

# Media360 workflow - Implementing a Workflow Engine Inside a Database

Carsten Blecken

Informix  
4100 Bohannon Drive  
Menlo Park, CA 94025  
USA  
Carsten.Blecken@informix.com

## Abstract

I describe a Workflow Management System, where the workflow engine is implemented inside a database. The focus here is on the architectural approach and on the lessons learned from a commercial implementation in the workflow component of the Informix Media 360 content management solution.

Conceptually, the two main primitives of workflow are activities, which represent the work items of a larger process (containing either predefined or user-defined logic) and transitions, which describe the flow logic – in essence in what order and under what conditions the activities are being executed.

Existing workflow systems are typically implemented with an n-tier architecture. The flow logic necessary to execute the workflow, user-defined logic and pre- and post-activity logic is all usually implemented in a middle tier, although the workflow-relevant data is stored in a back-end database.

However, it is rather straightforward to use features of a modern Object Relational Database Management Systems like Informix Internet Foundation.2000 to implement a workflow engine completely inside a database. The flow logic can be implemented using the trigger concept, which ‘fires’ an appropriate transition to execute the next activity. User-defined logic, and pre- and post-activity logic are implemented using database functions or routines, which are a common extensibility feature of any ORDBMS. There is a wide array of programming languages available to express user-defined logic.

The principal advantage of such an approach is that the workflow logic is kept directly where the workflow data resides, resulting in a more efficient, simpler and more compact system design. It also aids with the embedding of database-centric workflow into a larger framework application, since a DBMS is part of all enterprise applications.

Finally, I discuss the advantages and disadvantages of this conceptual approach, and show how additional common workflow features can be added to the current architecture of the Informix Media360 workflow component.

---

*Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.*

**Proceedings of the 26th International Conference on Very Large Databases, Cairo, Egypt, 2000**