

# Distributed Monitoring and Workflow Management for Goal-oriented Workflows

Kai Jander<sup>1\*</sup>, Lars Braubach<sup>1</sup> and Winfried Lamersdorf<sup>1</sup>

<sup>1</sup>*Distributed Systems and Information Systems Group, University of Hamburg, Vogt-Kölln-Str. 30, 22527 Hamburg, Germany*

## SUMMARY

Business process management has often focused on business processes with a production or administrative focus, resulting in a task-centric approach for modeling and execution of workflow based on the processes. However, the area of collaborative business processes exists, which includes important processes like research and development which have not been adequately addressed by the techniques offered for production workflows. For this, goal-oriented processes have been proven useful to allow for a more flexible approach in modeling and executing workflow. In addition, distributed workflow management allows for a more flexible approach regarding organizational structures. However, the actual recording of action taken in a process are still relevant for workflow analysis and reengineering. As a result, this paper presents a distributed approach for generating, gathering, distributing and storing events resulting from the execution of goal-oriented workflows in a distributed workflow environment while allowing the real-time association of occurring actions with business goals in the process using drill-down analysis. Copyright © 2015 John Wiley & Sons, Ltd.

Received . . .

**KEY WORDS:** distributed workflow management; business process management; goal-oriented workflows; GPMN

## 1. INTRODUCTION AND MOTIVATION

In our previous paper (see [13]), we presented an approach for monitoring a distributed workflow management system with enacted goal-oriented workflows. We've shown how distributed monitoring components can collect events generated within the system, transmit them efficiently and reconstruct a chain of causes to attribute specific actions of a workflow with their business goals.

In this paper we expand on this work by given a deeper background in this section on the specific types of business processes where goal-oriented workflows are used. In addition, we will demonstrate use of a persistent storage system for persisting the events and also as an alternative means of event distribution (see Section 5). Finally, the storage system is evaluated in terms of expected functionality and performance in Section 6.

Business process management deals with identifying business processes within organizations and shaping the resources of those organizations around those processes rather than optimizing specialized department with a narrow focus on providing a specific function to an organization like accounting. A major goal in this endeavor is the eventual partial or complete automation of such business process with the use of IT systems. The resulting executable models reflecting the real

---

\*Correspondence to: Distributed Systems and Information Systems Group, University of Hamburg, Vogt-Kölln-Str. 30, 22527 Hamburg, Germany, E-mail: jander@informatik.uni-hamburg.de

world business processes are called workflows and are executable on computer systems running an execution environment that enacts the workflow models called a workflow management system (see [26]).

Business processes and workflows can be categorized using a variety of factors. One approach suggest by Leymann and Roller (see [19]) uses the degree of business value and the degree of repetition and structuring for defining categories of workflows:

- *Degree of Repetition* describes the degree in which the same workflow model can be reused for multiple enactments (executions) of workflow instances. For example, mass-produced goods generally follow the same model of assembly everytime they are produced while custom-designed software is generally tailored towards a specific customer and follows a different development process for each one. A related concept is the *degree of structuring*, which describes how rigidly a process can follow a pre-determined plan for execution. The aforementioned mass-production rigidly follows the assembly plan, while the custom software is developed more flexibly to adapt to specific and possibly changing customer requirements.
- Business value describe the relationship of the business process with regards to the purpose of the organization. For example, the purpose of a bakery is to produce and sell bread, so processes and workflows that deal with this purpose are considered to have high business value while necessary but unrelated processes such as administrative tasks and file-keeping, while necessary for the business, are considered to have a low business value.

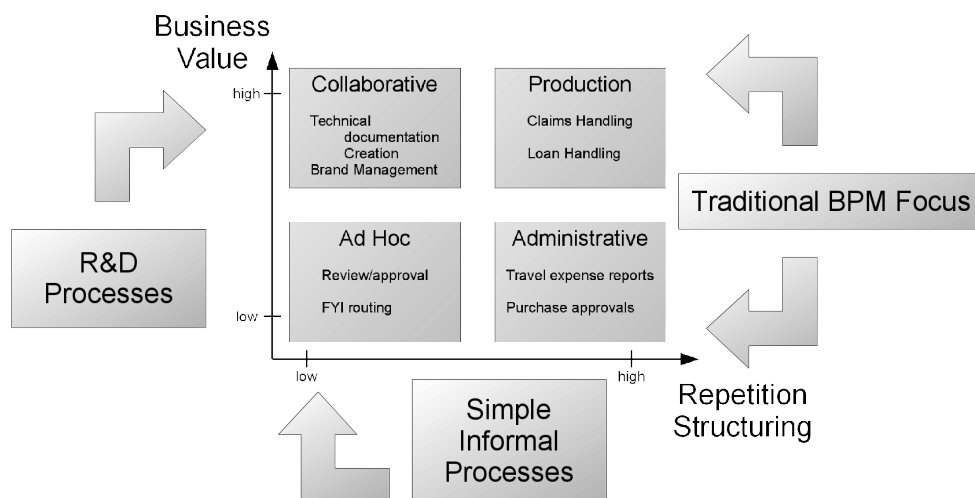


Figure 1. Different categories of workflows, the focus of the goal-oriented modeling approach centering around certain types of collaborative business processes and workflows, center chart adapted from [19]

Based on these scales, Leymann and Roller have proposed four different categories of workflows (see Fig. 1), which have been addressed by traditional business process management approach to a different degree:

- Ad-hoc business processes and workflows are often used to perform simple and low-priority tasks with low business value within organizations based on an unstructured ad-hoc basis. An example of this is for-your-information (FYI) routing. This process aims to distribute new information to people within the organization who require that information. This is done by sending the information to the people thought to be most interested in the information with the request to pass it onto additional people known to the recipient who might be interested. The approach is unstructured and does not require a lot of technical support to work as intended.
- The bureaucratic requirements of organizations are tackled with administrative business processes and workflows. This can involve task such as processing forms and approving purchase requests. They have a low business value since they are not directly involved in

the purpose of the organization, but generally are very repetitive and have a well-defined structure.

- Production business processes and workflows are the core of an organization and are involved in producing the product of the organization for its customers. In manufacturing, this can be the actual production of physical product but in service-oriented businesses with standardized services such as banks, the product can be a service provided to a customer. Since they are the most direct output of the organization, they have a high business value and since the product is standardized, the repetition and structure is high as well.
- Finally, collaborative business processes and workflows center around loosely-organized processes in which the participants are required to have a high degree of independence to perform their tasks and collaborate with the other participants in the process to further common goals. Typical example are research and development processes which are ultimately important to the organization's purpose and thus have a high business value but also necessitate a certain amount of creative freedom for the people involved in the process.

Leymann and Roller point out that the focus of business process management is typically centered on the production workflows. This is unsurprising, for three reasons: First, the high degree of structuring means that targeting this category is relatively straightforward while the high degree of repetition ensures that investments in the design of production workflows can amortize over the high number of repeated enactments. Finally, the high business value directly influences the competitiveness of the organizations. Furthermore, the structural similarities between the production and administrative processes means that techniques developed for one can also often be applied to the other.

As such, many workflow modeling languages and systems are geared towards such processes. However, this leaves the collaborative processes as a challenging area with a high potential due to their high business value. However, they differ structurally from both administrative and production processes and as a result, traditional workflow models such as the Business Process Model and Notation (BPMN) [20] often lack good support for flexibility in a multitude of business situations required by this category of processes, forcing the process designer to include numerous branches for every conceivable situation.

This deficit has led to a number of different approaches like ADEPT [18] attempting to attenuate those limitations. Part of the difficulties stem from the fact that most approaches to workflows are activity-based, which means they are focused on a specific order of actions with branching explicitly inserted at certain points. While this is a fairly intuitive approach for modeling workflows, it does not directly provide business reasons for activities, which can result in the inclusion of unnecessary or unwanted activities that merely exist for technical rather than business reasons in the workflow.

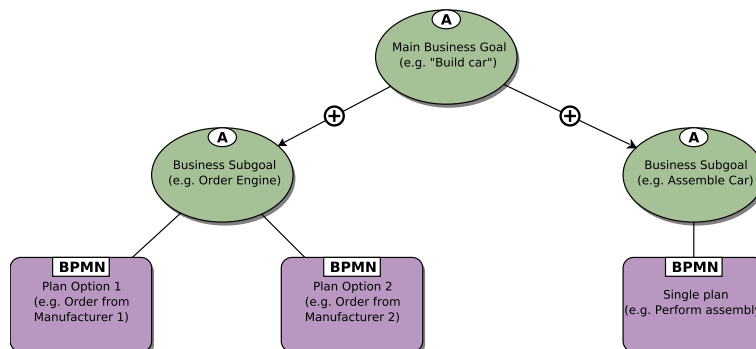


Figure 2. Example of a simplified goal-oriented process

An interesting concept for addressing both concerns are goal-oriented workflows [16]. In this approach, Belief-Desire-Intention (BDI) agents [3] are used to model workflows. The BDI approach allows the workflows to be based on business goals (see Figure 2), starting with the overall goal the

business aims to accomplish with the workflow which is then subdivided into multiple subgoals with various types of interactions such as priorities and inhibitions. These subgoals taken together represent the top-level goal by completely covering all aspects of that goal. The subgoals themselves can again be further decomposed until the goals are plain enough to be achievable by a simple action-based workflow. If multiple solutions for a goal are possible, multiple plans can be attached to the goal and a plan is chosen at runtime based on attached conditions and attributes. Key for both the conditions and attributes of goals and plans is the existence of a workflow context, which not only holds the internal state of the workflow but also contains information about the current business situation, allowing this information to be accessed by goal and plan conditions.

The methodology of generating GPMN models is similar to Hierarchical task network (HTN) [8] planning approaches, however, while the goal-structure of GPMN processes often end up being hierarchical, they do have to be. In addition, GPMN-modelling is not used for feasibility analysis or ahead-of-time planning. Rather, it allows process engineers to specify business contingencies and competing objectives for long-running processes in dynamic environments such as R&D processes at Daimler AG [15]. Instead, the resulting agent model can be directly executed by a BDI interpreter, which uses a BPMN interpreter for concrete plans. The reasoning of BDI is split between the goal deliberation cycle, which decides which of potentially several conflicting goals to follow and the means-end reasoning which chooses pre-defined plans to achieve selected goals. The BDI interpreter reasoning engine uses an approach called Easy Deliberation [23] to resolve this cycle.

While this top-down and business-driven approach towards workflow modeling is very intuitive and flexible, its adaptive behavior during runtime means that actions are not always attributable to the goals [14]. In addition, goal-based workflows are also often used in dynamic environments where not only the workflows but also the workflow management system are distributed to increase flexibility and robustness [17]. As a result, a distributed monitoring system is necessary that links the actions performed in the workflow with the goals and plans representing the business reasons.

Based on the requirements for a distributed workflow management system, distributed workflow execution and cohesion between goals and action, the proposed system should be able to meet the following objectives:

- Goal-Action Cohesion
  - The primary goal of the approach is providing event information that allow attribution of concrete business actions with business goals in a workflow.
  - The user must be able to perform a “drill-down” analysis to trace causes of events from the most detailed to the most high-level goals.
- Distributed Workflow Management
  - The system must be adaptable to a wide variety of business infrastructures, including transient and mobile systems and must respond in a robust fashion to changes in that infrastructure.
  - As a result, the system must be distributed, redundant, robust and adaptable. Communication must be kept low for efficiency.

In this paper we present a component- and service-based approach attempting to solve these objectives. It supports monitoring distributed goal-based workflows using a hierarchical structure of events combined with a distributable monitoring system for gathering and redistribution of the events to the workflow clients.

## 2. RELATED WORK

The approach presented in this paper primarily touches two important area of research. The first is the area of *Business Activity Monitoring (BAM)*. This area concerns itself how business transactions happening within a company can be recorded and processed, either retroactively or in realtime. The most common approach to this challenge is to record business transactions, either specifically for

monitoring purposes or incidentally as part of regular business record keeping, in a *data warehouse* [5]. In addition, *extract, transform, load (ETL)* processes can be used to extract additional data from other sources within the business [25]. The resulting data can be utilized in two ways: Complex analysis can be performed and long-term statistics can be gained by applying data mining techniques to the data available in the data warehouse [1]. This offers the user an in-depth and long-term perspective of the performance of the organization. However, it can require substantial time and computation to process the available data and thus may lag behind the current development of the business. For example, *online analytical processing (OLAP)* allows for multidimensional analysis of transactions and currently available data within the organization [2], but the data must already be available in a structured fashion, for example in a data warehouse. The data is processed to allow the user to approach it from multiple perspectives using multidimensional analysis [24] by forming structures such as OLAP cubes.

Alternatively, realtime monitoring can be achieved using *complex event processing (CEP)* [10]. This approach gathers and processes events as a stream and generates useful information for the BAM system. While aspects of this approach are similar to the approach presented here, it is focused on the processing of data, rather than providing a stronger coherence between the operational and strategic level by aligning actions and goals.

The resulting information can then be used to display information in a dashboard specific to the interest to the business user. Statistics and indicators are provided to allow the user to monitor the actual performance of the business and compare it with previously strategically defined *key performance indicators (KPI)*, which represent quantifiable values. However, the relationship between the processed data and the KPIs is implicit and cannot be derived from the data warehouse itself. As a result, the dashboard functionality of BAM systems often need to be customized to restore this relationship, which can only be partially compensated through standardization of common KPIs or interpretation of standardized charts on the dashboard. Furthermore, BAM focuses primarily on statistical data and thus does not provide an easy way for attributing transactions of the business with particular strategic goals. In contrast, the approach presented here focuses on attributing actions and tasks to business goals, while the generation of statistics is less of a concern. In addition, data mining solutions often focus on a centralized data warehouse solution while the monitoring system presented here provides distributed realtime monitoring.

The second area concerns itself with distributed event systems [21]. A common approach here is to employ *event brokers*, which receive published events from *event publishers* and then distribute them among the other brokers within the system, making them available to *event subscribers* throughout the system. Similar to this approach, we use one or more monitoring component instance to implement the broker functionality and distribute our goal-based events as part of a larger distributed workflow management system. The number of such brokers used can be chosen by the user of the system, with additional brokers increasing the redundancy and thus the resilience of the overall system and increasing the event processing performance by bundling event transfers.

### 3. EVENT STRUCTURE

The targeted goal-oriented workflows are based on the Jadex Active Components [22] approach. This approach offers a number of different active component types such as micro agents, BDI agents, BPMN workflows and the goal-oriented GPMN workflows, which are modeled as goal-plan hierarchies but are translated to BDI agents for execution.

Events generally denote an occurrence associated with an *event source*. An event source can be any entity or part of an entity within the distributed system such as an active component itself or some part of the component. In order to broadly distinguish events based on sources such as components, goals and plans, an event source can be attached to the event. A number of default categories are already provided for all active component types, including the component category for events emitted by component instances itself, the execution category for the execution of components steps, the service category for component services and the property category for component property. More specialized categories are available for action-oriented and goal-oriented workflows. Event

source categories for activities, goals and plans are available in addition to the workflow context fact category. The latter category is also an example of an event source category where the modification event applies.

While the exact nature of the event can vary wildly, a number of categories can be identified that not only define the relationship with the event source and between events over time. For example, BPMN offers three broad classes of event types: Start events, which mark the start of a process, end events, which denote the process end and finally intermediate events which can occur while the process is running.

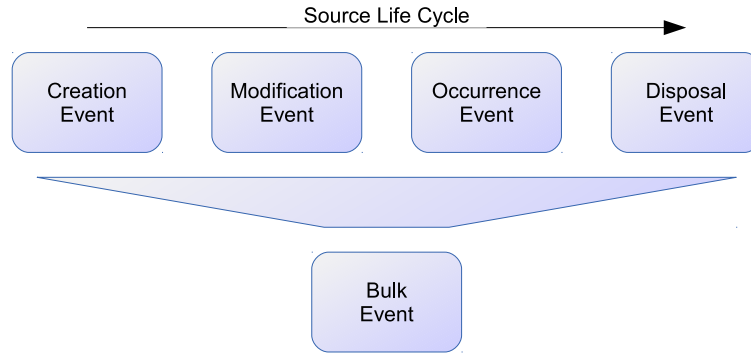


Figure 3. Event types used over the event source life cycle

Similar to this, we have based our system on five different types of events based on the life cycle of the source (see Figure 3). The first type are *creation events*, which are issued when the event source is first created, providing similar semantics as BPMN start events. The counterpart for end events are *disposal events*, issued when the life cycle of the event source has ended and the source has been disposed. During the life cycle, two additional types of events can occur: When the state of the source changes it results in a *modification event*, while the other event type is the *occurrence event*, which simply denotes a point in time when an action took place. For example, an occurrence event is generated when an agent receives an external message or an internal user event is triggered. Finally, a special type of event, the bulk event, only applies when multiple events are aggregated into a single event for efficiency when transferring events over the distributed system.




Figure 4. Events can have both sources and causes based on the static component model and their runtime instances

Since the active component model uses a static component model, the component hierarchy can be used to denote both the source and the *cause* of an event (cf. Fig. 4). For example, a goal-oriented workflow component, when executed, consists of a component instance. If this component instance adopts a new goal instance, a creation event is issued. The source of this creation event is the newly-created goal instance, however, the cause of the event is the component instance which adopted the goal. Once the means-end reasoning decides on a plan to perform in order to reach the goal, a plan instance is created and another creation event is issued. In this case, the source is the plan instance while the cause of the event is the goal instance that triggered the means-end reasoning. This chain can be traced down the model hierarchy down to tasks within BPMN-based plans.

Since instances have unique identifiers, the events merely have to include the cause and source identifiers, minimizing the size of each event. Nevertheless, once the events have been gathered, the

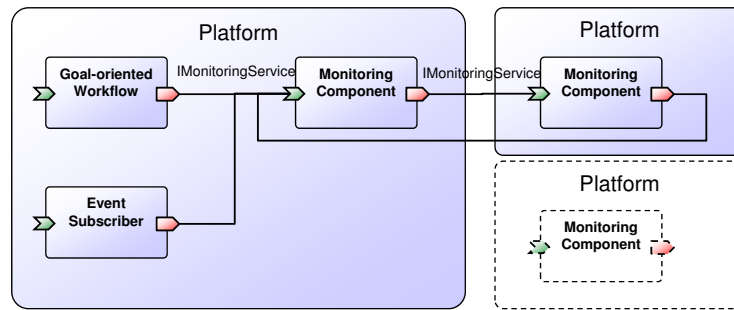


Figure 5. Monitoring infrastructure

cause and effect chain can be recreated by cross-referencing them with other events. In addition, events that are delayed due to the distributed nature of the system can be identified and estimates can be given. For example, if the creation event of the plan instance is missing, the lower boundary for the creation time of the plan instance is the creation time of its cause, i.e. the goal instance. This can at the very least provide the user with an adequate approximation until the missing events arrive.

Each event also requires a timestamp, identifying when the event occurred, especially in relationship to other events. In distributed systems, simple timestamps often pose an issue since it is difficult to reliably synchronize clocks between nodes. Other approaches such as vector clocks increase both complexity and data volume. However, in the case of goal-oriented workflows as presented here, the problem is reduced since each individual workflow instance runs on a single node and its events are therefore internally consistent. Only if the cause chains crosses node barriers, for example due to service calls, caution has to be applied regarding synchronization. Nevertheless, small inconsistencies can be corrected to a certain degree using the same approach as mentioned above regarding missing events.

#### 4. MONITORING ARCHITECTURE AND IMPLEMENTATION

The monitoring mechanism has been implemented using a monitoring service (`IMonitoringService`), which allows producers to publish events and consumers to subscribe for certain types of events. In Fig. 5 an overview of the architecture is depicted. It can be seen that in each platform a specific monitoring component realizes the monitoring service. Each component that creates (internally or intentionally) events, automatically publishes them to the corresponding service. For this purpose each component searches for an `IMonitoringService` when an event has to be published. In case a service is found the binding is fixed and the event is forwarded, if not the component stores the unavailability of the service and the event is dismissed. It will try to search for the service again when a new event occurs and after some specified time interval has elapsed.

Event consumers can subscribe at the monitoring service using an optional event filter. This mechanism allows for reducing the network load as only events are transferred which pass the filter test.

The global monitoring infrastructure is set up as a peer to peer infrastructure formed by multiple monitoring components realizing an information exchange protocol. Knowing that events are reported locally from the event sources, each monitoring service searches for all other monitoring services and forwards local events to all remote services. In this way all monitoring services internally build up a globally consistent event state. For scalability reasons, the monitoring components only hold a certain amount of event in memory and dismiss older events. If longer lasting book keeping is necessary they can also be configured using a distributed storage service (see Section 5) can be used to distribute the events within the system, which, in addition to providing its own synchronization, persistently stores older events.

#### 4.1. Workflow environment

The monitoring service is further used to supplement a distributed workflow management system [17] implemented as active components.

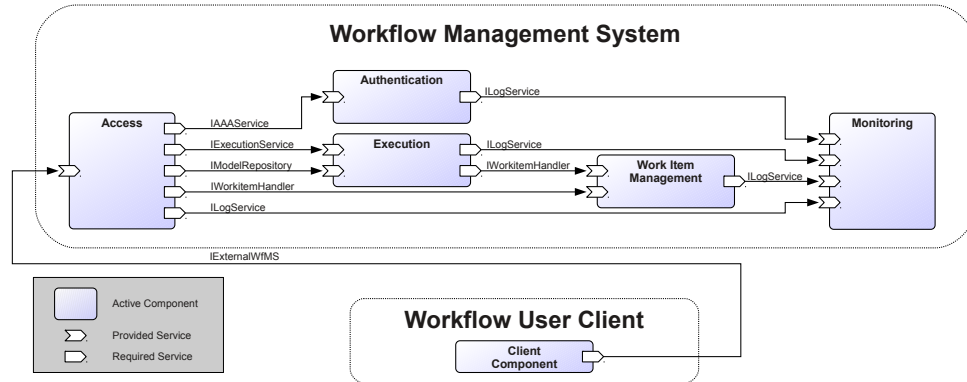


Figure 6. Event types used over the event source life cycle

The system is largely derived from the Workflow Management Coalition Reference Model [11], where the monitoring service represents the monitoring subsystems of a traditional workflow management system. The workflow management system consists of multiple active components similar to the monitoring service which use service calls to exchange information.

Aside from the monitoring, other components are available to provide the rest of the workflow management functionality. The access component manages the access to workflow management functionality for workflow clients. The authentication component verifies the credentials of users and specifies their access rights. The execution component stores workflow models and launches workflow instances. Finally, the work item component holds work items representing human tasks for users to perform. Each component can be replicated to provide robustness and scalability.

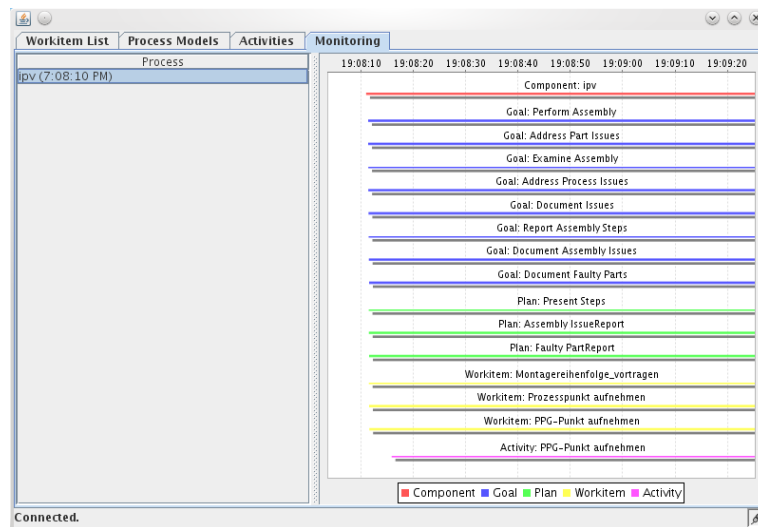


Figure 7. Gantt chart monitoring of a goal-oriented workflow in the workflow client

The information gained by the monitoring component can be used to provide an overview of a goal-oriented workflow instance. As shown in Fig. 7, users can select a running or past workflow instance and is presented with a Gantt chart of event sources. The user can click on individual sources and is provided with a view which includes the clicked item and sources that were created as a result of the item, providing a drilled-down view on parts of the process instance. If only partial



information is available, an estimate can be shown to the user based on other event sources in the hierarchy. However, due to the chaining of references a subtree may be omitted if the cause-source hierarchy is interrupted. While this is remedied as soon as the missing information arrives, it is still a limitation of the system.

## 5. EVENT PERSISTENCE AND DISTRIBUTION

An important aspect in such a dynamic and distributed workflow environment is ensuring that the events are distributed among the nodes and persisted in a manner that allows nodes in the system to disappear without events being lost. The disappearance of nodes can happen either through normal operation such as a shutdown or through exceptional occasions like an abnormal termination of a node or network outage.

The most extreme case is a complete shutdown of the system. After a subsequent recovery, users of the system will expect the events to recover as well. As a result, the events have to be persisted in permanent storage such as a hard disk.[7]

A straightforward solution for this issue is simply storing the event data locally into a database or file. While this approach is likely to be the fastest in terms of performance (provided distribution is done asynchronously), it replicates the events on every node. If there is a large number of nodes, this means that events are replicated excessively, beyond what is necessary for recovery after disappearance of nodes.

An alternative for storing the event data locally and using the default synchronization mechanism for monitoring components, a remote database could be used to store the events and to synchronize the monitoring services. However, this represents a single point of failure: The system is expected to retain functionality even in the event of unexpected failures of nodes in the system. As a result, a distributed storage solution is required.

This means that the solution for the problem is ultimately limited by the CAP-theorem (see [9]), which means that the system cannot guarantee consistency, availability and partition tolerance at the same time. However, since events are merely recorded and not further processed and the events are generated locally in a consistent manner, there is no possibility for a disagreement of event information due to partitioning in the scope of the monitoring service. The only exception of this would be a malicious monitoring service injecting fraudulent events, which is excluded by fact that the system being assumed to be a friendly environment (see [17] for a discussion on the security aspects of the workflow management system).

### 5.1. Existing Distributed Storage Solutions

A first approach could be a replicated relational database such as MySQL configured in a master-slave configuration could be used (see [12]). Two modes of consistency are available for this option, eager replication and lazy replication. Eager replication ensures the consistency of the data across all nodes but increases the overhead due to the additional synchronization required. In contrast, lazy replication employs a master-slave architecture where data can only be stored using the master node but can be read from any slave node. Commits performed on the master node are asynchronously transferred to the slave node, which may not always contain the most up-to-date data set.

However, when the master node fails, there is no easy mechanism for a failover. The system setup is rigidly pre-configured and the promotion of a slave node to a new master node as well as the addition of new nodes requires manual intervention. Furthermore, the data is always replicated on all slave nodes, regardless of the need for such replication, similar to the local storage approach.

Due to this, a more specialized solution outside the traditional strict relational approach could be used, for example by using a distributed NoSQL database which depart from the hard consistency requirements and structural expressiveness of the SQL-based approaches. Interesting candidates for storing the events in a distributed manner are Amazon Dynamo (see [6]), a distributed key-value store and Google Bigtable (see [4]), which uses a more structured data model called Column Family.

All of these approaches have some drawbacks when considering them for the distributed storage of events in a dynamic and distributed workflow management system. First, the configuration and setup is relatively complex and requires an explicit configuration on certain parameters, such as the number of nodes involved and in some cases designation of master nodes, a task requiring non-trivial expertise and later maintenance once the system has been configured. Since the configuration needs to be adapted to appearing and disappearing nodes, it does not lend itself well to a highly dynamic workflow environment where the nodes of the system may appear and disappear any time. In addition, while the systems can recover from the loss of nodes and maintain availability, node replacement generally must be configured to incorporate the additional nodes. This is often not an automatic process and requires the intervention of a system administrator with a centralized authority for the system. As mentioned in Section 1, such a centralized authority may not be available since the participants of the business processes targeted by the system are independent and autonomous actors who often maintain their own IT infrastructure.

### 5.2. Dynamic Distributed Storage Service

As a result, the monitoring component uses a distributed persistence service to store the events for each of the monitoring components. The system is based on the concept of a key-value store, similar to a distributed hash table but assumes a more dynamic environment. Application can store values in the storage system and request a certain replication level to ensure availability. This replication level is maintained when nodes disappear by replacing them with other nodes available in the system. Additional nodes that become available are automatically integrated in the system.

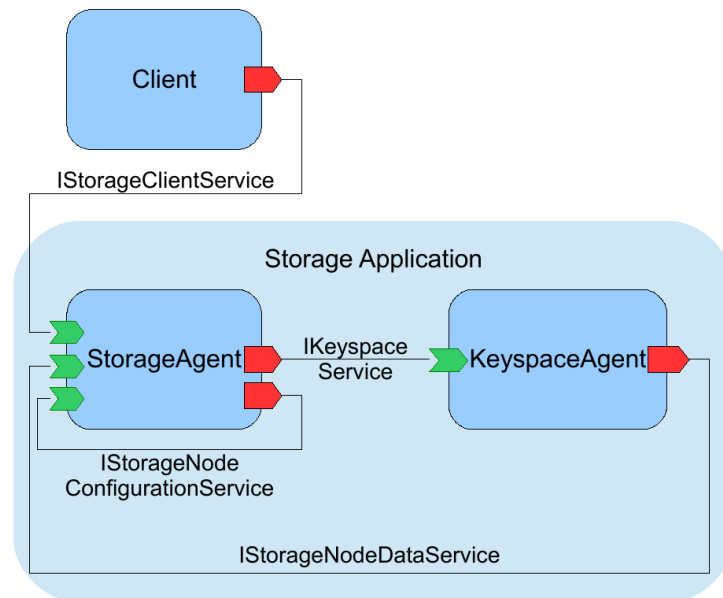


Figure 8. Architecture of the storage service, from [7]

Figure 5 shows the architecture of the system. The monitoring system can access the storage service by requesting a service implementing the `IStorageClientService` interface, which offers the relevant methods to store values. The distributed nature of the service is completely transparent for the monitoring component, it simply stores the event using the storage service which is responsible for further distribution. The storage is segmented using keyspaces in order to separate the storage of multiple applications or parts of applications.

Internally, the `StorageAgent` component performs the distribution of the data as necessary for the required replication level and maintains the records of this in case addition replication becomes necessary due to lost nodes.

The local storage for each node of the storage service is represented by the KeyspaceAgent, which uses an internal local database to store the values. If a StorageAgent receives a request for a certain value, it can attempt to request it from its local KeyspaceAgent, or, if it is locally unavailable, locate it by calling other StorageAgents, with caching allowing a quick lookup if the value had already been requested.

Since nodes in the system can disappear at any point, the storage service uses a peer-to-peer approach where every node is considered equal and can all be used to record data. As a result, the emphasis of the service is on the availability. In addition, the system is also resilient with regard to partitioning. Since this means that strict consistency cannot be guaranteed, the storage service uses vector clocks to determine inconsistent versions of recorded values in to allow a resolution of such a conflict at a later point. However, as noted before, the generated events are only assume a single value and as such cannot trigger conflicts in the storage service.

## 6. EVALUATION AND OUTLOOK

In Section 1 we provided two areas where the system had to fulfill a set of requirements. The first area involved the cohesion between business goals and workflow actions. Here, the monitoring system is required to establish a relationship between the actions of the workflow and the business goals and allow the user to “drill down” from individual goals down to specific actions. The modeling of the goal-oriented workflows allow the events to establish a cause-source hierarchy to attribute actions to goals.

The first set of requirements was due to the required flexibility of a distributed workflow management system. The events may be generated by workflows anywhere within the distributed system, forwarded to interested nodes and the system should be robust and tolerate disappearing nodes. These three requirements were achieved by implementing the monitoring system as a event broker, forwarding all events it receives for publication to corresponding services. Unless a node is permanently disabled before it can transmit new events, all events will eventually be available to the workflow client. Low communication overhead was achieved by minimizing the amount of information stored in the events, including only references to sources and causes and letting the receiver reconstruct the chain. Furthermore, bulk events are available to bundle multiple events for transfer.

The storage system has been tested regarding its functional requirements. This included tests for data access for single and multiple node storage services, adding additional nodes to the system, deliberate shutdown of nodes and deliberate termination of nodes simulating an unexpected disappearance of a node (see [7]). In each case the system performed as expected and maintained availability to the system.




Figure 9. Write performance for 1,000 write operations of the storage service compared to local database storage (from [7])

In order to ensure the performance of the service, measurements have been taken under different conditions (see Fig. 9) using 1000 write operations each. As a reference point, the performance of a local unversioned database system and a local versioned database system have been measured. The storage system was then tested with with a single node and multiple nodes, with the writes either

committed consecutively or in parallel. As the figure shows, while the performance does degrade compared to local storage, the overhead is not overly cumbersome and can be acceptable when the advantages of the distributed approach are considered.

However, a remaining concern is the potential loss of multiple events of a specific event source before they are recorded. Due to the events containing only a minimum of information, this would interrupt a link between the main hierarchy tree and one of its subtrees. While there is a balance involved with regard to the event sizes, this challenge could be address by including a more information about the chain in each event by including not only the current causes but also a number of previous causes, allowing the workflow client to include the subtree and only omit the missing source.

## REFERENCES

1. M. J. Berry and G. Linoff. *Data Mining Techniques: For Marketing, Sales, and Customer Support*. John Wiley & Sons, Inc., New York, NY, USA, 1997.
2. A. Berson and S. J. Smith. *Data Warehousing, Data Mining, and Olap*. McGraw-Hill, Inc., New York, NY, USA, 1st edition, 1997.
3. M. Bratman. *Intention, Plans, and Practical Reason*. Harvard University Press, 1987.
4. F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. *ACM Trans. Comput. Syst.*, 26(2):4:1–4:26, 2008.
5. S. Chaudhuri and U. Dayal. An overview of data warehousing and olap technology. *SIGMOD Rec.*, 26(1):65–74, March 1997.
6. G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: amazon’s highly available key-value store. *SIGOPS Oper. Syst. Rev.*, 41(6):205–220, 2007.
7. F. Demuth. Verteilter Speicherdienst für Nutzer-integrierende dynamische Cloud-Umgebungen. Diplomarbeit, Distributed Systems and Information Systems Group, Computer Science Department, University of Hamburg, November 2014. (in German).
8. K. Erol, J. Hendler, and D. S. Nau. Htn planning: Complexity and expressivity. In *AAAI*, volume 94, pages 1123–1128, 1994.
9. S. Gilbert and N. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, June 2002.
10. T. Greiner, W. Düster, F. Pouatcha, R. von Ammon, H.-M. Brandl, and D. Guschakowski. Business activity monitoring of norisbank taking the example of the application easycrredit and the future adoption of complex event processing (cep). In *Proceedings of the 4th International Symposium on Principles and Practice of Programming in Java*, PPPJ ’06, pages 237–242, New York, NY, USA, 2006. ACM.
11. D. Hollingsworth. *Workflow Management System Reference Model*. Workflow Management Coalition, 1995.
12. Oracle Inc. *MySQL 5.7 Reference Manual Replication*, 2014.
13. K. Jander, L. Braubach, and W. Lamersdorf. Distributed event processing for goal-oriented workflows. In *Intelligent Distributed Computing VIII*, Studies in Computational Intelligence, pages 49–58, Heidelberg, 2014. Springer International Publishing.
14. K. Jander, L. Braubach, A. Pokahr, and W. Lamersdorf. Jadex WfMS: Distributed Workflow Management for Private Clouds. In *Languages, Methodologies, and Development Tools for Multi-Agent Systems*, pages 39–55. Springer Berlin / Heidelberg, 2011.
15. K. Jander, L. Braubach, A. Pokahr, W. Lamersdorf, and K.-J. Wack. Goal-oriented processes with gpmn. *International Journal on Artificial Intelligence Tools (IJAIT)*, 20(6):1021–1041, 12 2011.
16. K. Jander and W. Lamersdorf. Gpmn-edit: High-level and goal-oriented workflow modeling. In *WowKiVS 2011*, volume 37, pages 146–157, 2011.
17. K. Jander and W. Lamersdorf. Jadex WfMS: Distributed Workflow Management for Private Clouds. In *Conference on Networked Systems (NetSys)*, pages 84–91. IEEE Xplore, 2013.
18. N. Jennings, T. Norman, and P. Faratin. ADEPT: An agent-based approach to business process management. *ACM SIGMOD Record*, 27(4):32–39, 1998.
19. F. Leymann and D. Roller. *Production workflow concepts and techniques*. Prentice Hall PTR, 2000.
20. Object Management Group (OMG). *Business Process Modeling Notation (BPMN) Specification*, version 2.0 edition, January 2011.
21. P. R. Pietzuch and J. Bacon. Hermes: A distributed event-based middleware architecture. In *Proceedings of the 22Nd International Conference on Distributed Computing Systems*, ICDCSW ’02, pages 611–618, Washington, DC, USA, 2002. IEEE Computer Society.
22. A. Pokahr and L. Braubach. The active components approach for distributed systems development. *International Journal of Parallel, Emergent and Distributed Systems*, 28(4):321–369, 2013.
23. A. Pokahr, L. Braubach, and W. Lamersdorf. A goal deliberation strategy for bdi agent systems. In T. Eymann, F. Klügl, W. Lamersdorf, M. Klusch, and M. Huhns, editors, *Proceedings of the 3rd German conference on Multi-Agent System TEchnologieS (MATES-2005)*. Springer, 2005.
24. P. Vassiliadis and T. Sellis. A survey of logical models for olap databases. *SIGMOD Rec.*, 28(4):64–69, December 1999.

25. P. Vassiliadis, A. Simitsis, and S. Skiadopoulos. Conceptual modeling for etl processes. In *Proceedings of the 5th ACM International Workshop on Data Warehousing and OLAP*, DOLAP '02, pages 14–21, New York, NY, USA, 2002. ACM.
26. M. Weske. *Business Process Management Concepts, Languages, Architectures*. Springer Verlag, 2007.