

WISE: An Infrastructure for E-Commerce

G. Alonso U. Fiedler A. Lazcano H. Schuldt C. Schuler N. Weiler

Swiss Federal Institute of Technology (ETH)

ETH Zentrum, Zürich CH-8092, Switzerland

E-mail: wise@ccic.ethz.ch

<http://www.inf.ethz.ch/department/IS/iks/research/wise.html>

May 7, 1999

1 Introduction

The Internet and the proliferation of inexpensive computing power in the form of clusters of workstations or PCs provide the basic hardware infrastructure for business to business electronic commerce in small and medium enterprises. Unfortunately, the corresponding software infrastructure is still missing. As part of the WISE project (Workflow based Interned SErvices), we have taken concrete steps to address this limitation: within WISE we have developed and deployed the software infrastructure necessary to support business to business electronic commerce in the form of virtual enterprises. The starting point was the idea to combine the tools and services of different companies as building blocks of a higher level system in which a process acts as the blueprint for control and data flow within the virtual enterprise. From here, the goal has been to build the basic support for an Internet trading community where enterprises can join their services to provide added value processes.

Following these ideas, WISE provides a working system capable of defining, enacting, and monitoring virtual enterprise business processes, as well as supporting related coordination activities. Such infrastructure includes an Internet workflow engine acting as the underlying distributed operating system controlling the execution of business processes, a process modeling tool for defining and monitoring the processes, a catalogue tool for virtual enterprise services in which to find the building blocks for the processes, and a collaborative multimedia communication environment. The project also incorporates in its design considerations about security, quality of service, execution guarantees, exception handling, high availability, and scalability, as well as diverse other aspects related to WWW based interaction, catalogue based information, catalogue search, and communication frameworks. In this regard, we have made a substantial effort to make WISE a complete solution, that is, a system incorporating all the necessary functionality to be used in practice. We firmly believe the real challenge in electronic commerce is how to provide a complete solution. In our case, this meant to develop a software tool capable of supporting the entire life cycle of a virtual business process. We see these business processes as valuable assets which need to be not only defined and enacted but also maintained, updated, and monitored. WISE supports all these chores, thereby avoiding the drawbacks of many existing products: ad-hoc and costly development,

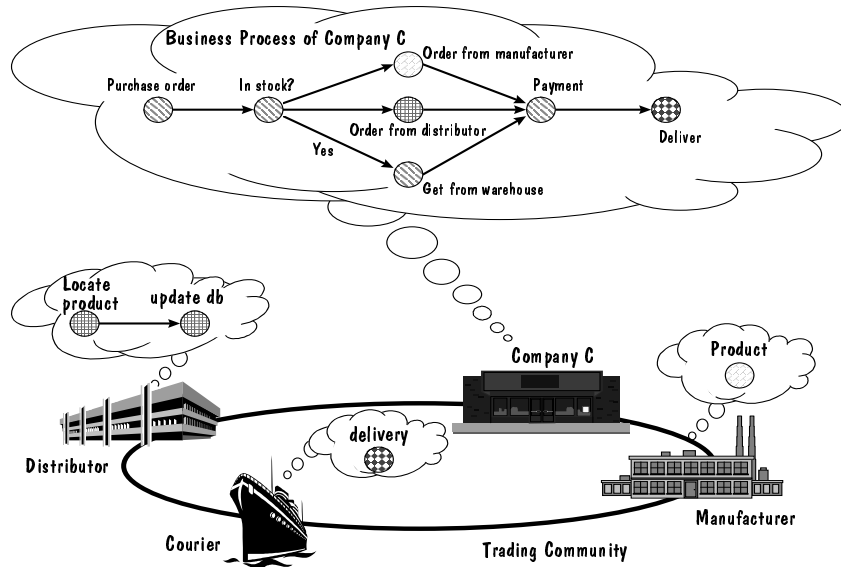


Figure 1: A company incorporating a virtual process as part of its own business processes

expensive maintenance, and of limited applicability.

2 Motivation

2.1 Virtual Enterprises

The most relevant activities within a corporation are often described in the form of business processes. This is not surprising since business processes model the procedures and rules followed in order to accomplish a concrete goal (open a new bank account, obtain a credit, purchase a computer, find out the current location of a parcel, resupply shops, etc.). Following this idea, we see electronic commerce as the incorporation of information and communication systems technology into the business process to expand it beyond the corporation boundaries. In this context, we define a *virtual business process* as a business process whose definition and enactment cannot be directly tied to a single organizational entity (be it a department or a company). From here, we define *virtual enterprises* as those whose business processes are virtual business processes. Given the trend towards decentralization, we also consider that virtual business processes linking together several departments of a single organization define as well a virtual enterprise. Finally, we refer to the set of companies participating in a virtual enterprise as a *trading community*. Each member of the trading community provides a number of services to be used as building blocks for the virtual process. Based on these services, the virtual enterprise can be created by defining a virtual process in which each individual activity corresponds to one of the services provided by the participants (Figure 1).

We believe trading communities, virtual enterprises and virtual processes are a very powerful approach to interpret and identify the needs of a wide range of electronic commerce practices. In particular, they are well suited to the four types of electronic commerce described above. For

instance, in the case of retailing, a company can provide a much more sophisticated product by outsourcing aspects of the operation which are not central to its activities. A common example are companies offering a product (books, CD, flowers) without actually handling (producing, storing or delivering) the product themselves. Most of the handling is left to companies providing specialized services, which allows to significantly reduce the operational costs. The virtual enterprise model naturally captures such scenarios by simply having the distribution and delivery services incorporated as activities within the business processes of the company selling the product as shown in Figure 1.

2.2 E-Comm Processes

In the example of Figure 1, independently of whether it involves mainframes and leased lines or a few PCs linked via Internet providers, an E-Comm application has many of the characteristics of a distributed computing environment. While the notions of trading community and virtual enterprise are conceptually useful, the real challenge is to use them in a software solution. Here is where the idea of process becomes relevant: the virtual business process can be seen as a distributed program running on some form of middleware linking together the resources of the trading community. These resources are the concrete applications or services offered to the virtual enterprise by the trading community and are used as the basic building blocks for the distributed program (the virtual business process). From here, the type of software to develop is the type of software that would be needed to support the definition and execution of such a coarse grained distributed program.

The analogy between an E-Comm process and a distributed program can be taken a step further. Any realistic solution to electronic commerce must take into account the true complexity of the problem. We see E-Comm processes as valuable assets needing to be properly specified, designed, developed, tested, debugged, and maintained in an effort not unlike software life-cycles. In order to do this, the language used to describe the processes must provide the necessary primitives, otherwise these tasks become extremely difficult and largely ad-hoc endeavors (as it is today).

2.3 Complete Solution

The WISE project is an integration effort with the final goal of providing a complete solution. Its architecture (Figure 2) is organized into four components (definition, enactment, monitoring, and coordination), each one of them with the role of addressing a particular issue. Thus, the *process definition* component allows virtual business processes to be defined using as building blocks the entries of a catalogue where companies within a trading community can post their services. Similarly, the *process enactment* component compiles the description of the virtual business process into a representation suitable for enactment and controls the execution of the process by invoking the corresponding services of the trading community. The *process monitoring and analysis* component is a tool keeping track of the progress made in the execution of the virtual business process and of the status of all active components in the system. The information produced by these tool is used to create an awareness model [7] used for load balancing, routing, and quality of service purposes as well as, later on, for analysis of the behavior of the process. Finally, the *coordination and communication* component supports multimedia conferencing and cooperative browsing of relevant

information between all participants in the trading community using the information produced by the business process as the main source for routing.

In WISE, these four components are tightly integrated reflecting an approach to electronic commerce based on transparency and ease of use. While there is innovation in each individual component, our measure of success is the degree of integration of the system as a whole.

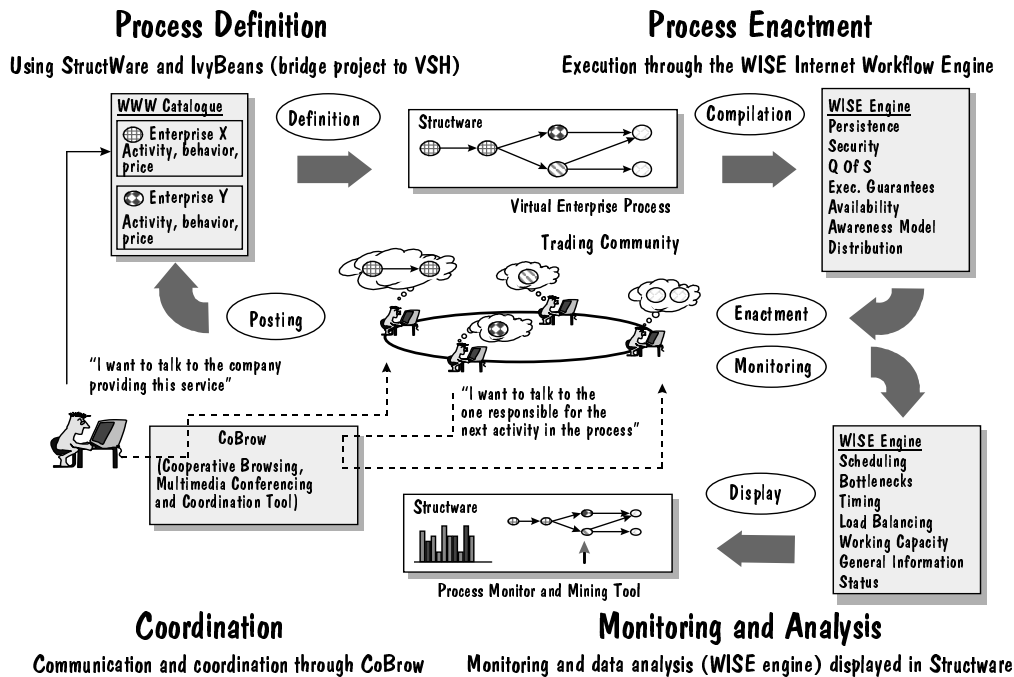


Figure 2: The different components of the WISE project

3 Architecture and Components

3.1 Modeling

In WISE, virtual business processes are constructed by using the services offered by different companies as building blocks. The virtual business process integrates the services of the different companies establishing the order of invocation, the control logic and the data flow between the participants in the same way a workflow process orchestrates business models within a single corporation. To make this idea a reality, there are two elements that WISE must provide. The first is a mechanism for the participants to publish their services. The second is a way to define a process based on such services. For these purposes, WISE uses a WWW catalogue and a business process modeling tool (Figure 2).

The WWW catalogue uses Java applet/servlet technology to allow companies in the trading community to advertise their services and to “see” the semantics of the services provided by other companies [15]. The catalogue contains objects encapsulating the behavior of a each service. A Java version of a business modeling tool supporting simulation and analysis (see below) is then

used to allow a company to see the exact characteristics of each entry in the catalogue. When a company wants to make an entry in the catalogue, it specifies the service using the modeling tool.

From the catalogue, a drag and drop type of interface is used to build the virtual business process. The tool we use for process definition is *Structware* [12], a product of IvyTeam, one of the partners in the project. Structware, which is internally based on Petri-nets, supports not only the modeling of business processes but also sophisticated analysis of its behavior (bottlenecks, average execution times, costs, delays, what if analysis, etc.). In terms of process definition, Structware supports the standard flow control primitives of a workflow tool. It is possible to define conditional branching, nested processes, and assign additional information to each task within the process. This last point is important from the point of view of WISE since it allows to use this additional information as the configuration information necessary to enact the process

We see this entire procedure as a form of high level, coarse grained programming. We have successfully applied this idea of “workflow programming” within WISE and other projects in order to provide sophisticated language primitives not available in commercial workflow tools. For instance, we can provide a complete exception handling capability [10], an event handling mechanism and inter-process communication [11]. This functionality is missing in current systems and we consider it to be crucial in realistic environments.

3.2 Enactment

The enactment of the virtual business processes is performed by the WISE engine, which is based on work done within the OPERA project [9, 2]. The WISE engine extends ideas from workflow management [6, 3], and uses known techniques for distributing this functionality [18, 13, 4]. In addition, a considerable amount of extensions have been introduced to make workflow a suitable foundation for electronic commerce (for a different approach to electronic commerce based on workflow technology see [16]). Among them, there are three that deserve special attention: security, quality of service, and execution guarantees.

Given the nature of the data exchanged between the different participants in the trading community, WISE incorporates the necessary security mechanisms in the form of encryption of data for transmission over the network as well as a complete set of authentication measures for both execution, access, and monitoring of the processes. Also, to make the notion of trading community viable given the current limitations of bandwidth, the WISE engine incorporates quality of service guarantees based on execution statistics and network characteristics. Our current approach is based on distinguishing different process categories (critical, important, normal) and providing for each of them a different quality of service. Finally, the WISE engine also incorporates execution guarantees, whereby a process is always guaranteed to finish in a consistent state either by removing all changes it has introduced or by forcing it to terminate following a sequence of actions with a pre-determined outcome [17]. The execution guarantees are based on the notion of spheres of atomicity and isolation [8, 5, 1, 14], which allow us to specify which parts of the business process need to be made atomic for recovery purposes and which parts of the process need to be isolated from interferences of other processes.

3.3 Audit and Monitoring

WISE provides tools to find out the status of any running process in the system in order to allow users to keep track and troubleshoot them when necessary. In addition, process design is a difficult task. In virtual enterprise environments it is difficult to foresee all possible eventualities until some example runs are available. Process design is an iterative procedure where WISE can be of great help by providing accurate measurements of all the characteristics affecting the execution of a process: overall duration, bottlenecks, relative duration of each task with respect to the duration of the entire process, loads at each participant site, deadlines missed, and so forth.

In order to provide this functionality, WISE incorporates the necessary modules within the execution engine to keep track of executing processes. In addition, it uses a history space where information about all already executed processes is stored and organized in a way that facilitates its analysis. For displaying this information, we plan to take advantage of the capabilities of Structware. In the same way that a Structware process is compiled and translated into notation understandable by the WISE engine, the information produced by the WISE engine will be translated into the appropriate format to be displayed using Structware's interface.

Finally, WISE will also include an awareness model [7] that will allow the engine to make decisions based on its own status and that of the participants. This awareness model is necessary for load balancing, increased availability, conflict resolution, notification mechanisms, and the handling of exceptions.

3.4 Coordination

Unlike in conventional workflow engines, WISE will operate in an environment where the different participants and the different elements of the process are not necessarily in a position to easily exchange information among them. Note that, as the concept of trading community implies, each participant could be not only on a different location but in an entirely different company. It is nevertheless important for the participants to be able to communicate in order to resolve the unavoidable inconsistencies and minor problems associated with any process (Figure 2). An essential aspect of this communication and collaboration is that it will be context based. That is, a user will not necessarily ask to communicate with a concrete person but, rather, with the person who played a given role in the execution of the process.

To achieve this goal, WISE uses the results of the CoBrow (Collaborative Browsing in Information Resources) project [19] to avoid having to develop the communication infrastructure as part of WISE.

4 The WISE system

The current version of WISE uses IvyFrame (a commercial product of IvyTeam) as front end, both for the definition and the monitoring of processes. WISE is platform independent (Server runs on UNIX, Clients on UNIX, OS/2, and Windows) and can interact with a variety of applications (existing interfaces include SAP R/3 and IBM FlowMark). From a practical point of view, WISE can be used as a generic workflow engine but its real potential lies as an engine for electronic

commerce. One possible scenario for the deployment of WISE is as the central tool for a company providing support for other companies wanting to engage in electronic commerce but not willing or able to do the necessary investments in resources and expertise. Another possibility is to use WISE as a tool for implementing value added business processes so that a company can offer new services by combining services provided by other companies. In particular, the application to virtual store fronts for generic customer services, computer equipment, bookstores, and appliances is immediate. We are currently working on supporting other possible scenarios related to payment protocols and electronic document exchanges.

5 Conclusions

In this extended abstract, we have presented a basic infrastructure for business to business electronic commerce. In this form of e-commerce, different companies join their services to form a virtual enterprise, which provides a business process that can be executed over the Internet. WISE includes different components to define, enact and monitor virtual enterprise processes, supporting also the communication and coordination between the participants.

WISE should be seen as an integration effort where several known technologies as well as new ideas are being brought together in order to provide a coherent technological solution. We expect that the results of the project will both enhance considerably the scope of application and expressive power of current workflow systems and open up significant opportunities in the area of electronic commerce.

Project Data

The WISE project is funded by the Swiss National Science Foundation. It started in December of 1997 and will have a duration of 29 months. There are three academic and two industrial partners in the project. On the academic side, the participants are the Database Research Group, the Computer Engineering and Networks Laboratory, and the Information and Communications Systems Research Group of ETH Zürich. The industrial partners are IvyTeam, and *online*SOLUTIONS.

References

- [1] G. Alonso, D. Agrawal, and A. El Abbadi. Process Synchronization in Workflow Management Systems. In *8th IEEE Symposium on Parallel and Distributed Processing (SPDS'96)*. New Orleans, USA., October 1996.
- [2] G. Alonso, C. Hagen, H.J. Schek, and M. Tresch. Distributed Processing over Stand-alone Systems and Applications. In *Proceedings of the 23rd International Conference on Very Large Databases (VLDB'97)*, Athens, Greece, August 1997.
- [3] F. Casati, P. Grefen, B. Pernici, G. Pozzi, and G. Sanchez. WIDE Workflow Model and Architecture. Technical Report 96-19, University of Twente, 1996.
- [4] Stefano Ceri, Paul W.P.J. Grefen, and Gabriel Sanchez. WIDE: A Distributed Architecture for Workflow Management. In *Proceedings 7th International Workshop on Research Issues in Data Engineering (RIDE'97)*, pages 76–79, Birmingham, UK, April 1997.
- [5] D. Georgakopoulos and M. Hornick. A Framework for Enforceable Specification of Extended Transaction Models and Transactional Workflows. *International Journal of Intelligent and Cooperative Information Systems*, 3(3), September 1994.

- [6] D. Georgakopoulos, M. Hornick, and A. Sheth. An Overview of Workflow Management: From Process Modeling to Workflow Automation Infrastructure. *Distributed and Parallel Databases*, 3(2):119–153, April 1995.
- [7] Dimitrios Georgakopoulos. Collaboration management infrastructure for comprehensive process and service management, May 1998. Presentation in International Symposium on Advanced Database Support for Workflow Management, Enschede, The Netherlands.
- [8] Dimitrios Georgakopoulos, Mark Hornick, Piotr Krychniak, and F. Manola. Specification and Management of Extended Transactions in DOMS. In *RIDE-IMS'93*, pages 253–257, Vienna, Austria, April 1993.
- [9] C. Hagen. Atomarität in Workflow- und Prozessunterstützungssystemen. In *9. GI-Workshop "Grundlagen von Datenbanken"*, Friedrichsbrunn, Germany, May 1997. In German.
- [10] C. Hagen and G. Alonso. Flexible exception handling in the OPERA process support system. In *Proc. of the 18th Intl. Conference on Distributed Computing Systems*, Amsterdam, The Netherlands, May 1998.
- [11] C. Hagen and G. Alonso. Beyond the black box: Event-based inter-process communication in process support systems. In *Proc. of the 19th Intl. Conference on Distributed Computing Systems*, Austin, Texas, USA, May 1999.
- [12] IvyTeam. Structware'98 Process Manager. Available through <http://www.ivyteam.com>, 1998.
- [13] S. Jablonski and C. Bussler. *Workflow Management*. International Thomson Computer Press, 1996.
- [14] Sushil Jajodia and Larry Kerschberg, editors. *Advanced Transaction Models and Architectures*, chapter 1, pages 3–34. Kluwer Academic Publishers, 1997.
- [15] H. Lienhard. IvyBeans - Bridge to VSH and the project WISE. In *Proceedings of the Conference of the Swiss Priority Programme Information and Communication Structures, Zürich, Switzerland*, July 1998.
- [16] P. Muth, J. Weissenfels, and G. Weikum. What Workflow Technology Can Do For Electronic Commerce. Technical report, University of the Saarland, Department of Computer Science, Saarbrücken, Germany.
- [17] H. Schuldt, G. Alonso, and H.-J. Schek. Concurrency Control and Recovery in Transactional Process Management. Technical report, ETH Zurich, 1998. Available at: http://www-dbs.ethz.ch/~schuldt/Publications/sas_tr98.ps.
- [18] H. Schuster, S. Jablonski, T. Kirsche, and C. Bussler. A Client/Server Architecture for Distributed Workflow Management Systems. In *Proc. of Third Int'l. Conf. on Parallel and Distributed Information Systems*, Austin, Texas, September 1994.
- [19] G. Sidler, A. Scott, and H. Wolf. Collaborative Browsing in the World Wide Web. In *Proceedings of the 8th Joint European Networking Conference, Edinburgh, Scotland*, May 1997.