damage of an insured object.




Figure 6: Relevant business concepts

Table 1 shows the set of activities and their accompanying pre- and post-conditions. These conditions can be added to the process model (Figure 2) or the set can be used as a process model on its own.

Activity	Pre-conditions	Post-conditions
Create Claim Case	/	ClaimCase cc exists and cc.state = opened
Identify Object Damage	cc.state = opened	cc.object exists and cc.object.damage filled out
Calculate Indemnity	cc.object.damage filled out	cc.indemnity filled out
Detect Possible Fraud	cc.state = opened	if fraud then cc.state = fraudulent else cc.state = genuine
Pay Indemnity	cc.indemnity filled out and cc.state = genuine	cc.state = paid
Close Claim Case	cc.state = paid or cc.state = fraudulent	cc.state = closed

Table 1: Claim handling activities, pre- and post-conditions

9. CONSEQUENCES

In most existing process modeling notations, the pre- and post-conditions on activities are often only implicitly specified as part of the control flow (arrows, messages and the conditions associated to gateway constructs). For stateless process enactment instead, one is obliged to model these conditions as they are used for deriving the process state at runtime. This clearly enhances the **executability** of the process since it is precisely defined when an activity is enabled or completed. When using typical process models, the workflow engine has no clue when an activity is completed until it receives an explicit notification. Moreover the engine can not verify the completion of the activity; it merely has to ‘believe’ the sender of the notification. This approach is error-prone especially for activities that require human input. On the other hand, explicitly modeled post-conditions enable the engine to verify them. Besides improving the workflow engine’s executability, allowing the specification of pre- and post-conditions also extends the modelers’ **expressiveness**.

The major drawback of this approach is the possibly intolerable drop in **performance** that may occur. Each time the workflow engine needs to derive state, all attributes that occur in the pre- or post-conditions must be queried in order to evaluate the conditions. Obviously, these operations can be extensively time-consuming if the process consists of many activities with many conditions to be checked.

The number of attributes to be retrieved can be decreased by adding “meta-state” attributes to business concepts. Those attributes explicitly indicate the state of the concept. For example, the state of an order (requested, shipped, paid, etc.) is often added as an attribute of order, or the claim case state in Figure 6 is a meta-state attribute of a claim case. Although such attributes simplify the reconstruction of process state, one should not be tempted to add attributes to business concepts only for the sake of a particular process. This obviously moves process information in the business concepts and consequently hinders separation of concerns. Note however that the addition of attributes to business concepts does not reintroduce the problems of the stateful approach since all state information stays part of the business concepts.

As indicated in the solution, it is possible to add pre- and post-conditions to an existing process model or it is possible to specify a set of activities only in terms of its pre- and post-conditions. This way a process can be seen as an unstructured set of activities in which an activity can be executed as soon as its pre-conditions are met. Such a specification offers the most **flexibility** since appropriate execution paths can be inferred at runtime. For example, according to the process model in Figure 2, it is possible to execute activity “Detect Possible Fraud” only after “Identify Object Damage” is executed, while according to Table 1 this is allowed as soon as the claim case is created. This way the claim handler has more freedom in deciding when to execute the activity. As another example, consider a claim handler who finds evidence to have the activity “Detect Possible Fraud” executed again at some arbitrary point in time after its first execution. Obviously, the process model in Figure 2 must be heavily extended to support this plausible scenario, whereas in the stateless approach, the claim handler can schedule this activity as long as its pre-conditions are satisfied.

In the stateful approach process progress is monitored and

enforced in the workflow engine by updating the state of the process after a completed activity returns control to the workflow engine. If particular activities become enabled, the engine has three options: (a) the engine waits for an actor to execute the activity; (b) the engine executes the activity if it requires no actor input; (c) the engine sends one or more tasks to actors who should execute the activity. The workflow engine can always send notifications when activity deadlines have to be respected.

In the stateless approach it is more complex to **monitor and enforce progress** if an activity that requires actor input is executed without assistance of the workflow engine. Indeed, in this case the workflow engine has no clue when that activity is completed. Despite this fact there are two possible approaches to make sure that progress is enforced. Firstly, the workflow engine can regularly calculate the state of a process to see if activities must be executed or tasks must be distributed. Given the previously discussed performance issues, it is indispensable to set an appropriate polling frequency. A second approach avoids this thorny task, namely by using a publish/subscribe mechanism: the workflow engine subscribes to particular data changes (published by business concept services) that influence activity pre- and post-conditions, so the workflow engine reacts only when necessary. A separate component in the workflow engine should (a) manage the subscriptions for each process and (b) possibly group multiple notifications before state is recalculated.

A control-flow based notation requires the explicit modeling of each process variant, which must be represented by a path in the model. For complex processes, this leads to gigantic models that are difficult to **comprehend**. On the contrary, processes with only a few variants can easily be represented and interpreted using graphical flow chart-like notations like BPMN. In addition, many tools offer animation to visualize and simulate process execution, which makes them indispensable for business people that need to assess as-is as well as to-be processes in a quick and efficient way. Finally, it is very hard to understand and verify the eventual behavior of a group of activities that are only specified by means of pre- and post-conditions. Even a simple set of rules, such as the ones represented in Table 1, are quite complex to understand.

10. KNOWN USES

As discussed in [4], only a few workflow management systems, such as FLOWer [2] and COSA [3], support the case handling paradigm. As we already mentioned earlier, the case handling paradigm embraces only part of this pattern: (a) only post-conditions of an activity can be expressed in terms of mandatory data objects that must exist and (b) a data object is the single fine-grained construct to express post-conditions. The open-source workflow engine, called con:cern (<http://con-cern.org>), fully applies this pattern.

This pattern is also part of the EM-BrA²CE framework, a unifying vocabulary and execution model for declarative process modeling [1]. Declarative process models represent business processes as trajectories in a state space, in which valid movements are described by declarative rules. In this pattern valid trajectories are constrained by pre- and post-conditions and the state space consists of business concepts. On the other hand, the EM-BrA²CE framework is more generic since it distinguishes ten additional rule types that

govern the execution of activities. Secondly, a state space does not only contain business concepts, but also other concepts such as activities, events and agents. The richer vocabulary facilitates the expression of more rule types to cover the contemporary business scenarios.

11. ACKNOWLEDGMENTS

The authors would like to express their gratitude to our shepherd António Rito Silva and the members of the ‘Centaur’ group of the Writers’ Workshop, whose suggestions greatly enhanced this pattern.

This pattern has been written as part of the KBC-Vlekho-K.U.Leuven research chair on ‘Service and Component Based Development’ sponsored by KBC Banking & Insurance.

12. REFERENCES

- [1] S. Goedertier, R. Haesen, and J. Vanthienen. EM-BrA²CE v0.1: A vocabulary and execution model for declarative business process modeling. FETEW Research Report KBL0728, K.U.Leuven, 2007.
- [2] Pallas Athena. Case handling with flower: Beyond workflow. Technical report, Pallas Athena, Apeldoorn, The Netherlands, 2002.
- [3] Software-Ley. Cosa activity manager. Technical report, Software-Ley GmbH, Pullheim, Germany, 2002.
- [4] W. M. P. van der Aalst, M. Weske, and D. Grünbauer. Case handling: a new paradigm for business process support. *Data and Knowledge Engineering*, 53(2):129–162, 2005.
- [5] WFMC. The workflow reference model. Technical Report WFMC-TC-1003, Workflow Management Coalition, 1995.